



Durham E-Theses

Managing the uncertainties in commodity trading: A Bayesian software implementation

Dow, Laura

How to cite:

Dow, Laura (2007) *Managing the uncertainties in commodity trading: A Bayesian software implementation*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/2857/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Managing the Uncertainties in Commodity Trading

A Bayesian Software Implementation

by

Laura Dow

The copyright of this thesis rests with the author or the university to which it was submitted. No quotation from it, or information derived from it may be published without the prior written consent of the author or university, and any information derived from it should be acknowledged.

A Thesis presented for the Degree of
Master of Science



Department of Mathematical Sciences
University of Durham

2007

17 OCT 2007



Managing the Uncertainties in Commodity Trading

A Bayesian Software Implementation

by

Laura Dow

Submitted for the Degree of Master of Science

2007

Abstract

A decision support software tool is developed to aid commodity traders with the management of uncertainty. A methodology to control the risk of losses is examined and this is likely to be highly desirable for practical use in the financial markets.

Bayesian techniques are used to forecast future prices and this is combined with the optimization of a utility function to support a trader's buy and sell decisions. The difficulties of eliciting both a prior belief and a utility function are discussed.

The effectiveness of the tool is explored using a practical application of the Design of Experiments (using simulated prices) and also by testing the software with traders. The trials of the tool with traders received positive feedback.

Declaration

The work in this thesis is based on research carried out at the Department of Mathematical Sciences, the University of Durham. No part of this thesis has been submitted elsewhere for any other degree or qualification and it all my own work unless referenced to the contrary in the text. Chapter 3 and Chapter 6 are largely based on research by David Wooff, Michael Goldstein, and Allan Seheult.

Copyright © 2007 by Laura Dow.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

Acknowledgements

All my thanks to Dr Robert Parish who provided me with the time and support I needed to complete this MSc. Without his generosity and commitment this thesis would not have been possible.

Thanks to David Wooff, my supervisor, who provided me with the guidance I needed to complete this MSc and made me feel welcome in Durham. This thesis is based on his research, in particular, the use of Bayes Linear Methods to forecast prices and the development of diagnostics to monitor these forecasts. I would like to express my gratitude for all the time David Wooff has spent on this project and for his efforts in giving feedback to my work.

Thank you to Klaas De Boer who worked on the trials with traders and made lots of valuable recommendations for improvement to the software tool.

Thanks to all the staff at Scitech who supported me throughout my time in the south-east. In particular, thanks to Dr John Heritage who helped me with Java and laughed at me, thanks to Dr David Browning for his expertise in the development of the software and for providing me with good advice, and thanks to Dr Sarika Pugla who recommended using WinEdt.

Finally, thanks to my family - Neil, Mum and Dad, Annie and Katie, Colin and Jo, and Maureen who provided me with the encouragement I needed to write and complete this thesis.

Contents

Abstract	2
Declaration	3
Nomenclature	11
List of Figures	12
List of Tables	13
1 Introduction	17
1.1 KTP project	17
1.1.1 The project objectives	17
1.1.2 Java	18
1.1.3 Energy Scitech - Background	19
1.2 Outline of contents	21
1.3 Thesis guide	22
2 Decision support	23
2.1 The commodity trader	23
2.2 Information	24
2.2.1 Technical analysis	25
2.2.2 Fundamental analysis	30

2.3	Decision support	31
2.4	Benefits	31
2.5	Difficulties	32
2.5.1	Capturing risk attitudes and prior expertise	32
2.5.2	Garbage in garbage out	33
2.5.3	Interpretation of output	33
2.5.4	Distrust	34
2.6	Existing decision support tools	34
3	Background	36
3.1	Trading to maximize utility	36
3.1.1	Notation and definitions	36
3.1.2	Constraining trading to limit losses	37
3.1.3	Further constraints	39
3.1.4	Maximum position	40
3.1.5	Requirement for a position on day n	40
3.1.6	Maximizing utility	41
3.1.7	Decision rule	43
3.1.8	Alternative maximization	45
3.2	The Bayesian methodology	45
3.3	Bayes linear methodology	47
3.4	The random walk model	47
3.5	Bayes linear updating of price forecasts given expert information	49
3.5.1	Example	51
3.5.2	Evolution of variance	53
3.5.3	Elicitation of variance	55
3.6	Bayes linear updating of price forecasts given actual prices . .	57

4	Utility analysis	60
4.1	Building a utility function	62
4.1.1	Example - Building your utility function	63
4.1.2	The expected utility theorem	65
4.1.3	Properties of utility functions	66
4.2	Standard utility functions	67
4.2.1	The logarithmic utility curve	67
4.2.2	The risk neutral utility curve	68
4.2.3	The quadratic utility curve	69
4.2.4	The exponential utility curve	70
4.3	Piecewise utility	73
4.3.1	Example	76
4.3.2	Problems	80
4.4	A utility function based on a massive loss	81
4.5	Loss aversion	81
4.6	Disallowing risk seeking attitudes	82
4.6.1	Objection to the risk seeking curve	82
4.6.2	Lottery players and gamblers	82
4.6.3	The existence of the insurance industry	84
4.7	Alternative investment selection models	85
4.7.1	Maximizing the geometric mean return	85
4.7.2	Safety first models	86
4.7.3	Value at Risk	86
4.7.4	Prospect theory	88
4.8	A Java tool for utility elicitation	91
4.9	Conclusions and recommendations	95

5	Elicitation of prior information	97
5.1	Introduction	97
5.2	Related work	100
5.2.1	Elicitation protocols	100
5.2.2	Existing software tools	101
5.2.3	BEEP	101
5.3	Cognitive biases	101
5.3.1	Individual bias	102
5.3.2	Group bias	105
5.4	The elicitation process	106
5.4.1	Obtaining prior beliefs	106
5.4.2	Preliminary training of experts	108
5.4.3	The Java prototype tool	108
5.5	Recommendations	119
6	Diagnostics	120
6.1	The size ratio	121
6.2	Gamma-probability	123
6.3	Discrepancy	123
6.4	Detection of systematic error	124
6.5	Comparison of results using different utility functions	126
6.6	Bayes linear variance learning	127
6.7	Example	129
7	Interactions and piloting with traders	137
7.1	User requirements	138
7.1.1	Screen displays	138
7.1.2	Option trading	138

7.1.3	Transaction costs	139
7.1.4	Price feeds	140
7.1.5	Market contract specifics	140
7.2	Trial with a large utility company	141
7.2.1	The prior belief	141
7.2.2	Utility function	143
7.2.3	Other parameters	143
7.2.4	Analysis of diagnostics	144
7.2.5	Comparison of trading decisions	148
7.3	Trial with a sugar broker	152
7.3.1	The prior belief	152
7.3.2	Utility function	153
7.3.3	Other parameters	154
7.3.4	Analysis of diagnostics	154
7.3.5	Comparison of trading decisions	158

8 Sensitivity analysis of the algorithms

	An application of design of experiments	161
8.1	Design of the experiment	161
8.1.1	Pseudo random walk	165
8.1.2	Interpretation of the Results	167
8.2	Results	167
8.2.1	Initial exploration: $P_t = p_{t-1} + Z_t$	167
8.2.2	Initial exploration: $P_t = p_{t-1} + 0.1 + Z_t$	176
8.2.3	Initial exploration: $P_t = p_{t-1} - 0.1 + Z_t$	183
8.2.4	Analysis of the experiment	183
8.2.5	Tables of results	189
8.3	Recommendations	190

9	Testing software	194
9.1	Code specification and planning tests	194
9.2	Limits of testing	195
9.3	The human machine interface	196
9.3.1	Handling user inputs	196
9.3.2	Error messages	197
9.3.3	Communication style	197
9.3.4	Work flow of the application	198
9.4	Testing internal processes	198
9.4.1	Java error handling	198
9.4.2	Interfacing between R and Java	201
9.5	Testing the installation procedure	203
10	Conclusion	205

Nomenclature

Bear Market	Market Prices are Declining
Bull Market	Market Prices are Rising
CRBI	Commodity Research Bureau Index - A benchmark index that is designed to indicate the broad directional trend of commodity prices
IDE	Integrated Development Environment - a tool to make it easier for programmers to build software.
KTP	Knowledge Transfer Partnership
OPEC	Organization of Petroleum Exporting Countries - A group of eleven oil producing nations.
Position	This is the number of lots the trader is holding.
Swing	Java Graphical User Interface tools
VaR	Value at Risk - A measure to control risk exposure which is used by traders.

List of Figures

1.1	The JBuilder IDE	20
2.1	Gold prices with moving averages	27
4.1	A Piecewise Utility Function	79
4.2	A value function	90
4.3	Java screen: Utility elicitation	92
4.4	Java screens: Refined utility elicitation	94
5.1	Java screen: Volatility estimate from historical prices	109
5.2	Java screen: Eliciting extreme price limits	110
5.3	Java screen: Doubting extreme price limits	111
5.4	Java screen: Eliciting price at time one - P_1	112
5.5	Java screen: Eliciting price at a end of section time - P_t	112
5.6	Java screen: Selecting a shape of the price curve	114
5.7	Java screen: Selecting the shape of the price curve	114
5.8	Java screen: Showing the selected shape of the price curve	115
5.9	Java screen: Showing the selected shape of a more complex price curve	115
5.10	Java screen: Ruler - Level of confidence about final price	116
5.11	Java screen: A summary of interval estimates assuming the Normal distribution	117

5.12	Java screen: Summary of the price outlook	118
6.1	The price belief with actual prices	129
6.2	The size ratio and Gamma-probability diagnostics	130
6.3	The Discrepancy diagnostic	131
6.4	Detection of systematic error	132
6.5	Variance learning	133
6.6	Sensitivity to the utility function	134
7.1	The initial price belief with actual prices	142
7.2	The updated price belief with actual prices	143
7.3	Size ratio and Gamma probability	146
7.4	Discrepancy	147
7.5	Systematic error	147
7.6	Variance learning	148
7.7	Comparison of trading decisions	150
7.8	Comparison of trading profits	151
7.9	The Price Belief with Actual Prices	153
7.10	Size ratio and Gamma probability	155
7.11	Discrepancy	156
7.12	Detection of systematic error	157
7.13	Variance learning	158
8.1	Plot of utility versus profit	168
8.2	Plot of the Dynamic Price Limits Versus Profit	169
8.3	Plot of the Volatility - σ^2 Versus Profit	171
8.4	Plot of the confidence versus profit	172
8.5	Plot of the prior trend versus profit	173
8.6	Plot of the dynamic price limits and volatility versus profit	174

8.7	Plot of the dynamic price limits and confidence versus profit . . .	175
8.8	Plot of the dynamic price limits and utility versus profit . . .	176
8.9	Plot of the utility function versus profit	177
8.10	Plot of the dynamic price limits versus profit	177
8.11	Plot of the prior volatility versus profit	178
8.12	Plot of the confidence about final price versus profit	179
8.13	Plot of the prior trend versus profit	180
8.14	Plot of the prior trend and volatility versus profit	181
8.15	Plot of the prior trend and utility versus profit	182
8.16	Plot of the prior trend and the dynamic price limits versus profit	183
8.17	Plot of the prior trend and the confidence about the final price versus profit	184
8.18	The interaction between the prior and the utility function . . .	187
8.19	The interaction between the prior and the volatility estimate . .	188
8.20	Normal QQ Plot: Residual analysis of the realizable profit model	189

List of Tables

4.1	Loss/Gain Table	85
6.1	CSCE coffee example:diagnostic results	136
7.1	Power trial: TradeABLE's results	149
7.2	Power trial: TradeABLE's diagnostic results	149
7.3	Power trial: Comparing an actual trader's results with Trade- ABLE's results	151
7.4	Sugar trial: TradeABLE's results	159
7.5	Sugar trial: TradeABLE's diagnostic results	160
8.1	A table illustrating the combinations of factors used in the design	162
8.2	Table of risk factors	164
8.3	A table illustrating the levels of factors used in the design . .	164
8.4	Table of summary results : Utility and distribution of profit .	168
8.5	Table of summary results : Dynamic price limits and distri- bution of profit	169
8.6	Table of summary results : Volatility - σ^2 - and distribution of profit	170
8.7	Table of summary results : Confidence about final price and distribution of profit	171

8.8	Table of summary results : Prior trend and distribution of profit	173
8.9	Denotation for the ANOVA model	185
8.10	Summary of ANOVA output for model where $P_t = p_{t-1} + 0.1 + Z_t$	186
8.11	Results for the simulated price with $P_t = p_{t-1} + Z_t$	191
8.12	Results for the simulated price with $P_t = p_{t-1} + 0.1 + Z_t$. . .	192
8.13	Results for the simulated price with $P_t = p_{t-1} - 0.1 + Z_t$. . .	193

Chapter 1

Introduction

1.1 KTP project

1.1.1 The project objectives

The thesis, which follows, is the result of the Knowledge Transfer Partnership between Energy Scitech Ltd and the University of Durham. The KTP project's main objective is to develop a software tool (called TradeABLE) which will help commodity traders and managers to understand and manage the uncertainty in commodity trading decisions and to improve their trading performance . The major tasks in the project (as stated in the KTP Partnership Proposal and Grant Application Form) are to:

1. Understand the trader's role and decision making process.
2. Translate into C++ statistical algorithms, based on linear Bayes approximations, to optimise the expected utility, to the trader/broker, of realisable trading profit.
3. Develop a software product based on these algorithms for commodity traders.

4. Provide a user friendly set of diagnostics that help the trader improve the use of the product.
5. Create a set of demonstrations using real commodity data to use in sales campaigns.
6. Produce technical documentation, user manuals and customer support mechanisms.
7. Develop company staff in product knowledge and expertise in the techniques and customer support.

The statistical algorithms will remain in a tool called R as this is the primary language used by David Wooff who develops the mathematical routines. Also, third party R libraries are used by the algorithms and these are not trivial to replicate.

1.1.2 Java

The software is written in Java and this software interfaces with R when the algorithm is required. This is not ideal. Nevertheless, it speeds up the process of delivering a prototype to show the concept of the decision support tool to potential customers. The early release versions will also have a Java to R interface and if the product proves to be a success, then coding the entire application in Java will be a worthwhile activity.

Java is chosen since it is a portable language. This means that an application that is written in Java can run on Linux, Unix, Apple, and Windows Operating Systems. Many existing trading packages are written in a language called Delphi or in C++. Today, C++ is the most popular language for developing professional software applications to be run on all platforms

(www.cplusplus.com). Nevertheless, C++ programs are not totally independent of the platform they are run on.

The disadvantage in using Java is that the user interface is awkward to finesse. For example, existing Java charting softwares' default plots often look plain and ugly, and hence it is often necessary to write extra code to make the screens aesthetically appealing. Also, it is often necessary to spend time manipulating Swing components, like JTable, to make them work as the user expects.

To develop software a programmer works inside an IDE (Integrated Development Environment). This is necessary to make basic programming tasks (such as debugging, documentation, and building user interfaces) efficient. To create TradeABLE an IDE called JBuilder (shown in Figure 1.1) is used.

1.1.3 Energy Scitech - Background

Energy Scitech Ltd started up in May 1999. The business was founded by the current directors who previously worked together in their oil and gas industry related careers.

Energy Scitech's flagship product is a successful software package called EnABLETM. This application was also developed in partnership with the University of Durham. EnABLETM is written in Java and Perl and it interfaces with the R language and can link to several different oil reservoir simulators.

EnABLETM helps engineers to understand oil reservoirs and to manage the uncertainty in oil well production levels. By applying statistical techniques, the software will aid decision makers with choosing the best wells to develop. This is done by predicting the production rate of oil wells over time and mapping this against the capital cost and the projected price of oil.

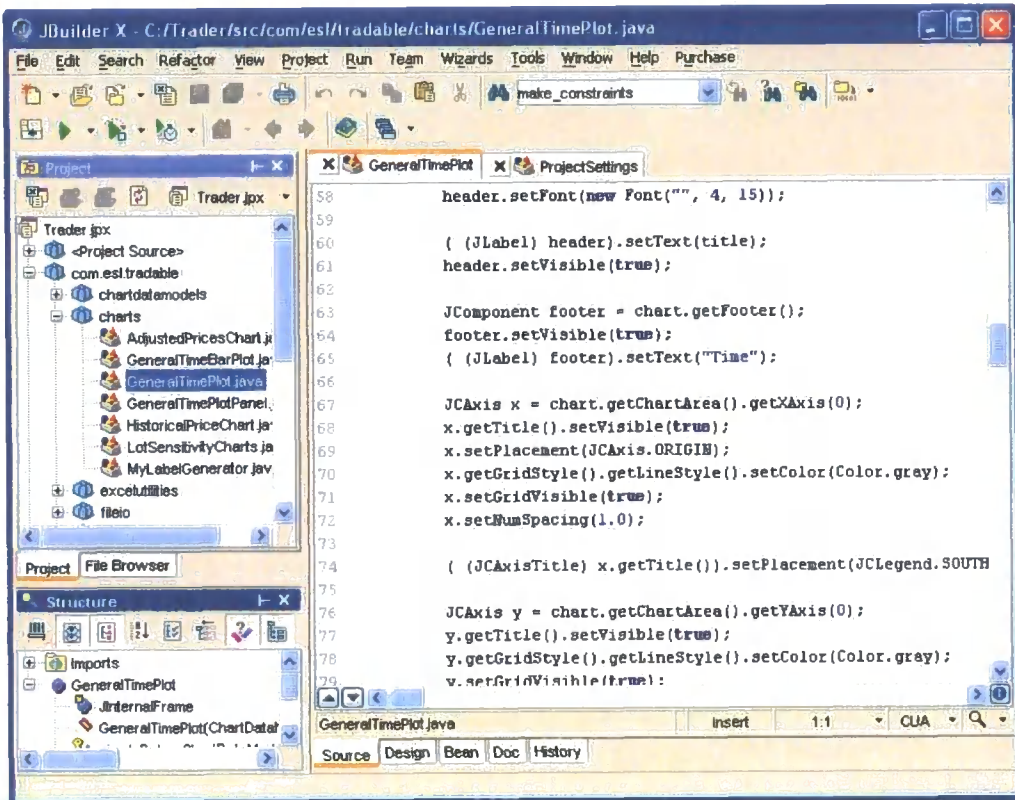


Figure 1.1: The JBuilder IDE

This helps the user to pinpoint the reservoirs that need to be developed to maximize profitability.

Energy Scitech is an innovative and proactive company. The company recently saw a market opportunity in the finance markets and acted on it by beginning the developing of TradeABLE. Essentially, TradeABLE is similar to EnABLETM since it manages uncertainty and optimizes profits. However, the new software is developed from scratch to meet the financial market sector's requirements. This thesis tells the story of the first steps taken to build TradeABLE.

1.2 Outline of contents

What is a commodity trader? What decisions do traders make? What information does the trader base his decisions on? What are the benefits and difficulties associated with providing decision support? In chapter 2, the answers to these questions are discussed in detail.

In chapter 3, the background mathematics behind the decision support tool are discussed. In particular, the maximization of expected utility of profit under an exponential utility function and the Bayesian methodology are explained. Most of this chapter is based on research by Michael Goldstein, Allan Seheult, and David Wooff from the University of Durham Mathematical Science department.

Chapter 4 explores utility analysis. Utility is used to support decisions made under uncertainty, and hence it is used as a risk management method. This part of the thesis looks at the history of utility, how a utility function can be built, and describes some Java utility screens that were designed to elicit a utility function. Alternative methods to utility for managing risk are also evaluated here.

Prior elicitation is dealt with in Chapter 5. How can precise price forecasts be obtained from the expert trader? What sort of mistakes are commonly made by attempting to obtain a person's probability assessment of an event happening? Java screens that were developed to obtain a price outlook from a trader are also shown here.

How are price forecasts monitored? When the price outlook turns out to be inaccurate it is essential to stop supporting decisions based on this imprecise information. Several diagnostics have been developed to flag the trader if the forecasting model does not correspond well with actual incoming prices. These diagnostics which were developed, or recommended, by David

Wooff are explained in chapter 6.

Chapter 7 looks at the interactions and trials of the tool carried out with commodity traders. It is necessary to obtain feedback and requirements for the tool from real potential customers. It is also essential to obtain data on how the decision support tool performs in a real life trading scenario. Trials with two major trading companies are investigated here.

An application of the Design and Analysis of Experiments is explained in chapter 8. Some formal statistical tests were carried out to gain a deeper understanding of the algorithm and how it responds to varying the input parameters.

Chapter 9 discusses the testing of software. The end result of the first phase of this project is a market ready professional software package. A critical step towards this goal is the software testing cycle. Some details of what needs to be considered in the testing process are discussed here.

Chapter 10 concludes the thesis including a critical evaluation of the solutions for trading decision making proposed in this project. What has been done? What has been learnt? What recommendations can be made for future research and development?

1.3 Thesis guide

- Random variables are denoted with capital letters and observed data are denoted with small letters.
- Code - written either in R or Java - is type set in this typeface.
- The definition of technical terms such as trading jargon can be found in the nomenclature.

Chapter 2

Decision support

Decision support is assisting with a verdict or a choice between alternative options. To support a decision, information and advice relevant to the decision is provided. Decision support is not the same as decision making. The commodity trader is a decision maker, since he chooses between alternatives using his judgement, given all of the information he has available. A person or tool that provides decision support is merely providing help.

The following chapter discusses the decisions that commodity traders make, and outlines why decision support for traders is important and useful, but also difficult to provide.

2.1 The commodity trader

A commodity trader is a person who buys and sells goods, or contracts for goods. There are many different goods they can trade - ranging from sterling to sugar, cocoa to copper, oil to oats, and T-Bonds to tin. Some traders are selling on behalf of farmers, and therefore aiming to sell goods at the highest price possible. Some traders are buying on behalf of corporations, such as Coca-Cola and Nescafé, and therefore aiming to buy at a low price. Others have no intention of selling or receiving goods and are only trading to take

advantage of profits that can be made with the fluctuation of price.

2.2 Information

A trader is provided with a sea of information. A trader typically has four computer screens at his desk, showing real-time information about prices - opening price, highs, lows, the last three trades, and the price change compared to yesterday. There may be prices for each of the contract months for any specialized commodity as well as the benchmark commodities which are gold, silver, and crude oil. Also on display is the Commodity Research Bureau Index, key exchange rates, information about volumes traded, and market closing times. They will also have access to technical analysis - which include point-and-figure graphs, candlestick charts, and chart patterns. Fundamentals are on show; a display of incoming news headlines which may have up to fifty stories - with a new headline coming in every minute. More important news, for example, American president assassinated, will be displayed as an alert.

Apart from the continuous flow of information directly in front of a trader, government reports and statistics, financial newspapers, internal and external research departments, weather reports, company end of year reports, factory outputs, word of mouth, and much more, are all possible places to base expertise about future price movement.

The trader needs to assess the value of information in predicting the development of prices in the futures market. For example, recent news items are more relevant than older news, and data relating to weather forecasts is more useful to sugar traders than to metal traders.

2.2.1 Technical analysis

Technicians analyze statistics generated from market price data as a foundation for their forecasting. There are four areas of technical analysis [15]:

1. Patterns on Price Charts
2. Trend-Following Methods
3. Character of Market Analysis
4. Structural Theories

The *weak-form efficient market hypothesis* states that all information available from past prices is reflected in today's price [4]. If the weak-form efficient market hypothesis is assumed to be true, it is argued that the information from the technician's analysis is already accounted for in the price because the technicians methods are widely available in the public domain and facilitated by specialist software providers.

Patterns on price charts

Patterns on price charts is an early and widespread technique used to analyze market prices. It includes creating bar charts that show lines displaying the open, close, and the high and low for each day in a time span of interest [15]. When looking at these lines for a number of consecutive days traders check for patterns. There is jargon terminology for types of pattern - some of these are 'head and shoulders', 'double top', 'selling climax' and 'triangle' [15]. One set of patterns are classed as continuation patterns, indicating price is likely to continue along the same trend [15]. The other set are reversal patterns - signifying a potential change in trend [15].

Point-and-Figure charts focus on the direction of price change, ignoring the volume of trading and the time component of the price series [15]. They are used in a similar way to bar charts - to assess the likeliness of a potential price reversal or continuation of trend.

Trend following methods

Many technicians believe that: “A price once established is more likely to continue than reverse” [15]. This implies that it is better to sell when there is an established downward trend, and buy when an upward trend is recognized.

Traders consider *moving averages* to find trends. The simplest n day moving average adds up the most recent n prices and divides this sum by n :

- At time n , the n day moving average is $\frac{\sum_{i=1}^n p_i}{n}$
- At time $n + 1$, the n day moving average is $\frac{\sum_{i=2}^{n+1} p_i}{n}$
- At time $n + 2$, the n day moving average is $\frac{\sum_{i=3}^{n+2} p_i}{n}$

The moving average smooths out local variation, and lags behind the prices. A 40 day moving average smooths out more variation than a 20 day moving average. Traders have various subjective conditions under which they expect different future price developments. For example, a trader might look for instances where the 10 day moving average crosses actual price to aid his trading since this is where he believes the start of a trend is beginning to form. A 10, 20, and 40 day moving average is shown for Gold prices in Figure 2.1.

Character of market analysis

Oscillators help the exploration of cyclical series. An oscillator measures the fluctuation between two things - in this case a high and a low price.

CMX Gold Closing Prices with Moving Averages

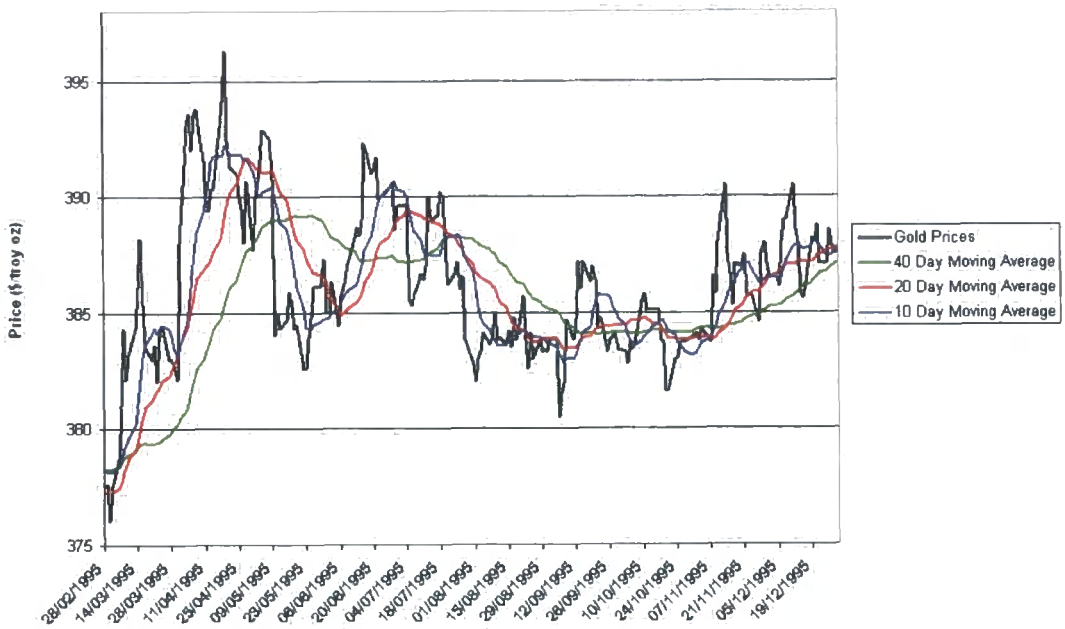


Figure 2.1: Gold prices with moving averages

A popular stochastic oscillator (called %K) was developed by Dr George Lane and is defined as follows [15]:

$$\%K = \frac{C_i - L_n}{H_n - L_n}$$

where

C_i = closing price in current time period

L_n = lowest low during the n time periods

H_n = highest high during the n time periods

i = specific time period

n = number of periods

%K lies between 0% and 100%. When the current price is equal to the lowest low %K = 0%, and when the current price is equal to the highest high %K = 100%.

%K is a useful indicator of when a market is about to reverse trends. If %K < 30%, then this is an indication that trend is about to go upwards. If %K > 70%, then trend is ready to go downwards. It is useful to use these sorts of tools to confirm market beliefs. For example, if a trader is confident that a bull trend will continue, he may change his mind if the stochastic oscillator is >> 70%.

Traders will assess the reliability of the oscillator given all other tools available. %K is a good monitoring tool for traders who believe that price is fluctuating within a range. However, approaching or passing the highest high is not a fool proof signal that price is more likely to dip than rise.

Volume is the total number of contracts traded over a defined time span, and "*open interest* is the total purchase or sales commitment outstanding. [15]" Technicians will use this data to make inferences about future market

behaviour by following general rules developed by the profession. For example, there is a belief that 'volume expands sharply as bottoms and tops are reached. [15]' Nevertheless, there is no empirical evidence suggesting that analyzing volume and open interest is beneficial to traders.

Contrary opinion is going against the view of the crowd if the mass belief is unsubstantiated. It is likely that there is an abnormality in the market when there are weak reasons for large groups to have strong beliefs about price development [15]. This has been shown in historical events such as the 'dot com bubble', and Holland's tulip bulb trading hysteria in 1634 to 1637 [15].

Structural theories

If a trader believes that patterns in historical price will repeat themselves at some time in the future he will be interested in structural theories.

A belief that there is a seasonal price cycle - for example, a price decrease at harvest time each year - is a structural theory. The thoughts underlying this process are derived from the way a commodity is produced and distributed. A supply greater than demand equates to reduced prices, and supply less than demand equates to inflated prices. This is a fundamental principle of economics - prices will vary to keep the market in equilibrium.

Another structural method is the *Elliot Wave Theory* [15] which is based on the Fibonacci series 0,1,1,2,3,5,8,13, This method uses patterns of waves to roughly forecast price development. There is confusion in how to apply the approach - it requires subjective judgement to define the amplitude and frequency of the waves.

2.2.2 Fundamental analysis

A fundamentalist looks at the impact of the supply and demand for a commodity upon its price. They have a lot of data to consider - the majority of commodities are produced in many countries and the economy of each individual country around the world will have an effect in the demand for a single good. Also, consider that a commodity has the potential to be transformed into hundreds of different consumer items.

For example, crude oil is used to make petrol, plastics, jet fuel, waxes, lubricant oil, and so on. The demand for these products is hard to estimate and depends on many factors including weather, government policy, and the time of year. Oil is produced in a wide spread of global locations - and the reliability of supply lines is difficult to assess. Pipelines, refineries, and wells are constantly at threat from terrorist activity or are a likely target for attack for nations at war. OPEC (Organization of Petroleum Exporting Countries), which is a group of eleven oil producing nations, attempts to control oil prices by agreeing how much oil they will provide to the world. Supply is also a function of refinery capacity and these production plants require effective management to work efficiently. Refineries are also at risk to explosion owing to the flammable nature of oil - for example, there was a large explosion at the Hemel Hempstead oil refinery in December 2005. A trader is unlikely to predict this sort of event and include it in their forecast model accurately. If they do know an event such as an explosion will happen, it is still very tough to say precisely how market prices will react.

The problem that the fundamentalist faces is that there is a vast amount of information to collate to arrive at a price belief. Additionally, under the efficient market hypothesis, the investigated data may already be reflected in the price. Nevertheless, the study of the cycle of supply and demand,

and economics, will help the trader understand more about the behaviour of markets.

2.3 Decision support

After collecting information to base expertise about future price movement, ultimately, the commodity trader's aim is buy low and sell high. The trader's quandary is getting the timing just right - buying at dips, and selling at peaks - and deciding how much capital he will risk in the process. The more risk he accepts, the greater the potential rewards, but the higher the chance of ruin.

2.4 Benefits

Among the benefits of making use of decision support is the audit trail for explaining why particular decisions were made - good or bad. This is a recommendation for use of tool by a consultancy firm called Contango Markets. If a trader makes his price beliefs explicit at all times it will be easy for him to communicate to stakeholders his thought processes at the time he made a decision.

The trader will also have assistance with understanding the volatility and uncertainties faced in commodity markets. The decision maker will be given advice on how to trade in these uncertain markets given his, or the trading company's, risk attitudes. Hence, the trader will become a more effective risk manager.

The tool pieces expert information and risk behaviour together to make decisions that lead to optimal profits. This tool will boost the trader's confidence to go ahead with buying or selling - and hence they will be able to

seize more opportunities to make money in the market.

Since the expert's beliefs are forced to be explicit, the support package can help the trader to improve forecasting of prices. For example, the trader may be consistently overestimating the level of prices, or underestimating the volatility of the market. The diagnostics will demonstrate that they are making these types of forecasting errors, and hence help them to adjust expertise to make them more accurate - and more precise beliefs will lead to increased profits.

2.5 Difficulties

2.5.1 Capturing risk attitudes and prior expertise

The approach that will be used to define a trader's risk behaviour is the concept of a utility function. Creating a coherent, and correct, utility function for a commodity trader, or any other person, is very challenging.

Traders will require some statistical understanding to manage risk using utility. At the moment, they are concerned mainly with VaR (Value at Risk) - which is the maximum amount they stand to lose given their holding position and the volatility of the market. They also need to consider the time it may take to unwind their position - that is, get rid of all the stock they hold, or buy stock to meet a future obligation to deliver. The definition and creation of the utility function is discussed in a later chapter.

Attempting to elicit beliefs via computer software is arduous. It is difficult to make sure that all of the experts' information is implicit in their prior specifications. For example, it has been observed¹ that traders have a feeling about price levels where the direction of trend will turn, but can not specify exactly what point in time this turning price will be reached. This means

¹Observation made by Klaas De Boer

that they may not be able to map their beliefs onto a time series chart with conviction about strict timing of peaks and troughs.

New information that affects beliefs will become available frequently. The trader will not have enough time, discipline, or motivation to insert all of this information into a computer system. The commodity trader will also act to buy or sell with gut instincts - a computer will not do this - and this could arguably be a disadvantage, or an advantage.

The process of eliciting prior information is explored more thoroughly in a later chapter.

2.5.2 Garbage in garbage out

The output from a process is a consequence of the input. The manner in which the inputs are retrieved from the user will have a massive impact on the success of any decision support software tool. If the inputs to the process are garbage, then the outputs from the process will also be garbage.

2.5.3 Interpretation of output

Output that is presented to the decision maker needs to be sensible, documented, and easy to understand.

Diagnostics will be presented to the traders which are based on deep statistical theory. It will be important to explain and document these diagnostics in layman's terms so that they are useful and meaningful. These diagnostics are designed to flag a trader if the forecasting model does not represent what is actually happening in the market. If price is consistently overestimated or underestimated, this will be shown, and hence the trader can learn how to forecast price with more precision in the future. The diagnostics are discussed fully in a later chapter.

2.5.4 Distrust

The trader is likely to distrust any new software product that claims it will increase profit. The market is saturated with packages that profess they will improve trading performance. Also, if a trader does not know how the black box creates its decision advice it will be unlikely to inspire confidence. The mathematics behind the tool will need to be articulated without disclosing intellectual property.

2.6 Existing decision support tools

There are many existing software tools for commodity traders on the market. Most of these software packages concern themselves with implementing technical and fundamental analysis - the two most widespread methods of trade selection. There are also several packages on the market that specialize in day trading - trading frequently throughout the day and liquidating any position before the market closes.

Among the most well known trading software are:

1. Metastock : Technical Analysis Software (www.metastock.ca).
2. OmniTrader : Technical Analysis Software (www.omnitrader.com).
3. k2TM and k3000TM from KWI: Risk Management and Asset Optimisation. Company bought by Global Energy Decisions (www.globalenergy.com).
4. Wave59: Technical Analysis allowing for a wide range of methodologies in trade selection (www.wave59.com).
5. Bloomberg (www.bloomberg.com) and Reuters (www.reuters.com): International news and data providers.

There does not appear to be any existing trading packages that use Bayes. Nevertheless, there are many educational reports and articles about the link between Bayesian methodology and commodity trading. There are several packages that use Bayesian methodology for non-trading decision making applications.

Successful Bayesian decision support applications include:

1. EnABLETM
2. BayesEngineTM Technology (www.bayes.com).

Chapter 3

Background

This chapter is based on the research work carried out by David Wooff, Michael Goldstein, and Allan Seheult. In particular, all of the mathematics discussed here stem from papers [9], [12], and [29].

3.1 Trading to maximize utility

3.1.1 Notation and definitions

Let p_t denote the closing price on day t of a futures contract. The trader will buy, sell, or hold on this contract for n consecutive days such that $t = 1, \dots, n$.

Let l_t denote a trading transaction on day t and let l_t be a positive number for lots bought, and let l_t be a negative number for lots sold. For example, if 20 lots are bought on day 5, $l_5 = 20$, and if 10 lots are sold on day 6, $l_6 = -10$. It is possible to sell lots without owning them.

The total stock held on any day is known as the position. It is possible for the position to be negative. Let the position at time t be denoted Pos_t , then:

$$Pos_t = l_1 + l_2 + \dots + l_t$$

The cumulative balance is the amount cash on the trader's book. If lots are bought there is a cash outflow and if lots are sold there is a cash inflow.

It is assumed that lots are bought and sold at the current market price with no transaction cost ¹. Let the cumulative balance at time t be denoted Bal_t , then:

$$Bal_t = -(l_1 p_1 + \dots + l_t p_t)$$

Realizable profit is the amount of money that will be made or lost if the trader cancelled out his position at the current price. The realizable profit profit on day $t + 1$, denoted RP_{t+1} , can be expressed (for convenience) as:

$$RP_{t+1} = Balance_{t+1} - l_t p_t + (Position_{t-1} + l_t) p_{t+1}$$

The objective is to choose l_t to maximize $E[RP_{t+1}]$.

3.1.2 Constraining trading to limit losses

Traders find it attractive to be able to specify a maximum acceptable loss since this is an easy to understand method of controlling risk.

Suppose that the trader states that he cannot lose more than K (for example, $K = \$250,000$ or $\pounds 100,000$). K will be in the same currency units as the listed prices for the contract traded. This means that it is required for $RP_{t+1} \geq -K$ for all p_{t+1} . That is:

$$RP_{t+1} = Bal_{t-1} - l_t p_t + (Pos_{t-1} + l_t) p_{t+1} \geq -K \quad (3.1)$$

p_{t+1} is not known as the buy or sell decision is made at time t . To protect against losing more than K an assessment about the extreme lower and upper limits for p_{t+1} , denoted p_L and p_U , is required at time t .

Setting p_{t+1} to p_L in 3.1,

$$Bal_{t-1} - l_t p_t + (Pos_{t-1} + l_t) p_L \geq -K$$

¹There is a transaction cost in the form of a bid-ask spread in practice.

Setting p_{t+1} to p_U in 3.1,

$$Bal_{t-1} - l_t p_t + (Pos_{t-1} + l_t) p_U \geq -K$$

With some simple rearrangement of the above equations, the following constraint is imposed on l_t to constrain for a maximum acceptable loss K .

$$\frac{K + Bal_{t-1} + Pos_{t-1} p_U}{p_t - p_U} \leq l_t \leq \frac{K + Bal_{t-1} + Pos_{t-1} p_L}{p_t - p_L} \quad (3.2)$$

If p_t is inside $[p_L, p_U]$, but is very close to either p_L or p_U , (3.2) could become unstable since a small change in p_t will have a big impact on the range of l_t allowed. This will seem strange and incoherent to a trader.

If $p_t \geq p_U$ or $p_t \leq p_L$ then the use of 3.2 is not sensible. Nevertheless, since p_t is known, and p_L and p_U are estimates, it is simple to force that p_L and p_U are on either side of p_t and at a reasonable distance. This feels sensible in practice and the distance will depend on market volatility.

Also, if p_{t+1} is observed and lies outside $[p_L, p_U]$ then the maximum loss barrier may have been breached. The user of the tool needs to be aware of this. Nevertheless, if the $[p_L, p_U]$ range is set too wide, the trading decisions will be too conservative, and hence opportunities for additional profits will be lost.

In the tool, it is difficult to capture and estimate for p_L and p_U for p_{t+1} at every time t . However, this could be functionality for a future version. Initially, extreme price limits were elicited for the entire period. That is, lower and upper prices that the trader was sure would not be passed on any of the n trading days. This turned out to be too conservative.

At the moment, the price limits for p_{t+1} are:

$$p_L = p_t - \psi\sigma$$

$$p_U = p_t + \psi\sigma$$

where σ is a volatility assessed from the market and ψ ranges from 1 to 10. 1 represents risky price limits which have a relatively high chance of being passed, and hence the maximum loss may exceed K . 10 represents very conservative price limits where there is a negligible chance that p_{t+1} will lie outside $[p_L, p_U]$. However, low risk trading to make sure K is upheld will see opportunities lost to make more money if the forecast model is accurate.

These constraints will be a good thing for a risk manager to impose on a risk neutral ² (or nearly risk neutral) trader as they will typically continue to buy or sell until they reach this maximum risk exposure in real life. This tool will also trade until constrained in this manner with a risk neutral function.

The objective is to choose l_t to maximize $E[RP_{t+1}]$ and to never lose more than a maximum acceptable loss.

3.1.3 Further constraints

Daily cap on lots

There is a constraint on the amount of lots that can be bought or sold in one day owing to:

- Market liquidity
- Capital constraints
- Restrictions imposed by the market exchange

Such that:

$$-a \leq l_t \leq b$$

where a is the maximum lots that can be sold in a day and b is the maximum lots that can be bought in a day. a and b are specified by the trader.

²Risk neutral is a type of utility function which is explained in a later chapter

3.1.4 Maximum position

The trader may wish to specify a maximum position that can be held on any one day such that:

$$-maxa \leq Pos_t \leq maxb$$

where $-maxa$ is a maximum short position specified by the trader, and $maxb$ is a maximum long position specified by the trader:

$$-maxa - Position_{t-1} \leq l_t \leq maxb - Position_{t-1}$$

Reasons why this constraint will be imposed include:

1. Market liquidity
2. Control on risk exposure
3. Position limits imposed by the market exchange

3.1.5 Requirement for a position on day n

It is common for a commodity trader to trade in contracts to deliver physical goods. They might need to end up with a non zero position on day n . In the tool, the trader can specify that:

$$Pos_n = closingstock$$

where $closingstock$ is the amount of lots the trader wants to hold on day n . This can be positive or negative amount.

It will also be common for a trader to start the maximization on day 0 with a non zero position, balance, and realizable profit level.

3.1.6 Maximizing utility

The objective is to choose a single lot size at time t , l_t , to maximize the expected utility of realizable profit one day ahead, $E[U(RP_{t+1}) | p_1, \dots, p_t]$, under the condition that an amount exceeding K is never lost given that K is specified by the trader.

The trader's utility for profit can take different forms (as described in a later chapter). The tool models the trader's utility function, $U(X)$, as:

$$U(X) = 1 - \exp\left(\frac{-X}{R}\right)$$

where X is RP_{t+1} ranging between $[-\infty, \infty]$, and R is called the risk factor. $R > 0$ and the utility function becomes more risk accepting as $R \rightarrow \infty$.

It is approximated that $X \sim N(\mu_X, \sigma_X^2)$, and hence the expected utility of X is:

$$\begin{aligned} E[U(X)] &= E\left[1 - \exp\left(\frac{-X}{R}\right)\right] \\ &= 1 - E\left[\exp\left(\frac{-X}{R}\right)\right] \\ &= 1 - M_X\left(\frac{-1}{R}\right) \\ &= 1 - \exp\left(\frac{-\mu_X}{R} + \frac{\sigma_X^2}{2R^2}\right) \end{aligned}$$

$M_X\left(\frac{-1}{R}\right)$ is the moment generating function of X . Since $X \sim N(\mu_X, \sigma_X^2)$, X has the moment generating function:

$$M_X(t) = E[\exp(tX)] = \exp\left(t\mu_X + \frac{1}{2}t^2\sigma_X^2\right)$$

This is shown by using the standard formula:

$$E[X] = \int_{-\infty}^{\infty} X.f(X)dX$$

$$\begin{aligned}
E[\exp(\frac{-X}{R})] &= \frac{1}{\sigma_X \sqrt{2\pi}} \int_{-\infty}^{\infty} \exp(\frac{-X}{R}) \exp(\frac{1}{2}(\frac{X - \mu_X}{\sigma_X})^2) dX \\
&= \exp(\frac{-\mu_X}{R} + \frac{\sigma_X^2}{2R^2})
\end{aligned}$$

where,

$$\begin{aligned}
\mu_X &= E[RP_{t+1}|p_1, \dots, p_t] \\
\sigma_x^2 &= Var[RP_{t+1}|p_1, \dots, p_t]
\end{aligned}$$

In particular, it is straightforward to show that:

$$\mu_X = Bal_{t-1} - l_t p_t + (Pos_{t-1} + l_t)E[P_{t+1}] \quad (3.3)$$

$$\sigma_x^2 = (Pos_{t-1} + l_t)^2 V[P_{t+1}] \quad (3.4)$$

Equation (3.4) is found since Bal_{t-1} , $l_t p_t$, and Pos_{t-1} are known at time t , and hence have a variance of 0.

$E[U(RP_{t+1})]$ is maximized with respect to l_t :

$$\begin{aligned}
&\frac{d}{dl_t} (1 - \exp(\frac{-\mu_X}{R} + \frac{\sigma_x^2}{2R^2})) \\
&= \frac{d}{dl_t} (1 - \exp(\frac{-Bal_{t-1} + l_t p_t - (Pos_{t-1} + l_t)E[P_{t+1}]}{R} + \frac{(Pos_{t-1} + l_t)^2 V[P_{t+1}]}{2R^2}))
\end{aligned}$$

Let

$$\phi = \frac{-Bal_{t-1} + l_t p_t - (Pos_{t-1} + l_t)E[P_{t+1}]}{R} + \frac{(Pos_{t-1} + l_t)^2 V[P_{t+1}]}{2R^2}$$

Since $\frac{d \exp(\phi)}{dl_t} = \frac{d\phi}{dl_t} \exp(\phi)$:

$$\frac{dU(RP_{t+1})}{dl_t} = (-\frac{p_t}{R} + \frac{E[P_{t+1}]}{R} - \frac{l_t V[P_{t+1}]}{R^2} - \frac{Pos_{t-1} V[P_{t+1}]}{R^2}) \exp(\phi)$$

The stationary point is found where $\frac{dU(RP_{t+1})}{dt} = 0$. This equation is solved as follows:

$$\begin{aligned} -\frac{p_t}{R} + \frac{E[P_{t+1}]}{R} - \frac{l_t V[P_{t+1}]}{R^2} - \frac{Pos_{t-1} V[P_{t+1}]}{R^2} &= 0 \\ \frac{l_t V[P_{t+1}]}{R^2} &= \frac{-p_t}{R} + \frac{E[p_{t+1}]}{R} - \frac{Pos_{t-1} Var[P_{t+1}]}{R^2} \\ l_t &= \frac{R(E[P_{t+1}] - p_t)}{Var[P_{t+1}]} - Pos_{t-1} \quad (3.5) \end{aligned}$$

3.1.7 Decision rule

The lot choice on day t , l_t , is chosen to maximize $E[U(RP_{t+1})]$ where the utility function is of the form:

$$U(X) = 1 - \exp\left(\frac{-X}{R}\right)$$

and to constrain trading such that no more than a monetary amount K is lost.

The decision is to choose:

$$l_t = \frac{R(E[P_{t+1}] - p_t)}{Var[P_{t+1}]} - Position_{t-1}$$

subject to:

$$\frac{K + Balance_{t-1} + Position_{t-1} p_U}{p_t - p_U} \leq l_t \leq \frac{K + Balance_{t-1} + Position_{t-1} p_L}{p_t - p_L} \quad (3.6)$$

and any additional constraint imposed by the trader.

The lot decision does not depend on realizable profit, RP_{t+1} , unless the lot choice has been constrained under (3.6). It is likely that attitude toward decision making will change as realizable profit changes, so it is advised that the trader re-tunes R , the risk factor, as his risk attitude changes. The exponential utility function displays a property of constant relative risk aversion.

This means that risk attitude is not dependent on wealth and this is perhaps an undesirable quality of the exponential function. This is discussed in detail in a later chapter.

If the risk factor is large, trading decisions will only be inhibited by the constraints specified by the trader. This reduces the potential in using the utility function to manage risk since the constraints will control the risk exposure. It may be better if the constraints were implicit in the utility function. Nevertheless, the traders and facilitators are inexperienced in defining their utilities, and hence it is important to state explicit constraints such as the maximum loss barrier.

p_t , $E[P_{t+1}]$, and $V[P_{t+1}]$ are in terms of price per lot above. To convert to listed prices the following is used:

$$l_t = \frac{R(E[P_{t+1}(L)] - p_{t(L)})}{Var[P_{t+1}(L)]F} - Position_{t-1}$$

where L denotes that price is in terms of listed prices and F is a conversion factor between price per lot and the listed price.

If the trader is confident in his price belief, there is less risk in buying more lots. For example, if $E[P_{t+1}]$ is much greater P_t relative to $V[P_{t+1}]$ this means that there is a high confidence that price is moving upward one day ahead, and hence more lots will be bought. $V[P_{t+1}]$ must always be greater than zero and this is certainly the case over all commodity markets.

l_t must be an integer lot size. If l_t is positive it is rounded down, and if l_t is negative it is rounded up. The direction of rounding keeps trading within the risk settings. This is a practicality of the market - trades can only be made in whole lots.

3.1.8 Alternative maximization

A slightly more complex maximization routine is applied in the tool. This is based on the same principles as explained in this section - maximizing an exponential utility of realizable profit. However, this method chooses l_t, l_{t+1}, \dots, l_n to maximize RP_n instead of just choosing l_t to maximize RP_{t+1} .

3.2 The Bayesian methodology

The Bayesian methodology begins with a *prior* which is the probability that a person will place on a particular event happening before seeing any data [1]. The approach also uses a *likelihood* function which is the probability density of a random variable X which has a dependence on a parameter θ . A Bayesian statistician combines the prior with the likelihood function to arrive at a *posterior* distribution.

The Bayesian framework allows the inclusion of an expert's *subjective* judgements about the probability of a future event. This has been objected to, by some, since personal opinion is not scientific. However, in areas where events are not repeatable, or little historical data is available, making assessments of future risks is difficult, and probably unreliable, using data only. Here, personal opinion about future uncertainty is required.

It was Thomas Bayes (1702 - 1761), a nonconformist minister, who devised Bayes Theorem. Bayes Theorem is:

$$\pi(\theta|x) = \frac{\pi(\theta)f(x;\theta)}{m(x)}$$

where:

- θ is a uncertain future event
- x is the observed data relevant to θ

- $\pi(\theta)$ is the prior distribution of an event θ .
- $f(x;\theta)$ is the likelihood of data x given parameter θ .
- $m(x)$ is the marginal distribution of x .
- $\pi(\theta|x)$ is the posterior distribution. A combination of $\pi(\theta)$ and $f(x;\theta)$.

Since $m(x)$ is not dependent on θ , the theorem is often written as:

$$\pi(\theta|x) \propto \pi(\theta)f(x;\theta)$$

or

$$\text{posterior} \propto \text{prior} \times \text{likelihood}$$

A trader will collate his information on price analysis with his knowledge and experience on external events that affect the market movement. As the knowledge base changes arguably the trader's forecast will change. This is a rational explanation of why two traders can have conflicting price beliefs. If they met and discussed the information on which they base their forecast, then it makes sense that they should arrive at the same belief after their talks are concluded.

Nevertheless, people interpret the same information differently. Some banks have traders with access to the same research and tools who have different forecasts for price ³. This may be owing to:

1. a stronger weighting to a particular fact or opinion.
2. different personal experiences.
3. missing out different parts of the knowledge base.
4. a different interpretation of the information available.

³Observation by Klaas De Boer

3.3 Bayes linear methodology

The Bayes linear methodology is similar to the Bayesian methodology described in the above section. However, there is one huge difference between the methods. The Bayes linear technique uses the *expected value* rather than using a full probability density function to express uncertainties. This is useful in *complex* problems where there are too many sources of information to be confident about making the full joint probability specification that is required for Bayesian inference [7].

3.4 The random walk model

The model that the algorithm uses assumes that prices follow a random walk, such that:

$$P_t = P_{t-1} + Z_t$$

That is, tomorrow's price is today's price plus an unknown step of Z_t . The expectation of Z_t is zero, such that, $E[Z_t]=0$, and the variance of Z_t is denoted σ_z^2 , such that, $V[Z_t] = \sigma_z^2$.

P_t can also be written as:

$$P_t = p_0 + (P_1 - p_0) + (P_2 - P_1) + \dots + (P_t - P_{t-1})$$

$$P_t = p_0 + \sum_{i=1}^t Z_i$$

This means that, $E[P_t] = p_0$ (since $E[Z_i] = 0$) where p_0 is the last observed price before a forecast is made. Also note that:

$$P_t = p_0 + Z_1 + Z_2 + \dots + Z_t$$

$$V[P_t] = V[p_0] + V[Z_1] + V[Z_2] + \dots + V[Z_t]$$

$$V[P_t] = t\sigma_z^2$$

σ_z^2 is the variance of the daily change in prices. Historical prices, either from the immediately preceding past, or a period in history that is expected to reasonably match volatility in the future trading period that interests us, are used to give an indication of σ_z^2 . If no suitable data exists, the expert will still be able to give an estimate of σ_z^2 .

The historical data does not need to be a specific length to calculate an estimate of σ_z^2 . However, it must be a reasonable sample size no less than ten data points and ideally more than twenty ⁴. If suitable price data of length n is available, then σ_z^2 is obtained via first differences:

$$\sigma_z^2 = \frac{\sum_{i=1}^{n-1} (P_{i+1} - P_i)^2}{2(n-1)} [29]$$

It is explained why the method of first differences is used to derive variance in section 6.4.

Example

Historical Data: 12.1, 12.7, 12.4, 13.1, 13, 12.8, 13.1, 13.1, 12.5, 12.3

$$\begin{aligned} \sigma_z^2 = & \frac{(12.7 - 12.1)^2 + (12.4 - 12.7)^2 + (13.1 - 12.4)^2 + (13 - 13.1)^2}{18} \\ & + \frac{(12.8 - 13)^2 + (13.1 - 12.8)^2 + (13.1 - 13.1)^2 + (12.5 - 13.1)^2}{18} \\ & + \frac{(12.3 - 12.5)^2}{18} = 0.08 \end{aligned}$$

This number means that there is a 70% chance that daily price change will be within ± 0.28 (one standard deviation, $\sigma_z = \sqrt{0.08}$), and there is a 95% chance that daily price change will be within ± 0.56 . The calculation of these

⁴Recommendation by David Wooff

confidence intervals assumes that prices are following a random walk with a normally distributed step with expectation of zero. This normal assumption is not made for the underlying model. Any trend in the historical prices is not taken into consideration in the calculation of the standard deviation (0.28) in this example.

If the expert does not think that the historical data gives a good estimation of volatility in the trading period of interest, he must adjust σ_z^2 to his belief. This is very important since the trader's subjective beliefs alone form the prior expertise. Vital information could be lost if the expert is allowed to think that historical prices will completely take care of the estimation of σ_z^2 . Nevertheless, volatility is difficult for the trader to estimate as the human brain tends to find probability assessments non-intuitive. Thus, an analysis of historical data is a good starting point for estimating future price volatility.

3.5 Bayes linear updating of price forecasts given expert information

Let P be a vector of a price series over the next n days. The model for this price series is the simple random walk explained in an earlier section such that $P_t = P_{t-1} + Z_t$.

Then a target vector X of p linear combinations of the n elements of P is identified (where $p < n$) [12]. X is the focus of the Bayesian analysis as the aim is to reduce uncertainty about its elements [12]. The example that follows considers X as, one element, P_n - the final price in vector P . A second element could be the average of a price series denoted as \bar{P} [12].

$X = LP$ where L is a known pxn matrix since L defines the p linear

combinations of P [12]. For example, when the focus of X is P_n L is $(0\ 0\ \dots\ 0\ 1\ 0\ 1)$ which has dimension $1 \times n$. If a second element \bar{P} was also included in the analysis L would be:

$$L = \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \\ \frac{1}{n} & \frac{1}{n} & \dots & \frac{1}{n} & \frac{1}{n} \end{pmatrix}$$

The expectation and variance of the target vector X is computed as follows [29] [12]:

$$E[X] = LE[P] \quad (3.7)$$

$$V[X] = LVar[P]L^T \quad (3.8)$$

$E[P]$ and $Var[P]$ are required to find these quantities. From the random walk model it can be shown that [29]:

$$E[P_t] = \mu, V[P_t] = t\sigma_z^2, Cov(P_t, P_{t+s}) = t\sigma_z^2$$

In vector/matrix notation this is [29]:

$$E[P] = \mu 1_n$$

$$Var[P] = \sigma^2 \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix}$$

When the expert (or observed data) supplies additional information that reduces the uncertainty about P further, $E[P]$ is revised to $E_w[P]$, and $V[P]$ is revised to $V_w[P]$ where w denotes new information. The Bayes updating formula that need to be applied when new information becomes available are [29] [12]:

$$E_w[P] = E[P] + Var[P]L^TVar[X]^{-1}[E_w[X] - E[X]] \quad (3.9)$$

$$Var_w[P] = Var[P] - Var[P]L^TVar[X]^{-1}[I - Var_w[X]Var[X]^{-1}]LVar[P] \quad (3.10)$$

3.5.1 Example

Suppose that a trader can only specify limited beliefs, namely, expectation and variance for the *final* price P_n . Let these beliefs be denoted as follows:

$$E_w[P_n] = \tilde{\mu}$$

$$V_w[P_n] = \tilde{\sigma}^2$$

The w subscript denotes that there is new information in the expectation or variance (here this new information was elicited from the expert).

L is known to be $(0 \ 0 \ \dots \ 0 \ 1)$ to correspond with $X = P_n$.

The expectation and variance of the target vector, X , are found using (3.7) and (3.8) as follows:

$$E[X] = (0 \ 0 \ \dots \ 0 \ 1) E \begin{pmatrix} P_1 \\ P_2 \\ \vdots \\ P_{n-1} \\ P_n \end{pmatrix} = E[P_n]$$

$$Var[X] = (0 \ 0 \ \dots \ 0 \ 1) \sigma_z^2 \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} = n\sigma_z^2$$

$E_w[P]$ can be found using (3.9) as follows:

$$E_w[P] = \begin{pmatrix} \mu \\ \mu \\ \vdots \\ \mu \\ \mu \end{pmatrix} + \sigma_z^2 \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \frac{1}{n\sigma_z^2} [\tilde{\mu} - \mu]$$

$$E_w[P] = \begin{pmatrix} \mu \\ \mu \\ \vdots \\ \mu \\ \mu \end{pmatrix} + \begin{pmatrix} \frac{1}{n} \\ \frac{2}{n} \\ \vdots \\ \frac{n-1}{n} \\ \frac{n}{n} \end{pmatrix} [\tilde{\mu} - \mu]$$

It follows that:

$$E_w[P_t] = \mu + \frac{t}{n} [\tilde{\mu} - \mu]$$

$Var_w[P]$ can be found using (3.10) as follows:

$$Var_w[P] = \sigma_z^2 \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix} - \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \frac{1}{n\sigma_z^2} \left[1 - \frac{\tilde{\sigma}}{n\sigma_z^2} \right] \begin{pmatrix} \sigma_z^2 & 2\sigma_z^2 & \dots & (n-1)\sigma_z^2 & n\sigma_z^2 \end{pmatrix}$$

For convenience, let:

$$\alpha = \frac{1}{n} \left[1 - \frac{\tilde{\sigma}}{n\sigma_z^2} \right], \quad (3.11)$$

then:

$$Var_w[P] = \sigma_z^2 \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix} - \begin{pmatrix} 1 \\ 2 \\ \vdots \\ n \end{pmatrix} \sigma_z^2 \alpha (1 \ 2 \ \dots \ n)$$

$$Var_w[P] = \sigma_z^2 \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & n \end{pmatrix} - \sigma_z^2 \alpha \begin{pmatrix} 1 & 2 & \dots & n \\ 2(1) & 2(2) & \dots & 2(n) \\ 3(1) & 3(2) & \dots & 3(n) \\ \vdots & \vdots & \vdots & \vdots \\ n(1) & n(2) & \dots & n(n) \end{pmatrix}$$

It follows that:

$$Var_w[P_t] = \sigma_z^2 t - \sigma_z^2 \alpha t^2 = (1 - t\alpha) t \sigma_z^2 \quad (3.12)$$

3.5.2 Evolution of variance

It seems sensible to assume that the expert is more certain about a price $t-1$ days ahead than a price t days ahead, such that:

$$V_w[P_t] \geq V_w[P_{t-1}]$$

and, given this assumption, there is a proof (see below) to show that:

$$\frac{t(t-1)}{2t-1} \sigma_z^2 \leq V_w[P_t] \leq t \sigma_z^2$$

If the expert specifies

$$V_w[P_t] = t \sigma_z^2,$$

there is no more information about P_t than the simple random walk model.

If the expert specifies:

$$V_w[P_t] > t \sigma_z^2,$$

then either:

1. σ_z^2 is too small.
2. The model in use is flawed.

If the expert specifies

$$V_w[P_t] < \frac{t(t-1)}{2t-1}\sigma_z^2,$$

then either:

1. σ_z^2 is too large.
2. The model in use is flawed.

Proof

Assuming that the trader has more or equal certainty about a future price $t-1$ days ahead than a future price t days ahead such that:

$$V_w[P_t] \geq V_w[P_{t-1}]$$

From previous definitions in (3.11) and (3.12) it is known that:

$$\alpha = \frac{1}{n} \left(1 - \frac{\tilde{\sigma}^2}{n\sigma_z^2}\right)$$

$$V_w[P_t] = (1 - t\alpha)t\sigma_z^2$$

Then,

$$(1 - t\alpha)t\sigma_z^2 = (1 - (t-1)\alpha)(t-1)\sigma_z^2$$

$$t\sigma_z^2 - \alpha t^2\sigma_z^2 \geq (1 - t\alpha + \alpha)(t\sigma_z^2 - \sigma_z^2)$$

$$t\sigma_z^2 - \alpha t^2\sigma_z^2 \geq t\sigma_z^2 - \sigma_z^2 - \alpha t^2\sigma_z^2 + \alpha t\sigma_z^2 + \alpha t\sigma_z^2 - \alpha\sigma_z^2$$

$$\sigma_z^2 - 2t\alpha\sigma_z^2 \geq 0$$

$$\sigma^2(1 - \alpha(2t-1)) \geq 0$$

Given that $\sigma_z^2 \geq 0$, then,

$$\begin{aligned}\alpha &\leq \frac{1}{2t-1} \\ \alpha &= \frac{1}{t} \left(1 - \frac{\tilde{\sigma}^2}{t\sigma_z^2}\right) \leq \frac{1}{2t-1} \\ -\frac{\tilde{\sigma}^2}{t\sigma_z^2} &\leq \frac{t}{2t-1} - 1 \\ \tilde{\sigma}^2 &\geq \sigma_z^2 \left(t - \frac{t^2}{2t-1}\right) \\ \tilde{\sigma}^2 &\geq \sigma_z^2 \left(\frac{t(t-1)}{2t-1}\right)\end{aligned}$$

Also, $\alpha \geq 0$,

$$\begin{aligned}\alpha &= \frac{1}{t} \left(1 - \frac{\tilde{\sigma}^2}{t\sigma_z^2}\right) \geq 0 \\ \tilde{\sigma}^2 &\leq t\sigma_z^2\end{aligned}$$

Hence,

$$\sigma_z^2 \left(\frac{t(t-1)}{2t-1}\right) \leq V_w[P_t] \leq t\sigma_z^2$$

3.5.3 Elicitation of variance

In the package the expert is asked for $V_w[P_n]$ only and the associated boundaries are forced upon the expert. To do this, the expert is asked how confident he is about his specification of the final price on a scale of one to ten. One is the least confident the expert can be and this corresponds to $0.9n\sigma_z^2$. Here, the coefficient of 0.9 is added for practical purposes. Ten is the most confident and corresponds to $\frac{n(n-1)}{2n-1}\sigma_z^2$. If an intermediate value is chosen, then $V_w[P_n]$ is scaled between $\frac{n(n-1)}{2n-1}\sigma_z^2$ and $0.9n\sigma_z^2$, and this is done by converting to standard deviations first. The equation for calculating the experts confidence is as follows:

$$SD_w[P_n] = \sigma_z \sqrt{0.9n} - \frac{X-1}{9} (\sigma_z \sqrt{0.9n} - \sigma_z \sqrt{\frac{n(n-1)}{2n-1}})$$

where X is the number on a scale of one to ten chosen by the expert.

This is converted back to the variance, and hence $V_w[P_n]$ is elicited for use in the R functions.

Example

Suppose that the following information is available:

$$\sigma_z = 0.1, X = 8, n = 20, E_w[P_{20}] = 14$$

Then

$$SD_w[P_{20}] = 0.1\sqrt{0.9(20)} - \frac{8-1}{9}(0.1\sqrt{0.9(20)} - 0.1\sqrt{\frac{20(20-1)}{40-1}}) = 0.33$$

$$V_w[P_{20}] = 0.33^2 = 0.109$$

In terms of beliefs this means that expert is 70% confident that price 20 days ahead will be 14 ± 0.33 and 95% confident it will be 14 ± 0.66 . These confidence bands assume that prices follow a random walk with a normally distributed step. Again, these are informal heuristics.

If the trader states that there is much more than a 70% chance that price 20 days ahead will lie between $E_w[P_n] \pm SD_w[P_n]$, then either:

1. σ_z^2 is too large.
2. The model in use is flawed.

If the trader states that there is much less than a 70% chance that price 20 days ahead will lie between $E_w[P_n] \pm SD_w[P_n]$, then either:

1. σ_z^2 is too small.

2. The model in use is flawed.

It is important that the trader sees the interval estimates so that he is aware of what picking a number on the ruler represents in terms of confidence levels. Nevertheless, if the confidence bands are derived using a Normal assumption this could be misleading for the expert since there is no Normal assumptions made to approximate future prices. These confidence bands are only informal guidelines for the trader.

For simplicity, the examples in this section have assumed a simple random walk with a normally distributed step with expectation zero to calculate confidence bands. The algorithm copes with an expert's assertion that prices follow a trend and that price changes do not follow a symmetric normal distribution. The decision rule in Bayes linear analysis is based on the first and second moment (expectation and variance) regardless of the underlying distribution.

3.6 Bayes linear updating of price forecasts given actual prices

An observed price is new information which is used to update the forecasts for future unknown prices.

There are two vectors of random quantities which are denoted B and D . On day t [29]:

B is the vector of future prices P_{t+1}, \dots, P_n

D is the vector of past prices P_1, \dots, P_t

At time t , prices p_1, p_2, \dots, p_t have been observed. The vector of these known prices is denoted d .

d is the vector of observed past prices p_1, p_2, \dots, p_t

At time 0, the expert made a statement about expectations of P_1, \dots, P_n , thus $E[P_1], \dots, E[P_n]$, and the variance and covariances between forecasted prices are known such that [29]:

$$E[B] = (E[P_{t+1}] \ E[P_{t+2}] \ \dots \ E[P_n])$$

$$E[D] = (E[P_1] \ E[P_2] \ \dots \ E[P_t])$$

$$d = (p_1 \ p_2 \ \dots \ p_t)$$

$Var[B]$ is the prior variance matrix for P_{t+1}, \dots, P_n .

$Var[D]$ is the prior variance matrix for P_1, \dots, P_t .

$Cov[B, D]$ is the covariance matrix which measures the strength of the correlation between P_1, \dots, P_t and P_{t+1}, \dots, P_n .

The Bayes linear update mechanism refreshes the expectation for future prices, $E[B]$, at each time $t = 1, \dots, t = n - 1$. (At time n there are no random quantities left to forecast). The adjusted expectation vector for B given that $d = (p_1 \ p_2 \ \dots \ p_t)$ is:

$$E_d[B] = E[B] + Cov[B, D]Var[D]^{-1}[d - E[D]] \quad [29]$$

where $Var[D]^{-1}$ is the standard inverse of $Var[D]$. Note that $E_d[B]$ depends on observed prices.

The adjusted variance matrix given that D is observed is:

$$Var_D[B] = Var[B] - Cov[B, D]Var[D]^{-1}Cov[D, B][29]$$

Again $Var[D]^{-1}$ is the standard inverse of $Var[D]$. $Var_D[B]$ does not depend on observed prices.

In the tool, these adjustments are calculated numerically within the R routines [29]. Checks are made to ensure that the variance matrices are non-negative definite [29].

Chapter 4

Utility analysis

A utility function for wealth is a mathematical measurement of a person's happiness for owning money. Within this research, utility theory is applied to a decision making process to choose between risky investment alternatives. It is adopted as a vital mechanism to manage the sell, hold, and buy decisions faced by the commodity trader.

Utility theory began with Jeremy Bentham (1748-1832), a philosopher and critique of the law, who was the first person to write about the utilitarian approach. He proposed that the law system should be structured to maximize the total utility of the country. Bentham's hypotheses on the utility for wealth are [26]:

1. Each portion of wealth has a corresponding portion of happiness.
2. Of two individuals with unequal fortunes, he who has the most wealth has the most happiness.
3. The excess in happiness of the richer will not be so great as the excess of his wealth.

This started the discussion on utility functions - questions were raised on its usefulness, applicability, and measurability. Many academics have ob-

jected to utility and alternative theories have been put forward. Debate still continues today - doubt about how people deal with losses under uncertainty and the elicitation of the utility function are among the unresolved problems.

In 1713, before Bentham was born, an insightful dilemma called the St. Petersburg Paradox was devised by Nicholas Bernoulli. This challenged the popular thought that the maximum people are willing to pay for an investment is the expected gain. The famous puzzle is as follows [27]:

Peter tosses a coin in the air repeatedly until it falls heads up. If this occurs on the n^{th} throw he pays Paul $\mathcal{L}2^{n-1}$. What is the maximum amount Paul should pay for this game?

The expected return from this game is:

$$\sum_{n=1}^{\infty} \frac{1}{2^n} \mathcal{L}2^{n-1} = \infty$$

Yet a reasonable Paul will not pay this infinite amount for the game. A solution was proposed by Daniel Bernoulli in which he advocated diminishing marginal utility of money - or a utility function that is bounded above, such that, Paul will pay:

$$\sum_{n=1}^{\infty} \frac{1}{2^n} U(\mathcal{L}2^{n-1}) = \mathcal{L}Finite$$

Daniel Bernoulli proposed that $U(\mathcal{L}X) = \log \frac{X}{c}$, where c is the amount required to live. Bernoulli also advised that it is reasonable that initial wealth or status will affect how much Paul is willing to pay for the game. He also 'suggested the assumption of constant marginal utility for small variations in wealth. [27]' This means that one is prepared to pay the expected gain of an investment if they are only putting a small amount of wealth at risk.

Progress, since the Bernoullis' valuable contribution, was mainly in the domain of measuring the utility function with solutions provided by Fisher and Pareto [27]. Otherwise, there were arguments for abandoning utility or the theory was completely ignored. Today, the theory is not overlooked, yet the ability to apply it to solve a practical decision problem remains difficult and subjective.

Utility theory is widespread in the academic world. However, it has yet to establish itself as a useful practical tool in business applications. This is partly owing to the challenges involved with building an individual's, or a company's, utility function. For example, there is no publicized method of eliciting a utility function that boasts empirical evidence to prove that it is effective for deriving a true representation of a person's risk attitude.

A prototype Java tool that elicits a trader's utility function has been developed within this study. Problems encountered, and questions raised, while creating this tool are discussed within this chapter.

4.1 Building a utility function

A technique for constructing a person's utility curve is to ask a series of structured risk-reward related questions. Preferably these questions are designed to make the process easily understood and intuitive for the interviewee(s). This is not a trivial task.

The utility function for wealth must satisfy a series of axioms as per the expected utility theorem and has defined properties (see below). The axioms, in particular, certainty equivalence, prove very useful for building the utility function. The property of non-satiation - one must prefer more to less - is easy to impose and the property of taste for risk is a great asset for visually interpreting a plot of a utility curve. If an investor is risk averse, the second

derivative of the curve is negative, if the investor is risk seeking the second derivative is positive, and if the investor is risk neutral the second derivative is zero.

To begin construction, boundary values on the axes of the utility function for wealth are predetermined. Arbitrary values are chosen on the y axis - the values that are chosen here are not important as it is the relative values of a measurement scale that are significant. On the x axis, which shows monetary values, the bounds are set such that the amounts that need to be considered by the decision maker are within the range.

A series of questions are now asked about a certain amount the individual is willing to exchange for an investment that holds an element of risk. Points on the utility function are calculated from the answers and eventually the person's risk behaviour emerges. The curve is completed by fitting a line through the derived co-ordinates.

4.1.1 Example - Building your utility function

The facilitator sets the boundary values to:

1. $U(\pounds 1000) = 1$
2. $U(\pounds 0) = 0$

He then asks you a question of the following format:

Suppose that you hold an investment that will pay out $\pounds 1000$ with probability 0.5 or $\pounds 0$ with probability 0.5. What is the least amount of money you will sell this investment for?

A coherent answer, which satisfies the axiom of non-satiation, will be somewhere between $\pounds 0$ and $\pounds 1000$. Say, that your answer is $\pounds K$ - an extra point

on the utility curve is calculated as follows:

$$U(\mathcal{L}K) = U(\mathcal{L}1000) \times 0.5 + U(\mathcal{L}0) \times 0.5$$

$$U(\mathcal{L}K) = 1 \times 0.5 + 0 \times 0.5 = 0.5$$

Your utility for $\mathcal{L}K$ is now set to 0.5.

A risk averse investor will sell for less than $\mathcal{L}500$, a risk neutral investor will sell for $\mathcal{L}500$, and a risk seeking investor will sell for more than $\mathcal{L}500$.

The process is continued by asking what the least amount of money for which you will sell an investment that pays out $\mathcal{L}K$ with probability 0.5, or nothing. Suppose that your answer is $\mathcal{L}L$, then:

$$U(\mathcal{L}L) = U(\mathcal{L}K) \times 0.5 + U(\mathcal{L}0) \times 0.5$$

$$U(\mathcal{L}L) = 0.5 \times 0.5 = 0.25.$$

Your utility for $\mathcal{L}L$ is now set to 0.25.

If the range of values to be considered is widened from $[\mathcal{L}0, \mathcal{L}1000]$, to say $[\mathcal{L}-1000, \mathcal{L}2000]$, then the process needs to start again from the beginning. However, the questions answered for the values between $\mathcal{L}0$ and $\mathcal{L}1000$ are kept as long as the facilitator ensures that the scaling on the y-axis is consistent and correct.

The number of questions asked will depend on the importance of the decisions to be made by the utility function. Enough points to show a general

shape will suffice in most circumstances.

There is a risk that the resultant function is a poor representation of the true utility curve owing to the interviewee misunderstanding and mistakes on the part of the facilitator. Nevertheless, a good evaluation of a utility will be an effective tool to optimize investment decisions.

4.1.2 The expected utility theorem

The expected utility theorem states that for an investor who satisfies the following four axioms, a function $U(w)$ can be constructed representing the investors utility for wealth [14].

1. Comparability

An investor can state a preference among all certain alternatives.

2. Transitivity

If an investor prefers A to B and B to C, then A is preferred to C.

3. Independence

Assume an investor is indifferent between two certain outcomes X and Y. Then the investor is also indifferent between:

- (a) X with probability p and Z with probability $(1-p)$
- (b) Y with probability p and Z with probability $(1-p)$

4. Certainty Equivalence

If A is preferred to B and B is preferred to C, then there is a unique probability p , such that the investor is indifferent between:

- (a) B with certainty
- (b) A with probability p , C with probability $(1-p)$

B is known as the certainty equivalent of gamble (b).

4.1.3 Properties of utility functions

The properties of utility functions are as follows [14]:

1. **Non-satiation** An investor always prefers more to less. $U'(w) > 0 \forall w$.
2. **Taste for risk**
 - (a) Risk averse investors reject fair gambles. $U''(w) < 0 \forall w$.
 - (b) Risk neutral investors are indifferent to fair gambles. $U''(w) = 0 \forall w$.
 - (c) Risk seeking investors accept fair gambles. $U''(w) > 0 \forall w$.
3. **Absolute risk aversion** Absolute risk aversion considers the amount of money invested in risky assets and is denoted $A(w)$.

$$A(w) = \frac{-U''(w)}{U'(w)}$$

<i>Condition</i>	<i>Definition</i>	<i>Property</i>
Increasing absolute risk aversion	As wealth increases invest less in risky assets	$A'(w) > 0$
Constant absolute risk aversion	As wealth increases invest the same in risky assets	$A'(w) = 0$
Decreasing absolute risk aversion	As wealth increases invest more in risky assets	$A'(w) < 0$

Most investors exhibit decreasing absolute risk aversion.

4. Relative Risk Aversion

Relative risk aversion considers the percentage of money invested in risky assets and is denoted $R(w)$.

$$R(w) = \frac{wU''(w)}{U'(w)} = wA(w)$$

<i>Condition</i>	<i>Definition</i>	<i>Property</i>
Increasing relative risk aversion	As wealth increases invest a smaller percentage of wealth in risky assets	$R'(w) > 0$
Constant relative risk aversion	As wealth increases invest the same percentage of wealth in risky assets	$R'(w) = 0$
Decreasing relative risk aversion	As wealth increases invest a greater percentage of wealth in risky assets	$R'(w) < 0$

4.2 Standard utility functions

Various standard mathematical functions have been used to model utility functions. These include power functions, quadratic functions, logarithmic functions, and the exponential function. These all have different properties with regards to the importance of the sign of the derivatives and these properties are described in this section. Using approximate functions for utility curves simplifies the problem of building bespoke utility curves to suit the individual and also simplifies the mathematics involved in optimizing decisions by maximizing utility.

4.2.1 The logarithmic utility curve

The logarithmic utility function has the following form:

$$U(w) = \ln(w) \quad \forall w > 0$$

The properties of this utility function are shown here:

$$U'(w) = \frac{1}{w} \quad \forall w > 0$$

Since $\frac{1}{w} > 0 \forall w > 0$, the function describes an investor who prefers more to less.

$$U''(w) = -\frac{1}{w^2} \quad \forall w > 0$$

Since $-\frac{1}{w^2} < 0 \forall w > 0$, the function describes a risk averse investor.

$$A(w) = \frac{-U''(w)}{U'(w)} = \frac{1}{w} \quad \forall w > 0$$

$$A'(w) = -\frac{1}{w^2} \quad \forall w > 0$$

Since $-\frac{1}{w^2} < 0 \forall w > 0$, the function describes an investor who displays decreasing absolute risk aversion.

$$R(w) = wA(w) = \frac{w}{w} = 1$$

$$R'(w) = 0$$

Since $R'(w)=0$, the function describes an investor with constant relative risk aversion.

If an investor wished to consider a wealth less than zero, the function $U(w) = \ln(w + x)$ will be used where w denotes wealth and x denotes an absolute value of the highest possible negative wealth to include in the utility model. The logarithmic function seems to be a very useful model of an investor's utility since the properties match a common investor's attitudes well.

4.2.2 The risk neutral utility curve

A risk neutral utility curve takes the following form:

$$U(w) = w$$

Since $U'(w) = 1$ and $U''(w) = 0$, the risk neutral investor prefers more to less and will accept a fair gamble. $A'(w) = 0$ and $R'(w) = 0$. However, it is not

possible to demonstrate both constant risk aversion and constant relative risk aversion, and hence the calculations for $A'(w)$ and $R'(w)$ do not make sense in this instance. The amount or proportion of wealth invested in risky assets is not implicit in the risk neutral function. The decision made using this function do not depend on current state of wealth. This is a disadvantage in using the risk neutral model as research has shown that persons of different status invest differently.

4.2.3 The quadratic utility curve

This type of function is used to model utility over a restricted range of wealth using the increasing part of the curve to conform with the investor typically preferring more to less. Let us evaluate a quadratic utility function of the form:

$$U(w) = w - bw^2$$

where b needs to be defined to facilitate the band of wealth to be considered.

$$U'(w) = 1 - 2bw$$

$1-2bw < 0$ iff $b < \frac{1}{2w}$. Here, the investor prefers more to less.

$$U''(w) = -2b$$

$-2b < 0$ iff $b > 0$. Here, the investor is risk averse.

$$A(w) = \frac{2b}{1 - 2bw}$$

$$A'(w) = \frac{4b^2}{(1 - 2bw)^2}$$

$\frac{4b^2}{(1-2bw)^2} > 0$ iff $b < \frac{1}{2w}$ and if this condition is satisfied the utility models an investor with increasing absolute risk aversion.

$$R(w) = \frac{2bw}{1 - 2bw}$$

This is differentiated using the product rule:

$$R'(w) = \frac{2b}{(1 - 2bw)^2}$$

$\frac{2b}{(1-2bw)^2}$ iff $b > 0$ and $b < \frac{1}{2w}$. This models an investor who displays increasing relative risk aversion. This means that the proportion of wealth invested in risky assets declines as wealth increases.

The parameters of the quadratic utility function will be tailored to suit the investors assets closely as their wealth changes if this is deemed a suitable process. Also the quadratic function is useful in mean-variance analysis as the expected utility can be defined in terms of means and variances when utility is quadratic - $E[U(w)] = E[w] - b\{\sigma_w^2 - (E[w])^2\}$ [4].

4.2.4 The exponential utility curve

In the approach used to model a trader's risk attitudes an assumption is made that the exponential function is suitable. The trader's utility function for wealth, W , is expressed in the following form:

$$U(W) = 1 - \exp\left(\frac{-W}{R}\right) \quad (4.1)$$

The exponential utility function describes a risk averse investor, who prefers more to less, and exhibits constant absolute risk aversion (as his wealth increases, he still invests the same amount of money in risky assets).

In the exponential utility function above, R , is known as the risk factor, and R must be greater than 0. The smaller the value of R , the more risk averse the utility function, and as $R \rightarrow \infty$, the exponential utility function approaches a risk neutral utility function.

The distinct objection in using the exponential utility is its property of constant absolute aversion. A real life investor is unlikely to behave in this

way. Nevertheless, the exponential function keeps the mathematics easy to manage. The optimization problem has not been solved using any other utility function (apart from risk neutral), and hence the exponential function is the best representation of a risk attitude that can be used at the moment. There is definitely scope for improvement in the future.

As $R \rightarrow \infty$

Consider the following *gamble*:

Y with probability p , or
 0 with probability $1-p$

The risk neutral investor is indifferent between this gamble and $pU(Y) = pY$. pY is also known as the risk neutral investor's certainty equivalent for the above gamble.

Suppose that an investor has an exponential utility function, such that, $U(W) = 1 - \exp(\frac{-W}{R})$, then their utility for the above gamble is:

$$U(\text{gamble}) = p(1 - \exp(\frac{-Y}{R})) + (1 - p)(1 - \exp(\frac{-0}{R}))$$

$$U(\text{gamble}) = p(1 - \exp(\frac{-Y}{R}))$$

Under the laws of the expected utility theorem the utility for the maximum certain amount that the investor is prepared to exchange for the gamble is equivalent to the utility for the gamble.

Let the Certainty Equivalent = CE , then, $U(CE) = U(\text{gamble})$:

$$1 - \exp\left(\frac{-CE}{R}\right) = p\left(1 - \exp\left(\frac{-Y}{R}\right)\right)$$

With some simple rearrangement, it is found that:

$$CE = -R \ln\left[1 - p\left(1 - \exp\left(\frac{-Y}{R}\right)\right)\right]$$

Let $X = \frac{-Y}{R}$, then:

$$CE = -R \ln[1 - p(1 - \exp(X))]$$

Since $\exp(X) = 1 + X + \frac{X^2}{2!} + \frac{X^3}{3!} + \dots$,

$$CE = -R \ln\left[1 - p\left(1 - \left(1 + X + \frac{X^2}{2!} + \frac{X^3}{3!} + \dots\right)\right)\right]$$

$$CE = -R \ln\left[1 + p\left(X + \frac{X^2}{2!} + \frac{X^3}{3!} + \dots\right)\right]$$

Since $\ln(1 + Z) = Z - \frac{Z^2}{2} + \frac{Z^3}{3} - \dots$,

$$CE = -R\left[p\left(X + \frac{X^2}{2!} + \frac{X^3}{3!} + \dots\right) - \frac{\left(p\left(X + \frac{X^2}{2!} + \frac{X^3}{3!} + \dots\right)\right)^2}{2} + \dots\right]$$

$$CE = -pRX + Rorder(X^2)$$

Substitute $X = \frac{-Y}{R}$,

$$CE = \frac{pRY}{R} + Rorder(X^2)$$

$$CE = pY + Rorder(X^2)$$

As $R \rightarrow \infty$, $CE \rightarrow pY$. This shows that as $R \rightarrow \infty$, the exponential utility of the form, $U(W) = 1 - \exp(-\frac{W}{R})$, tends towards a risk neutral utility function.

4.3 Piecewise utility

A piecewise function takes a different form over split sections of the X values.

Consider the following piecewise utility function which has 4 sections:

$$U(X) = \begin{cases} \alpha_1 + \beta_1 \exp^{XR_1} & \text{for } knot1 < X < knot3 \\ \alpha_2 + \beta_2 \exp^{XR_2} & \text{for } knot3 \leq X < knot5 \\ \alpha_3 + \beta_3 \exp^{XR_3} & \text{for } knot5 \leq X < knot7 \\ \alpha_4 + \beta_4 \exp^{XR_4} & \text{for } knot7 \leq X < knot9 \end{cases}$$

where $knot1, knot3, knot5, knot7$, and $knot9$ are predetermined X values such that $knot1 < knot3 < knot5 < knot7 < knot9$. There are nine knots altogether and the function of $knot2, knot4, knot6$ and $knot8$ is discussed later. $R_i \neq 0$.

If this piecewise function is used then it must be that $P(X < knot1) = 0$ and $P(X > knot9) = 0$. This is because all possible values of X need a corresponding utility value so that $E[U(X)]$ can be maximized. Since it is assumed that X is distributed normally, then these probabilities are only 0 if $knot1 = -\infty$ and $knot9 = \infty$. There is no problem in doing this since $\alpha_1 + \beta_1 \exp^{XR_1}$ can be used for all X values less than $knot3$ and $\alpha_4 + \beta_4 \exp^{XR_4}$ can be used for all X values greater than $knot7$.

It is also assumed that $X \sim N(\mu_X, \sigma_X^2)$ and X is realizable profit. This means that X has MGF $M_X(t) = E[\exp^{tX}] = \exp^{t\mu_X + \frac{1}{2}t^2\sigma_X^2}$

To choose the number of lots to maximize the expected utility of realizable profit the initial step is to find $E[U(X)]$.

Theorem

$$\begin{aligned}
 E[U(X)] = & p_1\alpha_1 + \beta_1 \exp^{\mu_X R_1 + \frac{1}{2}\sigma_X^2 R_1^2} p_1^\dagger \\
 & + p_2\alpha_2 + \beta_2 \exp^{\mu_X R_2 + \frac{1}{2}\sigma_X^2 R_2^2} p_2^\dagger \\
 & + p_3\alpha_3 + \beta_3 \exp^{\mu_X R_3 + \frac{1}{2}\sigma_X^2 R_3^2} p_3^\dagger \\
 & + p_4\alpha_4 + \beta_4 \exp^{\mu_X R_4 + \frac{1}{2}\sigma_X^2 R_4^2} p_4^\dagger
 \end{aligned}$$

where

$$\begin{aligned}
 p_1 &= P(\text{knot1} < X_A < \text{knot3}) & \text{and} & & X_A &\sim N(\mu_X, \sigma_X^2) \\
 p_2 &= P(\text{knot3} < X_B < \text{knot5}) & \text{and} & & X_B &\sim N(\mu_X, \sigma_X^2) \\
 p_3 &= P(\text{knot5} < X_C < \text{knot7}) & \text{and} & & X_C &\sim N(\mu_X, \sigma_X^2) \\
 p_4 &= P(\text{knot7} < X_D < \text{knot9}) & \text{and} & & X_D &\sim N(\mu_X, \sigma_X^2) \\
 p_1^\dagger &= P(\text{knot1} < X_A < \text{knot3}) & \text{and} & & X_A &\sim N(\mu_X + R_1\sigma_X^2, \sigma_X^2) \\
 p_2^\dagger &= P(\text{knot3} < X_B < \text{knot5}) & \text{and} & & X_B &\sim N(\mu_X + R_2\sigma_X^2, \sigma_X^2) \\
 p_3^\dagger &= P(\text{knot5} < X_C < \text{knot7}) & \text{and} & & X_C &\sim N(\mu_X + R_3\sigma_X^2, \sigma_X^2) \\
 p_4^\dagger &= P(\text{knot7} < X_D < \text{knot9}) & \text{and} & & X_D &\sim N(\mu_X + R_4\sigma_X^2, \sigma_X^2)
 \end{aligned}$$

Proof

$$\begin{aligned}
 E[U(X)] = & p_1\alpha_1 + \beta_1 E[\exp^{X R_1}] \\
 & + p_2\alpha_2 + \beta_2 E[\exp^{X R_2}] \\
 & + p_3\alpha_3 + \beta_3 E[\exp^{X R_3}] \\
 & + p_4\alpha_4 + \beta_4 E[\exp^{X R_4}]
 \end{aligned}$$

where

$$p_1 = P(\text{knot1} < X < \text{knot3})$$

$$p_2 = P(\text{knot3} < X < \text{knot5})$$

$$p_3 = P(\text{knot5} < X < \text{knot7})$$

$$p_4 = P(\text{knot7} < X < \text{knot9})$$

(4.2)

$$\begin{aligned} E[\exp^{XR_1}] &= \frac{1}{\sigma\sqrt{2\pi}} \int_{\text{knot1}}^{\text{knot3}} \exp^{R_1X} \exp^{-\frac{1}{2}\left(\frac{X-\mu_X}{\sigma_X}\right)^2} dX \\ &= \frac{1}{\sigma\sqrt{2\pi}} \int_{\text{knot1}}^{\text{knot3}} \exp^{R_1X - \frac{1}{2}\left(\frac{X-\mu_X}{\sigma_X}\right)^2} dX \end{aligned}$$

$$\begin{aligned} R_1X - \frac{1}{2}\left(\frac{X - \mu_X}{\sigma_X}\right)^2 &= -\frac{1}{2} \frac{(X - \mu_X)^2 - 2R_1X\sigma_X^2}{\sigma_X^2} \\ &= -\frac{1}{2} \frac{X^2 - 2X\mu_X + \mu_X^2 - 2R_1X\sigma_X^2}{\sigma_X^2} \\ &= -\frac{1}{2} \frac{X^2 - 2(X\mu_X + R_1X\sigma_X^2) + \mu_X^2}{\sigma_X^2} \end{aligned}$$

$$(X - (\mu_X + R_1\sigma_X^2))^2 + C = X^2 - 2(X\mu_X + R_1X\sigma_X^2) + \mu_X^2$$

$$C = -2R_1\mu_X\sigma_X^2 - R_1^2\sigma_X^4$$

$$\begin{aligned} R_1X - \frac{1}{2}\left(\frac{X - \mu_X}{\sigma_X}\right)^2 &= -\frac{1}{2} \frac{(X - \mu_X - R_1\sigma_X^2)^2 - 2\mu_X\sigma_X^2R_1 - \sigma_X^4R_1^2}{\sigma_X^2} \\ &= -\frac{1}{2}\left(\frac{(X - \mu_X - R_1\sigma_X^2)^2}{\sigma_X^2}\right) + \mu_XR_1 + \frac{1}{2}\sigma^2R_1^2 \end{aligned}$$

$$\begin{aligned}
E[\exp^{XR_1}] &= \frac{1}{\sigma\sqrt{2\pi}} \int_{knot1}^{knot3} \exp^{R_1X - \frac{1}{2}\left(\frac{X-\mu_X}{\sigma_X}\right)^2} dX \\
&= \frac{1}{\sigma\sqrt{2\pi}} \int_{knot1}^{knot3} \exp^{-\frac{1}{2}\left(\frac{X-\mu_X-R_1\sigma_X^2}{\sigma_X}\right)^2 + \mu_X R_1 + \frac{1}{2}\sigma^2 R_1^2} dX \\
&= \exp^{\mu_X R_1 + \frac{1}{2}\sigma^2 R_1^2} \frac{1}{\sigma\sqrt{2\pi}} \int_{knot1}^{knot3} \exp^{-\frac{1}{2}\left(\frac{X-(\mu_X+R_1\sigma_X^2)}{\sigma_X}\right)^2} dX \\
&= \exp^{\mu_X R_1 + \frac{1}{2}\sigma^2 R_1^2} Prob(knot1 < X_A < knot3)
\end{aligned}$$

where $X_A \sim N(\mu_X + R_1\sigma_X^2, \sigma_X^2)$

$E[U(X)]$ is now maximized with respect to l_t to choose the number of lots to buy at time t to maximize the expected utility of realizable profit. This needs to be done numerically given that:

$$\begin{aligned}
\mu_X &= Balance_{t-1} - l_t p_t + (Position_{t-1} + l_t)E[P_{t+1}] \\
\sigma_X^2 &= (Position_{t-1} + l_t)^2 V[P_{t+1}]
\end{aligned}$$

as explained in a previous chapter.

4.3.1 Example

To construct a nine knot piecewise utility function the X co-ordinate points of all the knots are fixed by the facilitator. Also, the corresponding utility values, $U(X)$, for knot 1, knot 5, and knot 9 are predetermined. In this example:

- knot 1 = (-100,000,-1)
- knot 2 = (-80,000, ?)
- knot 3 = (-50,000, ?)

- knot 4 = (-30,000, ?)
- knot 5 = (0,0)
- knot 6 = (40,000,?)
- knot 7 = (150,000,?)
- knot 8 = (350,000,?)
- knot 9 = (400,000,1)

The chance of X being less than knot 1 or greater than knot 9 needs to be negligible.

Suppose that an investor is indifferent between the following certain amounts and a gamble:

$$\begin{array}{l}
 \$150,000 \quad \text{and} \quad \left\{ \begin{array}{l} \$400,000 \quad \text{w.p. } 0.1 \\ \$0 \quad \text{w.p. } 0.9 \end{array} \right. \\
 \\
 \$40,000 \quad \text{and} \quad \left\{ \begin{array}{l} \$150,000 \quad \text{w.p. } 0.25 \\ \$0 \quad \text{w.p. } 0.75 \end{array} \right. \\
 \\
 \$350,000 \quad \text{and} \quad \left\{ \begin{array}{l} \$400,000 \quad \text{w.p. } 0.92 \\ \$150,000 \quad \text{w.p. } 0.08 \end{array} \right. \\
 \\
 -\$50,000 \quad \text{and} \quad \left\{ \begin{array}{l} -\$100,000 \quad \text{w.p. } 0.2 \\ \$0 \quad \text{w.p. } 0.8 \end{array} \right. \\
 \\
 -\$30,000 \quad \text{and} \quad \left\{ \begin{array}{l} -\$50,000 \quad \text{w.p. } 0.5 \\ \$0 \quad \text{w.p. } 0.5 \end{array} \right. \\
 \\
 -\$80,000 \quad \text{and} \quad \left\{ \begin{array}{l} -\$100,000 \quad \text{w.p. } 0.5 \\ -\$50,000 \quad \text{w.p. } 0.5 \end{array} \right.
 \end{array}$$

then, the following Y co-ordinate points on the investor's utility curve can be found (using the mathematics explained in section 4.1):

- knot2 = (-80,000, -0.6)
- knot3 = (-50,000, -0.2)
- knot4 = (-30,000, -0.1)
- knot6 = (40,000,0.1)
- knot7 = (150,000,0.4)
- knot8 = (350,000,0.95)

These knots are now used to fit a curve so that a utility can be given to every value of X between knot 1 and knot 9. This curve takes the form $\alpha + \beta \exp^{-XR}$.

Four curves are fitted:

- Fitting a curve using knots 1,2,and 3 gives $0.427 - 0.276 \exp^{1.64 \times 10^{-5} X}$
- Fitting a curve using knots 3,4,and 5 gives $0.157 - 0.157 \exp^{1.64 \times 10^{-5} X}$
- Fitting a curve using knots 5,6,and 7 gives $-2.12 + 2.12 \exp^{-1.15 \times 10^{-6} X}$
- Fitting a curve using knots 7,8,and 9 gives $1.11 - 2.164 \exp^{7.418 \times 10^{-6} X}$

The resultant utility function is shown in Figure 4.1.




Figure 4.1: A Piecewise Utility Function

4.3.2 Problems

There are some problems with this approach. Firstly, the utility for $U(0)$ was set to 0 when 0 is not a boundary X value. This has an impact on the shape of the utility curve without considering any answer to a risk-reward question. This can be corrected by only predetermining the utility for extreme X values.

It is unnecessary for the utility for \$0 to be 0. Nevertheless, it is desirable for $U(0) = 0$ since this is commonplace in utility theory literature. This can be achieved by eliciting the utility function first, say $U(X_1)$, and then shifting the curve by some value c such that $U(X_2) = U(X_1) - c$. c is found where $U(X_1 = 0)$.

Secondly, if the curve was fitted through different points a very different curve will emerge. If a line was fitted through knot 6, 7, and 8 in the above example, then the investor would be risk seeking for X values between knot 7 and 8 instead of risk averse. This is owing to limited information. The curve will become more accurate with more indifference points (more knots) identified by the investor.

Thirdly, the lack of detail and approximations can also cause sharp kinks in the curve and these are an undesirable feature. However, it is possible to fit a smoother curve such that the kinks are removed.

Finally, the piecewise function is inelegant and complex to use in optimization calculations. Nevertheless, it is not impossible to maximize the expected piecewise utility of profit.

4.4 A utility function based on a massive loss

If a person, business, or trader can state a loss of money that will have a big effect on their status, then this may help in building their utility function. For example, for a person, this massive loss of money will relate to losing their house, and for a business, the consequence of this massive loss will mean bankruptcy. For a trader, this massive loss will be the maximum loss they can withstand before their budget is gone, and this loss will probably also lose them their job.

One simple ad hoc method of defining a utility function, given a massive loss specified by the trader, is setting the parameters of the function so that¹:

$$U(M) = -Y$$

where M is the massive loss and,

Y is the lower bound of the y -axis, $U(M)$, which is preset.

4.5 Loss aversion

It is not clear how people behave when dealing with losses. This might be because the way people cope with losses is part of the personality trait which is unique to the individual.

Studies [17] [6] have found evidence that many people are risk accepting with gambling scenarios containing losses.

Nevertheless, in real life, people can be observed to be very risk averse to losses. For example, they buy insurance, invest in burglar alarms, and prefer secure employment.

There are others who will only buy insurance if required by law, start

¹Idea taken from David Wooff

up their own enterprises at their own risk, and enjoy betting money at the casino.

4.6 Disallowing risk seeking attitudes

4.6.1 Objection to the risk seeking curve

Perhaps, it is true that normal human behaviour dictates a risk accepting attitude when faced with losses or debts. This is supported by many studies [17] [16] [6] and led to the development of Tversky's and Kahneman's Prospect Theory [17].

If each trader becomes risk seeking when faced with losses, others will, on average, profit from this irrational behaviour.

A trader must ignore normal irrational tendencies and needs to be averse to losses to become successful. A successful trader will be aware of the importance of stopping losses from escalating. He will not trade with emotion or irrationality when he is faced with a loss.

Although there is evidence to suggest that people do have risk seeking tendencies when concerned with losses, does it make sense to mimic this behaviour to assist with decision making? The tool is not designed to support a trader in making incoherent decisions that increase his chances of ruin, and hence the utility elicitation component will force a risk averse or risk neutral attitude.

4.6.2 Lottery players and gamblers

A person who buys a UK National Lottery Ticket is accepting an unfair gamble. For each lottery ticket that is bought, the expected return is a loss of £0.55.

It is argued that there is utility gained by the National Lottery player that

outweighs the negative returns from the game - or there are other sensible reasons to explain why consumers participate in the lottery. These include:

1. Excitement from watching the National Lottery Draw.
2. £0.28 of the £1 stake goes toward good causes².
3. The consumer has been manipulated by advertising campaigns.
4. Regret Theory. If a person puts on the same numbers every week, they will be unable to cope with the feeling of regret if they neglected to put their line on one week, and their numbers came up.
5. The loss of £1 will have a negligible effect on a person's utility and the unhappy feeling from losing passes in a short space of time.
6. The consumer is misinformed about the risk associated with buying the ticket.
7. Gambling Addiction.

If there is a roll over lottery draw (the jackpot was not won in the last draw and hence it was carried forward), then it is feasible that the expected return on a ticket could be greater than £1³. In this scenario, some investors may perceive the lottery ticket as a good investment, even though the chance of winning the millions is small and may need to be shared with any other winners.

Other types of gambling, including sports betting and casino gaming, can be viewed in the same manner as taking part in lotteries. Perhaps there is a difference when placing a bet on sporting events as the punter may have

²Source: www.national-lottery.co.uk

³Observation by David Wooff

knowledge that makes a book maker's odds look attractive. Most of the time the book maker's odds are not favourable. For example, if the betting shop believes that a football team will win with odds 10 to 1, they will offer the customer unappealing odds, similar to 5 to 1, for the team in question to win.

There is empirical evidence that shows consumers do not spend a large proportion of their income on gambles where the expected return is negative. The small amounts of money people do spend on gambling, can be construed as rational decisions. A bet once a year on the Grand National is not perceived as an irrational choice or addictive gambling, it merely adds to a person's excitement of watching the race.

4.6.3 The existence of the insurance industry

Consider a policy that will pay out the market value of a house, say £100,000, if the house burns down. The policy customer and the insurance company both assess that the risk of the house burning down is 0.001.

The insurance company must charge more than $0.001 \times £100,000 = £100$ (and sell many policies) to have a good chance of making some profit.

To buy the policy the customers will need to prefer losing a fixed amount greater than £100 to risking the loss of £100,000 with probability 0.001.

Look at Table 4.1 to see the Loss/Gain table for the policy holder in all the possible scenarios. Here P denotes the premium charged by the insurance company. The premium needs to be greater than £100 for the insurance company to make a profit. This is an amount greater than the expected loss. A risk averse person will pay more than the expected loss to protect themselves against a massive loss, but a risk neutral or risk accepting consumer will not pay more than £100. Since there are many insurance companies

Table 4.1: Loss/Gain Table

	Insurance	No Insurance
House Burns Down (w.p. 0.001)	$\mathcal{L}100,000 - \mathcal{L}P$	$-\mathcal{L}100,000$
House Does Not Burn Down (w.p. 0.999)	$-\mathcal{L}P$	$\mathcal{L}0$

making profits, there must also be many risk averse people.

4.7 Alternative investment selection models

Competing alternatives to utility theory for selecting between investments, that hold an element of risk, include maximizing the geometric mean return, safety first models, stochastic dominance, Value at Risk, and Kahneman and Tversky's Prospect Theory[17], which are now discussed.

4.7.1 Maximizing the geometric mean return

To choose between investments, an investor may decide by selecting the option with the highest expected geometric return. The geometric mean, denoted \bar{R}_{Gj} , is defined as:[4]

$$\bar{R}_{Gj} = \prod_{i=1}^N (1 + R_{ij})^{P_{ij}} - 1$$

where

R_{ij} is the i^{th} possible return on the j^{th} portfolio, and

P_{ij} is the probability of the i^{th} outcome for portfolio j .

It is interesting to note that maximizing the geometric return is exactly the same as maximizing the expected value of a log utility function [4] - but not any other utility function.

By maximizing the geometric mean, an investor will never pick a strategy with a probability of ruin because if there is a possible outcome that will wipe out all capital (the return is -100%), then the geometric return will be zero⁴. If the probability of losing all capital invested is small, and the expected return is high, an investor may still be attracted to the investment - and in this scenario, maximizing the geometric return will be inappropriate.

4.7.2 Safety first models

Safety first models are based on reducing the risk of bad returns or losses. The first criterion of safety first, developed by Roy, is to choose the portfolio that has the smallest probability of producing a return (R_P) below a benchmark return (R_L) [4].

The second criterion (defined by Kataoka) states that an investor will maximize the amount R_L which return can fall below subject to $P(R_P < R_L) \leq \alpha$. α is a probability set by the decision maker [4].

The final criterion, by Tesler, states that an investor will maximize expected return, \bar{R}_P , subject to $P(R_P < R_L) \leq \alpha$. Again, α is a probability set by the decision maker [4].

First Criterion	minimize $P(R_P < R_L)$
Second Criterion	maximize R_L subject to $P(R_P < R_L) \leq \alpha$
Third Criterion	maximize $E[R_P]$ subject to $P(R_P < R_L) \leq \alpha$

4.7.3 Value at Risk

Value at Risk (VaR) is a measure of downside risk. Most banks and traders use VaR to evaluate their exposure to making a loss given their capital investments. How do they do this?

⁴ $(1+R_{ij}) = (1-1) = 0 \dots$

Risk managers need to estimate a distribution of return for the combination of all investments (their portfolio). This includes making probability assessments about how returns will be affected by adverse events.

The portfolio that is chosen will depend on the degree of acceptable risk. For example, a risk averse manager will not be prepared to accept a large risk of loss. Nevertheless, a more risk accepting manager will be prepared to accept a higher chance of loss, as this will also increase potential for higher returns.

A risk manager will analyze what the company can afford to lose, and adjust the investment portfolio to suit the company's risk profile. A risk manager with \$1 million will specify something like:

- Accept a 0.001 chance of losing more than \$100K (10% of capital)
- Accept a 0.01 chance of losing more than \$90K (9% of capital)
- Accept a 0.05 chance of losing more than \$75K (7.5% of capital)
- Accept a 0.3 chance of losing more than \$0 (0% of capital)

The manager is prepared to accept a 30% risk of loss for a 70% chance of a gain. Of course, if there is a risk of losing more than \$100K, there will also need to be a chance of being rewarded a gain of comparable value. A 0% return on capital is a loss, since the value of capital depreciates with time - this is owing to inflation and opportunity cost of capital.

To build an investment portfolio to suit risk preferences, a manager will typically diversify by using different types of asset classes. For example:

- Gilts. These are government bonds that pay a risk free, but low rate of interest. Increase investment in gilts to reduce risk.

- **Property.** Investing in property has potential for high returns. Nevertheless, property is not a liquid asset, and hence it is not possible to realize the profit instantaneously.
- **Equities.** Investors normally buy stocks for long term capital growth. Investing in equities will increase the risk of the portfolio. However, choosing stocks across a wide range of sectors can mitigate some risk.
- **Commodities.** Futures trading is typically high risk. However, commodities are negatively correlated with stocks (as commodity prices increase, stocks decrease in value). This is because the increased cost of raw material increases many listed companies costs thus reducing their profit. Commodities are a good diversification tool in a fund with a high proportion of capital invested in stocks - with careful planning spending capital in futures trading can reduce risk. Commodities as a sole investment are risky.

If the company becomes more risk averse its portfolio will need to change by putting more capital into low risk investments. Also, if a company is prepared to take more risk of loss to be rewarded with higher potential gains, the capital will need to be shifted to more risky ventures.

VaR is about controlling losses and analyzing worst case scenarios given that capital funds are invested in risky assets. Expected profits will be maximized subject to the implications of VaR.

4.7.4 Prospect theory

Tversky and Kahneman developed prospect theory in response to psychological research that uncovered how humans think when confronted with a problem containing a loss.

People will answer the same question differently if it is phrased in a different manner. This happens consistently when two identical risk-reward dilemmas are posed, but one question emphasizes the losses whilst the other emphasizes the gains.

For example, an experiment was carried out at Cornell university where undergraduate students traded coffee mugs under predetermined trading rules [16]. Half of the subjects were given a mug and they were asked what the minimum amount at which they would sell the mug. Those who did not have a mug were asked the maximum amount they would pay to buy the mug. The average selling price turned out to be more than the average buying price. This experiment was repeated with several different types of good to trade and the conclusions were the same - there is a discrepancy between the amount a person will accept for selling a product and the amount he will pay to buy the identical product. Although it is expected that a person should sell and buy at the same price this does not appear to happen in practice.

There is a strong link between psychology and investment decisions. Tversky and Kahneman put forward a solution to modelling people's behaviour towards risk, namely prospect theory. They propose the value function shown in Figure 4.2.

The origin is called the reference point, or current state of wealth. The function needs to be redefined each time wealth changes. As wealth increases it is common to become more risk accepting in the gains quadrant. This is because a unit of gain becomes less valuable with more wealth.

In the loss quadrant the curve is risk seeking, yet the curve is steeper for losses than gains. It places a drop in losses of $\mathcal{L}x$ to be a greater loss in value than the value gained if wealth increased by $\mathcal{L}x$ - this makes the curve

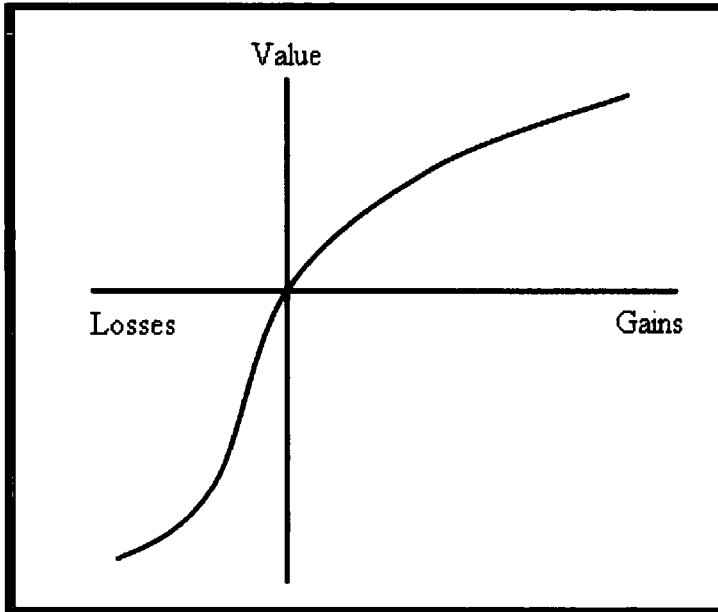


Figure 4.2: A value function

risk averse overall. “The aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount. [17]”

The derivation of the prospect function is more complicated than finding values to build the utility function. The loss of simplicity is owing to using decision weights. These weights measure the impact of events on the desirability of prospects, and not merely to the perceived likelihood of these events. Decision weights are not probabilities and hence they do not follow probability axioms and these weights are deduced from the choices between risky alternatives. The details are discussed in [17].

This theory suggests that traders are reluctant to sell at a loss owing to a perceived entitlement to the price they originally paid for it [16]. If a trader has lost money he is willing take greater risks to attempt to recover losses. If

the trader is faced with the same decision with trades at the moment making profit he will make a different choice.

4.8 A Java tool for utility elicitation

The elicitation process discovers the trader's risk attitudes, and hence provides the essential information for building the trader's utility function.

Several different designs for asking risk-related questions were presented to a bond trader and a financial market specialist. The main outcomes from these meetings were that the questions need to be quick to answer, clear and easy to understand, and be simple with no unnecessary and distracting information.

The initial design, which utilized this research, is shown in Figure 4.3. The trader clicks on either option A or option B (the option that he prefers) until he can no longer decide between them. The trader will stop clicking, and hence move to option C, when he reaches this indifference point.

Several problems were identified with the screen (Figure 4.3) from user feedback.

1. Many of the users found that the screens were non-intuitive and difficult to understand.
2. The user needed to click on an option endlessly to arrive at an indifference point.
3. In some cases, whole percentages did not allow enough precision for the user.
4. Many users would like to see a description of the risk behaviour - risk averse, risk neutral, or risk seeking.

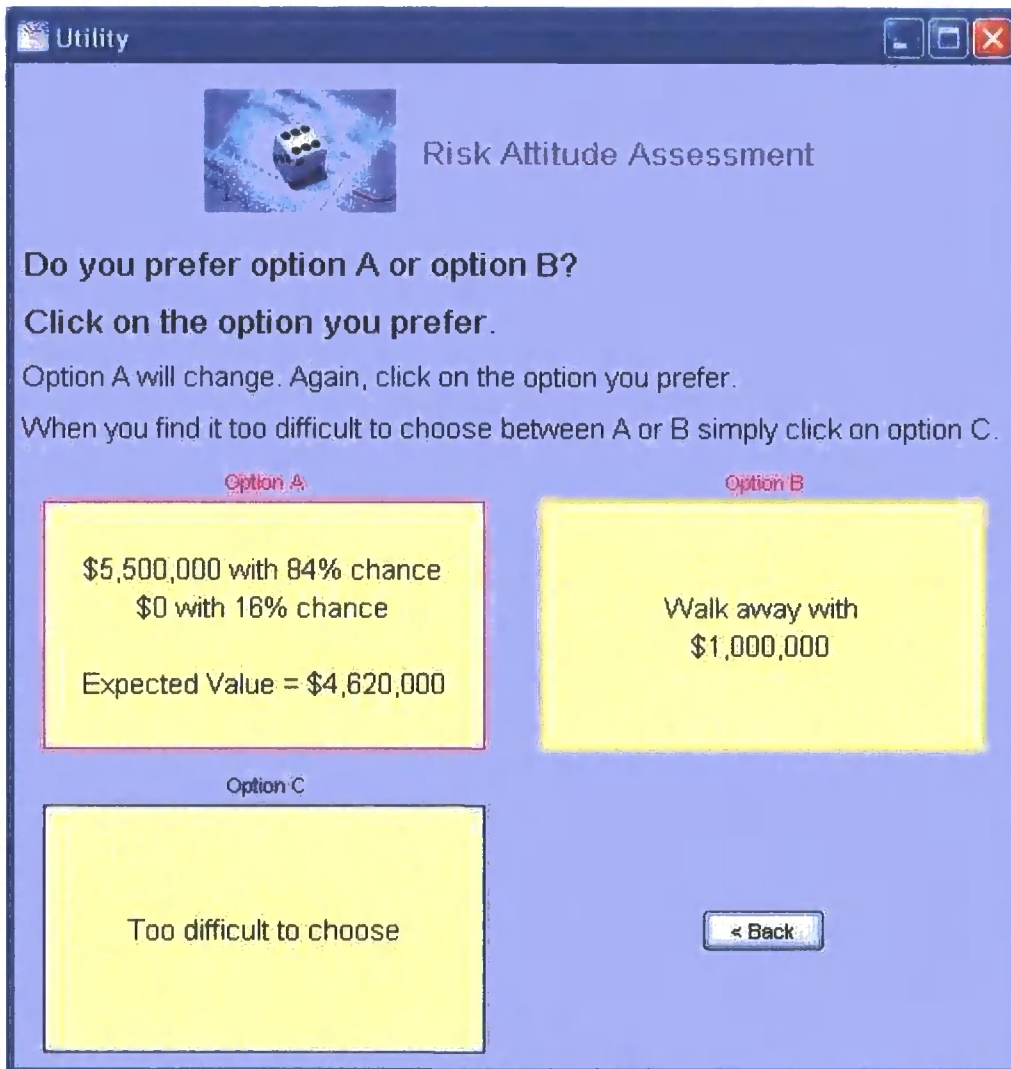


Figure 4.3: Java screen: Utility elicitation

5. The boxes with options are not normal action buttons, and this presented confusion.
6. The user would like a shortcut to risk neutrality.
7. Many users would like to be aided by way of pictures/diagrams - and interaction with the actual utility curve.
8. Users did not know how to proceed to the next screen.
9. After watching a number of users it was evident that the back button was in the wrong place, although nobody explicitly said that this was an annoyance.

A follow up design, shown in Figure 4.4, was created using this feedback. This screen solves several of the above problems by providing a slider to change the percentage chance on option A, and by allowing an extra decimal point in the percentage. It also provides a description of risk behaviour, action boxes for option A and B, back and next action buttons, and a button that resets the indifference point for a risk neutral investor. The new screen does not resolve the problem that people found the elicitation questions counter-intuitive, and would like to be aided by diagrams and pictures, or interact with the final curve that is built as a result of their answers. These difficulties are essentially in the area of behavioural psychology. It should be noted, that traders appear to be more at ease with these types of risk-reward questions than the Layman. Nevertheless, it is demonstrated from some of their incoherent answers that they do not fully understand or are uncomfortable with the approach used, so far, to elicit a utility function.

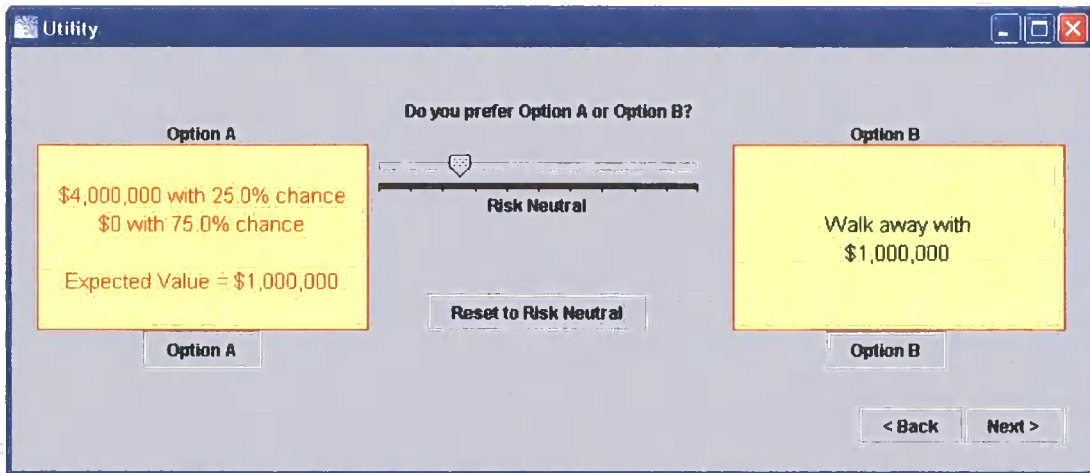
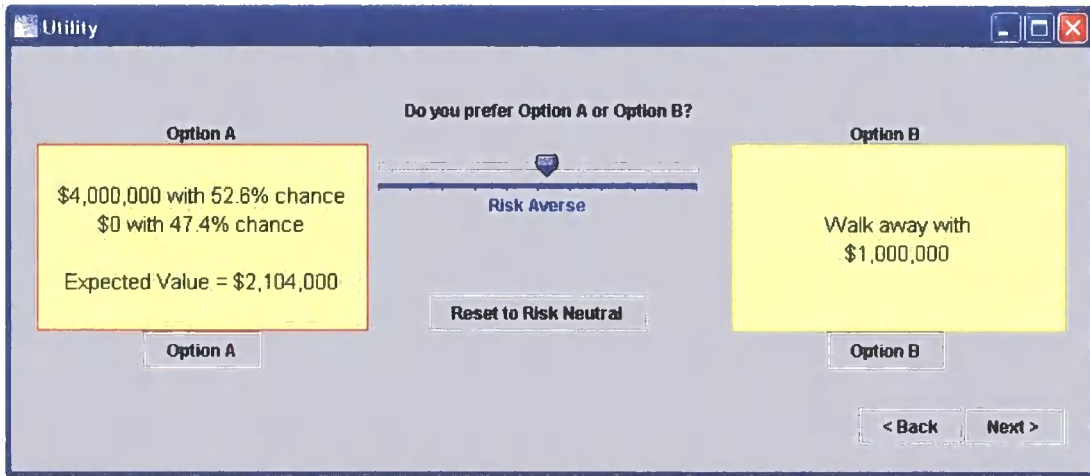


Figure 4.4: Java screens: Refined utility elicitation

4.9 Conclusions and recommendations

Utility theory is used to make decisions under uncertainty. In particular, an exponential utility function has been adopted to model a trader's risk attitude. With time it will become known if this approach to managing risk is suitable, and if not, a new solution will emerge.

The exponential utility is the best tool available, at the moment, to drive the decision support tool. Nevertheless, it can be criticized. First of all, a person with an exponential utility will invest the same amount of money in risky assets as his wealth increases. This type of behaviour is not readily observed in practice - most people will increase the amount of money they invest in risky assets as their wealth increases. Also, it is difficult to tune the risk factor of the utility function to define the trader's level of aversion.

To elicit a utility function it is essential to ask risk-reward type questions. It has been observed in practice that traders and people in general find these types of questions non-intuitive. This is a distinct disadvantage to using the utility approach since finding a person's certainty equivalent for a gamble is critical to defining the utility function. The Java tool for elicitation of utility is the best solution available at the moment. It will evolve and improve with time.

Also, the exponential function is very simple and lacks the actual detail involved in a person's real life behaviour when faced with decisions under risk. Utility is likely to be a function of several variables including money, time, and the investor's mood, confidence in price forecasts, and obligations.

For example, the trader behaves very differently when he has made enough profit to secure a bonus payment. He will not take risks that have potential to jeopardize this boost to his salary. For this reason, banks need to structure their incentive schemes carefully to avoid traders taking a back seat after

achieving some success.

Also, companies' risk profiles can change dramatically throughout the financial year. Consider that a company needs to publish end of year results at a set time each year. Traders can be told to do absolutely nothing for weeks as this company does not wish to blemish good year end results.

The investor's mood will have a large impact on their decision making. Positive, happy, and alert traders will trade more actively than cynical, stressed, and nervous traders.

Nevertheless, although traders have a risk attitude, they can be seen to take risks larger or smaller than they meant to simply because they are unaware of their exposure. Companies will take on more risk than they want to because they do not understand what their traders are doing and place too much trust in their traders' abilities.

The tool will minimize risk management errors with its systematic approach to a the decision problem by modelling risk averse behaviour. An approximation of a trader's risk attitude is better than none at all and forcing constraints on the utility function such as risk aversion is reasonable feature of the tool.

Chapter 5

Elicitation of prior information

5.1 Introduction

The aim of this chapter is to discuss the design of the prior elicitation tool, created to question expert traders, with the purpose of obtaining precise and accurate information about future commodity price movement.

The prior elicitation component of the software package will elicit the trader's expertise. The objective is to discover the trader's rational expectations or degrees of rational belief of price movement in the future, and model them.

The type of expertise, about a future time, that could be acquired include:

- **Expectation of price**
- **Interval estimates of price**
- **Extreme price limits**
- Any uncertain events that will cause future prices to change
- Probability of uncertain events occurring
- Expectation of peaks and troughs in the price series

- Confirmation that price will follow a random walk

The expertise highlighted in bold will be elicited by the prototype tool. The component will be designed to make it easy for the expert to enter his price outlook and to reduce biases developed from human cognitive thought processes.

When the prior information that the trader provides accurately reflects his opinion about price development the tool will be able to perform well. Here, performance is defined by exceeding the realizable profits that would be made by the trader without access to the tool.

If the prior is imprecise the model will either run the risk of unacceptably large losses, or lose the opportunity to make larger realizable profits. For example, if the trader is overconfident in his beliefs and therefore gives a narrow confidence interval, there will be a risk of unacceptably large losses [29]. If the trader is unsure about his beliefs, and hence provides a large confidence band, the tool will lose the opportunity to make a larger realizable profit [29].

As a result, the prior elicitation tool will have a large impact on the success of the software package. Therefore the design of the component must focus on eliciting an optimal amount of the genuine expertise available. That is, the ideal component will extract all of the information that an expert trader has about future price fluctuation without bias, and reflect this appropriately in his prior belief.

Perhaps the biggest challenge is ensuring that the tool copes with the traders specifications and interprets what they say correctly. Traders often use jargon to describe their beliefs - and terminologies used are not consistent across the world markets. For example, an expected downward trend could

be described as ‘a bear market’[15], or ‘the market is a dog’¹, or ‘looking topky’². It is essential to understand what these phrases mean in terms of building a prior distribution. Being precise about what the algorithm needs, and not confusing the interpretation of the information and data supplied by the expert is important.

The tool will perform checks to ensure that prior is consistent with rules defined in probability theory and the assumptions made within the model. The program will also include diagnostics that monitor the belief, and warn the expert of unusual forecast errors.

Many traders are likely to be less statistically naive than the Layman, and may already think in terms of probabilities. A fewer number will have formal training in statistical theory. Those who are learned in statistics will be more adept in providing a prior belief, but for a trader who has a poor understanding of probability, and what is required by the model in terms of estimating future prices, it is advised, as a first step, to give him basic training. Learning how to minimize common mistakes made when evaluating probabilities is part of this educational phase.

The Bayesian approach is favored over the Classical approach because the trader’s judgements about price development are personal and not necessarily based on sampling data. A Bayesian analysis can incorporate the subjective confidence that people have in a hypothesis about a single event and can be used to analyze the process by which probabilities and judgements of confidence are updated as new information arrives [1].

¹Observation by Klaas De Boer

²www.tradingfutures.biz/append4.html

5.2 Related work

5.2.1 Elicitation protocols

There is no universally accepted method for eliciting prior beliefs [1]. The method required to estimate reliable uncertainties on expert judgement remains elusive. There is little formal empirical evaluation of alternative approaches - a sparse amount of different approaches have been tried and tested and documented. Three recognized assessment protocols [19] [1] are :

1. The Stanford/SRI protocol
2. Morgan and Herion's protocol
3. The Wallsten/EPA protocol

The Stanford/SRI protocol and Morgan and Herion's protocol involve five main phases for probability elicitation. These are as follows [19] [1]:

1. Motivating the experts with the aims of the elicitation process.
2. Structuring the uncertain quartiles in an unambiguous way.
3. Conditioning the expert's judgement tool to avoid cognitive biases.
4. Encoding the probability distributions.
5. Verifying the consistency of the elicited distributions.

The Wallsten/EPA protocol [19] includes the preparation of a document that describes the objectives of the elicitation process, descriptions of cognitive heuristics and biases, and other related issues. The expert reads the documentation before the elicitation.

5.2.2 Existing software tools

Several attempts have been made to devise a software tool to elicit priors. These include:

- PROBES - a PRObability Elicitation System - “probing the complex minds of experts for their belief” written by Aik-Hiang Lau and Tze-Yun Leong [19].
- ELI: a program for the elicitation of uncertain knowledge written by van Lenthe [28].

5.2.3 BEEP

The BEEP (Bayesian Elicitation of Expert’s Probabilities)³, led by Tony O’Hagan, is working towards improving the way probabilistic beliefs are obtained from experts. Academics are pooling resources to develop the methodology to elicit priors.

5.3 Cognitive biases

It is well known to Bayesian statisticians that people do not generally cope well with Bayesian concepts, since statistical problems do not fit well with human cognitive thought processes [1]. Hence, it is important to understand human biases in probability judgments so that their effect can be mediated.

There is no existing elicitation method that fully eradicates human bias in a probability assessment.

³BEEP Home Page: www.shef.ac.uk/beep/



5.3.1 Individual bias

The biases that individual humans display when assessing probabilities are frequently owing to the following common heuristics:

Availability

The expert assesses probability by the ease at which the occurrence of the event is brought to mind [1]. Easily recalled events are assigned a higher likelihood of happening than events that are remembered with difficulty [3].

For example, a person is asked to assess the chances of their house being burgled. If friends or family have been burgled lately, or if they have seen anything about house robberies in the media in recent weeks, their evaluation of the probability is often too high.

Adjustment and anchoring

A person makes an initial assessment of probability then revises insufficiently in light of new information. The assessor's recalculation, with new data, is too close to his initial judgement because he assigns too much credibility to his first estimate [1].

To minimize the problem caused by this heuristic, the expert is asked to define the plausible range of the outcome before any central values [22]. If a middle values such as the median is elicited first, their assessment of the uncertainty level or spread tends to be too narrow. This is because the expert is anchoring too close to the central value. Also, the expert will be asked to carefully consider all information available, as a whole, before defining or redefining a belief.

Control

The expert is acting as if they can influence a situation over which they have no control [1]. For example, people will place more value on a lottery ticket with numbers that they have chosen than on a lottery ticket with a random selection of numbers even though the chances of winning are identical holding either ticket [1].

Representativeness

This is where the expert uses the similarity of two events to assess one's probability against the other [1]. By stereotyping one is demonstrating this type of bias. Often people assume (with too high a probability) that an individual belongs to a certain social group just by the way they look.

Over confidence

Over confidence results in elicited distributions that are too narrow [3] [1]. This type of bias is particularly unwanted in our elicitation tool because, in the situation where the distribution is too narrow, there is a risk of unacceptably large losses. Research shows that most individuals are observed to display overconfidence in their assessments and personal competencies [1] [22]. Nevertheless, if an expert is under confident in his assessment abilities, there will be a loss of opportunity to make larger realizable profits [29].

Conjunction fallacy

The probability axiom $P(A) \geq P(A \text{ and } B)$ is not adhered to. For example, inspect Kahneman and Tversky's problem which is mentioned in many prior elicitation papers:

Linda is 31, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and

social justice and also participated in antinuclear demonstrations. Is Linda more likely to be:

- a bank teller, or
- a bank teller and active in the feminist movement?

86% of people asked this question answered that there is more chance Linda is a bank teller and active in the feminist movement. This violates $P(A) \geq P(A \text{ and } B)$ - there is a higher probability that Linda is a bank teller alone than being both a bank teller and active in the feminist movement. This is common sense yet it evades many intelligent individuals.

It is necessary that the trader's prior belief is consistent with all probability axioms. The software package will disallow incoherent probability specifications and force the user to be rational with regards to complying with fundamental theory.

Hindsight bias

If the prior is assessed after seeing the data, the prior belief may be biased [3].

Base-rate neglect

Base-rate neglect is the inattention to the relative frequency with which events occur [1].

Gambler's fallacy

The belief that when a series of trials have all had the same outcome then the opposite outcome is more likely to occur in the following trials [1]. For example, at a roulette table, the outcome has been black ten times in a row. The punters falsely believe that there is more chance of a red than black in

the next game.

There are many additional heuristics used by humans to make judgements which are not mentioned here. For example, mood affects elicitation. A happy person tends to rely on pre-existing knowledge with little attention to precise details whereas a sad person will pay more attention to precise details than existing knowledge [1]. Also, the expert may demonstrate more than one bias complicating the problem even further.

To reduce errors in probability judgement the elicitation process will:

- use mental/visual imagery, such as representing information in a graphical or pictorial form.
- use consistent methods.
- use checks to identify errors such as the conjunction fallacy.
- elicit ranges first to minimize anchoring bias.
- educate the expert about cognitive biases in a bid to make his thought processes more rational.

5.3.2 Group bias

Group bias appears when attempting to elicit a probability from several experts. Mistakes and misjudgement made in the initial individual assessment of probability are compounded since they are communicated to the other experts [1].

There appears to be incentive to follow the crowd in group elicitation. Herding will result as a response to individuals' perception of their own ignorance. This herding is only rational if the other experts' judgements are

based on better information than their own [1]. Experts will incorporate other peoples' judgements into their own prior information set, and their posterior judgements may exhibit herding tendencies [1].

Stable outcomes will only be achieved if the herd can be led along a path of increasing the stock of common (real) knowledge [1]. In such cases, increases in the stock of reliable information will contribute to a convergence in posterior probabilities [1].

The main issue within the group decision problem is how to combine the probability estimates of several experts who have different background knowledge into a single probability distribution function. Two possible methods are to devise some way to join the individual estimates into one estimate, or to ask the experts to reach a consensus [3] [22] [1]. Joining the individual assessment, for example, taking an average, will lead to a probability density which is unequal to the density from a group consensus [1] [24].

Every trader has a different outlook on how price will move in the future, and a unique risk tolerance. This is shown by the wide variability in profit made by speculators across the world. Hence, traders are likely to disagree on their future price beliefs, and the maximum level of risk to accept. At this initial stage, the objective is to make the best decision for a single trader's risk attitude and expertise.

5.4 The elicitation process

5.4.1 Obtaining prior beliefs

The future price events in commodity markets are one-off uncertain events - price patterns do not repeat themselves systematically and conditions cannot be controlled - hence any predictions about the future are subjective. There are three main ways, that most traders combine, to calculate price forecasts

- fundamentals, technical analysis, and looking at the overall macroeconomic picture.

A fundamentalist believes that prices change in correspondence to supply and demand. They look at crop production, political news, weather forecasts, government programs, etc., when determining price beliefs. A fundamentalist will build a probability distribution of expected supply and demand of a commodity by estimating yields, or production rates, and consumer demand for next year. They will use this assessment as part of the process of constructing a price belief. If the fundamentalist expects the supply of cocoa to be less than the demand of cocoa, then he expects the price of cocoa to increase, and so on. There are many sources of uncertainty to be considered when estimating the effect of market forces on a commodity price.

A technician is a number cruncher and chart enthusiast concerned with time series analysis. His thoughts about price development are derived from the historical price data series rather than world events.

A macro-economist is concerned with the aggregate behaviour of the world economies. They have an interest in business cycles, unemployment rates, inflation, economic growth, and so on. They analyze government monetary and fiscal policies - studying their relationship with the relevant country's economic issues.

The causes of uncertainty that the trader believes will have the most impact on future prices, e.g., crop production, can be elicited from the trader and modelled using the software tool. These probability models of factors affecting prices can be used to aid the trader in the elicitation of his ultimate price beliefs. Eliciting beliefs about factors that affect future prices can be done in the same way that price beliefs are elicited.

5.4.2 Preliminary training of experts

Preliminary training of experts is recognized as an important part of elicitation. Training can improve experts' ability to provide judgements that are well calibrated [3]. Incorporating a computer aided tutorial into the elicitation component will be useful for this purpose. Perhaps this will entail an estimation of something like a country's population to familiarize the user with the process [22].

A document to educate and motivate the expert about the elicitation process, and its objective, can be produced for the use with the elicitation screens.

5.4.3 The Java prototype tool

Prototype screens were developed in Java to elicit prior information. These screens went through several iterations - even the type of information that needed to be obtained changed. This made the process arduous. It is unknown today what the final screens will look like in the market ready product. Nevertheless, the screens presented here are the latest screens developed for the prototype tool.

First, a volatility estimate is calculated from historical prices as per Chapter 3. These prices are uploaded by the trader from a .txt file or from an Excel file. Ideally these prices will be from a period immediately preceding the start of day one in the tool's trading period. Nevertheless, the trader can choose a period in history that they think will match future volatility more closely. For example, a user trading the Winter 06 Electricity for the month of May 06 may wish to use historical data from the Winter 05 Electricity prices in May 05 to estimate volatility.

Figure 5.1 shows the screen used to upload historical prices to estimate

volatility. The trader is also given the opportunity to revise the default estimate with his own subjective prediction.



Figure 5.1: Java screen: Volatility estimate from historical prices

Then, the extreme price limits are elicited. The extreme limits are where the trader is highly confident that price will not reach in the next n days. These limits are obtained at the start of the process to follow O'Hagan's advice that range should be elicited first to avoid over confidence bias [22]. The dialogue used to find the estimated extreme price limits is shown in Figure 5.2. In recent developments, the use of dynamic price limits is also being explored.

Figure 5.3, asks the trader how close price will need to be to the extreme price limits before he will begin to doubt the original price bounds for the

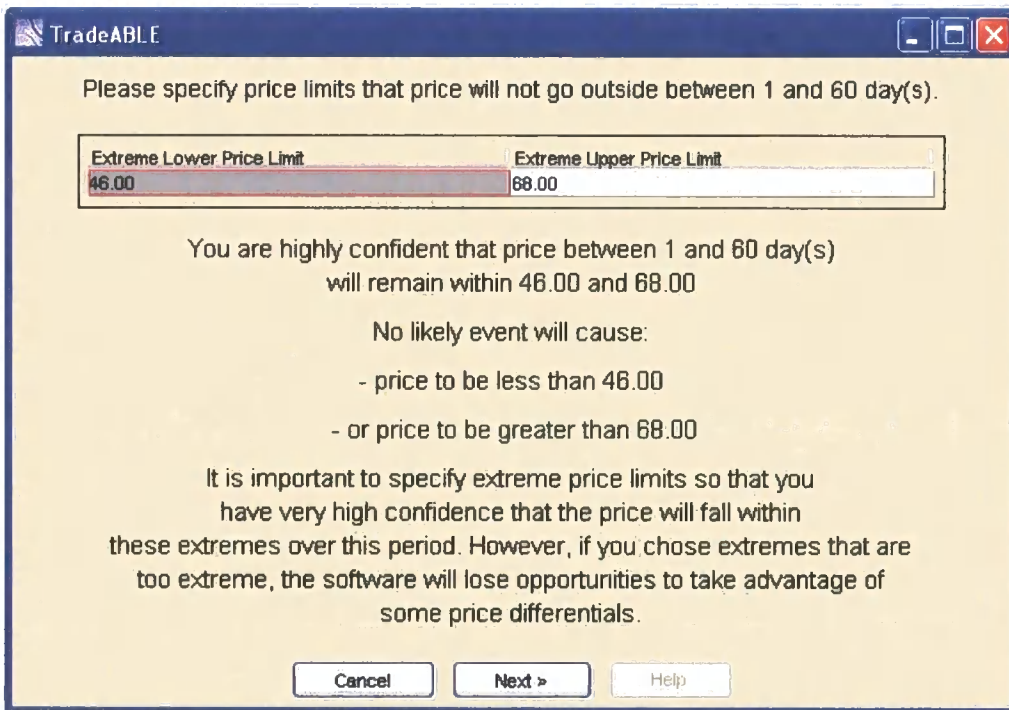


Figure 5.2: Java screen: Eliciting extreme price limits

rest of the period. This information is simply used to warn the trader that his price limits are being approached so he can refine his extreme price limits to avoid losing more than his maximum acceptable loss. This information is used internally by the Java application and is not utilized in the R algorithm.

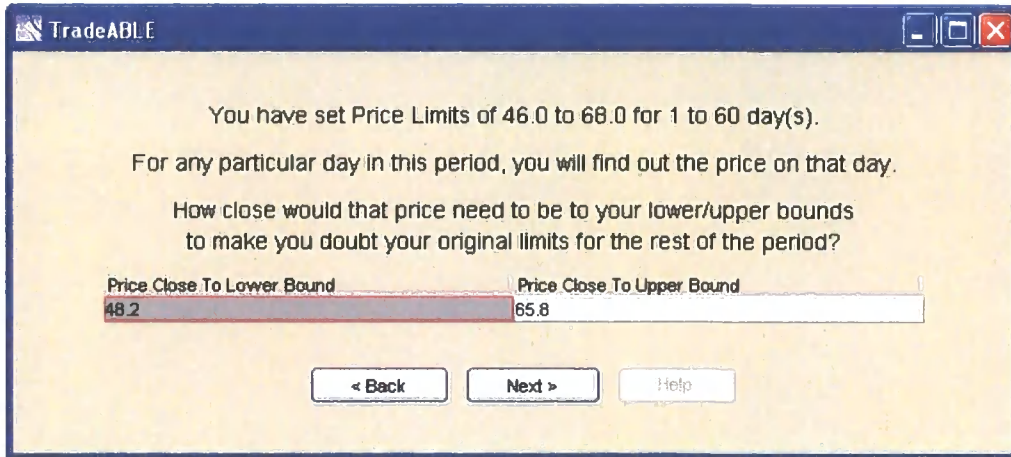


Figure 5.3: Java screen: Doubting extreme price limits

Now, the trader's expectation of prices are found. Figure 5.4 asks for the expectation of price on day one. The default value in the spinner box is the last price in the historical price data that the trader uploaded previously. Of course, this default price is likely to be completely unsuitable if the historical prices are not immediately preceding. The trader can modify $E_0[P_1]$ in the spinner box. The units of the listed prices are shown to remind the trader about the contract he is predicting prices for.

In Figure 5.5, the trader can split the period up into sections to give a more detailed prior. They can only split the period into equal sections, and up to a maximum of five sections.

In Figure 5.5, the trader is asked what they expect price to be at the end of the section. Now, since the estimated price at the end of the section




Figure 5.4: Java screen: Eliciting price at time one - P_1

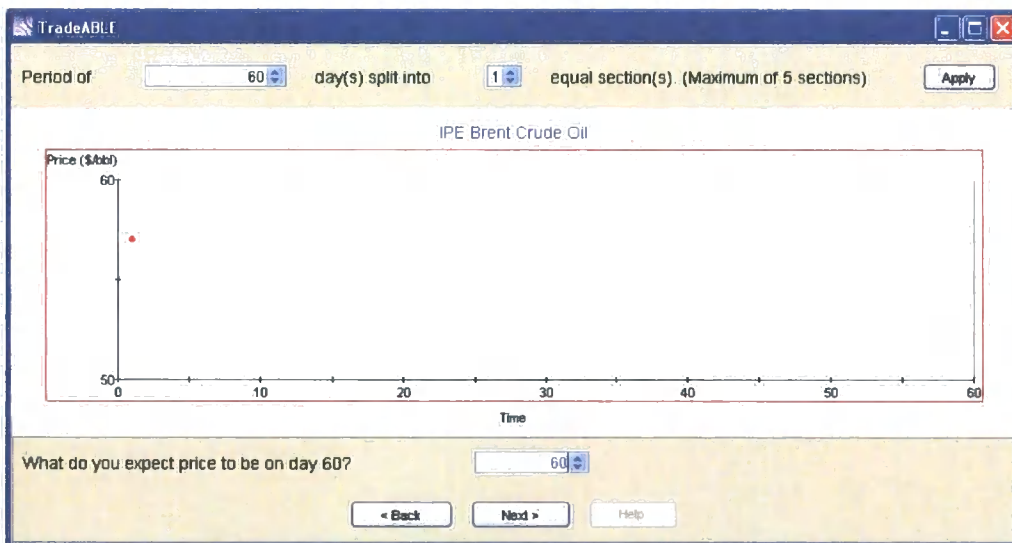


Figure 5.5: Java screen: Eliciting price at a end of section time - P_t

is $> E_0[P_1]$, the screen in Figure 5.6 shows three different upward trending curves. If the estimated price at the end of the section is $< E_0[P_1]$ then three different downward sloping curves will be displayed, and if the price at the end of the section $= E_0[P_1]$ Figure 5.6 will be skipped and a flat line will be drawn.

In Figure 5.6, the three upward trending curves are:

- Linear
- Exponential - increasing slowly at the start and quickly at the end.
- Exponential - increasing quickly at the start and slowly at the end.

These curves are expected to cover a trader's price outlook reasonably well. If not, they can give a more accurate prior by going back and increasing the number of sections in Figure 5.5 or ask Energy Scitech Ltd to support them in setting the tool up with a more detailed price outlook. The market ready tool could make this easier by providing table entry. However, it will take more time for a trader to enter his expertise in this format.

In Figure 5.6, the upward trending curve on the far right is outlined in yellow. When the cursor is hovering over one of the three boxes the yellow outline appears. When the mouse is clicked the curve is selected and a red outline appears as shown in Figure 5.7. Only one of the three curves can be selected (hence outlined in red) at one time and the default selection is the linear line.

Figure 5.8 shows the mean price curve elicited from the trader. Figure 5.9 shows a price curve with three sections - here the price expectation is a sharp rise, then slow rise, then flat, then a sharp rise at the end.

Figure 5.10 asks the trader to state how confident they are, on a scale of one to ten, about their estimate of the final price. To do this, the user




Figure 5.6: Java screen: Selecting a shape of the price curve

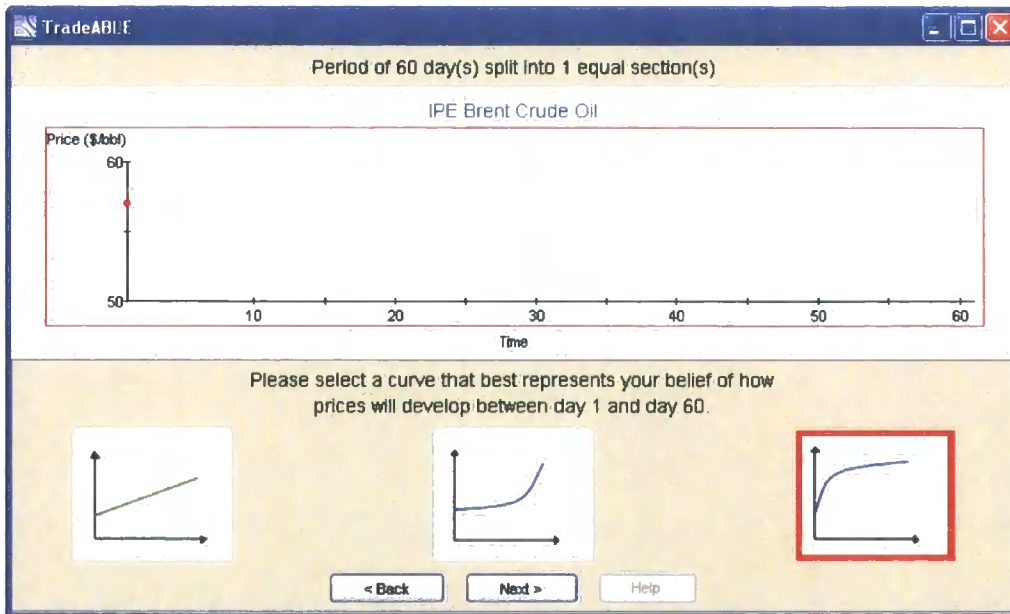


Figure 5.7: Java screen: Selecting the shape of the price curve

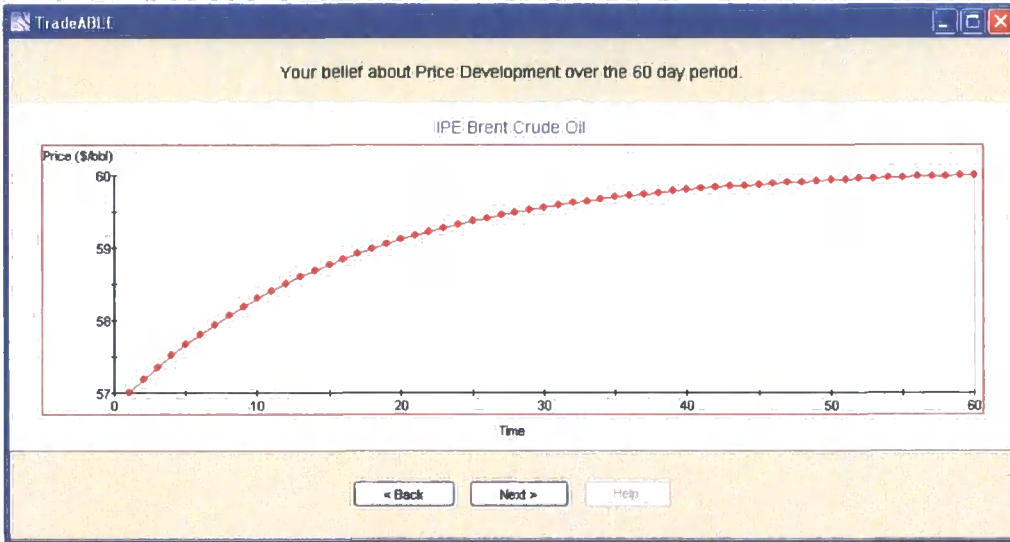


Figure 5.8: Java screen: Showing the selected shape of the price curve

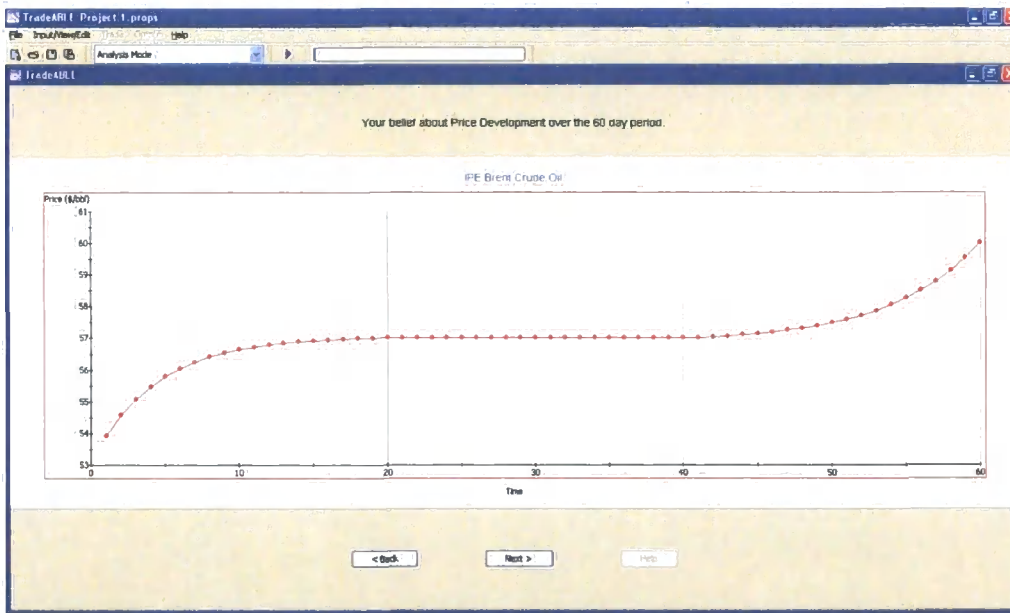


Figure 5.9: Java screen: Showing the selected shape of a more complex price curve

selects a number on the ruler. This information is used as per Chapter 3 to define $Var_W[P_n]$. This methodology is used to extract a confidence level such that the relationship between $Var_W[P_n]$ and σ^2 will suit the random walk model assumption. If the trader is unhappy with the resultant prior then it is possible that the trader does not assert that prices follow a random walk.

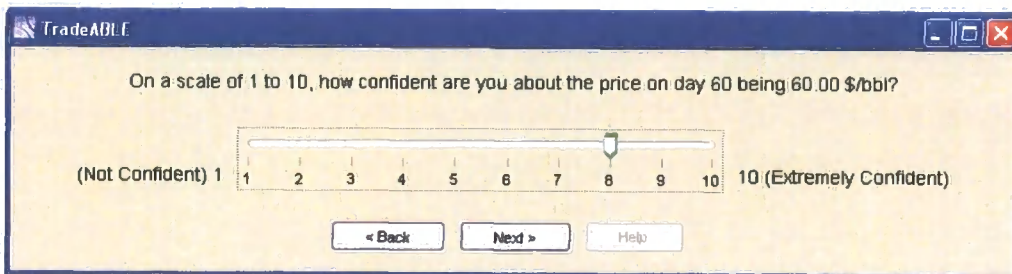


Figure 5.10: Java screen: Ruler - Level of confidence about final price

The expected final price and its corresponding expected standard deviation are given in a table on the screen shown in Figure 5.11. The 99%, 95%, 70%, and 50% interval estimates for the expected final price are given. These intervals assume a symmetrical Normal distribution $\sim N(E[P_n], SD[P_n])$. This is misleading for the user as no distributional assumptions are made in the approach used. It is also dangerous to allow the user to change $SD[P_n]$ as the trader may change it such that the random walk model assumption is invalid.

The summary of the price outlook is now presented to the trader. Here, the extreme price limits, the expected prices, and a 95% interval estimate are shown. The interval estimate is derived using the evolution of variance result explained in Chapter 3.

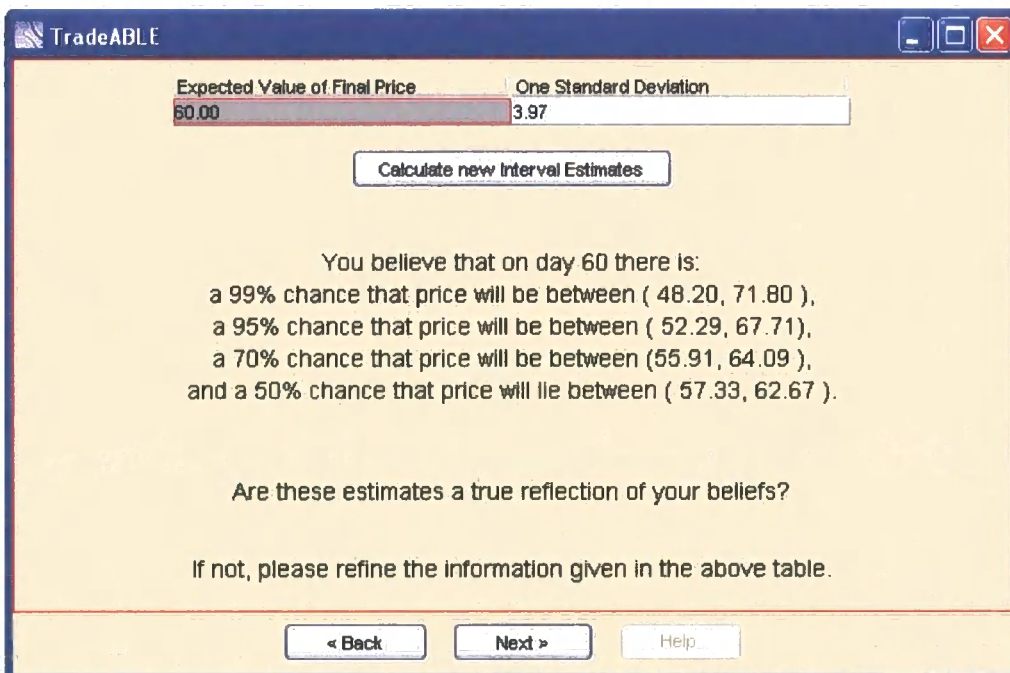


Figure 5.11: Java screen: A summary of interval estimates assuming the Normal distribution

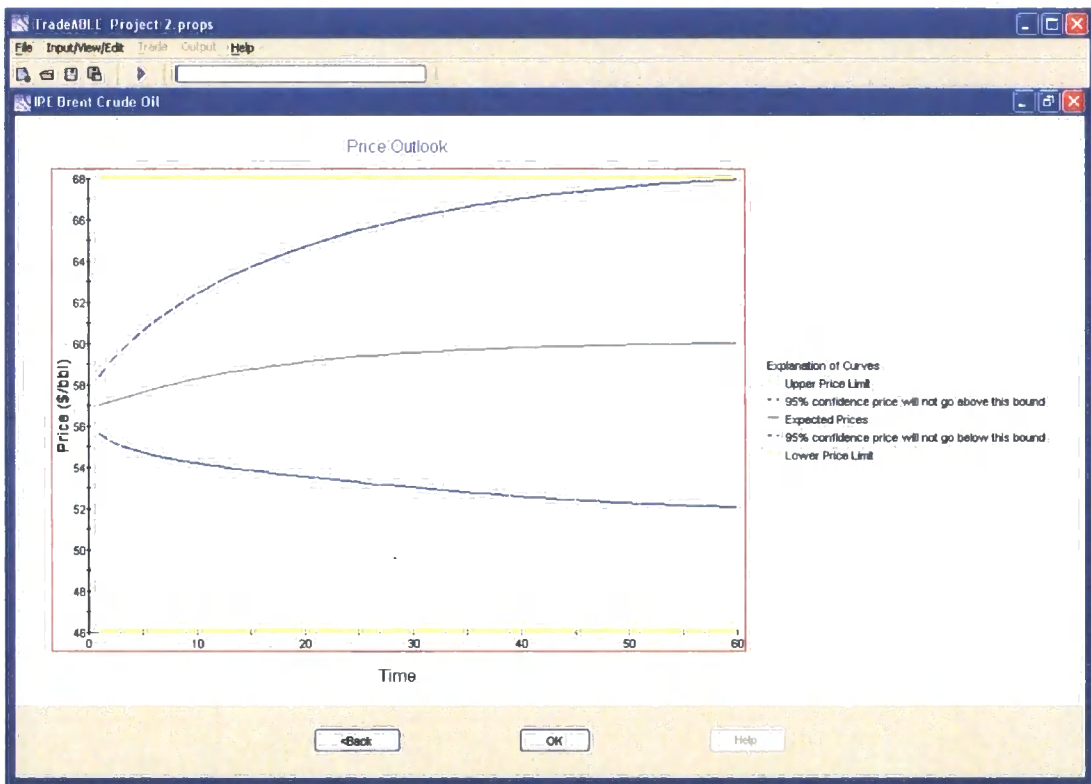


Figure 5.12: Java screen: Summary of the price outlook

5.5 Recommendations

The tool needs to obtain a precise price forecast from traders - who are typically stretched for time. This means that a prior needs to be elicited with a design that focuses on a minimal click approach. Does the tool described here suffice?

Of course, software development is ongoing and evolving. With the omission of Figure 5.2, 5.3, and 5.11 the elicitation routine becomes faster, but do the remaining screens provide the trader with the facility to give the detailed prior that is required?

This is the best solution available at the moment. Nevertheless, it is inevitable that the screens will improve with the lessons learnt to date. For example, the trader may be given the option to enter his price estimations in a tabular form, or interact with a chart to change price points.

User feedback from their interaction with prototype screens is the most important way to learn how to improve the software.

Chapter 6

Diagnostics

The main purpose of the diagnostics is to compare the forecast of prices with the observed market prices. The aim is to monitor the accuracy of the forecast and to flag the forecaster when observations occur that are unlikely under the forecast model.

If a flag is issued, the detail of the likely deficiency in the forecast will be explained. That is, the reason for the diagnostics indicating that an unusual observation has been recorded will be provided. The expert will want to know where this supposed deficiency in the model lies, for example, always overestimating prices. If an expert consistently has the same type of error in his forecasting model this monitoring process will prove extremely useful since the expert will be able to learn how to improve his forecasting technique.

Not all signals from the diagnostics will lead to flaws in the forecast model. Prices that have a small probability of occurring will happen. It is the forecasters responsibility to decide if he believes his forecast is still adequate by investigating the reason for the diagnostic alert.

At any time, the expert can continue with the Bayes forecast model or intervene by re-specifying the forecast.

In this chapter, various diagnostics that are used in the decision support

tool are discussed with an example from a fictional trial.

6.1 The size ratio

A size ratio is a standard Bayes linear technique for looking at the difference between the expected data and the observed data. Here, the expected data have been revised using Bayes linear methods to incorporate all information available including past prices.

If a price X has expectation $E[X]$ and variance $V[X]$, and X is eventually observed as x , then a standardized measure of discrepancy between the observation and its prior variance is given by [29] [13]:

$$SizeRatio = \frac{(x - E[X])^2}{Var[X]}$$

$(x - E[X])^2$ - the square of the difference between observed price and its forecast (made at a time before X is observed) - gives an indication of the distance of the forecast from the actual price.

To arrive at a meaningful comparison of these differences they are normalized by dividing by $Var[X]$. For example, a change of around 90 cents¹ in price in a day is common in the oil markets, but a daily change as large as 90 cents in the sugar market is very unlikely. Daily price fluctuations in the sugar market are typically less than 15 cents². In the oil market there is more tolerance for forecasting in terms of magnitude of prices than in the sugar market where fluctuations are relatively smaller. Hence, dividing by $Var[X]$ ensures that the diagnostic is standardized across all price series whether their typical volatilities are small or large.

The size ratio will, in theory, have an average value of one. When the

¹Based on 2004 Brent Oil Prices

²Based on 2004 New York Sugar No 11 Prices

size ratio calculation is close to one this means that the prediction of X is within realistic bounds of uncertainty.

When the size ratio records a value that is large this will concern the forecaster. Large size ratios are a warning that the forecasted price is very different to the latest observed price. There are three situations where a large diagnostic value will be recorded:

1. The prior estimate $Var[X]$ is too small. The observed volatility of the price series is higher than expected.
2. The forecast is unrealistic in comparison to the $Var[X]$ specified - or the distance from the observed price is large.
3. Owing to chance. It is expected that some discrepancies will be recorded.

A size ratio of around 7 implies an unusual actual price in relation to the forecast. This corresponds to an approximate 95% confidence threshold. A size ratio of around 14 implies a very unusual actual price in relation to the forecast. This corresponds to an approximate 99% confidence threshold. It is advised that the forecaster should begin to worry when the size ratio is larger than 3.5.

When the size ratio records a value that is less than one, the forecaster will not be concerned if this is a result of his forecast for price being very close to the actual price. Nevertheless, it is also possible for a small size ratio to indicate that $Var[X]$ is larger than the observed variance of X . In this case, a forecast that is assessed to be poor (with hindsight) will not be flagged via a large size ratio diagnostic. Also, excessively high estimates of $Var[X]$ will mean that there will be opportunities lost by the tool to make more profit.

6.2 Gamma-probability

The Gamma diagnostic is an approximate probability diagnostic for the size ratio based on assuming a Normal distribution for the standardized forecast errors. This diagnostic is more sensitive than the size ratio. However, it depends on the following assumptions [29]:

1. Prices are Normally distributed.
2. Approximating the size ratio to a Gamma distribution with the same mean and variance.

The diagnostic calculates one minus the probability of recording a given size ratio if the specified expertise is correct. This also relies upon the above assumptions which may not be satisfied.

The Gamma indicator is a probability and hence lies between zero and one. A value of 0.95 or 0.99 correspond to formal hypothesis tests at the 5% and 1% level respectively.

6.3 Discrepancy

The discrepancy diagnostic compares the distance between actual prices and their prior expectation relative to prior variation. This indicator checks if the observed prices, so far, match the prior expertise that was previously specified about them. The measure used is the Mahalanobis distance between the price data and their expected quantities.

Let us define D to be a vector of random quantities such that D is $(P_1 P_2 \dots P_t)$ and d to be a vector of observed past prices such that d is $(p_1 p_2 \dots p_t)$. If $E[D]$ and $Var[D]$ are known, then the Mahalanobis distance

(which is a single value) is calculated as [29] [13]:

$$Dis[d] = [d - E[D]]^T Var[D]^{-1} [d - E[D]]$$

where $Var[D]^{-1}$ is the standard inverse of $Var[D]$ where $Var[D]$ is a non-singular matrix.

The distance $Dis[d]$ has prior expectation $r = rank(Var[D])$ - for practical purposes rank, r , is the dimension of $Var[D]$ here. Then if an assumption is made that $(P_1 P_2 \dots P_t)$ are multivariate normal such that $Dis[D]$ has a chi-squared distribution with r degrees of freedom. The discrepancy indicator is finally defined as a probability [29] [13]:

$$P(Dis[D] \leq Dis[d] \mid Dis[D] \sim \chi_r^2)$$

This probability is the chance that $Dis[D]$ is less than $Dis[d]$ assuming random fluctuation. A large probability indicates that $Dis(d)$ is unlikely if the prior is precise, hence the expert will need to be concerned that his forecast is inconsistent with incoming data.

Thresholds of 0.95 and 0.99 correspond to a formal statistical test that the data received so far match with the stated prior expertise, at the 5% and 1% significance levels respectively. The 5% level implies there is a 1 in 20 chance that the hypothesis of no difference is rejected when it is true; a 1% test implying a 1 in 100 chance the hypothesis of no difference is rejected when it is true.

6.4 Detection of systematic error

This diagnostic identifies systematic discrepancy in the same direction. This is to find whether observed prices are consistently above or below the prior specification. First, the standardized one step forecast errors are denoted $\epsilon_2, \dots, \epsilon_n$ where [29] [13]:

$$\epsilon_i = \frac{p_i - E_{i-1}[P_i]}{\sqrt{\text{Var}_{i-1}[P_i]}}$$

$E_{i-1}[P_i]$ and $\text{Var}_{i-1}[P_i]$ are the forecasted expectation and variance for P_i (assessed at time $i - 1$) which have been updated using Bayes linear methods to include information about observed prices p_1, \dots, p_{i-1} and any other external information.

The standardized cumulative sum of the standardized one step forecast errors is defined as [29] [13]:

$$Q_t = \frac{1}{\sqrt{t-1}} \sum_{i=2}^t \epsilon_i, \quad t = 2, 3, \dots$$

Under some weak assumptions (and when the expertise stated is correct) $Q_t \sim N(0, 1)$. This means that the sequence Q_2, Q_3, \dots is expected to wander around zero. If a value of Q_t reaches or passes a threshold, e.g., 2 (representing 2 standard deviations), then it is possible that there are systematic, one-sided differences between the forecasts and the prices.

For example, it is expected that only around 5% of the Q_t values are outside the range $(-2, 2)$ given that the expertise is precise. Hence, if a Q_t is recorded outwith $(-2, 2)$ it is advised that the expert reviews his forecasts to check if there is a bias on one side of the observed prices.

Also, if the sequence of $Q_2, Q_3 \dots$ is consistently above or below zero there is a concern that Q_t is not following the $N(0, 1)$ distribution and in this situation it is recommended that the expert checks over his prior beliefs for forecasts which are persistently higher or lower than observed prices.

The choice of variance estimate, $\tilde{\sigma}^2$, is based on first differences, and this method was chosen to minimize the influence of a systematic bias in the forecast. This is discussed further at the end of section 6.6.

6.5 Comparison of results using different utility functions

Here, the difference in trading decisions and corresponding expected profit results are compared between using a risk neutral utility function and the trader's utility function which is typically risk averse. The risk averse utility is modelled by the exponential curve, $1 - \exp(\frac{-w}{R})$, which was explained in an earlier chapter. In the special case where the trader chooses to use a risk neutral utility this diagnostic is obsolete. However, it is possible that the risk neutral trader will wish to compare the risk neutral decision making with a more conservative risk averse decision mechanism.

The results from this sensitivity analysis include the consideration of the actual position (number of lots the trader is holding), actual balance, and the corresponding realizable profit. The outputs are:

1. The difference between the lot sizes recommended on day i for two utility functions.
2. The difference between the predicted expected profit at day n (the last day in the trading period) for the two utility curves.
3. The differences in lot strategies over the remaining period for the two utility functions.

When the risk factor, R , in the exponential function is large there will be little difference in the results since the portion of the utility curve in implementation is very close to a risk neutral utility. There will be a difference if the exponential utility models a risk averse trader properly by using a small risk factor.

The risk averse utility curve will trade more conservatively and have a smaller expected profit on day n than the risk neutral utility curve. Then,

the expert can look at how much profit is missed out on if the expertise is precise by taking a conservative approach. The expected variance of profit on day n for the risk neutral curve is wider than the risk averse curve - taking more risk by using the risk neutral curve increases the chances of a large loss.

6.6 Bayes linear variance learning

A prior variance, σ^2 , is assessed by the expert to form part of their price belief. This is a variance for a simple random walk, such that,

$$P_t = P_{t-1} + Z_t,$$

which has no trend.

A variance update is for learning about the volatility of prices via a Bayes linear technique. There is a method for revising the estimate of a volatility and this revision is fed back to the expert. The new volatility is not fed back into the algorithm automatically - the expert will need to change his inputs himself if he feels that his variance estimate needs to be revised.

To learn about variance Bayes linear methods are used. A simple way is to find an unbiased estimator of variance denoted $\tilde{\sigma}^2$ from a sample of n prices and to combine this with the prior variance σ^2 . The updated variance denoted $\hat{\sigma}^2$ is defined as follows [13]:

$$\hat{\sigma}^2 = \frac{m}{m+n}\sigma^2 + \frac{n}{m+n}\tilde{\sigma}^2$$

m is a notional sample size used to place a value on the prior information σ^2 . If m is a larger number than n then more weight is assigned to the prior belief than to the estimator $\tilde{\sigma}^2$. This is similar to Credibility Theory which is used by actuaries to assess risks in the general insurance industry. In the algorithm $m = 4$ - this is the smallest practical value for the theory to hold.

When five prices arrive more weight is given to $\tilde{\sigma}^2$ than to σ^2 and this weight increases with additional incoming prices such that the effect of σ^2 on $\hat{\sigma}^2$ becomes negligible with time.

$$\text{As } n \rightarrow \infty, \quad \hat{\sigma}^2 \rightarrow \tilde{\sigma}^2$$

A small value of m is regarded as appropriate since there is an interest in quickly detecting whether the volatility of incoming prices is changing in comparison to σ^2 .

To apply the above method a value for $\tilde{\sigma}^2$ is required. Now suppose that the price on day i follows a random walk such that [29],

$$P_i = \mu_i + \gamma + Z_i,$$

where μ_i is the expert's forecasted price for day i , Z_i represents the random fluctuation with $E[Z_i]=0$ and $Var[Z_i] = \sigma^2$, and γ represents possible systematic bias. γ will be zero unless the forecasted prices are systematically on one side of the observed prices.

The estimate, $\tilde{\sigma}^2$, is then found using first differences. Let $Q_i = P_i - \mu_i (= \gamma + Z_i)$ and let $V_i = Q_{i+1} - Q_i$ for $i = 1, \dots, n - 1$. Then $E[V_i] = 0$ and $V[V_i] = 2\sigma^2$ ³. The variance estimate is [29]:

$$\tilde{\sigma}^2 = \frac{1}{2(n-1)} \sum_{i=1}^{n-1} v_i^2$$

This estimate is robust in the situation where there is a undetected constant systematic bias γ .

Confidence intervals for $\tilde{\sigma}^2$ can be calculated. However, these are often wide and unreliable, and hence they are not useful in practice.

³ $V_i = \gamma + Z_{i+1} - (\gamma + Z_i) = Z_{i+1} - Z_i$. $E[V_i] = E[Z_{i+1}] - E[Z_i] = 0$. $V[V_i] = V[Z_{i+1}] + V[Z_i] = 2\sigma^2$.

6.7 Example

This example of diagnostic output is taken from a fictional prior belief about CSCE (Coffee Sugar Cocoa Exchange) coffee prices from 03/01/2005 to 30/03/2005. The price belief with a two standard deviation confidence envelope and actual prices (the blue line) for the example are shown in Figure 6.1. Figures that plot the diagnostics follow with an explanation of interpretation and Table 6.1 with the corresponding numerical data is also provided. In Table 6.1, alarming diagnostic outputs are indicated by bold values.




Figure 6.1: The price belief with actual prices

A large size ratio is recorded if the observed price is far away from its forecast relative to the forecasted variance. The first alarming size ratio is recorded at time three. With further investigation the listed price observed at time three was \$98.5 and the forecast, $E_2[P_3]$, was \$101.92. There is an error here of magnitude \$3.42 which is large considering that the standard

deviation of daily price change is estimated at \$1.30. Price at time two was \$101.45, and if a random walk with no trend is applied, then daily price change $\sim N(0, 1.3^2)$. Then an observed price at time three less than or equal to \$98.5 has a probability of 0.02⁴. This means an unlikely price (odds of 1 in 50) has been observed if the prior belief is correct. Any size ratio larger than 3.5 must be investigated with relation to the forecasting model and there are several alarming recordings in this example - see Figure 6.2.

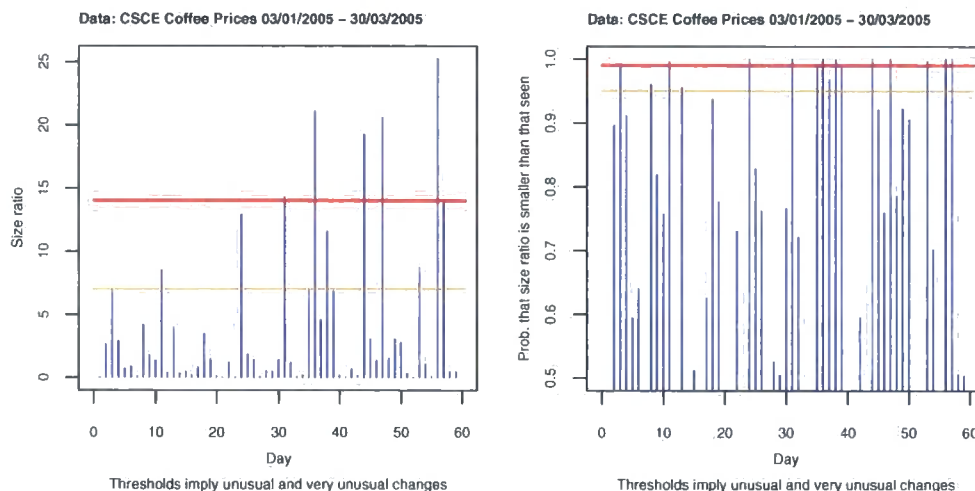


Figure 6.2: The size ratio and Gamma-probability diagnostics

The Gamma-probability diagnostic is derived from the size ratio, hence they both flag the same potential problem with the forecasting model. The Gamma output is shown in Figure 6.2. Alarming probabilities happen at the same time as thresholds are passed in the Size Ratio diagnostic.

The Discrepancy diagnostic records a value of 0.98 at time three. This

⁴ $P(\text{Price} \leq (98.5 - 101.45) / 1.3) = 0.02$




Figure 6.3: The Discrepancy diagnostic

indicates that the Mahalanobis distance $Dis[d]$ is large such that there is 98% chance $Dis[D]$ is less than $Dis[d]$ if forecasting is precise. This means the $Dis[d]$ recorded is unlikely if the forecasts provided are correct.

The Discrepancy indicator takes all past observed prices and past forecasts into consideration and this is the reason for the Discrepancy output remaining above alarming threshold values throughout the process. The observed prices dipped below the two standard deviation confidence envelope at time three and here lies evidence that the prior belief is inconsistent with the actual prices. The Discrepancy plot is shown as Figure 6.3.

The systematic error also flags at time three. The value of -3.01 indicates that there is a huge bias. The forecast for price, \$101.92, is a lot higher than the observed price, \$98.5, given the estimate of variance. This disturbing over bias on the top side of observed price continues until time eight and

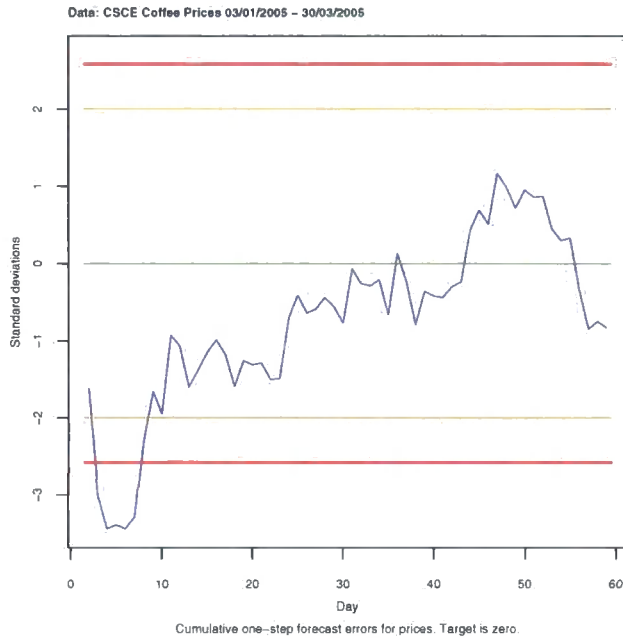


Figure 6.4: Detection of systematic error

the systematic error remains negative until time forty three. It is clear that this diagnostic is not fluctuating around zero as desired. This will present a point of concern to the forecaster and he should check that this bias - to overestimate on one side of prices - is not a habit inherent in his forecasting technique. The graph presenting the detection of systematic error is shown as Figure 6.4.

It appears from the variance learning diagnostic that the prior estimate, given by the flat line, of variance was too small. The volatility in the observed prices was higher than expected. The prior standard deviation estimate for daily price change was \$488.90 in price per lot which corresponds to a standard deviation of \$1.30 in terms of the listed coffee price. However, experienced volatility was over \$600 per lot (\$1.60 in terms of listed price). If a random walk with a normally distributed step is assumed this corresponds

Data: CSCE Coffee Prices 03/01/2005 – 30/03/2005

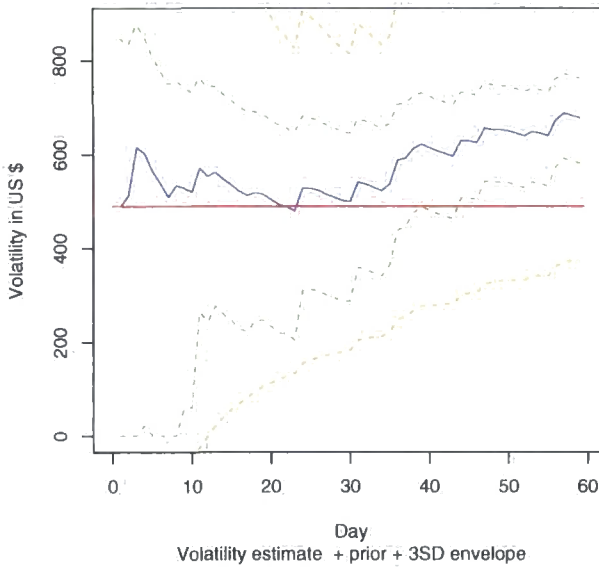


Figure 6.5: Variance learning

to 70% chance the daily price change was $\pm\$1.60$ rather than 70% chance the daily price change was $\pm\$1.30$.

The estimate of variance influences all other diagnostic readings. In this example the forecasted variance provided proved to be too small. The observed prices may have been a realistic distance from expected prices if the volatility estimate was more closely in line with the volatility of the observed prices. Variance is very difficult to estimate and hence a close watch will be kept on the variance learning to combine this with the other diagnostic readings. To reaffirm, if the variance estimate is too small, it is more likely that diagnostics will pass thresholds quickly. However, if the variance estimate is too large the diagnostics will present no alarm signals since a large distance between the observed price and its corresponding forecast will be expected. The variance learning plot is shown as Figure 6.5.

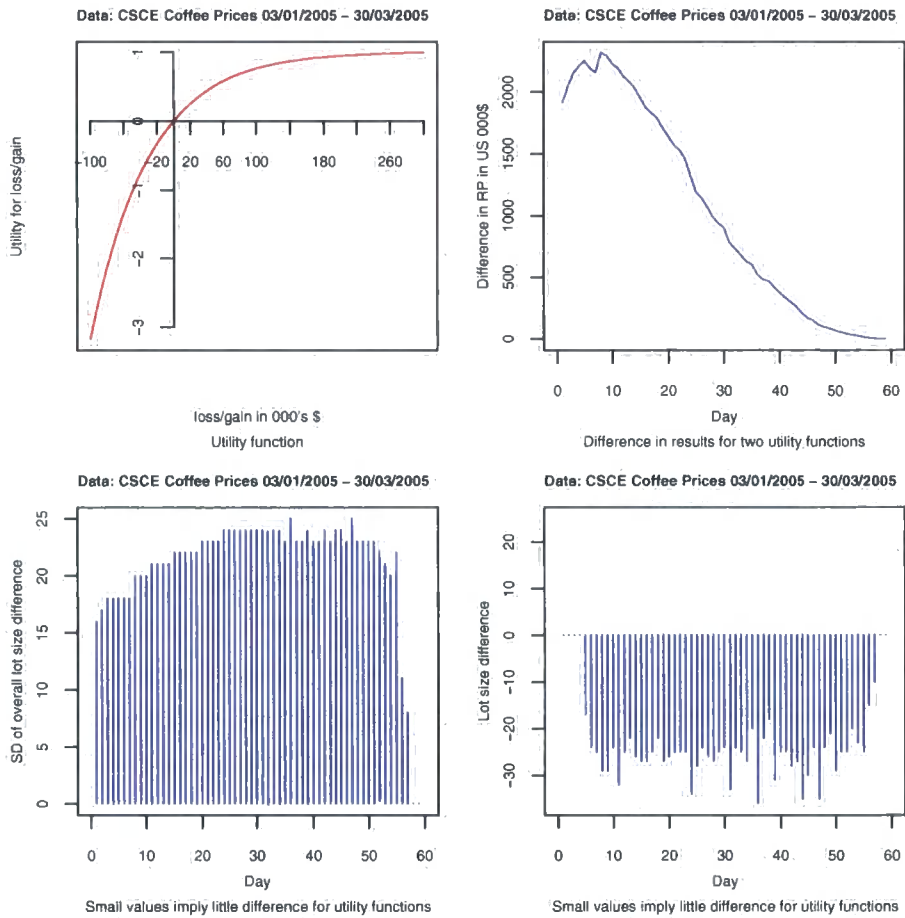


Figure 6.6: Sensitivity to the utility function

There is a big difference in the lot buy and sell strategy if a risk neutral utility function is used instead of the risk averse exponential utility function. The risk neutral investor takes on additional risk by holding a larger position than the risk averse investor. The top right-hand graph in Figure 6.6 shows the opportunity for profit that is lost - if the price forecasting is precise - by being more conservative than the risk neutral curve. The top left-hand chart in Figure 6.6 shows the exponential utility curve and the bottom right-hand chart shows the difference in lot decisions between the two utility curves.

Table 6.1: CSCE coffee example:diagnostic results

Time	Listed Prices	$E_t(P_{t+1})$	Size Ratio	Gamma	Discrepancy	Systematic Error	Variance Learning
1	103.1	103.56	0.00	0.000	0.00	0.00	488.9
2	101.45	101.92	2.63	0.896	0.73	-1.62	513.1
3	98.5	99.00	6.95	0.992	0.98	-3.01	615.0
4	96.8	97.32	2.88	0.911	0.99	-3.44	601.6
5	96.25	96.78	0.68	0.594	0.98	-3.39	564.0
6	95.6	96.14	0.83	0.640	0.97	-3.44	537.6
7	95.65	96.20	0.14	0.299	0.95	-3.29	509.0
8	98.85	99.37	4.18	0.960	0.98	-2.28	533.6
9	101.1	101.61	1.77	0.819	0.98	-1.66	528.6
10	100.1	100.62	1.35	0.757	0.98	-1.95	518.6
11	104.4	104.89	8.48	0.996	1.00	-0.93	570.4
12	104.1	104.59	0.37	0.459	1.00	-1.07	552.9
13	102	102.52	3.99	0.955	1.00	-1.60	561.6
14	103.25	103.76	0.32	0.432	1.00	-1.38	546.9
15	104.65	105.15	0.47	0.511	1.00	-1.15	534.5
16	105.75	106.25	0.21	0.359	1.00	-0.99	521.5
17	105.1	105.61	0.78	0.626	1.00	-1.18	512.4
18	103.2	103.73	3.44	0.937	1.00	-1.59	518.3
19	105.3	105.82	1.46	0.776	1.00	-1.26	514.8
20	105.35	105.87	0.13	0.284	1.00	-1.31	503.9
21	105.8	106.32	0.00	0.045	0.99	-1.29	493.2
22	104.9	105.44	1.20	0.730	0.99	-1.50	488.5
23	105.3	105.84	0.01	0.085	0.99	-1.49	479.0
24	110.5	110.99	12.93	1.000	1.00	-0.71	529.0
25	112.75	113.22	1.85	0.828	1.00	-0.41	527.2
26	111.7	112.19	1.37	0.762	1.00	-0.64	523.2
27	112.45	112.93	0.04	0.165	1.00	-0.59	514.6
28	113.85	114.32	0.50	0.525	1.00	-0.44	508.2
29	113.45	113.93	0.45	0.504	1.00	-0.56	501.8
30	112.4	112.90	1.40	0.766	1.00	-0.77	499.0
31	117.8	118.24	14.29	1.000	1.00	-0.07	540.7
32	116.85	117.31	1.16	0.721	1.00	-0.26	536.6
33	117.05	117.51	0.04	0.163	1.00	-0.29	529.2
34	118.05	118.51	0.17	0.326	1.00	-0.21	522.5
35	115.1	115.60	6.91	0.992	1.00	-0.66	536.0
36	121.55	121.98	21.09	1.000	1.00	0.13	587.3
37	119.2	119.66	4.59	0.968	1.00	-0.23	591.7
38	115.25	115.76	11.58	0.999	1.00	-0.79	612.3
39	119.15	119.62	6.83	0.991	1.00	-0.36	621.2
40	119.05	119.53	0.19	0.346	1.00	-0.42	614.3
41	119.35	119.83	0.02	0.113	1.00	-0.44	607.3
42	120.9	121.37	0.68	0.595	1.00	-0.30	602.1
43	121.85	122.31	0.14	0.296	1.00	-0.24	595.8
44	128	128.39	19.29	1.000	1.00	0.43	629.3
45	130.65	131.01	3.05	0.921	1.00	0.69	628.3
46	129.5	129.88	1.35	0.759	1.00	0.51	624.8
47	135.75	136.05	20.57	1.000	1.00	1.17	655.6
48	134.45	134.77	1.52	0.786	1.00	0.98	652.4
49	132.5	132.85	3.07	0.922	1.00	0.72	652.1
50	135	135.32	2.76	0.905	1.00	0.95	650.2
51	134.6	134.93	0.31	0.428	1.00	0.86	644.9
52	135.1	135.43	0.02	0.108	1.00	0.87	639.0
53	131.6	131.98	8.75	0.997	1.00	0.45	648.6
54	130.65	131.05	1.06	0.702	1.00	0.30	644.8
55	131.35	131.75	0.05	0.186	1.00	0.33	639.3
56	125.25	125.75	25.24	1.000	1.00	-0.35	673.0
57	120.9	121.48	14.07	1.000	1.00	-0.85	687.4
58	122.35	122.91	0.45	0.506	1.00	-0.75	682.6
59	122.05	122.63	0.45	0.503	1.00	-0.83	677.5
60	125	NA	NA	NA	NA	NA	NA

Chapter 7

Interactions and piloting with traders

Several pilots of the algorithm and the prototype tool were undertaken with real commodity traders from a variety of sectors. Overall, the feedback was positive, even in situations where a loss was realized owing to prior beliefs which turned out not to be consistent with the data. This was because the tool made substantial profits when the prior was accurate and showed the ability to control risk in sensible way. The diagnostics warned the traders from an early stage if their beliefs did not match with observed market prices. Two of these trials are discussed in the following sections.

To build a successful software decision tool for traders it is necessary to interact with traders to find out their requirements for this type of application. There are many questions that need to be answered by the experts to help make plans for the software specifications. Some of the considerations that need input from the experts are discussed in this chapter.

7.1 User requirements

7.1.1 Screen displays

To design a tool for a trader it is necessary to base it on minimizing the number of clicks of the mouse it takes to complete each task. This is because traders work in a fast paced environment and have no time to waste.

A typical trader watches and trades on more than one commodity market. This means that it is essential to have a software product where it is easy to switch between projects.

When changes need to be made - for example, updating a price belief - this needs to be simple for the trader and take a minimal amount of time to complete.

Although many traders have four screens at their workstation, it is necessary to design the software so that it is easy to use on one screen only. This is because the trader has several packages running at any one time.

7.1.2 Option trading

Traders also trade in Options as well as Futures. Options are contracts that present the opportunity to limit losses with potential for unlimited profits, or to limit profits with the risk of unlimited losses [15].

The Option buyer has the right to initiate delivery at any time while the seller has the obligation to meet the buyers choice at this time (With European Options the buyer can only initiate delivery at the expiration of the contract) [15]. An in depth definition of Options can be found in [15] and [4].

Adding an Option Trading decision making functionality will make Trade-ABLE more useful. This will be straightforward to implement. However, it is still an area for further investigation.

7.1.3 Transaction costs

A solution for handling transaction costs into the decision rule is not available yet. It is important to be aware of the impact of transaction costs as they take a bite into profits, and hence they cannot be ignored.

Margin requirements

When a trader takes on a position, he puts down a margin. A margin is like a holding deposit - typically between 5-20% of the total cost of the position. For example, 100 lots bought at \$1 with a 5% margin will mean putting down a deposit of \$5.

The trader is required to cover the margin and any realizable losses incurred. So, for example, if you bought 100 lots bought at \$1 with a 5% margin and tomorrow the value of the lots dropped to \$0.90 then holding the lots now means putting down an extra \$10 on top of the initial \$5 payment.

The margin is like a good faith deposit and it will be returned to the trader, less any losses, when he closes down his position. Nevertheless, the interest payment or opportunity cost of the capital costs the trader money.

Brokerage costs

To buy and sell commodities a transaction fee is paid to the broker. This is normally in the form of a bid-ask spread. The buying price is typically more than the selling price and the difference goes to the brokers and to the market exchange as a transaction fee. Different markets have different transaction costs and different trading companies negotiate different fees. Also, the bid-ask spread is variable.

7.1.4 Price feeds

The software must have the option to be connected to a live feed. Prices need to be available to the user within the package rather than have the user refer to another tool to find the price to enter into TradeABLE manually. Nevertheless, some traders will need to enter prices manually when a live data feed is unavailable, so a manual price entry option must also be available.

Several companies act as distributors of price data on behalf of trading markets. The market exchanges will also provide prices directly by themselves. Each company will have a slightly different protocol for their customers to pick up the data stream onto their computer. It is sensible to let the trader configure the set up to retrieve the data himself, for example, by providing a text box where he can enter the http address to locate the price feed.

The rate at which the data feed is sampled, for example, every 15 seconds, every 30 seconds, or every minute can also be configured by the user.

Sometimes there are anomalies in the data feeds. A “spike” or a bad price may come through. Also, there will be periods where the price feed is unavailable owing to technical faults. Most traders have experience to handle this.

7.1.5 Market contract specifics

There are many different market contracts available to trade. It is not possible to allow for all of these in the tool. Also, the contract specifications are subject to change, and hence details will need to be updated regularly.

Market specifications include information about the contract including:

- Price quotation (price quoted on publications - listed price)

- Contract size (quantity in one lot)
- Minimum price movement
- Contract months (e.g. Cocoa May 07 will be delivered at the end of May 07)
- Delivery/settlement terms

The tool will have a configuration file where the user can enter some of the more important details which are necessary for the tool to work. For example, the conversion factor between the price quotation and the price for one lot is essential.

7.2 Trial with a large utility company

7.2.1 The prior belief

The traders forecasted that UK Power Winter 06 closing prices would rise linearly from £60.15/MWh to £63/MWh over the 15 day trading period from 15/2/06 to 7/3/06. The random walk model was assumed.

The volatility estimate, $\sigma^2=1.188$, was derived from historical prices. However, this was overestimated by a factor of two. This is because it was believed that the formula for the volatility estimate was:

$$\sigma_z^2 = \frac{\sum_{i=1}^{n-1} (P_{i+1} - P_i)^2}{(n-1)},$$

instead of the correct estimate which is [29]:

$$\sigma_z^2 = \frac{\sum_{i=1}^{n-1} (P_{i+1} - P_i)^2}{2(n-1)}$$

This was not picked up as a mistake until a date after this trial and this was an important lesson learned.

Here $\tilde{\sigma}^2$ was set to 11.015 to correspond to a confidence of around 6 on the ruler (on a scale of 1 to 10 where 10 is very confident). The trial began with the price belief shown in Figure 7.1 with the actual prices shown in blue.




Figure 7.1: The initial price belief with actual prices

On the 22/2/06 the traders updated their prior beliefs. The actual closing price on 21/2/06 was £64.05/MWh, and here the traders expressed that they believed closing price would fall linearly from £64.05/MWh on 22/2/06 to a target of £63/MWh on 7/3/06. Here the volatility estimate was still the same ($\sigma^2=1.188$). The confidence was adjusted to $\tilde{\sigma}^2 = 6.6$ since there were fewer trading days left, and hence less uncertainty about the final price on 7/3/06. The updated price belief is shown with actual prices in Figure 7.2.




Figure 7.2: The updated price belief with actual prices

7.2.2 Utility function

The utility function used to model the trading risk behaviour was $1 - \exp(-X/721348)$.

The risk factor, R , of 721348 was derived from (see section 4.4):

$$1 - \exp(-MaximumLoss/R) = -1$$

where the maximum loss was specified at £500,000 by the traders.

7.2.3 Other parameters

In the utility function section it is stated that the traders expressed a maximum loss of £500,000 which is true. However, the maximum loss input used to derive R was £1,000,000 which has the effect of increasing the risk tolerance property of the utility curve.

The facilitator chose to do this, since the extreme price limits used at the time of the trial were very conservative to avoid the maximum loss. It was

observed that the traders did not exhibit the same level of risk aversion as the tool. This is one of the reasons why it was decided to move to dynamic price limits to constrain trading to ensure that losses do not pass a maximum acceptable level.

In this section, the trial ran with only the static extreme price limits functionality in place with the lower price limit set at £55 and the upper price limit set at £70, and the maximum loss input remains at £1,000,000.

The cap on lots (the most that can be bought or sold on one day) was 50. The starting stock was 30 lots, the starting balance was -£8,537,631, and the starting realizable profit at p_1 was -£387,369. Of course, when the price belief was updated, checks were made to ensure that the position, balance, and profit were carried forward correctly.

The contract traded was UK Baseload Power Winter 06. The conversion factor between listed price and lot prices was given by the traders as 4760. The facilitators relied on this conversion rate since the contract specification is not readily available in the public domain. In a further trial, the trader changed the specification of the conversion factor to 10Mw* 26 weeks* 7 days* 24 hours = 43680 MWh per lot. Here, the conversion factor is kept at 4760. However, in a real case scenario it is essential to make sure this conversion is correct as it is obviously important for balance and profit calculations and - more importantly - to control exposure to risk by recommending a magnitude of lots to buy or sell.

7.2.4 Analysis of diagnostics

The numerical diagnostic results are presented in Table 7.2 at the end of this section.

Size ratio and Gamma-probability

The size ratio and gamma probability plots are shown in Figure 7.3.

It is obvious on day 1 and day 2 that there is something wrong with the forecast model. The estimate for price on day 1 was 60.15 (the same as the price on day 0). However, the observed price, p_1 , was 62.50 - a 2.35 difference. This does not matter greatly as it is the forecasted prices for days 2 ... n that are of most importance since the first decision is made at the observed price, p_1 , of 62.50.

On day 2, the size ratio diagnostic records a very high value at 7.23. The corresponding gamma probability shows that the approximate probability of observing such a size ratio, given that expertise is precise, is 0.007.

$E_1[P_2]$ was 62.64 and the actual price observed on day 2, p_2 , was 59.75 - a discrepancy of 2.89. This is a poor forecast given that the assessed volatility is in the region of $SD_t[P_{t+1}] = 1.08$. This price movement wiped out approximately £380,000 of profit.

Also, on day 5, the size ratio is high at 5.90. The forecasted price for day 5 on day 4, $E_4[P_5]$ was 61.44, but p_5 was observed at 64.05. Since p_4 was 61.25 the trader was lucky that the current position was long and made £480,000. A trader can also make profits by chance under a poor forecasting model.

Discrepancy

On the first five days the discrepancy probabilities are very large. This indicates that the observed prices $p_1 \dots p_5$ do not match their predicted values. The expert would definitely be advised to investigate the mismatch, and if appropriate, change his price beliefs accordingly. The discrepancy probability is shown in Figure 7.4.

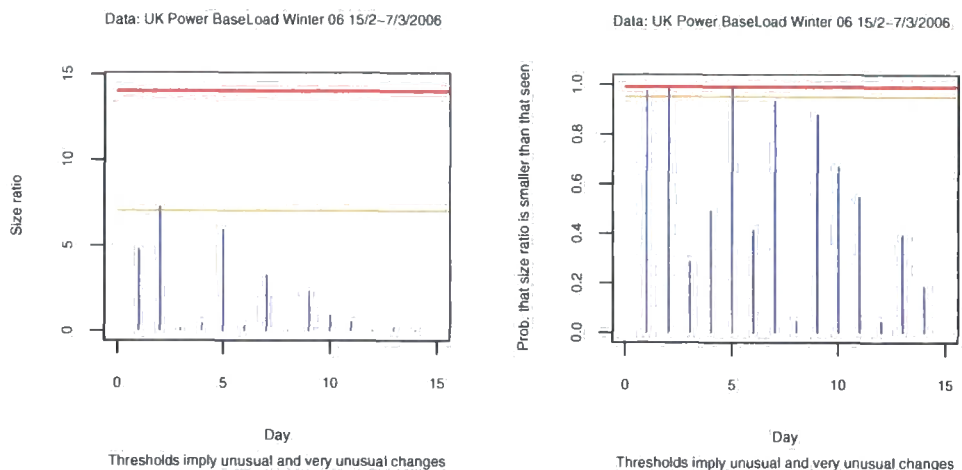


Figure 7.3: Size ratio and Gamma probability

After the price belief is updated (after day 6) the diagnostic values appear to be more reasonable. Here, there is no cause to be concerned about the forecast model owing to the discrepancy probability readings.

Systematic error

The majority of the cumulative one step forecast errors (plotted in Figure 7.5) lie below zero. This would not happen if the forecasting model was precise since the systematic error is supposed to wander around 0. It is clear from this diagnostic that the forecasts for price are consistently above the observed prices.

Variance learning

The volatility estimate drops over time. This is weak evidence that the volatility estimate for the UK Power market was incorrect (overestimated) at the start of the trial. Since it is known that the calculation to estimate σ^2 using historical data was calculated incorrectly by factor of two this is not a

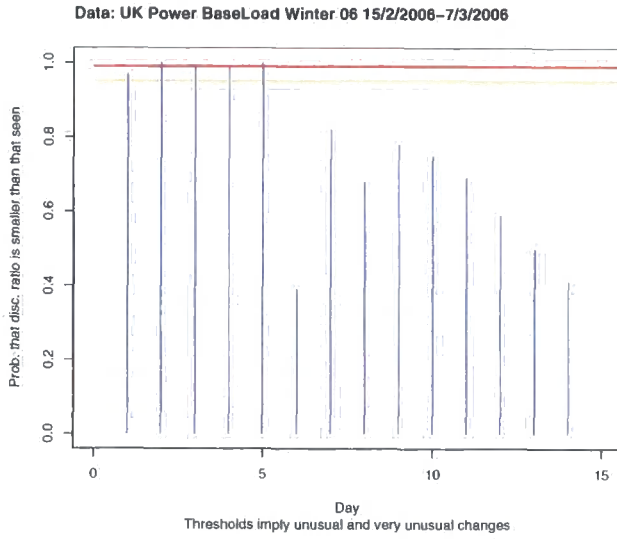


Figure 7.4: Discrepancy




Figure 7.5: Systematic error

surprise. View the variance learning diagnostic plot in Figure 7.6.

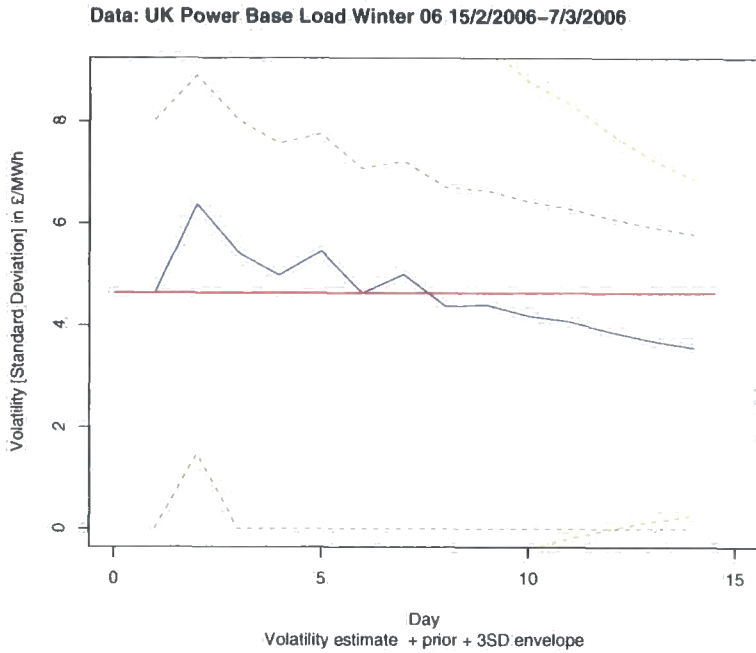


Figure 7.6: Variance learning

7.2.5 Comparison of trading decisions

The traders from the large utility company told the facilitators of the trial their trading decisions at every closing price for the fifteen days. They did not have access to TradeABLE's decisions. This was so that TradeABLE could be compared with actual traders in a meaningful way.

Every day the actual traders were prompted to communicate their decision. They were also asked if they were still comfortable with their price beliefs, and if not, they told the facilitators their revised prior.

The differences in trading decisions are plotted in Figure 7.7 and the differences in realizable profit are plotted in Figure 7.8.

As discussed before, the profit made from the fluctuation of price between

Table 7.1: Power trial: TradeABLE's results

Time	Price	$E_t[P_{t+1}]$	$SD_t[P_{t+1}]$	Position	Change in stock	Realizable Profit
1	62.50	62.64	1.08	29	-1	387369
2	59.75	59.97	1.08	43	14	7759
3	60.35	60.56	1.07	40	-3	130567
4	61.25	61.44	1.07	36	-4	301927
5	64.05	64.16	1.07	21	-15	781735
6	63.50	63.41	1.06	-20	-41	726747
7	61.50	61.50	1.06	0	20	917147
8	61.45	61.45	1.06	1	1	917147
9	59.85	59.94	1.06	17	16	909531
10	58.95	59.10	1.06	26	9	836703
11	59.85	59.95	1.06	18	-8	948087
12	60	60.10	1.05	16	-2	960939
13	59.60	59.73	1.05	19	3	930475
14	59.95	60.07	1.05	17	-2	962129
15	60.40	NA	NA	0	-17	998543

Table 7.2: Power trial: TradeABLE's diagnostic results

Time	Size Ratio	Gamma	Discrepancy	Systematic Error	Variance Learning (SD)
1	4.77	0.972	0.97	0.00	1.09
2	7.23	0.993	1.00	-2.69	1.50
3	0.13	0.287	0.99	-1.65	1.28
4	0.41	0.491	0.99	-0.98	1.17
5	5.90	0.986	1.00	0.37	1.28
6	0.27	0.414	0.39	0.00	1.09
7	3.22	0.932	0.82	-1.79	1.18
8	0.00	0.046	0.68	-1.30	1.03
9	2.28	0.878	0.78	-1.94	1.04
10	0.87	0.670	0.75	-2.14	0.98
11	0.50	0.547	0.69	-1.60	0.96
12	0.00	0.045	0.59	-1.44	0.91
13	0.22	0.394	0.50	-1.51	0.87
14	0.04	0.188	0.41	-1.34	0.84
15	NA	NA	NA	NA	NA

day 4 and day 5 is luck rather than profiting from an excellent forecasting model. The actual traders also benefit from this particular price movement.

In this trial, the tool makes more profit with less exposure to risk. TradeABLE's profit on day 15 was £998,543 which is larger than the actual trader's profit of £599,189. The highest position that the tool held was 43 lots in comparison to the trader's highest position of 70. Although the traders carried more risk on average they still did not make as much profit as the tool.




Figure 7.7: Comparison of trading decisions

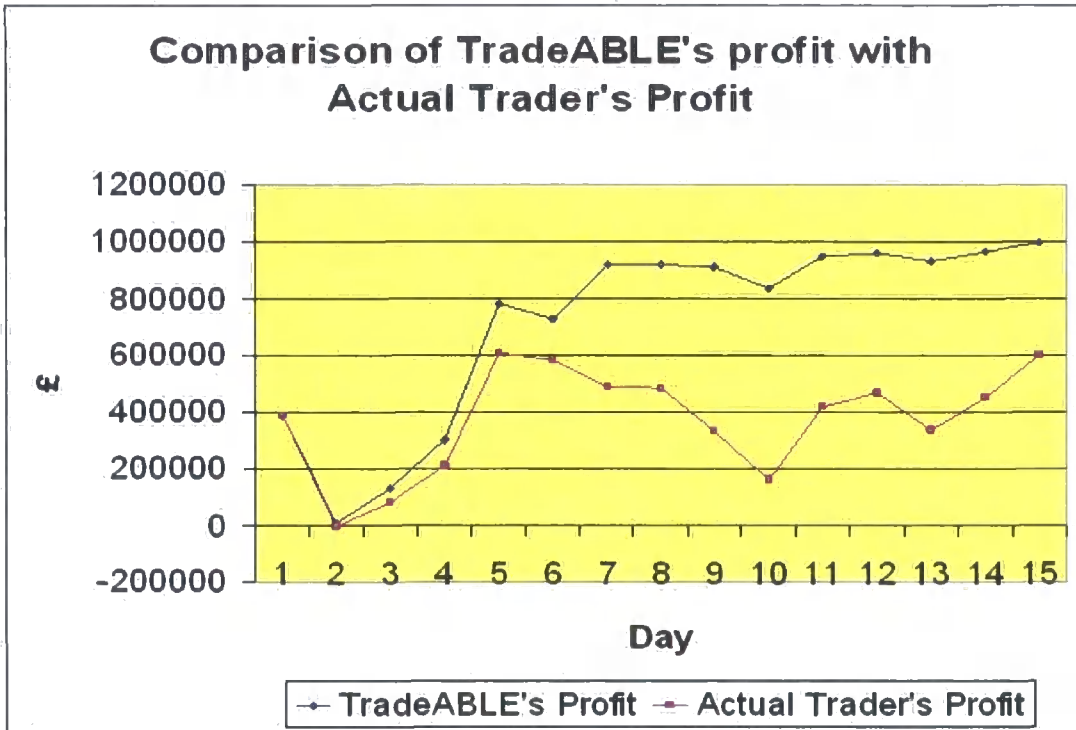


Figure 7.8: Comparison of trading profits

Table 7.3: Power trial: Comparing an actual trader's results with TradeABLE's results

Time	Price	Actual Trader			TradeABLE		
		Position	Change in stock	Realizable Profit	Position	Change in stock	Realizable Profit
1	62.50	30	0	387369	29	-1	387369
2	59.75	30	0	-5331	43	14	7759
3	60.35	30	0	80349	40	-3	130567
4	61.25	30	0	208869	36	-4	301927
5	64.05	10	-20	608709	21	-15	781735
6	63.50	10	0	582529	-20	-41	726747
7	61.50	20	10	487329	0	20	917147
8	61.45	20	0	482569	1	1	917147
9	59.85	40	20	330249	17	16	909531
10	58.95	60	20	158889	26	9	836703
11	59.85	70	10	415929	18	-8	948087
12	60	70	0	465909	16	-2	960939
13	59.60	70	0	332629	19	3	930475
14	59.95	70	0	449249	17	-2	962129
15	60.40	0	-70	599189	0	-17	998543

7.3 Trial with a sugar broker

7.3.1 The prior belief

The experts believed that closing prices would rise linearly from 17.85 cent/lb to 20 cents/lb over the 14 day trading period from 13/2/2006 to 3/3/2006. The assertion of a random walk around expected prices was assumed although this was not explicitly stated.

Normally $\tilde{\sigma}^2$ is deduced from a combination of:

1. the experts' confidence on a scale of one to ten about their expectation of the price on the final day of the trading period , and
2. the experts' initial estimate of σ^2 .

The experts chose to increase the size of the $\tilde{\sigma}^2$ that was fed back to them. A new volatility estimate $\sigma^2=0.4$ was then derived from the experts confidence envelope around their expectation of price on 3/3/2006 (20 ± 1.93 cents with around 70% confidence) to fit in with the random walk forecasting model. This was flagged internally as inappropriate by an email on 9/2/2006 - but no corrective changes were made. σ^2 and $\tilde{\sigma}^2$ are linked through the random walk component of the model and if σ^2 and $\tilde{\sigma}^2$ are not compatible then the model is wrong not σ^2 and $\tilde{\sigma}^2$. It is incorrect to manipulate σ^2 and $\tilde{\sigma}^2$ to satisfy implicit requirements.

Although there was an error in the elicitation process, the trial went ahead with the prior belief shown in Figure 7.9 - alongside the actual prices plotted in blue.




Figure 7.9: The Price Belief with Actual Prices

7.3.2 Utility function

The utility function used to model the trading risk behaviour was $1 - \exp(-X/721348)$.

The risk factor, R , of 721348 was derived from (see section 4.4):

$$1 - \exp(-MaximumLoss/R) = -1$$

where the maximum loss was specified at \$500,000 by the traders.

A concern was expressed via an email on 8/2/2006 that this was a poor representation of the risk attitude of the traders since it was not based on any risk-reward related questions. An exponential utility with risk factor 721348 is very close to a risk neutral utility function relative to the money the traders were dealing with. In a previous trial, an expert on utility theory elicited the same company's risk factor for an exponential curve at around 50000 to 100000, yet this information was ignored. No actions were taken

and the risk factor of 721348 was chosen.

7.3.3 Other parameters

In this example of the trial the cap on lots (the most that can be bought or sold) in one day is set at 50, the starting stock is 118 lots, the balance is -\$2,426,682 and the carried forward profit is \$0.

The extreme price limits were specified as (15,35). However, the dynamic price limits (caution level two) are used to protect against the maximum loss in the results given here.

The contract traded is New York Sugar No. 11. The conversion factor between listed prices (cents/lb) and lot prices (\$ /112,000lbs) is 1120.

7.3.4 Analysis of diagnostics

The numerical diagnostic results are presented in Table 7.5.

Size ratio and Gamma probability

The forecast $E_9[P_{10}]$ was poor - the prediction for price on day 10 was 18.01 cents; however, the observed price was 16.74 cents. This is an error of -1.27 cents. The size of the price drop from day 9 (17.8 cents) to day 10 (16.74 cents) was a big change for this sugar market given the volatility estimate. This problem was highlighted by the Size Ratio (4.14) and its corresponding Gamma-probability (0.96) on day 10. The approximate probability of experiencing a size ratio of magnitude greater than or equal to 4.14, given that the forecast model is precise, is 0.04.

10 out of the 13 size ratios recorded are less than the expected size ratio of magnitude 1. This feels peculiar and could indicate that the volatility estimate is too large. The size ratio and Gamma probabilities are shown as Figure 7.10.

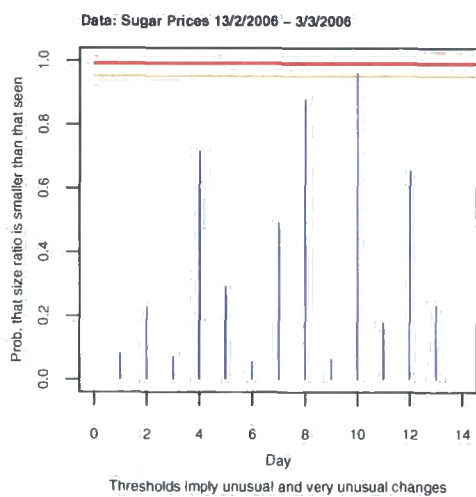



Figure 7.10: Size ratio and Gamma probability

Discrepancy

None of the Discrepancy probabilities are greater than 0.39, and there are several very small Discrepancy values. Here lies evidence that there is an inadequacy in the forecasting model as some larger probability values are expected. The Discrepancy indicator is shown in Figure 7.11.




Figure 7.11: Discrepancy

Detection of systematic error

There is a sequence of diagnostic values that are below zero. Since this indicator is expected to follow $N(0, 1)$ it will wander around zero when the forecasting model is good. Here, every value is negative and this means that the forecast for price was consistently above the observed prices. When this indicator records a value outside $(-1, 1)$ a warning will be flagged to the forecaster. This happens at trading day 10. This indicator provides quite

strong evidence that the forecasting model was inaccurate. Figure 7.12 shows a plot of the Systematic Error.




Figure 7.12: Detection of systematic error

Variance learning

The experienced volatility of observed price is significantly smaller than the prior volatility estimate. Nevertheless, the experts expressed that they felt their variance estimate was good. The variance learning is shown as Figure 7.13. A confidence envelope around the volatility estimate is provided. However, a reliable interpretation can not be derived from it.

Conclusion

There is quite strong evidence to suggest that the forecast for prices were biased on the upper side of observed prices. It is felt that the main inadequacy




Figure 7.13: Variance learning

in the prediction model was that it consistently overestimated price. Also, the volatility estimate may have been too large.

There were important lessons learned in the elicitation of information from the experts. The prior belief must not be perverted to suit the assumption of the random walk. Also, if the tool is to control risk successfully more care needs to be taken to model the trader's risk attitude via the utility function. The decisions output by the tool in this example were based on a high risk strategy which may have been impractical for the trader. The results vary dramatically with the choice of risk level.

7.3.5 Comparison of trading decisions

The traders supplied their decisions without seeing the results from the tool. Overall, they made a loss of \$185,427. The objective supplied was to close

Table 7.4: Sugar trial: TradeABLE's results

Time	Price	$E_t[P_{t+1}]$	$SD_t[P_{t+1}]$	Position	Change in stock	Realisable Profit
1	17.91	18.07	0.62	168	50	0.0
2	17.9	18.07	0.62	218	50	-1881.6
3	18.12	18.29	0.62	268	50	51833.6
4	17.63	17.81	0.62	318	50	-95244.8
5	17.59	17.78	0.62	314	-4	-109491.2
6	17.82	18.01	0.62	364	50	-28604.8
7	18.41	18.59	0.62	327	-37	211926.4
8	17.64	17.85	0.62	300	-27	-70078.4
9	17.8	18.01	0.62	250	-50	-16318.4
10	16.74	16.99	0.62	200	-50	-313118.4
11	17.12	17.36	0.62	150	-50	-227998.4
12	16.79	17.05	0.62	100	-50	-283438.4
13	16.88	17.15	0.62	50	-50	-273358.4
14	17.23	NA	NA	0	-50	-253758.4

with a position of 0. The traders closed with a long position and hence it is not meaningful to compare the actual trading versus the algorithm.

Table 7.5: Sugar trial: TradeABLE's diagnostic results

Time	Size Ratio	Discrepancy	Gamma	Systematic Error	Variance Learning (SD)
1	0.01	0.08	0.082	0.00	0.63
2	0.08	0.04	0.229	-0.28	0.52
3	0.01	0.01	0.071	-0.14	0.48
4	1.11	0.12	0.716	-0.72	0.49
5	0.13	0.07	0.292	-0.80	0.46
6	0.00	0.03	0.055	-0.69	0.44
7	0.41	0.03	0.491	-0.37	0.43
8	2.31	0.15	0.877	-0.92	0.46
9	0.01	0.09	0.065	-0.88	0.44
10	4.14	0.39	0.960	-1.51	0.49
11	0.05	0.31	0.179	-1.37	0.47
12	0.85	0.31	0.655	-1.58	0.46
13	0.08	0.24	0.232	-1.59	0.45
14	NA	NA	NA	NA	NA

Chapter 8

Sensitivity analysis of the algorithms

An application of design of experiments

The objective of the design of experiments is:

1. To understand how the algorithms work.
2. To provide information on the performance of the tool that the Layman can understand. This is of value for marketing the tool and for training staff at Energy Scitech Ltd
3. To find anomalies.

8.1 Design of the experiment

The design of experiments will provide information on how sensitive the algorithm outputs are to a number of varying factors. The output that is of main interest is the realizable profit - in this case, on the 60th day of a 60 day period. At the planning stage, one must consider which factors will be varied and to what levels, the required computing processing time, how to derive

Table 8.1: A table illustrating the combinations of factors used in the design

Prior	ψ	σ^2	σ^2	U(M)=-1	U(M)=-10	U(M)=-100
Flat	Low	Low	Low			
			High			
		High	Low			
			High			
	High	Low	Low			
			High			
High		Low				
		High				
Upward Trend	Low	Low	Low			
			High			
		High	Low			
			High			
	High	Low	Low			
			High			
High		Low				
		High				
Downward Trend	Low	Low	Low			
			High			
		High	Low			
			High			
	High	Low	Low			
			High			
High		Low				
		High				

the price series which will be run through the model, which results will be stored, and any error handling which is required in the execution program.

The plan involves carrying out a $3 \times 2 \times 2 \times 2 \times 3$ factor experiment. At each combination of factors 100 price series are simulated to represent actual market prices. Table 8.1 shows the design grid plan.

72 design points \times 100 simulated price sets = 7,200 runs \sim 10 hours CPU time. If more simulations are added to each design point this will improve the experiment and the testing can easily be extended in this way; however,

100 price sets per combination of factors is sufficient.

The 7,200 price series required are generated as random walks that emulate oil prices, where each price is the preceding price plus a pseudo random step. A description of how these simulated price series were created is given below. The design is repeated three times for:

1. $P_t = p_{t-1} + Z_t$
2. $P_t = p_{t-1} + 0.1 + Z_t$
3. $P_t = p_{t-1} - 0.1 + Z_t$

This can be analyzed as three separate designs or as a $3 \times 3 \times 2 \times 2 \times 2 \times 3$ factor experiment.

The factors that are varied are ψ , σ^2 , $\tilde{\sigma}^2$, the utility function, and the price belief.

Price at time $t+1$ is estimated to be within $p_t \pm \psi\sigma^2$. These price limits are in place to constrain trading to prevent losing more than a maximum amount specified by the expert. ψ is described in detail in chapter 3. The levels chosen for the experiment are $p_t \pm 2\sigma^2$ and $p_t \pm 6\sigma^2$ ($\psi=2$ and $\psi=6$).

σ^2 is the day to day volatility estimate. The levels chosen here are 0.17 and 1.57. These are on either side of the step variance of the simulated prices. 0.17 will on average be too small a volatility estimate and 1.57 will be too large.

$\tilde{\sigma}^2$ is the trader's confidence about the prior specified. The low and high levels correspond to very low confidence and very high confidence in expertise respectively. They are number one and number ten on the confidence ruler which is explained in chapter 3.

The priors are flat, upward trending, and downward trending. The gradients of the trends correspond to an expected step size of 0.1 as per the

Table 8.2: Table of risk factors

M	$U(M)=-1$	$U(M)=-10$	$U(M)=-100$
500,000	721,348	208,517	108,340

Table 8.3: A table illustrating the levels of factors used in the design

Factor	Level	Represented By:
Prior	Flat	40 to 40
	Upward Trend	40 to 45.9 (Linear)
	Downward Trend	40 to 34.1 (Linear)
ψ	Low	2
	High	6
σ^2	Low	0.17
	High	1.57
$\tilde{\sigma}^2$	Low	$0.9 \times 60\sigma^2$
	High	$0.5 \times 60\sigma^2$

simulated trending prices.

Utility is varied using $U(M)=-1$, $U(M)=-10$, and $U(M)=-100$. Here, M represents a maximum acceptable loss of 500,000 units and $U(w)$ is the exponential utility function. $U(M)=-1$ is approaching a risk neutral utility and $U(M)=-100$ is more risk averse. Table 8.2 shows the corresponding risk factors.

If factor levels were chosen differently, then the results obtained could be different. For example, if low and high levels for a factor were chosen to be too close, then the significance testing may show no influence on realizable profit owing to a change in the factor. With a restricted number of levels of a factor or with error in assessing the typical range of a factor the relationship between realizable profit and the factor could be completely missed.

8.1.1 Pseudo random walk

The code that follows is designed to return a vector of prices that are generated from a pseudo random walk of the form:

$$P_{t+1} = p_t + Z_t$$

where $Z_t \sim N(\mu, \sigma^2)$.

The arguments of the `simulateprices` function are defined as follows:

1. `startingprice` : the starting point of the price series.
2. `expectationofonestep`: μ
3. `varianceofonestep`: σ^2
4. `numberofsteps`: The number of steps in the price series. For example, `numberofsteps=60` will fill a vector with 60 prices.

An R function for simulating prices written by Laura Dow:

```
simulateprices<-function(startingprice,expectationofonestep,
                          varianceofonestep,numberofsteps){
  pseudorandomwalk<-numeric()
  pseudorandomwalk[1]<-startingprice
  for(i in 2:numberofsteps){
    newprice<-rnorm(1,startingprice+expectationofonestep,varianceofonestep)
    newprice<-round(newprice,2)
    startingprice<-newprice
    pseudorandomwalk[i]<-newprice
  }
}
```

```

return(pseudorandomwalk)
}

```

The prices are rounded to two decimal places since there are no real life markets prices that have more precision than this.

The design grid is run three times with all the required price series simulated as follows:

1. No trend in prices - `simulateprices(40,0,0.7,60)`
2. Upward trend in prices - `simulateprices(40,0.1,0.7,60)`
3. Downward trend in prices - `simulateprices(40,-0.1,0.7,60)`

The variance, 0.7, is chosen to emulate Brent Crude Oil prices, and the expected step is an arbitrary number chosen to match the trend in the prior belief factor. Each of the price series generated will be unique.

Now, for convenience, price series are added to an array as they are created using the following simple code:

```

arrayofsimulatedprices<-function(startingprice,expectationofonestep,
                                varianceofonestep,numberofsimulations,
                                numberofsteps){
  simulatedpricearray<-array(simulateprices(
                                startingprice,expectationofonestep,
                                varianceofonestep,numberofsteps),
                                dim<-c(numberofsteps,numberofsimulations))
  for(i in 2:numberofsimulations){
    simulatedpricearray[,i]<-
    simulateprices(startingprice,varianceofonestep,numberofsteps)
  }
}

```

```

    }
    return(simulatedpricearray)
}

```

8.1.2 Interpretation of the Results

It is difficult to test for the importance of different factors impacting on the results of computer simulations. Here, Analysis of Variance is used as an exploratory technique, in that it may be concluded that factors with larger mean-square values are more influential than factors with smaller mean-square values.

8.2 Results

8.2.1 Initial exploration: $P_t = p_{t-1} + Z_t$

The initial exploration involves an overview of the summary statistics and an evaluation of some simple exploratory plots. This section looks at the results obtained from simulating price with no trend.

Utility

Provided in Table 8.4 are the summary statistics for profit under the three different levels of utility. $U(M)=-100$ is more risk averse than $U(M)=-10$, and $U(M)=-10$ is more risk averse than $U(M)=-1$. As the utility function becomes more risk averse the spread of the distribution of realizable profit reduces. This is displayed pictorially in Figure 8.1.

The minimum realizable profit under each utility function is around -2,500,000. This loss exceeds the maximum acceptable loss of 500,000 by far. The runs where losses of more than 500,000 occur are owing to the optimization routine failing under the maximum loss constraints. Here, the

Table 8.4: Table of summary results : Utility and distribution of profit

Summary of Realisable Profit						
Factor	Minimum	1st Qu.	Median	Mean	3rd Qu.	Maximum
$U(M)=-1$	-2547000	-99040	-40	2738	140800	2881000
$U(M)=-10$	-2309000	-70740	2945	-479	84990	1668000
$U(M)=-100$	-2766000	-4760	1840	-1185	54560	1001000

maximum loss constraint is released to allow completion of the optimization. Of course, in a real life situation, the expert using the algorithm is informed that this has happened and is told to finish or proceed with extreme caution.

Exceeding the maximum loss has been allowed in these experiments because the runs are automated to complete overnight. Dumping and replacing runs that fail to keep within the maximum loss is tedious and probably unnecessary. There is a useful learning process about what happens or what needs to be done if the maximum loss constraint cannot be kept under the optimization routine.




Figure 8.1: Plot of utility versus profit

Table 8.5: Table of summary results : Dynamic price limits and distribution of profit

Summary of Realizable Profit						
Factor	Minimum	1st Qu.	Median	Mean	3rd Qu.	Maximum
$\pm 2\sigma$	-2547000	-90830	6240	-215	106300	2627000
$\pm 6\sigma$	-2766000	-62100	0	931	62950	2881000

Dynamic price limits - ψ

There are two levels of ψ . $\psi = 6$ is a more conservative approach than $\psi = 2$, and hence the distribution of profit will be narrower under $\psi = 6$ than $\psi = 2$. This is shown by the results of the experiment given in Table 8.5 and Figure 8.2.




Figure 8.2: Plot of the Dynamic Price Limits Versus Profit

Prior volatility - σ^2

The effect of the prior volatility estimate on profit can be seen immediately. If the volatility estimate is much lower than actual volatility of prices then

Table 8.6: Table of summary results : Volatility - σ^2 - and distribution of profit

Summary of Realizable Profit						
Factor	Minimum	1st Qu.	Median	Mean	3rd Qu.	Maximum
Low $\sigma^2=0.17$	-2766000	-240700	15040	3359	254500	2881000
High $\sigma^2=1.57$	-1216000	-40710	0	-2643	30280	799400

the distribution of profit is much wider than when the volatility estimate is far too high.

A good assessment of σ^2 is critical to the behaviour of the optimization. If markets are more volatile than expected there could be too much capital put at risk. However, if markets are less volatile than expected too little capital is put at risk and so an opportunity to make more profit (if the prior is precise) is forfeited.

Confidence about final price

When an expert expresses low confidence in his belief the tool trades more conservatively since there is less conviction about the trader's forecast. The lower the confidence about the final price, the tighter the distribution of profit. This can be seen in both Table 8.7 and Figure 8.4.

By expressing a high confidence in beliefs, the tool will place more faith in the expertise. In this situation, holding more lots will, in theory, carry the same risk as buying less lots when there is more uncertainty about the forecast of future prices.

There is a need for a more complex analysis to verify this hypothesis.




Figure 8.3: Plot of the Volatility - σ^2 Versus Profit

Table 8.7: Table of summary results : Confidence about final price and distribution of profit

Summary of Realizable Profit						
Factor	Minimum	1st Qu.	Median	Mean	3rd Qu.	Maximum
Low $\sigma^2=9.18$ or 84.78	-2547000	-60090	0	-1328	50430	2302000
High $\sigma^2=5.1$ or 47.1	-2766000	-81420	9830	2044	121700	2881000




Figure 8.4: Plot of the confidence versus profit

Prior trend

The distribution of profit under the upward and downward trend is more or less the same under the simulated prices that have no trending term. More interestingly, the distribution of profit when the prior is flat is noticeably narrower. This might be because $E[P_{t+1}]$ is often closer to p_t with a flat prior than with a trending prior.

The distance of $E[P_{t+1}]$ from p_t is important since it is an indicator of how confident the forecast model is in the price going up or down. The further $E[P_{t+1}]$ is from p_t the more lots are bought or sold.

If $E[P_{t+1}]$ is near to p_t on a more frequent basis with a non trending prior, then this could explain the tighter distribution of profit under the flat prior. However, this claim needs to be investigated further to check that it is true.

Table 8.8: Table of summary results : Prior trend and distribution of profit

Summary of Realizable Profit						
Factor	Minimum	1st Qu.	Median	Mean	3rd Qu.	Maximum
Flat	-2239000	-12430	4210	-1991	29530	1308000
Upward	-2766000	-96560	-5010	1586	140300	2627000
Downward	-2124000	-103700	-4080	1479	129900	2881000




Figure 8.5: Plot of the prior trend versus profit

A more complex analysis of dynamic price limits - ψ

A more complex analysis can be made to explore any interactions between factors. Here an investigation is made between ψ and all other factors. This is done to aid understanding and to look for any results that are unexpected.

In Figure 8.6, profit is plotted to look at both volatility σ^2 and ψ together. Again it is evident that the spread of profit reduces significantly when the volatility estimate becomes larger. This stresses the importance in estimating σ^2 accurately since the algorithm is very sensitive to its level. The impact of ψ (with comparison to the effect of σ) is small, although the spread of profit appears to narrow slightly with wider dynamic price limits.




Figure 8.6: Plot of the dynamic price limits and volatility versus profit

In Figure 8.7, profit is plotted to look at both ψ and confidence - $\tilde{\sigma}^2$. It is obvious that a lower confidence in the prior belief or wider dynamic price limits results in a tighter distribution of realizable profit.

In Figure 8.8, ψ and the utility function are explored. Here it appears that




Figure 8.7: Plot of the dynamic price limits and confidence versus profit

widening the dynamic price limits has a larger restrictive effect on the spread of profit when the utility function is more risk seeking. This is because a more risk averse utility has already constricted trading to some degree before the dynamic price limits are taken into account.

In the case where $U(M)=-1$ and $\psi = 2$ there is a very wide distribution of profit. Here, there is evidence of a perceivable high chance of the optimization routine failing under the maximum loss constraint. Considering that the mean profit is 9996 and the standard deviation is 652008 with $U(M)=-1$ and $\psi = 2$ in the experiment, the probability that profit is less than -500,000 is 0.22¹. Some traders will wish to avoid this scenario completely, and hence they will need to optimize with a more risk averse utility and wider dynamic price limits.

¹> pnorm(-500000,9996,652008) [1] 0.2170506




Figure 8.8: Plot of the dynamic price limits and utility versus profit

8.2.2 Initial exploration: $P_t = p_{t-1} + 0.1 + Z_t$

The following explorations look at the results where the trend of the simulated prices is upward.

Utility

As shown in Figure 8.9, the distribution of profit becomes narrower as the utility function becomes more risk averse. When $U(M)=-1$ there are more occasions where the maximum acceptable loss of 500,000 units is exceeded. This is owing to the maximum loss constraint being released when the optimization routine fails under the imposed constraints as explained previously.

Dynamic price limits

As the dynamic price limits become further apart the trading is more conservative and hence the variance of the profit results becomes smaller. This is evident in Figure 8.10.




Figure 8.9: Plot of the utility function versus profit




Figure 8.10: Plot of the dynamic price limits versus profit

Prior volatility

The estimation of volatility has a massive impact on the spread of the resulting realizable profit. This is shown in Figure 8.11. If the volatility estimate is too low the distribution of profit will be wide owing to under restrained trading such that there will be a significant chance that the maximum loss constraint will be breached. If the volatility estimate is too high then trading will be too conservative which results in a tight distribution of profit.




Figure 8.11: Plot of the prior volatility versus profit

Confidence about final price

Confidence about the final price also affects the spread of profit. The spread of profit for a trader who is highly confident about his forecast of the final price is wider than the trader who exhibits low confidence in his final price

forecast. This is shown in Figure 8.12. If a trader expresses high confidence in his beliefs than there is less risk in holding X lots than when the trader has low confidence in his price outlook.

For example, if a trader sources information such that he is sure price will go up tomorrow, he knows he is taking less risk by buying 10 lots than when he did not have this information and hence was more uncertain about tomorrow's price movement.




Figure 8.12: Plot of the confidence about final price versus profit

Prior trend

Here the simulated prices had an upward trend and it is seen in Figure 8.13 that more profit is made on average when that prior forecast featured an upward trend. With the upward trending prior the profit results are skewed to the side of positive profits. The profit distribution when the prior exhibited a downward trend show that there were more occasions when a

loss was incurred rather than a profit.

The size of the variance of profit under the trending priors is approximately the same, but it is noticeable that the variance of profit under the flat prior is significantly smaller. A possible explanation for this feature was explained earlier.




Figure 8.13: Plot of the prior trend versus profit

A more complex analysis of prior trend

The upward trending prior is a more accurate forecast than the flat prior and downward trending prior when the simulated prices contain an upward trend.

When the prior is trending upward the median profit declines as the volatility estimate increases from the low level ($\sigma^2=0.17$) to the high level ($\sigma^2=1.57$). See Figure 8.14. This is owing to taking on excess risk when the volatility estimate is too low and too little risk when the volatility estimate

is too high. If the volatility estimate is too high and the prior is accurate then an opportunity is lost to make more money.

When the prior is trending downwards the median profit increases as the volatility as the volatility estimate increases from the low to the high level. An excess risk is taken on board when the volatility estimate is underestimated and if the price outlook is imprecise it is expected that the trader will lose more money than when the volatility estimate is overestimated.

From 8.14, it is not clear what happens to median profit as the volatility estimate changes under the flat prior. It will be similar to the observation under the downward prior, but to a lesser extent.




Figure 8.14: Plot of the prior trend and volatility versus profit

If the prior is accurate (upward trending prior here), then the median profit increases as the utility function becomes more risk accepting.

The median profit decreases as the utility function becomes more risk accepting when the prior forecast is inaccurate (downward trending and flat

prior here).

This is a feature of utility functions - a more risk averse utility will protect more fiercely against losses at the expense of possible gains achievable when the prior is accurate. A more risk accepting utility will look for larger gains at the expense of realizing larger losses if the prior is inaccurate. The effect of utility and prior trend on profit is shown in Figure 8.15.




Figure 8.15: Plot of the prior trend and utility versus profit

As the dynamic price limits widen it can be seen from Figure 8.16 that the median profit decreases when the price forecast is accurate and that the median profit increases when the price forecast is inaccurate.

This is, again, a feature of what happens under different degrees of risk taking (the wide price limits, $p_t \pm 6\sigma$ are more conservative than $p_t \pm 2\sigma$).

Confidence about final price does not appear to have an obvious effect of the level of median profit with regards to the accuracy of the prior. All that can be inferred from Figure 8.17 is that the variance of profit widens as




Figure 8.16: Plot of the prior trend and the dynamic price limits versus profit confidence increases.

8.2.3 Initial exploration: $P_t = p_{t-1} - 0.1 + Z_t$

The results for the design grid where the simulated prices have a downward trend are approximately symmetrical to when the simulated price have an upward trend. It is not useful to repeat the conclusions that were explained in the previous section.

8.2.4 Analysis of the experiment

Introduction to ANOVA

ANOVA is shorthand for Analysis of Variance and it is used to investigate the impact of variables on data results. The application of ANOVA and the design and analysis of experiments was first developed by Fisher at Rothamsted Research station from 1919 to analyze agricultural data. For example,

Exploring the Effect of the Prior Trend and the Confidence About Final Price on Pr




Figure 8.17: Plot of the prior trend and the confidence about the final price versus profit

Fisher used ANOVA to investigate the effect of the application of Nitrogen to crop yields.

Experiments are used to gain knowledge and understanding in the majority of scientific fields. For example, tests are carried out frequently in medicine, environmental studies, and manufacturing processes. A good statistician can design these types of experiments efficiently and analyze the results using tools such as ANOVA to make reliable inferences from the data.

Application of ANOVA to explore the effect of the factors on the algorithms

The model used to look at the impact of the factors on the realizable profit results from the algorithm is quite complex with five factors that could all possibly interact with each other. The model is as follows:

Table 8.9: Denotation for the ANOVA model

Symbol	Factor	
α_i	σ^2	$i = \text{low or high level}$
ψ_j	Dynamic price limits	$j = \text{low or high level}$
β_k	σ^2	$k = \text{low or high level}$
γ_L	Utility function	$L = U(M) = -1, -10, \text{ or } -100$
δ_m	Prior trend	$m = \text{Flat, Upward, or Downward}$

$$\begin{aligned}
 y_{ijkLmn} = & \text{Intercept} + \alpha_i + \psi_j + \beta_k + \gamma_L + \delta_m + (\alpha\psi)_{ij} + (\alpha\beta)_{ik} + (\psi\beta)_{jk} + (\alpha\gamma)_{iL} \\
 & + (\psi\gamma)_{jL} + (\beta\gamma)_{kL} + (\alpha\delta)_{im} + (\psi\delta)_{jm} + (\beta\delta)_{km} + (\gamma\delta)_{Lm} \\
 & + (\alpha\psi\beta)_{ijk} + (\alpha\psi\gamma)_{ijL} + (\alpha\beta\gamma)_{ikL} + (\psi\beta\gamma)_{jkL} + (\alpha\psi\delta)_{ijm} \\
 & + (\alpha\beta\delta)_{ikm} + (\psi\beta\delta)_{jkm} + (\alpha\gamma\delta)_{iLm} + (\psi\gamma\delta)_{jLm} + (\beta\gamma\delta)_{kLm} \\
 & + (\alpha\psi\beta\gamma)_{ijkL} + (\alpha\psi\beta\delta)_{ijkm} + (\alpha\psi\gamma\delta)_{ijLm} + (\alpha\beta\gamma\delta)_{ikLm} \\
 & + (\psi\beta\gamma\delta)_{jkLm} + (\alpha\psi\beta\gamma\delta)_{ijkLm} + \epsilon_{ijkLmn}
 \end{aligned}$$

where y_{ijkLmn} is the response variable, realizable profit, and ϵ_{ijkLmn} represents the error terms ($\epsilon_{ijkLmn} \sim N(0, \sigma_e^2)$). The subscript n is $1 \dots 100$ and represents the 100 simulated price series at each combination of factors. The other terms represent the explanatory factors and are listed in Table 8.9.

R is used to perform the data analysis. The results for the ANOVA where the expected step of the simulated prices is greater than 0 is given in Table 8.10.

There are three significant three factor interactions and one significant four factor interaction. These are complex to analyze.

There are four significant two factor interactions and three significant main effects. The three main effects that can explain a lot of variation in the realizable profit model are the volatility estimate, the utility function, and the prior trend. These have already been explored to some extent in the

Table 8.10: Summary of ANOVA output for model where $P_t = p_{t-1} + 0.1 + Z_t$

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
v0	1	301	301	9.5870	0.0019672 **
psi	1	4	4	0.1173	0.7320062
vW	1	634	634	20.1973	7.095e-06 ***
riskfactor	1	6	6	0.1895	0.6633866
priortrend	1	9991	9991	318.3177	< 2.2e-16 ***
v0:psi	1	24	24	0.7767	0.3781731
v0:vW	1	169	169	5.3728	0.0204815 *
psi:vW	1	38	38	1.2241	0.2685982
v0:riskfactor	1	47	47	1.4872	0.2226971
psi:riskfactor	1	15	15	0.4920	0.4830675
vW:riskfactor	1	88	88	2.7880	0.0950188 .
v0:priortrend	1	4381	4381	139.5828	< 2.2e-16 ***
psi:priortrend	1	163	163	5.1958	0.0226704 *
vW:priortrend	1	74	74	2.3506	0.1252797
riskfactor:priortrend	1	359	359	11.4225	0.0007295***
v0:psi:vW	1	28	28	0.8969	0.3436585
v0:psi:riskfactor	1	1	1	0.0168	0.8969978
v0:vW:riskfactor	1	0.34790	0.34790	0.0111	0.9161541
psi:vW:riskfactor	1	0.01770	0.01770	0.0006	0.9810573
v0:psi:priortrend	1	75	75	2.4011	0.1212952
v0:vW:priortrend	1	171	171	5.4402	0.0197059*
psi:vW:priortrend	1	1	1	0.0190	0.8903157
v0:riskfactor:priortrend	1	4	4	0.1274	0.7211808
psi:riskfactor:priortrend	1	222	222	7.0809	0.0078084**
vW:riskfactor:priortrend	1	156	156	4.9720	0.0257910*
v0:psi:vW:riskfactor	1	2	2	0.0613	0.8044133
v0:psi:vW:priortrend	1	37	37	1.1911	0.2751367
v0:psi:riskfactor:priortrend	1	40	40	1.2758	0.2587216
v0:vW:riskfactor:priortrend	1	288	288	9.1706	0.0024681**
psi:vW:riskfactor:priortrend	1	2	2	0.0503	0.8224793
v0:psi:vW:riskfactor:priortrend	1	13	13	0.4151	0.5193900
Residuals	7168	224976	31		

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

early initial explorations. It is worth analyzing two of the highly significant two factor interactions with more detail.

If the prior is accurate (in this case an upward trending prior with prices that are simulated to have an upward trend) the more risk accepting utility curve makes a larger profit than a more risk averse curve. If the prior is inaccurate the more risk accepting utility makes a larger loss than a more risk averse curve. Evidence that supports this claim is illustrated in Figure 8.18.




Figure 8.18: The interaction between the prior and the utility function

When the prior is accurate, a low volatility estimate will make a higher mean profit than a high volatility estimate. However, when the prior is inaccurate, a low volatility estimate will incur a larger loss than a high volatility estimate. This means that it is important to estimate the volatility of the market accurately to control risk effectively. See Figure 8.19 for an illustration.




Figure 8.19: The interaction between the prior and the volatility estimate

Residual analysis

By using ANOVA it is assumed that the distribution of the error terms is $\sim N(0, \sigma_e^2)$. A Quantile-Quantile plot can be used to assess if data is from a particular distribution. A Normal Q-Q plot is used to assess if a set of data could be from a Normal distribution.

The Normal Q-Q plot is used here to assess if the residuals from the ANOVA model are from a Normal distribution. This is to check the validity of the assumptions that have been made.

Figure 8.20 shows the Normal Q-Q plot of the error terms from the model. If the data is normally distributed then there will be a straight line with gradient 1 and intercept 0. In Figure 8.20 this is clearly not the case since the line exhibits an obvious double curvature.

Double curvature is a feature of data that comes from a leptokurtic distribution. That is, it has fat tails and has a high peak in comparison to the

Normal distribution.

To stabilize the variance it is possible to try a transformation of the data like taking logs. Nevertheless, although it is evident that the residuals are not Normally distributed, and hence do not meet with the assumptions, the statistical tests are still quite robust in this situation.




Figure 8.20: Normal QQ Plot: Residual analysis of the realizable profit model

8.2.5 Tables of results

Here are the full results of the mean and standard deviation of the 100 runs of realizable profit on day 60 at each combination of factors. There are three tables for simulated prices with:

1. $P_t = p_{t-1} + Z_t$
2. $P_t = p_{t-1} + 0.1 + Z_t$
3. $P_t = p_{t-1} - 0.1 + Z_t$

The mean is colour coded as the blue number and the standard deviation is colour coded as the red number.

8.3 Recommendations

It is important to improve the understanding of the algorithm by creating new simplistic designs and then analyzing the results.

Consideration needs to be given to stabilize the variances in the model to comply with the ANOVA assumptions and this is not trivial to do since many of the investigated factors are designed to change the variance of the realizable profit results at different levels. Perhaps it is necessary to use another statistical tool to learn more about the algorithm.

In particular, it is advised to evaluate the performance of the algorithm in relation to the closeness of the price forecasts to the actual prices. Putting more effort into the planning stages of the statistical investigation will improve the validity and usefulness of any findings.

It is clear from the results how the different factors relate to each other, and this all seems to be intuitively supportable. Nevertheless, the application of ANOVA for significance testing in this situation, which involves computer simulations, is not rigorous. Hence, the results should be regarded as indicative only.

Table 8.11: Results for the simulated price with $P_t = p_{t-1} + Z_t$

Prior	ψ	σ^2	σ^2	U(M)=-1	U(M)=-10	U(M)=-100	
Flat	Low	Low	Low	13033 134517	-1208 42401	-1175 34928	
			High	-10887 578904	30994 252163	18115 175837	
		High	Low	-3362 19773	-248 4752	60.5 2265	
			High	-8682 105406	3528 37672	1882 18830	
	High	Low	Low	2353 111569	-7123 51638	-1852 21107	
			High	-45514 518753	-31394 346371	-9797 135568	
		High	Low	-596 14767	181 3412	384 2140	
			High	1298 51072	-2751 31882	4981 15031	
	Upward Trend	Low	Low	Low	57673 1066705	-59329 623891	10674 337236
				High	-70509 872626	-18916 724361	-90337 544417
			High	Low	131 267796	5552 73557	-5031 41983
				High	5468 367732	-3707 117050	7964 69324
High		Low	Low	9865 803832	43349 538603	24856 327767	
			High	47140 829953	105603 725484	454 523771	
		High	Low	1600 67981	-3659 60540	-2365 43755	
			High	-13503 62559	-7820 54364	-7080 52099	
Downward Trend		Low	Low	Low	-16809 865244	11753 632814	10984 363864
				High	86750 937094	76664 630659	-1707 506971
			High	Low	-22098 289453	-12371 74600	6763 41039
				High	-19600 386284	-10908 123056	1142 63355
	High	Low	Low	4250 889926	-65474 565309	-38234 335050	
			High	53727 839406	-58623 780106	41557 490032	
		High	Low	-5305 74149	-2384 62553	-2664 38533	
			High	-719 67782	-3196 58927	1980 5331	

Table 8.12: Results for the simulated price with $P_t = p_{t-1} + 0.1 + Z_t$

Prior	ψ	σ^2	σ^2	U(M)=-1	U(M)=-10	U(M)=-100	
Flat	Low	Low	Low	-133320 242285	-42227 95650	-15846 41656	
			High	-333059 756975	-329320 634014	-142307 309914	
		High	Low	-12312 27236	-2629 6623	-1129 3751	
			High	-65761 145408	-29152 56114	-22577 35543	
	High	Low	Low	-76798 180097	-38724 59080	-13809 32433	
			High	-321188 705155	-127798 327339	-109745 212847	
		High	Low	-12515 26709	-4302 8751	-1440 3242	
			High	-29925 49148	-24753 37152	-15987 28568	
	Upward Trend	Low	Low	Low	1173245 855292	565792 531460	423360 273661
				High	700059 1023111	715643 617625	452059 366228
			High	Low	238681 252095	80253 67645	39734 36310
				High	327350 262517	91048 96980	62390 43909
High		Low	Low	818219 796093	445886 603705	302749 312644	
			High	659250 853898	542195 642735	376098 388089	
		High	Low	113210 155272	80458 72799	43619 40481	
			High	94016 94388	73484 64731	45488 41789	
Downward Trend		Low	Low	Low	-860150 910732	-701979 605584	-337325 408927
				High	-758057 936087	-675249 982433	-611822 652102
			High	Low	-288168 256341	-62820 63070	-44846 37795
				High	-481047 489957	-139415 151675	-73635 82393
	High	Low	Low	-709189 754388	-569396 629897	-395460 354639	
			High	-628717 732785	-642297 753316	-549855 666585	
		High	Low	-51741 51713	-54002 44600	-40019 34340	
			High	-57413 35494	-50637 42436	-48391 40915	

Table 8.13: Results for the simulated price with $P_t = p_{t-1} - 0.1 + Z_t$

Prior	ψ	σ^2	$\bar{\sigma}^2$	U(M)=-1	U(M)=-10	U(M)=-100	
Flat	Low	Low	Low	-73639 228957	-43204 92478	-21999 45344	
			High	-387271 688631	-149652 394123	-144287 280387	
		High	Low	-13354 25897	-4354 8895	-1072 3084	
			High	-53499 147043	-33053 64783	-23945 38573	
		High	Low	-64155 159827	-28078 51588	-18122 38465	
			High	-336110 774293	-174926 422111	-45160 175669	
	High	Low	Low	-11003 22007	-3572 8047	-1319 3298	
			High	-13473 52109	-19379 40292	-19506 29580	
		High	Low	-1095193 860941	-613255 619795	-420461 394801	
			High	-887955 1029579	-777743 736594	-557837 603691	
		High	Low	-280613 274554	-64124 65393	-43690 38643	
			High	-467295 534135	-109417 132271	-75918 86063	
Upward Trend	Low	Low	Low	-928099 853566	-547158 612286	-380286 352731	
			High	-636413 867610	-832006 899860	-671016 623959	
		High	Low	-57496 36688	-50916 39900	-41795 33758	
			High	-55522 44257	-56969 43075	-48806 38032	
		High	Low	Low	671666 900550	554052 534016	373733 287981
				High	854541 968608	668812 619639	442918 339089
	High		Low	244379 254744	82989 78900	42733 32111	
			High	305427 260760	104019 82612	46544 57117	
	High		Low	783763 915930	543021 534068	354624 327403	
			High	626380 786429	597631 573695	444028 261715	
	Downward Trend	Low	Low	131084 162625	71742 68403	40272 41246	
			High	77192 99114	73139 61608	57358 32707	
High		Low	783763 915930	543021 534068	354624 327403		
		High	626380 786429	597631 573695	444028 261715		
High		Low	131084 162625	71742 68403	40272 41246		
		High	77192 99114	73139 61608	57358 32707		

Chapter 9

Testing software

“Testing is the process of executing a program with the intent of finding errors” [21]

9.1 Code specification and planning tests

To evaluate whether or not a program works properly, first of all, the specification of the code must be explicit [18]. The proposed end functionality of the code is needed so that it is possible for the tester to determine that the software does its intended job. During the test cycle, the tester will be encouraged to suggest improvements to the program and to challenge the original designs [18].

All test cases will be planned before implementation, and the entire process will be recorded. Basically, the test plan will show a Gantt chart - or a similar time management tool - that holds information about test cases to be performed. Schedules will include time for communication between the project team members. Records are vital for monitoring progress, improving the flow of information between the software team, and ensuring that time is not wasted. Masses of paper work are not useful to anyone, so streamlining the recording and planning documentation to a level that works best within

the overall software development process will be helpful.

Discoveries of errors will be passed to the manager with a clear problem summary, and instructions that detail how to reproduce the problem - alongside the suggested solution [18]. The manager decides if a bug will be fixed, the programmer will fix it, and the tester will test the fix.

It is important to allocate time for regression testing (repeating a set of standard tests) [18]. This is necessary since:

1. It is common for the programmer to have not fixed the bug properly owing to poor communication, incompetence, stress, or laziness.
2. It is possible that the fix introduced a bug to another area of the code.

9.2 Limits of testing

A test program will never confirm that there are no errors present in a program since [18]:

1. The design of the ideal user interface is subjective.
2. There is an extremely large number of paths through the code to test.
3. There is an infinite amount of possible user inputs to test.

To minimize this problem, an efficient tester designs a plan that will look at stress points where the program is most likely to fail [18].

The project manager must balance between, taking time and expense to test the software ensuring a good quality level, and the timing to the target deliverable date for driving the product into the market. The key issue to consider is economic viability - if the product remains in the test/development phase for too long, the project may have spent its budget before the software

has been shipped. Nevertheless, if the product is sent into the market full of bugs, it will undermine the credibility of the software, thus decreasing its value and damaging any potential long term success.

9.3 The human machine interface

9.3.1 Handling user inputs

User inputs will be handled effectively if the range and type of the values are planned properly, and hence are definitive for the programmer. For example, a field which defines a probability must be greater than or equal to zero, and less than or equal to one. With adequate preparation, it is simpler to ensure that invalid inputs are caught and handled effectively by the software. To test the resultant code efficiently, values are chosen where the code is most likely to fail. The tester will also try a correct input, to check that the program accepts it as a valid entry, and that the program handles the value as expected.

1. Try a character: A
2. Try the boundary values: 0 and 1
3. Try a negative number: -0.56
4. Try something between 0 and 1: 0.543

It is impossible to fully test given the extent of what a user can insert into a field. Nevertheless, the tester is free to use his creativity to design additional checks. Ultimately, the project manager will decide how much testing is sufficient.

9.3.2 Error messages

If the user enters an invalid input, an error message needs to be displayed to inform them what values are valid. This lets the user know the constraints on what they can enter on their next attempt.

Error messages tend to distress users. The communication of an error needs to be simple, written in an active voice, without any jargon terminology, or error codes [18]. Leading software will show a short explanation of the error, and give the user the option to view a longer description of the problem. Error displays need to be kept up to date and consistent - this is important to prevent the user from becoming confused. Exclamation marks need to be avoided as many users interpret their use as a scolding [18].

9.3.3 Communication style

The communication style of the screens is important. It needs to be simple for the users to understand what they are meant to do within the application.

Good software will display a progress bar, or moving icon, during a task that takes more than two seconds to complete [18]. This is so that the user does not think the system is in an infinite loop, or the computer has crashed.

Although superficial, spelling and grammar mistakes - on screen or in the documentation - will annoy customers. The client will think that if the software company is not worried about correcting bad English, perhaps they are also neglecting to fix more fundamental errors in the underlying software [18].

Keeping the layout, presentation style, and action of buttons consistent is recommended. This is so that the user does not become confused. For example, keep colours, font style and size, and the location of the OK, Cancel, and Help buttons the same throughout the application. Give the user the

option to change these by providing an option to customize the software tool.

People tend to prefer familiarity since this reduces learning time. This is a key argument for following Windows' user interface design guidelines as a large proportion of computers run on a Windows Operating System. An alternative view is that the competing operating systems - Linux and UNIX - have a rising market share, and are predicted to threaten Windows as the world dominant computer systems.

9.3.4 Work flow of the application

“A software error is present when the program does not do what its end user reasonably expects it to do.” [21]

All software users are lazy. They can not be bothered with reading thick support manuals, and get frustrated as soon as something does not happen as they expected. They also hate doing tasks that require a long series of steps.

Most software companies approach this hurdle by working closely with the customers to capture the best way to design the front end of their product. Other companies use rapid prototyping methods - building screens quickly, and asking the user to criticize. Perhaps it is best to mix both of the solutions.

9.4 Testing internal processes

9.4.1 Java error handling

Null pointer exceptions

Null pointers happen when the program is asked to access an object that does not exist.

The following code is an example of a Java `ArrayIndexOutOfBoundsException` where the program tries to retrieve an element of a vector which has not been defined.

```
import java.util.*; class MyClass{
    public static void main(String[ ] args){
        Vector prices = new Vector();
        Object myprice = prices.elementAt(0);
    }
}
```

The error given on running the code is:

```
Exception in thread "main"java.lang.ArrayIndexOutOfBoundsException:
0 >= 0
```

Parsing a value of the wrong type

The following example code will return a parsing exception as the computer does not know how to change the string "Laura" into a number. This type of error will also occur if the programmer attempts to convert "30." to an int value, or trying to parse a value that is not within the type range. For example, asking a Java program to change "20000000" to an int value will result in an exception.

```
class MyClass{
    public static void main(String[] args){
        String myname = "Laura";
        double lauratodouble = Double.parseDouble(myname);
    }
}
```



```
    }  
}
```

The error given on running the code is:

```
Exception in thread "main" java.lang.NumberFormatException: For  
input string: "Laura"
```

Dividing by zero

Attempting to divide by zero will return an `ArithmeticException` in Java. This type of error is easy to trace in a simple program. However, it may go unnoticed in a complex algorithm where a divide by zero scenario will only happen in a small number of cases. A tester will be more likely to find this error by walking through the code, and using logic to identify areas of computation where a divide by zero may happen.

```
class MyClass{  
    public static void main(String[ ] args){  
        int dividebyzero =4/0;  
    }  
}
```

The error given on running the code is:

```
Exception in thread "main" java.lang.ArithmeticException: / by zero
```

Input/Output failures

Input and output from a program carries many potential problems for the programmer. I/O involves writing or reading data, talking to devices such as printers, or communicating with another computer process. A few items

for consideration:

The coder needs to check that the data is written to file exactly where the user specified, and the user needs to be notified if they ask to save a file to a path name that already exists. This will be done by displaying a message similar to “File name X already exists. Do you wish to replace the file?”

If a file cannot be written to as requested, the user should be made aware that there was a failure, and if possible what action needs to take place so that the required file can be saved to disk. For example, an Excel file cannot be written to from a Java application, if the user has it open on their computer. They should be notified that the particular Excel workbook needs to be closed if they wish to overwrite the file. This is similar to the I/O failure message shown on Windows - “Another application is using this file. ...”.

File extensions will need to be handled properly. There are separate mechanisms required to read and write different file types, and many applications rely on the extension letters for recognition.

Software interfacing or integration is where two or more programs - which are not necessarily written in the same language - need to communicate with each other to perform a task. This type of work will present a lot of hazard for error to the programmer. Wrapping a Java application around a R process is an example of this, and is discussed in the next section.

9.4.2 Interfacing between R and Java

The interaction of two sets of code written in different languages is a difficult process to manage. The programs need to communicate with each other and this has the potential to cause many problems.

Firstly, there is a high risk of the communication part between the pro-

grams failing. The stream of data flowing between the two programs is vulnerable as it may break or become corrupt. Normally, the data is controlled by sending it through a buffer stream.

If there is a huge amount of data there is a larger risk of a problem occurring. For instance, the buffering process between R and Java can stop and cause the computer to hang. The common reasons for this are:

1. The R process has experienced a failure or produced an error and has not notified the Java program.
2. The Buffer is waiting for more data - nothing is being sent to it and hence the communication has stalled.
3. The R terminal quits before sending all the data required by Java and this means the Java program is missing essential results.

It is not trivial to catch an error in the secondary R process and display it back to the user in Java. For example, R prints out technical and unfriendly error messages and hence it is essential (initially) to be in the R environment to track down the source of the failure. It is feasible that the same error message could stem from several areas of the code complicating the problem even further.

If an R error occurs the programmer needs to know the cause so that this can be communicated back to the Java user. Errors are normally attributed to incomplete or incompatible input data, and hence a message needs to be displayed so that the user can rectify the problem.

If there happens to be a syntax error, the Java may only receive notification of an IO Exception. It is even worse if the R is waiting for a command - for example, an end of curly bracket - that was never sent. The Java is

waiting for the R to end which is not going to happen and the entire process halts.

This can only be tracked and fixed whilst in the development environment and hence it is essential the the R and Java interaction is tested thoroughly before the product is shipped.

The R process is run in a separate thread to the Java. This is so that the user can interact with the Java application while the R is running. Here are a few important items for consideration:

1. The user cannot be allowed to change anything in Java that will conflict with the R process or make parameters on the Java screen inconsistent with the results displayed.
2. It is feasible that the user will wish to run two or more R processes simultaneously. The threading and communication will need to be safe.
3. Priority of computing processing resources.

Keeping R and Java working together well requires a lot of effort.

9.5 Testing the installation procedure

Testing the installation procedure is not about finding errors in the application - it is to find problems in the installation procedure. The installation will be the customers' first experience with the product, and hence it is important to make it professional [18].

Software systems require the locations of files, libraries, and other applications, to work properly. The systems also needs a mechanism for launching.

There is often the added implication of licensing the software which inevitably causes an extra risk to a smooth installation.

Information technology systems have administrative and security settings, and these may conflict with any part of the system to be installed. The hardware specification - for example, memory space and processing power - must also be compatible.

It is normal, or common sense, to run through or recommend a test process after installation to check that all components work as expected. This test process will include visiting screens to ensure icons are accessed properly and the program code has not become corrupt. It will check the saving and loading of files several times - ensuring that the file I/O works at different locations, and if necessary different drives. All I/O processes need to be tested - this means checking printing devices, accessing databases, integration with other required processes, and so on.

A software company must ensure that the installation documentation is correct, easy to understand, and is in good English.

Chapter 10

Conclusion

Within this project, a prototype software application was developed to support commodity traders in making decisions under uncertain market conditions. To do this, a Bayesian approach is used for forecasting future commodity prices and this is combined with utility theory which is applied to control risk.

A Bayesian methodology is adopted since there is a focus on including the trader's expertise and knowledge and relevant market information into the forecasting model via the trader's prior expectations of the development of future prices. The forecast is then revised by applying Bayes Linear Updating techniques as prices become observed. The trader can intervene, at any time, to include new information - for example, nuclear power station catches fire - which a computer program can not incorporate into the price forecasting model itself.

Using Bayes linear methods means that expectations needs to be elicited from the trader. To do this, a software component was built. Here, the requirement is to obtain precise forecasts from the user whilst minimizing cognitive bias in the trader's price expectations. It is also critical to minimize the time it takes the trader to enter his price outlook. Creating an

effective way to do this is a difficult process with a sparse amount of academic literature on the subject and few practical examples of prior elicitation available. The screens that have been developed, so far, will inevitably be improved and it is expected that criticism and feedback from traders will be used as a starting point for better user interface designs.

Several diagnostics - mainly based on Bayes Linear methods - were developed. The purpose of these diagnostics is to warn the trader when the observed prices do not match with the price forecasting model. That is, when the price forecasting model is statistically unlikely to be precise given actual prices. Here, there are difficulties in assessing where to put the threshold values - the levels at which the diagnostics need to be recorded before the trader is warned that there is a possible problem. There is also a problem with the Layman's limited understanding of what the diagnostic readings indicate.

There is scope to develop more statistical diagnostics and to improve on the benefits of the existing diagnostics by producing documentation that a non-statistician can understand.

Utility theory has been adopted as the approach to manage risk. In particular, an exponential utility function is used to model a trader's risk attitude. Objections to the exponential utility function include its property of constant absolute risk aversion and its lack of detail and simplistic nature. Nevertheless, the exponential utility is easy to manipulate mathematically and it is a better approximation of a trader's risk behaviour than nothing at all.

Building an individual's utility function requires information about how the individual responds to risk-reward type questions and a prototype elicitation tool was developed to do this. Unfortunately, the majority of people

who used the tool thought that the risk-reward questions were non-intuitive, and hence found that the utility elicitation tool was difficult to use.

Utility elicitation is challenging since many people dislike probability type questions. It is not easy to say precisely what is needed to improve the approach to this utility elicitation problem to date. Observation of a trader's day to day risk management processes and feedback from the users of the tool have been used to some extent to recommend and make improvements to the utility component of the tool.

Writing clear and detailed documentation and providing software support will be necessary and useful for the end users of TradeABLE. The mathematics behind the tool are advanced, and while it unnecessary for the user to be an expert in the statistics used, it is essential that an overview of the statistical methods used is articulated to the trader.

Comprehension of the techniques used is also paramount for internal Energy Scitech Ltd staff. Perhaps, a training course for the software needs to be developed and this could be used as a internal induction to the tool. This training course could also be a service provided to external customers. A consultancy service - to show the trader how to optimize the use of the tool - could also be set-up when the tool is out in the marketplace.

For the future, there are endless options for adding functionality to the tool and making developments. Of course, fixing bugs, improving the look and feel of screens, and keeping up to date with technological advances is a normal part of the software development process. Innovating and progressing the underlying capabilities of the tool will also be necessary. Perhaps this will include adding diagnostics, commodity portfolio management, optimizing non-exponential utility functions, and considering new price forecasting models. Responding to criticism and feedback from traders will undoubtedly

take the forefront in all software development decisions.

It is time to manage uncertainty effectively. TradeABLE presents a solid approach to support the commodity trader who makes important buy, hold, or sell decisions in today's volatile market conditions.

Bibliography

- [1] Michelle C. Baddeley, Andrew Curtis, and Rachel A. Wood. An introduction to prior information derived from probabilistic judgements: Elicitation of knowledge, cognitive bias and herding. *Geological Society, London*, 239:15–27, 2004.
- [2] P Coolen-Schrijner and Michael Goldstein. Observational Uncertainty for Forecasts. University of Durham. Confidential.
- [3] D.J.Spiegelhalter, K.R.Abrams, and J.P.Myles. *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*, chapter 5, pages 139–149. John Wiley and Sons, Ltd, 2004.
- [4] Edwin J. Elton, Martin J. Gruber, Stephen J.Brown, and William N. Goetzmann. *Modern Portfolio Theory and Investment Analysis*. John Wiley and Sons, 2003.
- [5] Peter H. Farquhar. Utility Assessment Methods. *Management Science*, 30(11):1283–1300, 1984.
- [6] Hein Fennema and Marcel Van Assen. Measuring the Utility of Losses by Means of the Tradeoff Method. *Risk and Uncertainty*, 17(3):277–295, 1999.
- [7] Michael Goldstein. *Encyclopedia of Statistical Sciences*, volume 3, chapter Bayes Linear Analysis, pages 29–34. John Wiley and Sons.
- [8] Michael Goldstein. Utility and Risk. University of Durham. Confidential.
- [9] Michael Goldstein and Allan H.Seheult. A Simple Utility Trader. July 2004.
- [10] Michael Goldstein and Allan Seheult. Analysis of Trading Rules. University of Durham. Confidential.

- [11] Michael Goldstein and Allan Seheult. Viability Analysis of the - modified limited trading rule applied to the - closing prices. University of Durham. Confidential.
- [12] Michael Goldstein and Allan H. Seheult. Prior Viability Assessment for Bayesian Analysis. 2005 In Submission.
- [13] Michael Goldstein and David Wooff. *Bayes Linear Statistics*. Chichester:Wiley, 2006 Forthcoming.
- [14] Heriot-Watt University. *Heriot-Watt Course Financial Mathematics III Notes*.
- [15] Richard J. Teweles and Frank J. Jones. *The Futures Game Who Wins, Who Loses, and Why*. McGraw-Hill, 1999.
- [16] Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental Tests of the Endowment Effect and Coase Theorem. *The Journal of Political Economy*, 98(6):1325–1348, 1990.
- [17] Daniel Kahneman and Amos Tversky. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47(2):263–292, 1979.
- [18] Cem Kaner, Jack Falk, and Hung Quoc Nguyen. *Testing Computer Software*. John Wiley and Sons, 1999.
- [19] Aik-Hiang Lau and Tze-Yun Leong. PROBES: A Framework for Probability Elicitation from Experts. *The American Medical Informatics Association*, pages 301–305, 1999.
- [20] Harry Markowitz. The Utility of Wealth. *The Journal of Political Economy*, pages 151–158, 1952.
- [21] Glenford J. Myers. *The Art of Software Testing*. John Wiley and Sons.
- [22] Anthony O'Hagan. Elicitation. *Significance*, 2(2):84–86.
- [23] Anthony O'Hagan. Eliciting Expert Beliefs in Substantial Practical Applications. *The Statistician*, 47(1):21–35, 1998.
- [24] L. D. Philips. Group elicitation of probability distributions: are many heads better than one? *Bayesian Research Conference; Decision Science and Technology, Reflections on the Contributions of Ward Edwards, Los Angeles: Kluwer Academic, Boston and London.*, pages 313–330, February 1999.

- [25] Andy Pole, Mike West, and Jeff Harrison. *Applied Bayesian Forecasting and Time Series Analysis*. Chapman and Hall/CRC, 1994.
- [26] George J Stigler. The Development of Utility Theory I. *The Journal of Political Economy*, 58(4):307–327, August 1950.
- [27] George J Stigler. The Development of Utility Theory II. *The Journal of Political Economy*, 58(5):373–396, October 1950.
- [28] van Lenthe J. Scoring-rule feedforward and the elicitation of subjective probability distributions. *Organizational behaviour and human decision processes*, 59(2):188–209, 1994.
- [29] David Wooff. Trading Algorithms: Technical report, University of Durham, Confidential. 2006.

