

A Practical Guide to Forecasting
Financial Market Volatility

Ser-Huang Poon



John Wiley & Sons, Ltd

A Practical Guide to Forecasting
Financial Market Volatility

For other titles in the Wiley Finance series
please see www.wiley.com/finance

A Practical Guide to Forecasting
Financial Market Volatility

Ser-Huang Poon



John Wiley & Sons, Ltd

Copyright © 2005 John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England
Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The Publisher is not associated with any product or vendor mentioned in this book.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA
Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA
Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany
John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia
John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809
John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging-in-Publication Data

Poon, Ser-Huang.

A practical guide for forecasting financial market volatility / Ser Huang
Poon.

p. cm. — (The Wiley finance series)

Includes bibliographical references and index.

ISBN-13 978-0-470-85613-0 (cloth : alk. paper)

ISBN-10 0-470-85613-0 (cloth : alk. paper)

1. Options (Finance)—Mathematical models. 2. Securities—Prices—

Mathematical models. 3. Stock price forecasting—Mathematical models. I. Title.

II. Series.

HG6024.A3P66 2005

332.64'01'5195—dc22

2005005768

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN-13 978-0-470-85613-0 (HB)

ISBN-10 0-470-85613-0 (HB)

Typeset in 11/13pt Times by TechBooks, New Delhi, India

Printed and bound in Great Britain by TJ International Ltd, Padstow, Cornwall

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

I dedicate this book to my mother

Contents

Foreword by Clive Granger	xiii
Preface	xv
1 Volatility Definition and Estimation	1
1.1 What is volatility?	1
1.2 Financial market stylized facts	3
1.3 Volatility estimation	10
1.3.1 Using squared return as a proxy for daily volatility	11
1.3.2 Using the high–low measure to proxy volatility	12
1.3.3 Realized volatility, quadratic variation and jumps	14
1.3.4 Scaling and actual volatility	16
1.4 The treatment of large numbers	17
2 Volatility Forecast Evaluation	21
2.1 The form of X_t	21
2.2 Error statistics and the form of ε_t	23
2.3 Comparing forecast errors of different models	24
2.3.1 Diebold and Mariano’s asymptotic test	26
2.3.2 Diebold and Mariano’s sign test	27
2.3.3 Diebold and Mariano’s Wilcoxon sign-rank test	27
2.3.4 Serially correlated loss differentials	28
2.4 Regression-based forecast efficiency and orthogonality test	28
2.5 Other issues in forecast evaluation	30

3	Historical Volatility Models	31
3.1	Modelling issues	31
3.2	Types of historical volatility models	32
3.2.1	Single-state historical volatility models	32
3.2.2	Regime switching and transition exponential smoothing	34
3.3	Forecasting performance	35
4	Arch	37
4.1	Engle (1982)	37
4.2	Generalized ARCH	38
4.3	Integrated GARCH	39
4.4	Exponential GARCH	41
4.5	Other forms of nonlinearity	41
4.6	Forecasting performance	43
5	Linear and Nonlinear Long Memory Models	45
5.1	What is long memory in volatility?	45
5.2	Evidence and impact of volatility long memory	46
5.3	Fractionally integrated model	50
5.3.1	FIGARCH	51
5.3.2	FIEGARCH	52
5.3.3	The positive drift in fractional integrated series	52
5.3.4	Forecasting performance	53
5.4	Competing models for volatility long memory	54
5.4.1	Breaks	54
5.4.2	Components model	55
5.4.3	Regime-switching model	57
5.4.4	Forecasting performance	58
6	Stochastic Volatility	59
6.1	The volatility innovation	59
6.2	The MCMC approach	60
6.2.1	The volatility vector H	61
6.2.2	The parameter w	62
6.3	Forecasting performance	63
7	Multivariate Volatility Models	65
7.1	Asymmetric dynamic covariance model	65

7.2	A bivariate example	67
7.3	Applications	68
8	Black–Scholes	71
8.1	The Black–Scholes formula	71
8.1.1	The Black–Scholes assumptions	72
8.1.2	Black–Scholes implied volatility	73
8.1.3	Black–Scholes implied volatility smile	74
8.1.4	Explanations for the ‘smile’	75
8.2	Black–Scholes and no-arbitrage pricing	77
8.2.1	The stock price dynamics	77
8.2.2	The Black–Scholes partial differential equation	77
8.2.3	Solving the partial differential equation	79
8.3	Binomial method	80
8.3.1	Matching volatility with u and d	83
8.3.2	A two-step binomial tree and American-style options	85
8.4	Testing option pricing model in practice	86
8.5	Dividend and early exercise premium	88
8.5.1	Known and finite dividends	88
8.5.2	Dividend yield method	88
8.5.3	Barone-Adesi and Whaley quadratic approximation	89
8.6	Measurement errors and bias	90
8.6.1	Investor risk preference	91
8.7	Appendix: Implementing Barone-Adesi and Whaley’s efficient algorithm	92
9	Option Pricing with Stochastic Volatility	97
9.1	The Heston stochastic volatility option pricing model	98
9.2	Heston price and Black–Scholes implied	99
9.3	Model assessment	102
9.3.1	Zero correlation	103
9.3.2	Nonzero correlation	103
9.4	Volatility forecast using the Heston model	105
9.5	Appendix: The market price of volatility risk	107
9.5.1	Ito’s lemma for two stochastic variables	107
9.5.2	The case of stochastic volatility	107
9.5.3	Constructing the risk-free strategy	108

9.5.4	Correlated processes	110
9.5.5	The market price of risk	111
10	Option Forecasting Power	115
10.1	Using option implied standard deviation to forecast volatility	115
10.2	At-the-money or weighted implied?	116
10.3	Implied biasedness	117
10.4	Volatility risk premium	119
11	Volatility Forecasting Records	121
11.1	Which volatility forecasting model?	121
11.2	Getting the right conditional variance and forecast with the 'wrong' models	123
11.3	Predictability across different assets	124
11.3.1	Individual stocks	124
11.3.2	Stock market index	125
11.3.3	Exchange rate	126
11.3.4	Other assets	127
12	Volatility Models in Risk Management	129
12.1	Basel Committee and Basel Accords I & II	129
12.2	VaR and backtest	131
12.2.1	VaR	131
12.2.2	Backtest	132
12.2.3	The three-zone approach to backtest evaluation	133
12.3	Extreme value theory and VaR estimation	135
12.3.1	The model	136
12.3.2	10-day VaR	137
12.3.3	Multivariate analysis	138
12.4	Evaluation of VaR models	139
13	VIX and Recent Changes in VIX	143
13.1	New definition for VIX	143
13.2	What is the VXO?	144
13.3	Reason for the change	146
14	Where Next?	147

Appendix	149
References	201
Index	215

Foreword

If one invests in a financial asset today the return received at some pre-specified point in the future should be considered as a random variable. Such a variable can only be fully characterized by a distribution function or, more easily, by a density function. The main, single and most important feature of the density is the expected or mean value, representing the location of the density. Around the mean is the uncertainty or the volatility. If the realized returns are plotted against time, the jagged oscillating appearance illustrates the volatility. This movement contains both welcome elements, when surprisingly large returns occur, and also certainly unwelcome ones, the returns far below the mean. The well-known fact that a poor return can arise from an investment illustrates the fact that investing can be risky and is why volatility is sometimes equated with risk.

Volatility is itself a stock variable, having to be measured over a period of time, rather than a flow variable, measurable at any instant of time. Similarly, a stock price is a flow variable but a return is a stock variable. Observed volatility has to be observed over stated periods of time, such as hourly, daily, or weekly, say.

Having observed a time series of volatilities it is obviously interesting to ask about the properties of the series: is it forecastable from its own past, do other series improve these forecasts, can the series be modeled conveniently and are there useful multivariate generalizations of the results? Financial econometricians have been very inventive and industrious considering such questions and there is now a substantial and often sophisticated literature in this area.

The present book by Professor Ser-Huang Poon surveys this literature carefully and provides a very useful summary of the results available.

By so doing, she allows any interested worker to quickly catch up with the field and also to discover the areas that are still available for further exploration.

Clive W.J. Granger
December 2004

Preface

Volatility forecasting is crucial for option pricing, risk management and portfolio management. Nowadays, volatility has become the subject of trading. There are now exchange-traded contracts written on volatility. Financial market volatility also has a wider impact on financial regulation, monetary policy and macroeconomy. This book is about financial market volatility forecasting. The aim is to put in one place models, tools and findings from a large volume of published and working papers from many experts. The material presented in this book is extended from two review papers ('Forecasting Financial Market Volatility: A Review' in the *Journal of Economic Literature*, 2003, 41, 2, pp. 478–539, and 'Practical Issues in Forecasting Volatility' in the *Financial Analysts Journal*, 2005, 61, 1, pp. 45–56) jointly published with Clive Granger.

Since the main focus of this book is on volatility forecasting performance, only volatility models that have been tested for their forecasting performance are selected for further analysis and discussion. Hence, this book is oriented towards practical implementations. Volatility models are not pure theoretical constructs. The practical importance of volatility modelling and forecasting in many finance applications means that the success or failure of volatility models will depend on the characteristics of empirical data that they try to capture and predict. Given the prominent role of option price as a source of volatility forecast, I have also devoted much effort and the space of two chapters to cover Black–Scholes and stochastic volatility option pricing models.

This book is intended for first- and second-year finance PhD students and practitioners who want to implement volatility forecasting models but struggle to comprehend the huge volume of volatility research. Readers who are interested in more technical aspects of volatility modelling

could refer to, for example, Gouriéroux (1997) on ARCH models, Shephard (2003) on stochastic volatility and Fouque, Papanicolaou and Sircar (2000) on stochastic volatility option pricing. Books that cover specific aspects or variants of volatility models include Franses and van Dijk (2000) on nonlinear models, and Beran (1994) and Robinson (2003) on long memory models. Specialist books that cover financial time series modelling in a more general context include Alexander (2001), Tsay (2002) and Taylor (2005). There are also a number of edited series that contain articles on volatility modelling and forecasting, e.g. Rossi (1996), Knight and Satchell (2002) and Jarrow (1998).

I am very grateful to Clive for his teaching and guidance in the last few years. Without his encouragement and support, our volatility survey works and this book would not have got started. I would like to thank all my co-authors on volatility research, in particular Bevan Blair, Namwon Hyung, Eric Jondeau, Martin Martens, Michael Rockinger, Jon Tawn, Stephen Taylor and Konstantinos Vonatsos. Much of the writing here reflects experience gained from joint work with them.

Volatility Definition and Estimation

1.1 WHAT IS VOLATILITY?

It is useful to start with an explanation of what volatility is, at least for the purpose of clarifying the scope of this book. Volatility refers to the spread of all likely outcomes of an uncertain variable. Typically, in financial markets, we are often concerned with the spread of asset returns. Statistically, volatility is often measured as the sample standard deviation

$$\hat{\sigma} = \sqrt{\frac{1}{T-1} \sum_{t=1}^T (r_t - \mu)^2}, \quad (1.1)$$

where r_t is the return on day t , and μ is the average return over the T -day period.

Sometimes, variance, σ^2 , is used also as a volatility measure. Since variance is simply the square of standard deviation, it makes no difference whichever measure we use when we compare the volatility of two assets. However, variance is much less stable and less desirable than standard deviation as an object for computer estimation and volatility forecast evaluation. Moreover standard deviation has the same unit of measure as the mean, i.e. if the mean is in dollar, then standard deviation is also expressed in dollar whereas variance will be expressed in dollar square. For this reason, standard deviation is more convenient and intuitive when we think about volatility.

Volatility is related to, but not exactly the same as, risk. Risk is associated with undesirable outcome, whereas volatility as a measure strictly for uncertainty could be due to a positive outcome. This important difference is often overlooked. Take the Sharpe ratio for example. The Sharpe ratio is used for measuring the performance of an investment by comparing the mean return in relation to its 'risk' proxy by its volatility.

The Sharpe ratio is defined as

$$\text{Sharpe ratio} = \frac{\left(\text{Average return, } \mu \right) - \left(\text{Risk-free interest rate, e.g. T-bill rate} \right)}{\text{Standard deviation of returns, } \sigma}.$$

The notion is that a larger Sharpe ratio is preferred to a smaller one. An unusually large positive return, which is a desirable outcome, could lead to a reduction in the Sharpe ratio because it will have a greater impact on the standard deviation, σ , in the denominator than the average return, μ , in the numerator.

More importantly, the reason that volatility is not a good or perfect measure for risk is because volatility (or standard deviation) is only a measure for the spread of a distribution and has no information on its shape. The only exception is the case of a normal distribution or a lognormal distribution where the mean, μ , and the standard deviation, σ , are sufficient statistics for the entire distribution, i.e. with μ and σ alone, one is able to reproduce the empirical distribution.

This book is about volatility only. Although volatility is not the sole determinant of asset return distribution, it is a key input to many important finance applications such as investment, portfolio construction, option pricing, hedging, and risk management. When Clive Granger and I completed our survey paper on volatility forecasting research, there were 93 studies on our list plus several hundred non-forecasting papers written on volatility modelling. At the time of writing this book, the number of volatility studies is still rising and there are now about 120 volatility forecasting papers on the list. Financial market volatility is a 'live' subject and has many facets driven by political events, macroeconomy and investors' behaviour. This book will elaborate some of these complexities that kept the whole industry of volatility modelling and forecasting going in the last three decades. A new trend now emerging is on the trading and hedging of volatility. The Chicago Board of Exchange (CBOE) for example has started futures trading on a volatility index. Options on such futures contracts are likely to follow. Volatility swap contracts have been traded on the over-the-counter market well before the CBOE's developments. Previously volatility was an input to a model for pricing an asset or option written on the asset. It is now the principal subject of the model and valuation. One can only predict that volatility research will intensify for at least the next decade.

1.2 FINANCIAL MARKET STYLIZED FACTS

To give a brief appreciation of the amount of variation across different financial assets, Figure 1.1 plots the returns distributions of a normally

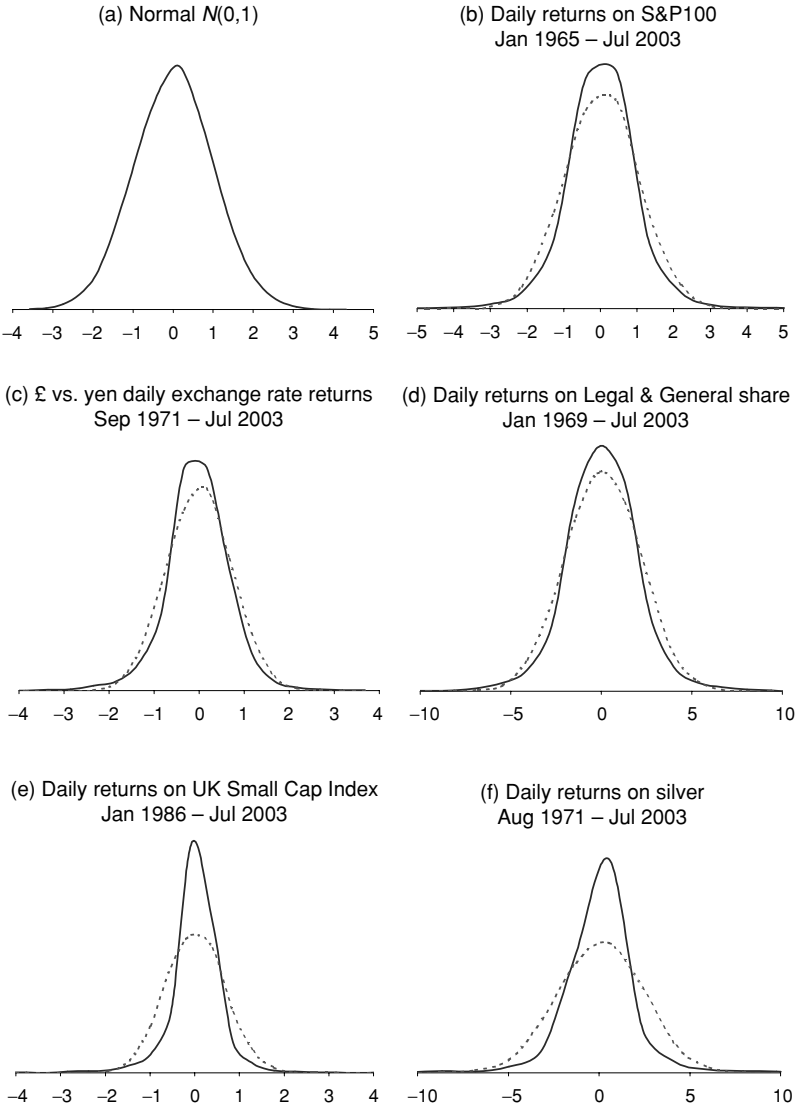


Figure 1.1 Distribution of daily financial market returns. (*Note: the dotted line is the distribution of a normal random variable simulated using the mean and standard deviation of the financial asset returns*)

distributed random variable, and the respective daily returns on the US Standard and Poor market index (S&P100),¹ the yen–sterling exchange rate, the share of Legal & General (a major insurance company in the UK), the UK Index for Small Capitalisation Stocks (i.e. small companies), and silver traded at the commodity exchange. The normal distribution simulated using the mean and standard deviation of the financial asset returns is drawn on the same graph to facilitate comparison.

From the small selection of financial asset returns presented in Figure 1.1, we notice several well-known features. Although the asset returns have different degrees of variation, most of them have long ‘tails’ as compared with the normally distributed random variable. Typically, the asset distribution and the normal distribution cross at least three times, leaving the financial asset returns with a longer left tail and a higher peak in the middle. The implications are that, for a large part of the time, financial asset returns fluctuate in a range smaller than a normal distribution. But there are some occasions where financial asset returns swing in a much wider scale than that permitted by a normal distribution. This phenomenon is most acute in the case of UK Small Cap and silver. Table 1.1 provides some summary statistics for these financial time series.

The normally distributed variable has a skewness equal to zero and a kurtosis of 3. The annualized standard deviation is simply $\sqrt{252}\sigma$, assuming that there are 252 trading days in a year. The financial asset returns are not adjusted for dividend. This omission is not likely to have any impact on the summary statistics because the amount of dividends distributed over the year is very small compared to the daily fluctuations of asset prices. From Table 1.1, the Small Cap Index is the most negatively skewed, meaning that it has a longer left tail (extreme losses) than right tail (extreme gains). Kurtosis is a measure for tail thickness and it is astronomical for S&P100, Small Cap Index and silver. However, these skewness and kurtosis statistics are very sensitive to outliers. The skewness statistic is much closer to zero, and the amount of kurtosis dropped by 60% to 80%, when the October 1987 crash and a small number of outliers are excluded.

Another characteristic of financial market volatility is the time-varying nature of returns fluctuations, the discovery of which led to Rob Engle’s Nobel Prize for his achievement in modelling it. Figure 1.2 plots the time series history of returns of the same set of assets presented

¹ The data for S&P100 prior to 1986 comes from S&P500. Adjustments were made when the two series were grafted together.

Table 1.1 Summary statistics for a selection of financial series

	$N(0, 1)$	S&P100	Yen/£ rate	Legal & General	UK Small Cap	Silver
Start date		Jan 65	Sep 71	Jan 69	Jan 86	Aug 71
Number of observations	8000	9675	7338	7684	4432	7771
Daily average ^a	0	0.024	-0.021	0.043	0.022	0.014
Daily Standard Deviation	1	0.985	0.715	2.061	0.648	2.347
Annualized average	0	6.067	-5.188	10.727	5.461	3.543
Annualized Standard Deviation	15.875	15.632	11.356	32.715	10.286	37.255
Skewness	0	-1.337	-0.523	0.026	-3.099	0.387
Kurtosis	3	37.140	7.664	6.386	42.561	45.503
Number of outliers removed		1			5	9
Skewness ^b		-0.055			-0.917	-0.088
Kurtosis ^{a,b}		7.989			13.972	15.369

^a Returns not adjusted for dividends.

^b These two statistical measures are computed after the removal of outliers. All series have an end date of 22 July, 2003.




Figure 1.2 Time series of daily returns on a simulated random variable and a collection of financial assets