


The background of the book cover is composed of several overlapping geometric shapes. On the left side, there is a vertical column of squares in shades of light blue, medium blue, dark purple, magenta, pink, and red. To the right of this column, there are horizontal and vertical stripes in black and white. The top right corner is a solid light blue rectangle. The bottom right corner is a solid orange rectangle. The central area is a mix of these colors and patterns.

ANTICIPATING CORRELATIONS

A New Paradigm for Risk Management

ROBERT F. ENGLE

Anticipating Correlations

This page intentionally left blank

Anticipating Correlations

A New Paradigm for Risk Management

Robert Engle

Princeton University Press
Princeton and Oxford

Copyright © 2009 by Princeton University Press

Published by Princeton University Press,
41 William Street, Princeton, New Jersey 08540

In the United Kingdom: Princeton University Press,
6 Oxford Street, Woodstock, Oxfordshire OX20 1TW

All Rights Reserved

ISBN: 978-0-691-11641-9 (alk. paper)

Library of Congress Control Number: 2008939934

British Library Cataloging-in-Publication Data is available

This book has been composed in LucidaBright using T_EX
Typeset and copyedited by T&T Productions Ltd, London

Printed on acid-free paper ∞

press.princeton.edu

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

Introduction	vii
1 Correlation Economics	1
1.1 Introduction	1
1.2 How Big Are Correlations?	3
1.3 The Economics of Correlations	6
1.4 An Economic Model of Correlations	9
1.5 Additional Influences on Correlations	13
2 Correlations in Theory	15
2.1 Conditional Correlations	15
2.2 Copulas	17
2.3 Dependence Measures	21
2.4 On the Value of Accurate Correlations	25
3 Models for Correlation	29
3.1 The Moving Average and the Exponential Smoother	30
3.2 Vector GARCH	32
3.3 Matrix Formulations and Results for Vector GARCH	33
3.4 Constant Conditional Correlation	37
3.5 Orthogonal GARCH	37
3.6 Dynamic Conditional Correlation	39
3.7 Alternative Approaches and Expanded Data Sets	41
4 Dynamic Conditional Correlation	43
4.1 DE-GARCHING	43
4.2 Estimating the Quasi-Correlations	45
4.3 Rescaling in DCC	48
4.4 Estimation of the DCC Model	55
5 DCC Performance	59
5.1 Monte Carlo Performance of DCC	59
5.2 Empirical Performance	61
6 The MacGyver Method	74

7	Generalized DCC Models	80
7.1	Theoretical Specification	80
7.2	Estimating Correlations for Global Stock and Bond Returns	83
8	FACTOR DCC	88
8.1	Formulation of Factor Versions of DCC	88
8.2	Estimation of Factor Models	93
9	Anticipating Correlations	103
9.1	Forecasting	103
9.2	Long-Run Forecasting	108
9.3	Hedging Performance In-Sample	111
9.4	Out-of-Sample Hedging	112
9.5	Forecasting Risk in the Summer of 2007	117
10	Credit Risk and Correlations	122
11	Econometric Analysis of the DCC Model	130
11.1	Variance Targeting	130
11.2	Correlation Targeting	131
11.3	Asymptotic Distribution of DCC	134
12	Conclusions	137
	References	141
	Index	151

Introduction

The Econometric Institute Lecture Series deals with topics in econometrics that have important policy implications. The lectures cover a wide range of topics and are not confined to any one area or subdiscipline. Leading international scientists in the fields of econometrics in which applications play a major role are invited to give three-day lectures on a topic to which they have contributed significantly.

The topic of Robert Engle's lectures deals with the dynamics of correlations between a large number of financial assets. In the present global financial world it is imperative both for asset management and for risk analysis to gain a better understanding of the changing correlations between a large number of assets and even between different financial markets. In Robert Engle's book several innovative models are proposed and tested with respect to their forecasting performance in the turbulent economic world of 2007.

As editors of the series we are indebted to the Erasmus University Trust Fund, the Erasmus Research Institute of Management, and the Tinbergen Institute for continued support for the series.

Philip Hans Franses and Herman K. van Dijk

*Econometric Institute
Erasmus School of Economics*

This page intentionally left blank

Anticipating Correlations

This page intentionally left blank

Correlation Economics

1.1 Introduction

Today there are almost three thousand stocks listed on the New York Stock Exchange. NASDAQ lists another three thousand. There is yet another collection of stocks that are unlisted and traded on the Bulletin Board or Pink Sheets. These U.S.-traded stocks are joined by thousands of companies listed on foreign stock exchanges to make up a universe of publicly traded equities. Added to these are the enormous number of government and corporate and municipal bonds that are traded in the United States and around the world, as well as many short-term securities. Investors are now exploring a growing number of alternative asset classes each with its own large set of individual securities. On top of these underlying assets is a web of derivative contracts. It is truly a vast financial arena. A portfolio manager faces a staggering task in selecting investments.

The prices of all of these assets are constantly changing in response to news and in anticipation of future performance. Every day many stocks rise in value and many decline. The movements in price are, however, not independent. If they were independent, then it would be possible to form a portfolio with negligible volatility. Clearly this is not the case. The correlation structure across assets is a key feature of the portfolio choice problem because it is instrumental in determining the risk. Recognizing that the economy is an interconnected set of economic agents, sometimes considered a general equilibrium system, it is hardly surprising that movements in asset prices are correlated. Estimating the correlation structure of thousands of assets and using this to select superior portfolios is a Herculean task. It is especially difficult when it is recognized that these correlations vary over time, so that a forward-looking correlation estimator is needed. This problem is the focus of this book. We must “anticipate correlations” if we want to have optimal risk management, portfolio selection, and hedging.

Such forward-looking correlations are very important in risk management because the risk of a portfolio depends not on what the correlations were in the past, but on what they will be in the future. Similarly, portfolio choice depends on forecasts of asset dependence structure. Many aspects of financial planning involve hedging one asset with a collection of others. The optimal hedge will also depend upon the correlations and volatilities to be expected over the future holding period. An even more complex problem arises when it is recognized that the correlations can be forecast many periods into the future. Consequently, there are predictable changes in the risk–return trade-off that can be incorporated into optimal portfolios.

Derivatives such as options are now routinely traded not only on individual securities, but also on baskets and indices. Such derivative prices are related to the derivative prices of the component assets, but the relation depends on the correlations expected to prevail over the life of the derivative. A market for correlation swaps has recently developed that allows traders to take a position in the average correlation over a time interval. Structured products form a very large class of derivatives that are sensitive to correlations. An important example of a structured product is the collateralized debt obligation (CDO), which in its simplest form is a portfolio of corporate bonds that is sold to investors in tranches that have different risk characteristics. In this way credit risks can be bought and sold to achieve specific risk–return targets. There are many types of CDOs backed by loans, mortgages, subprime mortgages, credit default swaps, tranches of CDOs themselves, and many other assets. In these securities, the correlations between defaults are the key determinants of valuations. Because of the complexity of these structures and the difficulty in forecasting correlations and default correlations, it has been difficult to assess the risks of the tranches that are supposed to be low risk. Some of the “credit crunch” of 2007–8 can probably be attributed to this failure in risk management. This episode serves to reinforce the importance of anticipating correlations.

This book will introduce and carefully explain a collection of new methods for estimating and forecasting correlations for large systems of assets. The book initially discusses the economics of correlations. Then it turns to the measurement of comovement and dependence by correlations and alternative measures. A look at existing models for estimating correlations—such as historical correlation, exponential smoothing, and multivariate GARCH—leads to the introduction (in chapter 3) of the central method explored in the book: dynamic conditional correlation. Monte Carlo and empirical analyses of this model document its performance. Successive chapters deal with extensions to the basic model, new

Table 1.1. Correlations of large-cap stocks from 1998 to 2003.

	IBM	MCD	GE	Citibank	AXP	WMT	SP500
IBM	1.000	0.192	0.436	0.419	0.387	0.283	0.600
MCD	0.192	1.000	0.308	0.238	0.282	0.303	0.365
GE	0.436	0.308	1.000	0.595	0.614	0.484	0.760
Citibank	0.419	0.238	0.595	1.000	0.697	0.439	0.740
AXP	0.387	0.282	0.614	0.697	1.000	0.445	0.715
WMT	0.283	0.303	0.484	0.439	0.445	1.000	0.584
SP500	0.600	0.365	0.760	0.740	0.715	0.584	1.000

estimation methods, and a technical discussion of some econometric issues. Many empirical studies are documented in particular chapters, including stock-bond correlations, global equity correlations, and U.S. large-cap stock correlations. Finally, in a chapter called “Anticipating Correlations,” these methods are used to forecast correlations through the turbulent environment of the summer and autumn of 2007.

The methods introduced in this book are simple, powerful, and will be shown to be highly stable over time. They offer investors and money managers up-to-date measures of volatilities and correlations that can be used to assess risk and optimize investment decisions even in the complex and high-dimensional world we inhabit.

1.2 How Big Are Correlations?

Correlations must all lie between -1 and 1 , but the actual size varies dramatically across assets and over time. For example, using daily data for the six-year period from 1998 through 2003 and the textbook formula

$$\hat{\rho}_{x,y} = \frac{\sum_{t=1}^T (x_t - \bar{x})(y_t - \bar{y})}{\sqrt{\sum_{t=1}^T (x_t - \bar{x})^2 \sum_{t=1}^T (y_t - \bar{y})^2}}, \quad (1.1)$$

it is interesting to calculate a variety of correlations. The correlation between daily returns on IBM stock and the S&P 500 measure of the broad U.S. market is 0.6 . This means that the regression of IBM returns on a constant and S&P returns will have an R^2 value of 0.36 . The systematic risk of IBM is 36% of the total variance and the idiosyncratic risk is 64% .

Looking across five large-capitalization stocks, the correlations with the S&P 500 for the six-year period range from 0.36 for McDonald’s to 0.76 for General Electric (GE). These stocks are naturally correlated with each other as well, although the correlations are typically smaller (see table 1.1).

Table 1.2. Correlations of small-cap stocks from 1998 to 2003.

	PVA	NSC	ARG	DRTK	MTLG	SP
PVA	1.000	0.159	0.050	0.063	0.014	0.185
NSC	0.159	1.000	0.253	0.006	0.034	0.445
ARG	0.050	0.253	1.000	0.068	0.081	0.326
DRTK	0.063	0.006	0.068	1.000	0.025	0.101
MTLG	0.014	0.034	0.081	0.025	1.000	0.080
SP	0.185	0.445	0.326	0.101	0.080	1.000

A more careful examination of the correlations shows that the highest correlations are between stocks in the same industry. American Express (AXP) and Citibank have a correlation of almost 0.7 and GE has a correlation with both that is about 0.6. During this period GE had a big financial services business and therefore moved closely with banking stocks.

Examining a selection of small-cap stocks, the story is rather different. The correlations with the market factor are much lower and the correlations between stocks are lower; table 1.2 gives the results. The largest correlation with the market is 0.45 but most of the entries in the table are below 0.1.

Turning to other asset classes let us now examine the correlation between the returns on holding bonds and the returns on holding foreign currencies (see table 1.3). Notice first the low correlations between bond returns and the S&P 500 and between currency returns and the S&P 500. These asset classes are not highly correlated with each other on average.

Within asset classes, the correlations are higher. In fact the correlation between the five- and twenty-year bond returns is 0.875, which is the highest we have yet seen. The short rate has correlations of 0.3 and 0.2, respectively, with these two long rates. Within currencies, the highest correlation is 45% between the Canadian dollar and the Australian dollar, both relative to the U.S. dollar. The rest range from 15% to 25%.

When calculating correlations across countries, it is important to recognize the differences in trading times. When markets are not open at the same times, daily returns calculated from closing data can be influenced by news that appears to be on one day in one market but on the next day in the other. For example, news during U.S. business hours will influence measured Japanese equity prices only on the next day. The effect of the news that occurs when a market is closed will be seen primarily in the opening price and therefore is attributed to the following daily return. To mitigate this problem, it is common to use data that is more time aggregated to measure such correlations.

Table 1.3. Other assets.

	T3M	T5YR	T20YR	CAD/USD	GBP/USD	AUD/USD	JPY/USD	SP500
T3M	1.000	0.329	0.206	0.011	0.076	0.025	0.031	−0.031
T5YR	0.329	1.000	0.875	−0.0007	0.136	0.007	0.005	−0.057
T20YR	0.206	0.875	1.000	0.007	0.103	−0.002	−0.049	−0.016
CAD/USD	0.011	−0.0007	0.007	1.000	0.117	0.415	0.145	0.015
GBP/USD	0.076	0.136	0.103	0.117	1.000	0.253	0.224	−0.018
AUD/USD	0.025	0.007	−0.002	0.415	0.253	1.000	0.269	0.040
JPY/USD	0.031	0.005	−0.049	0.145	0.224	0.269	1.000	−0.003
SP500	−0.031	−0.057	−0.016	0.015	−0.018	0.040	−0.003	1.000

Notes: “T3M” denotes three-month Treasury Bill returns; “T5YR” denotes five-year Treasury bond returns; “T20YR” denotes twenty-year Treasury bond returns; “CAD/USD” Canadian dollar/U.S. dollar returns; “GBP/USD” U.K. pound/U.S. dollar returns; “AUD/USD” Australian dollar/U.S. dollar returns; “JPY/USD” Japanese yen/U.S. dollar returns; “SP500” denotes Standard & Poor's 500 index of equity returns.

Cappiello et al. (2007) analyze weekly global equity and bond correlations. The data employed in their paper consist of FTSE All-World indices for twenty-one countries and DataStream-constructed five-year average maturity bond indices for thirteen, all measured relative to U.S. dollars. The sample is fifteen years of weekly price observations, for a total of 785 observations from January 8, 1987, until February 7, 2002. Table 1.4 shows a sample of global equity and bond correlations. The bond correlations are above the diagonal and the equity correlations are below the diagonal.

The equity correlations range from 0.23 to 0.73 with about a third of the sample above 0.5. The highest are between closely connected economies such as Germany, France, and Switzerland, and the United States and Canada. The bond return correlations are often much higher. France and Germany have a correlation of 0.93 and most of the European correlations are above 0.6. The U.S. correlation with Canada is 0.45, while the correlations with other countries hover around 0.2. Japanese correlations are also lower. Cappiello et al. also report correlations between equities and bonds that vary greatly. Many of these are negative. Typically, however, the domestic equity- and bond-return correlations are fairly large. This is partly due to the fact that both returns are denominated in U.S. dollars.

1.3 The Economics of Correlations

To understand the relative magnitude of all these correlations and ultimately why they change, it is important to look at the economics behind movements in asset prices. Since assets are held by investors in anticipation of payments to be made in the future, the value of an asset is intrinsically linked to forecasts of the future prospects of the project or firm. Changes in asset prices reflect changing forecasts of future payments. The information that makes us change these forecasts we often simply call “news.” This has been the basic model for changing asset prices since it was formalized by Samuelson (1965). Thus both the volatilities of asset returns and the correlations between asset returns depend on information that is used to update these distributions.

Every piece of news affects all asset prices to a greater or lesser extent. The effects are greater on some equity prices than on others because their lines of business are different. Hence the correlations in their returns due to this news event will depend upon their business. Naturally, if a firm changes its line of business, its correlations with other firms are likely to change. This is one of the most important reasons why correlations change over time.