

Evaluating “correlation breakdowns” during periods of market volatility

Mico Loretan and William B English¹

1. Introduction

In order to measure and manage market risks, financial analysts take account of the variability and correlation of the returns on assets held in their portfolios. One difficulty they encounter in doing so is that in periods of heightened market volatility, correlations between asset returns can differ substantially from those seen in quieter markets. The problem of “correlation breakdown” during periods of greater volatility is well known. For example, the former global risk manager of a major financial firm notes that, “during major market events, correlations change dramatically” (Bookstaber (1997)). A recent example of this phenomenon occurred following the Russian default in August 1998. One prominent money center bank attributed larger than expected losses in late summer and early fall last year to higher volatility, illiquidity and “breakdowns in historical correlations” (JP Morgan (1999)). Indeed, a comprehensive study found that the average correlation between five-day changes in yield spreads for 26 instruments in 10 economies rose from 0.11 in the first half of 1998 to 0.37 between mid-August and mid-September, but then fell back to 0.12 after mid-October (Bank for International Settlements (1999), Table A18).

It is tempting to explain the increased correlation of returns during hectic market periods as the result of a shift in the joint distribution of asset returns owing to contagion of some markets by others, the particular nature of the shocks, or changes in market structures and practices.² However, unless one has a model of when such periods are likely to arise, or at least how often, and what particular pattern of correlations will ensue, this approach makes it extraordinarily difficult to manage risk because the relationship between asset returns in some future situations is essentially unknown.³

Moreover, the inference that changes in measured correlations imply that the joint distribution of asset returns changes in volatile periods may not be warranted. Even if the behavior of asset returns is governed by an unchanged process, one would expect a link between volatility and measures of correlation. Indeed, a model of asset returns as simple as the bivariate normal can explain why periods of increased (sampling) volatility will also be periods of relatively high (sampling) correlations. The possible importance for economics and finance of this result has only been realized recently (Ronn (1995); Boyer, Gibson and Loretan (1999); Forbes and Rigobon (1999)).⁴ In this paper, we demonstrate that a significant portion of shifts in correlations over time – including those that occurred in the fall of 1998 – may reflect nothing more than the predictable effect of differences in sample volatilities on measured correlations, rather than breaks in the data generating process for asset returns.

To explore this possibility, we select three asset classes – equities, bonds and foreign exchange – in two representative countries and look at the quarterly correlations between daily returns over the

¹ The analysis and conclusions in this paper are those of the authors and do not indicate concurrence by other members of the research staff, by the Board of Governors, or by the Federal Reserve Banks. We thank Jim Clouse, Mike Gibson, Michael Gordy, Brian Madigan, Matt Prisker and Vince Reinhart for helpful comments and discussions. All remaining errors are ours.

² Recent discussions of possible routes for contagion include Drazen (1998), Eichengreen, Rose and Wyplosz (1996) and Gerlach and Smets (1995). The CGFS report on the events of the fall of 1998 (Bank for International Settlements (1999)) presents a narrative account of how the effects of shocks were reinforced and spread to other markets by market practices.

³ One advance along these lines is the structural model of contagion-like transmission of shocks presented in Kodres and Pritsker (1999).

⁴ Ronn (1995) attributes the insight to a conference discussion by Stambaugh.

1990s. Our calculations suggest that quarters with high correlations tend also to be quarters with higher than average volatility. Moreover, actual correlations during periods of relatively high volatility appear to be fairly close to the correlations one would expect conditional on the level of volatility and based on an unchanged process for asset returns. Our findings generalize the results reported for stock prices in Forbes and Rigobon (1999), and suggest that correlation breakdowns may reflect time-varying volatility of financial markets rather than a change in the relationships between asset returns.

Since the link between market volatility and in-sample correlations between asset prices appears to be empirically important, we go on to consider the implications of this link for risk management practices, the supervision of financial firms and the conduct of monetary policy. We note that the use of data for a relatively short period when calculating correlations for use in risk management models may lead to poor measurements of market risks. We also point to the need to use appropriate conditional correlations when examining the riskiness of a portfolio under high volatility scenarios. Finally, because monetary policy can affect the volatility of markets, monetary policymakers may find it useful to incorporate the effect of unexpected changes in policy on market participants' assessment of their risk exposures. Indeed, some monetary policymakers may, in practice, make this link: there is some evidence that monetary policy in the United States was initially adjusted relatively slowly in early 1994 because of concern that the long period of interest rate stability that preceded the rate hike had led some market participants to underestimate the riskiness of their positions.

2. The theoretical link between volatility and correlation

To see the link between volatility and correlation, consider the unrealistically simple case of two random variables, x and y , that are independently and identically distributed bivariate normal, with means equal to zero, variances equal to unity and a correlation of 0.5. A large sample of draws of such (x,y) pairs is shown in Figure 1. Now consider splitting the full sample into two subsamples based on the value of x : a “low volatility” subsample, including all (x,y) pairs with an absolute value of x less than 1.96; and a “high volatility” subsample, including all (x,y) pairs with an absolute value of x greater than or equal to 1.96.⁵ Intuitively, the effect of trimming the ends off the joint distribution in the low volatility subsample should be to reduce the sample correlation between x and y . By contrast, the correlation for the high volatility subsample should be enhanced because the support of its distribution is disjointed, with one portion picking up the large positive values of both variables while the other portion of the distribution picks up the large negative values of the variables. Indeed, as noted in the figure, the correlation for the high volatility subsample is 0.81, while that for the low volatility sample is 0.45. Note that the correlation in the latter subsample is close to the population value of 0.5; this latter result may not be surprising since the low volatility subsample includes 95% of the data.

2.1 Theoretical result

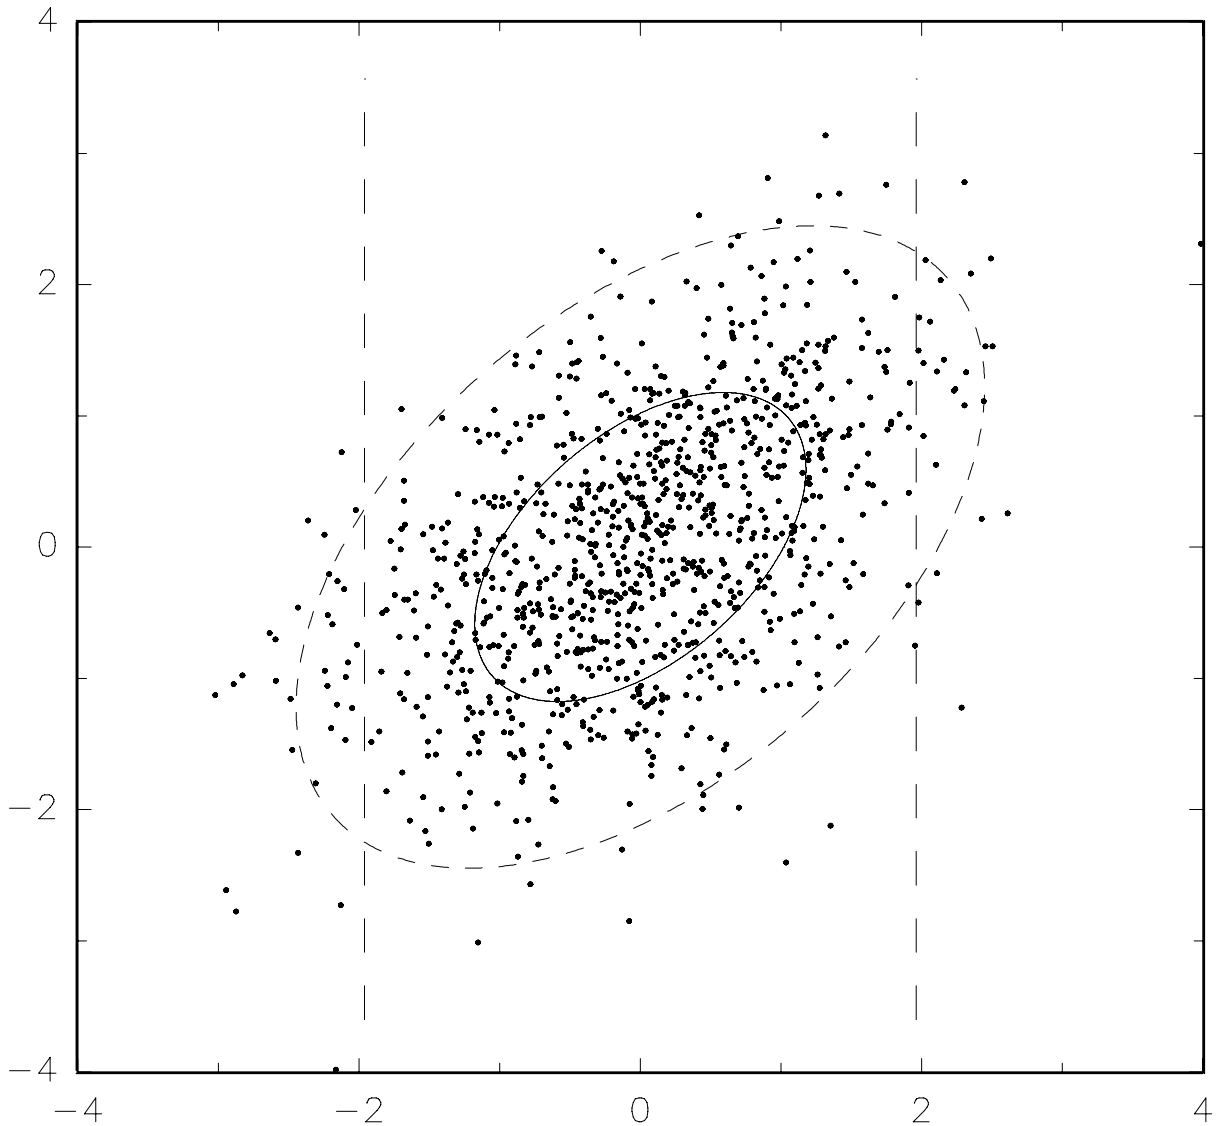
This intuitive result can be derived formally. Boyer, Gibson and Loretan (1999) provide the following theorem:

Theorem: consider a pair of bivariate normal random variables x and y with variances σ_x^2 and σ_y^2 , respectively, and covariance σ_{xy} . Let ρ ($= \sigma_{xy}/(\sigma_x\sigma_y)$) be the unconditional correlation between x and y . The correlation between x and y conditional on an event $x \in A$, for any $A \subset \mathbf{R}$ with $0 < \text{Prob}(A) < 1$, is given by:

$$(1) \quad \rho_A = \rho \left(\rho^2 + (1 - \rho^2) \frac{\text{Var}(x)}{\text{Var}(x|x \in A)} \right)^{-1/2}$$

⁵ The marginal distributions of x and y are (univariate) standard normal, hence $\text{Prob}(|x| > 1.96)$ is equal to 0.05.

Figure 1
 Bivariate normal random numbers, $\rho=0.5$



$$\rho((x,y) \mid |x| < 1.96) = 0.45, \quad \rho((x,y) \mid |x| \geq 1.96) = 0.81$$

Proof:⁶ Let u and v be two independent standard normal random variables. Now construct two bivariate normal random variables x and y with means μ_x and μ_y , respectively, variances σ_x and σ_y , respectively, and correlation coefficient ρ :

$$(2) \quad x = \mu_x + \sigma_x u$$

$$(3) \quad \begin{aligned} y &= \mu_y + \rho \sigma_y u + \sqrt{1 - \rho^2} \sigma_y v \\ &= \mu_y + (\rho \sigma_y / \sigma_x)(x - \mu_x) + \sqrt{1 - \rho^2} \sigma_y v \end{aligned}$$

⁶ A more detailed version of the proof can be found in Boyer, Gibson and Loretan (1999). Their proof depends on the fact that bivariate normal random variables can each be written as a weighted average of the other and an independent component that is also normally distributed. See Goldberger (1991), p. 75.

Consider any event $x \in A$ such that $0 < \text{Prob}(x \in A) < 1$. By definition, the conditional correlation coefficient between x and y , ρ_A , is given by:

$$(4) \quad \rho_A = \frac{\text{Cov}(x, y/x \in A)}{\sqrt{\text{Var}(x/x \in A)} \sqrt{\text{Var}(y/x \in A)}}$$

By substituting for u in (3) using equation (2), then substituting the resulting expression for y into (4), and using the fact that x and v are independent by construction, one can rewrite this as:

$$(5) \quad \rho_A = \frac{(\rho \sigma_y / \sigma_x) \text{Var}(x/x \in A)}{\sqrt{\text{Var}(x/x \in A)} \sqrt{(\rho^2 \sigma_y^2 / \sigma_x^2) \text{Var}(x/x \in A) + (1 - \rho^2) \sigma_y^2}}$$

which can, in turn, be simplified to yield the expression in (1).

Thus, the conditional correlation between x and y is larger (smaller) than ρ in absolute value if the conditional variance of x given $x \in A$ is larger (smaller) than the unconditional variance of x .⁷

2.2 Some generalizations

This theorem is based on several assumptions that are unrealistic in empirical practice, such as that the data are i.i.d. and are drawn from a bivariate normal distribution. In the present case, bivariate normality of x and y is used – see equation (3) in the proof above – only to re-express the variable y as an affine function of x , $\mu_y + (\rho \sigma_y / \sigma_x)(x - \mu_x)$, plus a component that is independent of x , $\sqrt{1 - \rho^2} \sigma_y v$. Therefore, the main result of the theorem is not limited to cases where the data are bivariate normal, but holds in any situation where y can be stated as a linear (or, more generally, an affine) function of x plus an independent component (the “error term”), a framework which encompasses the familiar bivariate linear regression model with independent (but not necessarily normally distributed) errors.

Economic data are often observed as time series, and it is of interest to understand how measured correlations are affected by sampling variability in the data. Time series also pose the question of how serial dependence in the data affects the correlations. Assuming the general linear regression model, $y = \beta x_t + u_t$, with $\text{Cov}(x_t, u_t) = 0$ and $t = 1, 2, \dots$, one may write the sampling correlation coefficient between x and y , for a sample of size n , as

$$(6) \quad \text{Corr}_n(x, y) = \frac{\text{Cov}_n(x, y)}{\sqrt{\text{Var}_n(x)} \sqrt{\text{Var}_n(y)}}$$

which can be rewritten as

$$(7) \quad \text{Corr}_n(x, y) = \frac{\beta \text{Var}_n(x) + \text{Cov}_n(x, u)}{\sqrt{\text{Var}_n(x)} \sqrt{\beta^2 \text{Var}_n(x) + \text{Var}_n(u) + \beta \text{Cov}_n(x, u)}}$$

where the subscript n denotes sampling, as opposed to population, moments.⁸

Of primary interest to our paper is how this correlation coefficient will vary across subsamples of time. It will differ from the population moment for two reasons: the sampling covariance between x ,

⁷ Henri Pagès has provided the following heuristic interpretation of the result given in (1): if conditioning on an event $x \in A$ brings about an increase in the conditional variance of x , $\text{Var}(x/x \in A)$, it also raises the population conditional correlation between x and y because of a concomitant increase in the ratio $\text{Var}(x/x \in A)/\text{Var}(x)$, which can be rewritten as $(\text{Var}(v)/\text{Var}(x))(\text{Var}(x/x \in A)/\text{Var}(v))$, where the first term is a constant and the second term is a signal to noise ratio; it is this increase in the signal to noise ratio that raises the conditional correlation.

⁸ The connection from equation (6) to (1) is made by setting $\beta = \rho \sigma_y / \sigma_x$ and $u_t = \sqrt{1 - \rho^2} \sigma_y v_t$, and by dividing the numerator and denominator on the right-hand side of (6) by $\text{Var}(x) = \sigma_x^2$, the population unconditional variance of x .

and the “error” term u_t may be non-zero, or the sampling variance of u_t may be time-varying (and correlated with $Var_n(x)$).

In financial time series, we often find that the mean of y_t is close to a linear function of x_t , and that $Cov_n(x_t, u_t) \approx 0$ is a reasonable assumption to make. However, even if the levels of x and u are approximately uncorrelated, their variances may well be serially, as well as contemporaneously, correlated. Hence, the term $Var_n(u)$ may move systematically across subsamples with the sampling variance of x . Suppose, for example, that $\beta > 0$ and that the contemporaneous volatilities of x and u (and hence the contemporaneous volatilities of x and y) are positively correlated. Then, time intervals exhibiting a high sampling variability of x will also tend to have larger than average values for $Var_n(u)$. As a result, the *sampling* correlation between x and y will tend to deviate *less* from its population value, on average, than would be suggested by equation (1). (The possible practical importance of contemporaneous dependence in volatilities is discussed below; see footnote 17.)

While the preceding discussion establishes that the simple expression stated in equation (1) for the relationship between conditional and unconditional correlations may have to be modified suitably when some of the maintained assumptions of the theorem are not met by the data, it is clear that conditioning on volatility will, in general, have strong systematic effects on the correlation between x and y .

2.3 A simple example

As an example of a time series application of that theorem, consider subdividing a bivariate time series (x_t, y_t) that is observed daily into equally sized subsamples (“quarters”) of 60 daily observations each, and then ordering the quarters by the level of the within-quarter variance of x . For each subsample, we may also calculate the correlation between x and y . Table 1 shows the results of such an exercise under the assumption that x and y are i.i.d. bivariate normal with unit variances and a population correlation coefficient equal to 0.5 (as in Figure 1). The rows of the table show ranges for the ratio of the quarterly variance in x to its population value, while the columns show the distribution of the values of the quarterly correlation given the ranges. For values of the within-quarter variance of x close to its population value (0.9 to 1.1), the median value of the correlation is 0.50, although the range of values is fairly wide, with a 90% confidence interval running from 0.34 to 0.64. However, for quarters with in-sample variance of x between 1.7 and 1.9 times the population value, the median correlation is 0.61, with the 90% confidence interval running from 0.48 to 0.72.

Table 1
In-sample correlations when conditioning on volatility

Range of variances of x relative to its population value	Conditional correlation of x and y		
	Bottom 5%	Median	Top 5%
0.3–0.5	0.17	0.36	0.53
0.5–0.7	0.24	0.43	0.58
0.7–0.9	0.29	0.47	0.61
0.9–1.1	0.34	0.50	0.64
1.1–1.3	0.38	0.54	0.67
1.3–1.5	0.41	0.57	0.69
1.5–1.7	0.45	0.59	0.71
1.7–1.9	0.48	0.61	0.72

Note: The values of x and y are i.i.d. bivariate normal with a population correlation of 0.5. Reported values for the $Corr(x,y)$ are based on 2.5 million random draws of 60 observations (“quarters”). There were too few observations with the variance of x less than 0.3 or greater than 1.9 times its population value for values of $Corr(x,y)$ to be reported with confidence.

3. An empirical application

3.1 The data

In order to assess the empirical importance of the relationship between volatility and correlation over time, we need to study specific pairs of asset prices. We consider a relatively broad set of assets, rather than focus on a single type of asset (e.g. equity prices, as in Forbes and Rigobon (1999)). As shown in Table 2, we settled on equities, bonds and foreign exchange. The Financial Times and DAX stock price indexes are as of the close of the stock markets in London and Frankfurt, respectively, and are taken from Bloomberg. The German and British government bond yields are those on 10-year instruments. They are also as of market close and are from Bloomberg. The dollar/mark and dollar/yen exchange rates are those prevailing around noon in the New York market and are collected by Federal Reserve staff. We selected these particular series because they represent large and liquid markets and the data reflect market conditions at roughly the same time, and so we do not have to be concerned about the implications of non-synchronous data collection.⁹

Table 2
Data series used

Asset prices	Quote time	Quote location	Source
<i>Equity prices</i>			
Financial Times 100	Market close	London	Bloomberg
DAX	Market close	Frankfurt	Bloomberg
<i>Government bond yields</i>			
UK 10-year note	Market close	London	Bloomberg
German 10-year note	Market close	Frankfurt	Bloomberg
<i>Exchange rates</i>			
US dollar / Japanese yen	About noon New York time	New York	Federal Reserve
US dollar / Deutsche mark	About noon New York time	New York	Federal Reserve

Note: The London and Frankfurt markets now close at the same time, although in the past their closing times have differed by one hour.

The returns on foreign exchange and equity holdings are simply the daily percentage changes in their respective prices. Calculating the returns on government bonds is more complicated, since the available data relate to bond yields rather than bond prices. The return we use is an estimate of the percentage daily gain or loss from holding the 10-year bellwether government bonds, based on the reported yields and under the assumption that the bonds have maturities of exactly 10 years, have coupon payments twice a year, and are trading at par.¹⁰

Figures 2–4 show time series plots of the within-quarter variances (panels a and b) and correlations (panel c) of daily returns for the three pairs of assets. In the case of stock prices (Figure 2), it is clear that last fall was a period of high volatility and, just as the theorem presented earlier would suggest, there was a high correlation between the two returns. The returns on bond investments were also somewhat volatile last fall (Figure 3), and the correlation between the two returns in 1998 Q3 and 1998 Q4 was at the high end of its range. By contrast, movements in the dollar/yen and dollar/mark

⁹ For a discussion of the problems associated with non-synchronous data collection, see RiskMetrics (1996), pp. 184–96.

¹⁰ The approximation is excellent for small changes in yields on bonds trading near par. Even for large movements in yields on bonds priced far from par, the approximation is fairly good. For example, on October 9, 1998 (the day of the largest one-day loss on UK bonds in the sample), the UK 10-year bond maturing in October 2008 was selling near 130. Based on bond price and yield data from the *Financial Times*, the actual loss on the bond for that day was 2.33%, while the approximation yields a loss of 2.55%.

exchange rates were not particularly highly correlated last fall (Figure 4), despite the extraordinarily high volatility in the dollar/yen rate in 1998 Q4.

Figure 2
Within-quarter variances and correlations, stock market indexes

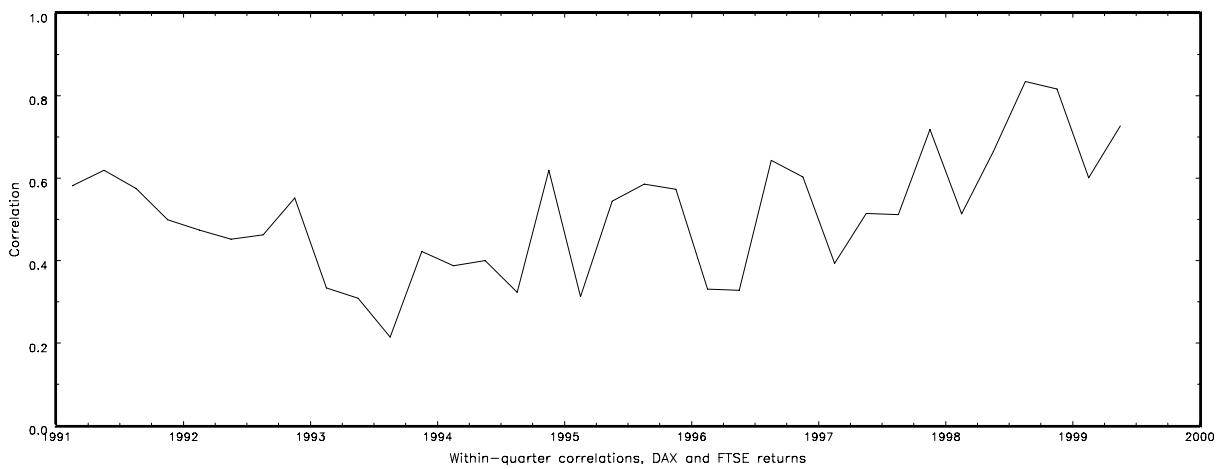
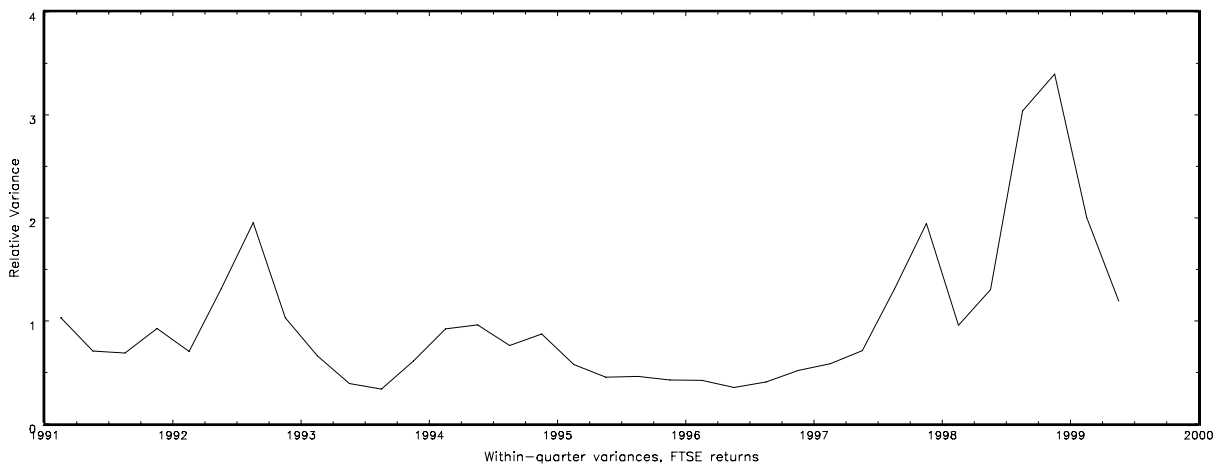
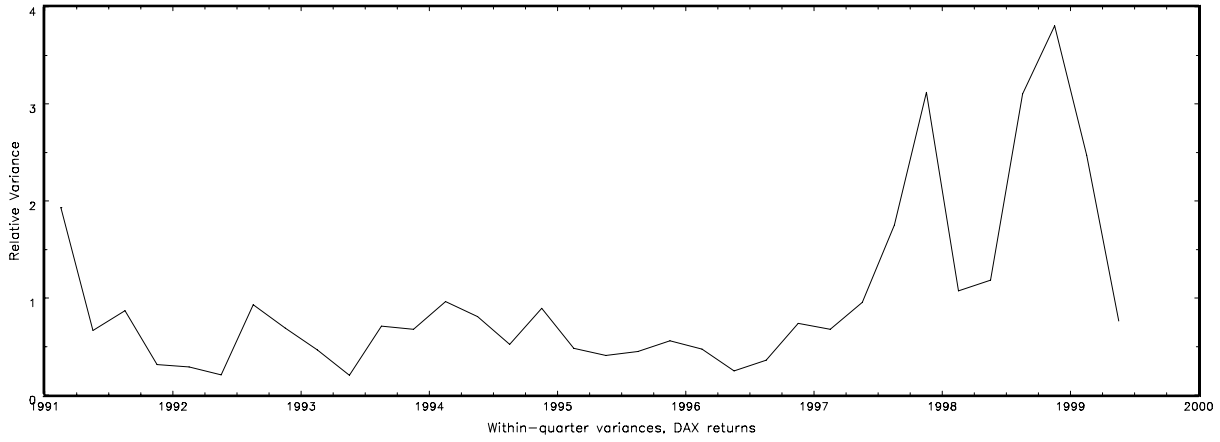


Figure 3
 Within-quarter variances and correlations, bond returns

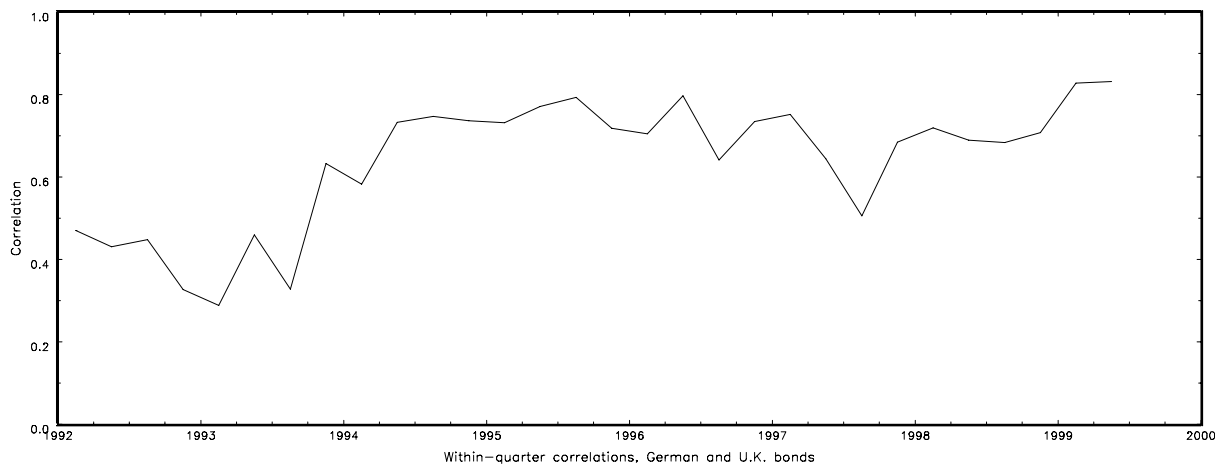
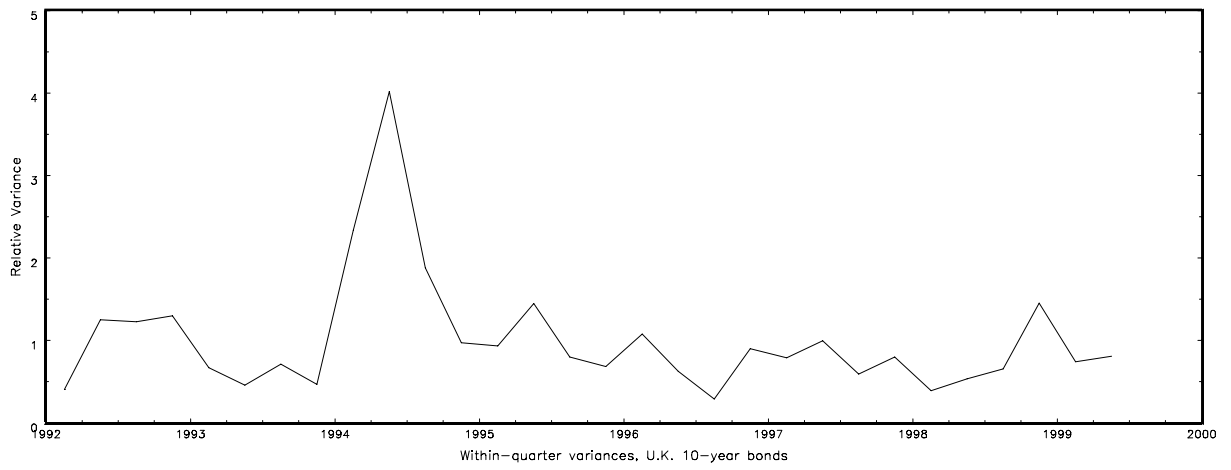
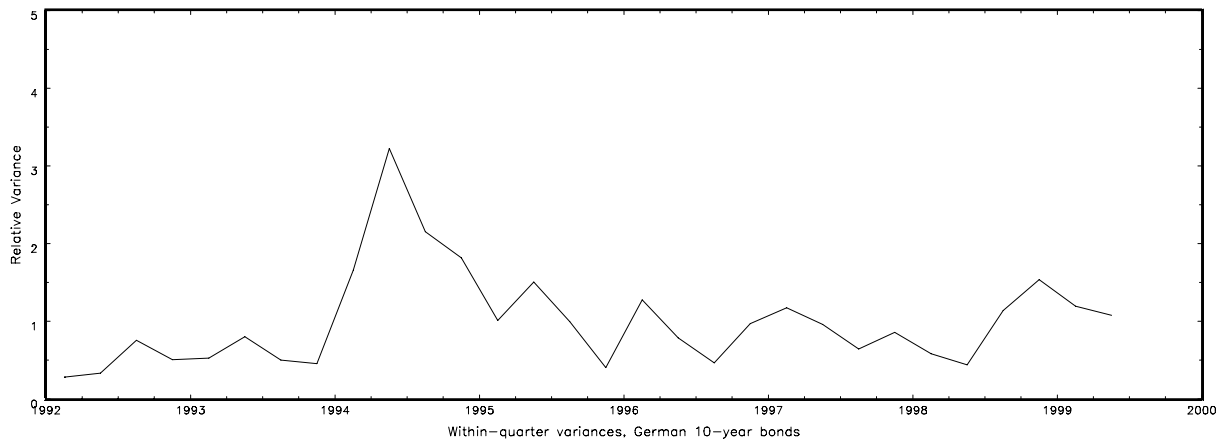
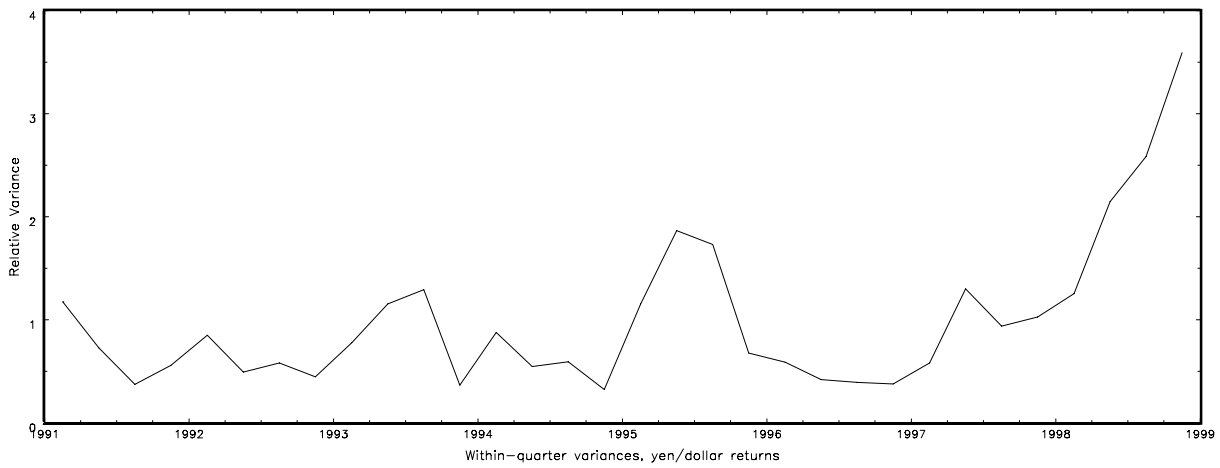
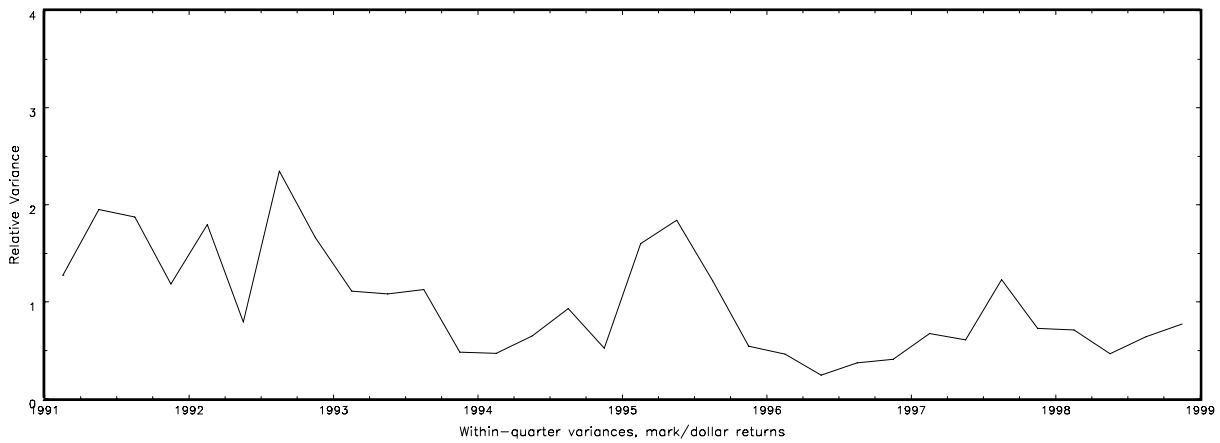


Figure 4
 Within-quarter variances and correlations, exchange rate returns



3.2 Using the theorem

Before using the theorem to try to explain the variation in quarterly correlations between returns, we need to examine whether the data satisfy its assumptions to ensure that its use is appropriate. Then, for each pair of returns, we need to determine the anticipated relationship between volatility and

correlation, based on the actual variances and correlations of the series, and compare the actual data to that anticipated relationship.

Figure 5
Nonparametric estimates of joint distribution of returns

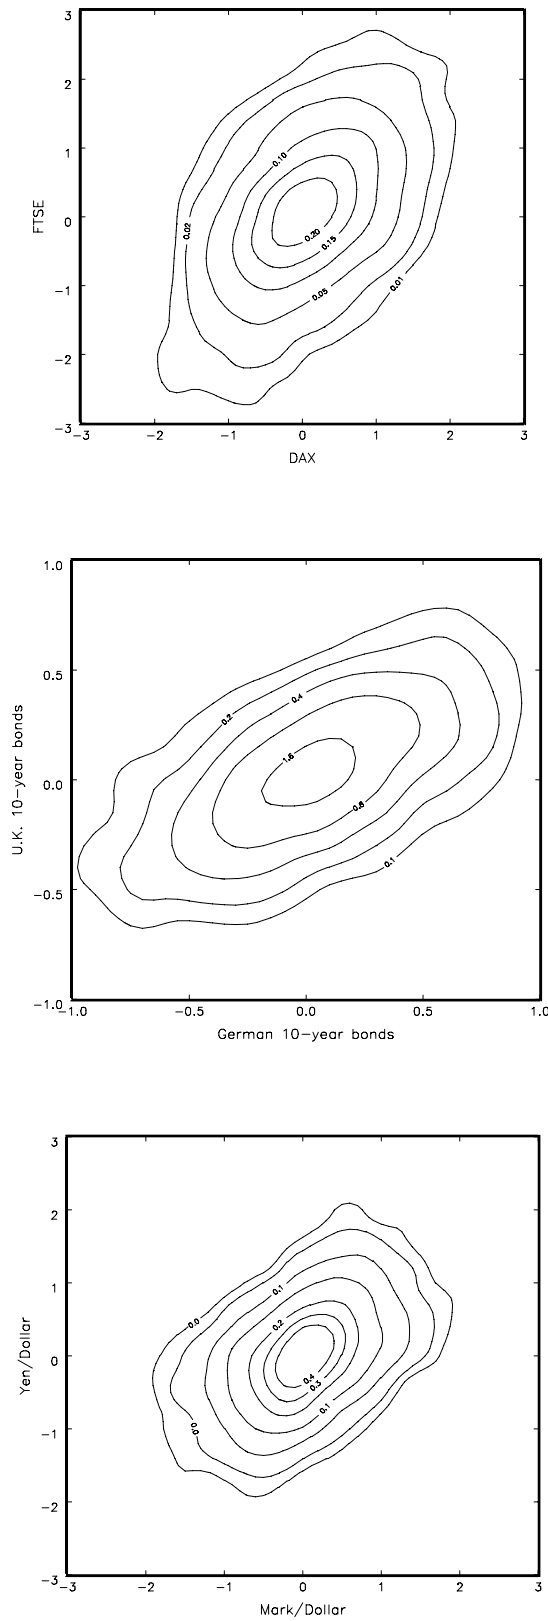


Figure 5 shows non-parametrically estimated level curves of the joint distributions of the pairs of asset returns.¹¹ These concentric curves are loci of constant height of the joint densities, and so they show the shape of the distribution of the observations. Recall that the theorem's main requirements are those of a linear regression model, which are satisfied if the joint distribution of the data is elliptic, that is, if the level curves of the joint density are ellipses (see Spanos (1986), pp. 457–8). The estimated level curves appear to be reasonably close to elliptic, although those for the exchange rates are somewhat more “rectangular” than the other two sets. We view these level curves as suggesting that the assumption of an elliptical distribution is reasonable for our empirical work, and hence that it is worthwhile to proceed comparing the empirical variance/correlation pairs with values predicted by that theorem.

To evaluate the importance of the theoretical link between volatility and correlation, we plot in Figure 6 the quarterly in-sample correlations against the in-sample volatility of one of the two asset returns.¹² In all three cases, a generally increasing relationship between conditional variances and conditional correlations is observable. However, the data also show a considerable dispersion in the in-sample correlation for a given level of volatility. In order to assess how close the points are to the locations indicated by the theorem, we also show in Figure 6 the theoretical relationship between asset return volatility and the correlation between asset returns, the dashed line (derived from equation (1)), and a 90% confidence contour around this theoretical locus, derived under the maintained assumptions that the data are i.i.d. bivariate normal, and that the population correlation is equal to the full-sample empirical correlation.¹³

The theoretical relationships appear to fit the data fairly well. In the case of equity prices, the increase in correlation late last year was very large and, as shown by the two points to the top right in the chart, these relatively high correlations were roughly consistent with what the theorem would lead one to expect given the increase in volatility at that time. The events of last fall left a much smaller imprint on the correlations between returns on government bonds and foreign exchange. In the case of foreign exchange, this may not be surprising given that the volatility of the dollar/mark exchange rate was not elevated. Over the entire period plotted, however, the empirical relationship between volatility and correlation seems to fit the theorem fairly well, although, in the case of the bond and foreign exchange returns, the fractions of observations falling outside the 90% confidence contours considerably exceed 10%.¹⁴

However, the stated 90% confidence intervals in Figure 6 will not be correct if the data are not well approximated by a bivariate normal distribution. For example, it is well known that the (unconditional) distribution of asset returns is strongly leptokurtic, resulting in far more “outlier” observations than one would expect under normality.¹⁵ One way to obtain improved confidence intervals is to use a bootstrap. We select random samples of a quarter's worth of observations (60 pairs of returns) from

¹¹ The non-parametric estimates of the joint densities are based on a bivariate normal kernel and a standard window width (Silverman (1986), pp. 86–7).

¹² The correlations could be plotted relative to the volatility of either return, but we have chosen the volatility of the German asset as the variable for the x -axis in all three cases. While this choice is immaterial theoretically, in the cases of the bond and foreign exchange returns, the results are somewhat better when the German asset volatility is selected. Note that the variances of the German asset returns have been expressed relative to their full-sample values.

¹³ Ronn (1995) plots average measures of comovement stratified by volatility for a number of interest rate measures and (separately) for a large number of US equity prices. He finds that, consistent with the theorem, large moves in asset prices are associated with a higher average degree of comovement. However, he does not try to match the empirical relationships with those implied by the theorem.

¹⁴ The confidence contours are shown over a relatively narrow range of quarterly variances because, under the assumption of an i.i.d. bivariate normal distribution of returns, there are too few observations outside this range (given our choice of 2.5 million repetitions) to allow the calculation of confidence intervals. The number of points on the left and right of the confidence contours shown reflects the fact that the actual distribution of returns has fatter tails than does the normal, resulting in greater dispersion of sampling variances.

¹⁵ See, for example, Mandelbrot (1963). For more recent discussions, see Jansen and de Vries (1991) and Loretan and Phillips (1994).

the actual data series (with replacement) and calculate the correlation and variance for each sample. By repeatedly sampling the data, we can trace out the median value of the correlation as well as 90% confidence bands

Figure 6
Scatterplots of quarterly variances vs. quarterly correlations,
normality-based confidence bands

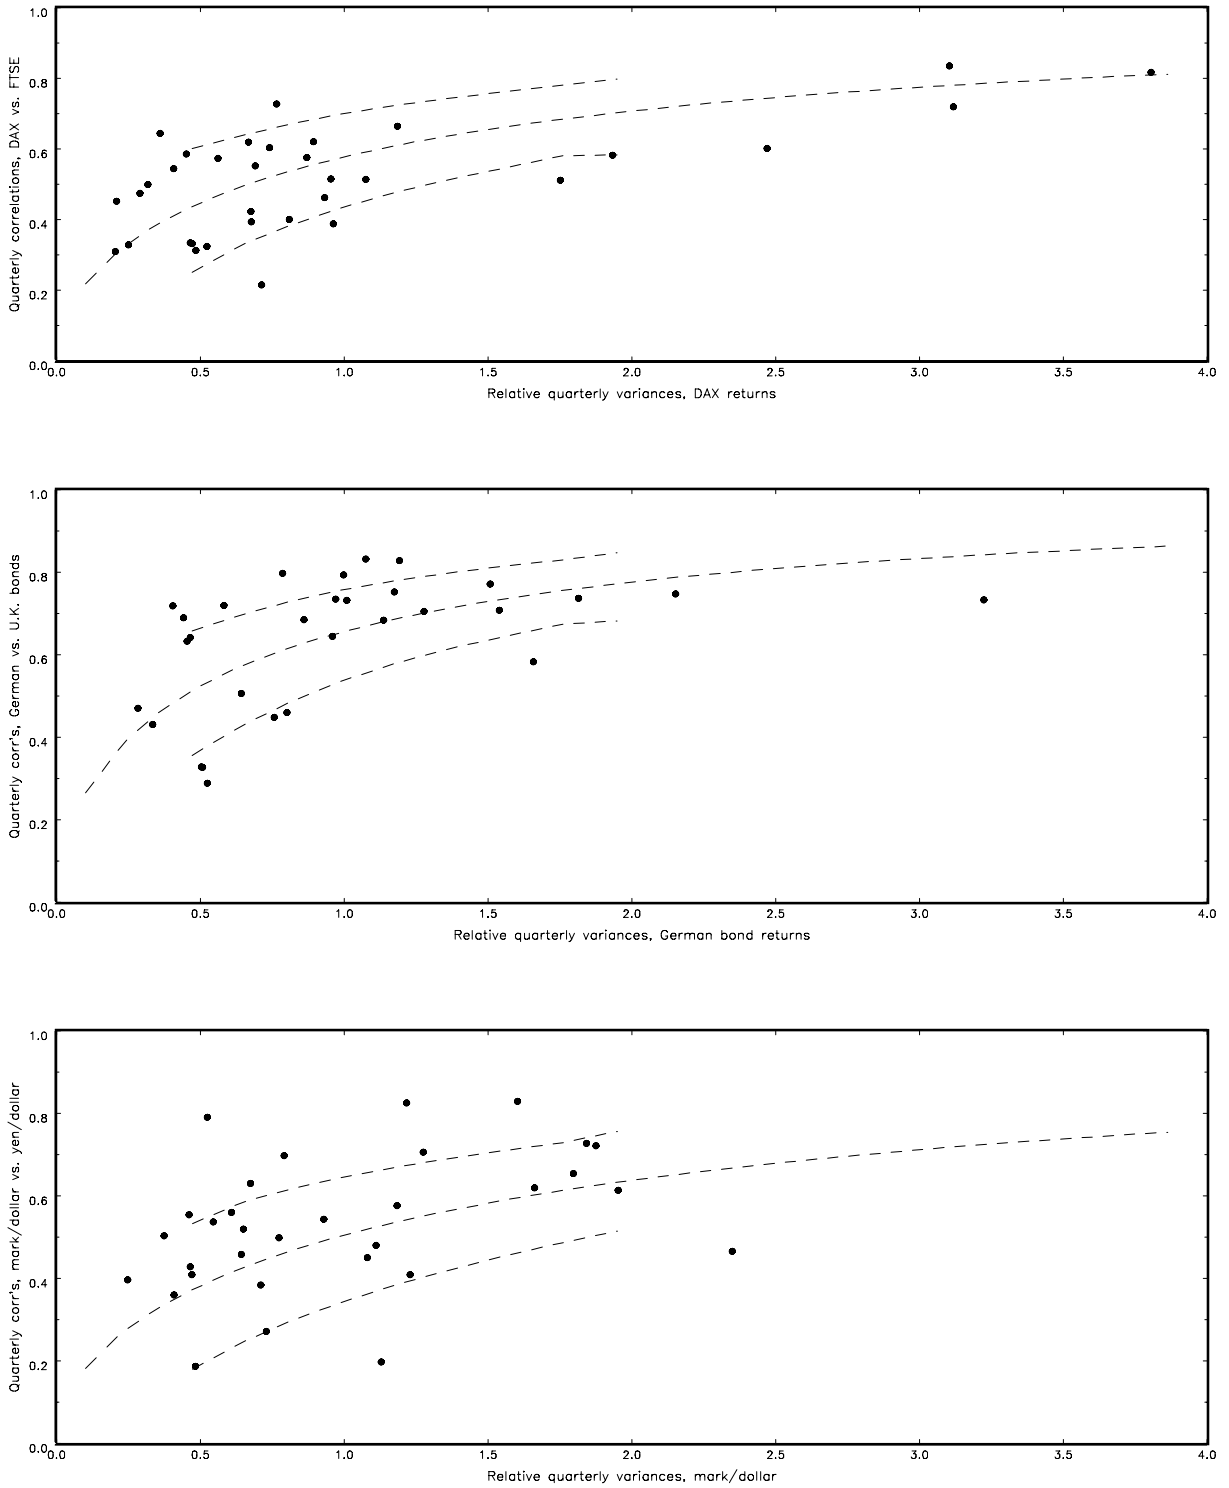
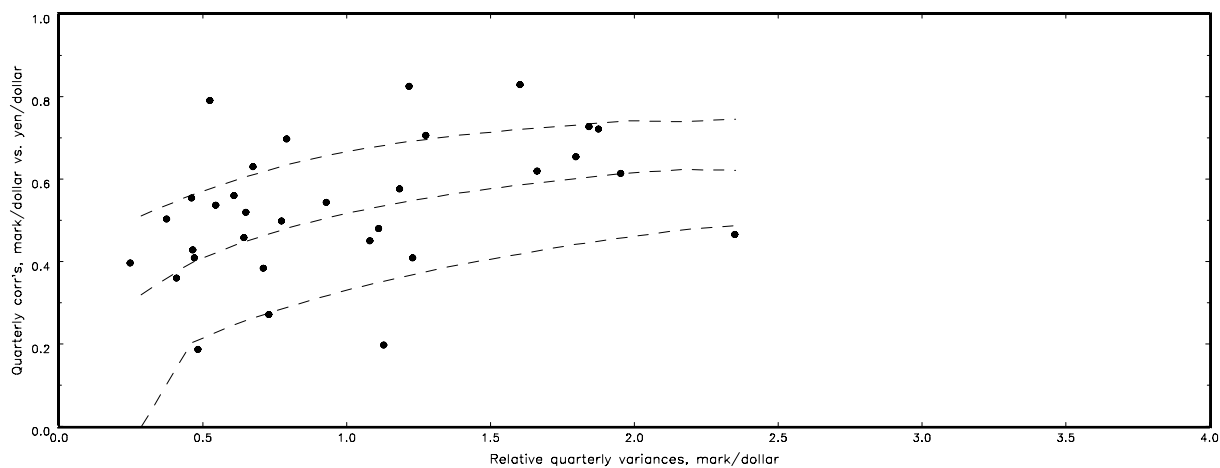
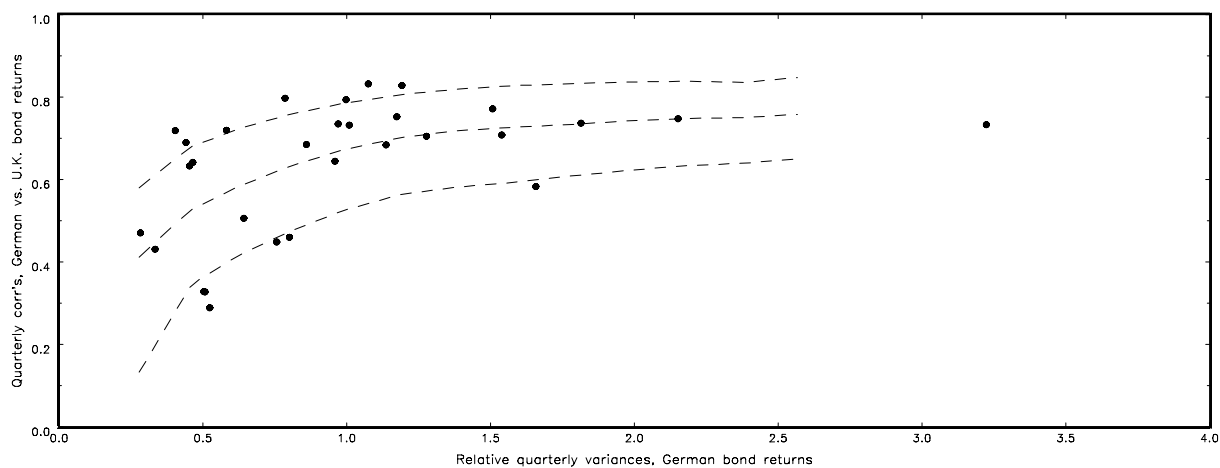
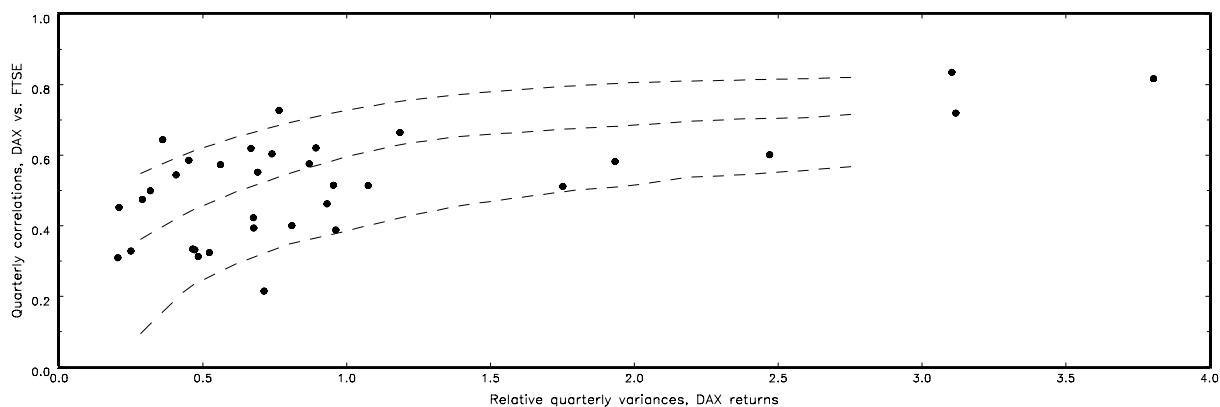


Figure 7
 Quarterly variances vs. quarterly correlations,
 bootstrapped confidence bands



confidence intervals based on the actual distributions of the data series.¹⁶ These Monte Carlo results are shown in Figure 7.

¹⁶ This bootstrap procedure preserves the contemporaneous correlation structure of the data, as well as the unconditionally heavy-tailed nature of the distributions, but it does not take account of serial dependence features such as GARCH which, as discussed in footnote 14, appear to be present in the data.

The improved methodology yields confidence contours that encompass a larger fraction of the quarterly data points.¹⁷ It appears that the equity data fit the theory relatively well, with only three of the 34 observations outside the 90% confidence contours. The bond return volatilities and correlations also appear to be roughly what one would expect given the theorem, although with some evidence of a larger than expected clustering of observations to the upper left of the chart. The foreign exchange returns generally match the upward slope predicted by the theorem, but the range of correlations still appears to be considerably wider than would be expected. The three outlier observations at the top of this panel are from 1994 Q4 to 1995 Q3, that is, the period during and after the Mexican crisis of December 1994. The location of these points may suggest that the sharp fluctuations of the dollar against the mark and yen that occurred following the Mexican crisis, and the associated increase in the correlation of these asset returns relative to long-term norms, may have been caused by a genuine, though temporary, change in the data generating process.¹⁸ For example, there were several episodes of concerted central bank intervention during this period, which would tend to boost daily correlations irrespective of changes in within-quarter volatility.

3.3 Summary

While there is some evidence in the case of foreign exchange suggesting the possibility of a temporary change in correlations following the Mexican crisis in late 1994, in general the link between high market volatility and high correlations between asset returns is remarkably close to what the theorem would suggest.¹⁹ While a more comprehensive test of the proposition of stationarity, including a broader array of asset prices, is beyond the scope of this paper, our results suggest that one should not be too quick to conclude that correlation breakdowns, including those observed last year, represent true changes in the distribution of asset returns. Rather, they may be nothing more than the predictable consequences of observing certain (low probability) draws from an unchanged distribution. These results need not imply that “contagion” does not occur; rather, they suggest that if one defines contagion to mean elevated correlations between asset returns, then contagion is a natural by-product of temporal variation in volatilities.

3.4 Implications

The statistical link between the volatility and correlation of asset returns discussed here has important implications for the evaluation of portfolio risk by market participants and investors, for the supervision of financial firms’ risk management practices, and for monetary policy.

In empirical practice, risk managers sometimes use data from a relatively short interval when calculating correlations and volatilities for use in risk management models. For example, one major banking company reports that they use the most recent 264 trading days’ changes in market prices in their calculations of value at risk, or VaR (Chase Manhattan (1999), p. 36). RiskMetrics uses 550 daily

¹⁷ The bootstrap also reveals that the location of the median correlation coefficient, for a given value of $Var_n(x)/Var(x)$, is (slightly) below the value predicted by equation 1 when $Var_n(x)/Var(x)$ is large, say greater than 2.5. This finding is consistent with our conjecture, advanced in Section 2 above, that a positive contemporaneous correlation between the volatilities of x and u could bring about lower bivariate correlations between the returns, on average, than the equation would suggest whenever $Var_n(x)$ – and hence $Var_n(u)$ – is large. We caution that the discrepancy we detect between the bootstrapped median value and the value predicted by equation (1) for the (subsample) correlations is quite small in all cases, and does not exceed 0.05. Hence, while the dependence between volatilities may bring about a slight deviation from the relationship stated in equation (1), the equation is still very useful for predicting the average location of the subsample correlations.

¹⁸ Recall also that the estimated joint distribution of returns on foreign exchange shown in Figure 5 appeared to be less elliptical than those of the equity and bond returns.

¹⁹ Our results are based on the volatility of asset returns with no distinction made between increases and decreases. In a related study, Longin and Solnik (1998) find that measured correlations between equity returns in different countries behave as the theorem would suggest when there are large positive stock market returns but are higher than the theorem would suggest when there are large negative returns. We leave an examination of this issue for future research.

return observations in the calculation of variances and correlations of returns, but because exponentially declining weights are applied, the effective number of daily returns employed is just 75 for estimates of daily correlations and volatilities and 150 for monthly values (RiskMetrics (1996), Table 5.9, p. 100). The use of a relatively short period for these calculations has some desirable features. Since financial markets can change over time, one may not want to depend on data from the distant past.²⁰ Moreover, the emphasis RiskMetrics puts on recent data allows it to take account of time-varying volatility, which appears to be a feature of actual returns (RiskMetrics (1996), pp. 55 and 79–80).

Nonetheless, the theoretical and empirical results presented here suggest that the use of a relatively short interval of data for estimating correlations and volatilities may be dangerous. If the interval is typically stable, then not only may the estimated volatilities be too low, depending on whether the assumed level of exponential decay captures the autocorrelation in volatility correctly, but, perhaps more importantly, the estimated correlations between returns will be lower than average. As a result, assessments of market risk may overstate the amount of diversification in a portfolio, leading the investing firm to take on excessive risk. Conversely, if the interval of data employed is a relatively volatile one, then the resulting estimates of correlations will be atypically high, and could lead the firm to take positions that are excessively risk-averse.

Another way in which the link between in-sample volatility and correlation could cause problems for risk managers is in the calculation of worst case scenarios and in stress testing. Put simply, risk managers should not consider the possible effects of a period of high volatility without also taking into account the likely effect that elevated volatility would have on the correlations between asset returns (see Ronn (1995), for a related discussion). Instead, risk managers may need to include information from historical periods of high volatility to form estimates of correlations conditional on the heightened volatility.²¹ These conditional correlations could then be used to evaluate the likely distribution of returns under a high volatility scenario. Put differently, the method used for stress testing a model must not (inadvertently) exclude the feature that periods of high volatility will also be periods of elevated correlations.

Supervisors of financial institutions also need to be aware of the link between volatilities and correlations when assessing firms' risk management. For example, under Federal Reserve regulations, an institution applying the market risk capital rules must hold capital based on its internal model's estimates of VaR. The internal model must be based on a minimum observation period of one year, and must be subjected to stress tests appropriate to the institution's particular situation. The VaR calculation should be based on a 99% (one-tailed) confidence level of estimated maximum loss (Federal Reserve (1999), pp. 11–2). In evaluating such internal models, supervisors need to keep in mind the difficulties noted earlier with relying on a relatively short interval of data for information on correlations and the need to form appropriate conditional correlations for stress tests.

Finally, monetary policymakers may also need to be aware of the link between volatility and correlations. Most obviously, correlation breakdowns during periods of financial market turbulence could lead policymakers to reassess the stance of policy in light of the apparent shift in correlations. So long as in-sample changes in correlations only reflect the movements in volatility, however, such reassessment would not be warranted unless private agents' adjustments to the volatility made a change in the stance of monetary policy desirable.

Perhaps more importantly, changes in the stance of monetary policy, particularly if they follow an extended period during which policy has remained unchanged, have the potential to cause volatility in financial markets to increase. Since such an increase in volatility is likely to be associated with

²⁰ Similarly, if the assets under consideration are firm-specific (rather than indexes), the behavior of firms can change over time as managers or business strategies are changed, making older information less useful.

²¹ Alternatively, firms might want to use actual data from earlier periods of high volatility to stress test their portfolios. Chase Manhattan indicates that they use asset price movements during the bond market sell-off in 1994, the 1994 Mexican peso crisis and the 1997 Asian markets crisis, as well as internally developed scenarios, when assessing the risk of their portfolio (Chase Manhattan (1999), p. 37).

increases in correlations between asset returns, a change in policy could impose substantial unanticipated risks on market participants if – as discussed earlier – their risk management were based on movements in asset prices during the previous period of relative stability. The resulting adjustments of portfolios by the affected firms could generate outside effects of policy on asset prices.²² These relatively large effects may lead policymakers to move more gradually than they otherwise would in such situations, providing a possible explanation for the interest rate smoothing behavior noted by Sack and Wieland (forthcoming). Moreover, uncertainties about the size and timing of such portfolio adjustments would make the likely effects of policy more difficult than usual to assess. This uncertainty would also suggest that a gradual approach would be desirable (Brainard (1967)).²³

Indeed, the Federal Reserve found that its tightening of monetary policy in 1994, which followed a substantial period of unchanged policy, led market participants to trim their risk profiles. It noted that uncertainty about the policy outlook, as well as “the capital losses following the tightenings, encouraged investors to shorten the maturity of their investments and reduce their degree of leverage” (Federal Reserve (1995), pp. 21–2). While this reaction may only have reflected the revision of overly optimistic beliefs about monetary policy and market volatility, it is also what one would expect to happen if the increase in volatility by itself caused correlations between asset returns to increase. There is some evidence that this was the case; for example, Bankers Trust indicated in its 1994 annual report that movements in interest rates in early 1994 (and also at the end of 1994, when the Mexican peso devaluation occurred) were “unusual in the degree to which interest rates across international markets moved together” (Bankers Trust (1995), p. 23).²⁴ It went on to note that “this phenomenon of increased correlation among interest rates reduced the risk management benefits derived from diversification across interest-sensitive instruments” (Bankers Trust (1995), p. 23). The bank responded to this unexpected situation by withdrawing from substantial market positions (Bankers Trust (1995), p. 24).

In fact, the possibility that a change in monetary policy could affect market volatility and correlations, and thereby influence the desired risk profiles of market participants, may have been taken into account by the Federal Reserve. In its discussion of the 1994 tightening, the February 1995 Humphrey-Hawkins report noted that:

“The FOMC, at its meeting in early February 1994, agreed that policy should be moved to a less stimulative stance. The pace at which the adjustment to policy should be made was less clear: A rapid shift in policy stance would minimise the risk of allowing inflation pressures to build, while a more gradual move would allow financial markets time to adjust to the changed environment. Although many market participants seemed to anticipate a firming move fairly soon ... some investors would undoubtedly reconsider their portfolio strategies.” (Federal Reserve (1995), p. 21).

Taking account of these factors, the FOMC agreed to move slowly at first. As a result, the first three policy moves were relatively small (25 basis points), and spread over three months. However, the Federal Reserve also reported that “once market participants seemed to have made substantial adjustments to the new direction of policy, a larger tightening move [of 50 basis points] was implemented in May” (Federal Reserve (1995), p. 21).

²² This mechanism could provide an alternative explanation for the larger effects on asset prices of monetary policy moves that represent a change in the direction of policy, noted by Goodfriend (1991).

²³ In addition, these effects could be reduced by increased transparency about the outlook for policy in advance of a policy move.

²⁴ The effect of increased volatility on conditional correlations is, of course, only one possible reason for the increase in correlations of asset returns at this time. For example, a structural model of asset returns might suggest that changes in monetary policy should generate relatively highly correlated movements in many asset prices.

Bibliography

- Bank for International Settlements, Committee on the Global Financial System (1999): *A Review of Financial Market Events in Autumn 1998*. Basel, October.
- Bankers Trust (1994): *Annual Report 1994*. New York.
- Bookstaber, Richard (1997): “Global Risk Management: Are We Missing the Point?” *Journal of Portfolio Management*, pp. 102–7.
- Boyer, Brian H, Michael S Gibson and Mico Loretan (1999): “Pitfalls in Tests for Changes in Correlations”. Federal Reserve Board, *IFS Discussion Paper*, No. 597R, March.
- Brainard, William (1967): “Uncertainty and the Effectiveness of Policy”. *American Economic Review, Papers and Proceedings*, Vol. 57, pp. 411–25.
- Chase Manhattan Corp. (1999): *Annual Report 1998*. New York.
- Drazen, Alan (1998): “Political Contagion in Currency Crises”. *Mimeo*, University of Maryland.
- Eichengreen, Barry, Andrew Rose and C Wyplosz (1996): “Contagious Currency Crises”. *NBER Working Paper*, No. 5681.
- Forbes, Kristin and Roberto Rigobon (1999): “No Contagion, Only Interdependence: Measuring stock market co-movements”. *NBER Working Paper*, No. 7267, July.
- Federal Reserve Board (1995): *Monetary Policy Report to the Congress Pursuant to the Full Employment and Balanced Growth Act of 1978*. February 21.
- Federal Reserve Board (1999): *Trading and Capital Markets Activities Manual: Capital Adequacy*. Washington, D.C.
- Gerlach, Stefan and Frank Smets (1995): “Contagious Speculative Attacks”. *European Journal of Political Economy*.
- Goldberger, Arthur S (1991): *A Course in Econometrics*. Cambridge, MA: Harvard University Press.
- Goodfriend, Marvin (1991): “Interest Rates and the Conduct of Monetary Policy”. *Carnegie-Rochester Conference Series on Public Policy*, Vol. 34, pp. 7–30.
- JP Morgan and Co. (1999): *Annual Report 1998*. New York.
- Jansen, D W and C G de Vries (1991): “On the Frequency of Large Stock Returns: Putting Booms and Busts into Perspective”. *Review of Economics and Statistics*, Vol. 73, pp. 18–24.
- Kodres, Laura E and Matthew Pritsker (1999): “A Rational Expectations Model of Financial Contagion”. *Mimeo*, Federal Reserve Board.
- Longin, François and Bruno Solnik (1998): “Correlation Structure of International Equity Markets During Extremely Volatile Periods”. *Mimeo*, Group HEC.
- Loretan, Mico and Peter C B Phillips (1994): “Testing the Covariance Stationarity of Heavy-Tailed Time Series”. *Journal of Empirical Finance*, Vol. 1, pp. 211–48.
- Mandelbrot, B B (1963): “The Variation of Certain Speculative Prices”. *Journal of Business*, Vol. 36, pp. 394–419.
- RiskMetrics (1996): *RiskMetrics–Technical Document*. 4th ed. New York: JP Morgan/Reuters.
- Ronn, Ehud I (1995): “The Impact of Large Changes in Asset Prices on Intra-Market Correlations in the Stock and Bond Markets”. *Mimeo*, Department of Finance, University of Texas at Austin.
- Sack, Brian and Volker Wieland: “Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence”. *Journal of Economics and Business* (forthcoming).

Silverman, B W (1986): *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall.

Spanos, Aris (1986): *Statistical Foundations of Econometric Modelling*. Cambridge: Cambridge University Press.