EUROPEAN CENTRAL BANK

WORKING PAPER SERIES

WORKING PAPER NO. 66

CAN SHORT-TERM FOREIGN EXCHANGE VOLATILITY BE PREDICTED BY THE GLOBAL HAZARD INDEX?

BY VINCENT BROUSSEAU
AND FABIO SCACCIAVILLANI

June 2001

WORKING PAPER SERIES

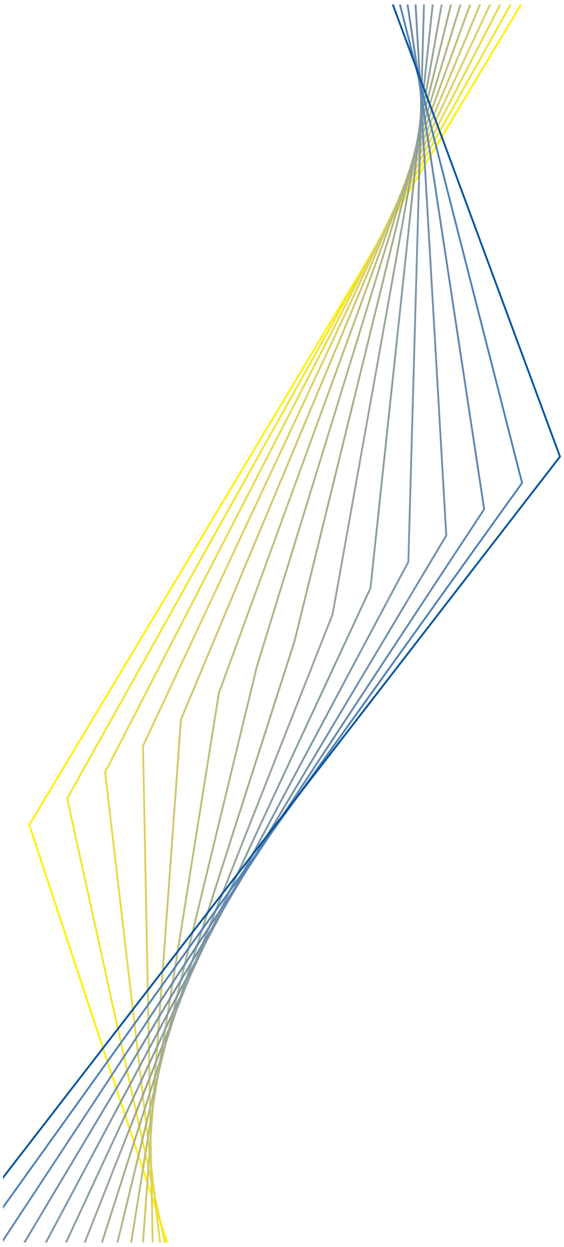

WORKING PAPER NO. 66

CAN SHORT-TERM FOREIGN EXCHANGE VOLATILITY BE PREDICTED BY THE GLOBAL HAZARD INDEX?

BY VINCENT BROUSSEAU AND FABIO SCACCIAVILLANI*

June 2001

* *Authors' e-mail address: vincent.brousseau@ecb.int; fabio.scacciavillani@gs.com*
We wish to thank Satu Jaatinen for providing some of the data used in this study. Furthermore we acknowledge the extremely useful comments made by Charles Goodhart, Carsten Detken, Philipp Hartmann, Benoît Mojon, an anonymous referee and the participants at the Conference on "The euro: what's in the future" in Venice on 8 to 9 June 2000. We accept sole responsibility for any remaining errors. The views expressed in this paper are not necessarily those of the European Central Bank.

Address	**Kaiserstrasse 29**
	D-60311 Frankfurt am Main
	Germany
Postal address	**Postfach 16 03 19**
	D-60066 Frankfurt am Main
	Germany
Telephone	**+49 69 1344 0**
Internet	**http://www.ecb.int**
Fax	**+49 69 1344 6000**
Telex	**411 144 ecb d**

ISSN 1561-0810

Contents

Abstract

This paper examines the predictive properties of risk indicators for the foreign exchange markets. In particular it considers the predictive properties of historical volatilities and implied volatilities for movements in various bilateral exchange rates and compares them with the analogous properties of a composite indicator of risk, the Global Hazard Index (GHI). The GHI is a function of the implied volatility derived from currency options on the three major exchange rates, i.e. the euro-US dollar, the US dollar-yen and the euro-yen. For the empirical analysis this paper employs the concept of kernel volatility, which, loosely speaking, expresses the volatility of one variable conditional on the level of another. Simple regressions show that the levels of all the indicators on a particular day have a strong link to the variance of the nominal bilateral exchange rate on the next day. A strong overall influence is displayed by the GHI, especially for the currencies of small open economies.

JEL classification codes: F01; F31
Keywords: exchange rates; currency options; market turbulence; implied volatility; risk; forecasting.

Executive summary

Brousseau and Scacciavillani (1999) devised the Global Hazard Index (GHI) as a forward-looking measure of global currency market risk. Indeed the GHI can be seen as an "overall multi-market index of anticipated exchange rates volatility", derived by extending the option currency formula of Garman and Kohlhagen (1983) to a multi-currency context. In essence, the GHI is a function of major exchange rate implied volatilities, expressing market expectations about the co-movement of different exchange rates.

This paper compares the predictive properties of the GHI as regards the short-term volatility of returns on foreign exchange holdings with those of the historical volatility and of the implied volatility of major exchange rates for a number of currency pairs.

Modern asset price models emphasise the positive link between risk and return in speculative markets, which generally holds true. Basically, the higher the expected return on an asset, the higher its risk. However, in foreign exchange markets an analogous line of argument is meaningless because the return on holding a particular currency is not linked solely to its risk. To calculate a return it is necessary to *choose or define a riskless currency*. But a riskless currency does not exist in reality, so that the relationship between a measure of risk in the global exchange rate market and the movements in a particular bilateral exchange rate is, from a theoretical standpoint, rather hazy. For this reason, this paper has a purely empirical focus aimed specifically at highlighting stylised facts about the relationship between indicators of risk and bilateral exchange rates. More precisely, it aims to examine whether and to what extent those indicators of risk contain information about the *subsequent* (some prefer to say *forthcoming*) realised volatility of a given currency pair.

The empirical analysis addresses two specific questions.

1) Can tomorrow's volatility of return in bilateral nominal exchange rates be predicted by today's level of hazard in the global foreign exchange market as represented by the GHI?

2) Is the predictive power of the GHI stronger than that of individual historical volatilities or individual implied volatilities?

In carrying out this analysis, we have developed the concept of "kernel correlation", which is essentially calculated by attaching a set of weightings to the observations so as to give greater weight to those correlating to a certain phenomenon. This "kernel volatility" or conditional volatility allows weight to be given to the correlation between an indicator and a particular exchange rate volatility, depending on the level of the indicator. We resort to the kernel volatility because standard regression analysis would not allow to find the particular pattern of the relationship between indicator and volatility. In contrast, the kernel volatility method allows you to visualise and possibly to recognise that pattern.

The main results are summarised in Tables 1 and 2 where the slopes of the weighted regressions and the conditional correlations are reported. The results we obtain indicate that the implied volatilities and the GHI tend to have a significant influence on the variance of major exchange rates next day returns, and on that of a number of smaller ones. The notable exceptions, however, are the Swiss franc and pound sterling. The overall findings indicate that, in general, the implied volatility of the own currency pair displays the strongest influence on the standard deviation of the next day return. Moreover, the GHI displays one of the strongest influences overall, and this influence is most significant in the sense that it is the least likely to stem from a spurious regression.

1 INTRODUCTION

Brousseau and Scacciavillani (1999) proposed the Global Hazard Index (GHI) as a forward-looking indicator of global currency market risk. The GHI can be seen as an "overall multi-market index of anticipated exchange rates volatility" (World Bank (1999)) which provides an implicit measure of global expected risk. It is derived in a multi-currency context under the hypothesis of no arbitrage and uses a combination of major exchange rate implied volatilities to provide an implicit measure of global expected risk, expressing market expectations about the shared portions of the movements of different currencies.

In essence, the GHI expresses a synthetic measure of risk in the world exchange rate market because the three major currencies account for a large proportion of world transactions and, furthermore, play a pivotal role in most cross exchange rates. In fact, since the markets between the currencies of small economies are illiquid, any amounts to be exchanged from, say, Mexican pesos into Polish zloty are first converted into a major currency. Therefore, higher levels of risk picked up by the GHI are likely to reverberate throughout the world currency markets, influencing the portfolio choices and asset allocations of international investors.

The GHI is defined as

$$\mathbf{GHI} = \frac{2\sigma_{\$M}\,\sigma_{\$¥}\,\sigma_{M¥}}{\sqrt{(\sigma_{\$M}+\sigma_{\$¥}+\sigma_{M¥})(-\sigma_{\$M}+\sigma_{\$¥}+\sigma_{M¥})(\sigma_{\$M}-\sigma_{\$¥}+\sigma_{M¥})(\sigma_{\$M}+\sigma_{\$¥}-\sigma_{M¥})}} \quad (1.1)$$

with σ representing the implied volatility taken from currency option data and the subscripts M, \$ and ¥ denoting the three major currencies, German mark (from January 1 1999, the euro), the US dollar and the Japanese yen, so that $\sigma_{M\$}$, for example, represents the implied volatility calculated from options on the dollar-mark exchange rate.[1] This construction ensures that the combination of the three implied volatilities[2] retains the same characteristics as any individual implied volatility, in particular the unit of measurement. The development of the GHI over the past few years is shown in Figure 1.[3]

1 The choice of the GHI as the diameter of the circle circumscribing the triangle formed by the three implied volatilities was motivated by a mathematical argument outside the scope of the present paper. The interested reader is referred to the original paper, Brousseau and Scacciavillani (1999).

2 The index is based on exchange rate volatilities and not on currency values, as is sometimes incorrectly stated.

3 In Brousseau and Scacciavillani (1999) we remarked that implied volatility could be seen as an indicator of what Mandelbrot (1997a) defines as wild risk, i.e. the risk deriving from erratic and possibly persistent shocks which eludes analysis based on the traditional Gaussian paradigm. In order to distinguish it from the traditional concept of risk, we resorted to the term hazard (hence the name GHI). In our view, this is why implied volatility calculated using the formula of Garman and Kohlhagen (1983) provides a unit of measurement for market turbulence. In a nutshell, the idea is similar to what physicists do when they introduce "effective" parameters. The real world is not Gaussian; but if it were, which value of the volatility parameter would drive the price of at-the-money options to be the same as the price observed in the markets? The answer to this question gives us the implied volatility. Precisely because the world is not Gaussian, the implied volatility will not be constant and will not provide information about market expectations.

Brousseau and Scacciavillani (1999) argued that the GHI could be useful as a diagnostic tool for market participants and policy-makers in comparing levels of risk across time and in monitoring the effects of policy actions on market behaviour and expectations.

In addition researchers at investment banks have been using the GHI

1) to evaluate the consequences of movements in a particular currency on the instability of the world foreign exchange market (Esposito and Laruccia, Banca Commerciale Italiana (1999));

2) to estimate the risk aversion degree of investors (Llewellyn, Lehman Brothers (2000), Redeker, BNP Paribas (2001));

3) to predict movements in risk-influenced currencies (Fornasari (2000), Bernhardsen and Røisland (2000), also see Lehman Brothers(1999), Norges Bank (1999).[4]) Those authors compare the development of their exchange rates with the GHI, as is shown in Figure 2; then derive the impact of the GHI on their (forthcoming) level.[5]

4) to forecast the dynamics of aggregate risk through a technical analysis of the trends in the GHI (Jaatinen, JP Morgan (2001)).

This paper extends these lines of research by comparing the predictive properties of the GHI for the next day volatility of returns on foreign exchange holdings with those of the historical volatility and of the implied volatility of major exchange rates for a number of currency pairs.

The question of the predictive power of implied volatilities has previously been addressed both by academics and by market analysts. Among the earliest was Jorion (1995), who performed linear regressions of realised volatility (over the remaining life of the currency contract) against the implied volatility. He found that the predictive power of implied volatilities outperforms statistical time-series models. Market analysts have approached the question in a similar way (see, for example, Lehman Brothers (2000)) and have found strong evidence of the significant predictive power of implied volatilities. However, the market analysts also looked at the relation between implied volatilities and the future realisation of historical volatility and found that the former consistently overestimate the latter. Other market economists have adopted more sophisticated models. Mezrich and Majewska (2000), for example, suggested using the whole term structure of implied volatilities as a predictor, rather than just one volatility.

[4] Economists at Norges Bank renamed the index "Global Risk Indicator" (GRI) in their paper of 1999.

[5] For example, a 1% increase of the GHI is said to mean a drop of the trade weighted NOK exchange rates in 2 months.

The purpose of this paper is to study and compare the predictive power of historical volatilities, implied volatilities and the GHI for the short-term volatility of selected bilateral exchange rates by means of the *smoothing kernel method*. As well as allowing the detection of relationships between predicting indicators and predicted variables, this method also provides visual information about the functional form of this relationship.

Indicators of risk can be considered as reflecting the information available to market participants and the inferences they make on the basis of that information, e.g. as the forecast from a complex (albeit unspecified) model. For extracting predictions from an indicator, several methodological alternatives may be adopted. These can be grouped into two rough categories: (i) an investigation of the relationships between the indicator and the variable in question over the entire period or (ii) a focus on particular episodes or periods characterised by a specific feature or pattern (case studies).

The smoothing kernel method, which is adopted here, can be seen as a synthesis of those two approaches. The level of the risk indicator (historical volatility, implied volatility, GHI) is used to determine whether the foreign exchange market is in a period of turbulence or stability. The relationship between the indicator and a specific exchange rate is then examined, "depending" (in a sense that will be clarified later) on whether the market is in a period of turbulence or stability.

The rest of the paper is divided into four sections: Section 2 provides an overview of some theoretical aspects of risk in the foreign exchange markets; Section 3 describes the methodology used in the empirical analysis; Section 4 discusses the main results; and Section 5 draws conclusions and proposes some avenues for further research.

2 A WINDOW ON THE INFLUENCE OF RISK

2.1 Risk and Return

Modern asset price models emphasise the positive link between risk and return in speculative markets, which generally holds true. Models such as the Capital Asset Pricing Model (CAPM) stress that returns on assets depend on their covariance with the return on the so-called market portfolio, i.e. a linear combination of all financial and non-financial assets in the economy. Basically, they give an empirical structure to the idea that the higher the expected return on an asset, the higher its risk.

However, in foreign exchange markets an analogous line of argument is meaningless because the return on holding a particular currency is not linked solely to this currency's risk. To calculate a return it is necessary to *choose or define a riskless currency*. But in reality a riskless currency does not exist; it is only the chosen (theoretical) currency of gains or losses.[6]

Furthermore, the return on holding foreign currency can hardly be seen as a function of a market portfolio, however defined. So the relationship between a measure of risk in the global exchange rate market and the movements in a particular bilateral exchange rate is, from a theoretical standpoint, rather hazy. Several strands of literature have focused on this risk relationship, emphasising elements such as contagion, flight to safety etc., but a fully satisfactory or widely accepted theoretical underpinning has not emerged yet.

For this reason, this paper has an eminently empirical focus, being specifically directed at highlighting stylised facts on the relationship between indicators of risk and bilateral exchange rates. Namely, we intend to examine whether and to what extent those indicators of risk contain information about the *subsequent* (some prefer to say *forthcoming*) realised volatility of a given currency pair.

2.2 Volatilities as indicator of risk?

The concept of volatility can be approached from two different standpoints, both relevant to the problem, but for different (and possibly opposite) reasons.

The historical volatility is the realised standard deviation of the (log) spot in a certain period of time. It is meaningful to use it as a predictor of currency movements, if and only if the spot process exhibits some regularity, which can be detected by statistical techniques, e.g. a positive auto-correlation of the squared returns.[7]

The implied volatility is derived from an option-pricing model (in the case of currency, the Garman and Kohlhagen formula) as the price of an at-the-money liquid plain-vanilla currency option. According to the conventional academic point of view, it is meaningful to use it as a predictor of future currency movements if and only if the participants in the options market act on information on the underlying exchange rates process over and beyond what can be extracted from past values.

---

6 Strictly speaking, the concept of holding returns applies to currency pairs rather than to currencies.

7 In recent years a host of empirical papers have been devoted to uncovering the properties of the exchange rate through a host of state of the art statistical techniques, such as ARCH, GARCH, EGARCH, neural network, non linear models, etc. Despite their intrinsic differences, all these methods provide evidence that the dynamics of exchange rates is far from being regular.

In other words if the set of information relevant for an accurate prediction of future developments is larger than that formed uniquely by the time series of exchange rates.

The two types of volatilities could reasonably be deemed informative. The reasons for which they might a priori have this predictive power differ, as we will examine more in detail in the next subsection.

2.2.1 Historical volatilities

The reason why historical volatilities might display predictive power has to do with a well known empirical regularity on any liquid exchange rate: the slowly decaying autocorrelation function of the squared returns. Figure 4 displays such autocorrelations for four exchange rates. At lag 0 the correlation is necessarily 1. For positive lags the correlation should not differ significantly from 0, otherwise some simple trading strategies would yield an unrealistically high Sharpe ratio. So, in essence, the auto-correlation should be represented by a Dirac function with mass at 0.

However the auto-correlation of the squared returns for positive lags ought not to be 0. In fact, from Figure 4, we observe that while a Dirac function, as expected, approximates the auto-correlation function of the returns, the auto-correlation function of the squared returns is not zero at positive lags, although it slowly decreases to 0 as the lags go to infinity. This second feature has already been documented in both academic and market research. It is referred as to the persistence of (historical) volatilities. "*Volatility also feeds on itself: small fluctuations can rapidly become larger as the result of "bandwagon" effects and the tendency of the foreign exchange markets to overshoot.*" (Patterson, Sienkiewicz and Avila (2000)). This persistence of volatility explains why the historical volatilities could be good predictors of the subsequent realised volatility.

2.2.2 Implied volatilities

It has been suggested (e.g. Jorion (1995)) that the possible predictive power of implied volatilities is linked to the fact that options market-makers are presumably better informed than the average market participant. However this explanation remains vague about the nature, and the content, of this supplementary information.

In our view, the predictive power of implied volatilities can be explained without any reference to any informational advantage by market makers. Instead, it would appear as the result of a rather mechanical phenomenon.

The foreign exchange option market comprises two types of players: market makers and their customers. Customers generally buy options from, but do not sell to, market makers. So the market makers are structurally short on volatility, while customers are structurally long. Furthermore, customers generally do not delta-hedge their option positions, while the market makers overwhelmingly do. Moreover, the market makers compute their delta-hedge through the Garman-Kohlhagen formula, regardless of its possible mis-specification. So the existence of open option positions results in a delta-hedging behaviour by market players who are short.

Simple computations show that gamma-short delta-hedging leads to buying the currency that has appreciated and selling the currency that has depreciated. Similarly, gamma-long delta-hedging would lead to selling the currency that has appreciated and buying the currency that has depreciated. But those who hedge are the market makers who are typically short on volatility. This results in a delta-hedging behaviour that exacerbates the fluctuations of the exchange rates. The more volatility they have sold, the more their hedging behaviour will exacerbate the fluctuation of the exchange rate. But also, the more volatility they have sold, the higher is that volatility. Thus a high volatility indicates that the movements of the spot will be amplified by the delta-hedging strategies.

3 THE SMOOTHING KERNEL METHOD

The empirical analysis in this paper addresses two specific questions.

1) Can tomorrow's volatility of return in the bilateral nominal exchange rates be predicted by today's level of hazard in the global foreign exchange market as represented by the GHI?

2) Is the predictive power of the GHI stronger than that of individual historical volatilities or individual implied volatilities?

The analysis we propose hinges on a "smoothing kernel" method, which can be interpreted as providing a measure of bilateral exchange rate volatility that is conditional on the level of the indicator. The description of the method involves two steps: in the first we define what we mean by conditional volatility, in the second we explain how the strength of the predictive power can be estimated through a weighted regression.

3.1 First step: the conditional volatility of the bilateral exchange rate

We describe the construction of the core concept of the method, which we call the "kernel volatility". First, we present the formula, second we explain the choice of an essential parameter.

3.1.1 Construction of the conditional volatility

Indicating the price of currency M expressed in terms of currency N at time t by $S_{MN,t}$, we define the return on holding a certain currency by $L_{MN,t} = \ln(S_{MN,t})$ (when confusion does not arise we will remove the subscript MN and/or t for ease of notation). Then we denote the value of a generic indicator at a discrete point in time t by Λ_t.

Our aim is to calculate the correlation between Λ_t and the next day's volatility of $L_{MN,t}$ "conditional" on the level of Λ_t. The simplest way to perform such task would be to divide the level of Λ_t into equal intervals, then calculate the correlation between the observations of Λ_t within these intervals and the volatility of $S_{MN,t}$ on the next day, i.e. the standard deviation of $L_{MN,t+1}$-$L_{MN,t}$. The drawbacks of such an approach are the arbitrary choice and the drastic separation of the intervals. A more refined method entails a weighting scheme designed to emphasise the observations with respect to Λ_t that fall close to a certain value (or interval) relative to those falling far from it.

A natural approach is to weight the observations on Λ through a Gaussian function. This method is sometimes called the "smoothing kernel method" in the literature, while the weighting scheme itself is called the "kernel function".[8] The reasons behind this choice are the symmetry of the Gaussian function, its simplicity and its smoothness.[9]

We first describe how these weights are defined and then explain how they can be used to construct the conditional standard deviation of a time series. The indicator Λ_t is a function mapping the risk in a particular currency market onto a positive subset of the real line $H \subset \Re^+$. One chooses, in the subset H, N equally spaced points $h_1, h_2, h_3 \ldots h_N$. Those points are called the reference levels of the indicator. Then one defines a matrix of weights **W** of dimension T x N (with T the number of observations on the indicator series) whose elements are defined as

$$w_{i,t} = \Phi_\xi \left[\ln\left(\frac{\Lambda_t}{h_i} \right) \right] \qquad (3.1)$$

where $\Phi_\xi[x]$ is the Gaussian density distribution with mean zero and standard deviation ξ.[10] We will refer to ξ as to the *width parameter*.

[8] For example, if this weighting scheme were just the indicator function of an interval, it would just give back the elementary equi-spaced intervals mentioned above.

[9] Of course this choice of the kernel function does not imply any assumption on the normality of the exchange rates or the indicator variable.

[10] The value of the parameter ξ is not specified a priori, but as will become evident it is one of the crucial elements of the analysis.

Therefore we define the conditional volatility, or *kernel volatility*, of an exchange rate for each column of the weight matrix as

$$\upsilon_i = \sqrt{\frac{\sum_{t=1}^{T} w_{i,t}(L_t - L_{t-1})^2}{\sum_{t=1}^{T} w_{i,t}} - \left(\frac{\sum_{t=1}^{T} w_{i,t}(L_t - L_{t-1})}{\sum_{t=1}^{T} w_{i,t}}\right)^2} \qquad (3.2)$$

where i is the index over the N intervals of the set H.[11] So, for each time series of a bilateral exchange rate, we will obtain as many υ_i's as the number of points h_1, h_2, h_3 … h_N that we select with each representing a level of risk. In that sense, each of these volatilities is conditional on a particular level of risk.

Note that, as ξ (the width parameter) diverges to infinity, the weights $w_{i,t}$ converge to a constant for any value of i, hence the kernel volatilities tend to the unconditional volatility of the exchange rates, (which implies that we lose all information about the influence of the indicator's level on the future return of the exchange rate). Conversely, as ξ approaches 0, we face a small sample problem, because only few data will be associated with a non-negligible weight. As a consequence, the kernel volatility would be more or less a random function of the level of the indicator, which serves no useful purpose. Hence, too small or too large values of the parameter ξ yield a series of υ_i's of no practical use.

3.1.2 A few words about the choice of the width parameter

Unfortunately, unless we make an hypothesis on the statistical nature of our variables, the choice of the standard deviation ξ in the kernel function cannot be based on theory. It must depend on heuristic considerations. A visual example can help to explain this point. Figure 3 shows the conditional volatilities for the euro-dollar (or dollar-mark for dates before 1 January 1999) when the indicator is the one-month implied volatility of the euro-dollar exchange rate. Each panel corresponds to three values of ξ, namely 0.0005 (too small a value), 0.05 (our choice) and 0.5 (too

[11] Expression (2.2) can be naturally interpreted from the textbook definition of standard deviation of discrete random variable X

$$s = \sqrt{\frac{\sum_{i=1}^{N} f_i X_i^2}{\sum_{i=1}^{N} f_i} - \left(\frac{\sum_{i=1}^{N} f_i X_i}{\sum_{i=1}^{N} f_i}\right)^2}$$

where f_i is the frequency of observation X_i. In essence here the frequencies are substituted by the weights defined by the kernel function.

large a value). For the value $\xi = 0.05$ one can observe that the kernel volatility displays a smooth (and overall increasing) pattern, which is robust to small changes in ξ. In fact the pattern would be virtually identical for values of ξ between 0.035 and 0.1. These features of the kernel volatility are roughly similar for all exchange rates and all the indicators considered in this paper. In essence, the choice of the value of ξ hinges on two criteria: the number of observations entering the calculation with a significant weight and the necessity to clearly separate the periods on the basis of the prevailing risk conditions. To balance these criteria an appropriate strategy would be to choose a value of ξ at the interval over which the changes in the overall pattern of the kernel volatilities remain stable.[12]

3.2 Second step: weighted correlations and linear regressions

In this section we analyse the relationship between the indicators and the kernel volatilities of the exchange rates. A visual inspection of numerous plots highlights this relationship as a power law. In other words, we find that the logarithm of the indicator and the logarithm of the kernel volatility are approximately affine functions[13] of each other. The next step is to estimate the slope of this affine function through of a weighted regression, aimed at filtering out the influence of some kernel volatilities with poor informational content.

3.2.1 Recognising the pattern of the relationship

Figure 4 highlights the relationship between the standard deviation of the next day return and the level of the indicator. The central panel in Figure 3 shows that, for conveniently chosen values of ξ,, to higher levels of the indicator correspond, the next day, large changes of the (log) spot exchange rate. The question is now how to *quantify* this relationship whose qualitative nature is *visually* captured by the graph.

One can see from the second graph in Figure 3 that the suitable statistical tool for this purpose must be some sort of linear regression method. In fact we observed the shape of several graphical plots of the kernel volatilities, and realised that they are closer to the power functions than to the affine function. So we take the logarithm of the kernel volatility on right hand side and the logarithm of the indicator on the left. In other words, we chose the logarithmic transformation on the basis of *heuristic*

[12] An alternative strategy would to compute an optimal value for the parameter ξ. But this entails making some assumptions on the statistical structure of the joint processes of exchange rate and indicator. In our view, the interest of the kernel method is precisely the ability to give hints about a relationship between two variables on a purely empirical basis.

[13] Affine functions are concave and convex at the same time, hence taking the form of constant plus linear functions.

considerations and a visual inspection.[14] Furthermore, we note that the variables on both the right and left hand sides are of the same nature, since the three indicators we consider (historical volatilities, implied volatilities, GHI) are expressed in the same unit of measure. This means that the dependent variable and the explanatory variable would have the same physical dimension, the same order of magnitude and could be expressed through the same conventions.

3.2.2 Filtering out unreliable observations

The next hurdle regards the small sample problem highlighted by the third panel in Figure 4. As we have noted, some kernel volatilities might be noisy, because too few observations fall within the intervals of the indicator. In other words, for those values of the indicator that are rarely reached, a small number of observations dominate. Furthermore, the smaller the standard deviation of the kernel ξ, the fewer the observations effectively associated with any reference level, because the values selected tend to cluster around the mean. So, in running a regression, it would be desirable to reduce the influence of those kernel volatilities that are actually calculated from fewer observations.

Again we resort to a weighting scheme, where the weights should be smaller for those kernel volatilities computed from fewer observations. A simple choice for this second-level weighting scheme would be the sum of the weights $w_{i,t}$ used in the computation of the kernel volatility.[15] In simple terms, if a kernel volatility corresponding to a point h_i is calculated mostly from observations "far apart" from the level h_i, then the sum of the corresponding weights will be small.

For ease of notation, we set

$$\begin{aligned} \mathbf{W}_i &\equiv \sum_{t=1}^{T} w_{i,t} \\ X_i &\equiv \mathrm{Ln}(h_i) \\ Y_i &\equiv \mathrm{Ln}(\upsilon_i) \end{aligned} \tag{3.3}$$

[14] Of course, the zero lower bound for volatilities was reflected in the logs, but not in the levels.

[15] Or it could have been the square root of the sum of the square of those weights. We also tried this possibility, without changing the pattern of our results.

Then we define the six sums:

$$
\begin{aligned}
W &\equiv \sum_{i=i}^{N} W_i \\
WX &\equiv \sum_{i=l}^{N} W_i X_i \\
WY &\equiv \sum_{i=l}^{N} W_i Y_i \\
WX2 &\equiv \sum_{i=l}^{N} W_i\, X_i^2 \\
WXY &\equiv \sum_{i=l}^{N} W_i X_i Y_i \\
WY2 &\equiv \sum_{i=l}^{N} W_i Y_i^2
\end{aligned}
\qquad (3.4)
$$

Then a weighted regression is expressed as

$$\upsilon_{t+1} = \exp\left(m * Ln\left(h_t\right) + b\right) = (\ h_t\,)^m * \exp(b) \qquad (3.5)$$

where m and b are parameters, being defined as

$$
\begin{aligned}
m &= \frac{WXY\ W - WX\ WY}{WX2\ \ W - WX\ WX} \\
b &= \frac{WY\ WX^2 - WX\ WXY}{WX2\ \ W - WX\ WX}
\end{aligned}
\qquad (3.6)
$$

where h_t is the value of indicator h at time t, and where v_{t+1} is the kernel standard deviation of the exchange rate return realised between times t and t+1. Furthermore, a measure of the goodness of fit, in the form of a weighted correlation can be expressed as

$$r = \frac{WXY * W - WX * WY}{\sqrt{\left(WX2 * W - WX * WX\ \right)\left(WY2 * W - WY * WY\right)}} \qquad (3.7)$$

The square of r is analogous to the familiar R^2 textbook goodness-of-fit measure for the weighted regression (3.5).

3.2.3 Measuring predictive powers

Our aim is to provide a measure of the predictive power of a given indicator (historical volatility, implied volatility, GHI) on a given exchange rate. To do so, we actually need to consider not only the value of r, but also the value of m.

The precise numerical value of ξ has in itself no absolute meaning. One can easily show (or numerically check) that for any two series of exchange-rate and indicator, setting ξ equal to a large enough number brings r^2 arbitrarily closed to 1. But we do not know at which speed this convergence takes place, and we certainly have no reason to rule out that this speed depends on the particular exchange-rate and indicator. So we cannot choose a precise ξ such that the r^2 can be compared for two different couples exchange-rate-and-indicator. (And this holds even if they are about the same exchange rate, but different indicators, or the same indicator, but different exchange rates.)

This problem can be addressed if, instead of looking at r^2 alone, we look at r^2 together with the slope m, and decide that a indicator dominates another one when it has a better r^2 and a no worse slope m than the other indicator, or when it has a better slope m and a not worse r^2 than the other indicator.

For this reason, we will report both r and m in the Tables 1 and 2 below in the text.

4 THE EMPIRICAL RESULTS

This section describes the data set and illustrates the results of the regression (3.5) for the risk indicators most commonly used in the currency market, i.e. historical volatilities, implied volatilities, and for the GHI.

4.1 Data and main results

The data set contains time series of 1715 daily data from March 1st 1993, to November 23rd 1999 from two sources:

- The exchange rates are taken from the Bank for International Settlements (BIS) database, which records synchronous observations taken daily at 2.15 p.m. C.E.T. Therefore they can be used to compute the cross rates accurately.

- The implied volatilities from currency options on the three major exchange rates have been kindly provided by Chase Manhattan Bank. The data are taken daily at London's close. We have used only one-month implied volatilities.

The GHI and the historical volatilities are calculated from these raw data. To account for the fact that London currency market closes after 2.15 p.m. C.E.T., we put the implied volatility (or the GHI) dated at t-1 in relation to the return realised between t and t+1 (otherwise the indicator would have reflected part of the information on the return). Actually, had we used London's close synchronous daily data for the exchange rates, the results would have been better.

We run the linear regression (3.5) for 12 exchange rates, namely (in alphabetical order and using the notation by Reuters) AUD=, AUDJPY=, AUDNZD=, CAD=, CHF=, EUR=, EURCHF=, EURGBP=, EURJPY=, GBP=, JPY= and NZD=. We have tried 18 indicators, namely the 12 corresponding one-month historical volatilities, the seven most liquid one-month implied volatilities (namely those on EUR=, JPY= , EURJPY= , CHF= , EURCHF= , EURGBP= , GBP=) and the one-month GHI.

The main results of the regressions are illustrated in Figures 5 to 15. The outcome of the empirical analysis is remarkably sharp: the regression lines provide an extremely good fit for the kernel volatility of exchange rates when the GHI is taken as the indicator of aggregate risk. In many cases bilateral implied volatilities also prove to have excellent predictive properties, while the picture for historical volatilities is decidedly less encouraging.

All in all, the results indicate that the best indicator, for a given currency pair, is generally its own *implied* volatility (when it is available, i.e. when it exists and is based on liquid contracts.) More surprisingly, they also indicate that the *aggregate* conditions of risk, as measured by the GHI, exert a strong influence on the volatility of the next day's return on most currency pairs. This influence, considering that the graphs depict the power-law relationship (3.5), increases with the aggregate risk; in most case they increase more than proportionally. The slope of the regression line differs among currencies; in particular, the sensitivity to global risk conditions is more marked for the small countrys' currencies. Hence, in periods of turmoil on the global foreign exchange market, the <u>currencies of especially smaller countries are likely to be swept by factors that have little relation to the changes in economic fundamentals directly affecting their economies.</u> More detailed results are presented in the next three subsections.

4.2 Discussion of the results

4.2.1 An overview on the major bilateral exchange rates

Figures 5 to 11 report the results of the weighted regressions for the three major exchange rates over the GHI for four values of the width ξ, i.e. the standard deviation of the kernel function. First, we note that the correlation coefficient is almost always remarkably high. Second, the slope of the regression line is always positive, but varies monotonically with changes in the value of the parameter

ξ. Nevertheless, we observe that the slope remains rather stable as ξ increases, so – in practice – a value comprised between 0.05 and 0.015 would be suitable for predictive purposes.

It is also noteworthy (Figures 5 to 7) that the bilateral exchange rates of the yen are more reactive to the GHI, with a slope around unity, i.e. double the slope of the regression involving the euro-dollar. The intercept is roughly similar in all graphs, especially those in Panels 2 and 3, a feature that stresses a common reaction to mild risk conditions.

A useful comparison can be made between the results in Figures 5 to 7 and those in Figures 8 to 10 where the implied volatility is taken as the indicator of risk. The main differences can be detected in the comparison between Figures 5 and 8 referring to the euro-dollar exchange rate, which is more reactive to its own volatility than to the GHI. Moreover, the values of ξ beyond 0.08 lead to sizeable changes in the values of the estimated parameter m. By contrast, the analogous comparison for the euro-yen and the dollar-yen (Figures 6 and 9 and Figures 7 and 10, respectively) does not display profound differences in the values of m and of the intercept (see especially Panels 3). This feature might be due to the fact that over the period examined the dollar-mark volatility influenced other volatilities rather than the opposite.

The last result referring to major exchange rates is shown in Figure 11, displaying the regression lines of the euro-dollar and euro-yen on their respective historical volatilities. One can note that the values of m are much smaller and the correlation somewhat weaker. In other words, the historical volatility does not have as strong predictive properties as the GHI and the implied volatilities that are *forward-looking measures of risk*. However, this peculiarity is not of a general nature: in fact, results for the dollar-yen historical volatility (not reported in the graphs) display similar features as those in Figure 7 (with m = 0.91 and m = 0.79 when ξ is equal to 0.08 and 0.15 respectively).

4.2.2 An overview on the exchange rates of selected small open developed economies

Tables 9 to 13 present a sample of the weighted regressions involving currencies of smaller open economies with a prominent role in the world foreign exchange market. The overall results are more variegated than those described in 4.2.1.

Both the exchange rate of the Swiss franc vis-à-vis the US dollar (referred to as the "Swissie") and the exchange rate of the pound sterling vis-à-vis the US dollar (referred to as the "Cable") display a relationship with the GHI that, although notable, is far weaker than that displayed by the major exchange rates. The slopes of the coefficients m are very close, especially for ξ = 0.15 and ξ = 0.35, but the intercept is sensibly higher for the Swissie. These two currencies play a crucial role in the portfolio diversification strategy of major international investors, because they are perceived as a safe haven (especially the Swiss franc) during times of turmoil. The low value of m confirms that, indeed, an increase in global risk is not likely to have a major impact on the volatility of those currency pairs. A second possibility, which is more relevant for the EUR/CHF volatility, is that the financial ties of UK and Switzerland with Japan are weak.[16] Actually, the curvilinear shape in Panels 2 and 3 of Figure 13 seems to indicate that an increase of the GHI towards extreme values is associated with a reduction in the volatility of the pound sterling on the next day. One must remember, however, the caveat made in Section 3.2.2, namely that this effect might result from the small sample bias: in other words, as the GHI rarely exceeded 17% between March 1993 and November 1999,[17] the last two kernel volatilities are calculated from few observations.[18]

The resilience of the returns on these two currencies is also confirmed by the regression involving the euro-dollar implied volatility which yields about the same coefficients as in Figures 12 and 13 (see Table 1).

In contrast to the Cable and the Swissie, both the Australian and the New Zealand dollar display a much sharper sensitivity to global hazard, as can be perceived in Figure 14, where the histograms fit a straight line almost perfectly and the slopes are almost of the same magnitude for the two currencies. The Canadian dollar, on the other hand, has a less pronounced, but rather significant relationship to the GHI (Figure 15).

It is important to highlight those cases where the GHI does not appear to have any influence on the next day's returns. Figure 16 reports the results of the weighted regression for the German mark – Swiss franc and the pound sterling – German mark. In the first case, we see that the slope of the line is practically zero, while the slope in the second, although positive, is rather small. We will see in the next section that these exchange rates represent an anomaly in the sense that they are the only ones whose return variance is little affected by the GHI.

Finally Figure 17 presents another interesting example: the exchange rate of the Australian dollar vis-à-vis the Japanese yen. The panels show how both the GHI and the implied volatility of the dollar-yen exchange rate exert a great influence, roughly of the same magnitude, on the variance of returns.

[16] This possibility was suggested to us by an anonymous referee.

[17] Note that, after November 1999, such GHI levels have become more frequent.

[18] The second weighting scheme introduced in 3.2.2 is indeed aimed at insuring that this bias does not affect the regression curves.

4.2.3 The overall results

The complete results referring to all currencies and all indicators are summarised in Tables 1 and 2 [19], reporting the estimated slopes and the correlation coefficients from regression (3.5) respectively for the entire set of currencies and indicators.

The main conclusion is that the GHI displays predictive properties over the volatility of the bilateral exchange rate returns that are at least comparable to the other indicators and are, in most cases, better. In particular for the three main currencies, we observe a number of stylised facts in Tables 1 and 2.

- For the dollar-yen and euro-yen, the GHI outperforms all other indicators, including the implied volatilities of the corresponding currency pairs. (Its results are better on Table 1, while being not worse on Table 2.)
- The historical volatilities are not particularly good predictors, except for the exchange rate of the US dollar vis-à-vis the Canadian dollar (and to a lesser extent the German mark vis-à-vis the Swiss franc).
- The implied volatility of the euro yen and the dollar yen has remarkably good predictive power over the volatilities of the Australian and New Zealand dollar.
- The GHI does not perform well vis–à-vis the German mark-Swiss franc, whose correlation is negative and low. The euro-Swiss franc is more highly correlated with its own implied volatility and to a lesser extent with the German mark-pound sterling implied volatility.
- The natural ranking in terms of forecasting properties can be indicated as follows: The best indicator tends to be the own implied volatility of the currency pair, (if it exists and is liquid), then the GHI, then some implied volatilities of related currency pairs, then the own historical volatility.

[19] The results of the 2 tables should be simultaneously taken into account. This is because the correct value of the parameter ξ is only guessed,, and because ξ big enough will always bring the correlation to 100% - so an apparently excellent result, if one considers Table 2 only.

Table 1- Estimated slope of the regression (with ξ = 0.05)

	AUD=	AUDJPY=	AUDNZD=	CAD=	CHF=	EUR=	EURCHF=	EURJPY=	GBP=	GBPEUR=	JPY=	NZD=
GHI	**1.1**	**0.68**	**0.63**	**0.66**	**0.35**	**0.57**	**-0.11**	**1.2**	**0.4**	**0.3**	**1.09**	**1.22**
CHF=	0.06	0.73	0.7	0.53	0.9	0.95	0.33	0.32	0.45	0.51	0.77	-0.01
EUR=	0.22	0.71	0.69	0.4	0.85	0.99	0.24	0.51	0.7	0.43	0.76	0.04
EURCHF=	0.74	0.47	0.33	0.64	0.25	0.24	0.95	1.12	0.44	0.46	0.83	0.51
EURGBP=	0.67	0.54	0.52	0.46	0.61	0.65	0.56	1	0.51	0.96	0.92	0.75
EURJPY=	1.03	0.42	0.3	0.53	-0.02	0.18	-0.05	1.15	0.37	0.05	0.82	1.14
GBP=	0.24	0.45	0.38	0.29	0.54	0.66	0.43	0.68	0.67	0.48	0.49	0.1
JPY=	0.94	0.58	0.57	0.51	0.25	0.42	-0.17	1.02	0.15	0.34	0.96	1.12
HV AUD=	0.61	0.29	0.12	0.33	0.11	0.2	-0.05	0.64	0.29	-0.02	0.47	0.75
HV AUDJPY=	0.26	0.56	0.28	0.32	0.31	0.38	0.14	0.7	0.14	0.16	0.69	0.05
HV AUDNZD=	-0.01	0.31	0.22	0.17	0.31	0.37	0	0.45	0.1	0.24	0.47	0.23
HV CAD=	0.4	0.43	0.26	0.64	0.18	0.13	0.14	0.49	-0.01	0.02	0.55	0.4
HV CHF=	0	0.4	0.09	-0.09	0.41	0.45	0.05	0.15	0.25	0.09	0.45	-0.31
HV EUR=	0.25	0.44	0.26	0.18	0.44	0.54	0.18	0.56	0.46	0.32	0.56	0.12
HV EURCHF=	0.29	0.09	-0.04	0.04	0.05	0	0.44	0.43	0.24	0.25	0.25	0.08
HV EURJPY=	0.54	0.3	0.18	0.37	0.04	0.08	0.16	0.77	0.26	0.08	0.55	0.51
HV GBP=	0.37	0.31	0.17	0.08	0.26	0.35	0.2	0.55	0.55	0.25	0.43	0.16
HV GBPEUR=	0.3	0.31	0.17	-0.03	0.32	0.33	0.31	0.24	0.23	0.56	0.33	0.14
HV JPY=	0.35	0.35	0.23	0.19	0.15	0.19	-0.06	0.51	0	0.14	0.58	0.34
HV NZD=	0.36	0.1	0.09	0.19	-0.13	-0.03	-0.24	0.36	0.01	0.15	0.17	0.61

Table 2 – The conditional correlations (with $\xi = 0.05$)

Currency	AUD=	AUDJPY=	AUDNZD=	CAD=	CHF=	EUR=	EURCHF=	EURJPY=	GBP=	GBPEUR=	JPY=	NZD=
GHI	**97%**	**98%**	**92%**	**93%**	**80%**	**93%**	**-30%**	**99%**	**82%**	**95%**	**96%**	**97%**
CHF=	12%	95%	83%	82%	95%	96%	63%	44%	82%	86%	84%	-2%
EUR=	46%	97%	88%	66%	95%	98%	57%	70%	85%	83%	83%	7%
EURCHF=	81%	84%	78%	73%	50%	53%	96%	87%	81%	92%	93%	65%
EURGBP=	85%	88%	93%	75%	80%	84%	75%	95%	94%	95%	95%	82%
EURJPY=	97%	76%	78%	93%	-7%	45%	-17%	99%	87%	16%	93%	96%
GBP=	41%	85%	81%	49%	89%	94%	60%	90%	95%	87%	73%	17%
JPY=	98%	96%	90%	86%	73%	91%	-54%	98%	40%	92%	96%	98%
HV AUD=	93%	76%	42%	83%	37%	55%	-18%	93%	86%	-6%	88%	92%
HV AUDJPY=	58%	92%	83%	74%	79%	86%	40%	99%	49%	65%	92%	20%
HV AUDNZD=	-3%	80%	68%	40%	78%	81%	-2%	75%	66%	61%	78%	55%
HV CAD=	75%	88%	73%	95%	53%	37%	48%	77%	-2%	13%	82%	73%
HV CHF=	0%	93%	37%	-33%	69%	80%	18%	48%	87%	20%	85%	-74%
HV EUR=	57%	88%	68%	41%	74%	91%	46%	88%	92%	86%	80%	27%
HV EURCHF=	58%	23%	-10%	8%	17%	1%	87%	62%	69%	78%	42%	18%
HV EURJPY=	92%	75%	62%	86%	16%	29%	58%	97%	79%	40%	90%	84%
HV GBP=	77%	87%	64%	20%	88%	94%	56%	85%	93%	81%	81%	37%
HV GBPEUR=	56%	68%	52%	-6%	90%	86%	83%	46%	74%	94%	58%	27%
HV JPY=	87%	85%	76%	64%	47%	56%	-35%	93%	0%	71%	93%	69%
HV NZD=	87%	29%	33%	70%	-36%	-8%	-69%	85%	2%	45%	38%	96%

4.3 How reliable are the results?

The last question regards the significance of the weighted regression estimates. In other words, how likely is it that the positive slopes of the lines result from spurious regressions?

To answer this question, we conduct an experiment: we generate 200 realisations of Brownian motion and we carry out exactly the same estimates as in (3.5) substituting the indicator with the synthetic series of Brownian motion. In theory, the slope of these regression lines should cluster around zero, but for some realisations of the Brownian motion, the deviations from zero can obviously be quite large.[20]

[20] In fact, in some cases the slope of the regression line obtained with the synthetic series exceed those of the regressions run with the indicators.

Table 3 - Significance test against a Brownian motion with (ξ = 0.05)

Currency	AUD=	AUDJPY=	AUDNZD=	CAD=	CHF=	EUR=	EURCHF=	EURJPY=	GBP=	GBPEUR=	JPY=	NZD=
GHI	**0%**	**0%**	**0%**	**3%**	**5%**	**0%**	**65%**	**0%**	**5%**	**1%**	**0%**	**1%**
CHF=	48%	0%	0%	5%	0%	0%	10%	33%	2%	2%	4%	50%
EUR=	41%	0%	0%	15%	0%	0%	12%	18%	1%	4%	5%	49%
EURCHF=	7%	2%	1%	4%	10%	12%	0%	2%	3%	1%	1%	31%
EURGBP=	7%	1%	0%	8%	1%	2%	4%	1%	0%	0%	1%	17%
EURJPY=	0%	3%	4%	2%	53%	25%	57%	1%	4%	35%	1%	0%
GBP=	42%	0%	3%	27%	2%	0%	6%	5%	0%	3%	9%	47%
JPY=	1%	0%	1%	4%	8%	1%	77%	1%	29%	1%	1%	0%
HV AUD=	5%	2%	24%	6%	28%	18%	69%	4%	4%	49%	4%	5%
HV AUDJPY=	34%	1%	2%	12%	2%	2%	27%	0%	23%	14%	3%	52%
HV AUDNZD=	45%	2%	7%	27%	5%	4%	45%	21%	16%	13%	6%	41%
HV CAD=	27%	3%	7%	0%	16%	30%	23%	17%	51%	41%	6%	27%
HV CHF=	43%	1%	29%	59%	3%	2%	37%	43%	4%	38%	4%	59%
HV EUR=	33%	1%	6%	30%	4%	2%	20%	6%	0%	7%	7%	49%
HV EURCHF=	31%	28%	53%	45%	34%	46%	2%	23%	12%	8%	35%	45%
HV EURJPY=	3%	4%	11%	4%	37%	34%	15%	1%	7%	31%	2%	21%
HV GBP=	20%	1%	9%	40%	2%	1%	12%	9%	1%	5%	6%	38%
HV GBPEUR=	34%	3%	16%	48%	1%	4%	4%	44%	8%	1%	23%	41%
HV JPY=	14%	1%	6%	22%	21%	15%	59%	5%	48%	11%	1%	36%
HV NZD=	13%	29%	26%	15%	68%	52%	81%	13%	44%	26%	33%	3%

Table 3 summarises the results of the experiment reporting the proportion of the 200 estimations run with the realisations of Brownian motion that results in a slope and a goodness-of-fit larger in absolute value than the slope of the corresponding regression line obtained with the data on the risk indicators. For example, the first entry in the first row of the table indicates that the regression on the simulated series never yields estimates of m larger than 1.10, i.e. the estimated coefficient in the regression AUD= on the GHI at the same time when estimates of the correlation were larger than 97%, i.e. the estimated correlation in the regression AUD= on GHI.

In general the estimates of (3.5) obtained with the GHI are again highly significant with the two exceptions of the Swiss franc and the pound sterling already highlighted above. In fact, while Table 1 and Table 2 gave the GHI a good, but not clearly dominant score, Table 3 actually does single the GHI out: No other indicator get so many 0% or 1%. In particular the three main currency pairs get a 0%. One cannot say that the GHI outperforms any other indicator since the GHI can be outperformed, for a given currency pair, by the implied volatility of that given currency pair. Even so, it is not by a lot. Table 3 indicates that the GHI indicator obtains excellent results for 8 or 9 currency pairs, while no non-global indicator does.

Not surprisingly, the least significant regressions are those performed on the historical volatility series. A visual impression of the results is contained in Figure 18, which provides the "topographic" view of Table 3. The dark areas represent the results that are most significant, i.e. those that are least likely to be produced by two completely independent time series.

5 CONCLUSIONS

Predictions of exchange rate fluctuation have always proved to be an extremely challenging endeavour. This study investigates whether the volatility of bilateral nominal daily exchange rates can be predicted from the level of historical volatility, implied volatility and the Global Hazard Index. The GHI is derived on the basis of the no-arbitrage hypothesis as a function of implied volatilities extracted from currency options on major exchange rates. It expresses a measure of wild risk in the sense of Mandelbrot (1999) in a multi-currency exchange market.

In carrying out this analysis, we have developed the concept of "kernel correlation", which is essentially calculated by attaching a set of weights to the observations so as to emphasise those in correspondence of a certain phenomenon. In practice, this "kernel volatility" or conditional volatility allows emphasis to be laid on the correlation between an indicator and a particular exchange rate, conditional on the level of the indicator. We resort to the kernel volatility because standard regression analysis would not highlight that the relationship can be stronger when the level of risk is high and weaker in "normal times". By contrast, the kernel volatility method allows examining if the influence of the indicators on the exchange rate changes as the risk level increases.

The main results are summarised in Tables 1 and 2 where the slope of the weighted regression and the conditional correlation are reported. The results we obtain indicate that the implied volatilities and the GHI tend to have a significant influence on the variance of major exchange rates and of a number of smaller currencies on the next day. We will summarise our results as follow: A main finding is that indicators based on implied volatilities, including the GHI, generally display a stronger influence on the variance of the exchange rates than the historical volatilities. For the dollar-yen and the euro-yen currency pairs the GHI displays a stronger influence than the own implied volatility. This is not the case for the other currency pairs investigated. Furthermore, it is interesting to observe that the influence of the GHI is the most significant in the sense that it is the less likely to come from a spurious regression.

References

Bernhardsen, T. and Ø. Røisland (2000): "Factors that influence the Krone exchange rate", Norges Bank, Economic Bulletin, No. 4, December, pp. 143-151

Brousseau, V. and F. Scacciavillani (1999) "A Global Hazard Index for the World Foreign Exchange Markets", European Central Bank, Working Paper Series No. 1, May.

Black, F. and M. Scholes (1973) "The Pricing of Options and Corporate Liabilities", Journal of Political Economy, Vol. 81, May/June, pp. 637-54.

Butler, C. and N. Cooper (1997) "Implied Exchange Rate Correlations and Market Perceptions of European Monetary Union", Bank of England Quarterly Bulletin, No. 4, November, pp.413-423.

De Rosa, D.F. (ed.) (1998) "Currency Derivatives", Wiley & Sons, New York, NY.

Duffie, D. (1992) "Dynamic Asset Pricing Theory", Princeton University Press, Princeton, NJ.

Esposito, M. and E. Laruccia (1999) "The General Contribution of the Euro to the General Riskiness of the Foreign Exchange Markets", Banca Commerciale Italiana, May, pp. 1- 2.

Fornasari, F. (2000) "Currencies And Global Risk", Lehman Brothers, Global Foreign Exchange Strategies, 20 April, pp 4-6.

Garman, M. B. and S. W. Kohlhagen (1983) "Foreign Currency Option Values", Journal of International Money and Finance, Vol. 2, pp. 231-237.

Isard, P. (1995) "Exchange Rate Economics", Cambridge University Press, Cambridge, England.

Jaatinen, S. (2001) "The 1M ECB Global Hazard Index is again trading at an important trendline resistance level", Invest Avenue, Canada, 16 March, http://www.investavenue.com/cgi-bin/enzine.cgi/art/918.html

Jorion, P. (1995) "Predicting Volatility in the Foreign Exchange Market", Journal of Finance, Vol. L, No 2, June.

Lehman Brothers (1999) "Measuring Global Foreign Exchange Risk", in Global Foreign Exchange Strategies, 23 July, pp. 8-9.

Lehman Brothers (2000) "Volatility Risk Premiums In Currency Markets", Analytical Research Series in International Fixed Income Research, March.

Llewellyn, J. (2000) "Global Letter", Lehman Brothers, Global Weekly Economic Monitor, 19 May.

MacDonald, R. and M. P. Taylor (1992) "Exchange Rate Economics: a Survey", IMF Staff Papers, Vol. 39, pp.1-57.

Mandelbrot, B. (1997a) "Fractals and Scaling in Finance: discontinuity, concentration and risk", Springer Verlag, New York.

Mandelbrot, B. (1997b) "Fractales, Hasard et Finance", Flammarion, Paris, France.

Mezrich, J.J. and A.M. Majewska (2000) "Forecasting Volatility to Predict Stock Markets and Currencies", Morgan Stanley Dean Witter, Global Equity and Derivative Markets, January, pp 1-11.

Norges Bank (1999) Inflation Report, Fourth Quarter

Patterson, B., D. Sienkiewicz and X. Avila (2000): "Exchange Rates and Monetary Policy", European Parliament, Working Paper No. 120, August.

Redeker, H.G. (2001) "Long US equity positions…", Invest Avenue, Canada, 20 March, http://www.investavenue.com/cgi-bin/enzine.cgi/art/927.html

Peters, E.E. (1991): "Chaos and Order in the Capital Markets", John Wiley, New York

World Bank (1999) Interest Bearing Notes Newsletter, Vol. 2, No. 3.

Figure 1 - The Global Hazard Index over the 4 last years
(Logarithmic scale. Source: Reuters TS1)

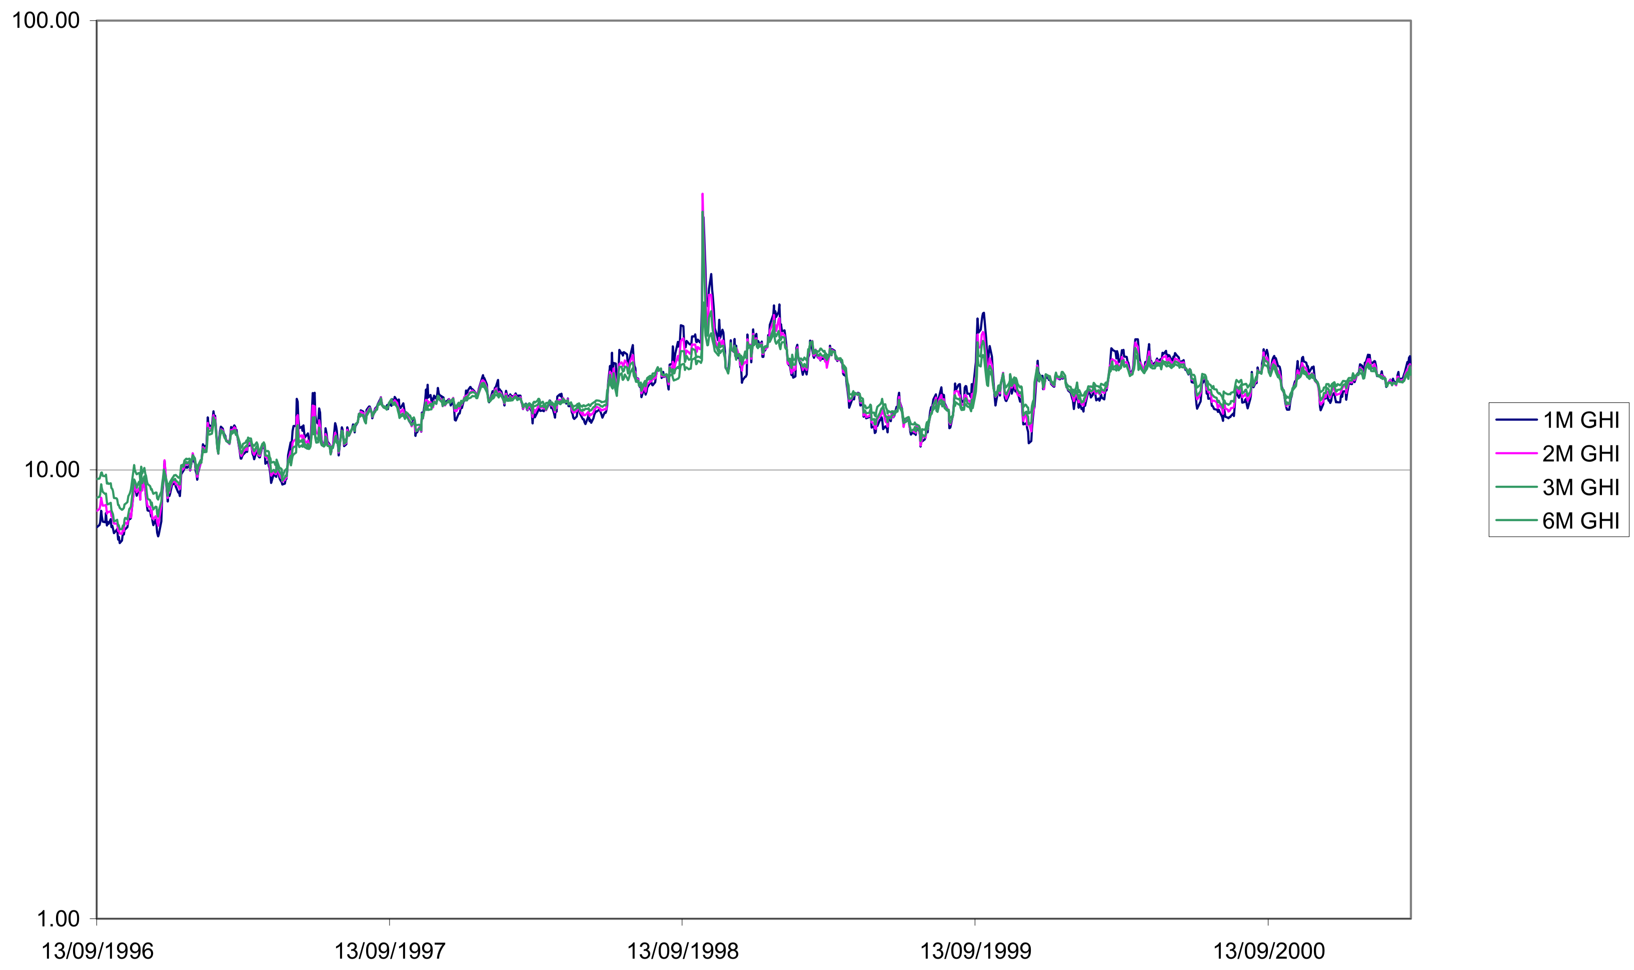

Figure 2 - The GHI and the exchange rates of small currencies
(Sources: Lehman Brothers (left, AUD vs. SDR) and Norges Bank (right, NOK vs. DEM))

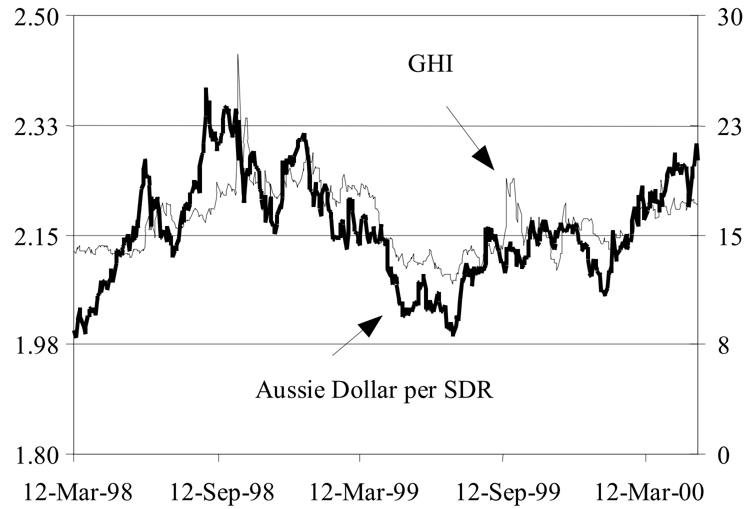

Chart 1 *Developments in GRI and the krone exchange rate*

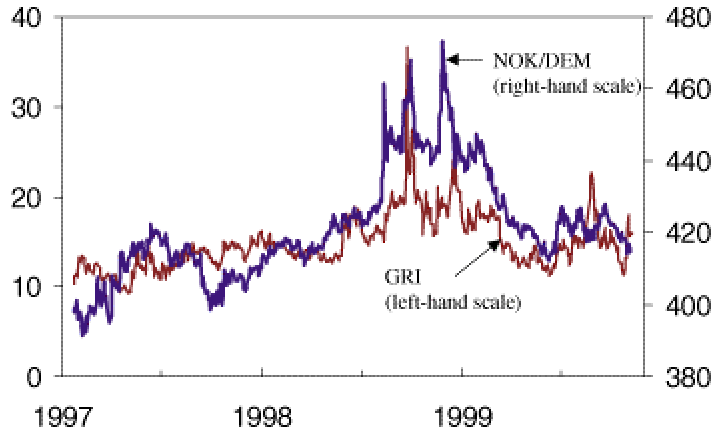

Figure 3 - Kernel volatilities. Spot: Dollar-mark   Indicator: Dollar-mark implied volatility.

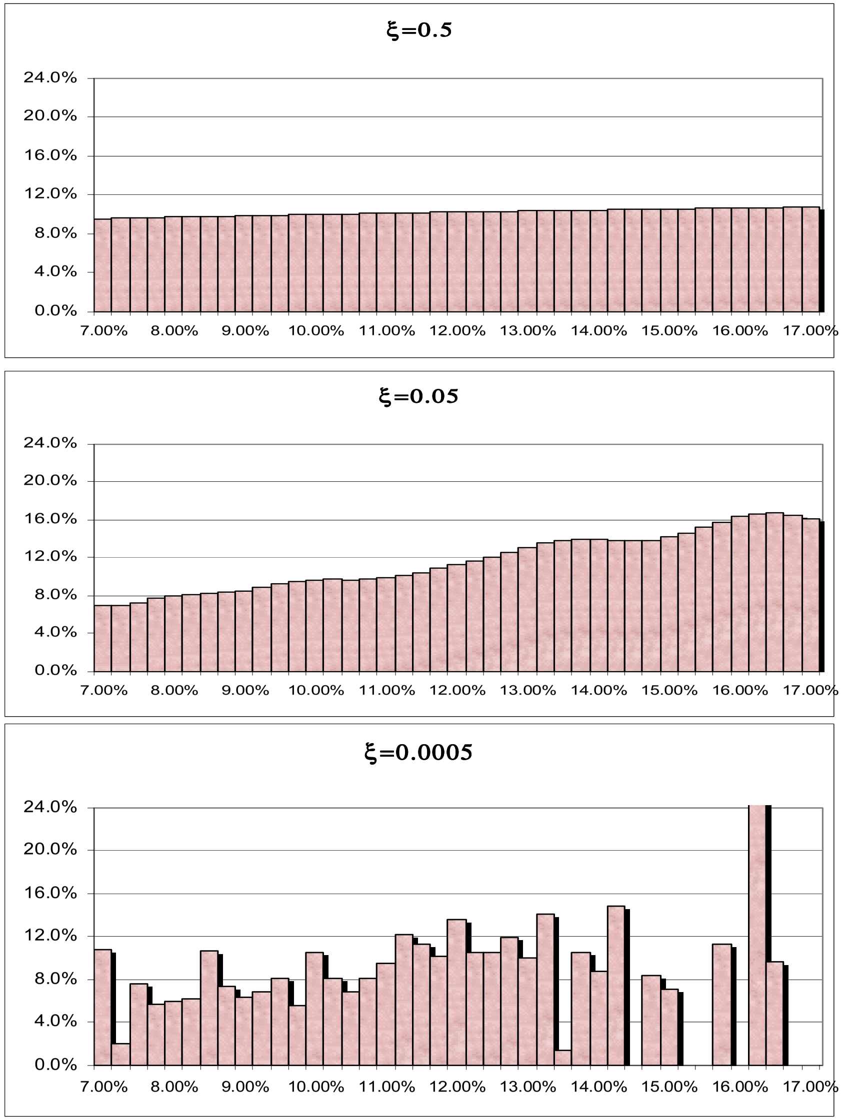

Figure 4 - Autocorrelations of returns and of squared returns of various exchange rates
(1st March 93 - 23 Nov. 99)
(Horizontal axis: Time lag, in working days)

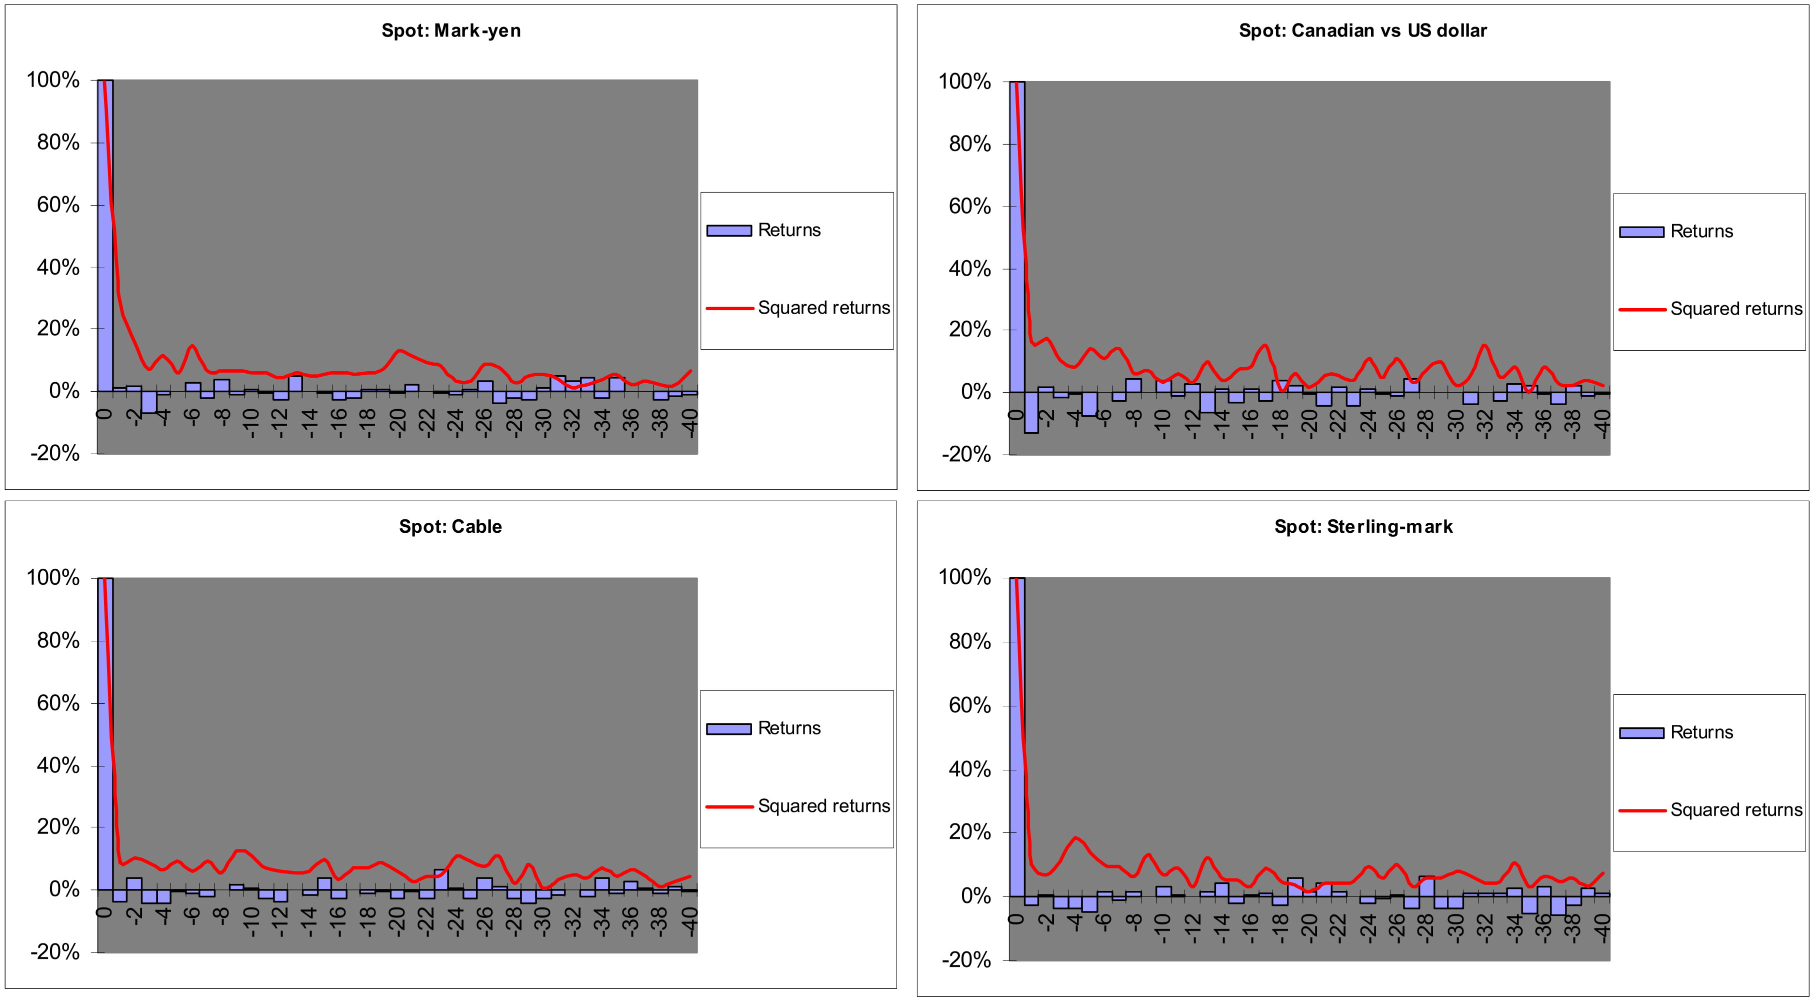

Figure 5 - Weighted regression line of kernel volatilities for the euro-dollar on the GHI with different ξ (in logs of daily data)

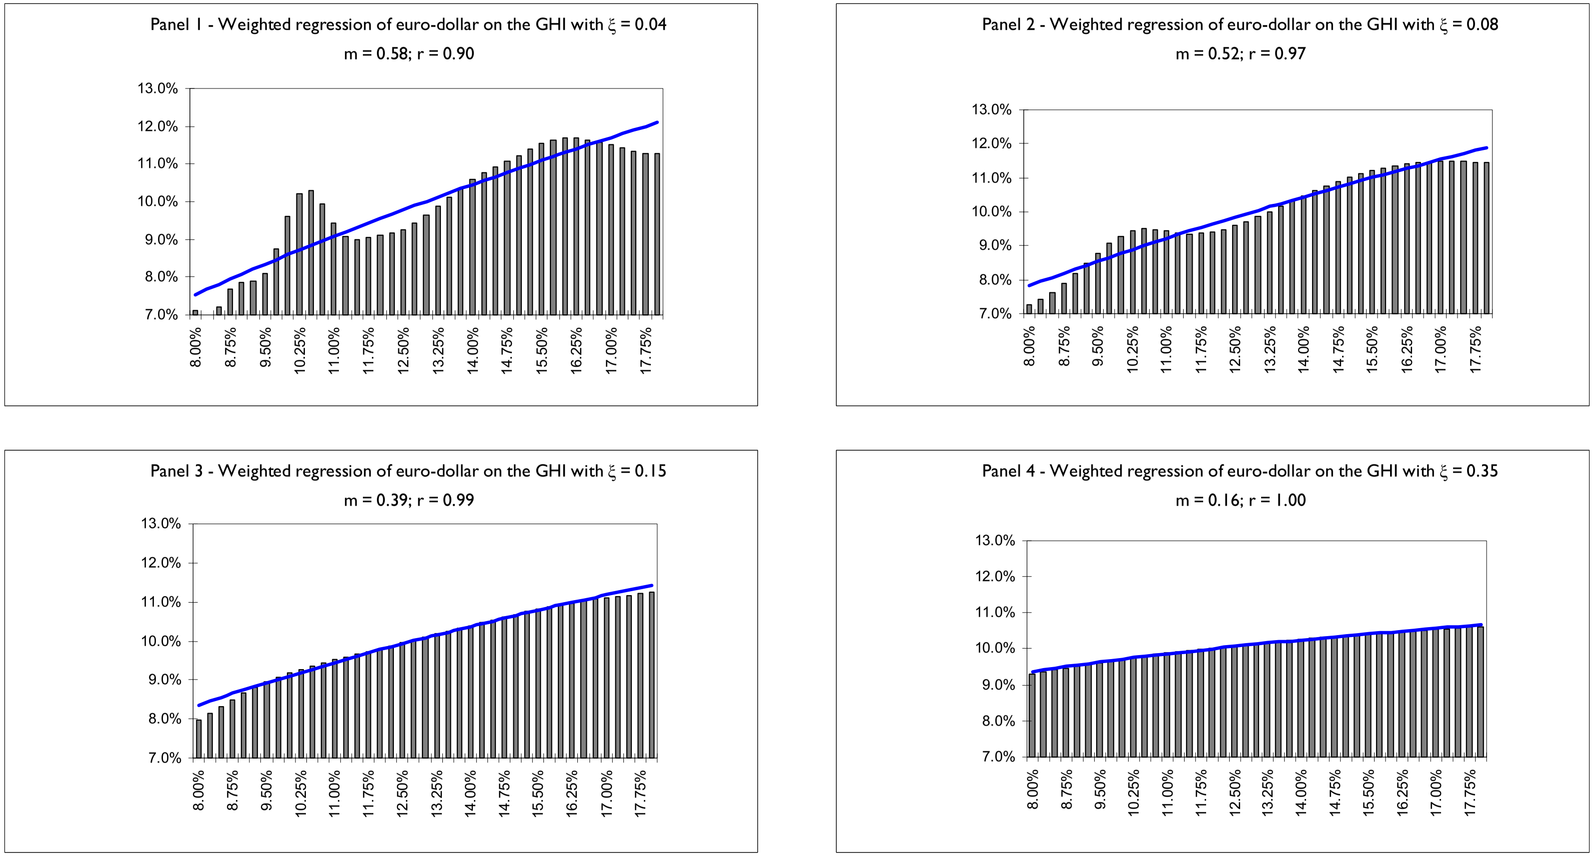

Figure 6 - Weighted regression line of kernel volatilities for the euro-yen on the GHI with different ξ (in logs of daily data)

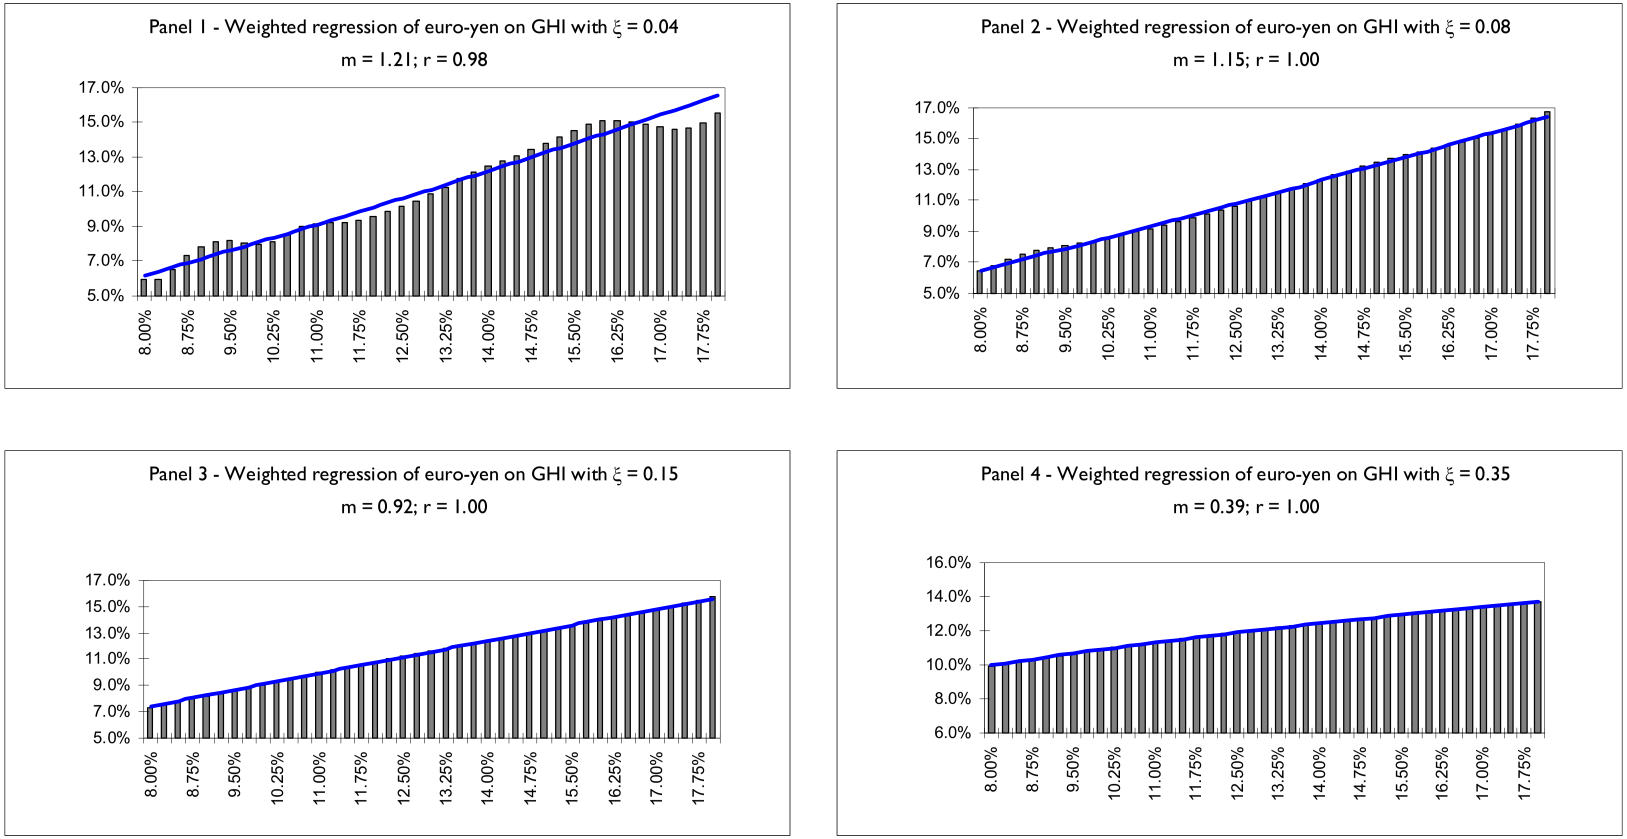

Figure 7 - Weighted regression line of kernel volatilities for the dollar-yen on the GHI with different ξ
(in logs of daily data)

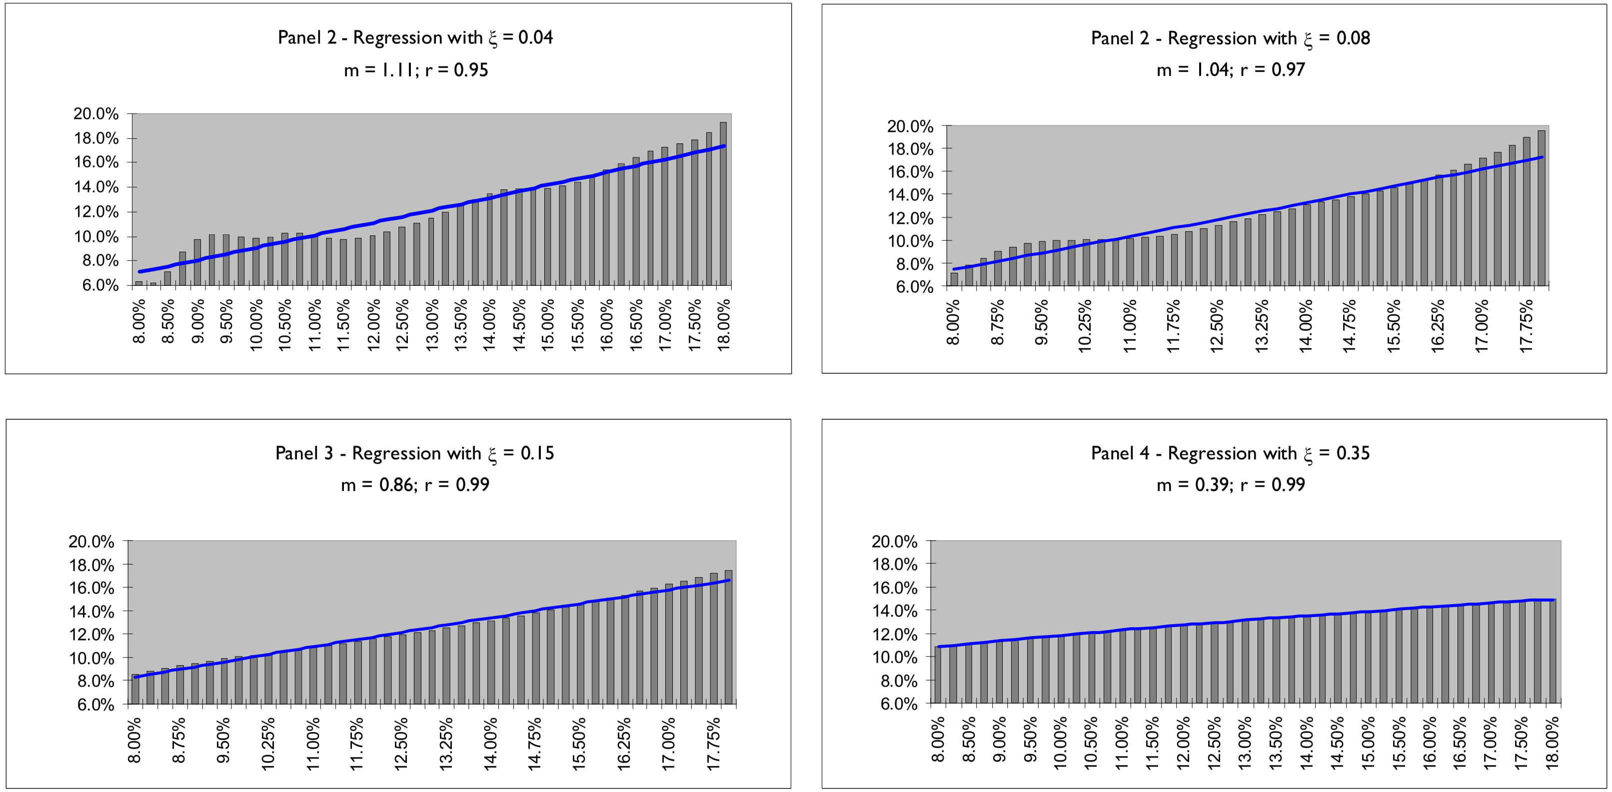

Figure 8 - Weighted regression line of kernel volatilities for the euro-dollar on its implied volatility with different ξ (in logs of daily data)

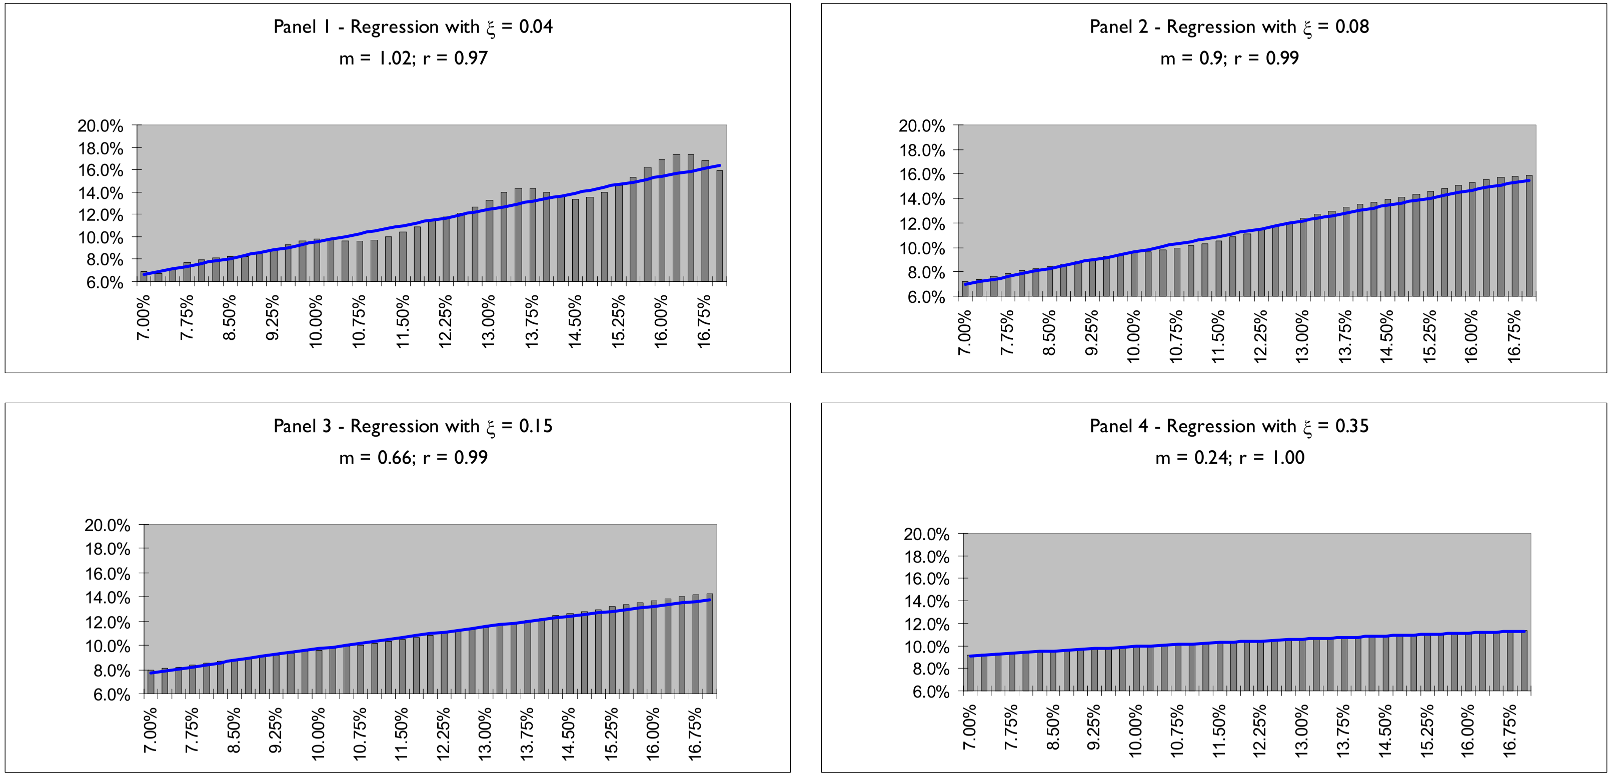

Figure 9 - Weighted regression line of kernel volatilities for the euro-yen on its implied volatility with different ξ (in logs of daily data)

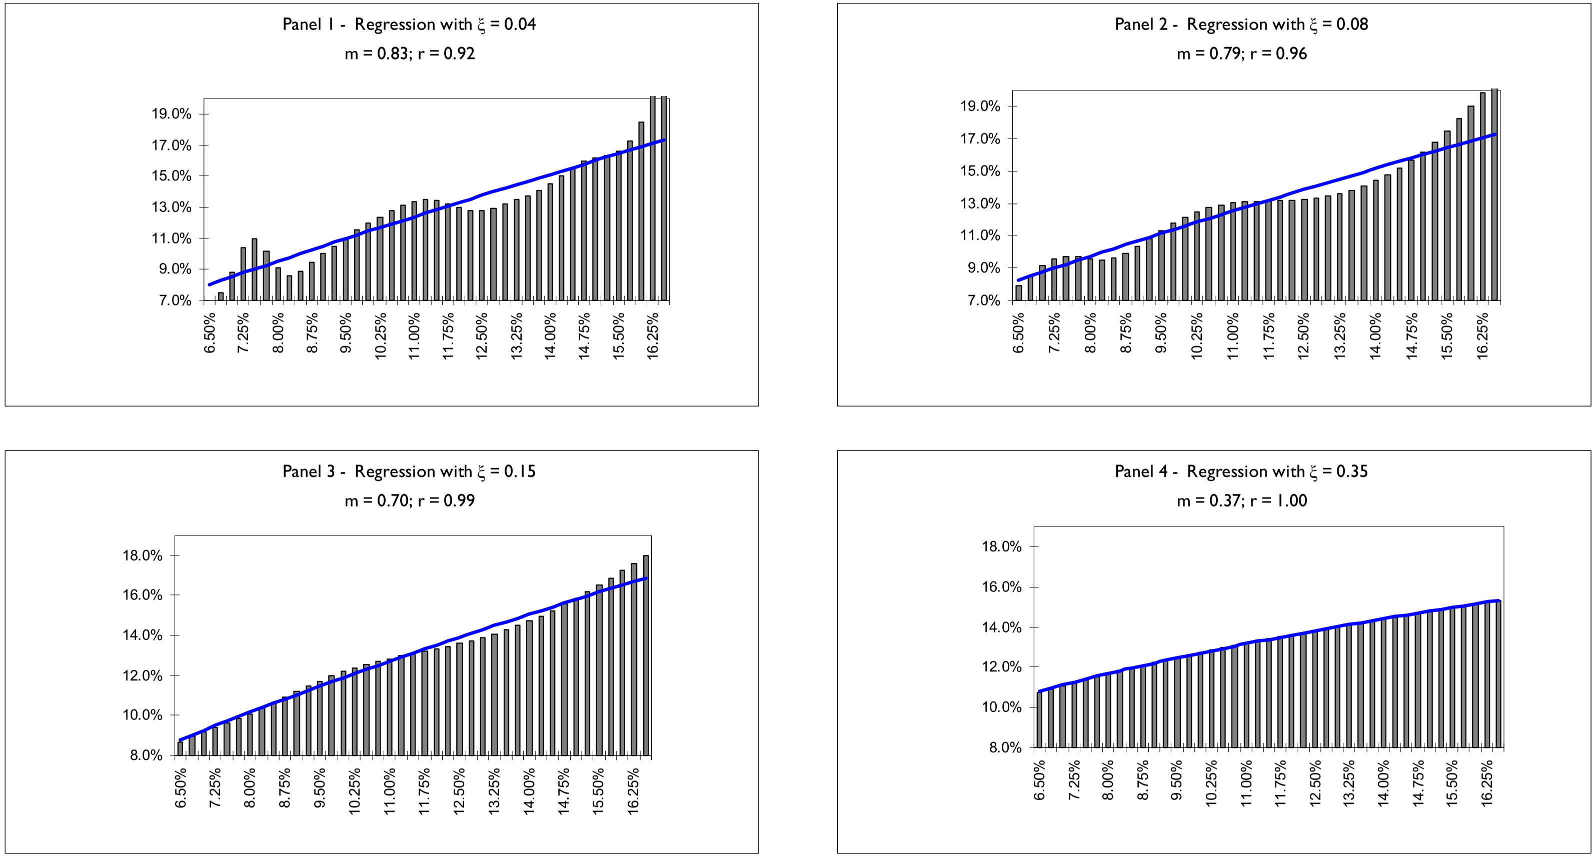

Figure 10 - Weighted regression line of kernel volatilities for the dollar-yen on its implied volatility with different ξ (in logs of daily data)

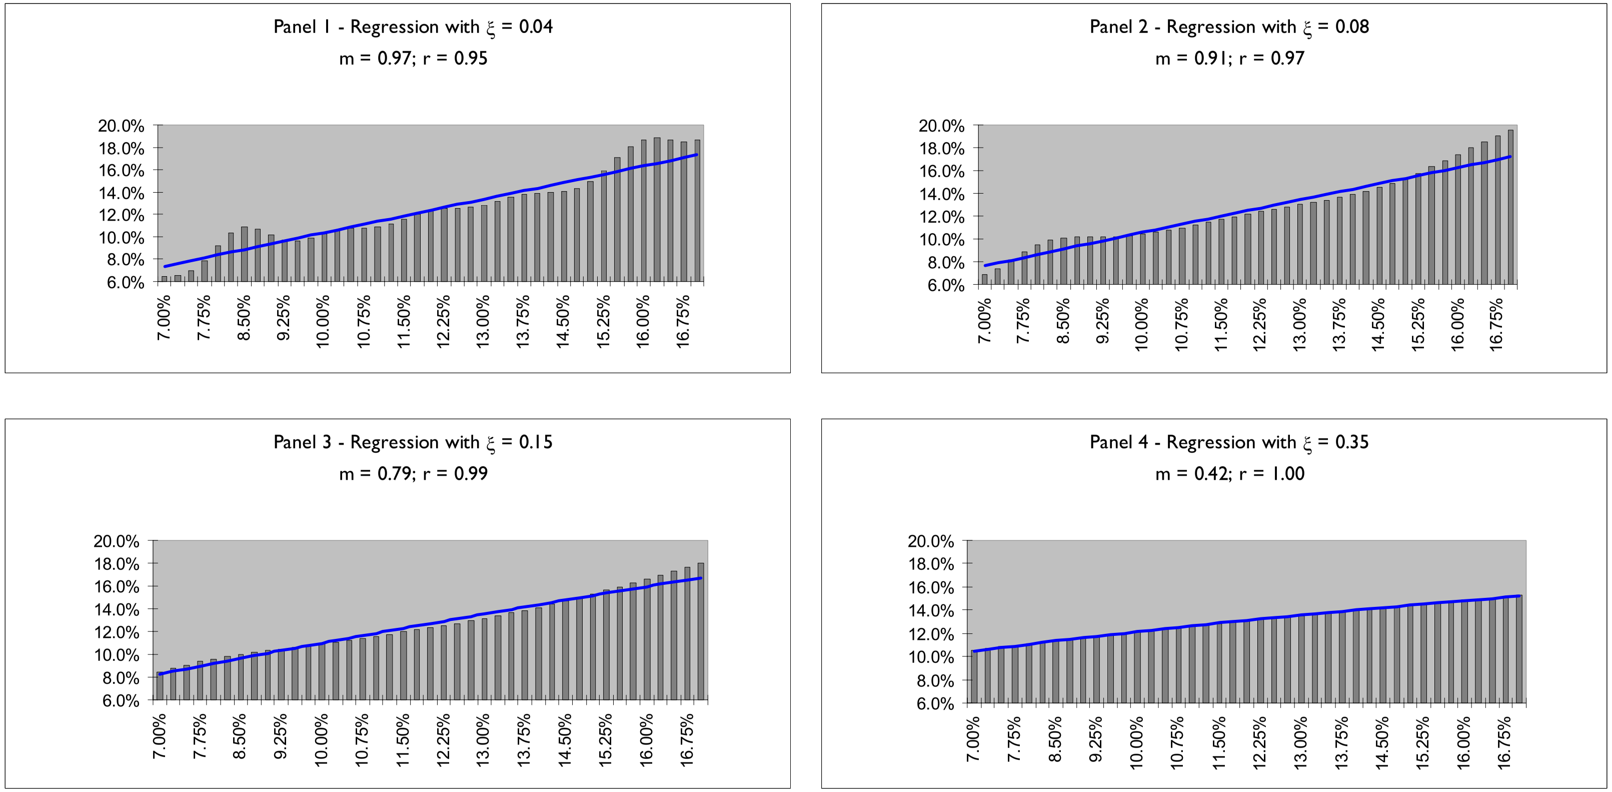

Figure 11 - Weighted regression line of kernel volatilities for the euro-dollar and the euro-yen on their respective historical volatilities with different ξ
(in logs of daily data)

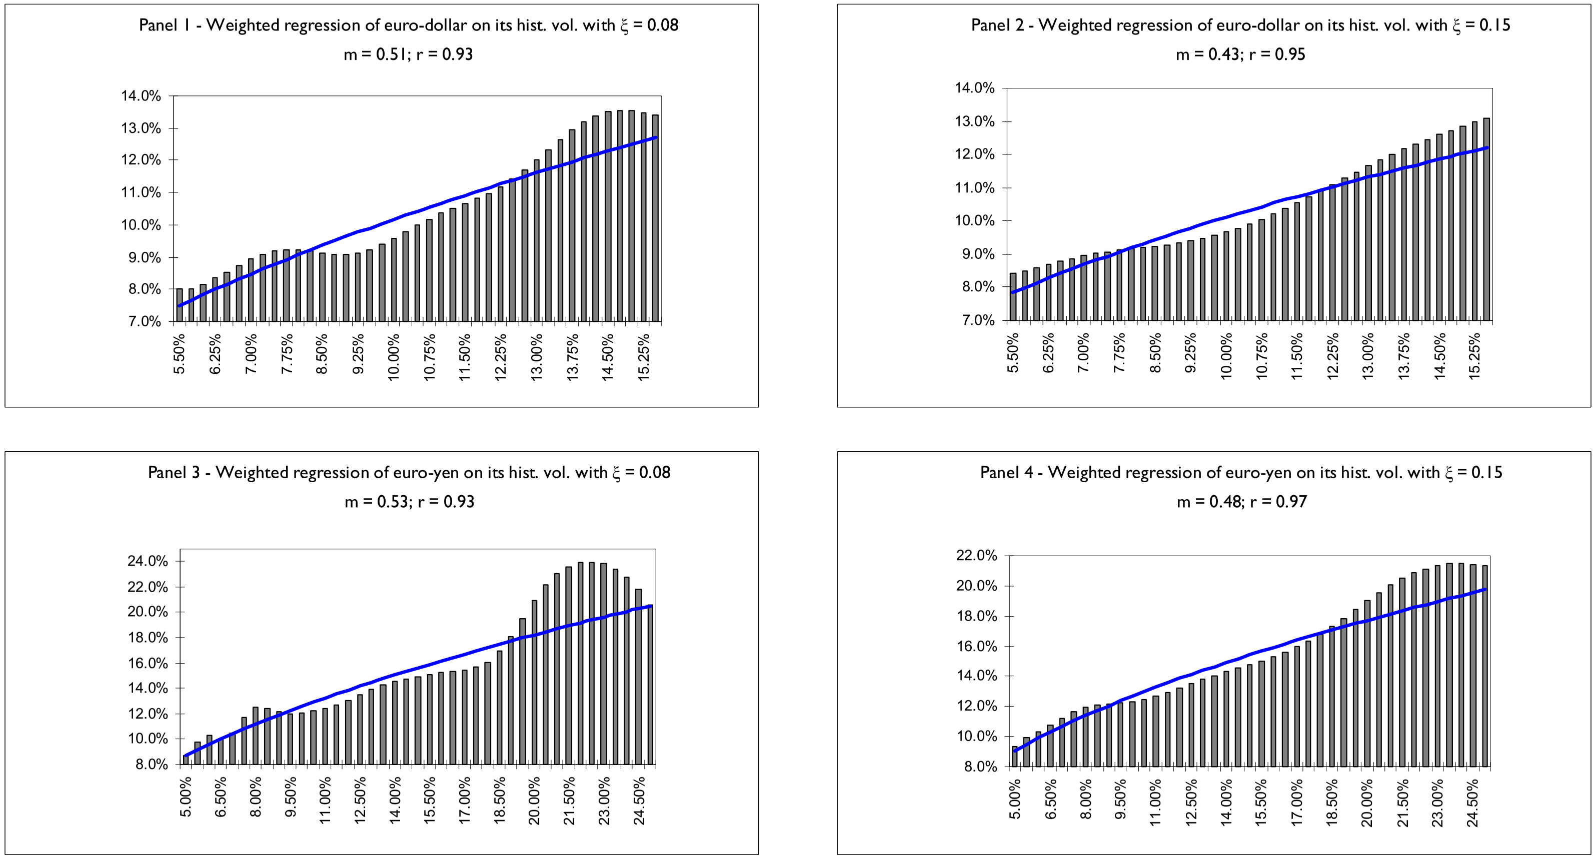

Figure 12 - Weighted regression line of kernel volatilities for the Swissie on GHI with different ξ
(in logs of daily data)

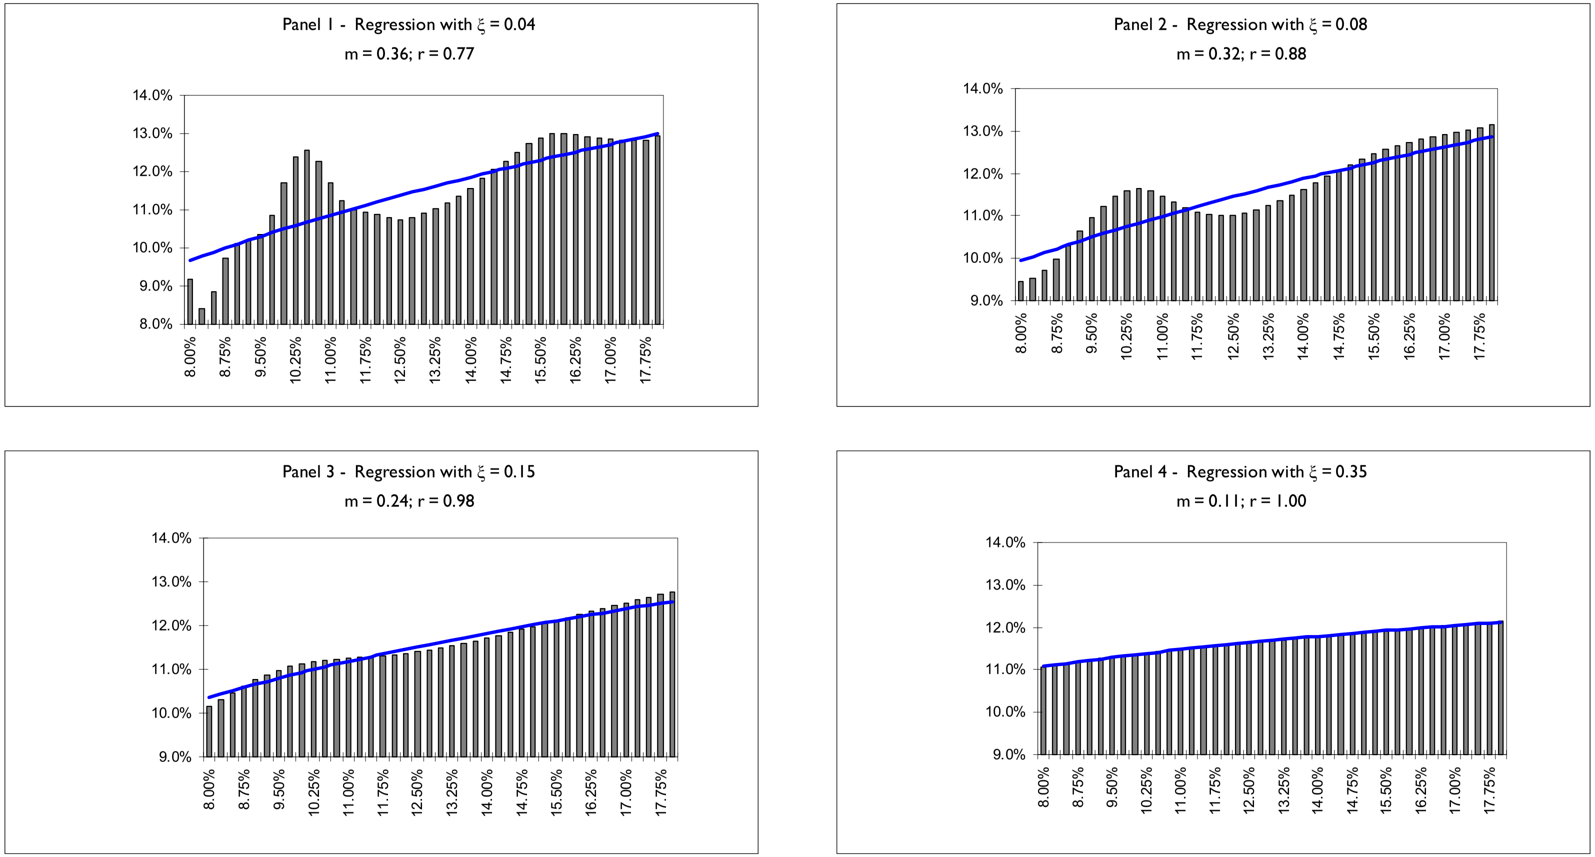

Figure 13 - Weighted regression line of kernel volatilities for the Cable on GHI with different ξ
(in logs of daily data)

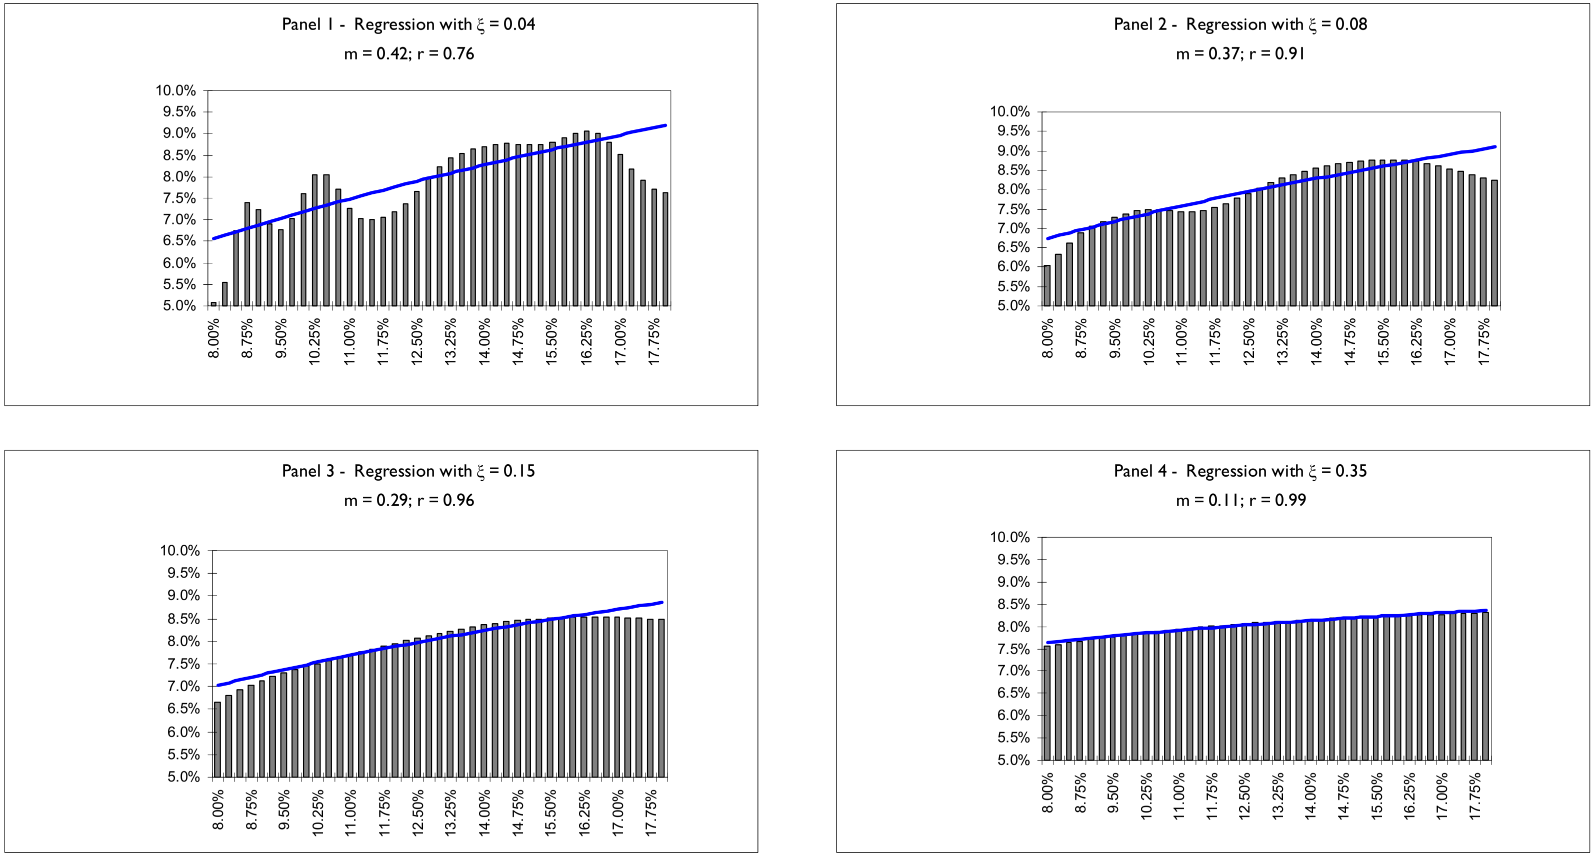

Figure 14 - Weighted regression line of kernel volatilities for the AUD and NZD on GHI with different ξ (in logs of daily data)

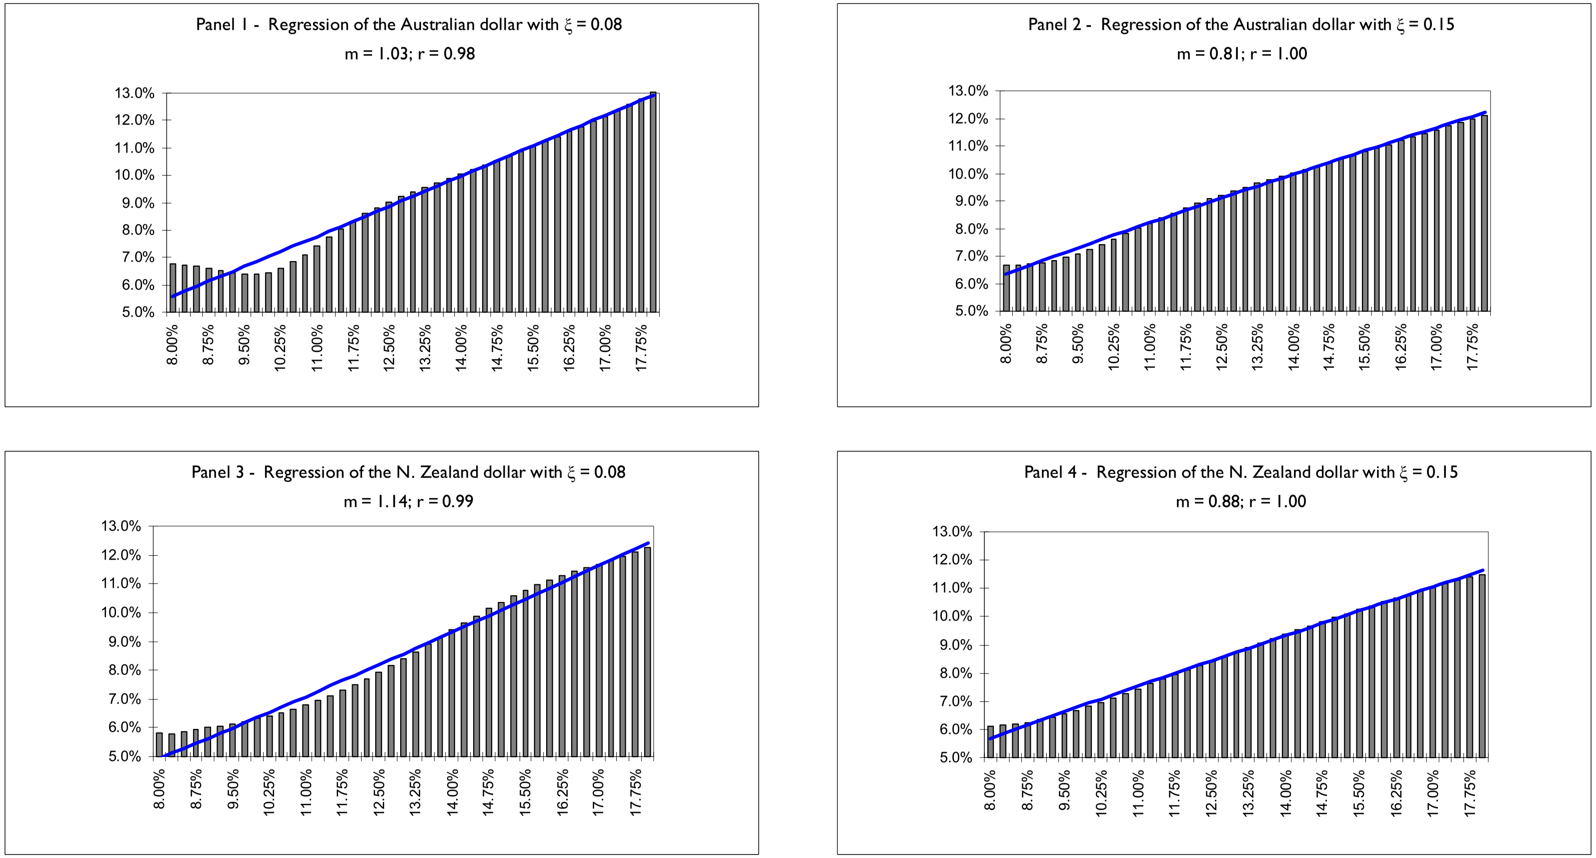

Figure 15 - Weighted regression line of kernel volatilities for the CAD with different ξ (in logs of daily data)

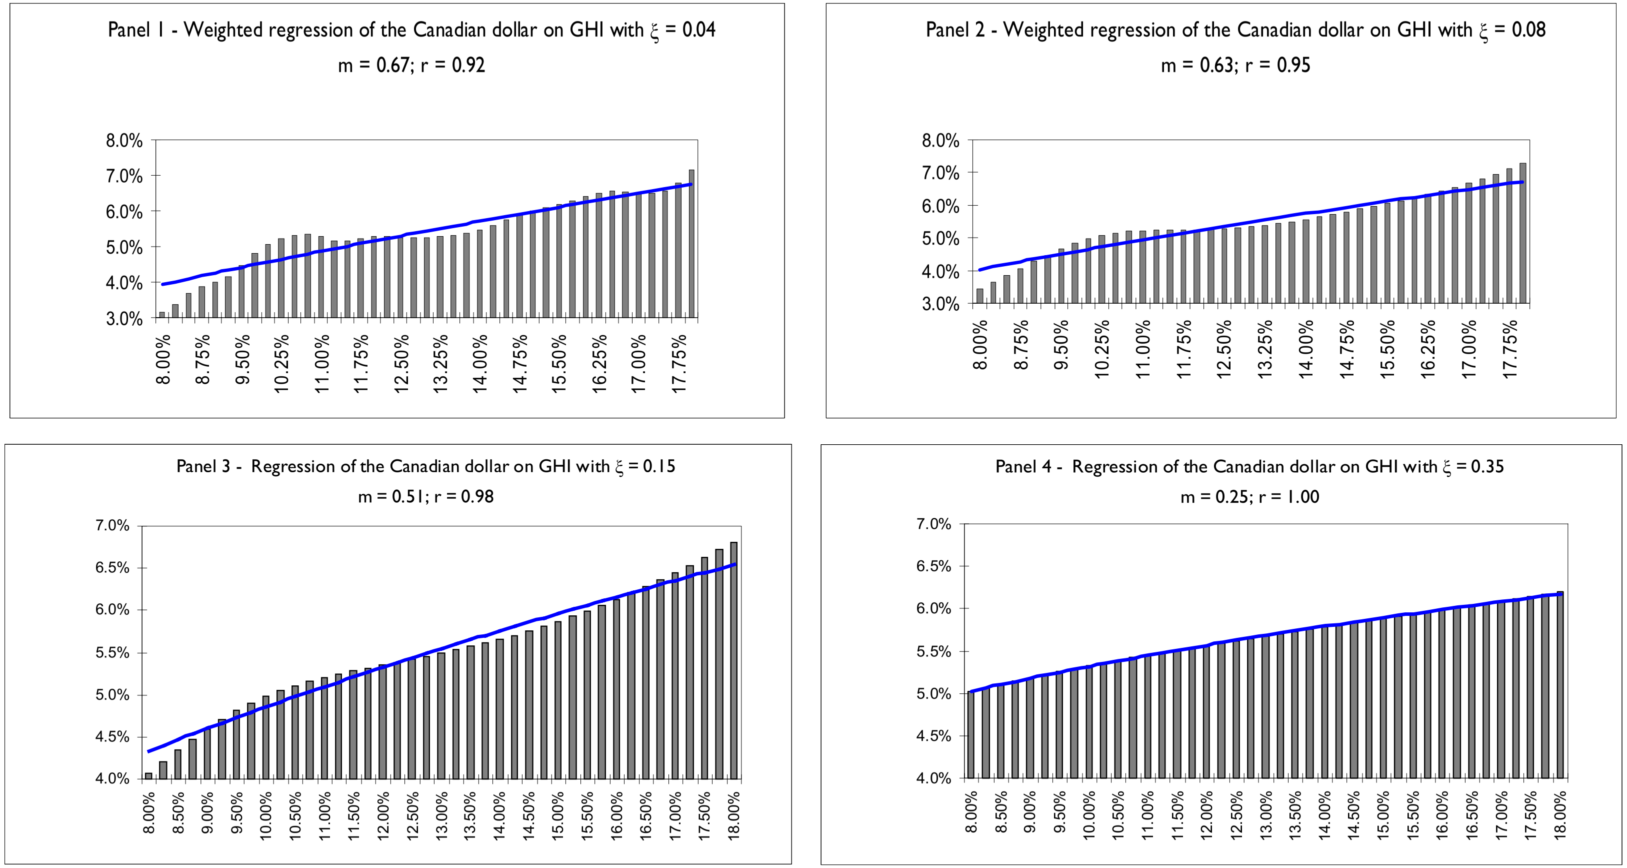

Figure 16 - Weighted regression line of kernel volatilities for DEMCHF and GBPDEM with different ξ (in logs of daily data)

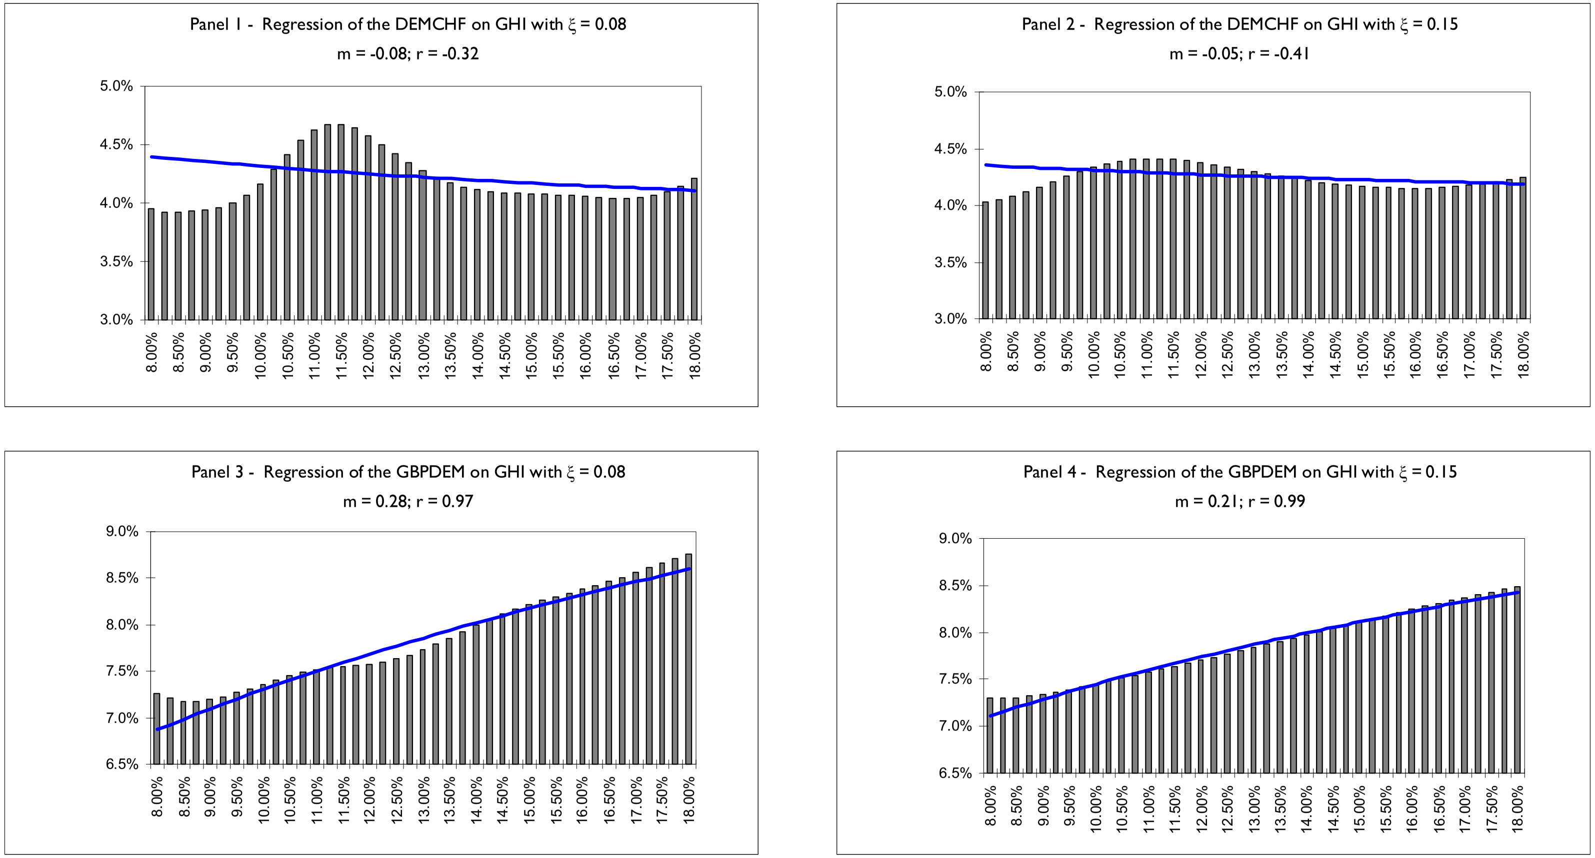

Figure 17 - Weighted regression line of kernel volatilities for the AUDJPY with different ξ
(in logs of daily data)

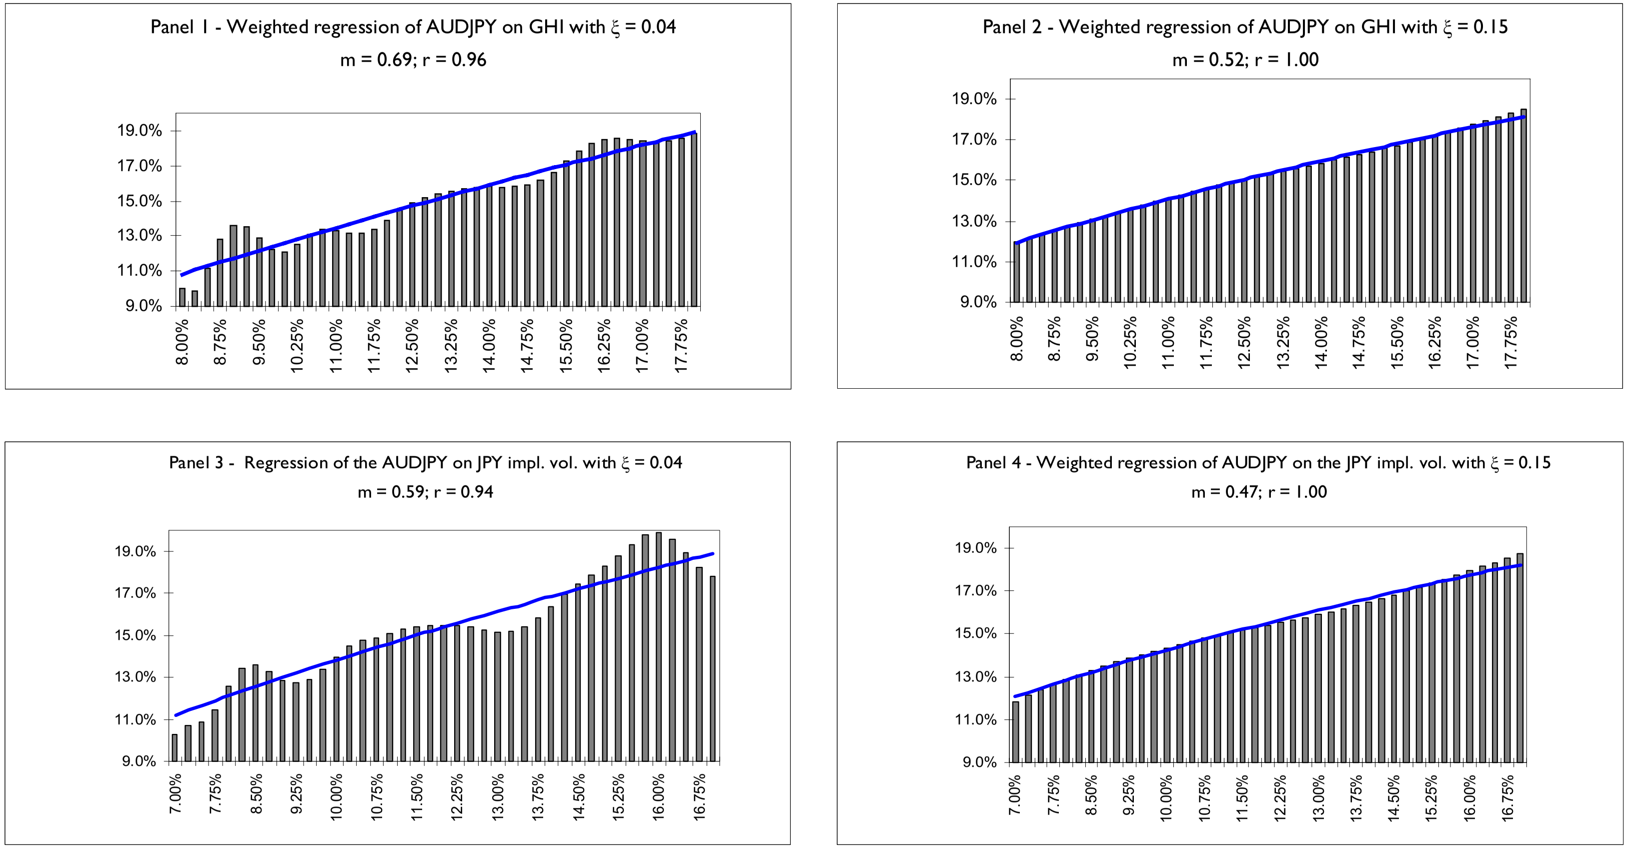

Figure 18 - Topographic view of the probability that the same results in the regressions could be generated by Brownian motion ($\xi = 0.05$)

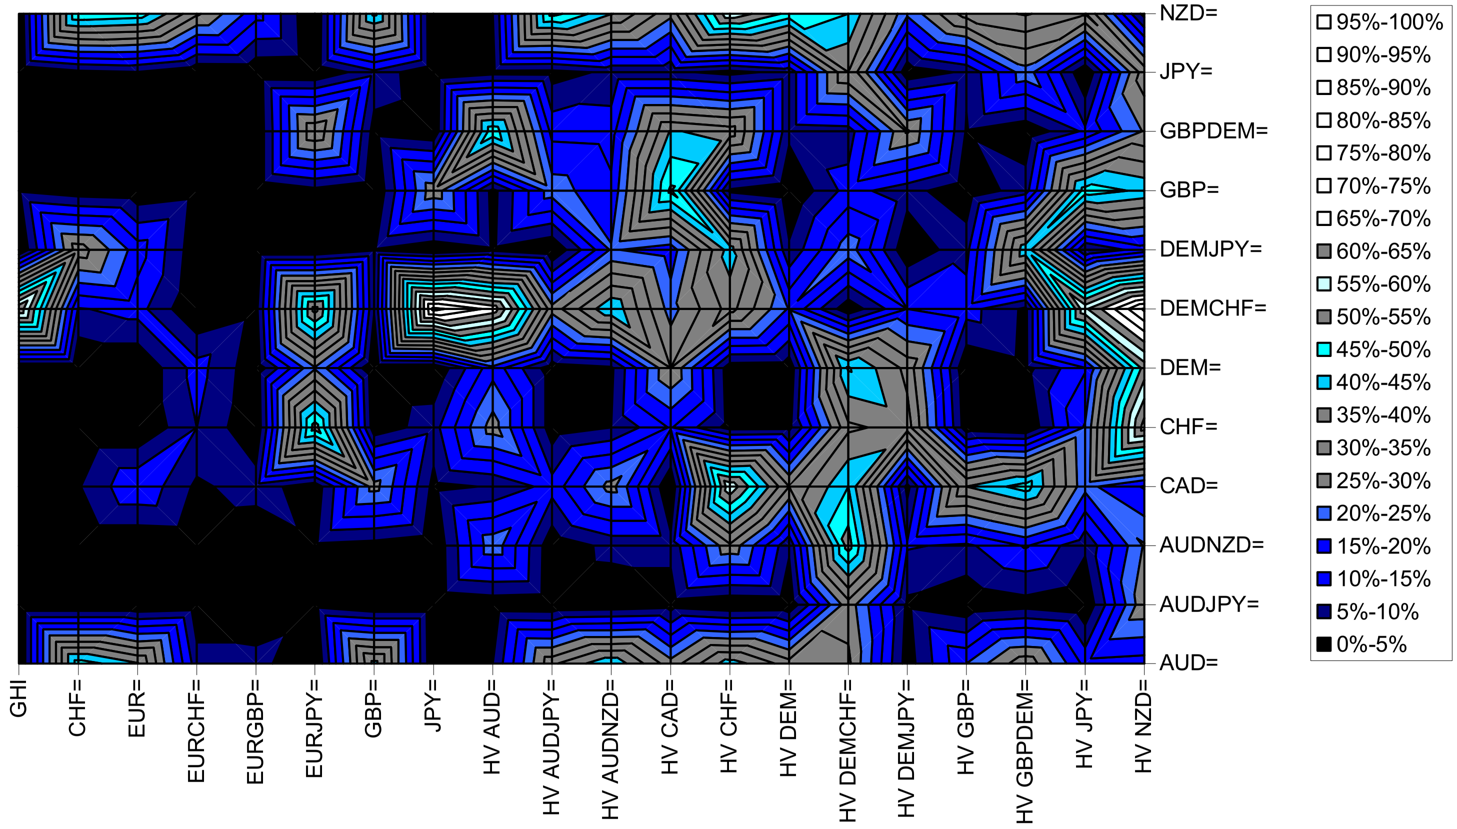

European Central Bank Working Paper Series

1 "A global hazard index for the world foreign exchange markets" by V. Brousseau and F. Scacciavillani, May 1999.

2 "What does the single monetary policy do? A SVAR benchmark for the European Central Bank" by C. Monticelli and O. Tristani, May 1999.

3 "Fiscal policy effectiveness and neutrality results in a non-Ricardian world" by C. Detken, May 1999.

4 "From the ERM to the euro: new evidence on economic and policy convergence among EU countries" by I. Angeloni and L. Dedola, May 1999.

5 "Core inflation: a review of some conceptual issues" by M. Wynne, May 1999.

6 "The demand for M3 in the euro area" by G. Coenen and J.-L. Vega, September 1999.

7 "A cross-country comparison of market structures in European banking" by O. de Bandt and E. P. Davis, September 1999.

8 "Inflation zone targeting" by A. Orphanides and V. Wieland, October 1999.

9 "Asymptotic confidence bands for the estimated autocovariance and autocorrelation functions of vector autoregressive models" by G. Coenen, January 2000.

10 "On the effectiveness of sterilized foreign exchange intervention" by R. Fatum, February 2000.

11 "Is the yield curve a useful information variable for the Eurosystem?" by J. M. Berk and P. van Bergeijk, February 2000.

12 "Indicator variables for optimal policy" by L. E. O. Svensson and M. Woodford, February 2000.

13 "Monetary policy with uncertain parameters" by U. Söderström, February 2000.

14 "Assessing nominal income rules for monetary policy with model and data uncertainty" by G. D. Rudebusch, February 2000.

15 "The quest for prosperity without inflation" by A. Orphanides, March 2000.

16 "Estimating the implied distribution of the future short term interest rate using the Longstaff-Schwartz model" by P. Hördahl, March 2000.

17 "Alternative measures of the NAIRU in the euro area: estimates and assessment" by S. Fabiani and R. Mestre, March 2000.

18 "House prices and the macroeconomy in Europe: Results from a structural VAR analysis" by M. Iacoviello, April 2000.

19 “The euro and international capital markets” by C. Detken and P. Hartmann, April 2000.

20 “Convergence of fiscal policies in the euro area” by O. De Bandt and F. P. Mongelli, May 2000.

21 “Firm size and monetary policy transmission: evidence from German business survey data” by M. Ehrmann, May 2000.

22 “Regulating access to international large value payment systems” by C. Holthausen and T. Rønde, June 2000.

23 “Escaping Nash inflation” by In-Koo Cho and T. J. Sargent, June 2000.

24 “What horizon for price stability” by F. Smets, July 2000.

25 “Caution and conservatism in the making of monetary policy” by P. Schellekens, July 2000.

26 “Which kind of transparency? On the need for clarity in monetary policy-making” by B. Winkler, August 2000.

27 “This is what the US leading indicators lead” by M. Camacho and G. Perez-Quiros, August 2000.

28 “Learning, uncertainty and central bank activism in an economy with strategic interactions” by M. Ellison and N. Valla, August 2000.

29 “The sources of unemployment fluctuations: an empirical application to the Italian case” by S. Fabiani, A. Locarno, G. Oneto and P. Sestito, September 2000.

30 “A small estimated euro area model with rational expectations and nominal rigidities” by G. Coenen and V. Wieland, September 2000.

31 “The disappearing tax base: Is foreign direct investment eroding corporate income taxes?” by R. Gropp and K. Kostial, September 2000.

32 “Can indeterminacy explain the short-run non-neutrality of money?” by F. De Fiore, September 2000.

33 “The information content of M3 for future inflation” by C. Trecroci and J. L. Vega, October 2000.

34 “Capital market development, corporate governance and the credibility of exchange rate pegs” by O. Castrén and T. Takalo, October 2000.

35 “Systemic risk: A survey” by O. De Bandt and P. Hartmann, November 2000.

36 “Measuring core inflation in the euro area” by C. Morana, November 2000.

37 “Business fixed investment: Evidence of a financial accelerator in Europe” by P. Vermeulen, November 2000.

38 “The optimal inflation tax when taxes are costly to collect” by F. De Fiore, November 2000.

39 “A money demand system for euro area M3” by C. Brand and N. Cassola, November 2000.

40 “Financial structure and the interest rate channel of ECB monetary policy” by B. Mojon, November 2000.

41 “Why adopt transparency? The publication of central bank forecasts” by P. M. Geraats, January 2001.

42 “An area-wide model (AWM) for the euro area” by G. Fagan, J. Henry and R. Mestre, January 2001.

43 “Sources of economic renewal: from the traditional firm to the knowledge firm” by D. R. Palenzuela, February 2001.

44 “The supply and demand for eurosystem deposits – The first 18 months” by U. Bindseil and F. Seitz, February 2001.

45 “Testing the Rank of the Hankel matrix: a statistical approach” by G. Camba-Mendez and G. Kapetanios, February 2001.

46 “A two-factor model of the German term structure of interest rates” by N. Cassola and J. B. Luís, February 2001.

47 “Deposit insurance and moral hazard: does the counterfactual matter?” by R. Gropp and J. Vesala, February 2001.

48 “Financial market integration in Europe: on the effects of EMU on stock markets” by M. Fratzscher, March 2001.

49 “Business cycle and monetary policy analysis in a structural sticky-price model of the euro area” by M. Casares, March 2001.

50 “Employment and productivity growth in service and manufacturing sectors in France, Germany and the US” by T. von Wachter, March 2001.

51 “The functional form of the demand for euro area M1” by L. Stracca, March 2001.

52 “Are the effects of monetary policy in the euro area greater in recessions than in booms?” by G. Peersman and F. Smets, March 2001.

53 “An evaluation of some measures of core inflation for the euro area” by J.-L. Vega and M. A. Wynne, April 2001.

54 “Assessment criteria for output gap estimates” by G. Camba-Méndez and D. R. Palenzuela, April 2001.

55 “Modelling the demand for loans to the private sector in the euro area” by A. Calza, G. Gartner and J. Sousa, April 2001.

56 “Stabilization policy in a two country model and the role of financial frictions” by E. Faia, April 2001.

57 “Model-based indicators of labour market rigidity” by S. Fabiani and D. Rodriguez-Palenzuela, April 2001.

58 “Business cycle asymmetries in stock returns: evidence from higher order moments and conditional densities” by G. Perez-Quiros and A. Timmermann, April 2001.

59 “Uncertain potential output: implications for monetary policy” by M. Ehrmann and F. Smets, April 2001.

60 “A multi-country trend indicator for euro area inflation: computation and properties” by E. Angelini, J. Henry and R. Mestre, April 2001.

61 “Diffusion index-based inflation forecasts for the euro area” by E. Angelini, J. Henry and R. Mestre, April 2001.

62 “Spectral based methods to identify common trends and common cycles” by G. C. Mendez and G. Kapetanios, April 2001.

63 “Does money lead inflation in the euro area?” by S. N. Altimari, May 2001.

64 “Exchange rate volatility and euro area imports” by R. Anderton and F. Skudelny, May 2001.

65 “A system approach for measuring the euro area NAIRU” by S. Fabiani and R. Mestre, May 2001.

66 “Can short-term foreign exchange volatility be predicted by the Global Hazard Index?” by V. Brousseau and F. Scacciavillani, June 2001.