Empirical Ross Recovery of International Market Expectations

Marie-Hélène Gagnon[†], Gabriel J. Power[‡], and Dominique Toupin[*]

This version: March 2018

The authors gratefully acknowledge financial support from SSHRC, FRQSC, la Chaire de recherche Industrielle-Alliance, and les Salles des Marchés FSA Jean-Turmel et Carmand-Normand. Any errors are our own. We thank seminar participants at the Norwegian University of Science & Technology (Industrial Economics) and conference participants at ASAC 2017.

[†] Gagnon is Associate Professor of Finance, Department of Finance, Insurance and Real Estate, FSA ULaval, and Research Fellow, CRREP, IHEIQ and CEPCI, P: (418) 656-2131, # 4742. F: (418) 656-2624. E: Marie-Helene.Gagnon@fsa.ulaval.ca

[‡] Power is Associate Professor of Finance, and Research Fellow, CRREP and CRIB, Université Laval. P: (418) 656-2131, # 4619. F: (418) 656-2624. E: Gabriel.Power@fsa.ulaval.ca

[*] Toupin is a PhD candidate in Finance, Université Laval. E: dominique.toupin.1@ulaval.ca

Empirical Ross Recovery of International Market Expectations

Abstract

This paper applies Ross's (2015) Recovery Theorem to international stock index options data to obtain the market's true expectations as well as the implied risk preferences. First, we document a reliable methodology to enable an empirical implementation of Ross's results. Then, we assess whether the quantities recovered are the true market expectations. Second, we show that the volatility of the recovered distribution is a significant and unbiased predictor of future realized volatility and that it improves the performance of volatility forecasts made using risk-neutral volatility (e.g. VIX). Third, we find evidence that recovered volatility and risk preferences can better explain index returns compared to the risk-neutral distribution. Our paper thus confirms the empirical validity, relevance and usefulness of the Recovery Theorem as way to obtain the true expectations and risk preferences of the market.

1 – Introduction

The price of an asset reflects the market's expectations about its future returns and the market's preference for risk. Prior to Ross's (2015) claims concerning his recovery theorem, we could not directly observe those expectations and risk preferences. In particular, constructing a historical distribution from past realized returns has shown to be unreliable, especially concerning the higher distributional moments (e.g. Conrad, Dittmar and Ghysels, 2013). Generally speaking, using past returns to construct a distribution is an approach that is backward-looking, while options data reveal the market's forward-looking views, as each option is linked to a specific horizon. The value of accessing the true distribution of expected returns cannot be overstated, whether it is for asset pricing, portfolio allocation, capital budgeting, or risk management.

It is well known that using option price data, one can get closer to the true market expectations by computing the risk-neutral distribution (RND) (e.g., Breeden and Litzenberger 1978). This RND corresponds to the expected returns distribution in a world where investors are risk-neutral, or equivalently it is the distribution of Arrow-Debreu state prices.

This does not mean that the RND does not contain useful information. The volatility of this distribution, called risk-neutral volatility (RNV) and, in the case of the US S&P500 index measured by the VIX index, has been shown to be a very useful tool. It has, for example, been shown to be a better predictor of future volatility than past returns (e.g. Christensen and Prabhala 1998; Szakmary et al. 2003).

Ross (2015) argues that one can determine the market's true expectations using only options price data and weaker assumptions than those used in the previous literature. Specifically, his recovery theorem proposes a way to identify separately the true expectations and the risk preferences that are implied in option prices. The intuition is that the recovery theorem exploits RND variation across return horizons (e.g. one month ahead, two months ahead, etc.). If the market has an idea about what future returns will look like, then those true expectations should be the best estimates available. This means that if the recovery theorem is valid, the true expectations that are computed using this approach might contain new information that was not present in the RND. More specifically, return and volatility forecasts, for example, could be improved when the RND is decomposed, using the recovery theorem, into two components namely the Ross-recovered distribution and the true, forward-looking pricing kernel (stochastic discount factor). To verify the validity of the theorem we ask: Does the Ross-recovered distribution contain new information compared to the RND? If so, the recovered market expectations should be closer to the true expectations of the market. It would also suggest that the recovered risk preferences accurately represent those of the market and it would give researchers a way to directly study those risk preferences.

Related literature

Our study is not the only one to investigate empirically the performance of Ross-recovered distributions, but it differs from other, related papers in its methodology, objectives, and research questions. First, in the theoretical literature, Carr and Yu (2012) show that Ross recovery can be extended to the continuous state-space setting for bounded diffusions. Walden (2017) extends the result to the case of unbounded diffusions. Relatedly, Qin and Linetsky (2016) show how to establish Ross recovery in the more general state-space of Borel right processes. Finally, Borovicka, Hansen and Scheinkman (2016) question whether Ross's theorem (2015) recovers uniquely the true physical distribution. Dubynskiy and Goldstein (2013) are also critical of Ross's bounding restrictions on state vector dynamics.

The empirical literature is smaller. Martin and Ross (2013) use the recovery theorem to study properties of the "long bond", namely the distant end of the yield curve. Bakshi, Chabi-Yo and Gao

(2018) test the recovery theorem using 30-year Treasury bond futures. The paper that is closest to ours is probably Jackwerth and Menner (2017), who investigate the performance of the Ross distribution to predict future realized returns for the US S&P 500 index. However, aside from the broad objective of empirically testing the performance of Ross distributions, our paper differs considerably from theirs.

Our study makes the following contributions to the literature. First, we propose a novel empirical design to implement Ross recovery, which yields more economically plausible outputs than what is obtained by naively applying Ross's theorem to the data. Second, we test empirically the performance of two distinct components recovered from Ross's theorem, namely the volatility of the physical distribution and its difference relative to risk-neutral volatility, in order to predict future realized volatility. Third, we look at how these components can better explain the returns on a stock index than what was already known from the risk-neutral distribution. Fourth and last, our study is not limited to the US S&P500 index but rather considers five major international equity index markets.

Briefly, our empirical analysis shows the recovered expected distribution is a significant predictor of future realized volatility and improves the performance of the volatility forecasts. It is unbiased, in contrast to the commonly used risk-neutral volatility (e.g. the VIX index). Second, the recovered volatility and risk preferences can better explain the returns of an index compared to the usual risk-neutral distribution. Our paper thus confirms the empirical validity, relevance and usefulness of the Recovery Theorem as way to obtain the true expectations and risk preferences of international markets.

2 – How to recover expectations and risk preferences

2.1 – Ross's "Recovery Theorem"

The stochastic discount factor (also called pricing kernel) is defined as a random variable m that relates the price p of an asset today to its final value (random variable x) in each possible future state of nature.

$$p = E(\tilde{m}\tilde{x}) \tag{1}$$

If the “law of one price” holds, the stochastic discount factor exists and if there is no arbitrage, it is strictly positive. This relation means that knowing the stochastic discount factor and the probability of getting to each future state of nature, any asset can be priced from its final values in each state of nature. This stochastic discount factor can be interpreted as the desirability of one dollar in each future state of nature. For example, a typical risk-averse investor might prefer having an additional dollar during an eventual recession instead of an additional dollar in a future state of nature where there is an economic boom. This lowers the variability of his expected future consumption and maximizes his utility today. This means that, by definition, the stochastic discount factor represents the aggregate preferences of all investors without imposing a particular structure on those preferences.

An Arrow-Debreu contingent claim is an asset with a final value of 1$ in state of nature j and 0$ in every other possible state. The price of this asset is commonly referred to as the price of state j. If the current state of nature is i, this state price is given by the relation:

$$p_{i,j} = E(\tilde{m}\tilde{x}) = \pi_{i,j} m_{i,j} \times 1 \tag{2}$$

where $p_{i,j}$ is the probability of going from state i to state j and $m_{i,j}$ is the stochastic discount factor for this state of nature. If this asset is traded, its price can be directly observed, but $m_{i,j}$ and $p_{i,j}$ cannot be separately identified. In a risk-neutral world, things are a bit simpler. Risk-neutral investors do not care about risk; they only take into account the time value of money. The previous pricing equation simplifies to:

$$p_{i,j} = E^*(\frac{1}{R_f}\tilde{x}) = \pi^*_{i,j}\frac{1}{R_f} \qquad (3)$$

where R_f is the risk-free interest rate and the asterisk (*) represents probabilities in a risk-neutral world. This time, the risk-neutral probability can be recovered directly from the observed price of the asset because the risk-free rate is known. This is also true of more complex assets like options, which can be represented as a sum of Arrow-Debreu contingent claims. This means that knowing any two of the following elements is enough to know the value of the third one:

- The physical probability distribution of attaining each state of nature. (The $\pi_{i,j}$ of equation (2))
- The stochastic discount factor of each state of nature.
- The risk-neutral probability distribution of attaining each state of nature OR the state price of each state.

State prices can be recovered directly from option prices (Breeden and Litzenberger 1978), leaving two unknown elements. The physical probability distribution and the stochastic discount factor are both of interest to finance researchers, but until now they have had to make assumptions about one of them to compute the other one. Jackwerth (2000), Aït-Sahalia and Lo (2000), and Rosenberg and Engle (2002) assume that the physical distribution of expected returns is equal to the historical returns distribution in order to study the stochastic discount factor. On the other hand, Bliss and Panigirtzoglou (2004) assume a specific form for the stochastic discount factor (time-separable power utility) which allows them to test the predictive power of options data for future realized returns. This means that the validity of the results from this literature relies on the correctness of the assumptions made about one of the three quantities previously discussed.

Ross (2015)'s "Recovery Theorem" aims to solve this problem by letting us recover all those elements directly from option prices. Using this theorem, the true expectations of the market and the pricing kernel representing its risk preferences can be computed by only looking at option prices. This approach has many potential applications. For example, it is well known that it is difficult to recover a physical distribution of asset returns from historical data (e.g., Conrad, Dittmar and Ghysels, 2013). Ross's approach is based on the idea of exploiting Arrow-Debreu state prices obtained from options data, and starts from a matrix P for which its element p_{ij} is the state price of state j when we are currently in state i. Those state prices are computed from assets that are contingent claims on the same Markovian state variable X. For example, X could be a stock index and P would be computed from the prices of options on this index at a particular time. In order to have only one possible matrix P that represents the evolution of X, Ross (2015) makes a first assumption:

Assumption 1

The process of X is time-homogenous on a finite state space.

This means that the matrix P represents the state prices from time 0 to t as well as those from time t to t+1, t+1 to t+2, … Another assumption is needed in order separately identify the physical probabilities and the stochastic discount factor:

Assumption 2

The stochastic discount factor is transition-independent and is of the following form:

$$m_{i,j} = \delta \frac{d_j}{d_i} \tag{4}$$

where δ is a positive constant that represents the market's average discount rate and d(·) is a function that depend only on the corresponding state of nature. This means that the stochastic discount factor is not dependent on intermediate states that are reached before attaining the final state j. A utility function for the representative agent that is intertemporally additive is an example of transition-independent stochastic discount factor. Epstein-Zin recursive preferences also lead to a transition-independent pricing kernel according to Ross (2015).

With those assumptions, equation (2) can be reformulated :

$$p_{i,j} = m_{i,j} p_{i,j} = d \frac{d_j}{d_i} p_{i,j} \qquad (5)$$

If the state space that can be reached by the Markovian variable X has a finite number of elements n, the state prices matrix P is nxn. The (physical) probability transition matrix F, with its elements $p_{i,j}$, is also nxn. Equation (5) can be rewritten:

$$P = dD^{-1} FD \qquad (6)$$

where $D_{n \times n}$ is a matrix with the elements d_i on the diagonal and zeros elsewhere. Knowing that F is a probability transition matrix, its rows must add up to 1 which means that $F\vec{1} = \vec{1}$ where $\vec{1}$ is a vector of ones. The previous equation ca now be written like this:

$$PD^{-1}\vec{1} = \delta D^{-1}\vec{1} \qquad (7)$$

Finally, we define the vector $z = D^{-1}\vec{1}$ in order to obtain the following familiar form:

$$P_{n\times n} z_{n\times 1} = \delta z_{n\times 1} \qquad (8)$$

This is a classical characteristic root problem where we are looking for the eigenvectors (z) and eigenvalues (δ) corresponding to a square matrix (P). Without additional assumptions, the solution would be expressed as complex numbers having a real and imaginary part. Since this does not make much financial sense, additional steps are needed to insure a solution without an imaginary part. The Perron-Frobenius theorem can help us here. It states that for a characteristic root problem like (8), if P is non-negative and irreducible, there exists only one strictly positive eigenvector and its corresponding eigenvalue is positive and real. This eigenvalue is the largest absolute in the possible eigenvalues. This means that if P is positive and irreducible, (8) has a unique positive solution. Since the elements of P are state prices, P is positive under the no arbitrage condition. To insure that it is also irreducible, we need one last assumption:

Assumption 3

The Markovian variable X can reach any state j from state i in a finite number of steps.

Now that (8) has a unique positive solution, we can compute the stochastic discount factor D, the discount rate δ and the physical probability transition matrix F from only the matrix of state prices P. This means that all those variables can be known only by looking at option prices for different strikes and maturities on a particular asset.

2.2 – From options to the matrix of state prices P

From Breeden and Litzenberger (1978), we know that state prices at a certain moment in time can be computed from the prices of call options C of a certain maturity T and for a continuum of strike prices K:

$$s(K,T) = \frac{\partial^2 C(K,T)}{\partial K^2} \tag{9}$$

where $s(K,T)$ is the state price corresponding to a state of nature where the value of the asset is K at time T. If we consider a finite state space of m maturities and n strike prices, those state prices can be grouped in the matrix $S_{n\,x\,m}$. This is an "implied state prices surface", a transformation of the commonly used "implied volatility surface" of options on an asset. But this is not yet the state prices matrix $P_{n\,x\,n}$ we need to use the "Recovery Theorem". The difference is that $S_{n\,x\,m}$ is composed of the prices (valued today) of the contingent claims that give 1$ if the underlying asset is equal to K at time T (for different values of (K,T)), while $P_{n\,x\,n}$ are the prices (valued in state i) of contingent claims that give 1$ if the final value of the asset is K (state j). In P, the maturity of those contingent claims is always the same. The matrix P could be computed for any choice of maturity, but since we have assumed that the underlying process is time-homogenous, this choice is arbitrary. This is analogous to the situation where the interest rate is constant and you can choose to use either a monthly interest rate or an effective annual rate and get the same result.

There is one row of P that is already known. For example, if we consider that P represents monthly transitions from state to state, the row of P where (i= today's state) and (j= every strike) is equal to the column of S where (K=every strike) and (T= 1 month). The remaining rows of P correspond to the state prices we would observe if today we were in a different state than *i*. Since we assumed that the underlying process is time-homogenous, *P* and *S* are related by:

$$S_{:,t}^{\mathsf{T}} P = S_{:,t+1}^{\mathsf{T}} \tag{10}$$

which means that if m>n, there are enough equations to find the n^2 unknown state prices.

It is important to note that by assuming that the future states of nature possible are completely defined by the strike prices of an asset, we have already made an empirical choice that is not inherent to the "Recovery Theorem". In theory, a state of nature could be defined by any number of variables. For example, state *i* could be defined as a state where the S&P500 is at 2200, market volatility is at 20%, the US economy is in a recession. A small change in any of those variables would mean that we are now in a different state of nature. Empirically, we are limited by the kind of contingent claims that are traded on the market. This mean that the pricing kernel recovered using this theorem will not be able to price any asset like its theoretical counterpart, but only those whose value depends only on the kind of states of nature that was chosen. This does not dampen in any way the relevance of this recovered pricing kernel but it should be kept in mind if we are to compare the results obtained with a theoretical economic model that assumed that the complete pricing kernel is known. For example, the recovered pricing kernel could be different from a theoretical model where the volatility corresponding to a state of nature is an important variable.

2.3 – Empirical considerations

2.3.1 – The implied state prices surface S

A standard method from the literature building on Birru and Figlewski (2012) to get a risk-neutral probability density function for a given maturity from option prices can be applied here. From equation

(3) we know that the state prices for a given maturity are directly related to the risk-neutral probabilities by the risk-free rate. This means that a method similar to Birru and Figlewksi (2012) can be repeated for all the *m* maturities needed and in order to get the implied state price surface *S*. If we consider *n* discrete strike prices, this gives us the S_{nxm} matrix of implied state prices.

2.3.2 – The transition state prices matrix P

According to equation (10), the state prices surface *S* has to be split in two overlapping submatrices in order to recover *P*. The lag time τ between those submatrices correspond to the horizon for the resulting transition matrix of state prices *P*. If we denote those submatrices by $A^{\mathrm{T}} = S_{:,[1:m-\tau]}$ and $B^{\mathrm{T}} = S_{:,[1+\tau:m]}$, finding the solution to equation (10) can be written as:

$$\min_{P \geq 0} \left\| AP - B \right\|^2 \tag{11}$$

This is equivalent to *n* separate least-squares problems:

$$\min_{p_j \geq 0} \left\| Ap_j - b_j \right\|^2, \quad j = 1, 2, \ldots, n \tag{12}$$

where p_j and b_j are the j^th columns of *P* and *B* respectively. There are algorithms that can easily solve this type of problem, but unfortunately the results might not be what we expect. Indeed, as noted by Audrino, Huitema and Ludwig (2015), the matrix *A* appears to be ill-conditioned, meaning that the solution *P* is highly sensitive to small perturbations of *A*. We can see more clearly the implications of this by looking at an example of P solved from (12). Graphically, this “unstable” solution for *P* looks like Figure 1.

As previously discussed, we already know what the values of one of the row of P should be. The row of P corresponding to today's asset price is the column of S at maturity τ. This would mean that in Figure 1, the row at a moneyness of 1 should look like Figure 2 but instead this row of P obtained from the least-square solution is quite different as we see in Figure 3.

Since the state prices transition matrix P is a slight transformation of the risk-neutral probability transition matrix, we would expect Figure 1 to show mainly a ridge following the diagonal. This would mean that for each possible beginning state of nature i, the most probable state at time τ is a state that is close to i. For example, if we begin in a state of nature where a stock index is at 2000, the most probable state at time τ would be at a level near 2000. This pattern is visible in parts of Figure 1, particularly on the diagonal from x,y coordinates (1, 1) to (1.8, 1.8). Nevertheless, there are many data points that are very far from this, especially on the left-side of Figure 1.

We now consider additional restrictions on the shape of P so that it would be closer to a typical transition matrix. Similarly to the methodology used in the brief empirical test of the recovery theorem in Ross (2015), unimodality is imposed for p_j of equation (12). After some experimentation, we found that this restriction is the simplest way to insure that the resulting rows of P "look like probability distributions". Ideally we wouldn't have to add this restriction, but it appears to be necessary in order to obtain meaningful results. In solving equation (12), we also impose the values for the known row of P. Finally, we use a relatively small number of interpolated strike prices. At first, it would be tempting to use a fine grid for the moneyness (strike prices) as it would result in a finer probability distribution for the true expectations of the market. This is problematic when we solve the problem numerically as it appears to exacerbate the sensibility of the solution to the input data. By using only 21 moneyness in P, the solution appears to be more in line with what we expect in a transition matrix. It is also faster to compute. This is not a very costly choice as we can still obtain a finer distribution for the true expectations of the market with some manipulations.[1]

3 – Data

Equity index excess returns are computed including dividends using data from Bloomberg. Country risk-free rates index are obtained from the FRED website of the Federal Reserve Bank of St. Louis. Realized variances computed as the sum of the 5-minute realized variance and squared overnight returns of the past month for each index are obtained from the Oxford-Man Institute's realized library's website (Gerd, Lunde, Shephard and Sheppard, 2009).

Index option data are obtained from OptionMetrics Ivy DB U.S. and Europe. The selected indexes are those with the most options available among the markets covered by this database. They are the S&P500 index for the United States, the DAX index for Germany, the SMI index for Switzerland, the CAC index for France, and the FTSE index for the United Kingdom. These markets are significant for the Eurozone, and have active index options markets. Strike prices, maturities, and implied volatilities are extracted for all available options on the selected indexes. Descriptive statistics for the raw option data is presented in table 1.

4 – Empirical tests

In this section we present our results for an investigation of the performance of measures on the Ross-recovered distribution for each of the five international equity index markets. Our objectives are twofold: (i) to assess whether a model forecasting future realized index volatility incorporating the information from the recovered distribution is improved compared to a model using only the risk-neutral distribution and (ii) to provide evidence that the market expectations and risk preferences revealed by

[1] Once we have the pricing kernel on the coarse grid of moneyness, we can fit it to a curve (e.g. a spline) and transform the fine RND with it to obtain a fine pdf of the market true expectations.

Ross recovery can help better explains the returns of an index. As we think that those objectives are met, this gives this empirical application of the Recovery Theorem the legitimacy needed in order to justify further empirical studies on what it can reveal about investors' expectations and preferences.

4.1 – Volatility of the physical distribution of expected returns

The "Recovery Theorem" is applied to option data for several international indexes. The implied state prices surface S is discretized in $m=100$ evenly spaced maturities from 1 month to two years. On the strike-axis it is discretized from a moneyness of 0.6 to 1.5 ($n=21$). Moneyness is the strike relative to the level of the index that day (*mn = strike / level of index*). The implied state prices matrix S is then split in two submatrices with a lag of $\tau=1$ month. This means that the recovered matrix of physical transition probabilities F represents the expectations for the 1-month-ahead returns. The relevant physical probabilities are the row of F that corresponds to the current state we are in (moneyness=1). The volatility of this physical probability distribution is computed and annualized for ease of comparison. This is also done for the risk-neutral distribution that was used to obtain the state prices. Those weekly time series are showed in Figures 4a to 4e. Descriptive statistics for all the explanatory variables considered are in Table 2.

4.2 – Volatility prediction

In this first empirical test, we predict realized volatility over the next h=21 business days (1 month). The explanatory variables at time t are the risk-neutral volatility (RNV), the volatility of the physical distribution of expected returns obtained from the Recovery Theorem (REV) and the difference between the two.

$$\text{RealizedVol}_{t,\ t+h} = \beta X_t + e_t \qquad (13)$$

In Table 3, we observe that for all indexes, both RNV and REV are significant predictors of future realized volatility. REV does not generally do a better job than the risk-neutral volatility for predicting future volatility when we use as a criterion the adjusted R^2 , which are usually smaller (except for the SMI). This might be explained by the additional noise that was introduced by applying the Recovery Theorem to the risk-neutral distribution. This additional noise would also explain the generally higher standard errors for the slope coefficients of REV compared to those of RNV for all indexes.

From the literature on volatility prediction, we know that risk-neutral volatility is a good predictor of future realized volatility, but that it is biased downward (e.g. Szakmary et al., 2003). We also test for the null hypothesis of β_{vol}=1 in order to investigate whether this bias remains when using the true expected volatility REV. We see that for all the studied indexes, except for the DAX, RNV is a biased predictor of future realized volatility, as its slope coefficient is significantly different than 1. This bias is downward, as has been previously observed in the literature. In contrast, the slope coefficient for REV the volatility of the true expectations of the market obtained with the Recovery Theorem is never significantly different than 1. This implies there is no observable bias, and thus it is a good indication that we have recovered the true expectations of the market.

In the third regression, we try to ascertain the contribution of the new information that was obtained by applying the Recovery Theorem. The information we already had is risk-neutral volatility, and the new information is measured by the difference between the two volatilities (RNV-REV). We note that for some indexes (S&P500, CAC), this additional information significantly improves the prediction of future realized volatility. This additional information also leads to higher adjusted R^2 compared to the first regression (what we already knew) in all indexes. The improvement in adjusted R^2 goes from 0.02% to 1.29%.

Those results imply that the volatility obtained by applying the Recovery Theorem is not only unbiased, but it also contains information that can help us better predict future realized volatility.

4.3 – Explaining returns using changes in expected volatility

In this second empirical test, we examine whether the recovered expectations can help us better explain the changes in the price of an index. The goal remains to find evidence that we have indeed recovered the true expectations of the market by applying the Recovery Theorem.

More specifically, we try to explain the weekly returns of an index using the changes in the measures of volatility during that period. Results are presented in Table 4.

$$\text{IndexReturn}_{t,\ t+1} = \beta^*(X_{t+1} - X_t) + e_t \tag{14}$$

In Regression 1, we observe that RNV always significantly explains index returns with adjusted R^2 of 0.208 to 0.605. The risk-neutral distribution contains information about both the true expectations of the market and its risk preferences. Since both quantities should affect the present value of a series of future cash flows, this is not a surprising result. The question remains however, whether we indeed extract the true expectations and risk preferences from the risk-neutral distribution by applying the Recovery Theorem. In this case, the risk preferences of the market are proxied by the difference between RNV and REV. All else being equal, when the difference between RNV and REV is greater, investors are farther away from risk-neutrality (e.g. they are more risk averse). In regressions 2 and 3 we see that those extracted components individually almost always significantly explain index returns. The adjusted R^2 in those regressions are smaller than for the RNV, which is plausible as the RNV contains information about both the true expectations (REV) and the risk preferences of the market. In regression 6, we see that if we define a new variable proxying for risk preferences ([RNV-REV]/RNV), the variable is still significant for most indexes.

In regression 5, we verify if the two recovered components perform better in explaining returns than does the RNV. We first observe that for most indexes, these components are both significant. The adjusted R^2 of regression 5 compared to regression 1 is higher for all indexes except the FTSE and CAC. For those indexes where the performance is improved, this appears to be explained by a larger coefficient on REV and a smaller coefficient on (RNV-REV), compared to the coefficient of RNV in regression 1. As an alternative to regression 5, regression 8 represents the (RNV-REV) component in a relative manner. With this specification, both recovered components are always significant but they lead to adjusted R^2 higher than regression 1 (what we already knew) in less cases. This result suggests that both components are relevant and that the preferences for risk recovered with Ross's theorem matter.

We investigate this question further with regressions 4, 7 and 9 where we assess directly whether the recovered components contain additional information compared to what is already known (RNV). We find evidence in support of this claim for some indexes. This is particularly true for the SMI index where in regression 9, the recovered components are significant, while RNV is not. For this regression on the S&P500, the recovered components remain significant even when we take into account what we already knew (RNV). These results suggest that decomposing the option-implied variable into its two components (true expectations and risk preferences) yields richer information. When we compare the adjusted R^2 of regression 1 (representing what is already known) to regressions that contain new information (regressions 2 to 9), we observe improvements for all indexes except the FTSE, going from 0.08% to as high as 10.4% in additional explanatory power. This evidence is another indication that by applying the Recovery Theorem we have recovered the market's true expectations as well as the true market risk preferences.

5 – Conclusion

In this paper, we developed an empirical framework to implement the Recovery Theorem of Ross (2015). We evaluate the empirical relevance of the resulting distribution to explain several quantities such as realized volatility. Applying the Recovery Theorem to actual data is not entirely straightforward and some additional steps had to be taken in order to get meaningful results.

Once we recovered the true expectations of the market and its preferences for risk, we first looked at the predictive power of the expected volatility over future realized volatility. In contrast to the risk-neutral volatility that had been used for this purpose in the literature, this true expected volatility recovered with Ross (2015) was not biased downward. While by itself it did not offer superior performance in terms of forecasting, this new information combined with the usual risk-neutral volatility can significantly outperforms a model containing only the risk-neutral volatility. This result holds internationally for several major international stock indexes.

In a second empirical test, we looked at how the changes in those variables could explain the returns of an index. There we also found evidence that applying the Recovery Theorem resulted in extracting the true expected volatility of the market and its preference for risk from the risk-neutral distribution that was already known. This decomposition of risk-neutral expectations resulted in factors that are significant in explaining returns. They also offered a superior explanatory power internationally for most indexes studied.

This paper is a first step in establishing the empirical validity and relevance of the Recovery Theorem as applied in the methodology we presented. As we consider that it was successfully done, this opens the doors for a closer study of the complete expected distribution of returns and the pricing kernel that is obtained from the Recovery Theorem. The superior forecasting performance of the new information in predicting future volatility could have practical applications in finance in terms of portfolio allocations and risk management which we will explore in following papers.

6 – References

Aït-Sahalia, Y. & A.W. Lo. (2000). "Nonparametric Risk Management and Implied Risk Aversion." Journal of Econometrics 94 (1-2): 9–51.

Audrino, F., Huitema, R., & Ludwig, M. (2015). An Empirical Analysis of the Ross Recovery Theorem. SSRN Electronic Journal.

Bakshi, G., Chabi-Yo, F., & Gao, X. (2018). "A Recovery that We Can Trust? Deducing and Testing the Restrictions of the Recovery Theorem." The Review of Financial Studies, 31(2), 532–555.

Bliss, R.R. & Panigirtzoglou, N. (2004). "Option-Implied Risk Aversion Estimates." Journal of Finance, 59 (1), 407-46.

Birru, J. & Figlewski, S. (2012). "Anatomy of a Meltdown: The Risk-Neutral Density for the S&P 500 in the Fall of 2008." Journal of Financial Markets 15, 2, 151-180.

Borovicka, J., Hansen, L. P., & Scheinkman, J. A. (2016). "Misspecified Recovery." Journal of Finance, 71 (6), 2493–2544.

Breeden, D.T. & Litzenberger, R.H. (1978). Prices of State- Contingent Claims Implicit in Option Prices. Journal of Business, 51 (4), 621–651.

Carr, P. & J. Yu. (2012). "Risk, Return, and Ross Recovery." Journal of Derivatives, 20 (1): 38–59.

Christensen, B., & Prabhala, N. (1998). "The relation between implied and realized volatility." Journal of Financial Economics, 50(2), 125–150.

Conrad, J., Dittmar, R., & Ghysels, E. (2013). Ex ante skewness and expected stock returns. The Journal of Finance, 68 (1), 85–124.

Dubynskiy, S. & Goldstein, Robert S. (2013). "Recovering Drifts and Preference Parameters from Financial Derivatives". Available at SSRN: https://ssrn.com/abstract=2244394

Gerd, H., Lunde, A., Shephard, N. and Sheppard, K. (2009). "Oxford-Man Institute's realized library v. 0.2." Oxford-Man Institute, University of Oxford.

Jackwerth, J.C. (2000). "Recovering Risk Aversion from Option Prices and Realized Returns." Review of Financial Studies 13 (2): 433-51.

Jackwerth, J.C. & Menner, M. (2017). "Does the Ross Recovery Theorem Work Empirically?". Available at SSRN: https://ssrn.com/abstract=2960733

Martin, I., & Ross, S. (2013). "The Long Bond." Working paper, Stanford GSB.

Qin, L., & Linetsky, V. (2016) "Positive Eigenfunctions of Markovian Pricing Operators: Hansen-Scheinkman Factorization, Ross Recovery, and Long-Term Pricing." Operations Research, 64(1), 99-117.

Rosenberg, J.V. & R.F. Engle. (2002). "Empirical Pricing Kernels." Journal of Financial Economics 64(3): 341-72.

Ross, S. (2015). "The Recovery Theorem." Journal of Finance, 70 (2), 615–648.

Szakmary, A., Ors, E., Kim, J. K., & Davidson, W. N. (2003). "The predictive power of implied volatility: Evidence from 35 futures markets." Journal of Banking & Finance, 27, 2151–2175.

Walden, J. (2017). "Recovery with unbounded diffusion processes." Review of Finance, 21(4), 1403–1444.

7 – Tables and Figures

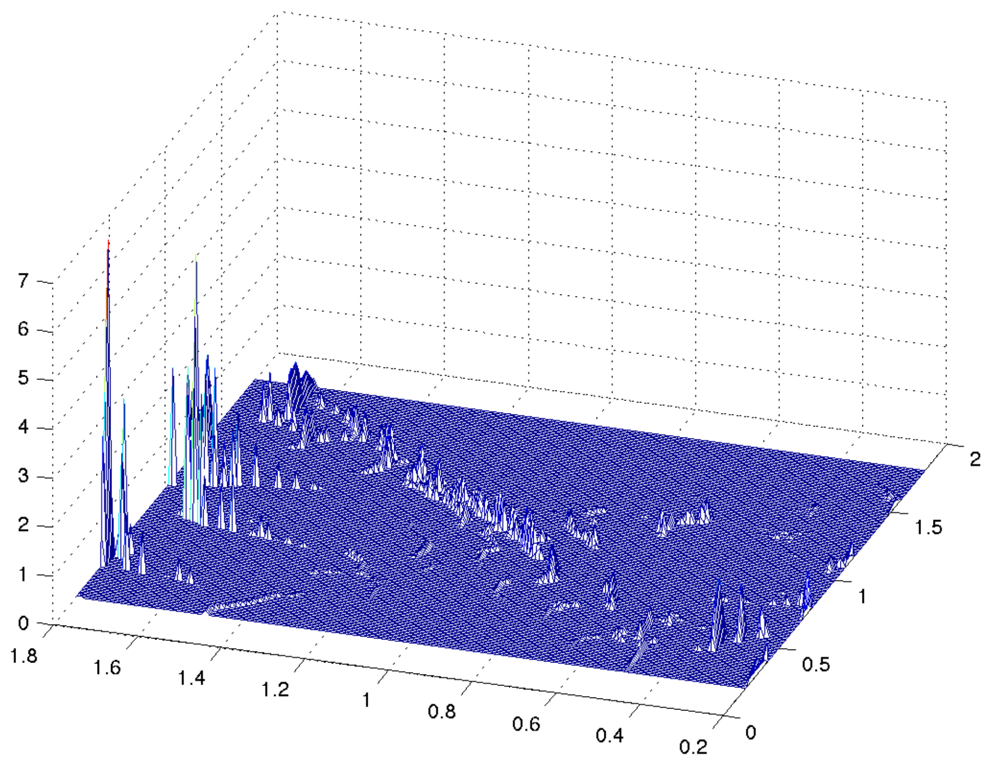

Figure 1. State prices transition matrix P obtained from least-square problem. From option data on the S&P500 index on 2000/01/03. X and Y axes are the moneyness (strike prices relative to the index level that day). The z-axis is the corresponding state price p_{ij}.

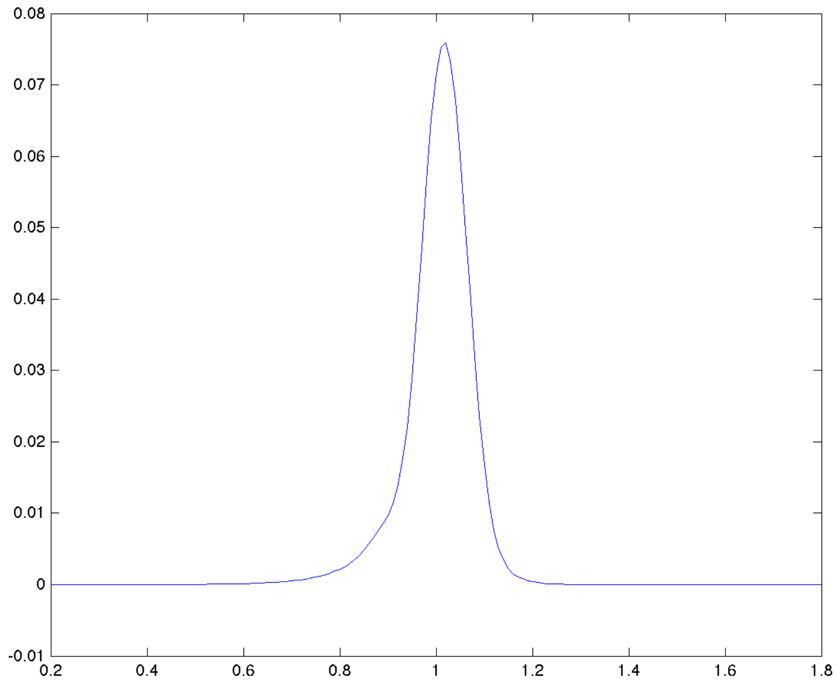

Figure 2. State prices for a one-month horizon. From option data on the S&P500 index on 2000/01/03. X-axis is the moneyness (strike prices relative to the index level that day). The y-axis is the state price.

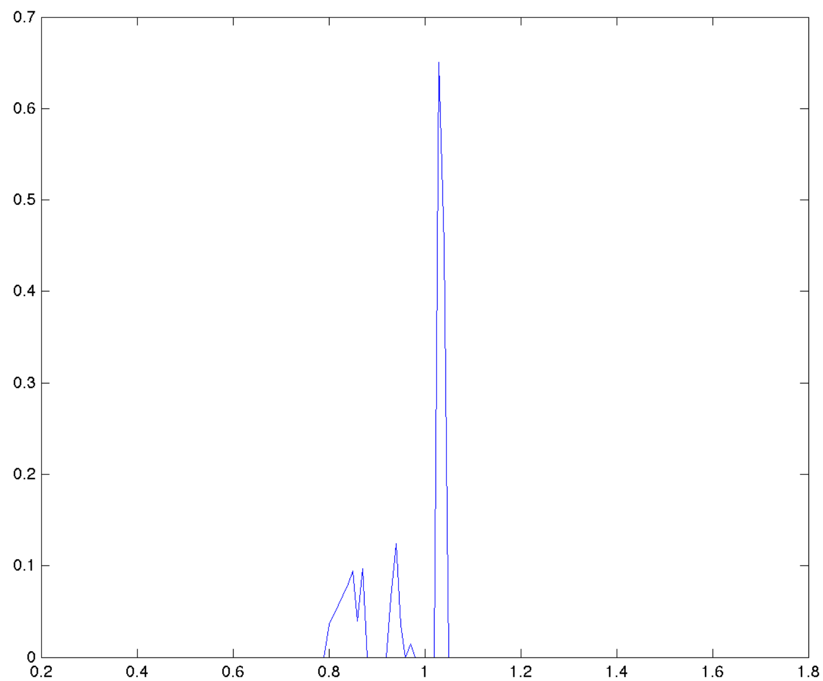

Figure 3. State prices for a one-month horizon according to the least-square solution. From option data on the S&P500 index on 2000/01/03. X axis is the moneyness (strike prices relative to the index level that day). The y-axis is the state price from the least-square solution of P.

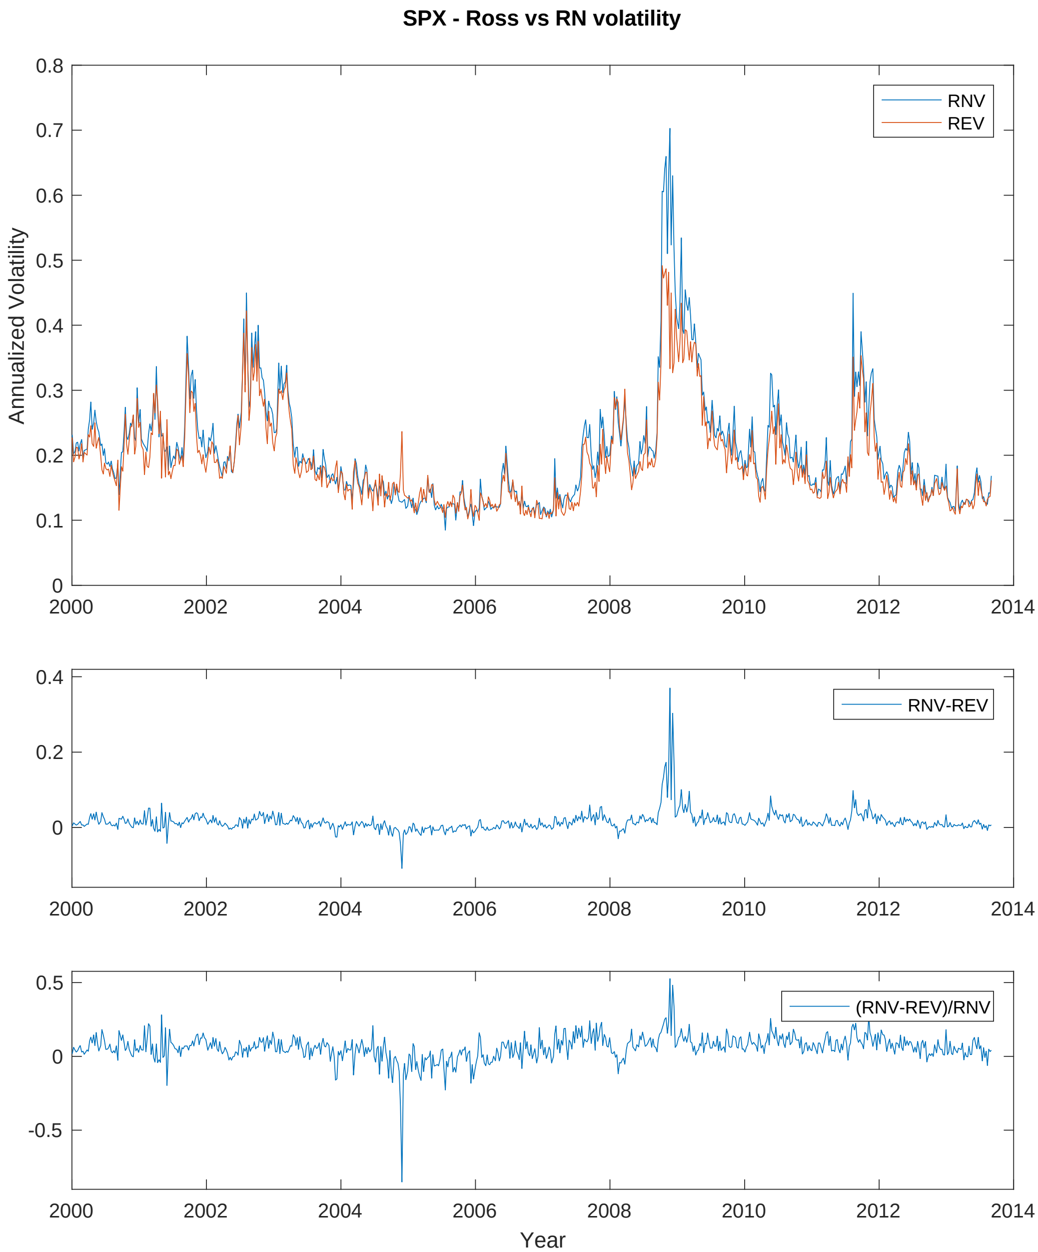

Figure 4a. S&P500 - Weekly time series of explanatory variables. The methodology of this paper is applied on option data every Wednesday from 2000/01/02 to 2013/8/31. *RNV* is the annualized volatility of the risk-neutral distribution while *REV* is the annualized volatility of the true expected distribution of future returns according to the Recovery Theorem of Ross (2015). The difference between those volatilities is presented in the second panel in absolute form and in the third panel in relative form.

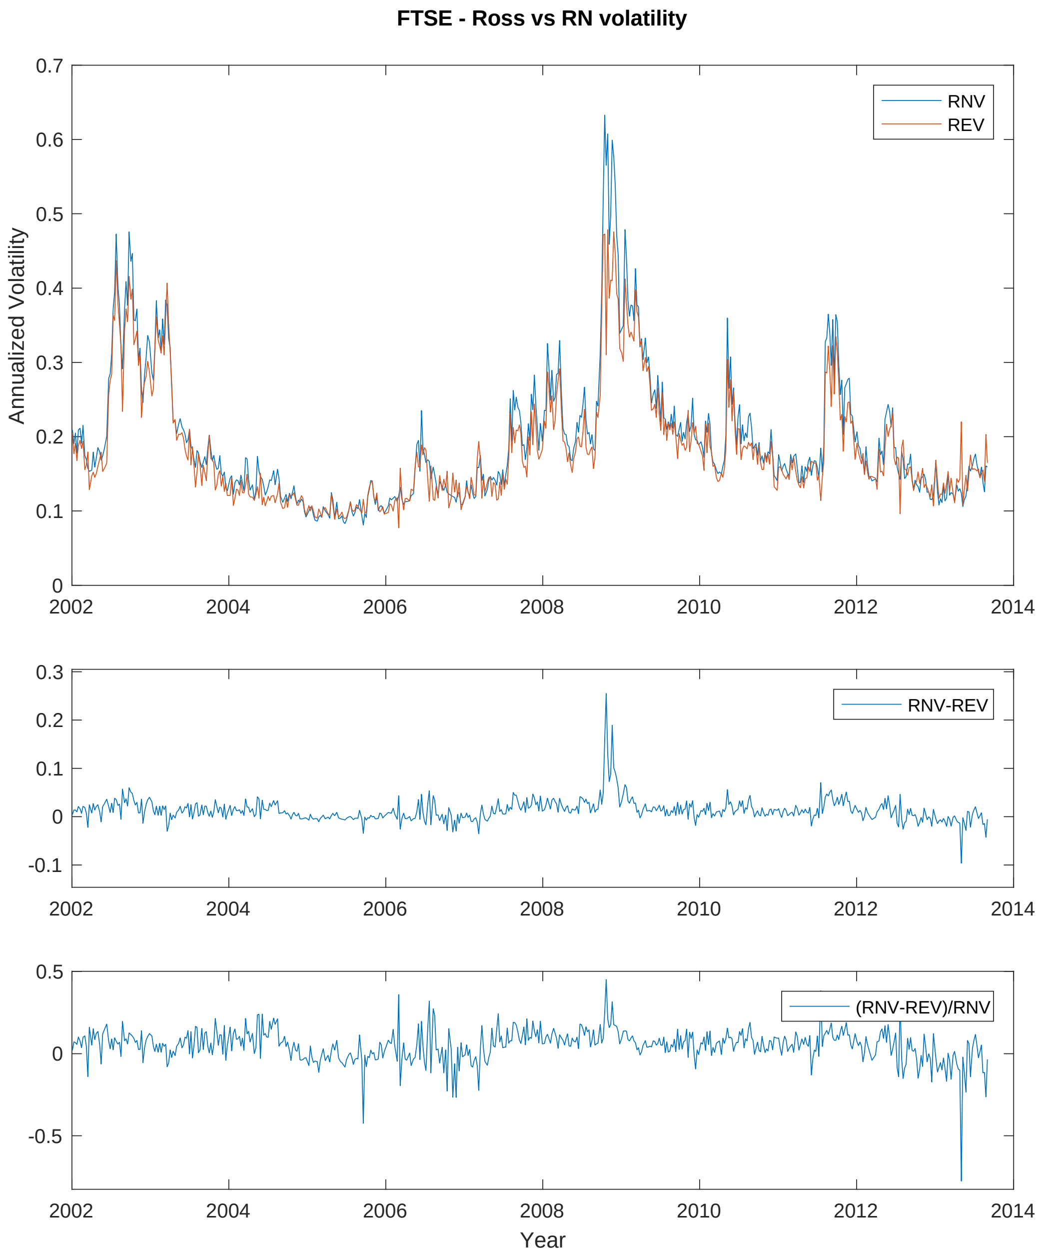

Figure 4b. FTSE - Weekly time series of explanatory variables. The methodology of this paper is applied on option data every Wednesday from 2002/01/02 to 2013/8/31. *RNV* is the annualized volatility of the risk-neutral distribution while *REV* is the annualized volatility of the true expected distribution of future returns according to the Recovery Theorem of Ross (2015). The difference between those volatilities is presented in the second panel in absolute form and in the third panel in relative form.

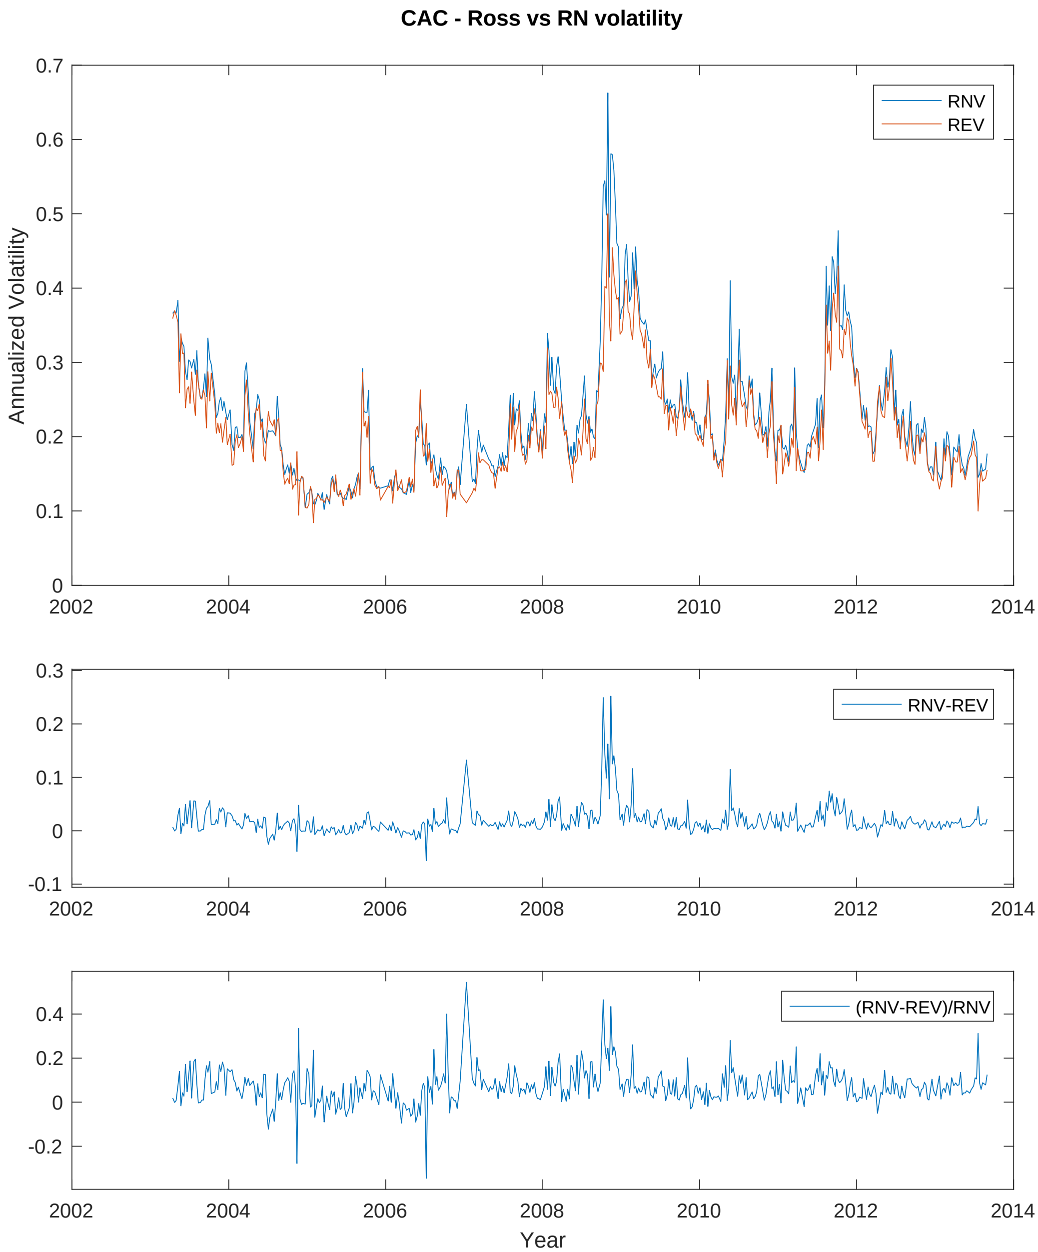

Figure 4c. CAC - Weekly time series of explanatory variables. The methodology of this paper is applied on option data every Wednesday from 2003/04/14 to 2013/8/31. *RNV* is the annualized volatility of the risk-neutral distribution while *REV* is the annualized volatility of the true expected distribution of future returns according to the Recovery Theorem of Ross (2015). The difference between those volatilities is presented in the second panel in absolute form and in the third panel in relative form.

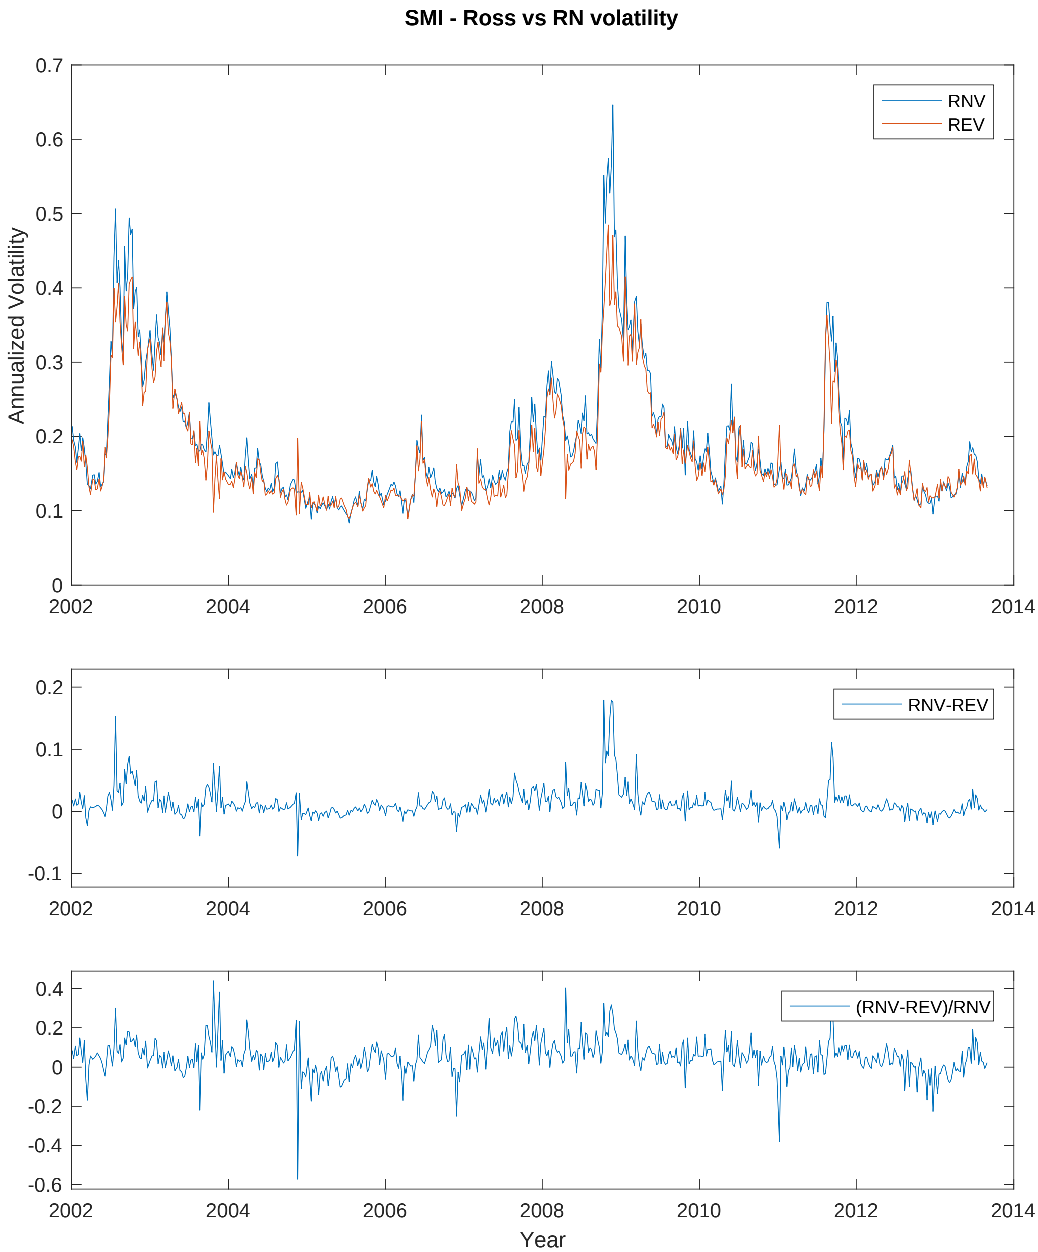

Figure 4d. SMI - Weekly time series of explanatory variables. The methodology of this paper is applied on option data every Wednesday from 2002/01/02 to 2013/8/31. *RNV* is the annualized volatility of the risk-neutral distribution while *REV* is the annualized volatility of the true expected distribution of future returns according to the Recovery Theorem of Ross (2015). The difference between those volatilities is presented in the second panel in absolute form and in the third panel in relative form.

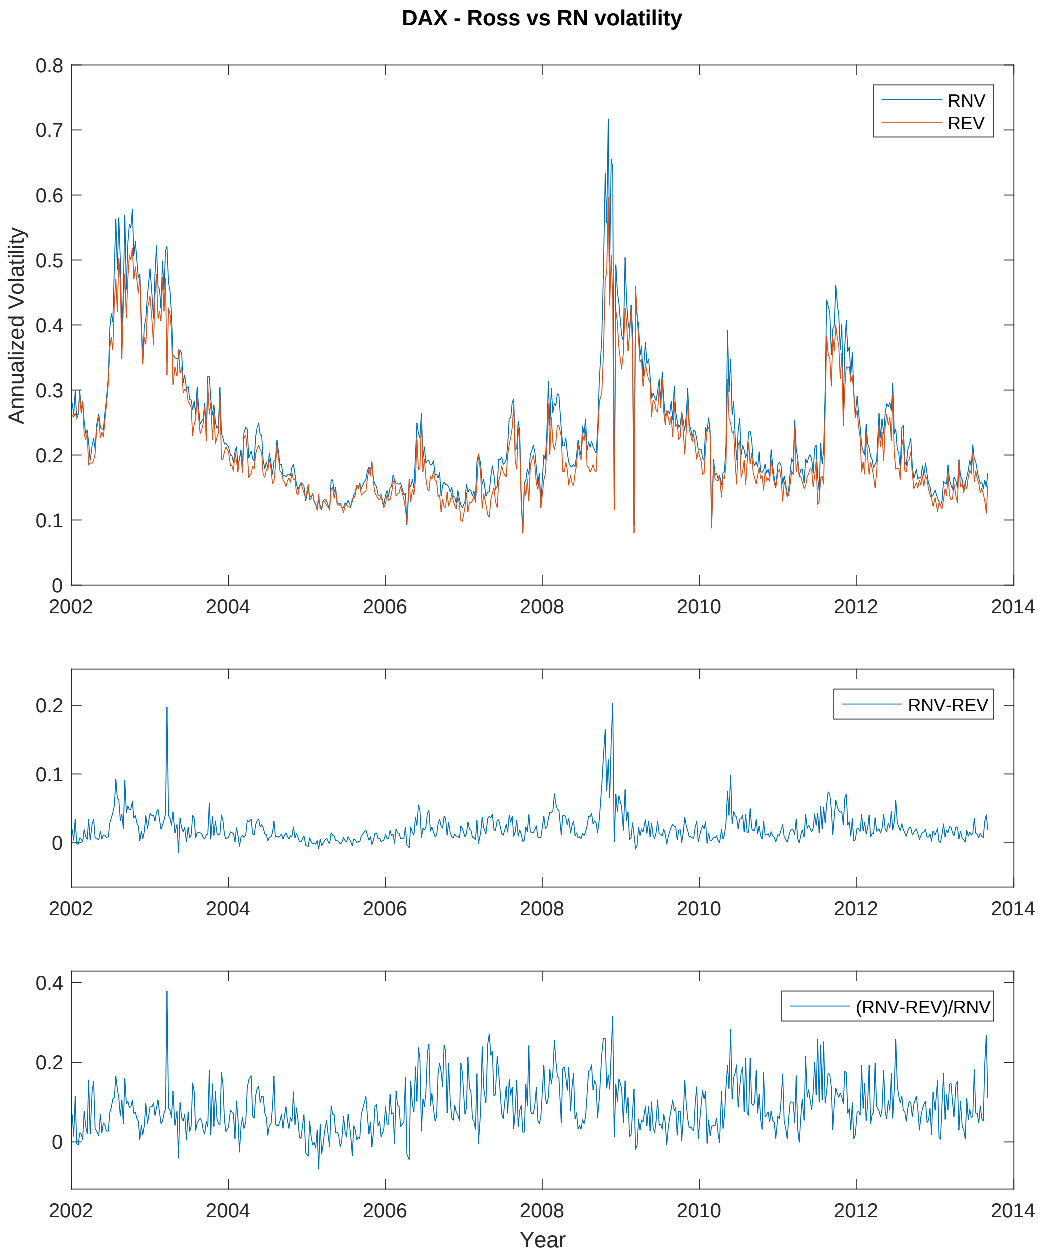

Figure 4e. DAX - Weekly time series of explanatory variables. The methodology of this paper is applied on option data every Wednesday from 2002/01/02 to 2013/8/31. *RNV* is the annualized volatility of the risk-neutral distribution while *REV* is the annualized volatility of the true expected distribution of future returns according to the Recovery Theorem of Ross (2015). The difference between those volatilities is presented in the second panel in absolute form and in the third panel in relative form.

Table 1 - Descriptive statistics for the raw option data. This table reports, for each equity index, the mean of each of the variables relating to characteristics of the raw option data. The source of these data is Optionmetrics Ivy DB USA and Europe. The starting date for the European indexes depends on data availability.

	S&P500	**DAX**	**CAC**	**FTSE**	**SMI**
Dates	**2000/01/02 to 2013/08/31**	**2002/01/02 to 2013/08/31**	**2003/04/14 to 2013/08/31**	**2002/01/02 to 2013/08/31**	**2002/01/02 to 2013/08/31**
Number of options per day	546.71	627.76	362.68	365.44	452.51
Implied volatility	.26	.28	.24	.23	.22
Strike prices	1172.42	5447.45	4044.97	5140.31	6388.85
Number of strikes per day	118.16	103.54	57.70	77.98	83.87
Time to maturity (days)	202.59	358.45	449.53	233.18	335.56
Number of maturities per day	10.98	13.59	13.26	9.85	11.74

Table 2- Descriptive statistics for the time series. The methodology of this paper is applied every Wednesday over the available sample of each index. In the left part of the table, the time series used for predicting the future realized volatility over the next month (RZV) are the risk-neutral volatility (RNV), the volatility of the true expected distribution of returns according the Recovery Theorem of Ross (2015) (REV), and their difference (RNV-REV). As this is a predictive regression, the first month (4 weeks) of RZV and the last month of RNV and REV available in the full sample are discarded. For the variables in the right part of the table, the weekly index returns over the full sample are explained by the weekly variations in RNV, REV, their difference in absolute form (RNV-REV) and in relative form (RNV-REV)/RNV.

Panel A : S&P500	**RZV**	**RNV**	**REV**	**RNV-REV**	**Index Return**	**ΔRNV**	**ΔREV**	**Δ(RNV-REV)**	**Δ[(RNV-REV)/RNV]**
Nb of obs.	715	715	715	715	718	718	718	718	718
Mean	0.1621	0.2065	0.1904	0.0161	0.0009	-0.0001	-0.0001	0.0000	0.0000
Median	0.1384	0.1889	0.1741	0.0115	0.0018	-0.0010	-0.0001	-0.0002	-0.0009
Min	0.07	0.09	0.10	-0.03	-0.15	-0.18	-0.15	-0.30	-0.39
Max	0.75	0.70	0.49	0.37	0.11	0.22	0.17	0.22	0.32
Std Dev.	0.09	0.09	0.07	0.03	0.02	0.03	0.03	0.02	0.08
Skewness	2.47	1.96	1.49	5.22	-0.36	0.56	0.15	-1.63	0.09
Kurtosis	12.40	8.95	5.59	54.32	6.95	13.54	7.93	52.72	5.75
Autocorr(1)	0.975	0.937	0.911	0.606	-0.07	-0.25	-0.40	-0.57	-0.47
P-Perron Pvalue	0.06	0.08	0.09	0.00	0.00	0.00	0.00	0.00	0.00

Panel B : FTSE	**RZV**	**RNV**	**REV**	**RNV-REV**	**Index Return**	**ΔRNV**	**ΔREV**	**Δ(RNV-REV)**	**Δ[(RNV-REV)/RNV]**
Nb of obs.	608	608	608	608	611	611	611	611	611
Mean	0.1684	0.1965	0.1834	0.0130	0.0014	-0.0001	-0.0001	0.0000	-0.0001
Median	0.1374	0.1716	0.1621	0.0091	0.0033	-0.0018	-0.0007	-0.0008	-0.0047
Min	0.06	0.08	0.07	-0.10	-0.12	-0.17	-0.20	-0.23	-0.65
Max	0.75	0.62	0.49	0.33	0.15	0.21	0.19	0.22	0.77
Std Dev.	0.10	0.09	0.08	0.03	0.03	0.03	0.03	0.02	0.11
Skewness	2.29	1.76	1.48	4.25	-0.21	0.74	0.08	0.26	0.87
Kurtosis	10.32	6.90	5.16	46.73	7.17	12.46	9.63	28.91	13.10
Autocorr(1)	0.972	0.934	0.908	0.549	-0.10	-0.26	-0.38	-0.41	-0.50
P-Perron(3)Pval	0.08	0.09	0.08	0.00	0.00	0.00	0.00	0.00	0.00

Panel C : CAC	RZV	RNV	REV	RNV-REV	Index Return	ΔRNV	ΔREV	Δ(RNV-REV)	Δ[(RNV-REV)/RNV]
Nb of obs.	533	533	533	533	537	537	537	537	537
Mean	0.2025	0.2270	0.2084	0.0186	0.0017	-0.0003	-0.0003	0.0000	0.0001
Median	0.1796	0.2097	0.1972	0.0121	0.0039	-0.0020	-0.0005	-0.0004	-0.0002
Min	0.08	0.10	0.00	-0.08	-0.14	-0.18	-0.18	-0.19	-0.89
Max	0.84	0.63	0.51	0.29	0.12	0.20	0.24	0.17	0.88
Std Dev.	0.10	0.09	0.07	0.03	0.03	0.04	0.04	0.03	0.12
Skewness	2.50	1.56	0.95	4.28	-0.60	0.74	0.43	-0.12	-0.04
Kurtosis	12.07	6.58	4.35	30.55	6.18	10.11	9.15	15.78	23.76
Autocorr(1)	0.966	0.911	0.854	0.629	-0.10	-0.26	-0.38	-0.46	-0.48
P-Perron(3)Pval	0.08	0.07	0.06	0.00	0.00	0.00	0.00	0.00	0.00

Panel D : SMI	RZV	RNV	REV	RNV-REV	Index Return	ΔRNV	ΔREV	Δ(RNV-REV)	Δ[(RNV-REV)/RNV]
Nb of obs.	605	605	605	605	609	609	609	609	609
Mean	0.1713	0.1902	0.1764	0.0138	0.0011	-0.0001	-0.0001	0.0000	-0.0001
Median	0.1363	0.1603	0.1511	0.0095	0.0021	-0.0012	0.0002	0.0003	0.0018
Min	0.07	0.08	0.08	-0.09	-0.13	-0.59	-0.23	-0.40	-0.69
Max	0.75	0.70	0.52	0.40	0.16	0.60	0.26	0.40	0.71
Std Dev.	0.10	0.09	0.08	0.03	0.03	0.05	0.04	0.03	0.11
Skewness	2.27	2.06	1.78	6.64	-0.03	0.42	0.41	0.03	0.30
Kurtosis	9.66	8.06	6.11	92.25	8.67	79.43	15.95	100.42	16.35
Autocorr(1)	0.970	0.852	0.893	0.292	-0.16	-0.42	-0.37	-0.48	-0.41
P-Perron(3)Pval	0.07	0.02	0.05	0.00	0.00	0.00	0.00	0.00	0.00

Panel E : DAX	RZV	RNV	REV	RNV-REV	Index Return	ΔRNV	ΔREV	Δ(RNV-REV)	Δ[(RNV-REV)/RNV]
Nb of obs.	607	607	607	607	610	610	610	610	610
Mean	0.2147	0.2380	0.2164	0.0217	0.0013	-0.0002	-0.0002	0.0001	0.0003
Median	0.1772	0.2070	0.1847	0.0162	0.0052	-0.0005	-0.0011	-0.0006	-0.0026
Min	0.09	0.08	0.07	-0.01	-0.15	-0.52	-0.32	-0.20	-0.30
Max	0.85	0.72	0.60	0.21	0.19	0.37	0.27	0.20	0.33
Std Dev.	0.12	0.11	0.09	0.02	0.03	0.05	0.04	0.02	0.07
Skewness	2.00	1.55	1.35	3.72	-0.38	-1.10	-0.30	0.10	0.12
Kurtosis	7.70	5.30	4.28	26.32	7.34	37.71	19.26	28.35	4.64
Autocorr(1)	0.974	0.895	0.903	0.478	-0.14	-0.38	-0.38	-0.44	-0.44
P-Perron(3)Pval	0.08	0.05	0.07	0.00	0.00	0.00	0.00	0.00	0.00

Table 3- Results of the predictive regression on future realized volatility. Weekly time series of the future realized volatility over the next month are regressed on the risk-neutral volatility (RNV), the volatility of the true expected distribution of returns according the Recovery Theorem of Ross (2015) (REV), and their difference (RNV-REV). The regression coefficients and their significance level are presented (*, **, *** for 10%, 5% and 1% respectively). Newey-West corrected standard errors for each regression coefficient are presented in parenthesis. The t-stat of the tests for the null hypothesis of Beta=1 are presented in italics with their level of significance (°,°°,°°° for 10%, 5% and 1% respectively).

	S&P500			FTSE			CAC		
Intercept	-0.009	-0.033*	-0.019	0.004	-0.005	0.009	0.011	0.012	0.027**
	(0.013)	(0.017)	(0.013)	(0.010)	(0.013)	(0.009)	(0.018)	(0.021)	(0.013)
RNV	0.83***		0.90***	0.84***		0.80***	0.84***		0.73***
	(0.07)		(0.08)	(0.06)		(0.06)	(0.09)		(0.07)
Beta =1	*-2.39°°*		*-1.27*	-2.59°°°		-3.35°°°	-1.72°		-4.04°°°
REV		1.02***			0.95***			0.91***	
		(0.10)			(0.09)			(0.12)	
Beta =1		*0.22*			-0.61			-0.72	
RNV-REV			-0.32**			0.21			0.49***
			(0.15)			(0.22)			(0.15)
Adj. R^2	0.627	0.609	0.631	0.629	0.574	0.630	0.514	0.405	0.526

	SMI			DAX		
Intercept	0.015	-0.001	0.003	-0.005	-0.007	-0.003
	(0.014)	(0.011)	(0.011)	(0.011)	(0.013)	(0.011)
RNV	0.82***		0.93***	0.92***		0.90***
	(0.08)		(0.08)	(0.05)		(0.06)
Beta =1	-2.16°°		-0.96	-1.41		-1.81°
REV		0.98***			1.02***	
		(0.08)			(0.07)	
Beta =1		-0.32			0.33	
RNV-REV			-0.58			0.20
			(0.39)			(0.25)
Adj. R^2	0.566	0.572	0.579	0.640	0.608	0.640

Table 4- Results of the regressions explaining index returns. The weekly index returns over the full sample are explained by combinations of the weekly variations in RNV, REV, their difference in absolute form (RNV-REV) and in relative form (RNV-REV)/RNV. Regression coefficient are presented with their level of significance (*, **, *** for 10%, 5% and 1% respectively). Newey-West corrected t-statistics are presented in brackets for each regression coefficient.

	Panel A : S&P500								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Intercept	0.0008	0.0008	0.0009	0.0008	0.0008	0.0009	0.0008	0.0008	0.0008
	[1.39]	[1.24]	[1.05]	[1.40]	[1.40]	[1.02]	[1.39]	[1.37]	[1.40]
ΔRNV	-0.63***			-0.66***			-0.63***		-0.49***
	[-18.27]			[-18.45]			[-18.36]		[-8.00]
ΔREV		-0.49***			-0.66***			-0.68***	-0.19**
		[-5.59]			[-18.45]			[-12.82]	[-2.52]
Δ(RNV-REV)			-0.31***	0.07*	-0.59***				
			[-4.56]	[1.85]	[-15.48]				
Δ[(RNV-REV)/RNV]						-0.05***	0.00	-0.16***	-0.04***
						[-2.85]	[0.41]	[-9.10]	[-2.69]
R^2 Adj.	0.579	0.309	0.085	0.582	0.582	0.023	0.578	0.514	0.586

	Panel B : FTSE								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Intercept	0.0013**	0.0013*	0.0014	0.0013**	0.0013**	0.0014	0.0013**	0.0013**	0.0013**
	[2.28]	[1.93]	[1.59]	[2.28]	[2.28]	[1.52]	[2.28]	[2.22]	[2.28]
ΔRNV	-0.59***			-0.59***			-0.59***		-0.59***
	[-20.05]			[-18.31]			[-19.51]		[-9.07]
ΔREV		-0.43***			-0.59***			-0.62***	0.00
		[-6.21]			[-18.31]			[-10.36]	[-0.02]
Δ(RNV-REV)			-0.32***	-0.02	-0.61***				
			[-5.04]	[-0.58]	[-15.77]				
Δ[(RNV-REV)/RNV]						-0.03***	0.00	-0.12***	-0.01
						[-3.14]	[-0.96]	[-15.62]	[-0.48]
R^2 Adj.	0.605	0.312	0.090	0.605	0.605	0.018	0.605	0.522	0.604

	Panel C : CAC								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Intercept	0.0015*	0.0015*	0.0017	0.0015*	0.0015*	0.0017	0.0015*	0.0015*	0.0015*
	[1.86]	[1.72]	[1.55]	[1.86]	[1.86]	[1.52]	[1.86]	[1.86]	[1.87]
ΔRNV	-0.57***			-0.57***			-0.57***		-0.46***
	[-12.43]			[-12.78]			[-12.49]		[-4.78]
ΔREV		-0.41***			-0.57***			-0.61***	-0.13
		[-7.69]			[-12.78]			[-11.19]	[-1.33]
Δ(RNV-REV)			-0.20***	0.02	-0.55***				
			[-3.16]	[0.69]	[-9.74]				
Δ[(RNV-REV)/RNV]						-0.02*	0.00	-0.12***	-0.03
						[-1.81]	[-0.12]	[-7.66]	[-1.47]
R^2 Adj.	0.526	0.304	0.035	0.525	0.525	0.003	0.525	0.486	0.526

	Panel D : SMI								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Intercept	0.0011	0.0011	0.0011	0.0011	0.0011	0.0011	0.0011	0.0011	0.0011
	[1.28]	[1.31]	[1.16]	[1.31]	[1.31]	[1.15]	[1.28]	[1.34]	[1.36]
ΔRNV	-0.25*			-0.36***			-0.26**		0.24
	[-1.95]			[-4.21]			[-2.08]		[1.22]
ΔREV		-0.37***			-0.36***			-0.42***	-0.74***
		[-4.42]			[-4.21]			[-3.63]	[-2.87]
Δ(RNV-REV)			-0.15	0.25*	-0.11				
			[-1.53]	[1.86]	[-0.70]				
Δ[(RNV-REV)/RNV]						-0.03**	0.01	-0.06**	-0.13***
						[-2.30]	[0.72]	[-2.20]	[-3.09]
R^2 Adj.	0.208	0.233	0.028	0.248	0.248	0.009	0.208	0.291	0.312

	Panel E : DAX								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Intercept	0.0013	0.0012	0.0013	0.0012	0.0012	0.0013	0.0013	0.0013	0.0013
	[1.19]	[1.15]	[1.08]	[1.19]	[1.19]	[1.05]	[1.18]	[1.20]	[1.20]
ΔRNV	-0.35***			-0.40***			-0.36***		0.02
	[-3.15]			[-3.29]			[-3.25]		[0.09]
ΔREV		-0.41***			-0.40***			-0.44***	-0.46
		[-3.70]			[-3.29]			[-3.27]	[-1.42]
Δ(RNV-REV)			-0.25*	0.20	-0.20				
			[-1.69]	[1.37]	[-1.49]				
Δ[(RNV-REV)/RNV]						-0.02	0.03	-0.08**	-0.08
						[-0.82]	[1.18]	[-2.16]	[-1.30]
R^2 Adj.	0.248	0.241	0.028	0.260	0.260	0.001	0.251	0.266	0.265