

RANDOM MATRIX THEORY AND FINANCIAL CORRELATIONS

LAURENT LALOUX, PIERRE CIZEAU and MARC POTTERS

Science & Finance

109-111 rue Victor Hugo, 92532 Levallois Cedex, FRANCE

JEAN-PHILIPPE BOUCHAUD

*Science & Finance , and Service de Physique de l'État Condensé, Centre d'études de Saclay
Orme des Merisiers, 91191 Gif-sur-Yvette Cédex, FRANCE*

We show that results from the theory of random matrices are potentially of great interest to understand the statistical structure of the empirical correlation matrices appearing in the study of multivariate financial time series. We find a remarkable agreement between the theoretical prediction (based on the assumption that the correlation matrix is random) and empirical data concerning the density of eigenvalues associated to the time series of the different stocks of the S&P500 (or other major markets). Finally, we give a specific example to show how this idea can be successfully implemented for improving risk management.

Empirical correlation matrices are of great importance for risk management and asset allocation. The probability of large losses for a certain portfolio or option book is dominated by correlated moves of its different constituents – for example, a position which is simultaneously long in stocks and short in bonds will be risky because stocks and bonds usually move in opposite directions in crisis periods. The study of correlation (or covariance) matrices thus has a long history in finance and is one of the cornerstone of Markowitz's theory of optimal portfolios^{1,2}: given a set of financial assets characterized by their average return and risk, what is the optimal weight of each asset, such that the overall portfolio provides the best return for a fixed level of risk, or conversely, the smallest risk for a given overall return?

More precisely, the average return R_P of a portfolio P of N assets is defined as $R_P = \sum_{i=1}^N p_i R_i$, where p_i ($i = 1, \dots, N$) is the amount of capital invested in the asset i , and $\{R_i\}$ are the expected returns of the individual assets. Similarly, the risk on a portfolio can be associated to the total variance $\sigma_P^2 = \sum_{i,j=1}^N p_i C_{ij} p_j$ where \mathbf{C} is the covariance matrix. The optimal portfolio, which minimizes σ_P^2 for a given value of R_P , can easily be found introducing a Lagrange multiplier, and leads to a linear problem where the matrix \mathbf{C} has to be inverted. In particular, the composition of the least risky portfolio has a large weight on the eigenvectors of \mathbf{C} with the smallest eigenvalues^{1,2}.

However, a reliable empirical determination of a correlation matrix turns out to be difficult. For a set of N different assets, the correlation matrix contains $N(N - 1)/2$ entries, which must be determined from N time series of length T ; if T is not very large compared to N , one should expect that the determination of the covariances is noisy, and therefore that the empirical correlation matrix is

to a large extent random. In this case, the structure of the matrix is dominated by measurement noise, therefore one should be very careful when using this correlation matrix in applications. In particular, as we shall show below, the smallest eigenvalues of this matrix are the most sensitive to this ‘noise’, the corresponding eigenvectors being precisely the ones that determine the least risky portfolios. It is thus important to devise methods which allows one to distinguish ‘signal’ from ‘noise’, i.e. eigenvectors and eigenvalues of the correlation matrix containing real information (which one would like to include for risk control), from those which are devoid of any useful information, and, as such, unstable in time. From this point of view, it is interesting to compare the properties of an empirical correlation matrix \mathbf{C} to a ‘null hypothesis’ purely *random* matrix as one could obtain from a finite time series of strictly independent assets. Deviations from the random matrix case might then suggest the presence of true information. The theory of Random Matrices has a long history since the fifties ³, and many results are known ⁴. As shown below, these results are also of genuine interest in a financial context (see ^{5,6,7}).

The empirical correlation matrix \mathbf{C} is constructed from the time series of price changes^a $\delta x_i(t)$ (where i labels the asset and t the time) through the equation:

$$\mathbf{C}_{ij} = \frac{1}{T} \sum_{t=1}^T \delta x_i(t) \delta x_j(t). \quad (0.1)$$

We can symbolically write Eq. (0.1) as $\mathbf{C} = 1/T \mathbf{M} \mathbf{M}^T$, where \mathbf{M} is a $N \times T$ rectangular matrix, and T denotes matrix transposition. The null hypothesis of independent assets, which we consider now, translates itself in the assumption that the coefficients $M_{it} = \delta x_i(t)$ are independent, identically distributed, random variables ^b the so-called random Wishart matrices or Laguerre ensemble of the Random Matrix theory^{8,10}. We will note $\rho_C(\lambda)$ the density of eigenvalues of \mathbf{C} , defined as:

$$\rho_C(\lambda) = \frac{1}{N} \frac{dn(\lambda)}{d\lambda}, \quad (0.2)$$

where $n(\lambda)$ is the number of eigenvalues of \mathbf{C} less than λ .

Interestingly, if \mathbf{M} is a $T \times N$ random matrix, $\rho_C(\lambda)$ is self-averaging and exactly known in the limit $N \rightarrow \infty$, $T \rightarrow \infty$ and $Q = T/N \geq 1$ fixed ^{8,9}, and reads:

$$\begin{aligned} \rho_C(\lambda) &= \frac{Q}{2\pi\sigma^2} \frac{\sqrt{(\lambda_{max} - \lambda)(\lambda - \lambda_{min})}}{\lambda}, \\ \lambda_{min}^{max} &= \sigma^2(1 + 1/Q \pm 2\sqrt{1/Q}), \end{aligned} \quad (0.3)$$

with $\lambda \in [\lambda_{min}, \lambda_{max}]$, and where σ^2 is equal to the variance of the elements of \mathbf{M} ⁹, equal to 1 with our normalisation. In the limit $Q = 1$ the normalised eigenvalue

^aIn the following we assume that the average value of the δx ’s has been subtracted off, and that the δx ’s are rescaled to have a constant unit volatility $\sigma^2 = \langle \delta x_i^2 \rangle = 1$.

^bNote that even if the ‘true’ correlation matrix \mathbf{C}_{true} is the identity matrix, its empirical determination from a finite time series will generate non trivial eigenvectors and eigenvalues, see ^{8,9}.

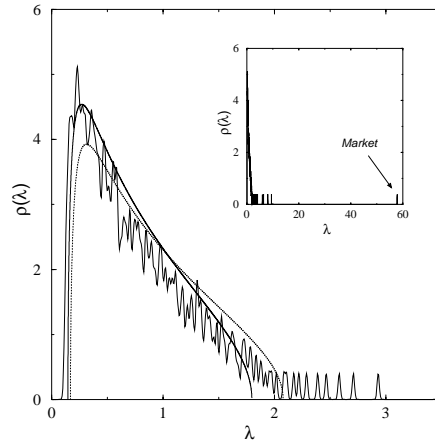


Figure 1: Smoothed density of the eigenvalues of \mathbf{C} , where the correlation matrix \mathbf{C} is extracted from $N = 406$ assets of the S&P500 during the years 1991-1996. For comparison we have plotted the density Eq. (6) for $Q = 3.22$ and $\sigma^2 = 0.85$: this is the theoretical value obtained assuming that the matrix is purely random except for its highest eigenvalue (dotted line). A better fit can be obtained with a smaller value of $\sigma^2 = 0.74$ (solid line), corresponding to 74% of the total variance. Inset: same plot, but including the highest eigenvalue corresponding to the ‘market’, which is found to be 25 times greater than λ_{max} .

density of the matrix \mathbf{M} is the well known Wigner semi-circle law, and the corresponding distribution of the *square* of these eigenvalues (that is, the eigenvalues of \mathbf{C}) is then indeed given by (0.3) for $Q = 1$. The most important features predicted by Eq. (0.3) are:

- the fact that the lower ‘edge’ of the spectrum is strictly positive (except for $Q = 1$); there is therefore no eigenvalues between 0 and λ_{min} . Near this edge, the density of eigenvalues exhibits a sharp maximum, except in the limit $Q = 1$ ($\lambda_{min} = 0$) where it diverges as $\sim 1/\sqrt{\lambda}$.
- the density of eigenvalues also vanishes above a certain upper edge λ_{max} .

Note that the above results are only valid in the limit $N \rightarrow \infty$. For finite N , the singularities present at both edges are smoothed: the edges become somewhat blurred, with a small probability of finding eigenvalues above λ_{max} and below λ_{min} , which goes to zero when N becomes large. The precise way in which these edges become sharp in the large N limit is actually known¹¹.

Now, we want to compare the empirical distribution of the eigenvalues of the correlation matrix of stocks corresponding to different markets with the theoretical prediction given by Eq. (0.3), based on the assumption that the correlation matrix

is purely random. We have studied numerically the density of eigenvalues of the correlation matrix of $N = 406$ assets of the S&P 500, based on daily variations during the years 1991-96, for a total of $T = 1309$ days (the corresponding value of Q is 3.22).

An immediate observation is that the highest eigenvalue λ_1 is 25 times larger than the predicted λ_{max} – see Fig. 1, inset. The corresponding eigenvector is, as expected, the ‘market’ itself, i.e. it has roughly equal components on all the N stocks. The simplest ‘pure noise’ hypothesis is therefore clearly inconsistent with the value of λ_1 . A more reasonable idea is that the components of the correlation matrix which are orthogonal to the ‘market’ is pure noise. This amounts to subtracting the contribution of λ_{max} from the nominal value $\sigma^2 = 1$, leading to $\sigma^2 = 1 - \lambda_{max}/N = 0.85$. The corresponding fit of the empirical distribution is shown as a dotted line in Fig. 1. Several eigenvalues are still above λ_{max} and might contain some information, thereby reducing the variance of the effectively random part of the correlation matrix. One can therefore treat σ^2 as an adjustable parameter. The best fit is obtained for $\sigma^2 = 0.74$, and corresponds to the dark line in Fig. 1, which accounts quite satisfactorily for 94% of the spectrum, while the 6% highest eigenvalues still exceed the theoretical upper edge by a substantial amount. Note that still a better fit could be obtained by allowing for a slightly smaller effective value of Q , which could account for the existence of volatility correlations ¹².

We have repeated the above analysis on different markets, including volatility markets, and found very similar results ⁶.

In a first approximation, the location of the theoretical edge, determined by fitting the part of the density which contains most of the eigenvalues, allows one to distinguish ‘information’ from ‘noise’. However, a more precise procedure can be applied, where the finite N effects are adequately treated, using the results of ¹¹, and where the effect of variability in σ^2 for the different assets can be addressed ^{9,12}.

This idea can be used in practice to reduce the real risk of optimized portfolios. Since the eigenstates corresponding to the ‘noise band’ are not expected to contain real information, one should not distinguish the different eigenvalues and eigenvectors in this sector. This amounts to replacing the restriction of the empirical correlation matrix to the noise band subspace by the identity matrix with a coefficient such that the trace of the matrix is conserved (i.e. suppressing the measurement broadening due to a finite T). This ‘cleaned’ correlation matrix, where the noise has been (at least partially) removed, is then used to construct an optimal portfolio. We have implemented this idea in practice as follows. Using the same data sets as above, the total available period of time has been divided into two equal subperiods. We determine the correlation matrix using the first subperiod, ‘clean’ it, and construct the family of optimal portfolios and the corresponding efficient frontiers. Here we assume that the investor has perfect predictions on the future average returns R_i , i.e. we take for R_i the *observed* return on the next subperiod.

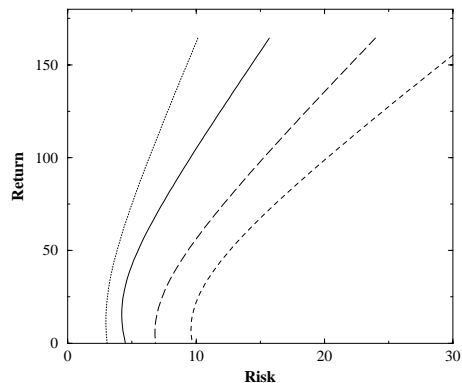


Figure 2: Efficient frontiers from Markovitz optimisation, in the return vs. volatility plane. The leftmost dotted curve correspond to the classical Markovitz case using the empirical correlation matrix. The rightmost short-dashed curve is the realisation of the same portfolio in the second time period (the risk is underestimated by a factor of 3!). The central curves (plain and long-dashed) represents the case of cleaned correlation matrix. The realized risk is now only a factor 1.5 larger than the predicted risk.

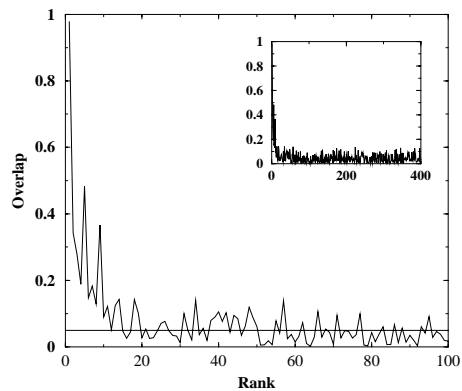


Figure 3: Eigenvector overlap between the two time periods, as a function of the rank n . After rank $n = 10$ (corresponding to the upper edge of the noise band λ_{max}), the overlap reaches the noise level $1/\sqrt{N}$. Inset: Plot of the same quantity in the full range.

The results are shown in Fig. 2 : one sees very clearly that using the empirical correlation matrix leads to a dramatic underestimation of the real risk, by over-investing in artificially low-risk eigenvectors. The risk of the optimized portfolio obtained using a cleaned correlation matrix is more reliable, although the real risk is always larger than the predicted one. This comes from the fact that any amount of uncertainty in the correlation matrix produces, through the very optimisation procedure, a bias towards low-risk portfolios. This can be checked by permuting the two subperiods: one then finds nearly identical efficient frontiers. (This is expected, since for large correlation matrices these frontiers should be self-averaging.) In other words, even if the cleaned correlation matrices are more stable in time than the empirical correlation matrices, there appears to be a genuine time dependence in the structure of the meaningful correlations. This is illustrated in Fig. 3 , where we have plotted the scalar product (or ‘overlap’) of the n^{th} eigenvector for the two subperiods, as a function of the rank n . One can see very clearly that this overlap decreases with n and falls below the noise level $1/\sqrt{N}$ when the corresponding eigenvalue hits the upper edge of the noise band, for $n \simeq 10$. It would be very interesting to study in greater details the dynamics of these eigenmodes out of the ‘noise-band’.

To summarise, we have shown that results from the theory of random matrices is of great interest to understand the statistical structure of the empirical correlation matrices. The central result of the present study is the remarkable agreement between the theoretical prediction and empirical data concerning both the density of eigenvalues and the structure of eigenvectors of the empirical correlation matrices corresponding to several major stock markets. Indeed, in the case of the S&P 500, 94% of the total number of eigenvalues fall in the region where the theoretical formula (0.3) applies. Hence, less than 6% of the eigenvectors which are responsible of 26% of the total volatility, appear to carry some information. This method might be very useful to extract the relevant correlations between financial assets of various types, with interesting potential applications to risk management and portfolio optimisation. We have given a concrete application of this idea, and have shown that the estimate of future risks on optimized portfolios is substantially improved. An interesting avenue is the study of the temporal evolution of the ‘meaningful’ eigenstates.

Acknowledgment

We want to thank J.P. Aguilar, M. Meyer, J.M. Lasry and S. Galluccio for discussions.

References

1. E.J. Elton and M.J. Gruber, *Modern Portfolio Theory and Investment Analysis* (J.Wiley and Sons, New York, 1995); H. Markowitz, *Portfolio Selection: Efficient Diversification of Investments* (J.Wiley and Sons, New York, 1959).

2. J.P. Bouchaud and M. Potters, *Theory of Financial Risk*, (Aléa-Saclay, Eyrolles, Paris, 1997) (in French, English translation to appear in Cambridge University Press, Spring 2000).
3. For a review, see: O. Bohigas, M. J. Giannoni, *Mathematical and computational methods in nuclear physics*, Lecture Notes in Physics, Vol. 209, Springer-Verlag (1983)
4. M. Mehta, *Random Matrices* (Academic Press, New York, 1995).
5. S. Galluccio, J.P. Bouchaud and M. Potters, *Physica A* **259**, 449 (1998).
6. L. Laloux, P. Cizeau, J.P. Bouchaud and M. Potters, *Phys. Rev. Lett.* **83**, 1467 (1999).
7. V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. Amaral, H. E. Stanley, *Phys. Rev. Lett.* **83**, 1471 (1999).
8. A. Edelman, *SIAM J. Matrix Anal. Appl.* **9**, 543 (1988) and references therein.
9. A.M. Sengupta and P.P. Mitra *Distribution of Singular Values for Some Random Matrices*, cond-mat/9709283 preprint, where the authors generalise the results of random Wishart matrices in presence of multivariate correlations both in space and time.
10. T.H. Baker, P.J. Forrester and P.A. Pearce, *J. Phys. A* **31**, 6087 (1998).
11. M.J. Bowick and E. Brézin, *Phys. Lett. B* **268**, 21 (1991), see also M.K. Sener and J.J.M. Verbaarschot, *Phys. Rev. Lett.* **81**, 248 (1998) and, J. Feinberg and A. Zee, *J. Stat. Phys.* **87** 473 (1997).
12. L. Laloux, P. Cizeau, J.P. Bouchaud and M. Potters, in preparation.
13. for a review, see: Hull, J. (1997) *Options, futures and other derivative securities*, Prentice-Hall. For a physicist approach, see ².