

Large dimension forecasting models and random singular value spectra

Jean-Philippe Bouchaud^{1,2}, Laurent Laloux¹
M. Augusta Miceli^{3,1}, Marc Potters¹

December 13, 2005

¹ Science & Finance, Capital Fund Management, 6 Bd Haussmann, 75009 Paris, France.

² Service de Physique de l'État Condensé Orme des Merisiers – CEA Saclay, 91191 Gif sur Yvette Cedex, France

³ Università La Sapienza, Piazzale Aldo Moro, 00185 Roma, Italy

Abstract

We present a general method to detect and extract from a finite time sample statistically meaningful correlations between input and output variables of large dimensionality. Our central result is derived from the theory of free random matrices, and gives an explicit expression for the interval where singular values are expected in the absence of any true correlations between the variables under study. Our result can be seen as the natural generalization of the Marčenko-Pastur distribution for the case of rectangular correlation matrices. We illustrate the interest of our method on a set of macroeconomic time series.

1 Introduction

Finding correlations between observables is at the heart of scientific methodology. Once correlations between “causes” and “effects” are empirically established, one can start devising theoretical models to understand the mechanisms underlying such correlations, and use these models for prediction

purposes. In many cases, the number of possible causes and of resulting effects are both large. For example, one can list an a priori large number of environmental factors possibly favoring the appearance of several symptoms or diseases, or of social/educational factors determining choices and tastes on different topics. A vivid example is provided by Amazon.com, where taste correlations between a huge number of different products (books, CDs, etc.) are sought for. In the context of gene expression networks, the number of input and output chemicals and proteins, described by their concentration, is very large. In an industrial setting, one can monitor a large number of characteristics of a device (engine, hardware, etc.) during the production phase and correlate these with the performances of the final product. In economics and finance, one aims at understanding the relation between a large number of possibly relevant factors, such as interest and exchange rates, industrial production, confidence index, etc. on, say, the evolution of inflation in different sectors of activity [1], or on the price of different stocks. Nowadays, the number of macroeconomic time series available to economists is huge (see below). This has lead Granger [2] and others to suggest that “large models” should be at the forefront of the econometrics agenda. The theoretical study of high dimensional factor models is indeed actively pursued [3, 4, 5, 6, 7, 1], in particular in relation with monetary policy [8, 1].

In the absence of information on the phenomenon under study, a brute force strategy would consist in listing a large number of possible explanatory variables and a large number of output variables, and systematically look for correlations between pairs, in the hope of finding some significant signal. In an econometric context, this is the point of view advocated long ago by Sims [9], who suggested to look at large Vector Autoregressive models, and let the system itself determine the number and the nature of the relevant variables. However, this procedure is rapidly affected by the “dimensionality curse”, also called the problem of sunspot variables in the economics literature [10]. Since the number of observations is always limited, it can happen that two totally unrelated phenomenon (such as, for example, stock prices and sunspots) appear to be correlated over a certain time interval T . More precisely, the correlation coefficient ρ , which would (presumably) be zero if very long time series could be studied, is in fact of the order of $1/\sqrt{T}$ and can be accidentally large. When one tries to correlate systematically N input variables with M output variables, the number of pairs is NM . In the absence of any true correlation between these variables, the largest of these NM empirical correlation coefficients will be, for Gaussian variables, of order

$\rho_{\max} \sim \sqrt{2 \ln(NM)/T}$, which *grows* with NM . For example, $\rho_{\max} \approx 0.25$ for $N = M = 25$ and $T = 200$. If the input and output variables are non Gaussian and have fat-tails, this number can be even larger: if two strongly fluctuating random variable accidentally take large values simultaneously, this will contribute a lot to the empirical correlation even though ρ should be zero for large T .

In this paper we want to discuss how recent results in Random Matrix Theory [11, 12] allow one to alleviate this dimensionality curse and give a precise procedure to extract significant correlations between N input variables and M output variables, when the number of independent observations is T . The idea is to compare the singular value spectrum of the empirical rectangular $M \times N$ correlation matrix with a benchmark, obtained by assuming no correlation at all between the variables. For $T \rightarrow \infty$ at N, M fixed, all singular values should be zero, but this will not be true if T is finite. The singular value spectrum of this benchmark problem can in fact be computed exactly in the limit where $N, M, T \rightarrow \infty$, when the ratios $m = M/T$ and $n = N/T$ fixed. As usual with Random Matrix problems [11, 12], the singular value spectrum develops sharp edges in the asymptotic limit which are to a large extent independent of the distribution of the elements of the matrices. Any singular value observed to be significantly outside these edges can therefore be deemed to carry some relevant information. A similar solution has been known for a long time for standard correlation matrices, for example the correlations of the N input variables between themselves that define an $N \times N$ symmetric matrix. In this case, the benchmark is known as the Wishart ensemble, and the relevant eigenvalue spectrum is given by the Marčenko-Pastur distribution [13, 14, 15]. Applications of this method to financial correlation matrices are relatively recent [16] but very active [17, 18]. Comparing the empirical eigenvalues to the correlation matrix to the theoretical upper edge of the Marčenko-Pastur spectrum allows one to extract statistically significant factors [16] (although some may also be buried below the band edge, see [17]). Similar ideas are starting to be discussed in the econometric community, in particular to deal with the problem of identifying the relevant factors in large dynamical factor models [19], and using them for prediction purposes (see also [6] for a different point of view). Here, we extend the Marčenko-Pastur result to general rectangular, non-equal time correlation matrices. We will first present a precise formulation of our central result, which we will then illustrate using an economically relevant data

set, and finally discuss some possible extensions of our work.

2 Mathematical formulation of the problem

We will consider N input factors, denoted as X_a , $a = 1, \dots, N$ and M output factors Y_α , $\alpha = 1, \dots, M$. There is a total of T observations, where both X_{at} and $Y_{\alpha t}$, $t = 1, \dots, T$ are observed. We assume that all $N + M$ time series are standardized, i.e., both X 's and Y 's have zero mean and variance unity. The X and the Y 's may be completely different, or be the same set of observables but observed at different times, as for example $N = M$ and $Y_{\alpha t} = X_{\alpha t+1}$. From the set of X 's and Y 's one can form two correlations matrices, C_X and C_Y , defined as:

$$(C_X)_{ab} = \frac{1}{T} \sum_{t=1}^T X_{at} X_{bt} \quad (C_Y)_{\alpha\beta} = \frac{1}{T} \sum_{t=1}^T Y_{\alpha t} Y_{\beta t}. \quad (1)$$

In general, the X 's (and the Y 's) have no reason to be independent of each other, and the correlation matrices C_X and C_Y will contain information on their correlations. As alluded to above, one can diagonalize both these matrices; provided $T > N, M$ – which we will assume in the following – all eigenvalues will, in generic cases, be strictly positive. In certain cases, some eigenvalues will however lie close to zero, much below the lower edge of the Marčenko-Pastur interval, corresponding to redundant variables which may need to be taken care of (see below). Disregarding this problem for the moment, we use the corresponding eigenvectors to define a set of uncorrelated, unit variance input variables \hat{X} and output variables \hat{Y} . For example,

$$\hat{X}_{at} = \frac{1}{\sqrt{T\lambda_a}} \sum_b V_{ab} X_{bt}, \quad (2)$$

where λ_a is the a^{th} eigenvalue of C_X and V_{ab} the components of the corresponding eigenvector. Now, by construction, $C_{\hat{X}} = \hat{X} \hat{X}^T$ and $C_{\hat{Y}} = \hat{Y} \hat{Y}^T$ are *exactly* identity matrices, of dimension, respectively, N and M . Using general property of diagonalisation, this means that the $T \times T$ matrices $D_{\hat{X}} = \hat{X}^T \hat{X}$ and $D_{\hat{Y}} = \hat{Y}^T \hat{Y}$ have exactly N (resp. M) eigenvalues equal to 1 and $T - N$ (resp. $T - M$) equal to zero.

Now, consider the $M \times N$ cross-correlation matrix G between the \hat{X} 's and the \hat{Y} 's:

$$(G)_{ab} = \sum_{t=1}^T \hat{Y}_{\alpha t} \hat{X}_{bt} \equiv \hat{Y} \hat{X}^T. \quad (3)$$

The singular value decomposition (SVD) of this matrix answers the following question [20]: what is the (normalised) linear combination of \hat{X} 's on the one hand, and of \hat{Y} 's on the other hand, that have the strongest mutual correlation? In other words, what is the best pair of predictor and predicted variables, given the data? The largest singular value s_{\max} and its corresponding left and right eigenvectors answer precisely this question: the eigenvectors tell us how to construct these optimal linear combinations, and the associated singular value gives us the strength of the cross-correlation. One can now restrict both the input and output spaces to the $N - 1$ and $M - 1$ dimensional sub-spaces orthogonal to the two eigenvectors, and repeat the operation. The list of singular values s_a gives the prediction power, in decreasing order, of the corresponding linear combinations.

3 Singular values from free random matrix theory

How to get these singular values? If $M < N$, the trick is to consider the matrix $M \times M$ matrix GG^T (or the $N \times N$ matrix G^TG if $M > N$), which is symmetrical and has M positive eigenvalues, each of which being equal to the square of a singular value of G itself. The second observation is that the non-zero eigenvalues of $GG^T = \hat{Y}\hat{X}^T\hat{X}\hat{Y}^T$ are the same as those of the $T \times T$ matrix $\mathcal{D} = D_{\hat{X}}D_{\hat{Y}} = \hat{X}^T\hat{X}\hat{Y}^T\hat{Y}$, obtained by swapping the position of \hat{Y} from first to last. In the benchmark situation where the \hat{X} 's and the \hat{Y} 's are independent from each other, the two matrices $D_{\hat{X}}$ and $D_{\hat{Y}}$ are mutually free [11] and one can use results on the product of free matrices to obtain the eigenvalue density from that of the two individual matrices, which are known. The general recipe [14, 11] is to construct first the so-called η -transform of the eigenvalue density $\rho(u)$ of a given $T \times T$ non negative matrix A , defined as:

$$\eta_A(\gamma) = \int du \rho(u) \frac{1}{1 + \gamma u} \equiv \frac{1}{T} \text{Tr} \frac{1}{1 + \gamma A}. \quad (4)$$

From the functional inverse of η_A , one now defines the Σ -transform of A as:

$$\Sigma_A(x) \equiv -\frac{1+x}{x} \eta_A^{-1}(1+x). \quad (5)$$

Endowed with these definitions, one of the fundamental theorems of Free Matrix Theory [11] states that the Σ -transform of the product of two free

matrices A and B is equal to the product of the two Σ -transforms. [A similar, somewhat simpler, theorem exists for *sums* of free matrices, in terms of “R-transforms”, see [11]]. Applying this theorem with $A = D_{\hat{X}}$ and $B = D_{\hat{Y}}$, one finds:

$$\eta_A(\gamma) = 1 - n + \frac{n}{1 + \gamma}, \quad n = \frac{N}{T}; \quad \eta_B(\gamma) = 1 - m + \frac{m}{1 + \gamma}, \quad m = \frac{M}{T}. \quad (6)$$

From this, one easily obtains:

$$\Sigma_{\mathcal{D}}(x) = \Sigma_A(x)\Sigma_B(x) = \frac{(1+x)^2}{(x+n)(x+m)}. \quad (7)$$

Inverting back this relation allows one to derive the η -transform of \mathcal{D} as:

$$\eta_{\mathcal{D}}(\gamma) = \frac{1}{2(1+\gamma)} \left[1 - (\mu + \nu)\gamma + \sqrt{(\mu - \nu)^2\gamma^2 - 2(\mu + \nu + 2\mu\nu)\gamma + 1} \right] \quad (8)$$

with $\mu = m - 1$ and $\nu = n - 1$. The limit $\gamma \rightarrow \infty$ of this quantity gives the density of exactly zero eigenvalues, easily found to be equal to $\max(1 - n, 1 - m)$, meaning, as expected, that the number of non zero eigenvalues of \mathcal{D} is $\min(N, M)$. Depending on the value of $n + m$ compared to unity, the pole at $\gamma = -1$ corresponding to eigenvalues exactly equal to one has a zero weight (for $n + m < 1$) or a non zero weight equal to $n + m - 1$. One can re-write the above result in terms of the more common Stieltjes transform of \mathcal{D} , $S(z) \equiv \eta(-1/z)/z$, which reads:

$$S_{\mathcal{D}}(z) = \frac{1}{2z(z-1)} \left[z + (\mu + \nu) + \sqrt{(\mu - \nu)^2 + 2(\mu + \nu + 2\mu\nu)z + z^2} \right]. \quad (9)$$

The density of eigenvalues $\rho_{\mathcal{D}}(z)$ is then obtained from the standard relation [11]:

$$\rho_{\mathcal{D}}(z) = \lim_{\epsilon \rightarrow 0} \Im \left[\frac{1}{\pi T} \text{Tr} \frac{1}{z + i\epsilon - \mathcal{D}} \right] = \lim_{\epsilon \rightarrow 0} \frac{1}{\pi} \Im [S_{\mathcal{D}}(z + i\epsilon)], \quad (10)$$

which leads to the rather simple final expression, which is the central result of this paper, for the density of singular values s of the original correlation matrix $G \equiv \hat{Y}\hat{X}^T$:

$$\rho(s) = \max(1-n, 1-m)\delta(s) + \max(m+n-1, 0)\delta(s-1) + \frac{\Re \sqrt{(s^2 - \gamma_-)(\gamma_+ - s^2)}}{\pi s(1 - s^2)}, \quad (11)$$

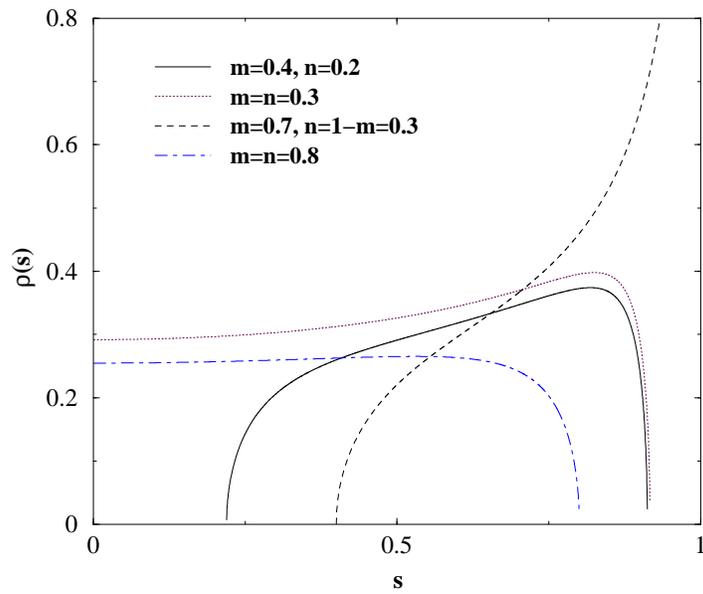


Figure 1: Continuous part of the theoretical random singular value spectrum $\rho(s)$ for different values of n and m . Note that for $n = m$ the spectrum extends down to $s = 0$, whereas for $n + m \rightarrow 1$, the spectrum develops a $(1 - s)^{-1/2}$ singularity, just before the appearance of a δ peak at $s = 1$ of weight $n + m - 1$.

where γ_{\pm} are the two positive roots of the quadratic expression under the square root in Eq. (9) above, which read explicitly:¹

$$\gamma_{\pm} = n + m - 2mn \pm 2\sqrt{mn(1-n)(1-m)}. \quad (12)$$

This is our main technical result, illustrated in Fig. 1. One can check that in the limit $T \rightarrow \infty$ at fixed N, M , all singular values collapse to zero, as they should since there is no true correlations between X and Y ; the allowed band in the limit $n, m \rightarrow 0$ becomes:

$$s \in [|\sqrt{m} - \sqrt{n}|, \sqrt{m} + \sqrt{n}]. \quad (13)$$

When $n \rightarrow m$, the allowed band becomes $s \in [0, 2\sqrt{m(1-m)}]$ (plus a δ function at $s = 1$ when $n + m > 1$), while when $m = 1$, the whole band collapses to a δ function at $s = \sqrt{1-n}$. When $n + m \rightarrow 1^-$, the inchoate δ -peak at $s = 1$ is announced as a singularity of $\rho(s)$ diverging as $(1-s)^{-1/2}$. Finally, when $m \rightarrow 0$ at fixed n , one finds that the whole band collapses again to a δ function at $s = \sqrt{n}$. This last result can be checked directly in the case one has one output variable ($M = 1$) that one tries to correlate optimally with a set of N independent times series of length T . The result can easily be shown to be a correlation of $\sqrt{N/T}$. A plot of the SV density $\rho(s)$ for values of m and n which will be used below is shown in Fig 2, together with a numerical determination of the SVD spectrum of two independent vector time series X and Y , after suitable diagonalisation of their empirical correlation matrices to construct their normalised counterparts, \hat{X} and \hat{Y} . The agreement with our theoretical prediction is excellent.

Note that one could have considered a different benchmark ensemble, where the independent vector time series X and Y are *not* diagonalized and transformed into \hat{X} and \hat{Y} before SVD. The direct SVD spectrum in that case can also be computed as the Σ -convolution of two Marčenko-Pastur distributions with parameters m and n , respectively (noted MP^2 in Fig. 2). The result, derived in the Appendix, is noticeably different from the above prediction (see Fig. 1). This alternative benchmark ensemble is however not well suited for our purpose, because it mixes up the possibly non trivial correlation structure of the input variables and of the output variables themselves

¹One can check that $\gamma_+ \leq 1$ for all values of $n, m < 1$. The upper bound is reached only when $n + m = 1$, in which case the upper edge of the singular value band touches $s = 1$.

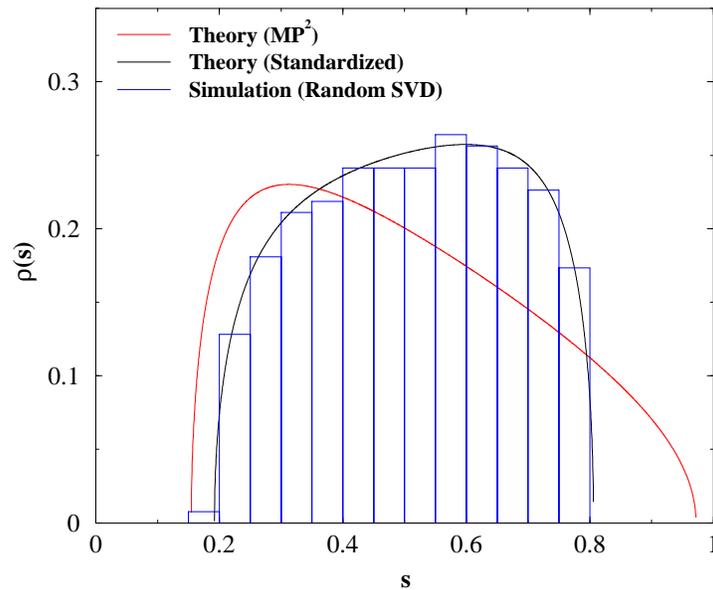


Figure 2: Random Singular Value spectrum $\rho(s)$ for $m = 35/265$ and $n = 76/265$. We show two possible theoretical calculations, corresponding either to bare random vectors X and Y , for which the singular value spectrum is related to the ‘square’ (in the free convolution sense) of the Marčenko-Pastur distribution MP^2 , or standardized vectors \hat{X} and \hat{Y} , obtained after diagonalizing the empirical correlation matrices of X and Y . We also show the results of a numerical simulation of the standardized case with $T = 2650$.

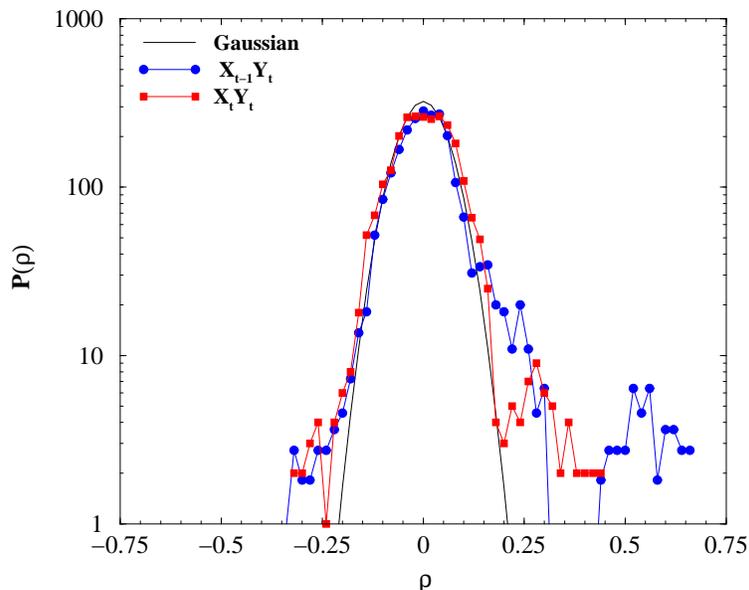


Figure 3: Histogram of the pair correlation coefficient ρ between X 's and Y 's, both at equal times and with one month lag. Note the ‘island’ of correlations around ≈ 0.6 for one-month lagged correlations, which corresponds to correlations between oil prices and energy related CPI's one month later. We also show a Gaussian of variance $1/T$, expected in the absence of any correlations.

with the issue at stake here, namely the *cross*-correlations between input and output variables.

4 Application: inflation vs. economic indicators

We now turn to the analysis of a concrete example. We investigate how different groups of US inflation indexes can be explained using combinations of indicators belonging to different economic sectors. As ‘‘outputs’’ Y , we use 34 indicators of inflation, the monthly changes of the Composite Price Indexes (CPIs), concerning different sectors of activity including and excluding commodities. These indexes were not selected very carefully and some are very redundant, since the point of our study is to show how the proposed method

is able to select itself the relevant variables. As explanatory variables, “inputs” X , we use 76 different macroeconomic indicators from the following categories: industrial production, retail sales, new orders and inventory indexes of all available economic activity sectors, the most important consumer and producer confidence indexes, new payrolls and unemployment difference, interest rates (3 month, 2 and 10 years), G7 exchange rates against the Dollar and the WTI oil price itself. The total number of months in the period June 1983-July 2005 is 265. We want to see whether there is any significant correlation between changes of the CPIs and of the economic indexes, either simultaneous or one month ahead. We also investigated two-month lag correlations, for which we found very little signal.

We first standardized the time series Y and X and form the rectangular correlation matrix between these two quantities, containing 34×76 numbers in the interval $[-1, 1]$. The distribution of these pair correlations is shown in Fig. 3, both for equal time $\hat{Y}_t \hat{X}'_t$ and for one-month lagged $\hat{Y}_t \hat{X}'_{t-1}$ correlations, and compared to a Gaussian distribution of variance T^{-1} . We see that the empirical distributions are significantly broader; in particular an ‘island’ of correlations around ≈ 0.6 appears for the one-month lagged correlations. These correspond to correlations between oil prices and energy related CPIs one month later. The question is whether there are other predictable modes in the system, in particular, are the correlations in the left and right flanks of the central peak meaningful or not? This question is *a priori* non trivial because the kurtosis of some of the variables is quite high, which is expected to ‘fatten’ the distribution of ρ compared to the Gaussian. Within the period of about thirty years covered by our time series, three major rare events happened: the Gulf War (1991-92), the Asian crisis (1998), and the Twin Towers Attack (2001). The kurtosis of the CPIs is the trace of the corresponding outliers, such as the food price index and its ‘negative’, the production price index excluding food, which are strongly sensitive to war events. Among economic indicators, the most responsive series to these events appear to be the inventory-sales ratio, the manufacturing new orders and the motor and motor parts industrial production indexes.

In order to answer precisely the above question, we first turn to the analysis of the empirical self-correlation matrices C_X and C_Y , which we diagonalize and represent the eigenvalues compared to the corresponding Marčenko-Pastur distributions in Fig. 4, expected if the variables were independent (see the Appendix for more details). Since the both the input and output variables are in fact rather strongly correlated at equal times, it is not surpris-

ing to find that some large eigenvalues λ emerge from the Marčenko-Pastur noise band: for C_X , the largest eigenvalue is ≈ 15 , to be compared to the theoretical upper edge of the Marčenko-Pastur distribution 2.358, whereas for C_Y the largest eigenvalue is ≈ 6.2 to be compared with 1.858. But the most important point for our purpose is the rather large number of very small eigenvectors, much below the lower edge of the Marčenko-Pastur distribution ($\lambda_{\min} = 0.215$ for C_X , see Fig. 4). These correspond to linear combinations of redundant (strongly correlated) indicators. Since the definition of \hat{X} and \hat{Y} include a factor $1/\sqrt{\lambda}$ (see Eq. (2)), the eigenvectors corresponding to these small eigenvalues have an artificially enhanced weight. One expects this to induce some extra noise in the system, as will indeed be clear below. Having constructed the set of strictly uncorrelated, unit variance input \hat{X} and output \hat{Y} variables, we determine the singular value spectrum of $G = \hat{Y}\hat{X}^T$. If we keep all variables, this spectrum is in fact indistinguishable from pure noise when \hat{X} precedes \hat{Y} by one month, and only one eigenvalue emerges ($s_{\max} \approx 0.87$ instead of the theoretical value 0.806) when \hat{X} and \hat{Y} are simultaneous.

If we now remove redundant, noisy factors that correspond to, say, $\lambda \leq \lambda_{\min}/2 \approx 0.1$ both in C_X and C_Y , we reduce the number of factors to 50 for \hat{X} and 16 for \hat{Y} ². The cumulative singular value spectrum of this cleaned problem is shown in Fig. 5 and compared again to the corresponding random benchmark. In this case, both for the simultaneous and lagged cases, the top singular values $s_{\max} \approx 0.73$ (resp. $s_{\max} \approx 0.81$) are very clearly above the theoretical upper edge $s_{ue} \approx 0.642$, indicating the presence of some true correlations. The top singular values s_{\max} rapidly sticks onto the theoretical edge as the lag increases. For the one-month lagged case, there might be a second meaningful singular value at $s = 0.66$. The structure of the corresponding eigenvectors allows one to construct a linear combination of economic indicators explaining a linear combinations of CPIs series. The combination of economic indicators corresponding to the top singular value evidences the main economic factors affecting inflation indicators: oil prices obviously correlated to energy production increases and electricity production decreases that explain the CPIs indexes including oil and energy. The second factor includes the next important elements of the economy: employment (new payrolls) affects directly the “core” indexes and the CPI indexes

²The results we find are however weakly dependent on the choice of this lower cut-off, provided very small λ 's are removed.

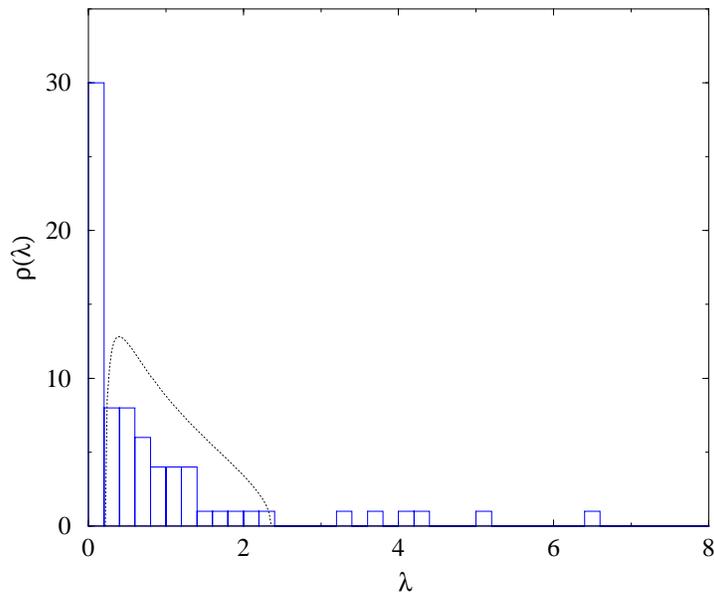


Figure 4: Eigenvalue spectrum of the $N \times N$ correlation matrix of the input variables C_X , compared to the Marčenko-Pastur distribution with parameter $n = 76/265$. Clearly, the fit is very bad, meaning that the input variables are strongly correlated; the top eigenvalues $\lambda_{\max} \approx 15$ is in fact not shown. Note the large number of very small eigenvectors corresponding to combinations of strongly correlated indicators, that are pure noise but have a small volatility.

excluding oil. New economy production (high tech, media & communication) is actually a proxy for productivity increases, and therefore exhibits a negative correlation with the same core indexes. We have also computed the inverse participation ratio of all left and right eigenvectors with similar conclusions [16]: all eigenvectors have a participation ratio close to the informationless Porter-Thomas result, except those corresponding to singular values above the upper edge.

Since Y_{t-1} may also contain some information to predict Y_t , one could also study, in the spirit of general Vector Autoregressive Models [3, 5, 1], the case where we consider the full vector of observables Z of size 111, obtained by merging together X and Y . We again define the normalised vector \hat{Z} , remove all redundant eigenvalues of $\hat{Z}\hat{Z}'$ smaller than 0.1, and compute the singular value spectrum of $\hat{Z}_t\hat{Z}_{t-1}'$. The size of this cleaned matrix is 62×62 ,

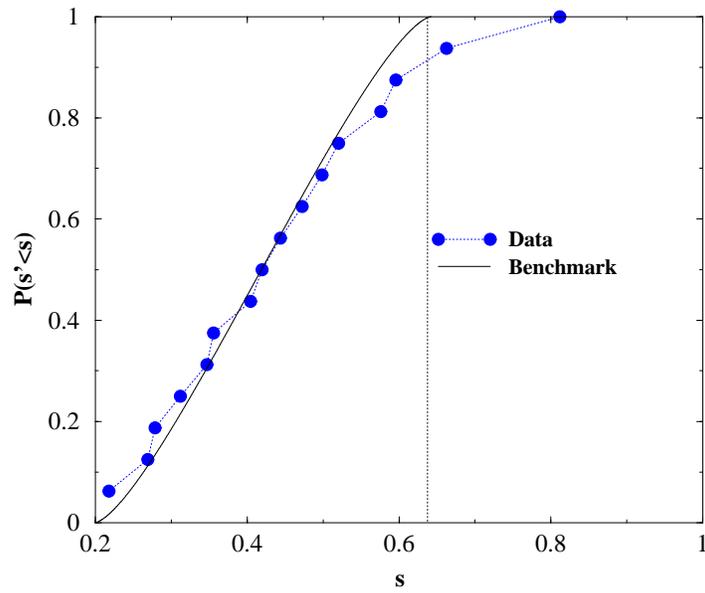


Figure 5: Cumulative singular value distribution for the “cleaned” problem, i.e. removing the factors with very small volatilities, leaving 50 factors in \hat{X} and 16 in \hat{Y} . The correlations we consider are lagged and correspond to $\hat{Y}_t \hat{X}_{t-1}^T$. The filled circles correspond to the 16 empirical singular values, and the plain line is the theoretical prediction in the purely random case with $n = 50/265$ and $m = 16/265$. Note that the top singular value $s_{\max} \approx 0.81$ clearly stands out of the noise band, the edge of which is at $s_{ue} = 0.642$. Finite T corrections are expected to smooth the edge over a region of size $T^{-2/3} \approx 0.025$ for $T = 265$.

and the upper edge of the random singular value spectrum is $s_{ue} \approx 0.84$. We now find that the top singular value is at $s_{\max} \approx 0.97$, and that ~ 8 factors have singular values above the upper edge of the random spectrum. The top singular value corresponds to sales and inventory/sales ratio, followed by CPIs that tend to be correlated over time. Further results are less intuitively simple. This analysis can of course be generalized to larger lags, by studying $\hat{Z}_t \hat{Z}_{t-n}^T$. We find that even for $n = 4$, there are still three singular values above the upper edge. The SVD results are therefore of great help to rank the importance of autocorrelations of degree n in the system; we will explore this point further in a future publication.

5 Conclusions and extensions

The conclusions of this illustrative empirical study are twofold: (i) in general, both input and output variables have a non trivial correlation structure, with many redundant factors which add a significant amount of noise in the problem. Therefore, in a first step, some data cleaning must be performed by eliminating these redundant variables; (ii) the singular value spectrum, compared to its purely random counterpart, allows one to answer precisely the question of the number and relevance of independent predictable factors in the problem under study. In the case considered, we have seen that although the number of pairs of apparently correlated factors is large (see Fig. 3), only a few modes can in fact be considered as containing useful information, in the sense that their singular value exceeds our analytical upper edge given in Eq. (11). When studying the full problem where all variables are treated together, we find that the effective dimensionality of the problem drops from 111 to eight or so independent, predictable factors. This compares quite well with the number seven quoted by Stock and Watson within their dynamical factor analysis of a similar data set [1]. A more thorough comparison of our results with those of the econometrics literature will be presented elsewhere.

What we mean by ‘exceed the upper edge’ should of course be specified more accurately, beyond the eye-balling procedure that we implicitly rely on. In order to have a more precise criterion, one should study the statistics of the top eigenvalue of \mathcal{D} , which is, in analogy with the known results for the Wishart ensemble, most probably given by a Tracy-Widom distribution, at least for Gaussian random variables (see [21, 22] for recent progress and references). For finite T , we expect the top eigenvalue of \mathcal{D} to ooze away

from the theoretical edge by a quantity of order $T^{-2/3} \approx 0.025$ for $T = 265$. Therefore, the difference between $s_{\max} \approx 0.81$ and the theoretical edge $s_{ue} = 0.642$ reported in Fig. 5 can safely be considered as significant when all variables are Gaussian. However, although the density of singular values is to a large degree independent of the distribution of the matrix entries, one should expect that the fuzzy region around the theoretical edge expands significantly if the input and output variables have fat tails. In particular, the Tracy-Widom distribution is expected to breakdown in some way that would be very interesting to characterize precisely. We leave this problem to future investigations.

In conclusion, we have presented a general method to extract statistically meaningful correlations between an arbitrary collection of input and output variables of which only a finite time sample is available. Our central result is derived from the theory of free random matrices, and gives an explicit expression for the interval where singular values are expected *in the absence of any true correlations* between the variables under study. Our result can be seen as the natural generalization of the Marčenko-Pastur distribution for the case of rectangular correlation matrices. The potential applications of this method are quite numerous and we hope that our results will prove useful in different fields where multivariate correlations are relevant.

Acknowledgments: We wish to thank Gérard Ben Arous, Florent Benaych-Georges and Jack Silverstein for most useful discussions on Random Matrix Theory.

Appendix: the MP^2 case

As indicated in the main text, one could have chosen as a benchmark the case where all (standardized) variables X and Y are uncorrelated, meaning that the ensemble average $E(C_X) = E(XX^T)$ and $E(C_Y) = E(YY^T)$ are equal to the unit matrix, whereas the ensemble average cross-correlation $E(G) = E(YX^T)$ is identically zero. However, for a given finite size sample, the eigenvalues of C_X and C_Y will differ from unit, and the singular values of G will not be zero. The statistics of the eigenvalues of C_X and C_Y is well known to be given by the Marčenko-Pastur distribution with parameters n

and m respectively, which reads, for $\beta = n, m < 1$:

$$\rho_{MP}(\lambda) = \frac{1}{2\pi\beta\lambda} \Re \sqrt{(\lambda - \lambda_{\min})(\lambda_{\max} - \lambda)}, \quad (14)$$

with

$$\lambda_{\min} = (1 - \sqrt{\beta})^2 \quad \lambda_{\max} = (1 + \sqrt{\beta})^2. \quad (15)$$

The Σ -transform of this density takes a particularly simple form:

$$\Sigma(x) = \frac{1}{1 + \beta x}. \quad (16)$$

Now, as explained in the main text, the singular values of G are obtained as the square-root of the eigenvalues of $D = X^T X Y^T Y$. Since $X^T X$ and $Y^T Y$ are mutually free, one can again use the multiplication rule of Σ -transforms, after having noted that the Σ -transform of the $T \times T$ matrices $X^T X$ and $Y^T Y$ are now given by:

$$\Sigma(x) = \frac{1}{\beta + x}. \quad (17)$$

One therefore finds that the η transform of D is obtained by solving the following cubic equation for x :

$$\eta^{-1}(1 + x) = -\frac{1 + x}{x(n + x)(m + x)}, \quad (18)$$

which can be done explicitly, leading to the following (lengthy) result. Denote $z = s^2$, one should first compute the following two functions:

$$f_1(z) = 1 + m^2 + n^2 - mn - m - n + 3z \quad (19)$$

and

$$f_2(z) = 2 - 3m(1 - m) - 3n(1 - n) - 3mn(n + m - 4) + 2(m^3 + n^3) + 9z(1 + m + n). \quad (20)$$

Then, form:

$$\Delta = -4f_1(z)^3 + f_2(z)^2. \quad (21)$$

If $\Delta > 0$, one introduces a second auxiliary variable Γ :

$$\Gamma = f_2(z) - \sqrt{\Delta}, \quad (22)$$

to compute $\rho_2(z)$:

$$\pi\rho_2(z) = -\frac{\Gamma^{1/3}}{2^{4/3}3^{1/2}z} + \frac{f_1(z)}{2^{2/3}3^{1/2}\Gamma^{1/3}z}. \quad (23)$$

Finally, the density $\rho(s)$ is given by:

$$\rho(s) = 2s\rho_2(s^2). \quad (24)$$

References

- [1] for a review, see: J. H. Stock, and M. W. Watson, *Implications of dynamical factor models for VAR analysis*, working paper, June 2005.
- [2] C. W. J. Granger, *Macroeconometrics, Past and Future*, Journal of Econometrics, **100**, 17 (2001).
- [3] J. Geweke *The Dynamic Factor Analysis of Economic Time Series* in D.J. Aigner and A.S. Goldberger eds. *Latent Variables in Social Economic Models*, North Holland: Amsterdam (1997).
- [4] J. H. Stock, M. W. Watson, *Forecasting Inflation*, Journal of Monetary Economics, **44**, 293-335 (1999), *Macroeconomic forecasting using diffusion indexes*, Journal of Business and Economic Statistics, **20**, 147-162 (2002), *Forecasting using principal components from a large number of predictors*, Journal of the American Statistical Association, **97**, 1167-1179 (2002).
- [5] M. Forni, M. Hallin, M. Lippi and L. Reichlin, *The Generalized Dynamic Factor Model: Identification and Estimation*, The Review of Economic and Statistics, **82**, 540-554 (2000); *The Generalized Dynamic Factor Model: Consistency and Rates*, Journal of Econometrics, **119**, 231-255 (2004), *The Generalized Dynamic Factor Model: One-Sided Estimation and Forecastings*, mimeo (1999).
- [6] J. Bai, *Inferential theory for factor models of large dimensions*, Econometrica, **71**, 135-171 (2003).
- [7] J. Bai, and S. Ng (2002) *Determining the number of factors in approximate factor model*, Econometrica, **70**, 191-221 (2002).
- [8] B. Bernanke, J. Boivin, *Monetary policy in a data rich environment*, Journal of Monetary Economics, **50** 525 (2003).

- [9] C. A. Sims, *Macroeconomics and Reality*, *Econometrica*, **48**, 1-48 (1980).
- [10] see, e.g. M. Woodford, *Learning to believe in sunspots*, *Econometrica*, **58**, 277-307, (1990); see also M. Wyart, J.P. Bouchaud, *Self referential behaviour, overreaction and conventions in financial markets*, to appear in JEBO, (2005).
- [11] for a recent review, see: A. Tulino, S. Verdù, *Random Matrix Theory and Wireless Communications*, *Foundations and Trends in Communication and Information Theory*, **1**, 1-182 (2004).
- [12] A. Edelman, N. Raj Rao, *Random Matrix Theory*, *Acta Numerica*, 1-65 (2005).
- [13] V. A. Marčenko and L. A. Pastur, *Distribution of eigenvalues for some sets of random matrices*, *Math. USSR-Sb*, **1**, 457-483 (1967).
- [14] J. W. Silverstein and Z. D. Bai, *On the empirical distribution of eigenvalues of a class of large dimensional random matrices*, *Journal of Multivariate Analysis*, **54** 175 (1995).
- [15] A. N. Sengupta and P. Mitra, *Distributions of singular values for some random matrices*, *Phys. Rev. E* **80**, 3389 (1999).
- [16] L. Laloux, P. Cizeau, J.-P. Bouchaud and M. Potters, *Noise dressing of financial correlation matrices*, *Phys. Rev. Lett.* **83**, 1467 (1999); L. Laloux, P. Cizeau, J.-P. Bouchaud and M. Potters, *Random matrix theory and financial correlations*, *Int. J. Theor. Appl. Finance* **3**, 391 (2000); M. Potters, J.-P. Bouchaud and L. Laloux, *Financial Applications of Random Matrix Theory: Old laces and new pieces*, *Acta Physica Polonica B*, **36** 2767 (2005).
- [17] Z. Burda, A. Görlich, A. Jarosz and J. Jurkiewicz, *Signal and Noise in Correlation Matrix*, *Physica A*, **343**, 295-310 (2004).
- [18] see the proceedings of the conference on *Applications of Random Matrix Theory*, published in *Acta Physica Polonica B*, **36** 2603-2838 (2005).
- [19] G. Kapetanios, *A new Method for Determining the Number of Factors in Factor Models with Large Datasets*, *mimeo Univ. of London* (2004).
- [20] see, e.g., W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, *Numerical Recipes in C : The Art of Scientific Computing*, Cambridge University Press (1992).

- [21] J. Baik, G. Ben Arous, and S. Peche, *Phase transition of the largest eigenvalue for non-null complex sample covariance matrices*, <http://xxx.lanl.gov/abs/math.PR/0403022>, to appear in Ann. Prob.
- [22] N. El Karoui, *Recent results about the largest eigenvalue of random covariance matrices and statistical application*, Acta Physica Polonica B, **36** 2681 (2005).