On the Statistical Properties of a Simple Index of Persistence: Theory and Empirical Applications

Giorgio Fagiolo*

September 2011

Abstract

The issue of estimating the persistence of a dynamic process arises frequently in many micro and macro applied exercises. In this note we study the statistical properties of a simple index of persistence. We show that the index is normally-distributed and we estimate its expected value and standard deviations under different benchmark null hypotheses for the underlying data-generating mechanisms. Finally, we present a macroeconomic application to exemplify how to statistically discriminate between high and low levels of the index according to the maintained null hypothesis.

JEL codes: C1, C4, F1, E3

Keywords: Persistence; Dynamic Processes; Statistical Properties; Output Growth; BIlateral Trade Imbalances.

*Institute of Economics, Sant'Anna School for Advanced Studies. Address: Piazza Martiri della Libertà 33, 56127 Pisa (Italy). Tel: +39-050-883359. Fax: +39-050-883344. Email: giorgio.fagiolo@sssup.it

1 Introduction

The issue of estimating the persistence of a dynamic process arises frequently in many micro and macro applied exercises. A commonly-employed practice is to start from a database recording the level of a (discrete or continuous) variable Y across C units and T time periods, and to estimate the transition probability matrix (TPM) of a time-homogeneous Markov chain defined on N classes (Anderson and Goodman, 1957). These may be absolute values or the quantiles of the Y-distribution in each given time period. A good persistence indicator is based on the idea that the larger the probability mass lying close to the main diagonal of the TPM, the more persistent the process, because the larger will be in that case the probability that a unit will be characterized tomorrow by a value of the variable Y closer to that of today.

Since the seminal works of Prais (1955) and Matras (1961) on social mobility —the complement of persistence— a lot of effort has been devoted to introduce persistence indicators that were able to satisfy some baseline axioms (e.g., immobility and monotonicity, see for example Bartholomew, 1973; Shorrocks, 1976; Geweke et al., 1986). In general, persistence indicators may be classified in "synthetic" ones (i.e., those fully or partially employing the information in the TPM to deliver a single number characterizing the persistence of the underlying process) or "profile-based" ones (i.e., those that associate to the process a persistence profile, cf. e.g. the persistence tents of Bartelsman and Dhrymes, 1994).

Despite many micro and macro applied works have been extensively employing persistence indicators in the last years[1], little is known about their statistical properties. In fact, to correctly assess whether a process is highly persistent or not, it is seldom sufficient to know the unconditional range of variation of the indicator that is being employed, as more on its distributional properties and moments under some null hypothesis is typically required.

In this note we address this issue using a very simple "synthetic" index of persistence, which makes use of the full information of the TPM. The index, which coincides with a particular case of (the complement to 1 of) the D index of mobility introduced by Bartholomew (1973, p.19, eq. 2.8), is shown to have, under general assumptions, a normal distribution. By means of simulation exercises, we therefore characterize its statistical properties under some simple null hypotheses about the underlying data-generation mechanisms, e.g. full-mobility case, auto-regressive of order one processes, etc.. Finally, we present a macroeconomic application to exemplify the foregoing methods and explore, for a panel of 110 countries, the persistence level of four macro-variables (country per-capita GDP, GDP country growth rates, absolute and relative bilateral trade imbalances). Our results suggest that the observed values of our simple persistence index can be easily compared to its distribution under some benchmark null hypotheses for the underlying process. This allows one to provide statistically-sound implications about observed persistence and to infer interesting insights about the relative persistence of different processes.

The note is organized as follows. Section 2 defines the persistence index and analyzes some of its simplest features. The statistical properties of the index are discussed in Section 3, whereas Section 4 contains an empirical application. Finally, Section 5 concludes.

[1] See among others Quah (1993), Bartelsman and Dhrymes (1994), Konings and Roodhooft (1997), Jarvis and Jenkins (1998), Cefis (2003), Jafry and Schuermann (2004), Ezcurra et al. (2006), Fiaschi and Lavezzi (2007), David and Rullani (2008), Zhang et al. (2010).

2 Theory

Let $M = \{m_{i,j}\}$ be a $N \times N$ transition probability matrix of an ergodic, time-homogeneous Markov chain, where $m_{i,j} \geq 0$ and $\sum_j m_{i,j} = 1$ all i. Suppose wlog that N is even and define $d_j^+(h) = m_{j,h+j-1}$ and $d_j^-(h) = m_{h+j-1,j}$, for $1 \leq j \leq N-h+1$, $h = 1, \dots, N$. Define also the vectors $D^+(h) = \{d_1^+(h), \dots, d_{N-h+1}^+(h)\}$ and $D^-(h) = \{d_1^-(h), \dots, d_{N-h+1}^-(h)\}$. It is easy to see that, when $h = 1$, $D^+(1) = D^-(1) = D$, where D is the main diagonal of M. When $h > 1$, $D^+(h)$ and $D^-(h)$ represent a diagonal list of entries of the TPM lying $h-1$ steps away from the main diagonal and, respectively, above and below it. In what follows, we will call $D^+(h)$ a h-upper diagonal of M and $D^-(h)$ a h-lower diagonal of M.

Let $S_h^+ = \sum_{j=1}^{N-h+1} d_j^+(h)$ and $S_h^- = \sum_{j=1}^{N-h+1} d_j^-(h)$ the total probability mass accounted for $D^+(h)$ and $D^-(h)$, respectively. A benchmark measure of persistence of the underlying process can be built starting from S_h^+ and S_h^- and defining, for each $k = 1, \dots, N$, the average mass of probability within a window of k steps above and below the main diagonal. More formally, let:

$$P_k = \frac{1}{N} \sum_{i=1}^{N} \sum_{j \in J_{i,k}} m_{i,j}, \tag{1}$$

where $J_{i,k} = \{j : i-k+1 \leq j \leq i+k-1, j \geq 1, j \leq N\}$. Notice that $P_N = 1$. Moreover, since $S_1^+ = S_1^- = S_1$, then $P_1 = \frac{1}{N} S_1^+ = \frac{1}{N} S_1^-$ and:

$$P_k = P_1 + \frac{1}{N} \sum_{h=2}^{K} (S_h^+ + S_h^-) \stackrel{def}{=} P_1 + \frac{1}{N} \sum_{h=2}^{K} S_h. \tag{2}$$

Note also that since for $k < N$:

$$P_{k+1} = P_k + \frac{1}{N} \sum_{j=1}^{N-K} (m_{j,k+j} + m_{k+j,j}), \tag{3}$$

then P_k is not decreasing in k.

For each given k, one can therefore assess the persistence of the process and plot persistence tents as done e.g. in Bartelsman and Dhrymes (1994). Here, however, we want to derive a "synthetic" index of persistence (in line with Prais, 1955; Bartholomew, 1973; Geweke et al., 1986). The simplest choice is to compute the arithmetic mean of all P_k's, that is:

$$\Pi(M; N) = \frac{1}{N} \sum_{k=1}^{N} P_k. \tag{4}$$

Straightforward computations show that Π can be also written as a weighted arithmetic average of probability masses contained in h-upper and h-lower diagonals of M, that is:

$$\Pi(M; N) = \sum_{h=1}^{N} \frac{N-h+1}{N^2} S_h. \tag{5}$$

Weights decrease as we move away from the main diagonal. This can be also seen by re-arranging the terms in (5). One indeed gets:

$$\Pi(M;N) = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}(N-|i-j|)m_{i,j}, \tag{6}$$

or, equivalently, exploiting the fact that $\sum_j m_{i,j} = 1$, $\forall i$:

$$\Pi(M;N) = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N-1}(N-|i-j|)m_{i,j} - \frac{N(N+1)}{2N^2}. \tag{7}$$

From a graphical perspective, Π can be shown to be proportional to the area of the surface lying below the curve P_k when plotted against k. Intuitively, the higher the persistence displayed by the underlying process, the larger the probability mass accounted by the main diagonal, the lower the mass lying in h-upper and h-lower diagonals for $h \geq 2$, the flatter the P_k curve, and the higher Π. More precisely, if the underlying process displays maximum persistence (i.e. $m_{i,i} = 1\ \forall i$), then $P_1 = 1$ and $S_h^+ = S_h^- = 0$ for all $h \geq 2$. Hence $P_1 = 1$, the curve stays flat for $k \geq 2$, and $\Pi_{max} = \Pi = 1$. Conversely, in the case of minimum persistence, one has $m_{i,i} = 0$, $\forall i$, $m_{i,N} = 1$ for $i \leq k \leq N/2$ and $m_{i,1} = 1$ for $N/2 < k \leq N$. In that case, one gets:

$$\Pi_{min}(N) = \Pi(M_{min};N) = \sum_{h=N/2+1}^{N} 2\cdot\frac{N-h+1}{N^2} = \frac{N+2}{4N}. \tag{8}$$

Therefore $\Pi(M;N) \in [\Pi_{min}(N), 1]$. Note that $\Pi_{min}(N)$ approaches $\frac{1}{4}$ as the number of classes of the Markov chain goes to infinity. $\Pi_{min}(N)$ can then be employed to appropriately re-normalize Π so as to have an index ranging in the unit interval. Another benchmark case is the "perfect mobility" setup, where all transition probabilities are the same (Prais, 1955; Bartholomew, 1973), i.e. $M = M_{PM} = N^{-1}\cdot\mathbf{I}_N$. In that situation, simple computations show that $\Pi(M_{PM};N) = \frac{2N^2+1}{3N^2}$. Note that $\Pi(M_{PM};N)$ is not the midpoint between $\Pi_{min}(N)$ and 1, and tends to 2/3 as $N \to \infty$.

From (6) and (7), it is easy to see that Π is equivalent to a particular case of (the complement to 1 of) the D index of mobility introduced by Bartholomew (1973, p.19, eq. 2.8), see also Prais (1955) and Matras (1961). More precisely, Π can be obtained from D by setting to N^{-1} all equilibrium (ergodic) probabilities, i.e. when the ergodic distribution is replaced in D by the "perfect mobility" case. As a result, we can interpret Π as a particular case of D when there is no bias in the way one weights the classes, or equivalently when the computation of the ergodic limit can be avoided or just disregarded (possibly because of a slow convergence rate).[2]

[2]Being a particular case of Bartholomew's D index, $\Pi(M;N)$ satisfies some interesting properties (Shorrocks, 1976), such as "immobility" ($\Pi(M;N) \leq 1$) and "strong immobility" ($\Pi(M;N) < 1$ unless $M = \mathbf{I}_N$). Note that, differently from D, it also satisfies "monotonicity" (i.e., the index decreases if one of the non-diagonal entries of M increases at the expense of diagonal terms) precisely because $\Pi(M;N)$ does not depend on the ergodic distribution.

3 Statistical Properties

In this Section we study some statistical properties of the $\Pi(M;N)$ persistence index, under suitable assumptions on the underlying randomness of the process.

Let us start from the hypothesis that we are only given a TMP M and we want to assess the extent to which the observed level of $\Pi(M;N)$ is far from the "perfect mobility" case. An extreme benchmark statistical model in that situation is a purely-random one where each $m_{i,j}$ is assumed to be drawn, independently for any i (i.e., row of the matrix), from a N-dimensional uniform distribution defined on the N-simplex. This ensures that each $m_{i,j}$, for each given i, and independently across different i's, is uniformly distributed in $[0,1]$ and each row of M is a probability distribution. In that case, of course, $E[m_{i,j}|N] = N^{-1}$. By plugging that into (6), one gets $E[\Pi;N] = \frac{2N^2+1}{3N^2}$, exactly as in the "perfect mobility" benchmark. Computing the standard deviation of Π is not straightforward: despite the index in this case is just a weighted sum of uniformly-distributed random variables, $m_{i,j}$ are not independent across j for any i. We can therefore resort to simulations in order to estimate $\sigma[\Pi|N]$. Table 1 reports the expected value and the standard deviation of Π for different Ns, computed over 100,000 independent Montecarlo samples. The values of N are chosen so as to match quantile classes, which are typically employed to define TPMs. Notice that $E[\Pi|N]$ decreases with N as N^{-2} and converges to $2/3$ as N increases. Simple computations show that also the standard deviation decreases with N, following $\sigma[\Pi|N] \cong 1/(5 * N^{0.98})$.

The last column of Table 1 also reports the p-value for the Anderson-Darling test for normality of Π (Anderson and Darling, 1954). The high p-values indicate that we cannot reject the hypothesis that Π is normally distributed, under the hypothesis that each row of the TPM is uniformly distributed over the N-simplex, independently of all other rows. In this case, this result is relatively straightforward, as Π is just a linear combination of independent random variables with bounded support. Indeed, the r.v.'s:

$$n_i(N,M) = \sum_{j=1}^{N-1} (N - |i-j|) m_{i,j} \tag{9}$$

are independent across i and are built as a linear combination of correlated r.v.'s uniformly distributed over the N-simplex. As a result, we should expect normality also for small N's. As we will see below this is a more general result, holding also when $m_{i,j}$ is estimated from the data. Normality of Π allows us to easily compute confidence bands to evaluate whether the observed persistence value is statistically different from the expected one under perfect mobility and uniformly-distributed TPM entries (see Fields, 2006, for an alternative procedure).

In more general settings, we do not simply face a given TPM, but we need to estimate it from the data. Suppose we are given the data set $Y_{H,T} = \{y_{h,t}, h = 1,\ldots,H, t = 1,\ldots,T\}$, recording the level of the relevant variable $y_{h,t}$ for unit h and time t, and that we want to estimate the persistence of the underlying generating process. No matter if the data are discrete or continuous, assume that we can reasonably approximate the process generating the data with a time-homogenous Markov chain defined over N classes (cf. Anderson and Goodman, 1957; Bickenbach and Bode, 2001, for statistical procedures testing for time-homogeneity in Markov processes). If the data are discrete, the classes are simply the values attained by the $y_{h,t}$'s. If the data are continuous, then one can use the quantiles of each time-t distribution as reference classes to build a proper $N \times N$ TPM. In that case, $m_{i,j}$ represents the probability that $y_{h,t}$ belongs to the j-th quantile of the distribution of y at time t given that it belonged to the i-th quantile of the distribution of y at time $t-1$. In either case, one can compute maximum-likelihood (ML) estimates $\hat{m}_{i,j}$ for $m_{i,j}$ following Anderson and Goodman (1957). In their seminal paper, they also show that $\hat{m}_{i,j}$'s have a limiting joint normal distribution with given covariance matrix. This means that also the persistence indicator $\Pi(\hat{M}; N)$, i.e. when estimated using the ML estimators $\hat{m}_{i,j}$ for $m_{i,j}$, is expected to be in the limit normally distributed, being a linear combination of limiting jointly normally-distributed r.v.'s.

In the following Section, we shall explore the behavior of $\Pi(\hat{M}; N)$ for two relevant macroeconomic datasets. In the rest of this Section, instead, we perform some Montecarlo simulations to study the statistical properties of Π when the index is computed starting from a continuous-valued process generated by an order-1 auto-regressive (AR) process of the form:

$$y_{h,t} = \beta y_{h,t-1} + \epsilon_t \tag{10}$$

where $\beta \in [0, 1]$ and $\epsilon_t \sim N(0, \sigma)$. This is an obvious benchmark, as it is customary in econometrics to study the persistence of the process using the first-order autocorrelation coefficient as proxied by an estimate of β (see also Tauchen, 1986, for a method to find a discrete Markov chain that approximates AR continuous-valued processes). For a given choice of the number of quantiles N and of the pair (β, σ), we can then generate a sufficiently large sample of statistically-independent datasets $Y_{M,T}$ and study the distribution of Π. In what follows, we set $M = 100$ and $T = 100$, and generate 1,000 datasets (for homogeneous initial conditions $y_{h,0} = 0$). Figure 1 shows estimates for the mean and standard deviation of Π as (β, σ) and N change. First notice that, as expected, Π reaches values close to its maximum when the underlying process is close to a random walk, and then quickly decreases towards 2/3 as β goes down. Conversely, the effect of N and σ on the mean of Π is not dramatic. A sharp threshold emerges instead in the standard deviation of Π for large values of β: when the autoregressive parameter is close to one, then one expects the distribution of Π to be very concentrated around its mean value. For larger values of β, the standard deviation remains instead almost constant.

It must be noted that AD tests (not reported here) confirm normality of Π Montecarlo distributions in all AR parameter setups and for all N. This implies that we can easily build confidence bands for Π and use them to evaluate the extent to which the empirically-observed values of Π are statistically-different from (or close to) a given AR benchmark (see Section 4). Of course, one can replicate the foregoing analysis for any class of underlying generating process that can serve as a null hypothesis (e.g., ARIMA processes) and, consequently, derive Montecarlo estimates for $E(\Pi; \cdot)$ and $\sigma(\Pi; \cdot)$.

4 Applications

We now exemplify the use of the index Π with some empirical applications. We employ two macroeconomics datasets. The first one is Penn World Tables (PWT) 6.1, from which we extract data for (real) country GDP —both in absolute terms and per-capita— from 1960 to 2007 (48 years). The second one is the Expanded-Trade dataset maintained by Kristian Gleditsch (2002), from which we get aggregate bilateral exports (in real terms) between world countries from 1981 to 2000 (20 years). For both data sources, we build a panel of $C = 110$ countries (see Table 2 for the complete list). We are interested in assessing the persistence of four macro variables: (i) (real) country per-capita GDP (YPC); (ii) (real) country GDP growth rates (YGR), defined as the difference between the logs of country GDP in two consecutive years; (iii) absolute trade imbalance (ATI) between any pair of countries (i, j) in the panel, defined as the absolute level of the difference between exports from i to j and exports from j to i; (iv) relative trade imbalance (RTI) between any pair of countries (i, j) in the panel, defined as the ratio between the correspondent ATI and the sum of total trade between i and j (i.e. exports from i to j plus exports from j to i).[3]. While the persistence of real GDP and their growth rates have been widely studied using many alternative techniques (see, among many others, Quah, 1993; Fatas, 2000; Chan et al., 2001), assessing the persistence of trade imbalances can shed some new light on the current international-trade debate (Claessens et al., 2010).

We therefore start from four data matrices: (i) $YPC_{c,t}$, where $c = \{1, \dots, C\}$ and $t = \{1960, \dots, 2007\}$; (ii) $YGR_{c,t}$, where $c = \{1, \dots, C\}$ and $t = \{1960, \dots, 2007\}$; (iii) $ATI_{i,t}$, where $i = \{1, \dots, C \cdot (C-1)/2\}$ and $t = \{1980, \dots, 2000\}$; and (iv) $RTI_{i,t}$, where $i = \{1, \dots, C \cdot (C-1)/2\}$ and $t = \{1980, \dots, 2000\}$. Note that i spans in $\{1, \dots, C \cdot (C-1)/2\}$ because one has to build all bilateral (symmetric) trade imbalances between C countries. We now employ the index Π to assess the extent to which each of these four processes are persistent, i.e. whether a unit (i.e. a country or a pair of countries) belonging to the k-th quantile of year-t distribution is likely to belong to the k-th or a nearby quantile in year $t+1$. We consider two choices for the number of quantiles ($N = 5, 10$). In each case, we compute the observed Π as in eq. (5), where $m_{i,j}$ are replaced by their ML estimate $\hat{m}_{i,j}$ (Anderson and Goodman, 1957).

Table 3 reports our main results. As expected, per-capita country GDP displays the highest persistence (observed values very close to 1, i.e. the maximum of Π). Trade imbalances also display a rather high persistence, whereas GDP growth rates are the less persistent variable. As a first comparison, for each dataset and $N = 5, 10$, we simulate a (time-homogeneous) Markov chain with TPM exactly equal to the observed one (estimated via ML). It can be seen that the expected persistence is almost equal to the observed one. What is interesting to note are the extremely narrow 95% Montecarlo bands.[4] This means that if the underlying process can be discretize and approximated by a time-homogeneous Markov chain, its expected persistence would fall 95% of the times very close to the empirically-observed value. The third row of Table 3 reports instead "perfect mobility" figures. Notice that all observed values, but those of GDP growth rates, fall definitely outside the correspondent 95% bands, meaning that all processes are statistically more persistent than in the "perfectly-mobile" case. This is not so for GDP growth rates, when approximated using quintiles ($N = 5$), which cannot be discriminated from our benchmark. In the $N = 10$ case,

[3] We set RTI=0 when total trade between i and j is zero

[4] Montecarlo confidence bands at 95% are computed using normality as $[\overline{\Pi} - 2s(\Pi), \overline{\Pi} + 2s(\Pi)]$, where $\overline{\Pi}$ is the Montecarlo mean of Π and $s(\Pi)$ is the Montecarlo standard deviation of Π.

however, the observed Π value falls outside the 95% "perfect mobility" confidence band, albeit probably would have fallen inside the 90% one. To further check the extent to which observed Π are statistically different from some alternative benchmark random model, we can perform a "reshuffling" exercise. Given any database $Y_{c,t}$, $c = \{1, \ldots, C\}$ and $t = \{1, \ldots, T\}$, we reshuffle a sufficiently large number of times the columns of each row of the matrix (independently across rows) to get, in each simulated instance, a new bootstrapped matrix where we destroy time dependence. Results are shown in the fourth row of the table. It is easy to see that observed persistence is statistically higher than in the reshuffled case, which exhibits quite narrow confidence bands. Finally, we employ simulated values for Π under AR processes to infer what is the most likely common AR model, i.e. the one that better matches observed values (see Figure 1). To find the most likely values for (β, σ) (for any given N), we simply find those (β, σ) delivering 95% bands that include the observed value. Of course, those values need not to be unique, and in those cases we simply report the most likely range of parameters. The last row of Table 3 shows that YPC series are consistent with a persistence level similar to that generated by a quasi random-walk model (with relatively small variance). Absolute and relative trade imbalances exhibit a medium persistence with $\beta \in [0.5, 0.7]$, whereas GDP growth rates display the lowest persistence, equivalent to a common AR process with $\beta = 0.3$.

5 Concluding Remarks

In this note, we have explored the statistical properties of a simple index of persistence. We have shown that the index is normally distributed both when the entries of the TPM are uniformly distributed and when they are estimated using some underlying sample of data. We have applied the index to study the persistence of four macroeconomic datasets for a panel of 110 countries and compared observed values to different statistical benchmarks. This may prove to be helpful to statistically discriminate between high and low levels of the index, according to the maintained null hypothesis.

Many extensions of the simple Π index studied here may be envisaged. Starting from Eq. (5), one can generalize the index by letting:

$$\Pi_G(M; N) = \sum_{h=1}^{N} f_h(N) S_h \tag{11}$$

where $f_h(N) \geq 0$ is a different weighting function satisfying $f_1(N) + 2\sum_{h=2}^{N} f_h(N) = 1$. Some straightforward examples are $f_h(N) = c > 0$ or $f_h(N) = (N - h + 1)^{-1}$. Similarly, an extension of Π that can be computed on transition probability kernels in the continuum might be considered (see Quah, 1993).

Finally, more extensive simulation exercises may be carried out to map the statistical properties of Π under alternative data generating processes, such as for example ARMA or ARIMA models.

References

Anderson, T. and D. Darling (1954), "A test for goodness-of-fit", *Journal of the American Statistical Association*, 49: 765–769.

Anderson, T. and A. Goodman (1957), "Statistical Inference about Markov Chains", *The Annals of Mathematical Statistics*, 28: 89–110.

Bartelsman, E. and P. Dhrymes (1994), "Productivity dynamics: U.S. manufacturing plants, 1972–1986", Finance and Economics Discussion Series 94-1, Board of Governors of the Federal Reserve System.

Bartholomew, D. (1973), *Stochastic Models for Social Processes*, John Wiley and Sons, London.

Bickenbach, F. and E. Bode (2001), "Markov or Not Markov This Should Be a Question", Kiel Working Paper 1086, University of Kiel.

Cefis, E. (2003), "Is there persistence in innovative activities?", *International Journal of Industrial Organization*, 21: 489–515.

Chan, L. K., J. Karceski and J. Lakonishok (2001), "The Level and Persistence of Growth Rates", NBER Working Papers 8282, National Bureau of Economic Research, Inc.

Claessens, S., S. Evenett and B. Hoekman (eds.) (2010), *Rebalancing the Global Economy: A Primer for Policymaking*, Centre for Economic Policy Research (CEPR), London.

David, P. and F. Rullani (2008), "Dynamics of innovation in an open source collaboration environment: lurking, laboring, and launching FLOSS projects on SourceForge", *Industrial and Corporate Change*, 17: 647–710.

Ezcurra, R., P. Pascual and M. Rapun (2006), "The Dynamics of Industrial Concentration in the Regions of the European Union", *Growth and Change*, 37: 200–229.

Fatas, A. (2000), "Do Business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth", *Journal of Economic Growth*, 5: 147–62.

Fiaschi, D. and A. Lavezzi (2007), "Nonlinear Economic Growth: Some Theory and Cross-Country Evidence", *Journal of Development Economics*, 84: 271–290.

Fields, G. (2006), "The many facets of economic mobility", ILR school site, Cornell university, Retrieved at: `http://digitalcommons.ilr.cornell.edu/articles/230/`, Cornell University.

Geweke, J., R. C. Marshall and G. A. Zarkin (1986), "Mobility Indices in Continuous Time Markov Chains", *Econometrica*, 54: 1407–1423.

Gleditsch, K. (2002), "Expanded Trade and GDP data", *Journal of Conflict Resolution*, 46: 712–24.

Jafry, Y. and T. Schuermann (2004), "Measurement, estimation and comparison of credit migration matrices", *Journal of Banking and Finance*, 28: 2603–2639.

Jarvis, S. and S. Jenkins (1998), "How Much Income Mobility is There in Britain?", *Economic Journal*, 108: 428–443.

Konings, J. and F. Roodhooft (1997), "Financial Ratio Cross-section Dynamics: A Non-Parametric Approach", *Journal of Business Finance and Accounting*, 24: 1331–1342.

Matras, J. (1961), "Differential Fertility, Intergenerational Occupational Mobility, and Change in the Occupational Distribution: Some Elementary Interrelationships", *Population Studies*, 15: 187–197.

Prais, S. (1955), "Measuring Social Mobility", *Journal of the Royal Statistical Society A*, 118: 56–66.

Quah, D. (1993), "Empirical cross-section dynamics in economic growth", *European Economic Review*, 37: 426–434.

Shorrocks, A. (1976), "Income mobility and the Markov assumption", *The Economic Journal*, 86: 566–578.

Tauchen, G. (1986), "Finite State Markov-chain Approximations to Univariate and Vector Auto-Regressions", *Economics Letters*, 20: 177–181.

Zhang, J., V. Avasarala and R. Subbu (2010), "Evolutionary optimization of transition probability matrices for credit decision-making", *European Journal of Operational Research*, 200: 557–567.

N	Mean	Std Dev	AD Test (pval)
4	0.6877	0.0522	0.8752
5	0.6801	0.0420	0.7065
10	0.6699	0.0215	0.7172
20	0.6676	0.0109	0.1815
50	0.6668	0.0044	0.4151
100	0.6667	0.0022	0.2376

Table 1: Mean and standard deviation of Π under the null hypothesis that each row of the TPM is uniformly distributed over the N-simplex, independently of all other rows. *Notes:* Mean and standard deviation computed over 100,000 Montecarlo replications. AD Test: Anderson-Darling test for normality (Anderson and Darling, 1954) of Montecarlo samples.

Algeria	Dem. Rep. of Congo	Haiti	Mauritius	Senegal
Argentina	Republic of Congo	Honduras	Mexico	Seychelles
Australia	Costa Rica	Hong Kong	Morocco	Singapore
Austria	Cote d'Ivoire	Iceland	Mozambique	South Africa
Bangladesh	Cyprus	India	Namibia	Spain
Barbados	Denmark	Indonesia	Nepal	Sri Lanka
Belgium	Dominican Republic	Iran	Netherlands	Sweden
Benin	Ecuador	Ireland	New Zealand	Switzerland
Bolivia	Egypt	Israel	Nicaragua	Syria
Botswana	El Salvador	Italy	Niger	Taiwan
Brazil	Equatorial Guinea	Jamaica	Nigeria	Tanzania
Burkina Faso	Ethiopia	Japan	Norway	Thailand
Burundi	Fiji	Jordan	Pakistan	Togo
Cameroon	Finland	Kenya	Panama	Trinidad &Tobago
Canada	France	Rep. of Korea	Papua New Guinea	Turkey
Cape Verde	Gabon	Lesotho	Paraguay	Uganda
Central African Rep.	Gambia	Luxembourg	Peru	United Kingdom
Chad	Ghana	Madagascar	Philippines	United States
Chile	Greece	Malawi	Portugal	Uruguay
China	Guatemala	Malaysia	Puerto Rico	Venezuela
Colombia	Guinea	Mali	Romania	Zambia
Comoros	Guinea-Bissau	Mauritania	Rwanda	Zimbabwe

Table 2: Countries in the panel.

	Per Capita GDP		GDP Growth Rates	
	N=5	N=10	N=5	N=10
Observed	0.9913	0.9888	0.7454	0.7371
Markov Chain	0.9912	0.9887	0.7453	0.7368
	(0.9900,0.9924)	(0.9878,0.9896)	(0.7390,0.7516)	(0.7304,0.7432)
Perfect Mobility	0.6800	0.6700	0.6800	0.6700
	(0.5955,0.7645)	(0.6271,0.7129)	(0.5955,0.7645)	(0.6271,0.7129)
Reshuffling	0.9217	0.9131	0.7016	0.6920
	(0.9192,0.9243)	(0.9108,0.9153)	(0.6957,0.7076)	(0.6862,0.6978)
AR Models	$\beta \in [0.99, 1]$	$\beta \in [0.99, 1]$	$\beta = 0.3$	$\beta = 0.3$
	$\sigma \in [0.01, 0.1]$	$\sigma \in [0.01, 0.1]$	$\sigma \in [0.1, 10]$	$\sigma = 0.1$

	Absolute Trade Imbalances		Relative Trade Imbalances	
	N=5	N=10	N=5	N=10
Observed	0.9321	0.8416	0.8590	0.7799
Markov Chain	0.9321	0.8416	0.8590	0.7798
	(0.9313,0.9329)	(0.8411,0.8422)	(0.8575,0.8605)	(0.7787,0.7810)
Perfect Mobility	0.6800	0.6700	0.6800	0.6700
	(0.5955,0.7645)	(0.6271,0.7129)	(0.5955,0.7645)	(0.6271,0.7129)
Reshuffling	0.8821	0.8010	0.6960	0.6857
	(0.8780,0.8863)	(0.7996,0.8025)	(0.6945,0.6975)	(0.6840,0.6874)
AR Models	$\beta = 0.7$	$\beta = 0.7$	$\beta = 0.5$	$\beta = 0.5$
	$\sigma = 0.01$	$\sigma = 0.01$	$\sigma \in [0.01, 0.1]$	$\sigma \in [0.01, 0.1]$

Table 3: The index Π and its statistical properties under different null hypotheses for four macroeconomic datasets. Notes: N=number of quantiles. 95% confidence bands in parentheses. Montecarlo sample size: 10,000.

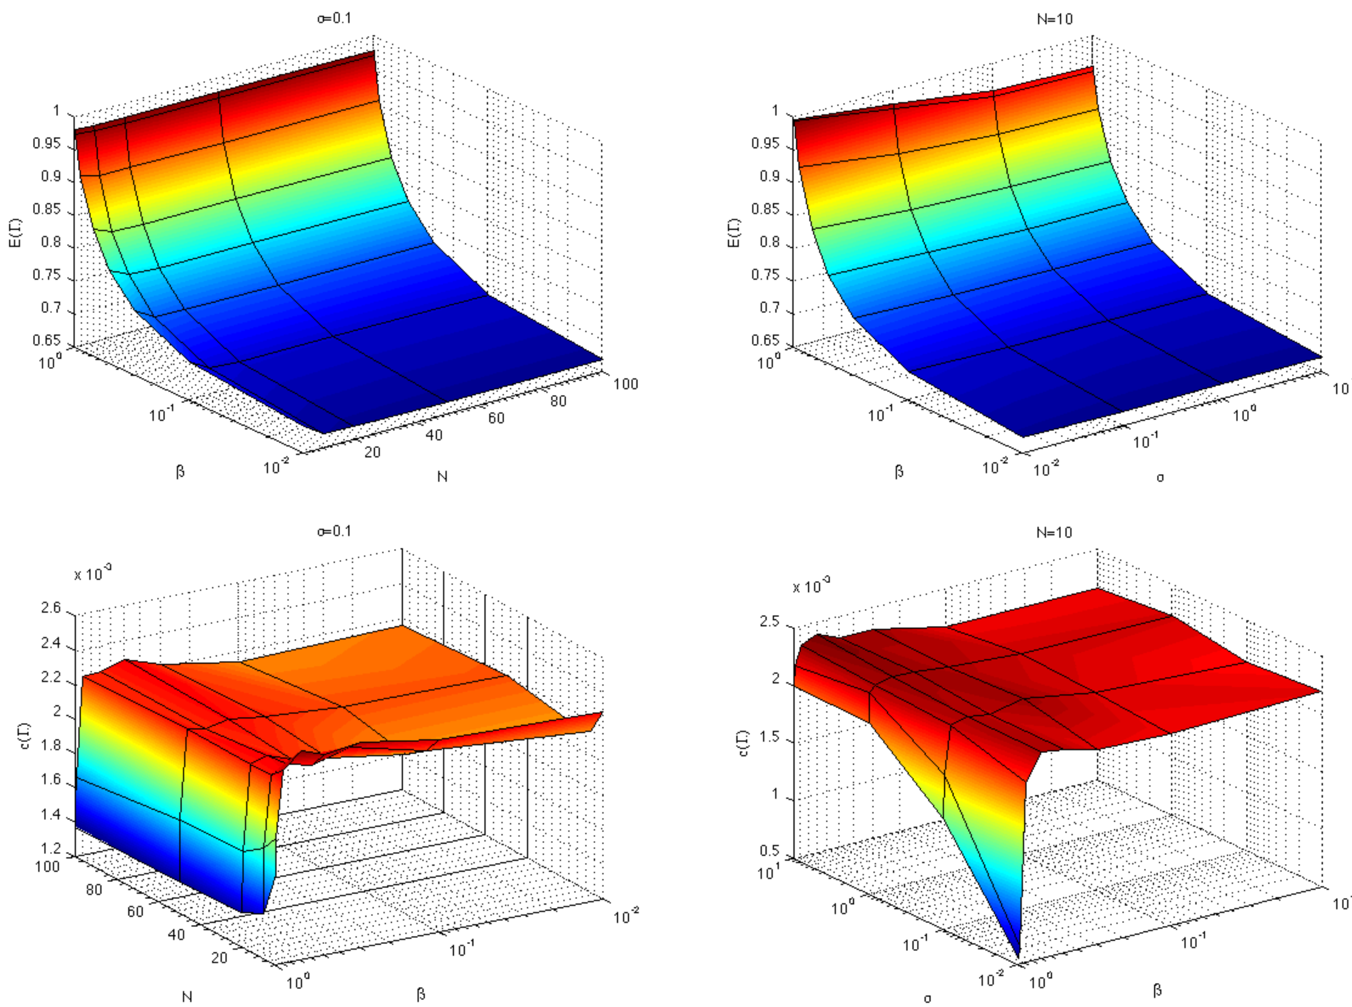

Figure 1: Mean and Standard Deviation of Π under the hypothesis that the underlying data generation process is an AR(1). Top Panels: Mean. Bottom Panels: Standard Deviation. Moments estimated using 1,000 datasets with 100 units and 100 time-periods. Note: Logs scales are employed for both X and Y axes.