



Are correlations of stock returns justified by subsequent changes in national outputs?

Bernard Dumas^a, Campbell R. Harvey^b, Pierre Ruiz^{c,*}

^a *INSEAD, 77305 Fontainebleau cedex, France*

^b *Duke University, Durham, NC 27708, USA*

^c *McGill University, Montreal, Quebec H3A 1G5, Canada*

Abstract

In an integrated world capital market, the same pricing kernel is applicable to all securities. We apply this idea to the stock returns of different countries. We investigate the underlying determinants of cross-country stock return correlations. First, we determine, for a given, measured degree of commonality of country outputs, what should be the degree of correlation of national stock returns. We propose a framework that contains a statistical model for output and an intertemporal financial market model for stock returns. We then attempt to match the correlations generated by the model with measured correlations. Our results show that under the hypothesis of market segmentation, the model correlations are much smaller than observed correlations. However, assuming world markets are integrated, our model correlations can be matched with observed correlations.

© 2003 Elsevier Ltd. All rights reserved.

Keywords: Business cycles; Stock markets; Correlation; Integration; Segmentation

1. Introduction

Why are the correlations of stock market returns between, say, the US and UK higher than between the US and Japan, but lower than between the US and Canada? Are stock markets correlations too high to be justified by the (lower) level of correlations between countries' economic fundamentals, such as output or consumption? Do stock market correlations tell us anything about the degree of international financial-market integration? In this paper, we use an equilibrium model

* Corresponding author. Tel.: +1-514-398-5394.

E-mail address: pierre.ruiz@mcgill.ca (P. Ruiz).

linking stock markets correlations to countries' output correlations to shed some light on these three important questions.

The first question really asks whether stock markets correlations across countries can be explained by economic fundamentals. Two main types of answers can be found in the literature. First, several researchers have examined empirically the determinants of international stock markets correlations (among the more recent investigations, see Longin and Solnik, 1995; Erb et al., 1994; Karolyi and Stulz, 1996; Ang and Bekaert, 2002). However, without an underlying economic model, it is difficult to understand why and how these determinants are relevant in explaining correlations.¹ Second, Hamao et al. (1990) and the many papers that followed this work study the spillover of information from one economy to another. While these studies are important in tracing the type of information that causes common movement in expected returns and volatility, they do not give us a starting point. That is, they do not tell us why stock markets correlations are different across countries in the first place. Proposing a theoretical explanation of the differences in the observed level of international stock markets correlations is the first goal of our paper.

The second question we address is motivated by Shiller's (1981) findings. Shiller shows with a simple asset pricing model that in the US "stock prices are too volatile to be justified by subsequent changes in dividends". This is confirmed by Campbell (1996) for a number of other countries. The "excess-volatility" puzzle is commonly ascribed to an "excessive" degree of volatility of the pricing kernel, where "excessive" is understood relative to the observed degree of volatility of consumption. In an integrated world capital market, the same pricing kernel is applicable to all securities. If the kernel is excessively volatile, this should translate into an equally excessive degree of correlation of world equity returns. That is, equity correlations should be too high to be justified by subsequent changes in national consumptions or outputs. Discovering whether there exists or does not exist an "excess-correlation puzzle" is the second goal of our paper.

Finally, we would like to infer from stock markets correlations the degree of international financial-market integration. But can one directly use stock market correlations as a measure of market integration? For instance, if one finds that correlations have been rising, one is tempted to conclude that financial markets are in the process of gradual integration.² Indeed, we shall demonstrate in the context of our particular model that, other things equal, over the relevant parameter range, correlations of stock returns are larger in an integrated market than they would be in a segmented market. In this statement, however, the "other things equal" caveat is crucial. For instance, one should not draw conclusions about integration from stock correlations if one has not controlled for the degree of correlation of economic fundamentals. The model that we develop can be used to control for "other

¹ For another underlying economic model, see Bansal and Lundblad (2000).

² Bekaert and Harvey (1995) link correlation with the degree of market integration. Freimann (1998) offers an alternative, entirely statistical procedure based on randomization of industrial sector returns, to compare country correlations to what they would have been under integration.

things equal” so that it provides us with a tool to diagnose the degree of integration of financial markets. If the market is segmented, the correlation of world equity returns should be excessively low relative to the model that assumes that markets are completely integrated. Admittedly, one could interpret a rejection (observed correlations lower than model correlations) as either evidence against the model or against the hypothesis of market integration. But we are able to apply the same model construction under the hypothesis of market segmentation and, in this way, provide a meaningful comparison. Drawing inferences about international financial-market integration is the third goal of our paper.

To answer the three questions posed, our basic intuition is simple: if most of the variation in economic activity in two countries is driven by the world business cycle, then the two countries should have high equity correlations. Hence, we combine a statistical model of the business cycle with a log-linear asset pricing framework applied to an exchange economy. After making some choices on the form of the utility function of the representative individual and some distributional assumptions, we are able to determine the model’s implied level of correlation for two countries’ realized rates of return, given the measured commonality in country outputs. The modeling of stock returns by means of an exchange economy is in the tradition of Lucas (1978) and Mehra and Prescott (1985). The assumption that world output equals world consumption is evidently simplistic and, worse yet, prevents us from drawing useful information from the physical investment time series.³ In the international context, however, the empirical behavior of each country’s consumption is, in fact, quite close to the behavior of the country’s production, much more so than it should be in an internationally integrated world.⁴ Furthermore, there is no doubt that the pure-exchange setting is a useful shortcut and that calibration exercises, in this tradition, have produced many insights concerning stock returns. Two of these insights, recognized as the “equity-premium puzzle” and the “excess-volatility puzzle” already mentioned, would actually be sufficient to reject the pure-exchange, representative-agent, form of modeling. In our paper, we ask the reader to suspend the disbelief that arises from these puzzles,⁵ in order to allow us to decide whether or not cross-moments also cause difficulties for this kind of modeling.

Our paper is thus intermediate between a macroeconomic real-business-cycle (RBC) paper and a finance paper.⁶ A real-business-cycle paper would attempt to explain observed facts such as the observation that outputs across countries are more highly correlated than consumptions. No equity returns would be measured. A finance paper would attempt to explain return correlations across countries with asset pricing model and variables like cash flows to equity that may not be directly or contemporaneously related to output.

³ See Cochrane (1991), Restoy and Rockinger (1994).

⁴ See Lewis (1999).

⁵ On that issue, please, refer to the calibration in Section 6.

⁶ See also Canova and de Nicolò (1995).

A purely financial approach of stock returns correlations has been proposed by Ammer and Mei (1996) and Campbell and Mei (1993). They decompose the innovations in stock returns into three or four components: news about future dividends, interest rates, possibly news about exchange rates and news about future excess returns, the latter being calculated as a residual only, as opposed to being determined by a dynamic pricing model. In Ammer and Mei, inferences about rising or falling financial-market integration are made on the basis of the rise or fall in the correlation between news about future excess returns. No theoretical benchmark is provided to indicate whether the observed value of that correlation is compatible with full integration. A similar question had been raised earlier by Shiller (1989) and Beltratti and Shiller (1993). Their asset pricing model was a “present-value” model with unspecified stochastic discount rates which were used to generate upper and lower bounds on the values of correlations. The upper and lower bounds were quite far apart. In our paper, the valuation equation is derived from optimal portfolio choices based on an explicit stochastic process for the payoff series.

Another purely financial approach uses the framework of a partial equilibrium model. Carriero et al. (2001) build on an asset pricing model derived under a known type of segmentation, the mild-segmentation, whereby some securities are restricted to some investors. It allows them to compute an “integration index” which captures the degree to which securities known to be inaccessible can be proxied by accessible securities. But the calculation of their index requires the specification of the set of securities which is not accessible to all investors. Our model, on the other hand, fully endogenizes stock returns in a general equilibrium framework so they only depend on the economic activity.

There have also been several previous attempts to understand the interplay between the real economy and stock returns. They have been mostly focused on the determination of conditional first moments of returns. A number of empirical papers (see, for example, Fama, 1990; Schwert, 1990; Choi et al., 1999) show that there is a relation between expected output and stock returns. Asset pricing tests offer another possible route of examination. These tests specify common factors which each country has sensitivity to (see Ferson and Harvey, 1993; Cheung et al., 1997). One can deduce from the estimated sensitivities to the common factors what the correlation of equity returns should be. Correlations are determined by a statistical model that captures the relative movement of each country’s return vs. some global benchmarks. They are not endogenized. The associated pricing model is not solved over time; it only provides expected returns from exogenously measured covariances. In principle, the resulting expected returns can be used to test the null hypotheses of integration or segmentation, as in Jorion and Schwartz (1986). Most often, however, expected returns are not estimated with sufficient precision to be able to reject either null hypothesis. We feel that correlation measurements hold the potential of providing more powerful tests of integration.

Our work is finally also related to, but unfortunately does not encompass, those studying time-varying correlations. We aim to understand why correlations are different across countries, not how they differ from year to year. In fact, the correla-

tions are modeled as being constant over time.⁷ Longin and Solnik (1995) show by means of a statistical model, how correlations change through time.⁸ Both Longin and Solnik (1995) and Erb et al. (1994) try empirically to explain, on the basis of economic variables, how correlations vary over time.⁹ The analysis in Erb et al. shows that while there is some time variation in the correlations of the G7 countries' equity returns through time, the ranking of the correlations rarely changes. That is, while there is variation in both the US–UK and US–Japan correlations through time, the US–UK correlation is always higher than the US–Japan correlation. Similarly, the US–Canada correlation is always higher than the US–UK correlation. While it is clear that correlations are not constant, our assumption should not interfere with the main point of our paper, which is to explain why correlations are different across different countries.

The paper is organized as follows. Section 2 explores the data and the phenomena that we are trying to explain. In Section 3, we develop the dynamic single-index model of Stock and Watson (1993) which we will use to define each country's business cycle. The log-linear pricing kernel of Restoy and Weil (1996) is explained in Section 4. Section 5 applies the log-linear pricing kernel to the dynamic single-index business cycle model to derive equilibrium security returns. We then examine, in Section 6, the correlations implied by the model and the actual correlations observed in the data. Section 7 develops a statistical test of the hypothesis of financial-market integration. Some concluding remarks are offered in the final section.

2. A first look at the data

In the empirical analysis below, we focus on the behavior of industrial production and stock returns in 12 OECD countries¹⁰ on a monthly basis from January 1970 to June 1996. The monthly frequency is dictated by the fact that the time period of the sample cannot be extended and that more frequent data do improve the precision of estimates of second moments. Fig. 1 shows a positive relation between the correlation of a country's output with OECD output and the correlation of that country's stock market returns with OECD stock returns. We report a country's correlation with the *other* countries, itself excluded. This is in order to abstract from the effect of a country's size. Had we calculated the corre-

⁷ We also execute calibrations (and, later, statistical tests) using unconditional moment conditions. But that choice is unrelated to the assumptions of the model since any model can be tested on the basis of its unconditional predictions.

⁸ Hodrick (1989) derives the multivariate GARCH process followed by stock prices when dividends themselves follow a multivariate GARCH process.

⁹ For instance, non-US stock returns tend to have a higher correlation with US stock returns while the US is in a recession than while it is in an expansion. Volatility of returns is also larger while the US is in a recession. See also Perez-Quiros and Timmermann (1996) and Ang and Bekaert (2002).

¹⁰ Austria, Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Spain, Sweden, UK and the US.

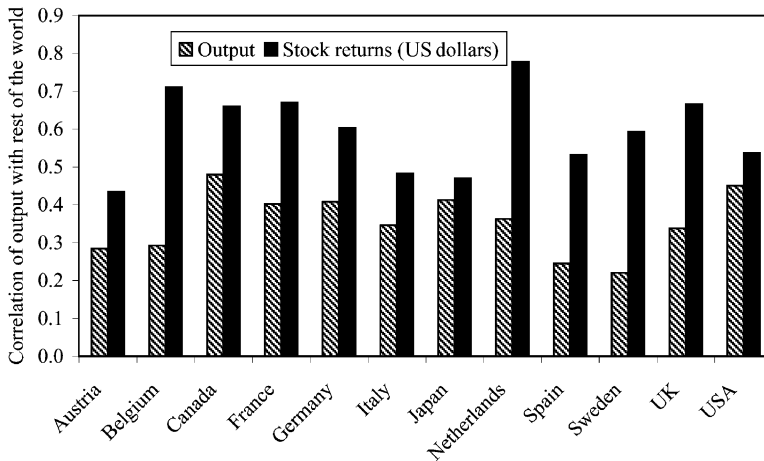


Fig. 1. Industrial production and stock returns correlations. Correlations of each country's output growth with the rest of the world and correlations of each country's equity returns (in U.S. dollars) with the rest of the world. Monthly data from January 1970 to June 1996.

lation of a country with the world, including the country itself, larger countries would automatically have shown a larger correlation. The index of the other countries' output is calculated with annually updated GDP weights.¹¹

The work in our paper is calibrated to industrial production growth. To be frank, the overriding reason for which we chose industrial production was a practical one. It was the only measure of economic activity that was available on a monthly basis for all the 12 countries in our sample. As will become clear later, in the context of the model, production proxies for the cash flows that are generated by stock securities. Furthermore, current and past production is the information variable that investors in the model use in their investment decision, as a way of predicating their decisions on the stage of the business cycle the economy is in.

There is a severe drawback to using output as a proxy for the cash flows generated by equity securities. In several countries, many of the companies listed in the stock exchange typically have levels of foreign activities markedly larger than the share of exports in the corresponding output series. To assess the extent of the problem, we measured the percentage of foreign sales for the companies in each of the 12 countries in 1997. Using the *Worldscope* universe, we constructed country aggregates by value-weighting these ratios by the total revenues of each firm. Belgium, Canada, and the Netherlands have the highest proportions (64.7%,

¹¹ To account for the possibility of lags, the correlations of a country's output with the OECD output is, in fact, the square root of the R^2 in a multiple regression of the country's industrial production on aggregate OECD industrial production (the country itself excluded), contemporaneous and with 11 monthly lags. The data period provides us with 315 monthly returns. Taking into account the 11 lags leaves us with 304 observations.

64.2% and 65.2%, respectively). The same ratio averages only 40.1% for the other countries in our sample. The share of exports in GDP is typically a lower number; in December 1997, it ranges from 7% for the US to 49% in the Netherlands. It is tempting to make a scale adjustment to each country's correlation based on the level of foreign activity, to bring that level down to equal the share of exports in GDP. This adjustment would have to change through time. For example, in 1991, the weighted proportion of foreign sales in Canada was 47.2% and it increases to 64.2% in 1997. But the proportion of sales is an imperfect measure because it only measures one part of earnings—the revenues. We have no information as to the extranational costs of the firms. We choose not to apply a scaling factor of that type. First, we felt that imposing the scaling factors based on these measures would be arbitrary as we do not observe for each country the composition of foreign trade by destination. Second, we were worried about introducing another level of estimation error.

Another possibility that we did not pursue is to use dividends or earnings for the countries that we study, as in [Bansal and Lundblad \(2000\)](#). While there are issues with the macroeconomic data, one encounters a different set of problems using dividends and earnings. Estimates of earnings are only available quarterly. Dividends are paid regularly on stock market indexes but their time pattern over the year may have nothing to do with fluctuations in economic activity. Instead, they have everything to do with conventions concerning the dates of payments, made by each individual firm, of a total yearly payment. This yearly payment has been declared at the end of the previous fiscal year; it contains only stale information and does not capture the latest news about economic activity.¹² In addition, from year to year, managers tend to smooth both dividends and earnings by means of accounting manipulations. While the macroeconomic data are far from ideal, we elected not to use the smoothed financial data.¹³

We must also evaluate industrial production as a proxy for output and for contemporaneous information on the business cycle. To address this issue, we collected real GDP data which were available on a quarterly basis for nine of the 12 countries and on an annual basis in the other three countries. We then calculated for each country their GDP correlations with the rest of the world and compared these to the industrial production correlations (these results are available on request). We found very similar patterns between the GDP and industrial production correlations. Both GDP and industrial production correlations are positively related to equity correlations. There were two countries, Germany and Japan, that had GDP growth correlations with the rest of the world that were far smaller than the industrial production correlations. Nevertheless, there seemed to be a reasonable corre-

¹² As an illustration of these problems, both the dividends and earnings available from Morgan Stanley Capital International are smoothed with a 12-month moving average.

¹³ Aside from the smoothing issue, there are two additional differences between our research and that of Bansal and Lundblad. First, they allow for heteroskedasticity and moving-average processes, which we do not. However, second, they use a static CAPM whereas we use a more appropriate intertemporal framework.

spendence between GDP and industrial production.¹⁴ All in all, industrial production appears to be the only measure of the pulse of the economy, that is contemporaneous with the business cycle and available at a high enough frequency.

The data we use for each country are expressed in US dollars. Randomly fluctuating exchange rates can cause a disconnection of realized returns expressed in local currency since, in theory, they ought to be linked by an equilibrium pricing relationship applicable to returns expressed in a common currency.¹⁵ In Section 6 below, we find that the calibrations we conduct are not markedly affected by the choice of unit. This is a reflection of the fact that stock returns expressed in dollars exhibit approximately the same measured correlations as do stock returns expressed in the respective local currencies.

Consistent with the observations by *Ammer and Mei (1996)*, the correlations of stock returns, in *Fig. 1*, are always higher than those of industrial productions. Our research has a simple objective. We aim to understand the empirical observations contained in *Fig. 1*. Specifically, for a given degree of commonality in real activity growth, our model will predict a level of correlation of market returns. We want to know what mapping connects the correlations of output to the correlations of stock returns.

3. The “dynamic single-index” model

We now describe the model that captures the evolution of the vector of national outputs. This will be the first component of our overall model. It is a purely statistical model of international business cycles. It represents a shortcut for RBC models, such as *Backus et al. (1992)*, which contain (i) a statistical model for productivity shocks and (ii) an explicit representation of the households’ consumptions and work decisions and the firms’ investment and production decisions. In our framework, we postulate a pure-exchange economy in which the dynamics of output is exogenous and people consume the entire output. It is hoped that not much will be lost by this short cut, since it is generally agreed that most of the dynamics in RBC models comes from the exogenous dynamics of productivity shocks and very little comes from the endogenous capital accumulation process.

The statistical model decomposes each country’s industrial output growth into two unobserved components: the “world” business cycle which is common to all and the “country-specific” business cycle. For reasons of parsimony, each of the cycles, whether common or specific, is assumed to follow an autoregressive process of order 2. We assume that the volatilities of the innovations for each cycle are

¹⁴ We also conducted sensitivity analysis on the number of lags and the frequency of measurement of industrial production. We found that the multiple correlation measures produced considerably higher correlations than using two lags or no lags. We also compared quarterly correlations with the monthly correlations and found broad similarity.

¹⁵ On that count, see the empirical results of *Dumas and Solnik (1995)*. Similar results were reached by *Dumas (1994)* who relates the international asset pricing relationship to business conditions.

constant (homoskedasticity) and that the innovations are independent of each other across cycle processes.

Throughout, $\Delta \mathbf{d}_t$ denotes a vector of output log-growth rates of a number of countries. We postulate a dynamic single-index model.¹⁶ The comovements at all leads and lags among the output variables are modeled as arising from a single common source c_t , a scalar unobserved variable that portrays the world business cycle. The idiosyncratic component, \mathbf{u}_t , which is the part not arising from leads and lags of c_t , is assumed to be stationary and uncorrelated across countries. Otherwise, it follows a general autoregressive process. We call this the idiosyncratic component the country-specific business cycle. The statistical representation of the system is:

$$\begin{aligned} \Delta \mathbf{d}_t &= \boldsymbol{\xi} \times c_t + \mathbf{u}_t, \\ \chi(L)c_t &= \eta_t, \\ \mathbf{D}(L)\mathbf{u}_t &= \boldsymbol{\varepsilon}_t \end{aligned} \tag{1}$$

where L is the lag operator, (ε_t, η_t) are serially uncorrelated with a diagonal covariance matrix and $\mathbf{D}(L)$ is diagonal.

The model is formulated in terms of demeaned log-growth rates. In terms of levels, the output series of all the 12 countries should be integrated (this is confirmed at the 10% level by a Dickey–Fuller (1979) test). Had it been written in terms of levels, the model would say that all 12 series are cointegrated with one common trend. This is a testable proposition. The null hypothesis in such a test (see Stock and Watson (1993)) is that the 12 series are not cointegrated. We performed the test, with ambiguous results; at the 10% level, the hypothesis could not be rejected.¹⁷ Had it been rejected, the estimation of the common trend could have been done in level form by cointegration methods; but such was not the case.¹⁸ The log-growth rates, on the other hand, are all stationary at the 1% level, a feature which validates the estimation method we use. One added advantage of this formulation is that the log-linear pricing kernel to be used below (Section 4) directly applies to log-growth rates.

The statistical model (1) is estimated by means of a linear Kalman filter. We use for the purpose the SCOREM algorithm of Raynauld et al. (1993).¹⁹ The program uses a few iterations of the “EM algorithm” (as described by Watson and Engle (1983), Shumway and Stoffer (1982), and Shumway (1988)) which is not very sensitive to initial values, in order to generate a first set of parameter estimates. That

¹⁶ See Sargent and Sims (1977), Geweke (1977), and Singleton (1980). This model is discussed at length in Stock and Watson (1989, 1991, 1993).

¹⁷ Kasa (1992) presents evidence of a single stochastic trend in GNP in a sample of five of the markets that we study.

¹⁸ For a similar approach to modeling the world business cycle and a comparison between estimations in terms of levels or growth rates, see Gregory, Head, and Raynauld (1997).

¹⁹ We are very grateful to Jacques Raynauld who generously provided us with the GAUSS code to run the algorithm, and to René Garcia who pointed out that the Raynauld–Simonato–Sigouin code would be useful to us.

Table 1
Coefficients of the single-index statistical model applied to international data

Country	Loadings on the world business cycle	Country-specific business cycle		
		AR1 coefficient	AR2 coefficient	Variance of innovations
Austria	0.158 (5.303)	-0.582 (-10.309)	-0.225 (-4.020)	0.696 (12.232)
Belgium	0.163 (5.466)	-0.554 (-9.968)	-0.289 (-5.240)	0.712 (12.209)
Canada	0.343 (6.654)	-0.122 (-1.993)	0.069 (1.150)	0.790 (11.459)
France	0.298 (8.072)	-0.494 (-8.017)	-0.233 (-3.843)	0.657 (11.321)
Germany	0.217 (6.917)	-0.553 (-9.838)	-0.312 (-5.574)	0.678 (11.907)
Italy	0.151 (5.202)	-0.586 (-10.488)	-0.273 (-4.932)	0.701 (12.256)
Japan	0.339 (8.260)	-0.463 (-7.296)	-0.139 (-2.188)	0.687 (11.190)
The Netherlands	0.127 (4.911)	-0.675 (-12.161)	-0.309 (-5.625)	0.643 (12.300)
Spain	0.098 (4.640)	-0.804 (-15.237)	-0.396 (-7.541)	0.543 (12.347)
Sweden	0.101 (2.934)	-0.410 (-7.343)	-0.147 (-2.625)	0.841 (12.457)
UK	0.176 (4.178)	-0.221 (-3.850)	-0.103 (-1.805)	0.894 (12.303)
USA	0.407 (6.629)	0.378 (5.167)	-0.017 (-0.247)	0.618 (9.421)
World		World business cycle		
		0.251 (2.866)	0.454 (5.000)	1.000

T-statistics are reported in parentheses beside the coefficients.

first set is then used as initial values for the “scoring algorithm” (described in Engle and Watson (1981)), which is fast and accurate.

The results for the countries in our sample are presented in Table 1. Practically all parameter estimates are significantly different from 0. The autoregressive behavior of the world business cycle is very different from that of the country-specific cycles. The world component is driven by positive coefficients which sum to 0.705 which implies no deterministic cycle but random shocks with a persistent behavior, whereas the country cycles mostly have negative coefficients implying a much more transient (strongly mean reverting) character. The persistent world cycle will play a driving role in the determination of stock returns.

Many alternative specifications of the statistical model could have been considered. For instance, a moving-average component might have been useful in representing the persistence of the world business cycle. The number of lags in the autoregressive specification could be varied.²⁰ A multi-index model, with a regional index for Europe, might be considered. Alternatively, the world business cycle could have been pre-specified as, for instance, a weighted average of country output growth rates. We cannot afford to try all specifications and some of them do not lead to convergence of the algorithm so that no comparison is possible. Furthermore, comparison of models, some of which are not nested, requires the use of somewhat ad hoc goodness-of-fit criteria. Ultimately, what we must demonstrate is our ability to capture almost all the common dynamic variation in output by

²⁰ A specification with three lags produced parameter estimates for the additional lags that were not significant.




Fig. 2. Actual and model simulated correlations of country output with the rest of the world.

means of the single index, while the residuals η_t and ε_t are almost uncorrelated. As a measure of the descriptive quality of the output model, we present Fig. 2 which compares the correlations of each country's output with the rest of the world to correlations obtained by a simulation performed under the assumptions of the model, including the zero-correlation of residuals assumption. The model does a good job of matching these correlations. For a number of countries (Canada, Germany, Italy, the Netherlands), the simulated correlations are somewhat lower than the actual ones. This indicates that our dynamic common factor does not completely capture the correlation between rates of growth in output. Some of it remains in the purely synchronous, static correlation of residuals. But the reader should keep in mind that, in subsequent analysis, we use the actual residuals rather than the simulated ones. This provides an even closer match. The statistical model really only serves to determine what component in the joint behavior of outputs is due to a long-lasting dynamic factor and what component can be captured by the synchronous, static correlation.

4. The log-linear pricing kernel

Restoy and Weil (1996) follow the lead of Campbell (1993) in log-linearizing the budget constraint of a household. They obtain an approximate pricing kernel for multiperiod securities that is based solely on consumption behavior. Their economy consists of many identical, infinitely lived consumers who are endowed with an intertemporal, recursive utility of the Epstein–Zin (1989) or Kreps–Porteus (1978) isoelastic form. This type of utility function allows a distinction, which we find useful (see below), between two behavioral parameters: (i) the relative risk

aversion on the one hand and (ii) the elasticity of intertemporal substitution (e.i.s.) on the other. The latter is a measure of the person's willingness to shift her consumption over time.

Let relative risk aversion be denoted by γ , elasticity of intertemporal substitution be denoted by $1/\rho$ and β denotes the discount factor of utilities, all of which are assumed to be constant and equal for all individuals. Let Δx_{t+1} denote the increment in the logarithm of the households' consumption. Epstein and Zin (1989, 1991) and Weil (1990) have shown that the increment, Δm_{t+1} , of the logarithm of the pricing kernel between time t and time $t + 1$ is given by:

$$\Delta m_{t+1} = \theta \ln \beta - \rho \theta \Delta x_{t+1} + (\theta - 1) r_{W,t+1} \quad (2)$$

where: $\theta = (1 - \gamma)/(1 - \rho)$ and $r_{W,t+1}$ denotes the logarithmic rate of return on aggregate wealth between times t and $t + 1$. This pricing kernel corresponds to an asset pricing model containing two risk premia: one based on the covariance with consumption, the other based on the covariance with wealth.²¹ But recall that, in our pure-exchange economy, consumption is equal to output.

Campbell (1993) and Restoy and Weil (1996) point out, however, that in this expression, Δx_{t+1} and $r_{W,t+1}$ are not independent quantities since wealth equals the present value of consumption: Δx represents changes in output and r_W captures changes in "discounted" future output. It is possible to derive an approximate relationship between these two quantities. Assuming that the households' consumption is one-step-ahead log-normal and conditionally homoskedastic, Restoy and Weil present an expression for the financial-market pricing kernel which does not involve the rate of return on wealth and, in fact, allows returns to be endogenous.²²

Restoy and Weil's (1996) work can be interpreted as meaning that the increment, Δm_{t+1} , in the logarithm of the pricing kernel is given by:

$$\Delta m_{t+1} = \ln \beta - (\rho - \gamma) \frac{1 - \gamma}{2} \text{var}_t[\Delta x_{t+1} + h_{t+1}] - \rho E_t[\Delta x_{t+1}] - \gamma S_{t+1}[\Delta x_{t+1}] + (\rho - \gamma) S_{t+1}[h_{t+1}] \quad (3)$$

where, because of homoskedasticity, the conditional variance $\text{var}_t(\Delta x_{t+1} + h_{t+1})$ is a constant, to be determined on the basis of the stochastic process for consumption, and δ is a linearization constant (equal to 1 minus the exponential of the unconditional expected value of the log-ratio of consumption over wealth) arising in the log-linear approximation to the budget constraint. In addition:

$$h_{t+1} = E_{t+1} \left[\sum_{j=1}^{\infty} \delta^j \Delta x_{t+j+1} \right], \quad (4)$$

²¹ See Epstein and Zin (1989, 1991) and Giovannini and Weil (1989).

²² Whereas Campbell (1993), by the same reasoning, derives an expression for the pricing kernel which does not involve consumption. Eq. (8) below reflects the approximate relationship between changes in output and changes in wealth, that is being used here.

$$S_{t+1}[\Delta x_{t+1}] = \Delta x_{t+1} - E_t[\Delta x_{t+1}], \tag{5}$$

$$S_{t+1}[h_{t+1}] = h_{t+1} - E_t[h_{t+1}]. \tag{6}$$

S_{t+1} is the “surprise” operator.

The above pricing kernel may be used to price any security in an exchange economy in which production and consumption are equal. For instance, the conditional expected value of the pricing kernel provides the one-period riskless rate of interest:

$$r_{f,t} = -\ln\beta + (\rho - \gamma) \frac{1 - \gamma}{2} \text{var}_t[\Delta x_{t+1} + h_{t+1}] + \rho E_t[\Delta x_{t+1}] - \frac{1}{2} \text{var}_t[-\gamma \Delta x_{t+1} + (\rho - \gamma)h_{t+1}]. \tag{7}$$

Because of homoskedasticity, both var_t terms are time invariant. Applying the kernel to an asset that pays aggregate consumption provides a value for the aggregate stock market return. This last task has also been undertaken by Restoy and Weil (1996) who show that:

$$r_{W,t+1} = \mu + \rho \Delta x_{t+1} + (1 - \rho)S_{t+1}[\Delta x_{t+1} + h_{t+1}], \tag{8}$$

where:

$$\mu = -\ln\beta - \frac{(1 - \gamma)(1 - \rho)}{2} \text{var}_t[\Delta x_{t+1} + h_{t+1}]. \tag{9}$$

Similarly, applying the pricing kernel to an asset which pays a dividend $d_{i,t}$ at time t , Restoy and Weil get the equilibrium rate of return on individual assets.

Stock market returns in country i are:

$$r_{i,t+1} = \pi_i + \rho \Delta x_{t+1} + S_{t+1}[\Delta d_{i,t+1} + f_{i,t+1}] - \rho S_{t+1}[\Delta x_{t+1} + h_{i,t+1}], \tag{10}$$

where:²³

$$\pi_i = -\ln\beta + (\rho - \gamma) \frac{1 - \gamma}{2} \text{var}_t[\Delta x_{t+1} + h_{t+1}] - \frac{1}{2} \text{var}_t[(\rho - \gamma)(\Delta x_{t+1} + h_{t+1}) + \Delta d_{i,t+1} + f_{i,t+1} - \rho(\Delta x_{t+1} + h_{i,t+1})], \tag{11}$$

$$f_{i,t+1} = E_{t+1} \left[\sum_{j=1}^{\infty} \delta_i^j \Delta d_{i,t+j+1} \right], \tag{12}$$

$$h_{i,t+1} = E_{t+1} \left[\sum_{j=1}^{\infty} \delta_i^j \Delta x_{t+j+1} \right], \tag{13}$$

²³ Recall that the var_t terms are assumed time invariant.

and δ_i is a Taylor-expansion coefficient arising from the log-linearization of the definition of a rate of return.²⁴ This constant is related to the unconditional expected value of the dividend yield of each security. In Eq. (10), observe the respective roles of country vs. world outputs. Out of the three random terms, two (the first and last one) stand for the current and future behavior of world output whereas only the center term refers to the future behavior of asset i 's specific output stream. The terms related to world output reflect the movement in the pricing kernel applicable to all assets worldwide.

A fascinating result falls out of Eq. (10). Whereas the conditionally expected return on assets ($\pi_i + \rho E_t[\Delta x_{t+1}]$) depends on risk aversion γ , the e.i.s. $1/\rho$, and on the impatience parameter β , the second moments (volatilities and correlations) of the asset return *depend on only one utility parameter: the elasticity of intertemporal substitution*. This result would hold exactly in the case of Epstein–Zin utilities with constant risk aversion and e.i.s., and identically, independently distributed returns (see Epstein (1988)). In the case of our output model, the result holds under the log-linear approximation made by Restoy and Weil. This remarkable property is *the* reason why we have chosen to adopt this type of utility function. It will prove most convenient in what follows.

Here is the intuitive reason for which elasticity of intertemporal substitution governs the price response to shocks. Suppose that output undergoes a positive shock; if consumers are willing to absorb this shock into consumption without further ado, there is no need to adjust asset prices. But, if their e.i.s. is low, they will have to be induced to consume the increased current output by the device of higher market prices of assets relative to current consumption. This intuition is well known. It is identical to the one Lucas (1978) gives on this issue,²⁵ except for the fact that Lucas ascribes to risk aversion the role actually played by the elasticity of intertemporal substitution (while pointing out in a footnote that, in his case of time-additive utility, “the term ‘risk aversion’ is perhaps misleading, since the curvature of [the utility function] also governs the intertemporal substitutability of consumption”).

A multivariate rendition of the same intuition goes as follows. Suppose that output in any given country undergoes a positive shock; if world consumers are willing to absorb this shock into world consumption without further ado, there is no need to adjust asset prices. But, if their e.i.s. is low, they will have to be induced to consume the increased current output by the device of higher market prices of *all* assets relative to current consumption. That is why the e.i.s. (inverse of ρ) plays a crucial role in this model. An increase of the parameter ρ (lower elasticity) implies both more volatility and more covariance across assets. The resulting effect on correlations remains ambiguous at this point. It was asserted in Section 1 that, were the kernel excessively volatile, this can translate into an equally excessive degree of

²⁴ The Campbell–Shiller linearization of return begins with the observation that log return equals log dividend growth minus the log of the current price–dividend ratio plus the log of one plus the next-period price–dividend ratio. That identity is then linearized.

²⁵ Lucas (1978, p. 1439).

correlation of world equity returns. Fig. 3 illustrates that this is true in the context of our model over the relevant range of values of the parameter ρ .

It may seem surprising that the second moments of rates of return do not depend on risk aversion. In fact, the “level” of rates of returns does change when people become more risk averse, but, in a homoskedastic world, it changes by a time invariant amount. As a consequence, the time-series “volatility” of each individual return is not affected by changes in risk aversion. The effect of risk aversion is time invariant in a homoskedastic world.²⁶

One more observation should be made on the basis of Eq. (10). The model induces correlations of stock returns through two channels. There is a common movement in the world interest rate caused by consumption growth, and there is a forecast of movements of future dividends from the common component of world growth. Despite the homoskedasticity assumption, it is *not* true that all the common movements in stock returns come from the interest rate. The second channels still causes ex post equity premia (or excess returns) to fluctuate with a common component (in Section 6 below, we examine the common behavior of excess returns). But it *is* true that conditionally expected excess returns are constant over time so that there is no equity return variability caused by variability in equity risk premia.²⁷

5. The log-linear pricing kernel combined with the dynamic single-index model of output

Rodriguez et al. (2002) specialize (8) to the case in which the growth rate in aggregate consumption is AR(2):

$$(1 - \phi_1 L - \phi_2 L^2) \Delta x_{t+1} = \varepsilon_{t+1} \tag{14}$$

This particular autoregressive process implies that:²⁸

$$S_{t+1}(\Delta x_{t+1} + h_{t+1}) = \frac{1}{1 - \phi_1 \delta - \phi_2 \delta^2} \varepsilon_{t+1} \tag{15}$$

and, therefore:

$$\text{var}_t(\Delta x_{t+1} + h_{t+1}) = \left[\frac{1}{1 - \phi_1 \delta - \phi_2 \delta^2} \right]^2 \text{var}(\varepsilon). \tag{16}$$

Our approach is similarly to apply the pricing kernel to the dynamic single-index model in Eq. (1) and obtain the behavior of individual stock returns where stocks are defined as claims on individual output series. Since we have made the assump-

²⁶ We are grateful to Philippe Weil for a helpful discussion on this point.

²⁷ Heteroskedasticity of output would produce time-varying expected returns which would contribute, to a small extent, to an explanation of the high volatility of stock returns.

²⁸ The result can be extended trivially to an autoregressive process of any order.

tion of an exchange economy, aggregate consumption growth is equal to the weighted sum of output growth rates of individual countries: $\Delta x_{t+1} = \sum_j w_{j,t} \Delta d_{j,t+1}$.

For the dynamic single-index model, the terms of (10) can be particularized as follows:

$$S_{t+1} [\Delta d_{i,t+1} + f_{i,t+1}] = A_{i,i} \eta_{t+1} + B_{i,i} \varepsilon_{i,t+1}, \quad (17)$$

$$S_{t+1} [\Delta x_{t+1} + h_{i,t+1}] = \sum_j [A_{j,i} \eta_{t+1} + B_{j,i} \varepsilon_{j,t+1}], \quad (18)$$

$$\text{var}_t [\Delta x_{t+1} + h_{t+1}] = \text{var} \left\{ \sum_j [A_{j,0} \eta + B_{j,0} \varepsilon_j] \right\}, \quad (19)$$

$$A_{j,i} = \zeta_j \frac{1}{\sum_{s=0}^{\infty} \chi_s \delta_i^s}; \quad A_{j,0} = \zeta_j \frac{1}{\sum_{s=0}^{\infty} \chi_s \delta^s}, \quad (20)$$

$$B_{j,i} = \frac{1}{\sum_{s=0}^{\infty} D_{j,s} \delta_i^s}; \quad B_{j,0} = \frac{1}{\sum_{s=0}^{\infty} D_{j,s} \delta^s}, \quad (21)$$

$$\begin{aligned} & \text{var}_t [(\rho - \gamma)(\Delta x_{t+1} + h_{t+1}) + \Delta d_{i,t+1} + f_{i,t+1} - \rho(\Delta x_{t+1} + h_{i,t+1})] \\ &= \text{var} \left\{ (\rho - \gamma) \sum_j [A_{j,0} \eta + B_{j,0} \varepsilon_j] + A_{i,i} \eta + B_{i,i} \varepsilon_i - \rho \sum_j [A_{j,i} \eta + B_{j,i} \varepsilon_j] \right\}. \end{aligned} \quad (22)$$

Our next goal is to determine whether the second moments of observed stock returns can be matched with those of the theoretical model above.

6. Calibration of the model

The system of Eq. (1), coupled with Eqs. (10)–(13), (17)–(22), provides a strong set of restrictions on the output and stock returns series. The unknown parameters are: those of the dynamic single-index model ζ , ξ , \mathbf{D} , χ , the initial value for the common cycle c_0 , the variances of the residual terms ε and η , the utility parameters β , ρ and γ (and the linearization constants δ , δ_i , which, however, are not in principle independent entities).

While it is possible (but extremely difficult) to estimate the full system of equations, it is not clear how useful such an exercise would be. We already know that a model of an exchange economy, even if coupled with a generalized utility formulation such as the Epstein–Zin utility function, has difficulties accounting for observed mean rates of return on equity in the US. This is the “equity-premium” puzzle of Mehra and Prescott (1985).²⁹ Like Kocherlakota (1996) and Campbell

²⁹ Jorion and Goetzmann (1999) have pointed out that most countries are not like the US and have argued that US equity rates of return presumably represent a repeated sequence of surprises, not to be confused with high expected returns.

and Koo (1997), we use Epstein–Zin utility functions. Previous research shows that the freedom separately to choose the risk aversion and the elasticity of intertemporal substitution of the representative individual allows a somewhat better fit of first moments (equity returns and interest rate) than standard time-additive utilities. Nonetheless, it is clear that, at generally accepted levels of risk aversions, our model would not account for the US equity premium.

In what follows, therefore, we employ the result obtained above (see Eq. (10)) that, to an approximation, the elasticity of intertemporal substitution alone determines second moments, while once second moments have been determined, the two taste parameters jointly determine the first moments. We rely on this argument to separate out the debate on the equity premium and we focus exclusively on second moments.

Further, we know from Shiller (1981) that stock return volatilities tend to be larger in reality than can be explained by a simple pricing model based on dividends. In our model, however, dividends, or even earnings, are not the basis for the determination of stock returns—output is. In the real world, two layers of leverage (operational and financial) should normally tend to magnify the volatility of dividend growth rates relative to that of output growth rates. In our dataset, the average taken across countries of the volatilities of industrial production growth rates is equal to 2.02% per month whereas the average volatility of dividend growth rates is equal to 4.7% per month and the average volatility of earnings is equal to 8.4% per month.³⁰ Without the extensive amount of smoothing that is built into the dividend and earnings series, the increased volatility created by leverage would be even more apparent. Under the simplifying (probably simplistic) assumption that dividends and output are exactly linearly related, leverage magnifies volatilities but leaves correlations unchanged.³¹

Below, we report the correlations and the volatilities produced by our model, knowing, however, that actual volatilities are magnified by leverage, which is unobserved. For this reason, we focus mostly on correlations. The correlations, viewed as moment conditions, provide us with a calibrated version of the model. In a later section (Section 7), we use the generalized method of moments to test the validity of the moment conditions.

6.1. Calibration under integration

The calibration is carried out in a simple way. We have already observed that, according to Eq. (10), the second moments are all dictated by the elasticity of intertemporal substitution of the market participants. Based on the dynamic single-index statistical model, which we estimate in a first stage, we select at a second stage the degree of e.i.s. that will best match the levels of a number of correlations

³⁰ The average of earnings growth rates given here is based on eight countries only. In the US, these numbers are 0.8%, 2.4% and 3.6%, respectively. In all cases, earnings and dividends are measured as 12-month moving averages.

³¹ We are grateful to Huntley Schaller for helping us articulate this distinction.

between stock returns and output. Once that is done, we have pinned down all the parameters of the model.³² We then calculate the equilibrium stock returns for the history of shocks which we have identified statistically and compute their correlations.³³

The range of reasonable values for the coefficient ρ is dictated by measurements of e.i.s. ($=1/\rho$) that have been conducted in the past. Regressing the rate of growth of aggregate consumption on changes in the rate of interest, Hall (1988) finds an elasticity equal to 0.1 which is lower than most previous estimates which range as high as e.i.s. = 1. Epstein and Zin (1991), in their test of the CAPM implied by their preferences and applied to stock returns, find values for the elasticity ranging from 0.2 to 0.8. In short, values for ρ ranging from 1 to 10 seem reasonable.³⁴

Ultimately, three kinds of “moments” will have to be matched: the correlations of a country’s stock return with the country’s own output, the correlations of a country’s stock returns with the rest of the world output and the correlations of a country’s stock returns with the rest of the world stock returns. In the hypothesis tests of Section 7, we evaluate the fit of all three.

But in the calibration, we would like to proceed differently. We choose to calibrate the model (i.e. pick the value of ρ) on the basis of correlations with the rest of the world output. Then, we show how the other two, “overidentifying” moments are matched. Using an objective function which downweights the countries with more volatile correlation estimates and searching for the best fitting value of ρ , we find an optimal value equal to 2.1.³⁵

Fig. 3 then illustrates the calibration trade-offs that we are facing. The figure displays the straight arithmetic average across countries of the theoretical correlations and the average level of the corresponding observed correlations that we try to match.³⁶ The figure confirms that the value of ρ that would best match the average correlation with world output is about 2.1. That value of the parameter ρ explains remarkably well the average level of cross-country stock market correlations. But the figure also shows that it does not explain the within-country correlation with

³² Parameters β and γ are ignored since they play no role in determining correlations. The same is true for the linearization constant δ . For δ_i , we use the mean dividend yields calculated over the entire sample.

³³ A two-stage procedure raises the problem that first-stage parameter values are not known for certain. This problem is explicitly taken into account in the hypothesis test (Section 7) but not in the calibration.

³⁴ In order to carry out this type of analysis, it would be best to examine disaggregated consumption according to social categories (borrowers vs. lenders, old vs. young, employed vs. unemployed, etc.). See Deaton (1992). Attanasio and Weber (1989), placing themselves explicitly in the Epstein–Zin paradigm, and using a single cohort of household, found a value $\rho = 0.514$ or an e.i.s. approximately equal to 2.

³⁵ Standard errors of estimates will be provided in the section containing the hypothesis tests (Section 7). The confidence intervals will prove to be narrow.

³⁶ As mentioned earlier, given that we may not observe in the real-world the exact synchronicity between stock returns and output shocks that is postulated in the model, the correlations in question are not simple correlations. They are actually square roots of the R^2 of multiple regressions of stock returns on contemporaneous, plus 11 lagged, output growth rates.

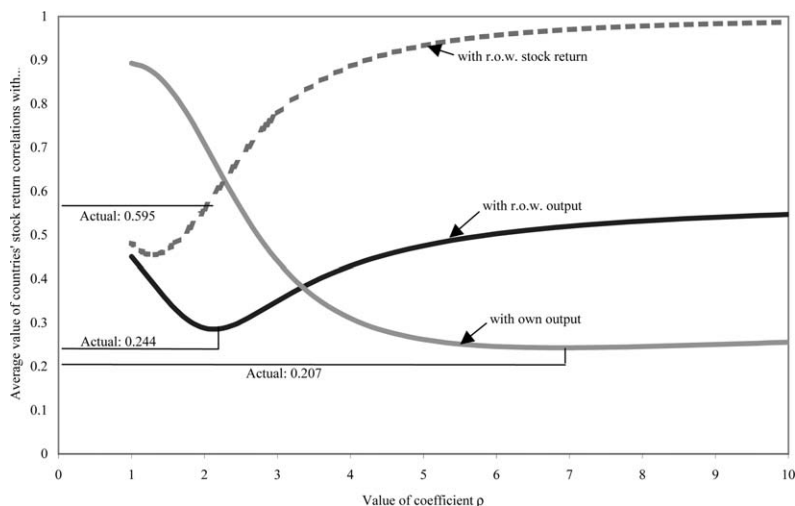


Fig. 3. The calibration trade-off under integration. For each of the correlations, the level marked “Actual” is the target (observed) correlation that is to be matched.

output. In fact, the value of ρ that would best match the average within-country correlation with output is about equal to 6.9.

For each country, Fig. 4 shows the comparison between actual and model stock market correlations resulting from the value $\rho = 2.1$.

There are two factors that play a role in the derivation of model correlations. First, the world pricing kernel, which applies to all securities by construction, has been set in such a way as to match the observed correlations of stock returns with rest of the world output, as we just explained. Second, in our dataset, the world




Fig. 4. Stock returns correlations with the rest of the world under integration.

business cycle (see Table 1) has been found to be fairly persistent. A component of stock returns fluctuations comes from the anticipation of discounted future dividends (see Eq. (10)). If a time series is persistent, any movement occurring today is the harbinger of a lasting movement in future realizations and produces a large immediate effect on returns. In our statistical model, this large component is common to all countries since it originates in the world business cycle.

The model correlations turn out to be of a magnitude similar to the observed ones.³⁷ Fig. 3 shows that, if the pricing kernel were excessively volatile, it would be the case, if the same kernel applied worldwide, that international stock returns would be excessively correlated. However, no such excess correlation appears in our results.

6.2. Discussion

If the above calibration errs, the direction in which it errs is perfectly clear; the correlations observed in the data can only be viewed as being lower than (or equal to) those of the model, for two reasons. First, the model correlations that we have calculated are lower bounds. By choosing a relatively low value of ρ , we can only have *understated the model correlations*. As has been mentioned, if, instead of what we did, we had attempted to match a country's stock return correlation with its own output, we would have picked a value of ρ approximately equal to 6.9 (see Fig. 3). For that value of ρ , the model would have given us an average cross-country correlation of stock returns equal to 0.97.

What if we had tried instead to match volatilities? It is evident from Fig. 5 that the model, at the chosen value $\rho = 2.1$, explains only about 25% of the actual volatilities. Based on the leverage argument that we have put forward, we do not intend to take the measured volatilities literally. Our analysis of the relative volatilities of output vs. dividend or earnings growth rates is sufficient to explain away the discrepancy in stock return volatilities. Had we tried to match the observed volatilities, however, the needed value of ρ would have been slightly above 9, a value at which the model would indicate an average cross-country correlation of stock returns equal to 0.98. So, at the level of e.i.s. that matches the excessive volatility, there would be no evidence whatsoever of excessive correlation in the data.

The second reason why our calibration can only err in the direction we have stated is that the correlations observed in the data are, if anything overstated. This is because, in several countries, many of the companies listed in the stock exchange typically have levels of foreign activities markedly larger than the share of exports in the corresponding output series. When an overstated share of profits originates abroad, one would surmise, the correlation of the market value of these profits relative to the rest of the world is also overstated. This conjecture is buttressed by the fact that the Netherlands is a clear outlier in Fig. 4; the Amsterdam stock mar-

³⁷ Canova and de Nicolo (1995), in the context of a full-blown, calibrated model of international business cycles find much larger model correlations. The difference is due to the fact that their choice of parameters was dictated by first and own second moments.




Fig. 5. Actual and model stock return volatilities.

ket covers an industrial base which, in fact, is a world-wide one. Besides the Netherlands, Belgium and Canada are the countries for which actual correlations fall above their full-integration levels. We have seen in Section 2 that these countries have the highest proportions of foreign sales.

Other countries that stand out are Japan, Germany and Austria because their full-integration correlation is much higher than the observed correlation. But we know from previous research that Japan was segmented from the world market by regulation until 1981 at least (see Gultekin et al. (1989)). We have no explanation for the Austrian and German deviations.

Except for the cases of the Netherlands, Belgium and Canada, it seems perfectly reasonable to conclude that there is no evidence of excess correlation in the data. If anything, the cross-correlations may be lower than what they should be under full integration,³⁸ this being especially true if we insist on trying to match the correlation between a country's stock return and its own output (see Section 7 below).

6.3. Calibrating local-currency returns, deflated returns and excess returns

We need to check whether our results depend on the way we measure the stock returns. Table 2 compares the results of the calibration, conducted as above, depending on whether the returns are measured in US dollars (as has been done so far), in US dollars in excess of the US interest rate, in local currency but deflated by the local consumer price index or in local currency in excess of the local rate of interest.

³⁸ Within an integrated country such as France, the correlation of stock returns across industrial sectors is equal to 0.818 whereas the cross-country correlation, as we saw, is equal to 0.587. Freimann (1998), as mentioned, offers an alternative, entirely statistical procedure based on randomization of industrial sector returns, to compare country correlations to what they would have been under integration. He finds that cross-correlations are lower than they should be under full integration.

Table 2
The calibration of the integration hypothesis using various units

Country	US\$ return		US\$ excess return		Local return		Local real return		Local excess return	
	$\rho = 2.10$		$\rho = 1.67$		$\rho = 2.10$		$\rho = 2.08$		$\rho = 1.77$	
	Actual	Fitted	Actual	Fitted	Actual	Fitted	Actual	Fitted	Actual	Fitted
Austria	0.435	0.547	0.470	0.197	0.358	0.545	0.364	0.537	0.683	0.192
Belgium	0.710	0.508	0.723	0.183	0.668	0.506	0.672	0.499	0.705	0.321
Canada	0.660	0.531	0.649	0.263	0.702	0.529	0.696	0.524	0.685	0.258
France	0.671	0.675	0.696	0.253	0.645	0.674	0.646	0.670	0.602	0.362
Germany	0.603	0.726	0.615	0.388	0.575	0.725	0.577	0.721	0.470	0.433
Italy	0.483	0.568	0.506	0.451	0.452	0.567	0.458	0.565	0.474	0.333
Japan	0.470	0.663	0.491	0.349	0.456	0.661	0.462	0.657	0.474	0.333
The Netherlands	0.778	0.542	0.789	0.307	0.743	0.541	0.741	0.534	0.772	0.328
Spain	0.531	0.534	0.532	0.366	0.529	0.533	0.535	0.530		
Sweden	0.593	0.571	0.600	0.256	0.553	0.568	0.557	0.559		
UK	0.666	0.598	0.684	0.227	0.662	0.596	0.662	0.589	0.739	0.166
USA	0.537	0.576	0.551	0.657	0.619	0.576	0.627	0.574	0.649	0.705

The first set of results in Table 2 show that, whether returns are measured in dollars, in local currency or in local currency deflated makes very little difference.

The results in terms of excess returns are markedly different from those in terms of other units. The fit of excess stock return correlations is poorer. It is clear that the model does not explain actual 1-month interest rates very well. While interest rates make little difference to the variance of returns, they have a clear impact on correlations.

When examining returns measured in local currency in excess of the local rate of interest, we have had to reduce the number of countries and reduce the length of the sample period because some data on 1-month Euro-rates of interest were not available.³⁹ Some of the differences in the results for “local excess returns” are accounted for by the change in sample. An additional discrepancy arises because the output process was not refitted to the shorter sample. But most of the difference is the result of interest-rate behavior.

The model should probably not have been expected to explain the behavior of local-currency excess returns since it has not been designed to distinguish interest rates denominated in different currencies. But its failure to explain the correlations of US dollar excess returns is more disappointing.

6.4. Calibration under segmentation

With the same estimated “dynamic single-index” business cycle model as in Section 3, we now modify the log-linear pricing kernel, taking each national stock market as a stand-alone financial market. The required change in the pricing kernel is straightforward: Eq. (8), where x now stands for each country’s output, instead of Eq. (10), is used to obtain individual country stock returns. In this formulation, each country lives in autarky. The correlation in output behavior which happens to exist statistically is the only source of common behavior in stock returns. The pricing kernel is a different one in each country although the pricing kernels of different countries do exhibit some degree of cross-correlation since outputs are cross-correlated.

The calibration trade-offs that we face in this case are displayed in Fig. 6. It is immediately apparent from this picture that no value of ρ will allow us to match the actual between-country stock return correlation of 0.595; the model values for these correlations barely reach the value 0.47 when ρ is as high as 10. As far as the correlations with output are concerned, the correlations of a country’s stock returns with its own output are very large in this model. Even with a value of $\rho = 10$, the correlation falls to 0.67 which is still far greater than the observed correlation of 0.207.

Focusing on the only correlation that can reasonably be matched, we choose the value of the e.i.s. to get the best possible match of correlations of stock returns with world output, weighted by the reliabilities of correlation estimates: $\rho = 1.42$. The combined result of the calibration exercise for each country is shown in Fig. 7.

³⁹ The reduced sample starts in October 1978.




Fig. 6. The calibration trade-off under segmentation. For each of the correlations, the level marked “Actual” is the target (observed) correlation that is to be matched.




Fig. 7. Stock return correlations with the rest of the world under two hypotheses: integration and segmentation.

With the single exception of the US, we find that theoretical stock market correlations now fall far below actual ones.⁴⁰ This suggests that the observed levels of

⁴⁰ Even if the value of ρ selected under integration had been maintained, the correlations under segmentation would have been similar to what they are in Fig. 7.

international stock returns correlations are inconsistent with the hypothesis of market segmentation.

In short, we find that correlations are about equal to (or lower than) their full-integration levels and markedly above their complete-segmentation levels, given the common behavior of outputs. This gives us some reason to try and see whether we can construct a test of these hypotheses.

7. A statistical test of the integration hypothesis

It is evident from the international calibration exercise of Section 6 that reality is very much at variance with the full segmentation hypothesis. Two of the three correlation categories that we have chosen to look at, present no prospect of coming reasonably close to their measured counterparts, no matter what value of the crucial parameter ρ we choose. Hence, we focus in this section exclusively on the design and implementation of a test of the full-integration hypothesis.⁴¹

7.1. Test design

We construct deviations between model and reality u_t in the following way. Define y_t as a variable which has been regressed on a set of explanatory variables and call ε_t the residuals of that regression. Let

$$\xi_t = \frac{(\varepsilon_t)^2}{\text{var}(y)} \tag{23}$$

Notice that $\sum_t \xi_t$ is equal to 1 minus the R^2 of the regression. In our application, y_t is each country's monthly rate of stock return and the regressors are the contemporaneous and 11 lagged values of either the country's own output growth rate, or the rest of the world output growth rate or the rest of the world stock return, as the case may be. We calculate this variable in two versions: one denoted ξ_t is based on realized observations; the other $\hat{\xi}_t$ is calculated from model outputs and is, therefore, a function of the unknown parameter ρ .

The deviations between model and reality are calculated as:

$$u_t = \xi_t - \hat{\xi}_t \tag{24}$$

Since we have 12 countries and three categories of residuals that we try to match, we have 36 such deviations at each point in time. We stack them in a 36-element vector which we then use to construct a weighted objective function in the manner of the generalized method of moments. The moments form a vector $g = \sum_t u_t$ and

⁴¹ Under the integration hypothesis, it is not easy to allow for a different value of the parameter ρ for each country's subpopulation of investors. This is because the aggregate of a world population of investors with recursive utility is not a representative investor with recursive utility (see Dumas et al. (2000)). Under the segmentation hypothesis, it would have been possible, of course, to allow such a difference from country to country.

the weighting matrix is the inverse of the variance–covariance matrix of u . This objective function can be minimized to obtain an estimate of the single unknown parameter ρ .

Asymptotically, the minimized objective function is χ^2 distributed with 35 degrees of freedom. We do not know the distribution of the minimized objective function in a finite sample but that is not a problem in our case. The total number of numerical items to be estimated, equal to the number of parameters plus the elements of the weight matrix, is equal to 631 (i.e. $1 + 36 \times 35/2$), while with 304 monthly observations, the total number of stock returns, domestic output and rest-of-the-world output observations is equal to 10,944 (i.e. $304 \times 12 \times 3$). The adjustment for finite-sample size suggested by Ferson and Foerster (1994, 1995), albeit in a somewhat different context, implies that estimated variances of estimates should be multiplied by 1.06 (i.e. $10,944/(10,944 - 631)$), or that the standard error of the estimate of ρ should be increased by 3% only.

We make two amendments to the procedure. The first one aims to take account of the possible serial dependence of the vector u . The adjustment involves an optimal number of lags of the vector u . It follows the method proposed by DenHaan (1996). We allow a maximum lag of 15 months. We only perform a univariate correction: a series' own past values only are considered in this correction for serial dependence. The lag length is chosen on the basis of the Schwarz Bayesian information criterion. Call w the inverse variance–covariance matrix of u after adjustment for serial correlation.

Our second adjustment takes first-stage parameter uncertainty into account. Recall that, in the first stage of our procedure, the dynamic single-index model has been estimated to model output behavior while, in the second stage of the estimation, we now estimate the preference parameter ρ . This is acceptable because the output model is independent of the financial model.⁴² While the structure of the procedure is sound, the parameter uncertainty of the first stage must be taken into account at the second stage. This is easily achieved by first computing the Jacobian matrix $\partial g/\partial\theta$ where θ stands for all the first-stage parameters, and then adjusting the weighting matrix w as follows:

$$w_1^{-1} = w^{-1} + [\partial g/\partial\theta]\Omega[\partial g/\partial\theta]' + [\partial g/\partial\theta]\Gamma + \Gamma'[\partial g/\partial\theta]' \quad (25)$$

where Ω is the 50×50 variance–covariance matrix of the first-stage parameter estimates and Γ is the 50×36 matrix of covariances between the first-stage parameter estimates and the moments.

The GMM iterates over the choice of the parameter ρ and over the choice of the weighting matrix w . Once that is done, the matrix w is replaced by the matrix w_1 and one more iteration series is performed over the choice of ρ .

⁴² Some improvement in efficiency could still be achieved if the two model components were estimated jointly but that is not feasible.

7.2. Test results

Table 3 presents the results of the tests conducted over the 36 moment conditions of the 12 countries of our sample. In the first row, we reject the hypothesis of financial-market integration for all the 12 countries against an unspecified alternative. It is notable that our study of correlations has produced a test powerful enough to reject the integration hypothesis, whereas extant tests based on the first moment and a partial equilibrium model such as an international capital asset pricing model, for the most part, have had too little power to reject.⁴³

This result is subject to an important caveat. It is possible that our rejection is a result of imposing the assumption that ρ is the same for all countries. An extension of the test to accommodate the possibility of country-specific ρ 's will await future research.

7.3. Robustness: breaking the sample and other variations

The model rejection is very much dependent on the fact that we tried to match three types of correlations: correlations of stock returns with each country's output growth, with the rest of the world output and with the rest of the world stock returns. The rejection of the integration hypothesis is a rejection of the adequate match of these three moments. If the integration model is tested with only the two moment categories involving the rest of the world,⁴⁴ we no longer reject the hypothesis at the usual 1% level (even though we still reject it at the 5% level) in the second row of Table 3. A test of the pair of moment conditions involving cross-country stock returns and stock return correlation with own output would lead to a rejection (p -value = 0.003; not shown) but the moment pair involving each country's stock return correlation with world output and with own output would not lead to a rejection (p -value = 0.99; also not shown). It is the cross-country correlation of stock returns which is not compatible with the within-country stock/output correlation.

We are also able to use our test to determine whether the world financial market has evolved over time towards more integration. We split the sample into two halves. Taking output behavior as given and independent of the workings of the financial market, we estimate the second-stage financial component of the model over two subsamples. The third row of Table 3 shows that we reject the integration hypothesis for all the 12 countries over both subsamples.

It is also interesting to see which country, if any, causes the full-integration hypothesis to fail. While commenting on the calibration results, we had second thoughts about including the Netherlands in the sample and we recognized that Japan may have been segmented from the world financial market. However,

⁴³ For an exception, see Jorion and Schwartz (1986). Note that the equity-premium puzzle would not cause partial-equilibrium models to be rejected, except if they are based on consumption behavior.

⁴⁴ Namely, the correlation of stock returns with the rest-of-the-world stock returns and the correlation of output with the rest-of-the-world output.

Table 3
Tests of overidentifying restrictions

Model	Sample	ρ	χ^2	p -value
Test of market integration	January 1970–June 1996	2.88 (0.070)	103.47	<0.001
Two moment test: correlation of stock returns with ROW and correlation of output with ROW output	January 1970–June 1996	2.26 (0.074)	39.42	0.018
Test of market integration over two time periods	January 1970–March 1983	2.66 (0.076)	83.78	<0.001
	April 1983–June 1996	3.84 (0.142)	105.47	<0.001
Test of market integration in local currency	January 1970–June 1996	2.99 (0.083)	99.54	<0.001
Test of market integration in local currency and deflated	January 1970–June 1996	2.94 (0.082)	99.37	<0.001
Test of market integration in excess of the risk free rate	January 1970–June 1996	3.42 (0.091)	94.72	<0.001

Standard errors in parentheses.

excluding one country at a time, we still reject the integration hypothesis at the 1% level (the complete results are available on request). One thus finds no evidence in favor of the idea that one country, being perhaps segmented financially from the rest of the world, would have caused the overall integration test to fail. The calibration exercise that we performed provides us with a reason to try and exclude one pair of countries that may both have been at some point segmented away from the world financial market, namely Japan and Austria. But excluding these two countries together continues to lead us to reject the hypothesis at the 1% level.

Finally, we checked in the calibration section that our results are not very sensitive to the way we measure stock returns. In order to confirm this intuition, we conducted a test for each convention (local currency, real local currency, excess of the risk free rate) we considered. The results in [Table 3](#) indicate that the conclusions remains the same: we reject the integration hypothesis at the 1% level.

8. Conclusion

In this paper, we use the framework of an economic model to link the correlations of international stock markets to those of countries' outputs. We have pursued three goals.

Our first goal is to propose a theoretical explanation for the differences in the observed levels of international stock market correlations. The model has two components. We represent the behavior of outputs (actually, industrial productions) by means of a “dynamic single-index” statistical model, designed to capture the world business cycle as well as country-specific fluctuations. Second, assuming that output and securities' payoffs are closely linearly related to each other, we apply a dynamic representative-agent asset-pricing kernel to the estimated behavior of output. We consider the two polar cases: full integration of international markets, where the different countries are viewed as one country, and complete segmentation, where each country lives in autarky. In the end, we are able to express international stock market correlations as a function of output correlations and of only one preference parameter (the elasticity of intertemporal substitution), under both hypotheses of full integration and complete segmentation.

The second goal is to empirically determine whether the observed level of international stock market correlations is too high to be justified by subsequent changes in national outputs. If this is the case, it leads to an “excess-correlation puzzle”, in the same way [Shiller \(1981\)](#) has documented an “excess-volatility puzzle”.

Using 12 OECD countries and data from January 1970 to June 1996, our estimation method consists of two steps. Initially, we estimate a statistical model for outputs. The values we find seem reasonable and produce a common world business cycle that is fairly persistent. Next, we calibrate the overall model to select a value for the elasticity of intertemporal substitution that best matches the levels of correlations between stock returns and output. Again, the value we find seems reasonable. Finally, with all these parameters, we compute the stock market correlations generated by our equilibrium model. The main result is that, under the hypothesis of integrated financial markets, the international stock market correlations

generated by our model can be matched to the levels of the actual ones. Hence, we argue that there is no “excess-correlation puzzle”.

The last goal of our paper is to use stock market correlations to draw inferences about the degree of international financial-market integration. One often hears the assertion that higher global stock market correlations indicate increased global financial-market integration. Such an assertion is problematic because it does not control for the economic fundamentals within each country. However, our model is ideally suited for this task because it links fundamentals to financial variables. We are able to compare the polar cases of integration and segmentation. The hypothesis of financial-market integration yields stock market correlations in line with the actual ones. In contrast, the hypothesis of financial-market segmentation produces correlations markedly lower than the observed ones. We argue that this is consistent with world stock markets being reasonably well integrated. One type of correlations, however, has not been explained satisfactorily by our model under integration. We find that the theoretical correlation of each country’s stock returns with the own-country industrial production are higher than the observed ones. This is the reason for which, in formal tests of all the moment conditions, the full-integration model is rejected by the data. In the end, our research has demonstrated the potential of focusing on correlations combined with economic data to construct powerful tests of market integration.

Acknowledgements

Dumas and Ruiz acknowledge gratefully the support of the HEC Foundation. Ruiz also acknowledges gratefully the support of IFM₂. Dumas is also affiliated with the University of Pennsylvania (as an Adjunct Professor), the NBER and the CEPR. Harvey is also affiliated with the NBER. The comments of Enrique Arzac, Pierre Batteau, Geert Bekaert, Sudipto Bhattacharya, Bruno Biais, Martin Boyer, William Brainard, John Campbell, Jacques Crémer, Harris Dellas, Vihang Errunza, Christian Gouriéroux, Robert Hodrick, Hayne Leland, Jacques Olivier, Jayendu Patel, Franck Potier, Jean-Charles Rochet, Sergei Sarkissian, Jay Shanken, Huntley Schaller, Robert Shiller, Alan Stockman, and Philippe Weil are gratefully acknowledged. We also appreciate the comments of participants at the NBER Asset Pricing Workshop at the University of Pennsylvania, the French Finance Association, the SCSE Conference, and the meeting of Inquire (Europe), as well as seminar participants of the City University of London, the European Central Bank, the European University Institute in Florence, the University of Rochester, the London School of Economics, Rice University, Universitat Pompeu Fabra, the University of Toulouse, Yale University, the University of Bern, the Centre de Recherche en Economie et Statistique (CREST), the University of Virginia and McGill University. We thank Jacques Raynauld who gave us the SCOREM code to estimate the “dynamic single-index” model and Wouter DenHaan for keeping statistical programs available on his website at the University of California, San

Diego. Ben Zhang was very helpful in developing the econometric test in the final subsection of the paper.

Appendix A. Data appendix

The data used in this article are monthly time series covering the industrial production and the stock returns of a subset of OECD countries. The data are from two different sources, both available on DATASTREAM: OECD for industrial production series and Morgan Stanley Capital International (MSCI) for stock return data.

We selected the following 12 countries based on data availability and on the joint sizes of their economy and their stock market during the last 20 years: Austria, Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Spain, Sweden, the UK and the US. There exist no monthly output series of any kind for Switzerland and Australia.

Some summary statistics are presented in [Table A1](#). Panel A shows the equity capitalization to gross domestic product (GDP) ratios for the 12 countries in 1971, 1980, 1990 and 1995. The proportional size of the equity market increases in every country in our sample except for Canada and Spain. Panel B shows the equity capitalization as a ratio of the MSCI world. The 12 countries in our sample encompass 81% of world market capitalization in 1995. Finally, we examine the GDP as a proportion of OECD GDP. The 12 countries we choose have 91% of OECD GDP (reported in Panel C) in 1995.

Appendix A.1. Industrial production

We have used the monthly time series of real industrial production with a 1990 basis year, deseasonalized, as published for each of the 12 countries by the OECD.

The series codes of the series in DATASTREAM are: OEOCIPRDG, BGO-CIPRDG, CNOCIPRDG, FROCIPRDG, BDOCIPRDG, ITOCIPRDG, JPO-CIPRDG, NLOCIPRDG, ESOCIPRDG, SDOCIPRDG, UKOCIPRDG and USOCIPRDG.

For the weighting of each country in the world aggregate economy, we have used the yearly values of GDP, with a 1990 basis year for the prices and for the exchange rates, as published by the OECD.

The series codes of the series in DATASTREAM are: OEGDP90, BGGDP90, CNGDP90, FRGDP90, BDGDP90, ITGDP90, JPGDP90, NLGDP90, ESGDP90, SDGDP90, UKGDP90 and USGDP90.

Appendix A.2. Stock returns

We have used the monthly time series of MSCI indices, measured in US dollars, with dividends reinvested, with a 1970 basis year, for the 12 countries.

For the weighting of each country in the world stock market, we have used the yearly values of GDP.

The average dividend yields of each country also come from MSCI. We should caution that in some countries, the dividend yield of the index is not available in the early years of the sample period. We have assumed the yield to be constant over the period with missing data.

The consumer price indices (CPI) used to deflate the stock returns and the risk free rates for the excess returns also come from DATASTREAM. The interest rates are the 1-month Euro-dollar deposit rates.

References

- Ammer, J., Mei, J., 1996. Measuring international economic linkages with stock market data. *Journal of Finance* 51, 1743–1763.
- Ang, A., Bekaert, G., 2002. International asset allocation with regime shifts. *Review of Financial Studies* 15, 1137–1187.
- Attanasio, O.P., Weber, G., 1989. Intertemporal substitution and the Euler equation for consumption. *The Economic Journal* 99, 59–73.
- Bansal, R., Lundblad, C., 2000. Fundamental values and asset returns in global equity markets. Working Fuqua School of Business, Duke University.
- Bekaert, G., Harvey, C.R., 1995. Time-varying world market integration. *Journal of Finance* 51, 403–444.
- Beltratti, A.E., Shiller, R.J., 1993. Actual and Warranted Relations between Asset Prices. *Oxford Economic Papers* 45, 387–402.
- Backus, D., Kehoe, P., Kydland, F., 1992. International business cycles. *Journal of Political Economy* 100, 745–775.
- Campbell, J.Y., 1993. Intertemporal asset pricing without consumption data. *American Economic Review* 83, 487–512.
- Campbell, J.Y., 1996. Consumption and the stock market: interpreting international evidence. *Swedish Economic Policy Review* 3, 251–299.
- Campbell, J.Y., Koo, H.K., 1997. A comparison of numerical and analytic approximate solutions to an intertemporal consumption choice problem. *Journal of Economic Dynamics and Control* 21 (2–3), 273–295.
- Campbell, J.Y., Mei, J., 1993. Where do betas come from? Asset price dynamics and the sources of systematic risk. *Review of Financial Studies* 6, 567–592.
- Canova, F., de Nicolo, G., 1995. Stock returns and real activity. *European Economic Review* 39, 981–1015.
- Carriero, F., Errunza, V., Hogan, K., 2001. Characterizing world market integration through time. Working paper, McGill University.
- Cheung, Y.-W., He, J., Ng, L., 1997. What are the global sources of rational variation in international equity returns. *Journal of International Money and Finance* 16, 821–836.
- Choi, J.-J., Hauser, S., Kopecky, K.J., 1999. Does the stock market predict real activity? Time series evidence from the G-7 countries. *Journal of Banking and Finance* 23 (12), 1771–1792.
- Cochrane, J.H., 1991. Production-based asset pricing and the link between stock returns and economic fluctuations. *Journal of Finance* 46, 207–234.
- Deaton, A., 1992. *Understanding Consumption*. Oxford University Press, New York.
- DenHaan, W., 1996. Inferences from parametric and non-parametric covariance matrix estimation procedures. NBER Technical paper #0231. Substantially revised April 1996.

- Dumas, B., 1994. A test of the international CAPM using business cycles indicators as instrumental variables. In: Frankel, J.A. (Ed.), *The Internationalization of Securities Markets*. University of Chicago Press, Chicago.
- Dumas, B., Solnik, B., 1995. The world price of foreign exchange risk. *Journal of Finance* 50, 445–479.
- Dumas, B., Uppal, R., Wang, T., 2000. Efficient allocations with recursive utilities. *Journal of Economic Theory* 93, 240–259.
- Engle, R.F., Watson, M., 1981. A one-factor multivariate time series model of metropolitan wage rates. *Journal of the American Statistical Association* 76, 774–781.
- Epstein, L., 1988. Risk aversion and asset prices. *Journal of Monetary Economics* 22, 179–192.
- Epstein, L., Zin, S., 1989. Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57, 937–969.
- Epstein, L., Zin, S., 1991. Substitution, risk aversion and the temporal behavior of consumption and asset returns: an empirical framework. *Journal of Political Economy* 99, 263–286.
- Erb, C.B., Harvey, C.R., Viskanta, T.E., 1994. Forecasting international equity correlations. *Financial Analysts Journal* November/December, 32–45.
- Fama, E.F., 1990. Stock returns, expected returns and real activity. *Journal of Finance* 45, 793–808.
- Ferson, W.E., Foerster, S.R., 1994. Finite sample properties of the generalized method of moments in tests of conditional asset pricing models. *Journal of Financial Economics* 36, 29–55.
- Ferson, W.E., Foerster, S.R., 1995. Further results on the small sample properties of the generalized method of moments: tests of latent variable models. In: Chen, A.H. (Ed.), *Research in Finance*, vol. 13. JAI Press, Greenwich, Conn. and London, pp. 89–142.
- Ferson, W.E., Harvey, C.R., 1993. The risk and predictability of international equity returns. *Review of Financial Studies* 6, 527–566.
- Freimann, E., 1998. Economic integration and country allocation in Europe. *Financial Analysts Journal* September–October, 32–41.
- Geweke, J., 1977. The dynamic factor analysis of economic time series. In: Aigner, D.J., Golderger, A.S. (Eds.), *Latent Variables in Socioeconomic Models*. Amsterdam, North Holland.
- Giovannini, A., Weil, P., 1989. Risk aversion and intertemporal substitution in the capital asset pricing model. NBER No. 2824.
- Gregory, A.W., Head, A.C., Raynauld, J., 1997. Measuring World Business Cycles. *International Economic Review* 38, 677–701.
- Gultekin, B., Gultekin, M., Penati, A., 1989. Capital controls and international market segmentation: the evidence from the Japanese and American stock markets. *Journal of Finance* 44, 849–870.
- Hall, R.E., 1988. Intertemporal substitution in consumption. *Journal of Political Economy* 96, 221–273.
- Hamao, Y., Masulis, R.W., Ng, V., 1990. Correlations in price changes and volatility across international stock markets. *Review of Financial Studies* 3, 281–308.
- Hodrick, R.J., 1989. Risk, uncertainty and exchange rates. *Journal of Monetary Economics* 23, 433–460.
- Jorion, P., Goetzmann, W.N., 1999. Global stock markets in the Twentieth Century. *Journal of Finance* 54 (3), 953–980.
- Jorion, P., Schwartz, E., 1986. Integration vs. segmentation in the Canadian stock market. *Journal of Finance* 41 (3), 603–614.
- Karolyi, G.A., Stulz, R.M., 1996. Why do markets move together? An investigation of US–Japan stock return comovements. *Journal of Finance* 51, 951–986.
- Kasa, K., 1992. Common stochastic trends in international stock markets. *Journal of Monetary Economics* 29, 95–124.
- Kocherlakota, N.R., 1996. The equity premium: it's still a puzzle. *Journal of Economic Literature* XXXIV, 42–71.
- Kreps, D., Porteus, E., 1978. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46, 185–200.
- Lewis, K., 1999. Trying to explain home bias in equities and consumption. *Journal of Economic Literature* XXXVII, 571–608.
- Longin, F., Solnik, B., 1995. Is the correlation in international equity returns constant: 1970–1990? *Journal of International Money and Finance* 14, 3–26.

- Lucas, R., 1978. Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.
- Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. *Journal of Monetary Economics* 15, 145–161.
- Perez-Quiros, G., Timmermann, A., 1996. On business cycle variation in the mean, volatility and conditional distribution of stock returns. Working paper, University of California, San Diego.
- Raynauld, J., Simonato, J.-G., Sigouin, C., 1993. SCOREM 2.11: A Program for the Estimation of General State-Space Models with the EM and SCORING Algorithms. Ecole des HEC, Montréal.
- Restoy, F., Rockinger, G.M., 1994. On stock market returns and returns on investment. *Journal of Finance* 49, 543–556.
- Restoy, F., Weil, P., 1996. Approximate equilibrium asset prices. Working paper ECARE, Brussels.
- Rodriguez, R., Restoy, F., Peña, J.I., 2002. Can output explain the predictability and volatility of stock returns? *Journal of International Money and Finance* 21, 163–182.
- Sargent, T., Sims, C., 1977. Business cycle modeling without pretending to have too much a priori economic theory. In: Sims, C. (Ed.), *New Methods in Business Cycle Research*. Federal Reserve Bank of Minneapolis, Minneapolis.
- Schwert, G.W., 1990. Stock returns and real activity: a century of evidence. *Journal of Finance* 45 (4), 1237–1257.
- Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.
- Shiller, R.J., 1989. Comovements in stock prices and comovements in dividends. *The Journal of Finance* XLIV, 719–729.
- Shumway, R.H., 1988. *Applied Statistical Time Series Analysis*. Prentice Hall, Englewood Cliffs, NJ.
- Shumway, R.H., Stoffer, D.S., 1982. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis* 3, 253–264.
- Singleton, K., 1980. Latent time series model of the cyclical behavior of interest rates. *International Economic Review* 21, 559–575.
- Stock, J.H., Watson, M.W., 1989. New indexes of leading and coincident economic indicators. *NBER Macroeconomics Annual* 4, 351–394.
- Stock, J.H., Watson, M.W., 1991. A probability model of the coincident economic indicators. In: Lahiri, K., Moore, G.H. (Eds.), *Leading Economic Indicators, New Approaches and Forecasting Record*. Cambridge University Press, Cambridge, England.
- Stock, J.H., Watson, M.W., 1993. A procedure for predicting recessions with leading indicators: economic issues and recent experience. In: Stock, J.H., Watson, M.W. (Eds.), *Business Cycles, Indicators, and Forecasting, Studies in Business Cycles*, vol. 28. University of Chicago Press, Chicago, pp. 95–156.
- Watson, M., Engle, R., 1983. Alternative algorithms for the estimation of dynamic MIMIC factor and time varying coefficient regression models. *Journal of Econometrics* 2, 395–406.
- Weil, P., 1990. Non-expected utility in macroeconomics. *Quarterly Journal of Economics* CV, 29–42.