INTERNATIONAL ECONOMIC REVIEW
Vol. 37, No. 3, August 1996

INTERNATIONAL CONSUMPTION RISK SHARING*

BY FABIO CANOVA AND MORTEN O. RAVN[1]

This paper formally examines the implications of international consumption risk sharing for a panel of industrialized countries. We theoretically derive the international consumption insurance proposition in a simple setup and show how to modify it in more complicated models. We analyze the implications of the theory for pairs of countries and find that aggregate domestic consumption is almost completely insured against idiosyncratic real, demographic, fiscal and monetary shocks over short cycles, but that it covaries with these variables over medium and long cycles. The cross equation restrictions imposed by the theory are rejected. The policy implications are discussed.

"Quisiera saber por qué me complace tanto", le dijo. "Porque los ateos nos acertamos vivir sin los clerigos", dijo Abrenuncio. "Los pacientes nos encomiendan sus cuerpos, pero no sus almas, y andamos come el diablo, tratando de disputarselas a Dios".

—Del Amor y Otros Demonios, Gabriel Garcia Marquez

1. INTRODUCTION

The idea that agents attempt to insure their consumption streams against individual income fluctuations is a pervasive feature of many modern dynamic macro models. It arises in closed economy setups where there is heterogeneity in income or preferences (see Mace 1991, or Marcet and Marimon 1992), in open economy setups where countries with heterogeneous income streams borrow and lend internationally in order to bear only aggregate world-wide risk (see Backus, Kehoe and Kydland 1992), and in models where consumption insurance may be a trigger (or a deterrent) for long term growth (see Devereux and Smith 1994, and Obstfeld 1994). As emphasized by Cochrane (1991), the basic idea of consumption insurance is the cross-sectional counterpart of the permanent income hypothesis. With complete insurance consumption of the individual units (agents, families or countries) should not vary in response to idiosyncratic income shocks while if the permanent income

* Manuscript received July 1994; final revision October 1995.

[1] Part of this research was undertaken while both authors were associated with the European University Institute, Florence. Canova acknowledges the financial support of the European Community through a European University Institute research grant. We would like to thank three anonymous referees, Dave Backus, Bernard Dumas, Alex Kane, Karen Lewis, Jane Marrinan, Maurice Obstfeld, Linda Tesar, Eric van WinCoop and the participants of seminars at Aarhus, Brown, Magdeburg, Southampton, University of California San Diego, University of Southern California, University of California Santa Barbara, the 1993 Society for Economic Dynamics and Control Meeting, Nauplio, and of the Conference on International Asset Markets, London, for comments and suggestions.

hypothesis holds consumption of an individual unit should not vary in response to transitory shocks.

Full consumption insurance obtains in a competitive equilibrium when financial markets are complete or when there are institutions implementing optimal allocations. Some of these institutional arrangements do exist in the real world. At the individual agent level, unemployment or medical insurance schemes, welfare and social government programs or even intergenerational transfers may help in reducing the effect of individual specific shocks. At a country level, charities or disaster relief programs, international lending agreements or direct foreign aid may help to insure national consumption from catastrophic, idiosyncratic income fluctuations. Completeness of financial markets or the presence of institutions implementing first best allocations are, however, only sufficient conditions for full consumption insurance to take place. For example, Duffie and Huang (1985) show that continuous trading of a few long-lived securities may implement optimal allocations. Similarly, Marcet and Singleton (1992) and Baxter and Crucini (1995) show that close to full insurance obtains even when financial markets are strongly incomplete and no mechanism for implementing optimal allocations exists as long as agents have similar preferences (like more consumption to less and less variability to more) and differ only in their income streams.

Although the theoretical mechanics of consumption insurance are very simple, casual empiricism suggests that perfect international risk sharing is hardly a feature of the real world. It is often claimed that national aggregate consumption *does* react to country specific shocks. In support of this idea it is typically reported that poor underdeveloped countries starve when droughts or civil wars cut their own food supplies. For these countries, insurance markets seem to be imperfect or partially inaccessible, no institutions implementing optimal allocations exist and borrowing opportunities are limited (see Atkeson 1991).

Furthermore, many authors (see Backus, Kehoe, Kydland 1992, or Devereux, Gregory and Smith 1992) have informally suggested that cross country consumption correlations among developed countries are less than perfect and, in general, lower than cross country correlations of national outputs. This evidence, combined with simulation results obtained in models where perfect risk sharing occurs, have led Backus, Kehoe and Kydland (1995) to label the empirical cross country relationship of consumptions and incomes a major puzzle.

One additional piece of often cited evidence indicating the lack of complete insurance is the fact that the portfolio of developed countries is composed in large part of domestic assets (see Tesar and Werner 1993) and that agents would be better off by changing the composition of their portfolio (see French and Poterba 1991, or Van Wincoop 1994). While the question of why portfolios are so little diversified is a subject of a debate, several authors have provided a simple rationalization of this phenomenon. Stockman and Dellas (1989) and Tesar (1993) show that nondiversification may occur if nontraded goods account for a large fraction of total consumption. Alternatively, Cole and Obstfeld (1991) and Obstfeld (1992) have suggested that the gains from international diversifications may be small either because there are large informational costs involved in predicting future payoffs of foreign assets or because the cyclical properties of national incomes are alike.

The purpose of this paper is to formally examine the validity of the theory of international consumption insurance. We begin in Section 2 by formulating a simple setup and showing that for consumption insurance to hold a monotonic transformation of the marginal utility of aggregate consumption must be highly correlated across countries. Given a functional form for utility, we next derive three testable implications of the theory. We then describe how to modify the conclusions when some of the restrictive assumptions are removed and derive exact restrictions in two cases of interest: when the planner can not relocate factors affecting the utility of agents (e.g., non-tradable goods or government expenditure) and when she can (e.g., leisure).

Next, we address issues of crucial importance in empirically verifying the implications of the theory. In particular, we discuss problems connected with the choice of instruments and the presence of measurement errors in the variables. Because our theoretical setup disregards possible reallocations of wealth across countries over time, we need to abstract from redistributive considerations in conducting our tests. In addition, because our testing procedure requires stationarity, we will focus attention only on stationary fluctuations and this brings up the issue of which stationary inducing transformation to employ. Canova (1993) shows that different detrending methods leave cycles of different length in the data. Therefore, by appropriately selecting the detrending procedure, we can examine the extent of international risk sharing by frequencies. Such an analysis may shed light on three important issues. First, whether consumption insurance primarily covers temporary or more permanent types of disturbances (see Marcet and Marimon 1992). Second, whether international insurance can be achieved without permanent income redistribution across countries (see Melitz and Vori 1992). Third, whether measurement error has different properties at different frequencies. Finally, we compare our testing approach to others existing in the literature. With the exceptions of Obstfeld (1989), (1993), Lewis (1993) and Kollman (1995) the literature has only informally tested international risk sharing primarily by computing consumption correlation between the U.S.A. and some other country (typically Japan or Canada). The reported evidence, however, concerns primarily the U.S.A. and one may expect closer ties between the E.E.C. or between countries where labor migrations have created a self-insurance mechanism of transfers to chronically depressed areas. In addition, since no standard errors are provided, point estimates of international consumption correlations do not provide conclusive evidence on the issue.

In Section 3 we formally study the properties of pairs of aggregate consumption for a panel of nine fairly homogeneous countries over the sample 1970 through 1990. Since the theory requires that aggregate domestic consumption should be unpredictable once we control for aggregate foreign consumption, natural orthogonality conditions emerge by interacting domestic variables (including variables proxying for real, fiscal, monetary and demographic factors) with the residuals of the regression between pairs of aggregate consumptions. The theory in its simplest form also imposes two additional restrictions: a set of cross-equation constraints across the moments we test and the restriction that cross country consumption correlations are perfect. We examine all these implications using a GMM technique. As a byproduct of our analysis, we also provide novel estimates and standard errors of the ratios of relative risk aversion coefficients of the representative agent of pairs of countries which are of independent interest for researchers engaged in calibrating international business cycle models.

Four major results stand out from our analysis. First, aggregate consumption appears to be fully insured against high frequency fluctuations in real, fiscal, monetary and demographic variables. Second, aggregate consumption covaries with some lagged demographic and labor market variables over medium-long cycles. These two results are robust to the presence of measurement errors, modifications of the empirical specification of the model, changes in the set of instruments and alterations of the functional form for the utility function. Third, the remaining two implications of international risk sharing are rejected regardless of the type of fluctuations we considered. Fourth, there is little evidence of risk sharing in the long run, both in the sense that consumption patterns tend to diverge and that the more permanent component of domestic consumption responds to fluctuations in domestic variables.

Two important conclusions emerge from our investigation which we discuss in Section 4. These conclusions have implications for the design of institutions intended to implement optimal allocations, for the opening of new financial markets and for questions concerning fiscal independence of national governments in a unified Europe. First, because existing arrangements appear to provide substantial insurance against business cycle fluctuations in the short run, both the opening of international financial markets or the introduction of government institutions intended to augment the efficiency of existing markets (see Persson and Tabellini 1992) would only marginally affect the welfare of the representative agent of these countries. This result is particularly important for European countries since one of the foreseeable tasks of the central authority of a unified Europe is to provide short run consumption insurance to member countries (see Padoa-Schioppa 1987). Second, because consumption insurance is harder to obtain for low frequency fluctuations in consumption, international institutions (old and new) should be more concerned with shielding consumers from this type of domestic fluctuations. However, this type of insurance may involve semi-permanent redistribution of income. Therefore, to fully evaluate the effects of these programs on the well being of agents, it is necessary to examine them in the context of a specific model of risk sharing and redistribution.

2. THE THEORY OF INTERNATIONAL RISK SHARING

2.1. *The Basic Setup.* We assume that there is a representative consumer in each of the J countries of the world, $j = 1, \dots, J$, with preferences defined over a homogeneous aggregate nondurable consumption good. Let s^t be the realization of the state at time t, $s = 1, \dots, S < \infty$ and assume that $s_t = \{s^t, s^{t-1}, \dots s^1\}$, the history of realizations of s^t, is observed by the agents of all countries at time t, that the ex-ante probability (at time $t = 0$) of a particular history realization is $\pi(s_t)$ and that $\sum_{s_t} \pi(s_t) = 1$. Each country is endowed with an exogenous stochastic amount $y_t(s_t)$ of the good at each t. The lifetime expected utility of the representative consumer in each country is given by:

$$(1) \qquad V \equiv \sum_{t=1}^{\infty} \beta^t \sum_{s_t} \pi(s_t) U\left[c_t^j(s_t), b_t^j(s_t)\right]$$

where $0 < \beta < 1$ is the discount factor common to all j's and is independent of s_t, $c_t^j(s_t)$ is the consumption of country j at time t, if the history s_t occurs and $b_t^j(s_t)$ represents factors other than consumption affecting agents' utility (such as leisure, government consumption, consumption of nontradables or home production effects, as in Benhabib, Rogerson and Wright 1992, etc.). In our simplest setup b_t^j captures all those factors which are uninsurable from the planner's point of view (see Obstfeld 1993 for a similar interpretation). That is, the planner computes optimal allocations of the consumption goods only, taking everything else as given. This restrictive assumption will be relaxed in the next subsection where the planner is allowed to allocate some components of b_t^j (e.g., leisure) in addition to the consumption good.

The optimality condition for the planner problem is:

$$(2) \qquad \Phi_j U_{c_t}^j = \mu_t$$

where $U_{c_t}^j = (\partial U / \partial c_t^j(s_t))$, $\mu_t = (\lambda_t / \beta^t \pi(s_t))$, λ_t is the Lagrange multiplier for the resource constraint and $\Phi_j = (\Psi_j / \chi_j)$ is a set of time and state invariant weights which are related to the size χ_j and the initial wealth Ψ_j of each country. Implicit in the formulation of the problem are the assumptions that along the equilibrium path the planner does not reallocate wealth and that the size of the country (in terms of population) does not change.

The essence of the theory is embodied in (2) and it is given by the fact that μ_t is the same for all j. Consequently, (2) implies that, apart from a scale factor, the marginal utility of consumption is equalized across countries. We can use (2) and the fact that μ_t is independent of j, to derive the following two conditions, which must hold for any pair $j, k \in J$:

$$(3) \qquad \log U_{c_t}^j - \log U_{c_t}^k = \xi_{jk}$$

$$(4) \qquad \log U_{c_t}^j - \log U_{c_t}^a = B_j$$

where $U_{c_{\tau t}}^a = 1/J \sum_{j=1}^J U_{c_{\tau t}}^j$, and where $\xi_{jk} = \log \Phi_k - \log \Phi_j$, $B_j = \log \Phi_a - \log \Phi_j$, and $\Phi_a = 1/J \sum_{j=1}^J \Phi_j$ are independent of the date and the state. Both equations must hold for all countries, all histories s_t and all periods t, they are not distinct (taking the difference between any two indices j and k in (4) we obtain (3)) and the first is stronger than the second since there is a loss of information in (4) due to the averaging of consumption across countries. That is to say, if (3) holds for all j, k, then (4) must be satisfied, while the reverse is not true.

Both (3) and (4) impose strong restrictions on the data since they equate the marginal utility of consumption for any j, k and for all histories and dates. Because c_t is not observable for all histories, and in testing risk sharing we have to face the problem of uneven quality of international consumption data, we next derive weaker restrictions based on the moments of the logarithm of the marginal utility of aggregate consumption. These implications are easily testable and more robust to measurement errors. From (3) it is straightforward to obtain

$$E\left[\log U_{c_t}^j\right] = E\left[\log U_{c_t}^k\right] + \xi_{jk} \tag{5}$$

$$\operatorname{var}\left[\log U_{c_t}^j\right] = \operatorname{var}\left[\log U_{c_t}^k\right] \tag{6}$$

$$\operatorname{cov}\left[\log U_{c_t}^j, \log U_{c_t}^j\right] = \operatorname{cov}\left[\log U_{c_t}^j, \log U_{c_t}^k\right] \tag{7}$$

Note that (6) and (7) hold regardless of the welfare weights employed in the planner problem as long as they are constant. Therefore, there are two reasons why (5)–(7) are weaker than (3). First, because they imply that international consumption insurance holds for moments of the data instead of point by point in time and state. Second, because these relationships can be examined even when the planner weights are unknown.

There are many decentralization schemes which support the optimal risk sharing allocations. To achieve optimality in decentralized economies it is typical to assume the existence of equity markets in each country (as in Lucas 1982) or of contingent claim markets where countries trade claims to each others income (as in Backus, Kehoe and Kydland 1992). However, even when contingent claims markets are missing, richness in the set of trading opportunities may partially compensate for the lack of variety in the available securities (see Duffie and Huang 1985). At the polar extreme, the lack of international financial markets is not sufficient to prevent risk sharing. For example, Persson and Tabellini (1992) describe a cooperative political equilibrium arrangement which implements optimal allocations even when international financial markets are absent. Because of these considerations and since the implementation of the tests is unaffected, we decided to leave the financial instruments available to agents unspecified (see also Townsend 1994 on this point).

To give empirical content to the theory, we use a homothetic utility function of the type

$$U\left(c_t^j, b_t^j\right) = \frac{1}{1-\sigma_j}\left[\left(c_t^j b_t^j\right)^{1-\sigma_j} - 1\right] \quad \text{if } \sigma_j \neq 1 \quad \forall j \tag{8}$$

$$U\left(c_t^j, b_t^j\right) = \log\left(c_t^j\right) + A^j \log\left(b_t^j\right) \quad \text{otherwise} \quad \forall j \tag{9}$$

where σ_j is the coefficient of constant relative risk aversion (CRRA) of country j. Although we adopted a CRRA specification, any member of the class of HARA utility functions would serve the purpose (see Brennan and Solnick 1989). We selected a CRRA utility function because it is consistent with a balanced growth path. Devereux, Gregory and Smith (1992), who have studied the implications of risk sharing in an international model of the business cycle where leisure choices are included, have used a utility function of the form:

$$U\left[c_t^j(s_t), n_t^j(s_t)\right] = \log\left(c_t^j - \gamma n_t^j\right) \tag{10}$$

where n_t is hours. For these preferences, the marginal utility of consumption is not independent of labor supply choices but the income elasticity of leisure is zero. However, these preferences are not of the HARA class and are incompatible with balanced growth.

Given (8)–(9) and setting $b_t^j = b$, $\forall j, t$ for the moment, we let

$$X_{1t} = \log c_t^j - \left(\frac{\sigma_k}{\sigma_j}\right)\log c_t^k + \xi_1 \tag{11}$$

$$X_{2t} = \left(\log c_t^j\right)^2 - \left(\frac{\sigma_k}{\sigma_j}\right)^2 \left(\log c_t^k\right)^2 \tag{12}$$

$$X_{3t} = \sqrt{\left(\log c_t^j\right)^2\left(\log c_t^k\right)^2} - \left(\frac{\sigma_k}{\sigma_j}\right)\operatorname{corr}\left(\log c_t^j, \log c_t^k\right)\left(\log c_t^j\right)^{-2} \tag{13}$$

where ξ_1 is a constant function of the welfare weights and of the CRRA coefficient of country j and let $X_t = [X_{1t}, X_{2t}, X_{3t}]'$. The first testable implication of the theory, given our setup and the auxiliary assumptions, is that X_t should be unpredictable. That is, for any variable Z_t belonging to the information set available at t, $E_t[X_t | Z_t] = 0$. Since s_t is observed in all countries, Z_t may in principle include variables of all countries in the panel. However, to make tests of the theory more stringent, we will consider only Z_t variables which are specific to country j.

The theory of international risk sharing also imposes cross equation constraints across (11)–(13). A second implication of the theory is therefore that a transformation of the slope coefficient of the three moment conditions equals the ratio of the coefficients of relative risk aversion of the representative agents of the two countries. These restrictions are distinct from the basic orthogonality conditions, independent of the chosen functional form for utility (see (5)–(7)) and provide a further and more stringent test of the theory. Finally, in our restricted setup, risk sharing implies that consumption correlations for any pair of countries should be perfect. This can be easily seen from (6) and (7) and the fact that $\operatorname{cov}[\log U^k_{c_{\tau t}} \log U^k_{c_{\tau t-q}}] = \operatorname{cov}[\log U^j_{c_{\tau t}} \log U^j_{c_{\tau t-q}}]$, $\forall q$. This restriction is independent of the functional form for utility, but not of the other restrictions imposed by the theory.

The three implications we have derived are robust to modifications of several assumptions of the model. For example, we can relax the assumption that there is only one good. As long as there is complete specialization, that is country j receives an endowment of good j only, pairs of aggregate consumption still satisfy (3), and given our CRRA specification for utility, the implications we derived still hold. However, when there are multiple goods country specific disturbances change the terms of trade and these changes may automatically pool national risk. For example, in Cole and Obstfeld (1991) a technology shock induces, other things being equal, a negative relationship between terms of trade and output. We could further relax the exchange economy setup, assume that production requires capital goods from one or more countries and still generate the same three empirical implications obtained in the simplest setup (see Cole and Obstfeld 1991).

The empirical implications we have derived are also qualitatively robust to the assumption of a representative agent in each country. If there is some heterogeneity within each country one can assume that the planner maximizes a weighted average of the expected utility of the median voter of each country (as in Persson and Tabellini 1992) and obtain, once again, condition (3). If (8) and (9) are the instantaneous utility function of the median voter with $b_t^j = b$, $\forall t, j$, then (11)–(13) still hold but now ξ_1 is a function of the characteristics of the median voter of countries j and k.

It is useful to ask whether it makes any difference for the restrictions we have derived if y_t contains both an idiosyncratic and an aggregate component. Since (3) is obtained without any reference to the stochastic process for income, it holds under alternative assumptions on the properties of the idiosyncratic and of the aggregate component of income. Therefore, in our simple setup and given the auxiliary assumptions made, risk sharing implies that domestic consumption should react to total world income not to individual income or to any of the two components of individual income (see also Crucini 1994).

Finally, since (3) is independent of the functional form for utility selected, it holds when preferences do not display time separability in consumption (e.g., when there is habit persistence or durability in consumption, see Mace 1991). Furthermore, because versions of the Porteus-Kreps preferences are observationally equivalent to von Neumann-Morgenstern preferences with a particular form of habit persistence (see Constantinides 1991), (3) holds even when preferences satisfy non-expected utility axioms. If the instantaneous utility function is specified as in (8) and (9), where now $c_t^j = C_t^j - \sum_{p=1}^{P} \alpha_p C_{t-p}^j$ and $\alpha_p > 0$ with habit persistence and $\alpha_p < 0$ with durability, the three testable implications we derived still hold with c_t now representing a linear combination of current and past consumption. There are two consequences of this. First, the theory predicts that X_t, constructed with linear combinations of current and past consumptions, is unpredictable. Second, if only current consumption is used to generate orthogonality conditions, $E(X_t|Z_t) \neq 0$, unless Z_t are uncorrelated with past p lags of consumption.

2.2. *Extensions.* The setup we have used so far is appealing because it is simple but it imposes restrictive assumptions on the type of effects that b_t^j may capture and on the activities of the planner. Here we examine first what happens when some of the components of b_t^j are stochastic but still beyond the control of the planner (e.g., if there is domestic stochastic government expenditure or stochastic non-traded or home produced goods in the utility) and, second, what happens when some of the stochastic components of b_t^j are under the control of the planner (e.g., if b_t^j includes leisure choices).

To see exactly what the first extension entitles we consider the special case where b_t^j represents non-tradables. In this case (3) still holds. However, if there are tradable and non-tradable goods and the instantaneous utility function is of the form:

$$(14) \qquad U\left[c_t^j(s_{\tau t}), b_t^j(s_{\tau t})\right] = \frac{1}{1-\sigma_j}\left(\left(tr_t^j\right)^{\gamma_j}\left(nt_t^j\right)^{1-\gamma_j}\right)^{1-\sigma_j} \quad \text{if } \sigma_j \neq 1\ \forall j$$

$$(15) \qquad U\left[c_t^j(s_{\tau t}), b_t^j(s_{\tau t})\right] = \log\left(tr_t^j\right) + A^j \log\left(nt_t^j\right) \quad \text{if } \sigma_j = 1\ \forall j$$

where nt_t^j is non-tradable consumption, tr_t^j is tradable consumption and γ_j is the share of tradables in total consumption in country j, the optimality condition implies that

$$(16) \quad x_{1t}^* = \log\left(tr_t^j\right) - \frac{\gamma_k(1-\sigma_k)-1}{\gamma_j(1-\sigma_j)-1}\log\left(tr_t^k\right) - w_{1t} - w_2$$

$$(17) \quad x_{2t}^* = \left(\log tr_t^j\right)^2 - \left(\frac{\gamma_k(1-\sigma_k)-1}{\gamma_j(1-\sigma_j)-1}\right)^2\left(\log\left(tr_t^k\right)\right)^2 - w_{1t}^2$$

$$(18) \quad x_{3t}^* = \sqrt{\left(\log tr_t^j\right)^2\left(\log tr_t^k\right)^2} - \left(\frac{\gamma_j(1-\sigma_j)-1}{\gamma_k(1-\sigma_k)-1}\right)\left(\operatorname{corr}\left(\log\left(tr_t^j\right)\log\left(tr_t^k\right)\right)\left(\log\left(tr_t^j\right)^{-2} - \sqrt{\left(\log tr_t^k\right)\right)^2\left(\log tr_t^j\right)^2} w_{1t}\left(\log tr_t^j\right)^{-1}\right)$$

are unpredictable where $w_{1t} = (1/(\gamma_j(1-\sigma_j)-1))[(1-\gamma_k)(1-\sigma_k)\log nt_t^k - (1-\gamma_j)(1-\sigma_j)\log nt_t^j]$, $w_2 = (1/(\gamma_j(1-\sigma_j)-1))(\log \Phi_k - \log \Phi_j + \log \gamma_k - \log \gamma_j)$. From (16)–(18), it is clear that if the marginal utility is nonseparable in traded and non-traded consumption, fluctuations in the endowment of domestic and foreign non-tradables affect the optimal allocation of the traded good. In particular, the optimal allocation depends on the share parameter (γ_j), the intertemporal elasticity of substitution $(1/\sigma_j)$ and the joint processes for traded and non-traded goods (see Tesar 1993).

Similarly, if b_t^j is identified with home production or with exogenous government consumption expenditure affecting the utility of the agents of their own country, the optimality condition for the planner problem still requires the equalization of the marginal utility of private consumption across countries. However, when private and government consumption expenditure (or home production) are nonseparable in utility and because these factors are country specific and cannot be reallocated by the planner, the logic of these cases is identical to the case of non-tradables: fluctuations in government expenditure or home produced output affects the optimal allocations of the market consumption good.

There are several implications of these results. First, while the optimality condition requires that $E(X_t^* | Z_t) = 0$, it is not necessarily the case that $E(X_t | Z_t) = 0$ where X_t is computed neglecting that b_t^j are stochastic. In other words, unless Z_t are uncorrelated with b_t^j, tests conducted using consumption data only will generally reject unpredictability. Second, while the cross equation restrictions still hold, the estimated slope parameter measures a nonlinear combination of the ratio of the relative risk aversion coefficients between pair of countries and the ratio of shares of private (tradable) consumption in the utility. Third, it is no longer true that international consumption correlations are perfect, unless the utility function is separable in c_t^j and b_t^j. Therefore, finding that consumption correlations are less than perfect is not necessarily an indication of the lack of full risk sharing. Conversely, even if risk sharing is perfect, the presence of factors entering in a nonseparable way with consumption in utility drives a wedge between the profile of domestic and foreign consumption (see also Marrinan 1994 on this point).

Next, consider the case where the planner has control over the cross country allocation of b_t^j. To make the exposition concrete, we let b_t^j represent leisure. To make the problem meaningful we abandon the pure exchange setup and assume that y_t^j is produced with domestic labor according to $y_{jt} = f_t^j(1 - L_{jt})$ where $(1 - L_{jt})$ is the amount of labor used in country j, f_t^j is time varying because of a country specific technology disturbance and satisfies standard regularity conditions. Let the instantaneous utility function be:

$$(19) \qquad U\left[c_t^j(s_{\tau t}), b_t^j(s_{\tau t})\right] = \frac{1}{1 - \sigma_j}\left(\left(c_t^j\right)^{\gamma_j}\left(L_t^j\right)^{1-\gamma_j}\right)^{1-\sigma_j}, \quad \text{if } \sigma_j \neq 1 \ \forall j$$

$$(20) \qquad U\left[c_t^j(s_{\tau t}), b_t^j(s_{\tau t})\right] = \log\left(c_t^j\right) + A^j \log\left(L_t^j\right) \quad \text{if } \sigma_j = 1 \ \forall j$$

where γ_j represents the share of consumption in the utility function of country j. Manipulation of the two first order conditions, yields that $\hat{X}_t = [\hat{x}_{1t}, \hat{x}_{2t}, \hat{x}_{3t}]$, defined as:

$$(21) \quad \hat{x}_{1t} = \log c_t^j - \log c_t^k - \log L_t^j + \log L_t^k - \log f'_{jt} + \log f'_{kt} - B_{jk}$$

$$(22) \quad \hat{x}_{2t} = \left(\log c_t^j\right)^2 - \left(\log c_t^k\right)^2 - \left(\log L_t^j\right)^2 + \left(\log L_t^k\right)^2 - \left(\log f'_{jt}\right)^2 + \left(\log f'_{kt}\right)^2$$

$$(23) \quad \hat{x}_{3t} = \sqrt{\left(\log c_t^j\right)^2\Big)\Big(\log c_t^k\Big)^2\Big)} - \operatorname{corr}\left(\log\left(c_t^j\right), \log\left(c_t^j\right)\right)\left(\log c_t^j\right)^{-2}$$
$$- \log\left(c_t^j\right)^{-1}\left[\left(\log L_t^j\right) - \left(\log L_t^k\right) + \left(\log f'_{jt}\right) - \left(\log f'_{kt}\right)\right]$$

should be unpredictable given the information available at t, where f'_{jt} is the marginal product of labor in country j and $B_{jk} = \log((1 - \gamma_k)/\gamma_k) - \log((1 - \gamma_j)/\gamma_j) + \log(\Phi_k/\Phi_j)$. The modified condition for international insurance states that differences in the moments of consumption profiles across countries should be matched by differences in the moments of leisure and productivity profiles. No other variable should be important in predicting differences in aggregate consumption across countries. Once again, while there are cross equation constraints, it is no longer true in this setup that cross country consumption correlations are perfect.

2.3. *Empirical Issues.* In both extensions we have presented and when utility is time nonseparable, testing international consumption risk sharing using only private aggregate consumption data to construct X_t is valid as long as the instruments are judiciously chosen. In other words, if an econometrician omits variables which are stochastic and either nonseparable with consumption in the utility function, enter in the computation of the optimal planner rule or help to predict these variables, she may reject full risk sharing even if it holds true. For example, when tradables and non-tradables are nonseparable in the utility function, variables like domestic income or domestic production may be significant in predicting X_t because they may be correlated with consumption of non-tradables in country j (as noted by Dellas and Stockman 1989). Similarly, domestic income is a poor instrument if government consumption enters the utility function, in particular, when government expenditure accounts for a large portion of GDP; lagged consumption may help to predict cross country differences if preference are time nonseparable in consumption; finally, variables correlated with the marginal product of labor may be significant in explaining differences in private consumption across countries if leisure is a choice variable.

Because we want to examine the implications of the theory in its simplest form and interpret deviation from full risk sharing when they arise we proceed in two steps: first, we exclude all these variables from the Z_t vector and examine whether the basic implications of the theory hold. Second, if the theory is rejected we add them to the instruments to see if their omission is responsible for the results. If not, misspecification of the nonseparable effects in preferences is unlikely to be the cause of the rejection of the theory.

There are three further empirical issues, all having to do with the selection of the Z_t vector, which need to be addressed before the tests are undertaken. The first concerns the presence of measurement error in consumption. It is very likely that the properties of the measurement error in consumption differ across countries given the variety of collection procedures employed and this may generate serially correlated measurement errors in X_t. In this case, care must be exercised in selecting the Z_t's because testing may be invalid when the instruments are mismeasured and the error is correlated with the measurement error in X_t. For example, if measurement errors in consumption are highly serially correlated but uncorrelated across countries, lagged consumption may not be a valid instrument even when the theory in its simplest form holds. Similarly, current GDP may not be a valid instrument since its measurement error may be highly correlated with the measurement error in domestic consumption (see also Cochrane 1991). To reduce the extent of the serial correlation in the measurement errors one can consider instruments sufficiently lagged in the past. This procedure has however the disadvantage of also reducing the power of the tests since the correlation of the instruments with the current marginal utility of consumption will also be smaller. An alternative approach, which will reduce the importance of measurement error without necessarily affecting the power of the tests, is to use as instruments the component of Z_t which is unexplained by past values of Z_t. To see why this may solve the problem suppose that measurement error is the same for all variables in a given country, it is stationary and it is serially correlated over time. Then $T^{-1}\Sigma_t X_t Z_{t-m} \to 0$ only as $m \to \infty$. However, $T^{-1}\Sigma_t X_t \hat{Z}_{t-m}$ where $\hat{Z}_{t-m} = Z_{t-m} - E[Z_{t-m} | I_{t-m}]$, where I_t is the information set at t, is likely to be zero even for small m since the constant and common component of the measurement error is purged, to some extent, from the instruments. Using $\hat{Z}_t$ in place of Z_t may also be more consistent with the idea that the logarithm of the marginal utility of domestic consumption is insensitive to idiosyncratic *shocks* (see Mace 1991, or Townsend 1994). Comparing the results obtained with $\hat{Z}_t$ and Z_t as instruments may therefore shed light on some of the properties of the measurement errors.

As we have already stressed, in our setup the planner does not reallocate wealth along the equilibrium path. To examine the empirical counterpart of the theory we therefore need to abstract from redistributive considerations. That is to say, variables correlated with wealth differences across pairs of countries should be excluded from the Z_t vector as they are likely to be correlated with cross country differences in consumption profiles. As this exclusion may substantially reduce the dimensionality of the Z_t vector, one can alternatively restrict attention to those Z_t's which may reflect idiosyncratic *cyclical* fluctuations. One way to insure that this is the case is to transform the data into a stationary form. Apart from theoretical considerations, stationarity of X_t and Z_t is also needed by our estimation and testing methodology.

The third issue concerns how to induce stationarity in the data. Since there are many detrending procedures which make a time series stationary, one may consider whether the choice of detrending affects the conclusions. There are various reasons for why this could happen. For example, it may be the case that short-term fluctuations are easier to insure than fluctuations that have a more permanent nature since the latter may signal structural imbalances and the possibility of debt repudiation. Alternatively, it may well be the case that measurement errors are more important at some frequencies than others (e.g., growth rates are easier to measure than levels). For this reason we detrend the data with three methods: linear detrending (LT), Hodrick and Prescott filtering (HP) and first order differencing (FOD). Canova (1993) has demonstrated that these detrending methods leave stationary cycles in the data with an average periodicity of less than 3 years for FOD, between 4 to 6 years for HP and 7 to 10 years for LT. By using three different methods we are therefore able to investigate whether risk sharing occur at some particular frequencies of the spectrum.[2] To illustrate these characteristics of the filters, we plot in Figure 1 the three cyclical components obtained by applying the three detrending methods to U.S. consumption data. An alternative view about this topic is that certain detrending methods are optimal given a particular generating

[2] It is useful to stress that an approach that separates the information by frequency shares some similarities with "band spectrum regression" technique of Engle (1974), and with the procedure used in Canova and Dellas (1993), where one computes the average coherence between pairs of series in a selected frequency band and checks whether the value of the statistic over the band can be related to interesting variables.

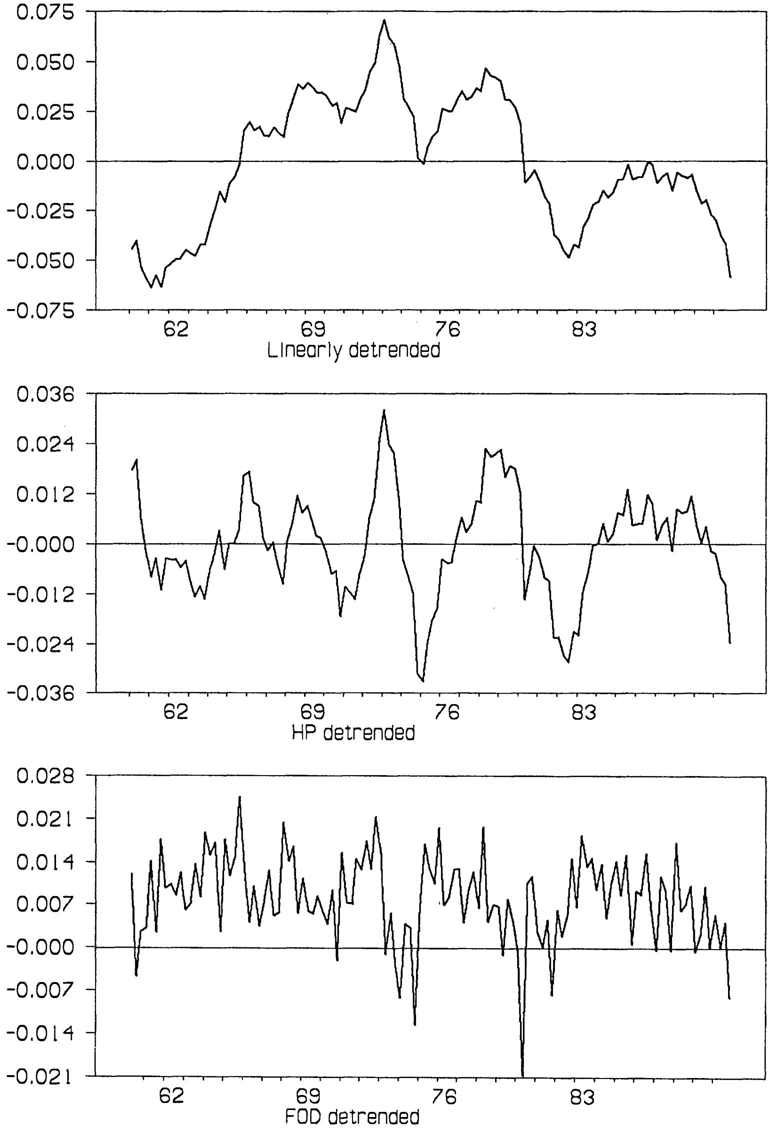

FIGURE 1

process for the data (see King and Rebelo 1993). Accordingly, the use of a particular filter may lead to distortions if applied to a series for which it is not optimal (e.g., the application of HP filter to random walk data in Harvey and Jeager 1993). However, to proceed this way one has to assume that the investigator knows the DGP of the series, an assumption we are reluctant to make given the size of the available data and existing econometric techniques.

To test the three implications of the theory we will employ a GMM methodology. To check for the unpredictability of X_t we examine whether the coefficients of a regression of X_t on the instruments are all zero using a J-test. To test the validity of the cross equation constraints, we take unrestricted GMM estimates of the slope coefficients of the three moment conditions and test their equality using a Wald test. Finally, to examine whether cross country consumption correlations are perfect we employ a t-test on the Fisher Z-transform of the estimated correlations implicit in the slope coefficient of X_{3t} and their standard deviations.

We conduct our investigation using a panel consisting of nine O.E.C.D countries: Australia, Canada, France, Germany, Italy, Japan, Switzerland, United Kingdom, and U.S.A. and test the restrictions on a country-pair basis. This strategy minimizes the chance of incorrect inferences due to anomalies associated with one specific country and gives more power to the tests. The countries were selected because they are structurally homogeneous and because they are the only ones with consistent quarterly data for a sufficiently long span of time.

As a measure of consumption we use the natural logarithm of total real aggregate consumption per-capita. Total consumption may not be an ideal measure for our exercises, especially if there are persistent cross country differences in the components of aggregate consumption (such as services, which account for a large portion of nontraded goods). Nevertheless, there are serious problems in obtaining disaggregated data and if one wants to restrict attention to only nondurables and services or exclude nontradables from consumption data, the panel of countries must be drastically reduced. In addition, there is no a priori reason to expect that the use of more disaggregated data would reduce misspecification. If risk sharing is associated with consumption of nondurables and services or tradables and the utility function is separable between various categories of consumption goods, it is appropriate to use a consumption measure that excludes durables and nontradables. However, if the utility function is nonseparable between various types of goods, then excluding durables or nontradables would cause distortions. Finally, it has to be kept in mind that cross-country differences in measurement errors are likely to be more serious the more disaggregated is the consumption data used. As instruments in the tests we use proxies for real, fiscal, monetary and demographic variables. The sources of the data, its definition and its sample availability are described in the appendix.

2.4. *A Comparison with the Existing Literature.* Although the theory of consumption insurance is a quarter of a century old (see Wilson 1968) relatively few *formal* empirical investigations of the implications of the insurance proposition have been conducted. Mace (1991), Cochrane (1991) and Townsend (1994) have investigated whether agents respond to aggregate but not to idiosyncratic shocks. Mace and Townsend use individual income as a measure of idiosyncratic shocks and investigate whether agents react to individual income using pooled time series cross sectional regressions. Cochrane employs a number of individual characteristics to capture idiosyncratic shocks and tests whether variations in consumption are related to changes in these individual characteristics using a cross-section approach. All authors find some evidence of pooling of idiosyncratic shocks, but they also find that there are many sources of individual risk (such as long illnesses or long unemployment spells) which cannot be insured. There are two major differences between our approach and theirs. First, we test risk sharing employing time series observations for a sufficiently long time span (Townsend (1994) attempted to do the same but had only 10 annual observations per individual household). Second, we use different estimation and testing techniques.

Asdrubali, Sorensen and Yosha (1995) have examined the extent of regional risk sharing in the U.S.A. in an ingenious way. Instead of testing whether allocations are optimal, their work aims at examining how much of the time series variance of state income is insured by factor mobility, government transfers and credit markets. They find that the first two channels insure a substantial portion of statewide GDP fluctuations. Crucini (1994), on the other hand, examines the extent of risk sharing in U.S.A. and in Canadian provinces by allowing consumers to invest in a world mutual fund. The regression coefficient of provincial consumption on this mutual fund measures the extent of risk pooling needed to achieve full insurance, with zero indicating perfect risk sharing and one complete autarky. As do Atkeson and Bayoumi (1993), he finds that there is insurance but that it is far from perfect.

At an international level, and to the best of our knowledge, only Obstfeld (1989), (1993), Lewis (1993) and Kollman (1995) have formally examined the implications of risk sharing. Obstfeld (1989) tests uncovered interest parity in an attempt to measure the extent of capital mobility across countries and Obstfeld (1993) tests implication (11) of risk sharing as a way to gauge the completeness of the international financial markets. Kollman (1995), on the other hand, examines the implications of complete and incomplete markets for consumption and the real exchange rate in the short and long run. Because completeness of financial markets is only a sufficient but not necessary condition for full insurance to take place their tests are not completely comparable to ours. Finally, Lewis exploits cross-sectional information to examine whether domestic consumption varies with domestic income once world income is included in the regression. The results of all these investigations are mixed: they detect some evidence of financial market incompleteness but they also find that risk sharing has increased after 1973. Because our model specification, estimation technique, testing approach, data set and panel employed differ from theirs, our investigation provides an alternative and complementary attempt to evaluate the optimality of international consumption allocations.

3. THE RESULTS

3.1. *Business Cycle Implications.*

3.1.1. *Testing the orthogonality conditions.* Table 1 presents the basic results of our investigation. It reports the significance level (in percentage terms) of a J-test for the overidentifying restrictions implied by risk sharing when we jointly use the three moment conditions (11) to (13). The first panel of the table presents results for LT detrended data, the second for HP detrended data and the third for FOD detrended data. The instruments we employ are a constant, lagged domestic outputs, lagged domestic population and lagged domestic prices. This means that when we consider, for example, the Australia-Canada pair, the instruments are a constant and lagged Australian output, population and prices. We only use lagged values of the instruments in order to minimize endogeneity problems, which are primarily present with output data, and mismeasurement problems, which may arise when the measurement error in the instruments is contemporaneously correlated with the measurement error in consumption. In all cases we apply the same detrending transformation to consumption data and to the instruments. That is, if we linearly detrend consumption, we also linearly detrend the instruments. This makes the cross frequency comparison more appropriate and allows us to address the question of the comovements between consumption and the instruments at various frequencies.

TABLE 1
P-VALUES (IN PERCENTAGES) FOR THE TEST THAT $E[X_t|Z_t]=0$:
DETRENDED INSTRUMENTS*

	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
Linearly Detrended data									
AUS		0.0143	0.0396	0.1008	0.0791	0.9778	0.0003	0.6023	6.3237
CAN			1.9655	0.0730	0.0131	2.9627	0.0001	0.0000	0.0079
FRA				1.7262	1.9468	6.2842	0.1165	0.5232	0.7088
ITA					0.2206	0.4928	0.8878	2.9091	0.2089
JAP						0.9912	0.0000	0.0161	0.9182
SWI							1.4268	7.5635	1.9050
UK								0.0041	0.0205
USA									0.0090
HP Detrended data									
	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
AUS		0.8781	1.1076	12.965	10.040	25.763	7.2284	1.3623	14.541
CAN			15.918	3.3472	7.6228	4.2277	2.5303	0.8536	0.3313
FRA				8.8903	10.872	5.5677	7.3689	15.560	18.961
ITA					12.452	7.4112	3.4685	6.1406	74.793
JAP						16.439	5.1440	1.5182	40.020
SWI							27.766	14.668	14.614
UK								0.1582	4.2776
USA									2.4461
FOD Detrended data									
	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
AUS		27.133	24.361	17.586	10.810	34.740	6.0590	51.467	15.893
CAN			74.914	48.487	48.610	15.277	34.339	92.790	18.617
FRA				15.146	49.769	38.609	62.368	43.603	59.375
ITA					18.040	12.192	6.2284	28.422	51.058
JAP						42.652	41.906	49.538	3.8762
SWI							26.881	54.242	39.708
UK								79.449	52.388
USA									20.179

* The table reports P-values for testing the hypothesis that X_t as defined in (11)–(13) cannot be predicted using a vector Z_t consisting of a constant, domestic output, domestic prices and domestic population lagged one and two periods. Values in excess of 5.00 indicate failure to reject the null using a χ^2 test with 17 degrees of freedom.

AUS, Australia; CAN, Canada; FRA, France; ITA, Italy; JAP, Japan; SWI, Switzerland; UK, United Kingdom; USA, United States of America; WG, West Germany. FOD, first order differencing; HP, Hodrick and Prescott filtering.

Two main results stand out from Table 1. First, aggregate domestic consumption covaries with some of the instruments. Second, the strength of the association depends on the length of the cycle considered. In particular, the three orthogonality conditions are never rejected for short cycles (FOD detrended data), are rejected in approximately one-third of the cases for medium cycles (HP detrended data) and are rejected in all but three cases for long cycles (LT detrended data). Therefore, there appears to be more support for the hypothesis that consumption allocations across countries are optimal the shorter are the fluctuations in consumption we examine. When we ask which of the instruments is responsible for the rejection of the orthogonality conditions over medium-long cycles, we find that domestic consumption covaries primarily with demographic variables. This is not entirely surprising since the tests are performed with per-capita consumption data and population, which is highly serially correlated within each country, does not comove over the medium long run across the nine countries.

Since casual empiricism suggests that full risk sharing should be rejected, one may be inclined to believe that the results obtained with FOD and the majority of HP detrended data are dubious. In support of this one can claim that (the log of) aggregate consumption is well approximated by a random walk. That is, the cyclical component of aggregate consumption extracted with a FOD filter is, for most practical purposes, unpredictable given available information. Therefore, tests conducted with FOD detrended data tell us nothing about risk sharing. Similarly, since consumption is very smooth over medium cycles, variables with large cyclical fluctuations like output and prices are unlikely to be significant in explaining the time series properties of X_t. In this case the tests conducted with HP detrended data are weak: the orthogonality conditions will not be rejected frequently, since the instruments are not useful in predicting consumption movements over medium cycles.

We do not give much credence to these arguments for at least three reasons. First, the spectrum of FOD detrended consumption in many countries does not resemble the flat spectrum of a white noise. Second, the estimated coefficient on foreign aggregate consumption is significant in almost half of the cases when FOD detrended data is used so that the growth rate of consumption is not completely unpredictable. Third, foreign aggregate consumption always has a significant coefficient and both output and prices are significant in some cases when HP detrended data is used. In conclusion, it is incorrect to establish a direct relationship between the rejection of the risk sharing hypothesis and the univariate time series properties of consumption.

A second and related argument which may lead us to doubt the results obtained with FOD detrended data has to do with the type of consumption data we are using. Recall that because of data constraints, we are forced to use total aggregate consumption in testing international risk sharing. If consumer durables respond more to permanent fluctuations in the instruments, the high frequency results we obtained are spurious. In other words, if important components of aggregate consumption have some durability aspect, aggregate consumption will be acyclical over short cycles but will react to the state of the economy when the length of the fluctuations increases. While this fact constitutes a matter of concern, we believe it does not satisfactorily explain the phenomena because the pattern of cross-frequency results is robust to the choice of country pair despite the fact that the durable component of total consumption is not very homogeneous across countries.

A third argument which may lead us to doubt the results obtained with FOD and HP detrended data is the possibility that measurement errors are more of a problem at the higher frequencies of the spectrum. This could be the case, for example, if the collection of higher frequency data involves substantial costs. In this situation the results we obtain with FOD (and to some extent HP) detrended data are spurious since X_t may be highly contaminated by large and unpredictable measurement errors. We find this explanation somewhat more appealing but also difficult to quantify. As mentioned in Section 2.3 we can get some grasp of the properties of the measurement errors present by employing $\hat{Z}_t$, the unpredictable component of Z_t, in testing the orthogonality conditions. Because shocks in Z_t are exogenous to the agents' information set, they are more likely to be uncorrelated with the measurement error in X_t.

Table 2 presents the significance level of the J-test for the overidentifying restrictions using the three moment conditions (11) to (13) and a constant, current and lagged *shocks* to real, demographic and nominal domestic variables as instruments. Shocks are constructed as the residuals of a VAR(4) on domestic output, domestic population and domestic prices. The message of Table 2 is very strong: no matter how we detrend the data we always fail to reject the orthogonality conditions. Therefore, if measurement error exists it is likely to be stronger at low frequencies and more correlated across domestic variables than across consumption of different countries.

3.1.2. *Assessing the relevance of misspecifications.* To try and further understand why the orthogonality conditions are rejected at some frequencies but not others, we have conducted several experiments designed to check for three possible sources of misspecification. First, we want to study whether misspecification of the empirical relationship can explain the pattern of results of Table 1. The tests we have conducted so far examine the implications of the theory in its simplest form. As we argued, in more complicated setups international risk sharing implies that consumption across countries should be proportional after accounting for factors which are nonseparable in utility with consumption. If these factors are correlated with the instruments, the rejections we observe at low and medium frequencies may be the result of the misspecification of the orthogonality conditions. Prime candidates to enter the computation of the orthogonality conditions are government expenditure, if it partially substitutes for private consumption in utility, employment, if utility is nonseparable in consumption and leisure, and lagged consumption, if utility is time nonseparable in consumption. We investigate their relevance for patterns of rejections by adding them to the set of instruments. If empirical misspecification is the reason for the rejection of the orthogonality conditions, some of these new instruments should have significant coefficients. Surprisingly, none of the instruments appears to be relevant. Particularly important are the results concerning government consumption expenditure, which, for the nine countries we

TABLE 2
P-VALUES (IN PERCENTAGES) FOR THE TEST THAT $E[X_t | Z_t] = 0$:
SHOCK INSTRUMENTS*

	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
Linearly Detrended data									
AUS		81.9882	46.2925	97.6432	73.4789	65.6436	67.3339	65.1412	79.8557
CAN			65.6527	59.7208	42.5670	68.9452	21.3315	26.3213	42.7275
FRA				66.5922	44.5308	34.3820	24.6531	55.7311	31.3766
ITA					39.1124	53.3856	60.1932	37.9021	22.1858
JAP						74.5787	19.1094	93.1501	98.0725
SWI							25.9047	51.7344	10.1648
UK								2.9673	5.8525
USA									3.0361
HP Detrended data									
	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
AUS		26.8775	44.3816	81.6847	2.9654	45.9652	29.9638	33.5965	20.1338
CAN			62.5038	75.9179	71.3667	65.4429	45.6255	50.0499	22.6597
FRA				58.9114	60.5837	32.4957	93.5456	56.5847	32.8338
ITA					60.0372	17.886	45.1379	33.8598	70.8611
JAP						52.9332	63.5244	15.1089	76.8317
SWI							54.5571	65.6172	76.2277
UK								50.3640	15.0141
USA									10.7206
FOD Detrended data									
	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
AUS		67.8751	43.2496	52.1166	25.2209	68.4601	95.7350	83.7958	99.0131
CAN			7.4336	34.0989	79.2133	31.9034	56.5786	53.2041	42.6540
FRA				73.5163	11.8144	16.1363	54.6077	30.4766	55.0956
ITA					4.0657	18.2037	69.9286	66.8692	12.6185
JAP						60.3460	89.4216	67.3000	46.6399
SWI							21.0858	21.2397	18.6343
UK								46.1325	28.9546
USA									86.2626

* The table reports P-values for testing the hypothesis that X_t as defined in (11)–(13) cannot be predicted using a vector Z_t consisting of a constant, the current and one period lagged residuals of a VAR(4) on domestic output, domestic prices and domestic population. Values in excess of 5.00 indicate failure to reject the null using a χ^2 test with 17 degrees of freedom.
See Table 1 for definitions.

consider, is essentially acyclical. This may signal the lack of direct programs aimed at insuring private consumption through automatic movements in government expenditure and is in contrast with some of the results of Christiano and Eichenbaum (1992), who suggest that government expenditure is an important determinant of cyclical fluctuations in aggregate consumption in the U.S.A. In conclusion, misspecification of the X_t vector does not appear to be the reason for rejections obtained at the low frequencies of the spectrum.

Second, we are interested in knowing whether the results we have obtained depend on the choice of instruments or on some of the auxiliary assumptions made. To examine these possibilities we checked whether the timing of the instruments is important, conducting tests using current and lagged instruments and adding income to the instruments. Since current values of the instruments may be endogenous, it is possible that this may amplify the significance of measurement error. Therefore, this set of tests may further shed light on whether endogeneity and measurement errors are the primary reasons for the rejections. In addition, to examine whether the choice of variables is important, we have modified our proxies for real, demographic and monetary instruments and included net exports among the instruments. To check whether a country specific effect mattered for the results we also omitted a constant in constructing X_{1t}. Finally, to examine whether the sample period mattered, we excluded from the tests observations prior to 1974. The results of Table 1 appear to be very robust: none of the main qualitative conclusions is affected by any of our modifications and, even quantitatively, results are fairly similar.

Third, we would like to know if the pattern of rejections is sensitive to the choice of functional form for the utility of the representative agent. So far we have used a CRRA specification and it may be that a (locally) quadratic utility function is more appropriate (see Townsend 1994). With quadratic utility, international consumption risk sharing implies that the cyclical component of consumption (as opposed to the cyclical component of the logarithm of consumption) should be proportional across countries. In contrast to Mace (1991), we find that the pattern of rejections of the orthogonality conditions is not sensitive to such respecification.

In summary, none of the modifications we tried changed the essence of the results. With caveats concerning possible differential features of measurement error at different frequencies, the results indicate that risk sharing is strong at high frequencies of the spectrum, exists to some extent at standard business cycle frequencies and appears to be lacking at low frequencies. The question of interest is whether this finding can be rationalized with economic arguments. One could be that the longer is the average period of the fluctuation in consumption, the harder countries will find it to insure them since longer fluctuations may signal structural imbalances which would require wealth redistribution rather than risk sharing. Marcet and Marimon (1992) showed that, in an economy with participation constraints, agents will be able to insure only short run fluctuations so that domestic consumption will track permanent income in the medium-long run. It may be argued that the countries of the panel may not really qualify as countries for which the participation constraint is binding. However, debt repudiation need not ever occur as contracts are written so that countries will never find it optimal to repudiate. A version of the informationally constrained model analyzed by Phelan and Townsend (1991) may also lead to risk sharing in the short but not long run. Therefore, it may well be the case that, given existing institutions, short run fluctuations are much easier to insure than more persistent ones.

3.1.3. *Testing the cross-equation restrictions.* As discussed in Section 2, the theory of international risk sharing implies, apart from orthogonality conditions, a set of cross-equation constraints on the slopes of (11) to (13). These constraints can be used to perform an alternative and somewhat stronger test of the theory: the previous subsections examined if X_t can be predicted by the instruments. Here we check whether the three moment conditions contain the same information.

The results of the Wald tests are very strong. Regardless of the instruments used, the country pair and the length of the cycles considered, the cross equation constraints are always rejected. Hence, although there were instances where X_t

TABLE 3
ESTIMATES OF THE RATIO OF CRRA COEFFICIENTS
USING THE U.S. AS NUMERAIRE*

	Detrended Instruments			Shock Instruments		
Country	LT	HP	FOD	LT	HP	FOD
AUS	1.006	1.282	0.819	0.947	1.486	0.859
	(0.032)	(0.049)	(0.087)	(0.052)	(0.062)	(0.104)
CAN	0.739	0.857	0.786	0.715	1.049	0.777
	(0.036)	(0.059)	(0.082)	(0.075)	(0.087)	(0.094)
FRA	1.940	1.621	1.170	1.808	1.677	1.138
	(0.027)	(0.016)	(0.046)	(0.040)	(0.030)	(0.068)
ITA	1.912	1.179	1.632	1.305	1.117	1.526
	(0.056)	(0.042)	(0.037)	(0.090)	(0.063)	(0.067)
JAP	0.472	1.020	0.695	0.546	1.045	0.678
	(0.054)	(0.044)	(0.073)	(0.117)	(0.055)	(0.120)
SWI	2.581	1.492	1.132	2.879	1.521	1.303
	(0.017)	(0.039)	(0.039)	(0.035)	(0.071)	(0.072)
UK	0.940	0.689	0.660	0.940	0.689	0.625
	(0.038)	(0.073)	(0.089)	(0.094)	(0.106)	(0.114)
WG	0.571	0.830	0.509	0.571	0.830	0.618
	(0.022)	(0.042)	(0.049)	(0.034)	(0.058)	(0.048)

* The estimates are obtained from the variance restriction (12). Standard errors of the estimates are in parentheses.
See Table 1 for definitions.

appears to be unpredictable, we reject the stronger implication that the information contained in (11) to (13) is the same.

To gain some intuition on the reasons for this sound rejection, we examine whether estimates of the slope coefficient are sensitive to measurement errors in the instruments. Table 3 presents point estimates of the slopes obtained with (12) and their standard errors using basic Z_t or their unpredictable component. We present estimates obtained with (12) because they are more reasonable as they are constrained to be positive. Point estimates of these slopes obtained with the other two moment conditions appear in an appendix available on request. To provide an intuitive meaning to the estimates we present values obtained when one of the countries is the U.S.A. Because for the specification we tested the slopes should measure the ratio of CRRA coefficients of the representative agent of the two countries, and because more is known about the CRRA parameter for U.S.A. consumers, this approach may help to gauge the risk characteristics of the representative consumer of various countries. For example, looking at the first line of Table 3, the Australian representative agent is as risk averse as the American one with LT detrended data (1.006), slightly more risk averse with HP data (1.282) and slightly less with FOD data (0.819) when standard instruments are used to estimate the slope.

The impression that one gets from Table 3 is that, in general, the representative agent of half of the countries are less risk averse than the representative agent in the U.S.A. Moreover, estimates obtained with the two types of instruments are not statistically different. On the other hand, estimates appear to vary with the length of the fluctuation considered. Since the CRRA coefficients are among the deep parameters of the model, one would hope that slope estimates are invariant across frequencies. We tested this formally using a Wald test with mixed results. Except for France, Switzerland and the U.K., estimates of the ratio of CRRA coefficients obtained with LT and FOD are not significantly different. On the other hand, except for the case of U.K. and Switzerland, estimates of the ratio of CRRA coefficients obtained with HP and FOD are significantly different. Finally, estimates of the ratio of CRRA coefficients obtained with HP and LT are significantly different in all cases. As a qualification, one should note that even though the null hypothesis is rejected, the range of estimates obtained is economically small, except perhaps for Switzerland.

Obstfeld (1989) has also provided estimates of the risk aversion parameter for the U.S.A., Germany and Japan and found that over the sample 1961–1985 they are significantly different from each other, while over two subsamples (1961–1972, 1973–1985) they are not. The point estimates he reports are comparable to ours but the standard errors are very different: Obstfeld's estimated standard errors are in fact 100 times larger than ours. This indicates that the instruments he chose may contain little information to gauge the risk characteristics of the three countries.

3.1.4. *Examining the magnitude of consumption correlations.* As we have discussed at length in Section 2, cross country correlations need not be equal to one even when risk sharing is perfect. However, the results we have obtained so far imply that nonseparabilities between consumption and other factors in utility are not particularly important so that, apart from measurement errors, consumption profiles should be perfectly correlated across countries. Clearly, the rejection of the cross equation restrictions suggests that consumption correlations are unlikely to be perfect. However, it may still be of interest to examine the size of the correlations implicit in the estimated slope of (13) first, to formally test whether they are different from one and, second, to assess whether their magnitude may explain some of the differences between our results and existing work.

Table 4 presents the estimated correlation and their standard errors when Z_t includes a constant, lagged domestic output, prices and population. Three major facts stand out. First, estimated correlations are significantly different from one in almost all cases, regardless of the length of cycle considered. This finding confirms the informal analyses of Backus, Kehoe and Kydland (1992) and Devereux, Gregory and Smith (1992). Second, consumption correlations are stronger for pairs of E.E.C. countries, a result which suggests that countries with closer economic ties may also have more efficient risk sharing mechanisms. Third, consumption correlations obtained when the U.K. is one of the countries are negative with LT detrended data, a result which suggests the presence of high durability content in U.K. consumption data (see also Blackburn and Ravn 1992).

3.2. *Long-Run Implications.* The risk-sharing proposition is an appropriate idealization to describe a mechanism which insures cross-country aggregate consumption from domestic *cyclical* fluctuations. However, it is possible to investigate the risk-sharing proposition even when country specific fluctuations have a permanent component, so long as it is common to all countries. In this case the theory

TABLE 4
ESTIMATED CONSUMPTION CORRELATIONS*

	AUS	CAN	FRA	ITA	JAP	SWI	UK	USA	WG
				Linearly Detrended data					
AUS		0.767	0.387	0.246	0.809	0.154	−0.514	0.726	0.783
		(0.027)	(0.042)	(0.047)	(0.014)	(0.063)	(0.041)	(0.038)	(0.223)
CAN			0.638	0.393	0.629	0.560	−0.226	0.624	0.809
			(0.044)	(0.047)	(0.035)	(0.056)	(0.071)	(0.032)	(0.020)
FRA				0.671	0.409	0.472	−0.184	0.292	0.586
				(0.042)	(0.054)	(0.061)	(0.061)	(0.055)	(0.057)
ITA					0.194	0.449	−0.112	−0.103	0.274
					(0.023)	(0.050)	(0.041)	(0.057)	(0.041)
JAP						0.367	−0.225	0.443	0.826
						(0.066)	(0.031)	(0.036)	(0.015)
SWI							0.232	0.426	0.383
							(0.064)	(0.052)	(0.053)
UK								−0.113	−0.331
								(0.059)	(0.044)
USA									0.551
									(0.035)
				HP Filtered data					
AUS		0.271	0.237	0.250	0.232	0.024	0.115	0.047	−0.194
		(0.065)	(0.083)	(0.087)	(0.072)	(0.083)	(0.061)	(0.067)	(0.068)
CAN			0.401	0.244	0.039	0.455	0.392	0.607	0.060
			(0.069)	(0.051)	(0.058)	(0.059)	(0.049)	(0.041)	(0.071)
FRA				0.172	0.428	0.439	0.508	0.580	0.258
				(0.075)	(0.061)	(0.091)	(0.054)	(0.090)	(0.097)
ITA					0.317	0.468	0.395	0.060	0.205
					(0.068)	(0.058)	(0.056)	(0.080)	(0.096)
JAP						0.287	0.552	0.408	0.229
						(0.080)	(0.065)	(0.065)	(0.080)
SWI							0.429	0.451	0.511
							(0.053)	(0.059)	(0.061)
UK								0.366	0.125
								(0.058)	(0.080)
USA									0.340
									(0.074)
				FOD Filtered data					
AUS		0.419	0.403	0.070	0.251	0.269	0.046	0.313	0.070
		(0.068)	(0.068)	(0.097)	(0.073)	(0.093)	(0.078)	(0.066)	(0.066)
CAN			0.297	0.133	0.071	0.344	0.089	0.531	0.133
			(0.068)	(0.112)	(0.068)	(0.103)	(0.080)	(0.059)	(0.066)
FRA				0.065	0.299	0.355	0.169	0.162	0.305
				(0.071)	(0.068)	(0.073)	(0.092)	(0.057)	(0.064)
ITA					0.090	0.355	0.094	−0.038	0.048
					(0.094)	(0.118)	(0.078)	(0.101)	(0.098)
JAP						0.035	0.077	0.267	0.091
						(0.068)	(0.048)	(0.046)	(0.053)
SWI							0.154	0.172	0.122
							(0.061)	(0.084)	(0.053)
UK								0.189	0.149
								(0.057)	(0.068)
USA									0.233
									(0.077)

* Standard errors are in parentheses.
See Table 1 for definitions.

requires (i) that the aggregate consumption profile of pairs of countries should move together in the long run (i.e., the vector X_t must be stationary) and (ii) that the error in predicting the logarithm of the marginal utility of consumption of country j, given the logarithm of the marginal utility of consumption of country k, should be unpredictable using long run characteristics of country j. Note that, even if the long-run component of consumption differs across countries, a certain type of risk sharing may still take place, but, as pointed out in the literature (see Melitz and Vori 1992), it involves a permanent system of transfers across countries.

An analysis of the long-run implications of international risk sharing may also shed some light on a related issue which is of interest to macroeconomists. Following the permanent income tradition, it is typical to identify the permanent component of income with the long-run behavior of consumption (see Quah 1990). Our analysis may therefore shed light on the relationship among permanent incomes of different countries and on the issue of convergence of national incomes, a topic which has received substantial attention in the current growth literature (see Barro and Sala-i-Martin 1992).

We first examine the implication that X_t is stationary. This is a very weak implication since it does not involve the additional restriction of unpredictability given country specific characteristics. To test this restriction we retain the assumption that the instantaneous utility is of CRRA type, that the marginal utility of consumption is independent of factors other than consumption and examine whether X_t contains a unit root. We conduct tests for integration and cointegration in aggregate consumption data using Phillips and Perron's (1986) Z-test[3] and we find that we cannot reject the hypothesis that aggregate consumption is an I(1) variable, except for Switzerland. For each pair of countries for which consumption is I(1), we run a regression like $\log c_{jt} = a + b * \log c_{kt} + e_{jt}$ and check if the residuals are stationary.[4] Table 5, panel A reports the largest (in absolute value) t-statistics over possible lag augmentations (up to a maximum of 10 lags), and the lag value at which it occurs. The results show that the regression error appears to contain a unit root in many cases. The exceptions are the pairs Australia–Japan, Japan–France, Australia–France in addition to few borderline cases. Hence, consumption patterns tend to permanently diverge and even this weaker restrictions implied by risk sharing in the long run is generally unsupported.

For those pairs of countries for which X_t is stationary we next examine whether it is unpredictable. Table 5, panel B reports the results of the J-test: risk sharing is

[3] For sensitivity we also conducted tests using the Dickey and Fuller (1979) and Stock and Watson (1989) tests. The three procedures produce the same answer for all series except for French private consumption where the Stock and Watson test rejects the null of a unit root. We also experimented with different lag augmentations and also excluded the linear trend from the regression but this change did not alter in any way the essence of the results. Tables with the results are available on request from the authors.

[4] Because of the poor quality of consumption data and the relative short sample available we also considered an approach which imposes that pairs of consumption series are cointegrated with cointegrating vectors $[1, -1]$ (so that all countries have the same attitude toward risk) and examine whether or not the residuals display a unit root. Although some of the results change, the basic message of the exercise is maintained.

TABLE 5a
t-STATISTICS OF PHILLIPS AND PERRONS'S Z-TEST*

	AUS	CAN	FRA	ITA	JAP	UK	USA	WG
AUS		−2.37 (10)	−3.88 (6)	−3.31 (6)	−4.40 (6)	−1.69 (10)	−2.88 (8)	−3.39 (8)
CAN			−1.89 (10)	−2.15 (10)	−2.37 (10)	−1.37 (10)	−2.16 (10)	−3.30 (10)
FRA				−3.39 (6)	−3.66 (10)	−1.93 (10)	−2.32 (10)	−3.46 (10)
ITA					−2.98 (6)	−1.63 (10)	−2.02 (10)	−3.19 (10)
JAP						−2.97 (2)	−2.29 (10)	−3.07 (10)
UK							−1.59 (10)	−1.65 (10)
USA								−2.23 (10)

TABLE 5b
P-VALUES (IN PERCENTAGES) FOR THE TEST THAT $E[X_t|Z_t]=0$:
LONG RUN CYCLES*

	AUS	CAN	FRA	ITA	JAP	UK	USA	WG
AUS			0.2636	0.0000	0.1043			0.0123
CAN								0.0000
FRA				0.2837	0.1265			0.2567
ITA								0.0000
JAP								0.0000
SWI								0.9871

* Phillips and Perron's Z-test checks if the residuals of a cointegrating regression between pairs of consumptions are stationary. The table reports the largest t-statistics over possible lag augmentations of the test. The number in parentheses refers to the lag augmentation used (maximum value is 10). The test is run with no deterministic variables. Panel B reports p-values for testing the hypothesis that X_t as defined in (11) is unpredictable when Z_t consisting of a constant, domestic output, domestic prices and domestic population lagged one and two periods. Values in excess of 5.00 indicate failure to reject the null using a χ^2 test with 17 degrees of freedom. Empty cells indicate that no test is undertaken for that pair.
See Table 1 for definitions.

rejected in the long run in all cases and domestic consumption covaries in the long run with all domestic instruments. Two interesting conclusions can be derived from these results. First, contrary to the implications of the basic neoclassical growth model and, consistent with the findings of Canova and Marcet (1995), permanent incomes across countries show little signs of convergence. Second, if any insurance scheme is in place in the long run, it must involve a permanent redistribution of income across countries.

4. CONCLUSIONS

In this paper we have formally analyzed international aspects of the theory of risk sharing. We showed that risk sharing is a property of optimal international consumption allocations in a wide variety of theoretical setups and empirically examined three testable implications of the theory. We find that aggregate domestic consumption covaries with demographic variables once foreign consumption is taken into account over long but not over medium and short cycles. In addition, we find evidence that measurement errors present in both consumption data and in the instruments may be responsible for these results. We strongly reject the other two implications of the theory, namely the cross equation constraints and the implication that consumption correlations are perfect and show that little risk sharing appears in the long run. As a by-product of the analysis, we have provided novel estimates of the CRRA coefficients for the representative agent of the countries in the panel (relative to the U.S.A.), which can be used by researchers interested in calibrating international business cycle models.

Our results are in agreement with those of Obstfeld (1993), Atkeson and Bayoumi (1993), Lewis (1993) and Kollman (1995). We find some evidence of risk sharing in international consumption data, but contrary to them, we are also able to distinguish the implications of the theory which are at odds with the data. Our conclusions are also in line with those of Backus, Kydland and Kehoe (1992) and Devereux, Gregory and Smith (1992): consumption correlations are both statistically and economically different from one but that they are higher among European countries suggesting that more risk sharing is taking place within the E.E.C.

Although the analysis we have conducted did not specify the market structure supporting the optimal allocations, there are at least two implications of our results which are useful in designing mechanisms intended to implement optimal allocations. Our analysis has shown that, whatever they are, existing market structures are able to shield domestic consumption from business cycle shocks in the short run (recall that government consumption expenditure does not play a major insurance role). This result seems to question the need of government intervention both in terms of providing automatic stabilizers in the economy and new institutions to help the economies to achieve optimal allocations. However, this statement needs two qualifications. First, the countries we examined are among the most industrialized of the world and for some of them temporary labor migrations, semi-permanent remittance programs from emigrants or other individual schemes may have created strong insurance which need not to be present for other OECD countries or even less for LDC countries. Second, because domestic aggregate consumption covaries with some domestic variables over cycles of 8–10 years, there is room to improve the quality of consumption allocations by designing institutions (or opening markets) which insure agents against this type of fluctuation. However, insuring fluctuations of longer average period may have marked redistributive effects.

A second issue which is of interest concerns the integration of domestic financial markets into a world market and its effect on domestic business cycle fluctuations. Our results suggest that, because that existing structures offer sufficient insurance against the most interesting sources of domestic business cycle fluctuations, the welfare gains obtained by opening foreign financial markets to domestic consumers may be small. Hence, the prospective integration of European financial markets and the current globalization of security markets need not bring substantial changes in the features of domestic consumption fluctuations.

Universitat Pompeu Fabra, Spain, Universita' di Modena, Italy, and CEPR, U.K.
University of Southampton, U.K.

APPENDIX

The data we use is taken from Datastream. Consumption measures aggregate private consumption expenditure on nondurables, durables and services. It is transformed into a per-capita series by dividing the original series by population. Because data on population is annual, quarterly data are obtained by taking the predicted values of an AR(3) regression fitted to a dummy quarterly series, constructed assigning the annual value to each of the four quarters. Government data measures current government expenditure except in the case of Australia where also gross government fixed investment is included. Income data measures gross domestic product (GDP) except for Japan, U.S.A. and West Germany where it measures gross national product (GNP). All data is in real terms. The base year however, differs across countries. For Australia, Italy, Japan, U.K., and West Germany the base is 1985, for France and Switzerland the base is 1980, for Canada the base is 1986 and for the U.S.A. the base is 1987. All variables are measured in annual rates. Employment data is not completely compatible since it measures different aggregates in different countries. The series used measure total employment in Canada, Italy, Japan, Switzerland, U.K., U.S.A., West Germany, employment on the payroll in Australia and civilian employment in France. Because some of these series are non-seasonally adjusted, we deseasonalize them by using an exponential smoothing procedure. U.K. employment data are the same as those employed by Blackburn and Ravn (1992). Finally, the price data measures the implicit price deflator for GNP (or GDP).

The sample covers the period 1960,1–1991,4 for Australia, Canada, United Kingdom, U.S.A. and Germany; the period 1965,1–1991,4 for Japan; the period 1967,1–1991,4 for Switzerland; and the period 1970,1–1991,4 for France and Italy. The tests are constructed using the shortest data samples for each pair.

REFERENCES

ASDRUBALI, P., B. SORENSEN, AND O. YOSHA, "Channels of Interstate Risk Sharing: US 1963–1990," working paper 95-13, Brown University, 1995.

ATKESON, A., "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica* 59 (1991), 1069–1084.

——— AND T. BAYOUMI, "Do Private Capital Markets Insure Regional Risk? Evidence from the United States and Europe," *Open Economies Review* 4 (1993), 303–354.

BACKUS, D., P. KEHOE, AND F. KYDLAND, "International Real Business Cycles," *Journal of Political Economy* 101 (1992), 745–775.

———, ———, AND ———, "International Business Cycles: Theory and Evidence," T. Cooley, ed., *Frontiers of Business Cycle Research* (Princeton: Princeton University Press, 1995), 331–356.

BARRO, R. AND X. SALA-I-MARTIN, "Convergence," *Journal of Political Economy* 101 (1992), 233–252.

BAXTER, M. AND M. CRUCINI, "Explaining Savings/Investment Correlations," *American Economic Review* 83 (1993), 416–436.

——— AND ———, "Business Cycles and the Asset Structure of Foreign Trade," *International Economic Review* 36 (1995), 821–854.

BENHABIB, J., R. ROGERSON, AND R. WRIGHT, "Homework in Macroeconomics II: Aggregate Fluctuations," *Journal of Political Economy* 99 (1992), 1166–1187.

BLACKBURN, K. AND M. RAVN, "Business Cycles in the UK: Facts and Fictions," *Economica* 59 (1992), 383–401.

BRENNAN, M. AND B. SOLNICK, "International Risk Sharing and Capital Mobility," *Journal of International Money and Finance* 8 (1989), 359–373.

CANOVA, F., "Detrending and Business Cycle Facts," CEPR working paper 781, 1993.

——— AND A. MARCET, "The Poor Stay Poor: Non-Convergence Across Countries and Regions," mimeo, Universitat Pompeu Fabra, 1995.

CHRISTIANO, L. AND M. EICHENBAUM, "Real Business Cycle Theories and Aggregate Labor Market Fluctuations," *American Economic Review* 82 (1992), 430–450.

COCHRANE, J., "A Simple Test of Consumption Insurance," *Journal of Political Economy* 99, (1991), 957–976.

COLE, H. AND M. OBSTFELD, "Commodity Trade and International Risk Sharing: How Much Do Financial Markets Matter?" *Journal of Monetary Economics* 28 (1991), 3–24.

CONSTANTINIDES, G., "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98 (1991), 519–543.

CRUCINI, M., "On the International and National Dimension of Risk Sharing," mimeo, Ohio State University, 1994.

DEVEREAUX, M., A. GREGORY, AND G. SMITH, "Realistic Cross Country Consumption Correlations in a Two-Country, Equilibrium, Business Cycle Model," *Journal of International Money and Finance* 11 (1992), 3–16.

——— AND G. SMITH, "International Risk Sharing and Economic Growth," *International Economic Review* 34 (1994), 535–550.

DICKEY, D. AND W. FULLER, "Distribution of the Estimators for a Autoregressive Time Series with a Unit Root," *Journal of The American Statistical Association* 74 (1979), 424–431.

DUFFIE, D. AND C. F. HUANG, "Implementing Arrow-Debreu Equilibria by Continuous Trading of Few Long Lived Securities," *Econometrica* 53 (1985), 1137–1356.

ENGLE, R., "Band Spectrum Regression," *International Economic Review* 15 (1974), 1–11.

FRENCH, K. AND J. POTERBA, "Investors Diversification and International Equity Markets," *American Economic Review, Papers and Proceedings* 81 (1991), 222–226.

HARVEY, A. AND A. JEAGER, "Detrending, Stylized Facts and the Business Cycle," *Journal of Applied Econometrics* 8 (1993), 231–247.

KING, R. AND S. REBELO, "Low Frequency Filtering and Real Business Cycles," *Journal of Economic Dynamics and Control* 17 (1993), 207–231.

KOLLMAN, R., "Consumption, Real Exchange Rates and the Structure of International Asset Markets," *Journal of International Money and Finance* 14 (1995), 191–211.

LEWIS, K., "What Can Explain the Apparent Lack of International Consumption Risk Sharing?" mimeo, University of Pennsylvania, 1993.

LUCAS, R., "Interest Rates and Currency Prices in a Two Country World," *Journal of Monetary Economics* 10 (1982), 335–360.

MACE, B., "Full Insurance in the Presence of Aggregate Uncertainty," *Journal of Political Economy* 99 (1991), 928–956.

MARCET, A. AND R. MARIMON, "Communication, Commitment and Growth," *Journal of Economic Theory* 58 (1992), 219–249.

——— AND K. SINGLETON, "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints," mimeo, Universitat Pompeu Fabra, 1992.

MARRINAN, J., "Government Consumption and Private Consumption Correlations," mimeo, Universitat Pompeu Fabra, 1994.

MELITZ, J. AND S. VORI, "National Insurance Against Unevenly Distributed Shocks in a European Monetary Union," mimeo, INSEE, 1992.

OBSTFELD, M., "How Integrated Are World Capital Markets? Some New Tests," in Calvo, G., et. al., eds., *Debt, Stabilization and Development: Essays in the Memory of Carlos Diaz Alejandro*, (Oxford: Basic Blackwell, 1989).

———, "International Risk Sharing and Capital Mobility: Another Look," *Journal of International Money and Finance* 11, 1992, 115–121.

———, "Are Industrial Country Consumption Risks Globally Diversified?" in L. Lederman and A. Razin, eds., *Capital Mobility: The Impact on Consumption, Investment and Growth*, (Cambridge: Cambridge University Press for the CEPR, 1993).

———, "Risk Taking, Global Diversification and Growth," *American Economic Review* 84 (1994), 1310–1330.

PADOA-SCHIOPPA, T., *Efficiency, Stability and Equity*, (Oxford: Oxford University Press, 1987).

PERSSON, T. AND G. TABELLINI, "Federal Fiscal Constitutions. Part I: Risk Sharing and Moral Hazard," CEPR working paper 728, 1992.

PHELAN, C. AND R. TOWNSEND, "Computing Multiperiod Information Constraint Optima," *Review of Economic Studies* 58 (1991), 853–881.

PHILLIPS, P. C. AND P. PERRON, "Testing for Unit Roots in Time Series Regressions," *Biometrika* 70 (1986), 335–346.

QUAH, D., "Permanent and Transitory Movements in Labor Income: An Explanation Based on Excess Smoothness in Consumption," *Journal of Political Economy* 98 (1990), 449–475.

STOCK, J. AND M. WATSON, "Interpreting the Evidence on Money Income Causality," *Journal of Econometrics* 40 (1989), 161–181.

STOCKMAN, A. AND H. DELLAS, "International Portfolio Nondiversification and Exchange Rate Variability," *Journal of International Economics* 26 (1989), 271–289.

TESAR, L., "International Risk Sharing and Nontraded Goods," *Journal of International Economics* 35 (1993), 69–90.

——— AND I. WERNER, "Home Bias and High Turnover," *Journal of International Money and Finance* 14 (1995), 236–251.

TOWNSEND, R., "Risk and Insurance in Village India," *Econometrica* 62 (1994), 539–591.

VAN WINCOOP, E., "Welfare Gains from International Risk Sharing," *Journal of Monetary Economics* 34 (1994), 175–200.

WILSON, R., "The Theory of Syndicates," *Econometrica* 36 (1968), 119–132.