

Crucini, Mario J; Hess, Gregory D

Working Paper

International and Intranational Risk Sharing

CESifo Working Paper, No. 227

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Crucini, Mario J; Hess, Gregory D (1999) : International and Intranational Risk Sharing, CESifo Working Paper, No. 227, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/75610>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CESifo *Working Paper Series*

INTERNATIONAL AND INTRANATIONAL RISK SHARING

Mario J. Crucini
Gregory D. Hess*

Working Paper No. 227

December 1999

CESifo
Poschingerstr. 5
81679 Munich
Germany
Phone: +49 (89) 9224-1410/1425
Fax: +49 (89) 9224-1409
<http://www.CESifo.de>

* This paper has been prepared for *Intranational Macroeconomics*, edited by Gregory D. Hess and Eric van Wincoop, Cambridge University Press, NY, forthcoming, 2000. We would like to thank Marco del Negro and Eric van Wincoop for extremely helpful comments.

INTERNATIONAL AND INTRANATIONAL RISK SHARING

Abstract

This paper reviews three important issues in the literature on international and intranational risk sharing. First, we establish a comprehensive set of stylized facts for consumption risk sharing within and across countries. Consistent with the findings in the literature, we find that the correlation of our consumption measure across countries is much less than that for output. As pointed to by Backus, Kehoe and Kydland (1993), this constitutes a violation of complete international risk sharing, which they labeled the 'quantity anomaly'. The findings using international data are contrasted with those using intranational data for Japanese prefectures, U.S. states and Canadian regions. Consistent with the findings by Crucini (1998) and Hess and Shin (1997, 1998), intranational data continues to demonstrate the quantity anomaly for the U.S., but not for Japan and Canada. Second, following the work by Crucini (1998), we estimate an econometric specification which allows us to quantify the extent of risk sharing within and across countries. The results indicate that, while still incomplete, a larger fraction of risk is shared within countries than across them. Finally, using these estimates of the current extent of risk sharing, we calculate the potential welfare benefits from additional international and intranational risk sharing.

JEL Classification: E20, E31, E32

Mario J. Crucini
Vanderbilt University
Department of Economics
Box 1819, Station B
Nashville, TN 37235
USA

email: Mario.J.Crucini@Vanderbilt.edu

Gregory D. Hess
Oberlin College
Department of Economics
233 Rice Hall
Oberlin, OH 44074
USA

email: gregory.hess@oberlin.edu

1 Introduction

According to the theory of aggregate risk sharing, if individual households have access to a complete market for financial assets then they can, by pooling together their risk, insure fully against the idiosyncratic uncertainty in their resources. See, among others, Townsend (1987), Altug and Miller (1990), Cochrane (1991) and Mace (1991). In such an environment, theory predicts that marginal utility is perfectly correlated across agents. Under quite general assumptions the theory also predicts that household changes in consumption should move one-for-one with aggregate changes in consumption absent idiosyncratic fluctuations in preferences or measurement error. Importantly, differences in consumption changes across households should not be correlated with changes in a household's resources.¹

Recently, dynamic general equilibrium macroeconomic models have been extended to the open economy to study the international co-movement of aggregate variables. The first generation of these models were quite successful in explaining two aspects of the data: a countercyclical trade balance (Backus, Kehoe and Kydland (1994)), and a high correlation between savings and investment despite perfect capital mobility (Baxter and Crucini (1993)). One startling rejection of one of the model's predictions, identified by Backus, Kehoe and Kydland (1993), spawned a large literature as to whether and by how much countries share risk.² The rejection of the theory that these authors point to is as follows: In a two-country standard dynamic general equilibrium model where only exogenous productivity shocks drive fluctuations in output, since individuals in each country can share risk internationally and thereby diversify their nation specific risk, the cross-country correlation of consumption is predicted to be much higher than the cross-country correlation of output or productivity shocks. In the data, however, they find that the opposite is true, and labeled this phenomena the *quantity anomaly*.

More recently, Obstfeld (1994) studied the cross correlations of consumption and output growth rates between an individual country and the rest of the world as well as between OECD countries. He reported that although the cross correlation of consumption growth rates have risen since 1973, they are always lower than those for output. Bayoumi and McDonald (1995) found using consumption and income based tests that only Japanese consumption was fully integrated with the rest of the world. See also Canova and Ravn (1995) and Hess and Shin (1997) for recent international evidence for the OECD. Lewis (1996) found that consumption growth had a smaller common component in countries that imposed capital controls compared to countries that did not, suggesting that some risk

¹This risk sharing can take place through formal market institutions, or through less formal ones that may apply to families or rural villages. Risk sharing within families has been explored by Hayashi, Altonji and Kotlikoff (1996) and at the village level by Townsend (1995).

²The second major discrepancy between theory and data that these authors point out is that the terms of trade is dramatically more volatile in the data than the models suggest. The relationship between intranational and international price movements is discussed in the chapter by Engel and Rogers and in Hess and Shin (1997,1998).

sharing is taking place at least among countries that do not explicitly restrict international capital mobility.

In order to better understand the reasons why economic regions such as countries may only share risk incompletely, researchers have attempted to re-examine this important issue across economic regions within a country. So-called intranational evidence provides a natural experiment for understanding international puzzles since regions within a country have reduced barriers to trade in goods and assets while also sharing a common currency and legal framework. Therefore, by using within country data, we can help to identify whether the lack of consumption risk sharing pointed to by Backus, Kehoe and Kydland is due to international factors, such as capital restrictions, and exchange rate uncertainty, or is merely a matter of geography independent of national boundaries. The intranational approach to risk sharing has been used by Atkeson and Bayoumi (1993), and Asdrubali, Sorensen and Yosha (1996), Crucini (1998), Hess and Shin (1997,1998,2000) and Athanasoulis and van Wincoop (1998) to explore the quantity anomaly within regions and states of the United States. Bayoumi and McDonald (1994), Bayoumi and Klein (1995), and Crucini (1998) also analyze the quantity anomaly for Canadian provinces, while van Wincoop (1995) examines Japanese regions. Collectively, these studies suggest that the quantity anomaly may be reversed in the case of the Canadian provinces, but remains in force for the Japan prefectures (marginally), and for the U.S. states.

The goal of this survey paper is to re-examine consumption risk sharing, within and between countries, using a number of data sets from recently published papers. We provide empirical findings using the following data sets: for the U.S. – Asdrubali, Sorensen and Yosha (1997); for Canada – Crucini (1998); for Japan – van Wincoop (1995); and for the OECD – Hess and Shin (1997). These data sets are discussed in more detail below (see Section 3.1).

We examine three important issues in the risk sharing literature. First, in Section 3.2 we establish a comprehensive set of stylized facts for intranational risk sharing and international risk sharing. Second, in Section 3.3, following the work of Crucini (1998), we estimate an econometric model which helps us to pin down the extent of risk sharing. This is an important step given that a simple rejection of perfect risk sharing based on stylized facts does not by itself help researchers to uncover the extent to which risk is not being shared. Finally, as the ultimate reason for risk sharing is to reduce risk and thereby enhance expected welfare, in Section 3.4 we present calculations for the welfare implications of risk sharing and the welfare implications of regional business cycles.

2 Risk Sharing Theory

In this section we derive testable implications of risk pooling across regions from two models. The first is an endowment economy with complete Arrow–Debreu markets analogous

to that studied by Mace (1991). The second model follows Crucini (1998), allowing for intertemporal smoothing at a constant real interest rate but varying degrees of risk pooling across regions. Many of the qualitative predictions of these models hold up in more elaborate production economies such as in Backus, Kehoe, and Kydland (1993) and Baxter and Crucini (1995).

2.1 Complete Risk Sharing

In this section we derive testable implications from a model in which risk is shared across R individuals, each a representative agent for her region of residence, taken to be either a country (international risk sharing) or a region within a country (intranational risk sharing). The model considered below closely follows the Arrow-Debreu complete markets model studied by Mace (1991). Each region i is assumed to consist of a representative household. Uncertainty is summarized by the state variable s_τ , $\tau = 1, 2, \dots, S$ that takes on S different values in each time t . The term $\pi_t(s_\tau)$ captures the probability of state s_τ occurring at t with $\sum_{\tau=1}^S \pi_t(s_\tau) = 1$. Expected lifetime utility of the representative household in each country or region is given by:

$$\sum_{t=0}^{\infty} \beta^t \sum_{\tau} \pi_t(s_\tau) U(c_{it}(s_\tau), b_{it}(s_\tau)) \quad (1)$$

where $c_{it}(s_\tau)$ is consumption for the representative household in region i at time t when the state is s_τ , and $0 < \beta < 1$ is the subjective discount factor, assumed to be the same across households. $b_{it}(s_\tau)$ is an exogenous preference shock to region i , given the state of nature. Each region i receives the endowment of $y_{it}(s_\tau)$ at time t under state (s_τ) , so that the feasibility constraint for this economy is $\sum_{i=1}^R y_{it}(s_\tau) = \sum_{i=1}^R c_{it}(s_\tau)$.

Suppose that risk is shared completely across all R regions. Then the resulting allocation is obtained by the following Lagrangian problem solved by the social planner at each time t :

$$\mathcal{L} = \sum_{i=1}^R \omega_i \sum_{t=0}^{\infty} \beta^t \sum_{\tau} \pi_t(s_\tau) U(c_{it}(s_\tau), b_{it}(s_\tau)) + \mu_t(s_\tau) \left[\sum_{i=1}^R y_{it}(s_\tau) - \sum_{i=1}^R c_{it}(s_\tau) \right] \quad (2)$$

where the μ 's are the Lagrangian multipliers. The planner assigns time-invariant weights ω_i such that $0 < \omega_i < 1$ and $\sum_{i=1}^R \omega_i = 1$. The first order conditions are:

$$\hat{\mu}_t(s_\tau) = \omega_i U_c(c_{it}(s_\tau), b_{it}(s_\tau)) \quad i = 1, \dots, R, \quad (3)$$

where $\hat{\mu}_t(s_\tau) \equiv \mu_t(s_\tau) / \beta^t \pi_t$ and μ_t is the Lagrange multiplier associated with the aggregate feasibility constraint. To understand the implications of the theory, notice that the left hand side of (3) is only a function of time and the aggregate state, and hence is the same for all regions. This multiplier is the current shadow value of an extra unit of resources at time t , given state s_τ .

To understand the broader empirical predictions of the model, assume that the representative individual has the following CRRA utility function:

$$U(c_{it}(s_\tau), b_{it}(s_\tau)) = \exp((1 - \gamma)b_{it}) \left(\frac{1}{1 - \gamma} \right) (c_{it})^{(1 - \gamma)}, \gamma > 0;$$

where γ is the coefficient of constant relative risk aversion. The first order condition for consumption of individual i can then be simplified by taking logs, aggregating and first differencing so that household consumption behaves according to:

$$\Delta \log(c_{it}) = \Delta \log(c_t^a) + \left(\frac{1 - \gamma}{\gamma} \right) \Delta(b_{it} - b_t^a), \quad i = 1, \dots, R \quad (4)$$

where $b_t^a = (1/T) \sum_{i=1}^R b_{it}$, and $\log(c_t^a) = \frac{1}{R} \sum_{i=1}^R \log(c_{it})$. Hence for CRRA utility, changes in consumption of each individual move one-for-one with the change in average consumption holding fixed the preference shock. To note, with exponential (CARA) utility the relationship holds in levels rather than log-levels – see Mace (1991).

Under perfect risk sharing, consumption growth should move one-for-one across regions, *ceteris paribus*. Of course, exogenous shocks to preferences could drive regional changes in consumption as well as measurement error even with perfect risk sharing. Importantly though, under complete regional risk sharing the differences in consumption changes across regions should be unrelated to their individual resources as the benevolent social planner is assumed to share resources across all regions. This simple intuition has lead to the following important empirical proposition for testing regional risk sharing (both within and across countries).

Proposition 1

- Let the observed change in consumption growth, $\Delta \log(c_{it})$, be the sum of true consumption growth, $\Delta \log(c_{it})^*$, measurement error, $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$, and an idiosyncratic preference shock, $\nu_{it} \sim N(0, \sigma_\nu^2)$; namely, $\Delta \log(c_{it}) = \Delta \log(c_{it})^* + \epsilon_{it} + \nu_{it}$
- Let the observed change in income growth, $\Delta \log(y_{it})$, be the sum of true income growth, $\Delta \log(y_{it})^*$, and measurement error, $\eta_{it} \sim N(0, \sigma_\eta^2)$, such that $\Delta \log(y_{it}) = \Delta \log(y_{it})^* + \eta_{it}$.
- Assume that observed consumption growth and income growth have variances that are constant across regions and time, $\sigma_{\Delta \log(c)}^2$ and $\sigma_{\Delta \log(y)}^2$, respectively.

The theory of perfect risk sharing predicts that true consumption growth should be more correlated across countries than true income growth:

$$\text{Corr}(\Delta \log(c_{it})^*, \Delta \log(c_{jt})^*) > \text{Corr}(\Delta \log(y_{it})^*, \Delta \log(y_{jt})^*) \quad i \neq j. \quad (5)$$

In the presence of measurement error and idiosyncratic preference shocks, support for this prediction can be inferred from observed correlations if the volatilities of measurement error and preference shocks satisfy:

$$\frac{\text{Corr}(\Delta \log(c_{it}), \Delta \log(c_{jt}))}{\text{Corr}(\Delta \log(y_{it}), \Delta \log(y_{jt}))} > \frac{[1 - \{(\sigma_\epsilon^2 + \sigma_\nu^2)/\sigma_c^2\}]}{[1 - (\sigma_\eta^2/\sigma_y^2)]} \quad i \neq j \quad (6)$$

The major thrust of Proposition 1, that consumption should be more correlated across countries than output, is the fundamental prediction from the theory that was posed by Backus, Kehoe and Kydland (1993). This is stated in expression (5). However, since measurement error and preference shocks can affect the observed correlations, expression (6) clarifies how ‘cleanly’ must measured consumption growth be relative to income growth – the right hand side – in order for any observed ratio of cross correlations in consumption and income – the left hand side – to provide evidence of (5). Notice that the right hand side numerator of (6) is the ratio of the true volatility of consumption to its observed volatility. This ratio will rise if there is more measurement error (σ_ϵ^2) in consumption, or if preference shocks become larger (σ_ν^2). This expression also indicates why it is so difficult to determine whether measurement error or preference shocks are behind the finding of the low cross correlation of observed consumption as they both enter symmetrically. Similarly, the right hand side denominator of (6) is the ratio of the true volatility of income to its observed volatility. This ratio will rise if there is more measurement error in income (σ_η^2).

The role of preference shocks and measurement error in interpreting consumption and income correlations to assess the theory of complete risk sharing has been explored in a number of papers. In a two-country model with tradable and non-tradable goods, Stockman and Tesar (1995) find that by introducing taste shocks which are 85 percent as large as productivity shocks, the model is more consistent with data. However, the trade balance is no longer countercyclical in their model. Moreover, using exogenous taste shocks in addition to exogenous, large and persistent productivity shocks does not appear to leave much room for theory to explain open economy fluctuations. Devereux, Gregory and Smith (1992) show that when preferences are non-separable between consumption and leisure, leisure operates like a non-tradable good, and low cross-country consumption correlations can result. The formalization of the role of measurement error is discussed in Hess and Shin (1998).

2.2 Incomplete Risk Sharing

Models that feature less than complete risk sharing are actually the rule, rather than the exception, in international macroeconomics. The most obvious example is the work-horse small open economy model with the marginal rate of substitution equated to the world real interest rate (often assumed to be constant) while the opportunity set changes over time, possibly due to changes in income, productivity, government spending, or taxes. In this simple model, as there is only assumed to be one asset rather than a more complete set of financial instruments (e.g., equity markets for domestic and foreign securities), consumption growth will contain both aggregate and idiosyncratic components even in the absence of taste shocks or measurement error.

Based on this insight, Baxter and Crucini (1995) demonstrate that the observed deviations from complete risk sharing will be larger the more persistent are the output shocks that each nation (or region) faces. Using an open economy model with a risk free bond as the only asset (i.e. incomplete markets), they show that the more persistent are each countries' shocks the less cross-correlated consumption will be. The intuition is that highly persistent shocks with incomplete markets will make the shadow value of each country's resources (i.e. their permanent incomes) diverge more so that, in turn, consumption will have a smaller common component.

In an attempt to gauge the extent or scope of risk sharing, Crucini (1998) and Hess and Shin (2000) consider specific environments of incomplete risk sharing. We follow Crucini's (1998) methodology since, using regional level data, it nests the Mace (1991) consumption specification with minimal modification. Specifically, agents have access to borrowing and/or lending opportunities at a fixed interest rate, r , that allows them to shift resources across time periods. The fraction of their stochastic income stream that the representative agent in each region chooses to pool in turn determines the flow of income that is available for intertemporal smoothing. Proposition 2 describes the assumptions and the resulting consumption function.

Proposition 2

- *If the representative agent in each region only pools a fraction, λ , of their income with the remaining fraction, $(1 - \lambda)$, of their income is not pooled; and*
- *If agents are able to borrow or lend at a risk free rate $r = \beta^{-1} - 1$ to smooth their consumption paths according to the permanent income hypothesis; and*
- *There are no region specific preference shocks.*

Then consumption changes by region will approximately follow the process:³

$$\Delta \log(c_{it}) = \lambda \Delta \log(c_t^a) + (1 - \lambda)(1 - \beta) \sum_{k=0}^{\infty} \beta^k [E_t \log(y_{it+k}) - E_{t-1} \log(y_{it+k})] \quad (7)$$

The interpretation of this equation is straightforward using permanent income reasoning. Agents change consumption one-for-one with changes in permanent income which is a λ -weighted average of changes in their own permanent income and changes in the permanent income of others engaged in the risk sharing arrangement. At one extreme, if regions completely share risk then $\lambda = 1$ and consumption changes move one-for-one with the aggregate consumption change. At the other extreme, if regions do not risk share at all then $\lambda = 0$, and each region's consumption change will be driven by the revisions in their own permanent income. We utilize a variant of this specification in the next section to estimate the extent of risk sharing across regions using the various panel datasets.

2.3 Welfare Implications

Finally, while Propositions 1 and 2 are helpful to understanding the degree of risk sharing across regions, an essential question to ask is what are the potential welfare benefits from risk sharing. This question is fundamental to placing the earlier results in perspective since if the potential gains from risk sharing are small, then results which suggest that there is only limited risk sharing observed would be of less interest since the implied loss in welfare would be insignificant.

In an influential study on the costs of business cycles, Lucas (1987) asked the question: What fraction of a constant stream of consumption would a representative agent be willing to pay to avoid fluctuations in consumption that are associated with the business cycle? Adopting a CRRA utility function: $U(c_{it}) = \left(\frac{1}{1-\gamma}\right) (c_{it})^{(1-\gamma)}$, he demonstrated that the welfare cost of consumption variability is equal $g(\sigma_z^2) \approx \frac{1}{2} \gamma \sigma_z^2$, where σ_z^2 is a measure of the variability of consumption measured as the variance of an i.i.d. shock around a deterministic trend. He calculated the welfare losses to be less than one tenth of a percentage point of consumption assuming consumption growth was trend stationary, even with very low rates of intertemporal substitution (i.e. very high values for constant relative risk aversion).

Lucas's approach has been criticized for only allowing temporary shocks to consumption's level rather than permanent ones as well (e.g. see Cole and Obstfeld (1991) and van Wincoop (1994)).⁴ van Wincoop (1994) demonstrated that if consumption growth were

³See Crucini (1998) for the derivation.

⁴van Wincoop (1994,1998) also demonstrates that the issue of the correct underlying statistical model for the shocks to output (i.e. the underlying risk that each region wishes to diversify) and the specification of preferences can dramatically affect the calculations for the extent of potential gains from risk sharing. See also Dolmas (1998) on this last point.

a random walk, as the simple permanent income hypothesis would suggest, then the gain measured in terms of the percent change in welfare associated with the complete elimination of consumption variability is $-0.5 \cdot \gamma \cdot \sigma_{\Delta \log(c)}^2 / (r - \bar{\mu})$, where $r = (\beta^{-1} - 1) + \gamma \bar{\mu}$ is the risk free rate, $\bar{\mu} = \mu - (\gamma/2)\sigma_{\Delta \log(c)}^2$ is the economy's risk adjusted growth rate, and μ is the economy's expected growth rate.

These welfare exercises provide the basis for a useful benchmark for the extent of risk sharing faced by individuals within the context of a representative agent economy. Following the specification of welfare gains embodied in van Wincoop (1994), we wish to understand the potential gain over an infinite horizon from sharing risk across regions and the cost of regional business cycles based on a utility derived calculation in the spirit of Lucas (1987). These issues have already been examined in a number of other papers. At the international level, Cole and Obstfeld (1991) found only small gains from international risk sharing (.2 percent of output per year), comparable to those found by Lucas for the cost of U.S. business cycles. Using a more general setup, van Wincoop (1994) found much larger benefits from international risk sharing: 1.8 - 5.6 percent welfare gains were still obtainable. Recently, for U.S. states, Athanasoulis and van Wincoop (1998) calculate the permanent increase in income associated with a welfare improvement by which all regional output risk is shared. They find that over a five and fifteen year horizon, the welfare gain is .7 percent and 1.4 percent, respectively.

3 Results

This section is sub-divided into four sub-sections. The first describes the data sets from published studies that we use in our study. The second provides descriptive statistics of this data, while the third provides estimates of the extent of risk sharing across regions. The fourth sub-section provides estimates of the welfare gains from exhausting the as yet unexploited gains from regional risk sharing.

3.1 The Data

In this sub-section we describe the international and intranational data sets we use to test our propositions on risk sharing. These data sets have been used in recently published studies and cover a range of countries (OECD) as well as regions within the following countries: Canada, the United States and Japan. The usefulness of using these alternative data sets is that we can re-explore the extent to which regions share risk in a common framework and better gauge the robustness of currently published findings.

The Canadian income and consumption data is from Crucini (1998). The original source is the Provincial Economic Accounts published by Statistics Canada. The income data is

gross domestic product while the consumption data is total consumption of goods (both durables and non-durables) and services by the private sector. Provincial deflators for final demand are used to deflate both gross domestic product and consumption. The Canadian panel contains data for the 10 provinces and a Canadian aggregate. The sample period is from 1971 to 1991.

For the USA, we use the state level income and consumption data from Asdrubali, Sorensen, and Yosha (1996). The income measure for the United States is also gross domestic product but the consumption data are based on retail sales. The original source of the output data is the Bureau of Economic Analysis – see Beemiller and Dunbar (1993) for details. The consumption data is from the Survey on Buying Power published by Sales Management. Asdrubali, Sorensen, and Yosha gross-up the original consumption data by the difference between national consumption and retail sales measure (both durables and non-durables). State and local government consumption is added to this figure to arrive at state consumption. The U.S. aggregate CPI is used to convert the data to constant dollar terms.⁵ The panel includes all 50 states plus the District of Columbia. The sample period is from 1963 to 1990.

The Japanese data, from van Wincoop (1995), includes gross domestic product and total private consumption (for non-durables plus durables) for 47 prefectures. The main source of the data is the Annual Report on Prefectural Accounts produced by the Economic Planning Agency. Gross domestic product deflators are available for each prefecture but consumption deflators are available at a more aggregative level (10 districts). The sample period is from 1975 to 1990.

Finally, the income and consumption data employed by Hess and Shin (1997) is used for the international dimension of our study. The original source is the World Data CD-ROM 1995, produced by the World Bank. Output is nominal gross domestic product divided by the gross domestic product deflator. Consumption is private consumption of goods (both durables and non-durables) and services, also deflated by the gross domestic product deflator. While Hess and Shin (1998) found for U.S. states and regions that the inclusion of durable goods raises the cross correlation of consumption as well as its time series volatility, we do not further explore this issue in this paper due to data availability.

3.2 Descriptive Statistics

We begin with a few simple descriptive statistics. The raw data have been transformed to real per capita terms and we use the log-differences of the variables in what follows. To ensure maximal comparability across panels we use the following sample periods in our

⁵See Del Negro (1998) for issues in constructing state level CPI indices from city level reports.

subsequent analysis: Canada, 1972 to 1991; Japan, 1975 to 1990; United States, 1972 to 1990; and OECD countries, 1972 to 1990.

To get at the issue of regional comovement we compute two cross-regional correlations. The first is the correlation of regional consumption or income growth with *aggregate* consumption or income growth. The second is the correlation of regional income or consumption growth with a *reference* region. We compute an aggregate national variable for the Japanese prefectures and U.S. states by taking a sum of the real levels of domestic product (or consumption) across regions. The aggregate is then transformed to real per capita growth using the same transformation applied to the regional data. The Canadian provincial data includes a Canadian aggregate so we use this construct for Canada. Note that the regional data and national aggregates are not reconciled with national income and product accounts so these aggregates can differ somewhat from the OECD data for Canada, Japan, and the United States. The second measure of regional comovement is the correlation of the growth rate in each region with a reference region. We chose Ontario, Tokyo, Ohio, and the United States as our reference regions for the Canadian, Japanese, U.S., and OECD panels, respectively. Table 1 reports these regional comovement measures averaged across all regions in the panel along with panel averages of the standard deviation of regional consumption and output, and the cross correlation of regional consumption with regional output.

[– Table 1 Around Here –]

Beginning with the standard deviation of output we see similar values of 2.28 and 2.30 for the Japanese prefectures and OECD countries while both the U.S. states and Canadian provinces have much larger, but comparable, standard deviations: 4.60 and 4.67, respectively. Turning to the volatility of consumption, the U.S. is the clear outlier with an average standard deviation of consumption equal to 5.63 percent per year compared to less than 2.5 for the other panels. Recall that the U.S. consumption data differs from the other panel data in that it is measured by retail sales as opposed to actual consumption. The greater volatility of retail sales is consistent with the finding in Hess and Shin (1998) that for aggregate U.S. data, total retail sales is about 10% more volatile than total consumption of durable and non-durable goods. Moreover, the volatility of the consumption measure falls if one includes only the retail sales of non-durable goods, which is available only for 19 states since 1978.⁶ We are at a loss to explain the differences in output volatility across the panels.

⁶Hess and Shin (1998) exclusively analyze these states and find that while the volatility of non-durables is much lower than that for durables, the cross region correlation is much lower for non-durables than it is for durables. They report that non-durable retail sales are approximately one-quarter less volatile than total retail sales. As non-durable consumption data for regions and countries is generally not available, we do not explore this issue further in this chapter.

The third column of Table 1 reports the cyclicalities of regional consumption as measured by its contemporaneous correlation with regional output. Consumption is procyclical with respect to regional economic activity in all cases, though the magnitudes differ. The average correlation of consumption and output growth across the OECD is 0.69, quite close to the values of 0.64 for the Canadian provinces, and 0.53 for the U.S. states. The typical Japanese prefecture, however, has a consumption–output correlation of only 0.30.

The comovement across regions for both output and consumption are presented in the final four columns of Table 1. We first discuss the comovement of regional output and consumption with their aggregate counterparts which are reported in columns four and five. We see that the output comovement is highest across the U.S. states, averaging 0.74, and lowest across the Japanese prefectures, where the correlation averages 0.56. Given that we are only using about 15 years of data, these differences are not statistically significant. The consumption comovements are more distinct across panels and do not line up with the income comovement patterns. Canadian provincial consumption growth is very highly correlated, averaging 0.80, while the average correlation of the Japanese prefectures with the aggregate is only 0.52.

Figure 1 presents the full set of bilateral comparisons with the correlation of consumption growth of each pair (a region and the aggregate) plotted against the corresponding correlation of income growth. According to the theory of aggregate risk sharing (see Proposition 1), the data points should lie about the 45-degree line. We see that consumption correlations lie uniformly above income correlations in the Canadian panel (all points lie on or above the 45-degree line) while for the United States and the OECD countries income correlations often exceed consumption correlations (the majority of points lie below the 45-degree line). The Japanese prefectures give a more ambiguous picture.

The correlation patterns using the reference regions in the final two columns of Table 1 typically yield lower cross–regional correlations which is what one would expect given that the aggregate contains the regional variable by construction. Tokyo was chosen based on its economic size and appears more idiosyncratic than Ontario, Ohio, or the United States. In terms of the *quantity anomaly* pointed to by Backus, Kehoe and Kydland (1993), we see that the average correlation of output with aggregate or reference output is distinctly higher than the average correlation of consumption with aggregate or reference consumption for the OECD and U.S. cases, but the reverse is true of the Canadian provinces. The ranking of the correlations for the Japanese prefectures (at least based on averages across regions) is unclear, as the relative ranking of the magnitudes of the correlations switches depending on whether the aggregate or reference region is used in constructing the correlation.

From expression (6), one can calculate how large must be the magnitude of measurement error and preference shocks in order for the true correlation of consumption to be greater than that for output. If we assume that output is measured without error, i.e. $\sigma_\eta^2 = 0$, then consumption growth would have to be approximately fifteen percent noise and preference

shocks $(1.0 - (.63/.74))$ in order to make the true correlation of consumption equal to that for output across states within the U.S for durable plus non-durable goods. Hess and Shin (1998) report that measurement error and preference shocks would have to be over fifty percent of measured consumption for only non-durable goods.

An additional factor to consider in the analysis of cross correlations of consumption and output is that we would expect to observe comovement in consumption and income across countries even in the absence of international trade and financial market linkages to the extent that technological changes and policy changes move in a synchronous fashion. An initial attempt to consider the simultaneity in the data proceeds as follows. Let z_{it} be either regional consumption or regional income growth and let z_t^a is either the growth rate for the aggregate variable or the reference variable in the panel. We then estimate the following regression as a time series for each individual region within the panel:

$$z_{it} = \alpha_i + \beta_i z_t^a + \epsilon_{it} \quad (8)$$

We refer to the ratio of the explained sum of squares to total sum of squares from equation (8) as the common component and the remainder (unexplained variation) as the idiosyncratic component.⁷ Table 2 reports the results, again averaging across the regions in each panel.

[– Table 2 Around Here –]

Consistent with the earlier findings reported in Table 1, the states within the U.S. have (on average) an exceptionally strong common component in output, and a smaller common component in consumption. For both panels A and B, this relative ranking of common output versus common consumption component is also shared with the OECD countries, but is reversed for Canadian regions. Once again, for Japan, the relative rankings are ambiguous.

Figure 2 presents the full complement of consumption–income correlations between each region and the aggregate using raw and common component adjusted data. The x-coordinates are raw correlations of income growth and consumption growth within each region or country while the y-coordinates are correlations of income and consumption after removing a common component. After controlling for possible common movements in income and consumption across regions we see a decrease in the income and consumption comovements within countries as compared to across OECD countries – i.e. the within country correlation pairs generally lie below the 45 degree line for Canada, Japan and the

⁷See Del Negro (1997) for a much more sophisticated decomposition into common and idiosyncratic components based on factor analysis.

U.S. while the across country correlation pairs for the OECD lie nearly on the 45 degree line. However, even within regions and after the adjustment, the correlations tend to be strongly positive.

3.3 Risk Sharing Specification

We extend the results of estimating the risk sharing specification to the regional Japanese data to see if the conclusion that the extent of risk sharing is greater within than across countries that Crucini (1998) found for the U.S. states and Canadian provinces compared to the G-7 is also true of Japanese prefectures. However, we also employ different data and countries in the estimation process for the U.S. regions and country level data as a robustness check (the results for the Canada provinces are taken directly from Crucini (1998)).

The U.S. data employed by Crucini (1998) differs in a number of respects from the data employed by Asdrubali, Sorensen, and Yosha (1996), which we use below. Crucini (1998) deflates consumption by the aggregate U.S. deflator for non-durable goods and services deflates and uses state personal income as the income measure deflated by the aggregate U.S. GNP deflator. Asdrubali, Sorensen and Yosha deflate both consumption and income by the aggregate U.S. consumer price index and their income measure is gross domestic product. The last difference between the estimation conducted here and in Crucini (1998) is that we include 21 OECD countries as opposed to only the G-7.

To estimate the incomplete risk sharing equation (7) from Crucini (1998), we need to make an assumption as to the process for the log of permanent income. Accordingly, we assume that innovations to the log of permanent income are reasonably proxied by the change in income (i.e. annual income is close to a random walk).⁸ We estimate the following equation separately, region-by-region as a time series, for each region in the panel.

$$\Delta \log(c_{it}) = \alpha_i + \lambda_i \Delta \log(c_t^a) + (1 - \lambda_i) \Delta \log(y_{it}) + u_{it} \quad (9)$$

where the variables are now in logarithms and the variable $\Delta \log(c_t^a)$ is a simple average of the consumption growth rates of all the regions in the panel. The error term is interpreted as a combination of measurement error and/or taste shocks, the latter of which would have to be uncorrelated with the determinants of aggregate consumption or regional permanent income.

[– Table 3 Around Here –]

⁸Crucini (1998) examines the robustness of this assumption and finds that the risk sharing parameter estimates are somewhat sensitive to the estimated process of regional income.

The estimation results for equation (9) are reported in Table 3. The first three columns report, for each panel, the average estimated value for λ , the average standard error of the estimate for λ , and the average R^2 , respectively. The final four columns of the table give an indication of the distribution of the estimates of λ within each panel, which may be more informative than simply the statistical mean. We see that the risk sharing parameters average more than 0.8 across the Canadian provinces, U.S. states, and Japanese prefectures, compared to about one-half this magnitude for the OECD countries. These results reaffirm Crucini's (1998) findings that the degree of risk sharing is higher across regions within countries than across industrialized countries, although it is less than perfect in all cases. The complete risk sharing model cannot be rejected in 6 of 10 provinces, 29 of 51 states, and 41 of 47 Japanese prefectures.⁹ In contrast, the null can be rejected in all but 3 of the 21 OECD countries. The results for Japan are even stronger than for the United States in the sense that while both have a similar number of economic regions, a significant number – 12 – of the states have risk sharing parameter estimates with 95% confidence intervals overlapping both 0 and 1. Japan has only one such region.

The estimated risk sharing parameters for the United States and the Japanese prefectures appear somewhat in contrast to what the simple correlation patterns of consumption suggest about the extent of risk sharing across regions. It appears likely that measurement error in consumption (and/or taste shocks) are partly responsible for the differences since these components of regional consumption variance would be cast into the residual term in the regression equation (assuming they are uncorrelated with regional income).¹⁰ Alternatively, specification error in the construction of the revision of permanent income $(1 - \beta) \sum_{k=0}^{\infty} \beta^k [E_t y_{it+k} - E_{t-1} y_{it+k}] = \Delta \log(y_{it})$ under the random walk assumption, will lower the contribution from the region-specific component and increase the contribution from the unexplained component. Measurement error in the construction of regional income data will also lower the region-specific component. Finally, further specification error could be due to the presence of a fraction of consumers who neither share risk nor consume out of permanent income, but rather consume according to a 'rule of thumb' such that consumption deviates from permanent income considerations.

In a complementary study, Hess and Shin (2000), using household food consumption data for the U.S. from the Panel Study for Income Dynamics, estimate a regression equation similar to (9). They estimate the fraction of risk that is shared within regions, within industries, and across regions and industries (i.e. aggregate). Importantly, rather than using income data to proxy regional and/or industry changes in permanent income, they use, according to the theory, constructed consumption changes by region and industry for this purpose. They report that only about 45 percent of risk is shared across regions and industries while the remainder is shared within regions and industries.

⁹The finding that a large proportion of U.S. states do not perfectly share risk accords well with Hess and Shin (1998) in their results for only 19 states with consumption measured by the retail sales of non-durables.

¹⁰This view is also emphasized by Sorensen and Yosha (1998).

3.4 Welfare Gains

We explore the welfare implications of regional and national fluctuations in consumption in this section. Adopting a CRRA utility function: $U(c_{it}) = \left(\frac{1}{1-\gamma}\right) (c_{it})^{(1-\gamma)}$, we follow the approach taken by van Wincoop (1994), assuming consumption is a random walk. This leads to an increase in the welfare implications of fluctuations relative to what Lucas (1987) found since he assumed consumption reverted to a deterministic trend. As discussed in section 2.3, the formula for a welfare gain associated with a drop in the standard deviation of consumption growth by $d\sigma_{\Delta c}^2$ is:

$$-0.5\gamma d\sigma_{\Delta c}^2 / (r - \bar{\mu}) \quad (10)$$

where $\bar{\mu} = \mu - (0.5\gamma\sigma_c^2)$ is the risk adjusted growth rate and $r = (\beta^{-1} - 1) + (\gamma \cdot \bar{\mu})$. We set $\gamma = 4$, $\bar{\mu} = 0.017$, and $\beta = .99$, which are the same values used by van Wincoop (1994) except for the coefficient of relative risk aversion (he used unity). The implied level of r is 7.8%.

To calculate the welfare gain from removing the remaining region specific variability of consumption growth we need a good measure of the latter. Fortunately, the incomplete risk sharing specification, equation (9), provides a theoretically appealing variance decomposition of consumption growth that relates to the portfolio choices implied by the estimated risk sharing parameters. Using equation (9) we apportion the variability of regional consumption into a common part, $(\hat{\lambda}_i^2 \text{var}(\Delta \log(c_t^a)))$, a part related to variation in regional income, $((1 - \hat{\lambda}_i)^2 \text{var}(\Delta \log(y_{it})))$, the covariance between the two, $2\hat{\lambda}_i(1 - \hat{\lambda}_i) \text{cov}(\Delta \log(c_t^a), \Delta \log(y_{it}))$, and an unexplained part which is just the variance of the regression residual ($\sigma_{u_i}^2$). Importantly, the final three components can be summed to obtain that part of consumption variability that is not accounted for by aggregate consumption movements (i.e. regional consumption growth net of the volatility of shared consumption growth— $\text{var}(\Delta \log(c_{it})) - \hat{\lambda}_i^2 \text{var}(\Delta \log(c_t^a))$), and hence it is these terms which we will focus on below in our welfare calculations.

Table 4 reports the welfare results using the variance decomposition implied by the risk sharing regression. Specifically we ask how much agents in each region would be willing to give up in terms of average consumption to eliminate the variance of regional consumption growth that is not accounted for by movements in average consumption. The first component we identify is the welfare gain from removing the variance associate with regional income growth variation (see column denoted ‘Regional PI’). The second term in the calculation is the welfare term is proportional to the covariation between regional income growth and aggregate consumption (see column denoted ‘Covariance’), and the final

welfare gain is from the variance of the residual (see column denoted ‘Unexplained’).¹¹ The final column, denoted ‘Total’, reports the sum of the three components. The total should be interpreted as an upper bound on the welfare gains since it includes the regression residual.¹²

Each column of Table 4 reports the welfare implications (in percentages of steady-state consumption) of moving each of these variances to zero. Due to the fact that we have regions that are heterogeneous with respect to their risk sharing parameters (λ_i) and underlying variances of permanent income, we calculate the welfare measures separately for each region and then report the mean welfare cost (across all regions in a panel) and the mean welfare cost for the quintile with the highest or lowest total welfare benefit from eliminating consumption risk not already shared.

[– Table 4 Around Here –]

Three things should be noted in comparing our results with Lucas’s. First, he considered only the welfare gains from removing aggregate consumption variability, whereas we make our calculations for removing region specific consumption variability after conditioning on aggregate consumption. Second, we assume that the fluctuations to consumption’s level are permanent, whereas Lucas assumed that fluctuations to consumption were i.i.d. around a deterministic trend. Most importantly, however, Lucas was interested in the reduction of consumption variability brought about by a reduction in output variability. To be consistent with Lucas, the increased welfare brought about by the removal of the regional output cycle would be presented under the column ‘regional PI’ in Table 4. In contrast, the risk sharing literature emphasizes the reduction in idiosyncratic consumption variability. This is related to the final column of Table 4 presented under the header ‘Total’ as in this case we have controlled regional consumption only for aggregate consumption.¹³

The results in Table 4 suggest that the average Canadian and Japanese regions would be willing to give up about 1% of consumption to eliminate their total regional risk in consumption. The estimate rise to about 1.5% for the mean OECD nation. The mean U.S. state would be willing to pay 7% of consumption, though almost all of this (5.73%) is attributable to the error term in the estimated risk sharing regression. Some of this might

¹¹The columns headed by ‘Regional PI’, ‘Covariance’ and ‘Unexplained’ are calculated as $-0.5 \cdot \gamma \cdot d\sigma_{\Delta c}^2 / (r - \bar{\mu})$ multiplied by $((1 - \hat{\lambda}_i)^2 \text{var}(\Delta \log(y_{it})))$, $2\hat{\lambda}_i(1 - \hat{\lambda}_i)^2 \text{cov}(\Delta \log(c_t^a), \Delta \log(y_{it}))$, and the variance of the residual, $(\sigma_{u_i}^2)$, respectively.

¹²In effect, when using the first column to compute welfare changes, we are assuming that all of the residual variance is diversifiable risk when in reality it is some unknown combination of regional permanent income variation (that is not captured by regional income growth), taste shocks, and measurement error. The number does take into account risk already shared however since the variance attributable to aggregate consumption growth is excluded from the calculation.

¹³The welfare calculations when $\lambda = 1$ is imposed provides slightly higher values than those reported in Table 4, since those reported in the Table are calculated based on estimates of λ that minimize the unexplained sum of squares.

well be captured by a more carefully specified permanent income process but some is also likely measurement error in which case it should not be included as a potential welfare improving change.

Focusing on the variance that is attributed to regional income variation the welfare measures are quite modest for average Canadian and Japanese region and exactly equal to one another at 0.15%. However, there is a greater variance of welfare benefits cross-sectionally in Japan than in Canada with the highest quintile prepared to give up about 0.5% of consumption compared to only 0.32% for the top quintile among the Canadian provinces.

Turning to the U.S. states we see comparable welfare implications of regional income variability to what is observed internationally. The large possible gains in the OECD are easy to understand because they have a low risk sharing parameter which gives regional income movements a larger influence in regional consumption variability. The larger possible gains in the U.S. case suggests that while risk sharing is more extensive, in the sense of pooling of resources, the benefits to additional pooling are greater because state level income is more variable than national income. The puzzle that remains is why more risk sharing does not take place in the United States if the costs of doing so are comparable to what exists in Canada and Japan. We discuss some possibilities in our concluding section.

4 Conclusions

In this survey, we have provided evidence that while both intranational and international risk sharing are imperfect, the data suggests that there is more of the former than of the latter. We have not discussed, however, the channels by which this intranational risk sharing has taken place. While using different methodologies, though the same data sources for U.S. states, Asdrubali, Sorensen and Yosha (1996), Del Negro (1997) and Athanasoulis and van Wincoop (1998) measure the extent of risk sharing that takes place through financial markets via assets, through fiscal policy tax and transfer schemes, and the component of risk that is left unshared.¹⁴ The range of estimates are remarkably similar. Approximately 10 - 20 percent of shocks to gross state product are shared through fiscal tax-transfers, approximately 45 - 60 percent of risk sharing is achieved through financial markets, and 25 -35 percent is left ‘unsmoothed’ or unshared.

¹⁴The methodologies they employ differ quite substantially. Asdrubali, Sorensen and Yosha (1996) implement their idea of ‘smoothing’ to calculate the percentage of a shock that is smoothed through capital markets (income smoothing), savings adjustment (consumption smoothing), the federal tax-transfer system, and the remainder is left unsmoothed. Del Negro (1997) uses factor analysis to decompose output and income fluctuations and then relates these to consumption growth. Athanasoulis and van Wincoop (1998) base their calculations on a 26 year horizon that distinguishes between risk shared through financial markets and those through fiscal net-transfers.

The fact that we observe incomplete risk sharing across geographic regions has also lead researchers to uncover the source of this incompleteness. One approach is that due to private information or merely preference, individuals simply prefer to hold local assets even within a country. Coval and Moskowitz (1997) and Huberman (1997) find strong evidence of a local preference for equities based on portfolio decisions both for fund managers and individuals. Alternatively, some shocks to resources may be hard to diversify. A good candidate for these types of shocks are those that affect human capital which may be embodied in the industry that individuals are attached to. Using state level data for the U.S., Sorensen and Yosha (1997) explore the basis of industry related movement in resources. Using family level evidence, Hess and Shin (2000) find that industry related movements in resources are not perfectly shared across regions. We anticipate that future work in the area of risk sharing will attempt to further our understanding of how market incompleteness manifests itself both within and across countries. The issue of whether unexplained idiosyncratic fluctuations in measured consumption are better interpreted as measurement error or preference shocks also requires additional research.

REFERENCES

- Altug, S. and Miller, R., 1990, Household Choices in Equilibrium, *Econometrica*, 1990, 82, 1177-1198.
- Asdrubali, P., Sorensen, B.E. and Yosha, O., 1996, Channels of interstate risksharing: United States 1963-1990, *Quarterly Journal of Economics* 111, 1081-1110.
- Athanasoulis, S. and van Wincoop, E., 1998, Risksharing within the United States: What have financial markets and fiscal federalism accomplished?, Federal Reserve Bank of New York, Research Paper 9808.
- Atkeson, A. and Bayoumi, T., 1993, Do private capital markets insure regional risk? Evidence from the United States and Europe, *Open Economies Review*, 4, 303-324.
- Backus, D.K., Kehoe, P.J. and Kydland, F.E., 1992, International real business cycles, *Journal of Political Economy* 84, 84-103.
- Bayoumi, T. and McDonald, R., 1995, Consumption, income and international capital market integration, *IMF Staff Papers*.
- Baxter, M. and Crucini, M., 1993, Explaining saving/investment correlations, *American Economic Review*, 83, 416-36.
- Baxter, M., and Crucini, M.J., 1995, Business cycles and the asset market structure of foreign trade, *International Economic Review*, 36, 821-854.
- Beemiller, R. and Dunbar, A., 1993, Gross state product, 1977-90, *Survey of Current Business*, 73, 1993, 28-49.
- Canova, F. and Ravn, M., 1996, International Consumption Risk Sharing, *International Economic Review*, 37, 573-601.
- Cochrane, J. H., 1991, A simple test of consumption insurance, *Journal of Political Economy* 99, 957-76.
- Cole, H., and Obstfeld, M., 1988, Commodity Trade and International Risk Sharing: How much do financial markets matter?, *Journal of Monetary Economics*, 28, 3-24.
- Coval, J.D. and Moskowitz, T., 1997, Home Bias at Home: Local Equity Preferences in Domestic Portfolios, mimeo, University of Michigan Business School.
- Crucini, M.J., 1998, On international and national dimensions of risk sharing, *Review of Economics and Statistics*, forthcoming.
- Del Negro, Marco , 1997, Aggregate Risksharing Across US States, mimeo.

- Devereux, M., Gregory, A., and Smith, G., 1992, Realistic cross country consumption correlations in a two country, equilibrium, business cycle model, *Journal of International Money and Finance* 11, 3-16.
- Dolmas, J., 1998, Risk Preferences and the Welfare Cost of Business Cycles, *Review of Economic Dynamics* 1, 646–676.
- Hayashi, F., Altonji, J. and Kotlikoff, L., 1996, Risk Sharing Between and Within Families, *Econometrica*, 64, 261-294.
- Hess, G.D. and Shin, K., 1997, International and Intranational Business Cycles, *Oxford Review of Economic Policy*, 13, 93-109.
- Hess, G.D. and Shin, K., 1998, Intranational business cycles in the United States, *Journal of International Economics*, 44, 289-313.
- Hess, Gregory and Shin, K., 2000, Risk sharing by households within and across regions and industries,” *Journal of Monetary Economics*, forthcoming.
- Huberman, G. , 1997, Familiarity Breeds Investment, mimeo, Columbia Business School.
- Lewis, K., 1996, What can explain the apparent lack of international consumption risk sharing, *Journal of Political Economy*, 104, 267–297.
- Lucas, Robert E. Jr., (1987), “Models of Business Cycles” Basil Blackwell: Oxford, United Kingdom.
- Mace, B., 1991, Full insurance in the presence of aggregate uncertainty, *Journal of Political Economy*, 99, 928–956.
- Obstfeld, M., 1994, Are industrial-country consumption risks globally diversified?, in *Capital Mobility: the Impact on Consumption, Investment and Growth*, eds. Leonardo Leiderman and Assaf Razin, Cambridge University Press, Cambridge, U.K.
- Sorensen, B. and Yosha, O., 1997, Income and consumption smoothing among U.S. states: regions or clubs, mimeo.
- Sorensen, B. and Yosha, O., 1998, International risk sharing and european monetary unification, *Journal of International Economics*, forthcoming.
- Statistics Canada, 1991, Provincial Economic Accounts, Annual Estimates 1961 to Date. Ministry of Supply and Services Canada.
- Stockman, A. and Tesar, L., 1995, Tastes and technology in a two-country model of the business cycle: explaining international comovements, *American Economic Review* 85, 168-85.

- Townsend, R., 1995, Consumption insurance: an evaluation of risk-bearing systems in low-income countries, *Journal of Economic Perspectives*, 9, 83–102.
- United States Department of Commerce, Bureau of Economic Analysis. State Personal Income: 1929 – 1987: A Statement of Sources and Methods.
- United States Department of Commerce, Bureau of the Census. Statistical Abstract of the United States. Washington D.C.: U.S. Government Printing Office.
- van Wincoop, E., 1994, Welfare gains from international risksharing, *Journal of Monetary Economics*, 1994, 34, 175-200.
- van Wincoop, E., 1995, Regional risksharing, *European Economic Review* 37, 1545-1567.
- van Wincoop, E., 1998, How big are the potential gains from international risksharing? *Journal of International Economics*, forthcoming.

Table 1. – Properties of Output and Consumption Across Panel Data Sets

Panel	Standard Deviation		Correlation of Output and	Correlation with Aggregate		Correlation with Reference Region	
	Output	Cons.	Cons.	Output	Cons.	Output	Cons.
Canada	4.67	2.26	0.64	0.63	0.80	0.57	0.75
Japan	2.28	2.00	0.30	0.56	0.52	0.27	0.42
United States	4.60	5.63	0.53	0.74	0.63	0.73	0.59
OECD	2.30	2.35	0.69			0.44	0.37

Notes: Each cell reports the average moments across all regions within the panel. The first two columns present the average standard deviations. Column three reports the contemporaneous correlation between output and consumption. Columns four and five report the average cross-correlations of output and consumption, respectively. The final two columns report the average regional-reference correlations for output and consumption, respectively. The reference regions are: Ontario, Tokyo, Ohio, United States for the four panels, respectively.

Table 2. – Variance Decomposition: Simple Projection

Panel	Output		Consumption	
	Common	Region-Specific	Common	Region-Specific
Panel A: Projection on Aggregate Variable				
Canada	0.44	0.56	0.65	0.35
Japan	0.35	0.65	0.32	0.68
United States	0.65	0.35	0.45	0.55
OECD				
Panel B: Projection on Reference Variable				
Canada	0.37	0.63	0.59	0.41
Japan	0.12	0.88	0.24	0.76
United States	0.64	0.36	0.39	0.61
OECD	0.25	0.75	0.19	0.81

Notes: Estimation results for equation (8) averaged across regions within each panel. Common (idiosyncratic) refers to the explained (unexplained) sum of squares.

Table 3. – Risk Sharing Parameter Estimates

Panel	Mean Statistics			Number of Regions for Which a 95 % Confidence Interval Encompasses a Risk Sharing Parameter Equal to:			
	$\hat{\lambda}$	$se(\hat{\lambda})$	R^2	$\hat{\lambda} = 1$	$1 > \hat{\lambda} > 0$	$\hat{\lambda} = 0$	$\hat{\lambda} = 0 \text{ or } 1$
Canada	0.88	0.08	0.75	6	4	0	0
Japan	0.93	0.24	0.39	41	1	2	3
United States	0.81	0.32	0.50	29	6	4	12
OECD	0.40	0.19	0.50	3	10	7	1

Notes: Parameter estimates of equation (9), where the equation is estimated for each region in the panel as a time series and the average estimate of λ , standard error of λ (denotes ‘se’), and R-square are reported in the first three columns. The last four columns report information associated with the distribution of the estimated λ ’s.

Table 4. – Welfare Gains From Removing Region Specific Consumption Growth Variability ($\gamma = 4$)

Panel	Statistic	Region PI	Covariance	Unexplained	Total
Canada	Highest quintile	0.32	0.56	0.64	1.52
	Average	0.15	0.34	0.46	0.94
	Lowest quintile	0.12	0.27	0.12	0.50
Japan	Highest quintile	0.47	-0.36	2.00	2.12
	Average	0.15	-0.06	0.85	0.94
	Lowest quintile	0.04	-0.04	0.35	0.35
United States	Highest quintile	2.18	0.48	9.71	12.37
	Average	1.01	0.25	5.73	6.99
	Lowest quintile	0.68	-2.62	3.71	1.78
OECD	Highest quintile	1.55	0.00	1.50	3.06
	Average	0.74	0.13	0.65	1.52
	Lowest quintile	0.09	0.04	0.30	0.44

Notes: From regression (9), we decompose the region specific variability of consumption, $var(\Delta \log(c_{it})) - \hat{\lambda}_i^2 var(\Delta \log(c_t^a))$, into three components: $((1 - \hat{\lambda}_i)^2 var(\Delta \log(y_{it}))$, $2\hat{\lambda}_i(1 - \hat{\lambda}_i)^2 cov(\Delta \log(c_t^a), \Delta \log(y_{it}))$, and the variance of the residual, $\sigma_{u_i}^2$, respectively. The columns headed by ‘Regional PI’, ‘Covariance’ and ‘Unexplained’ are calculated as $-0.5 \cdot \gamma \cdot d\sigma_{\Delta c}^2 / (r - \bar{\mu})$ multiplied by these three terms. The final column, denoted ‘Total’, reports the sum of the three components. As we have λ_i for each separate region or country, we report the mean, highest and lowest quartile welfare changes.

Figure 1. - Cross-National and Cross-regional Income and Consumption Correlations




Figure 2. - Consumption-Income Correlations Within Regions and Countries

