Trade Costs, Asset Market Frictions and Risk Sharing*

Doireann Fitzgerald†

July 2010

Abstract

I use bilateral import data to test for the role of trade costs and asset market frictions in impeding international consumption risk sharing. Trade costs play a significant role. I do not reject the null of optimal risk sharing within OECD countries, though I do reject for the world as a whole. I calculate the impact on ex-post welfare of moving from historical asset market frictions to optimal consumption risk sharing, holding trade costs fixed. The gains to OECD countries are small while the gains to non-OECD countries are larger. These gains are dwarfed by those from eliminating trade costs.

JEL Classification: F41

1 Introduction

In a world where there are no frictions in goods markets, and a full set of contingent claims can be traded, the growth rate of the marginal utility of consumption should be perfectly correlated across countries. This prediction is rejected by the data (e.g. Backus, Kehoe and Kydland [1992]). Popular explanations for the failure of perfect consumption risk sharing include costs of trading goods internationally and deviations of international asset markets

*I am grateful for financial support from the University of California and the NSF under grant #0647850. Expert research assistance was provided by Daniel Beltran. I thank two anonymous referees, Manuel Amador, Jonathan Eaton, Pierre-Olivier Gourinchas, Sam Kortum, Kenneth Rogoff, Iván Werning and seminar participants at Harvard, Stanford, the Spring meetings of the NBER ITI program, the NBER Summer Institute and the Second Annual CEPR Workshop on Global Interdependence for comments and suggestions. The usual disclaimer applies.

†Department of Economics, Stanford University, Stanford, CA 94305, dfitzger at stanford.edu

from the Arrow-Debreu benchmark (e.g. Obstfeld and Rogoff [2000] and Heathcote and Perri [2002]). Trade costs make risk sharing costly, so it is optimal not to share consumption risk fully. Asset market frictions limit countries' ability to write and honor the contracts necessary to implement the optimal amount of consumption risk sharing conditional on trade costs. There is a substantial body of evidence that trade costs are large (Anderson and van Wincoop [2004]). There is also evidence that frictions such as limited commitment affect asset trade betwen developed and developing countries. But somewhat surprisingly, it appears that frictions in asset markets also play an important role in impeding optimal consumption risk sharing between OECD countries (see Backus and Smith [1993], Kollmann [1995] and Ravn [2001]).

In this paper, I test for and quantify the importance of trade costs and asset market frictions in explaining the failure of perfect consumption risk sharing. In contrast with the previous literature, I find that frictions in international asset markets significantly impede optimal consumption risk sharing between OECD and non-OECD countries, but not necessarily within OECD countries. Trade costs, in contrast, significantly impede risk sharing for all countries.

To motivate the tests and to provide a framework for the welfare calculations, I first present a theoretical framework that nests both trade costs and asset market frictions. In the way I introduce trade costs, I borrow heavily from the literature on the theoretical foundations of a gravity model of intra-temporal trade. Gravity models do a good job of matching the pattern of bilateral trade, and have become a workhorse for welfare analysis in the trade literature (see Arkolakis, Costinot and Rodriguez-Clare [2009]). Several different static models of specialization and trade yield the same gravity structure. For simplicity, I follow Anderson and van Wincoop [2003] in assuming Armington specialization. Together with CES demands and iceberg costs of trade, this assumption yields predictions about the matrix of bilateral trade which are observationally equivalent to the predictions of the Ricardian model of Eaton and Kortum [2002], or to those which would result from an increasing returns story as in Krugman [1980].

I nest the gravity model inside a standard DSGE model, where countries produce output using elastically supplied labor, accumulable capital and stochastic productivity. The structure of international asset markets determines the extent to which countries can engage in trade across states and over time. Irrespective of what goes on in asset markets, all types of trade - within and across states and periods - are limited by the presence of resource costs of trade. I do not take a stand what the asset market looks like. Nevertheless, exactly as in a world with no trade costs, a key implication of complete and frictionless asset markets is that a country's inverse marginal utility of wealth (relative to other countries) is constant over time. Just as in a world without trade costs, a natural metric for the salience of asset market frictions in impeding consumption risk sharing between groups of countries is the degree to which the relative inverse marginal utility of wealth moves around across states of the world and over time.

The model provides a natural setting for testing for the role of trade costs and asset market frictions in limiting consumption risk sharing both within and across states and over time. Empirical gravity models have long been used to infer the presence of trade costs from the trade-reducing effects of distance. The key contribution of this part of the paper is to show how bilateral trade data can also be used to test for the role of asset market frictions in impeding perfect consumption risk sharing. This relies on the appearance of the utility-consistent consumption price, which is equal to the ratio of the marginal utility of consumption and the marginal utility of wealth, in the model-based gravity equation. The test is then based on comparing the ability of a restricted gravity model that imposes a constant relative inverse marginal utility of wealth across countries to explain the variation in bilateral trade data with that of an unrestricted gravity model. In addition, I show how the same structure can be used to test the null hypothesis of financial autarky against the alternative of some international asset flows, by placing a different restriction on the empirical gravity model.

I implement the tests using bilateral trade data for a sample of 88 developed and developing countries from 1970-2000. Unsurprisingly, given what we know from the literature on estimated gravity models, the null hypothesis of no trade costs is overwhelmingly rejected for all countries. The null hypothesis of financial autarky is rejected, while the null hypothesis of optimal risk sharing between OECD and non-OECD countries conditional on trade costs is also rejected. However the results for OECD countries are strikingly different from the previous literature: the test does not reject the null hypothesis of optimal consumption risk sharing within OECD countries.

Next, I use the model to provide a more continuous metric of relative distance from optimal consumption risk sharing. More precisely, I reallocate the output historically devoted to consumption consistent with optimal risk sharing conditional on historical trade costs, and then calculate the implied changes in trade and ex-post welfare relative to the historical benchmark. I do this separately for optimal risk sharing within the OECD and optimal risk sharing in the world as a whole. I also examine the incremental effect of moving from optimal risk sharing in the world as a whole under historical trade costs to optimal risk sharing with zero trade costs.

These exercises require estimates of historical trade costs and output. The obvious way to obtain estimates of trade costs is by estimating a gravity model. To make everything internally consistent, I also use the gravity model to obtain estimates of the output prices which I use to deflate nominal output. More precisely, I use the gravity model to estimate both of these variables up to the order of an exponent which depends on the elasticity of substitution. In the estimation, I make use of the full set of restrictions implied by the structural model. In particular, I impose the restriction that the predicted ratio of net exports to GDP for each country exactly matches that in the data. The fitted model does a good job of matching the pattern of bilateral trade and the evolution of ratios of trade to GDP. Based on the estimates from the gravity model, and an assumption about the elasticity of substitution from the previous literature, I can then calculate both the real output historically devoted to private consumption and the level of historical trade costs. The additional assumption of a value for risk aversion allows me to estimate the relative inverse marginal utility of wealth for each country in each year, and to calculate trade and welfare under counterfactual time-invariant values for this variable. The estimates of the historical inverse marginal utility of wealth are used to choose "plausible" points on the Pareto frontier.

Under optimal consumption risk sharing within OECD countries alone, ex-post welfare in the median OECD country is almost unchanged from the benchmark. In contrast, under optimal risk sharing in the world as a whole, the increase in ex-post welfare in the median non-OECD country is on the order of 3%, while there is a reduction of 1% in welfare in the median OECD country. This suggests that over the sample period, OECD countries were much closer to optimal risk sharing (with each other) than with non-OECD countries. This is reflected in the behavior of trade and net exports in the counterfactual exercises compared with the fitted model. Under optimal risk sharing within OECD countries alone, total intra-OECD trade as a share of OECD GDP is 6% higher than in the fitted model, while the median absolute value of the trade balance as a share of GDP for OECD countries widens from 2% to 4%. In contrast, under optimal risk sharing in the world as a whole there are substantial increases in trade between OECD and non-OECD countries - this trade as a share of non-OECD GDP is 23% higher than in the fitted model. This reflects an increase in the median absolute value of the trade balance for non-OECD countries from 6% to 15%, and for OECD countries from 2% to 5%. Finally, the incremental effect of eliminating trade costs on welfare and trade is enormous, which is not surprising given the estimated size of trade costs.

This paper is related to several different literatures. The test for the presence of frictions in international asset markets is related to Lewis [1996] who tests for perfect consumption risk sharing in a large sample of countries in a framework that does not have trade costs. She does not reject the null of risk sharing for countries classified as having more unrestricted asset trade. It is also related to Backus and Smith [1993], Kollmann [1995] and Ravn [2001] among others, who test for optimal consumption risk sharing among OECD countries, conditional on frictions in goods markets. Relative to this latter literature, I innovate by giving a unified treatment to asset markets and trade costs, and by comparing OECD with non-OECD countries. The role of trade costs in explaining international macro puzzles has been previously explored by Backus, Kehoe and Kydland [1992], Dumas [1992], Obstfeld and Rogoff [2000], Dumas and Uppal [2001], Heathcote and Perri [2004], Kose and Yi [2006], Mazzenga and Ravn [2004] and Fitzgerald [2008] among others. I innovate relative to much of this literature by undertaking a quantitative analysis for a large number of countries based on a structural gravity model. Importantly, in contrast to much of this work, my analysis does not require me to take a stand on what exactly are the frictions in asset markets. The paper is also related to the substantial literature on the specification and estimation of structurally-based static gravity equations of bilateral trade, including Eaton and Kortum [2002], Anderson and van Wincoop [2003] and Alvarez and Lucas [2007]. More recently, a growing literature uses calibrated models of this type for welfare analysis (see Arkolakis, Costinot and Rodriguez-Clare [2009] for citations), generally under the assumption of balanced trade. An exception is Dekle, Eaton and Kortum [2008], who examine the effect on world trade and welfare of eliminating trade balances in the context of a calibrated gravity model. Of this latter literature, theirs is the most closely related paper. Finally, this paper is related to the literature on measuring the potential welfare gains from international risk sharing, as summarized by van Wincoop [1999].

The next section lays out the theoretical framework. The third section describes the test for the role of frictions. The fourth section presents the data and test results, and discusses their interpretation. The fifth section describes the welfare analysis and results. The final section concludes.

2 Theoretical framework

Summary

There are N countries in the world, indexed $i = 1, \ldots, N$. Each country produces a distinct intermediate good (also indexed i) using capital, labor and materials. Capital is accumulable. Labor is elastically supplied. Productivity in the production of each country's intermediate good is stochastic. No restrictions are placed on the joint process for productivity in all countries. The intermediate goods are tradeable at some cost which takes an iceberg form. They are combined using a CES aggregator, the same in all countries, to produce a non-tradeable final good used for private and public consumption, investment and materials. Asset markets are complete and frictionless within countries, but there may be exogenous or endogenous limitations on the contracts that can be written betweeen agents from different countries.

Uncertainty

In each period t, the world economy experiences one event, $s_t \in S$. Denote by s^t the history of events from date 0 to date t. The probability of history s^t at date t is given by $\pi_t\left(s^t\right)$.

Preferences and technology

For simplicity, the problem is described as if country i had a single agent with expected utility given by:

$$U^i = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t\left(s^t\right) L_t^i u\left(\frac{C_t^i\left(s^t\right)}{L_t^i}, \frac{H_t^i\left(s^t\right)}{L_t^i}\right) \tag{1}$$

where L_t^i is the (deterministic) population of country i, $C_t^i\left(s^t\right)$ is total consumption and $H_t^i\left(s^t\right)$ is the total number of hours worked.

Country i produces intermediate good i by combining capital, labor and materials using a constant returns to scale production function:

$$Y_t^i\left(s^t\right) = \left[F\left(K_t^i\left(s^{t-1}\right), A_t^i\left(s^t\right) H_t^i\left(s^t\right)\right)\right]^{\sigma} M_t^i\left(s^t\right)^{1-\sigma} \tag{2}$$

where $A_t^i\left(s^t\right)$ is the realization of productivity, $K_t^i\left(s^{t-1}\right)$ is the (predetermined) capital stock available for use in production in country i at time t and $M_t^i\left(s^t\right)$ is materials used up in production.

The production function for the final good, X is:

$$X_t^i\left(s^t\right) = \left[\sum_{k=1}^{N} Z_t^{ik}\left(s^t\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{3}$$

where $Z_t^{ik}\left(s^t\right)$ is absorption in country i of intermediate good k at s^t. If $\eta \to \infty$, we are in the case of a one-good world. The resource constraint for the (non-traded) final good in i is:

$$X_t^i\left(s^t\right) = C_t^i\left(s^t\right) + I_t^i\left(s^t\right) + M_t^i\left(s^t\right) \tag{4}$$

where $I_t^i\left(s^t\right)$ is investment in country i. Given the rate of depreciation δ, capital in country i accumulates according to:

$$K_{t+1}^i\left(s^t\right) = (1-\delta)\, K_t^i\left(s^{t-1}\right) + I_t^i\left(s^t\right)$$

Resource costs of trade

Intermediate goods trade may be costly: in order for one unit of j's good to arrive in i, $\tau_t^{ij}\left(s^t\right)$ units must be shipped, with $\tau_t^{ii}\left(s^t\right) = 1$, $\tau_t^{ij}\left(s^t\right) \geq 1$ and $\tau_t^{ij}\left(s^t\right)\tau_t^{jk}\left(s^t\right) \geq \tau_t^{ik}\left(s^t\right)$. It is not in general required that $\tau_t^{ij}\left(s^t\right) = \tau_t^{ji}\left(s^t\right)$. The intermediate goods resource constraints must take account of the resource cost of trade:

$$Y_t^i\left(s^t\right) = \sum_{k=1}^{N} \tau_t^{ki}\left(s^t\right) Z_t^{ki}\left(s^t\right) \tag{5}$$

Goods market

Producers of intermediate goods are assumed to be atomistic price takers. But due to trade costs, intermediate goods prices differ across countries:

$$Q_t^{ki}\left(i, s^t\right) = \tau_t^{ki}\left(s^t\right) Q_t^{ii}\left(s^t\right) \tag{6}$$

where $Q_t^{ii}\left(s^t\right)$ is the spot price of intermediate i in country i at s^t and $Q_t^{ki}\left(s^t\right)$ is its spot price in country k. In what follows, $Q_t^{ii}\left(s^t\right)$ is abbreviated to $Q_t^{i}\left(s^t\right)$, and $Q_t^{ki}\left(s^t\right)$ is replaced by $\tau_t^{ki}\left(s^t\right) Q_t^{i}\left(s^t\right)$.

Asset market

At s^t, country i enters with a vector of asset holdings $\mathbf{B}_t^i(s^{t-1})$ that pays dividends with value given by $\mathbf{D}_t(s^t) \cdot \mathbf{B}_t^i(s^{t-1})$. The vector of asset prices, taken as given by each country, is $\mathbf{R}_t(s^t)$. The value of asset holdings is given by $\mathbf{R}_t(s^t) \cdot \mathbf{B}_t^i(s^{t-1})$. The country can choose to re-optimize its holdings by spending $\mathbf{R}_t(s^t) \cdot \mathbf{B}_{t+1}^i(s^t)$ to purchase a new vector with $\mathbf{B}_{t+1}^i(s^t) \in \mathcal{B}_i(s^t, \mathbf{K}_t(s^{t-1}), \mathbf{B}_t(s^{t-1}))$. Assets are distinguished by their dividend vectors, while the asset market structure determines $\mathcal{B}_i(s^t, \mathbf{K}_t(s^{t-1}), \mathbf{B}_t(s^{t-1}))$, the set of assets available to country i at s^t. Assets are defined such that they are in zero net supply:

$$\sum_{i=1}^{N} \mathbf{B}_t^i\left(s^t\right) = 0 \tag{7}$$

This setup is general enough to encompass frictionless asset markets, financial autarky, and a variety of different types of frictions that allow for partial risk sharing through asset markets.[1]

Sequential competitive equilibrium

At each point in time, country i chooses $C_t^i(s^t)$, $H_t^i(s^t)$, $\mathbf{Z}_t^i(s^t)$, $M_t^i(s^t)$, $K_{t+1}^i(s^t)$ and $\mathbf{B}_{t+1}^i(s^t)$ to maximize expected utility (1) subject to its aggregate good resource constraint:

$$\left[\sum_{k=1}^{N} Z_t^{ik}\left(s^t\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} = C_t^i\left(s^t\right) + K_{t+1}^i\left(s^t\right) - (1-\delta) K_t^i\left(s^{t-1}\right) + M_t^i\left(s^t\right) \tag{8}$$

and its budget constraint, expressed in nominal terms:

$$\sum_{k=1}^{N} \tau_t^{ik}\left(s^t\right) Q_t^k\left(s^t\right) Z_t^{ik}\left(s^t\right) - Q_t\left(i, s^t\right) F\left(K_t^i\left(s^{t-1}\right), A_t^i\left(s^t\right) H_t^i\left(s^t\right)\right)^{\sigma} M_t^i\left(s^t\right)^{1-\sigma} =$$

$$\left[\mathbf{D}_t\left(s^t\right) + \mathbf{R}_t\left(s^t\right)\right] \cdot \mathbf{B}_t^i\left(s^{t-1}\right) - \mathbf{R}_t\left(s^t\right) \cdot \mathbf{B}_{t+1}^i\left(s^t\right) \tag{9}$$

where the left hand side is the difference betwen the value of expenditure and the value of output, and the right hand side is the difference between the value of wealth entering s^t, and the value of wealth sent forward to the future. $\mathbf{R}_t(s^t)$, $\mathbf{D}_t(s^t)$ and $\mathbf{Q}_t(s^t)$, are taken as given, and $\mathbf{B}_{t+1}^i(s^t) \in \mathcal{B}_i(s^t, \mathbf{K}_t(s^{t-1}), \mathbf{B}_t(s^{t-1}))$. A competitive equilibrium is a vector of prices $(\mathbf{Q}^*, \mathbf{R}^*)$ and a vector of quantities for each country $\left(\mathbf{C}^{i*}, \mathbf{H}^{i*}, \mathbf{Z}^{i*}, \mathbf{M}^{\mathbf{i}*}, \mathbf{K}^{i*}, \mathbf{B}^{i*}\right)$ such that each country solves the above problem at every realized state s^t and (5), (7) and (8) hold for all realized s^t (i.e. all markets clear).

[1] For example, it nests the case where the only internationally traded asset is a risk-free bond. It also nests the cases of limited commitment with endogenous and exogenous incompleteness.

First order conditions

Assume the existence of a competitive equilibrium. Then at any s^t, given $\mathbf{K}^*$ and $\mathbf{B}^*$, the first order conditions with respect to absorption of each intermediate and aggregate consumption are necessary for competitive equilibrium (i.e. the problem of country i at s^t is convex).[2]

The first order condition with respect to $Z_t^{ik}(s^t)$ is:

$$P_t^i\left(s^t\right) X_t^i\left(s^t\right)^{\frac{1}{\eta}} Z_t^{ik}\left(s^t\right)^{-\frac{1}{\eta}} = \tau_t^{ik} Q_t^k\left(s^t\right) \tag{10}$$

The first order condition with respect to $C_t^i(s^t)$ is:

$$\lambda_t^i\left(s^t\right) u_c\left(\frac{C_t^i\left(s^t\right)}{L_t^i}, \frac{H_t^i\left(s^t\right)}{L_t^i}\right) = P_t^i\left(s^t\right) \tag{11}$$

where $\lambda_t^i(s^t)$ is the inverse of the marginal utility of current per capita nominal wealth for country i evaluated at s^t (i.e. the inverse of the multiplier on the budget constraint), $u_c(\cdot)$ is the marginal utility of consumption per capita, and $P_t^i(s^t)$ is the nominal price of the consumption aggregate in i at s^t, i.e.:

$$P_t^i\left(s^t\right) = \left[\sum_{k=1}^{N}\left(\tau_t^{ik} Q_t^k\left(s^t\right)\right)^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{12}$$

As usual, the $\lambda_t^i(s^t)$'s have an alternative interpretation as the per capita Pareto weights for each country in a squence of intra-temporal planning problems. A useful normalization of the λ's, and consequently, prices, suggested by this interpretation is $\sum_i \lambda_t^i(s^t) = 1$.

Complete markets

In general, the relative marginal utility of nominal wealth, and hence its inverse, $\lambda_t^i(s^t)$, varies across states and over time in a way that depends jointly on the process for shocks, trade costs, and on the exact specification of the asset market. The evolution of $\lambda_t^i(s^t)$ over time is a convenient way of summarizing the combined impact of all of these factors on ex-post consumption risk sharing. The value of this summary statistic lies in the fact that it is straightforward to characterize some aspects of the behavior of $\lambda_t^i(s^t)$ under complete and frictionless asset markets.

When asset markets are complete and frictionless, the sequential competitive equilibrium of this economy is Pareto optimal. The competitive equilibrium allocation can therefore be recovered as the solution to a planning problem where for appropriate $\left\{\lambda^1, \ldots, \lambda^N\right\}$ the

[2] It is not necessary to make use of the other necessary conditions.

planner chooses sequences **C**, **H**, **Z**, **M**, and **K** to maximize the ex-ante weighted sum of expected utilities:

$$\sum_{i=1}^{N} \lambda^i U^i = \sum_{i=1}^{N} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t \left(s^t\right) \lambda^i L_t^i u \left(\frac{C_t^i \left(s^t\right)}{L_t^i}, \frac{H_t^i \left(s^t\right)}{L_t^i} \right) \tag{13}$$

subject to (for all i and s^t) the resource constraints (5) and (8). The first order conditions for this problem with respect to choices of **C** and **Z** are exactly (11) and (10), with constant λ^i.

3 A test for consumption risk sharing

It is natural to base tests for consumption risk sharing on the first order condition for aggregate consumption, (11). Given (11), the relative marginal utility of consumption is (dropping the state-contingent notation for the sake of brevity):

$$\frac{u_c \left(\frac{C_t^i}{L_t^i}, \frac{H_t^i}{L_t^i} \right)}{u_c \left(\frac{C_t^j}{L_t^j}, \frac{H_t^j}{L_t^j} \right)} = \left[\frac{\lambda_t^i}{\lambda_t^j} \right]^{-1} \left[\frac{P_t^i}{P_t^j} \right] \tag{14}$$

This expression illustrates the fact that the relative marginal utility of consumption can vary over time and across states for two reasons. First, if asset markets are not frictionless, λ_t^i/λ_t^j need not be constant. In particular, relative wealth may respond to current output. Of course, the converse is not true, as depending on the exact nature of frictions, the process for shocks and the elasticity of substitution η, λ_t^i/λ_t^j could be constant even if asset markets are not frictionless. Second, in order for consumption risk sharing to take place, goods must be shipped internationally. If shipping is costly, agents will optimally choose not to smooth consumption perfectly. The effect of trade costs is captured by the (consumption) real exchange rate, P_t^i/P_t^j, which in general will vary with current output when trade is costly.

One strand of the literature on testing for international consumption risk sharing (e.g. Lewis [1996]) examines whether relative traded goods consumption growth rates are correlated with relative output growth rates. One disadvantage of this approach is that if traded goods prices differ across countries, the test cannot distinguish between failures of risk sharing due to frictions in goods markets and failures of risk sharing due to frictions in asset markets. In contrast, the literature on consumption-real exchange rate correlations (e.g.

Backus and Smith [1993], Kollmann [1995] and Ravn [2001]) infers the presence of frictions in international asset markets from an examination of the relationship between real exchange rates (the inverse of relative prices) and the relative real consumption implied by those prices. This literature implicitly conditions on (but does not test for or estimate) frictions in goods markets. It finds that for OECD countries, the correlation between relative consumption and real exchange rates is zero, or has the opposite sign to that predicted under the null hypothesis of frictionless asset markets and concave utility. The conclusion drawn from this finding is that there must be frictions in asset markets impeding consumption risk sharing between OECD countries.

3.1 Tests based on bilateral import flows

I now describe a unified framework for testing for the role of asset market frictions and trade costs in impeding perfect consumption risk sharing. The tests are based on estimating restricted and unrestricted gravity models of bilateral imports. As such, the role of trade costs is identified principally from the trade-reducing effect of distance. The test for optimal risk sharing (conditional on trade costs) is, like the tests described above, based on the first order condition for consumption. But in contrast with the literature on consumption-real exchange rate correlations, which infers the presence of asset market frictions from the correlation between relative consumption growth and the change in real exchange rates, the role of frictions in asset markets is inferred from the degree to which restricting the marginal utility of wealth to be constant over time rather than allowing it to vary in an unrestricted way reduces the ability of the gravity model to fit the bilateral trade data. Additionally, the test uses the prices consistent with the gravity model of trade (rather than measured prices) to decompose the value of consumption into price and quantity terms. Finally, it is also possible to test the null hypothesis of financial autarky against the alternative of some asset trade within the context of this framework. Because this test is based on the ability of a (different) restricted gravity model to match the same bilateral trade data, it provides a way of measuring whether international asset markets are further from complete markets or from financial autarky.

The model described above yields an expression for bilateral imports that takes the gravity form. Combining (10), the first order condition for intermediate goods, with the resource constraints for intermediate goods we get:[3]

$$\frac{IM_t^{ik}}{EXP_t^i OUT_t^k} = \frac{\tau_t^{ik} Q_t^k Z_t^{ik}}{\left(P_t^i X_t^i\right)\left(Q_t^k Y_t^k\right)} = \left(\frac{P_t^i \Pi_t^k}{\tau_t^{ik}}\right)^{\eta-1} \tag{15}$$

where

$$\left(\Pi_t^k\right)^{1-\eta} = \sum_{j=1}^{N} \left(\frac{P_t^j}{\tau_t^{jk}}\right)^{\eta-1} P_t^j X_t^j \tag{16}$$

and IM_t^{ik} is the value of imports into country i from country k, EXP_t^i is the value of total expenditure by country i and OUT_t^k is the value of gross output of country k. This is a relationship between the value of imports into i from k, the expenditure of the importer, the output of the exporter, the iceberg trade cost between the two countries, and two terms that are known in the gravity literature as "multilateral resistance" terms.[4] This expression is valid for any asset market structure, and under the assumption that trade costs take the iceberg form.[5]

There are two points to be noted about the appearance of P_t^i in (15). First, since P_t^i is related to λ_t^i through (11), expresssion (15) can be used to test hypotheses about the behavior of λ as well as τ. Second, while gravity equations of the form (15) do a good job of fitting the data on bilateral imports (see Anderson and van Wincoop [2004] for extensive citations on the literature on estimating gravity models), the relative prices they imply behave somewhat differently from measured real exchange rates. This can be at least partially attributed to the fact that the model-consistent prices implicitly value variety in a way that measured price indexes do not. The model-based approach to valuing variety is the basis for a substantial literature that measures the gains from intra-temporal trade (see the citations in Arkolakis, Costinot and Rodriguez-Clare [2009]), and the treatment of gains from inter-temporal trade here is consistent with that literature.

To see the implications for (15) of assuming optimal risk sharing, substitute in the first order condition for aggregate consumption, (11), and impose $\lambda_t^i = \lambda^i$:

$$\frac{IM_t^{ik}}{EXP_t^i OUT_t^k} = \frac{\left(\lambda^i\right)^{\eta-1}\left(u_c\left(\frac{C_t^i}{L_t^i}, \frac{H_t^i}{L_t^i}\right)\right)^{\eta-1}\left(\Pi_t^k\right)^{\eta-1}}{\left(\tau_t^{ik}\right)^{\eta-1}} \tag{17}$$

[3]See the online Appendix for the full derivation.

[4]This terminology is due to Anderson and van Wincoop [2003].

[5]Helpman, Melitz and Rubinstein [2008] develop a generalization of this expression under fixed and per unit costs of trade. This complicates the gravity expression, but does not affect the methodological point made here.

Of course, as already mentioned, while the vector λ is constant under complete and frictionless asset markets, this is not an if-and-only-if relationship. A constant λ may be consistent with frictions in asset markets under certain circumstances.[6] For this reason, I say that (17) holds under the null hypothesis of optimal risk sharing, rather than the null hypothesis of complete and frictionless asset markets.

To see what is implied by the assumption of zero trade costs, note that the model predicts that in the absence of trade costs, the composition of the expenditure basket is identical across countries, with expenditure shares given by exporter shares in world output. This does not depend on the nature of frictions in the asset market, which dictate the size of the basket across periods and states. This restriction on (15) can be easily imposed by setting $\tau_t^{ik} = 1 \; \forall i, k$, $P_t^i = 1 \; \forall i$ (a convenient normalization) and noting that in this case, $\Pi_t^k = \Pi_t$.

$$\frac{IM_t^{ik}}{EXP_t^i OUT_t^k} = \Pi_t^{\eta-1} \tag{18}$$

Finally, as shown by Anderson and van Wincoop [2003],[7] under financial autarky and symmetric trade costs, (15) reduces to an expression where $\Pi_t^k = P_t^k$:

$$\frac{IM_t^{ik}}{EXP_t^i OUT_t^k} = \left(\frac{P_t^i P_t^k}{\tau_t^{ik}}\right)^{\eta-1} \tag{19}$$

Clearly the fact that countries have non-zero trade balances indicates that they are not in financial autarky. But the fit of this expression provides a metric of relative distance from the polar opposites of optimal risk sharing and financial autarky.

Since (15) nests (17), (18) and (19), I test the null hypotheses of (17), (18) and (19) in turn against the alternative of (15).

3.2 Implementation

Marginal utility of consumption

Some assumptions must be made about the form of the marginal utility of consumption in order to implement the test of the null of optimal risk sharing against the alternative. I

[6] For example, Cole and Obstfeld show how in an endowment economy, risk sharing can be achieved through terms-of-trade movements alone, if $\eta = 1$.

[7] The derivation is reproduced in the online Appendix.

assume that marginal utility can be written:

$$u_{ct}^{i}\left(\cdot,\cdot\right)=\left(\frac{C_{t}^{i}}{L_{t}^{i}}\right)^{-\rho}\left(\frac{H_{t}^{i}}{L_{t}^{i}}\right)^{\psi} \tag{20}$$

Given that a contribution of the test I propose is to estimate P_t^i rather than to use measured prices, to be internally consistent, I do not want to use any information on aggregate prices or real exchange rates in implementing it. The assumed form for marginal utility allows the first order condition for aggregate consumption under the null to be rewritten as a function of the *value* of consumption:

$$P_{t}^{i}=\left(\lambda_{t}^{i}\right)^{\frac{1}{1-\rho}}\left(\frac{VC_{t}^{i}}{L_{t}^{i}}\right)^{\frac{-\rho}{1-\rho}}\left(\frac{H_{t}^{i}}{L_{t}^{i}}\right)^{\frac{\psi}{1-\rho}} \tag{21}$$

where $VC_t^i = P_t^i C_t^i$.

Trade costs

The standard assumption in the empirical gravity literature on the form of bilateral trade costs is (see Anderson and van Wincoop [2004]):

$$\left(\tau_{t}^{ik}\right)^{1-\eta}=\prod_{n=1}^{J}\left(D_{n}^{ik}\right)^{\gamma_{n}},\ D_{n}^{ik}=1\text{ if }i=k,\ D_{n}^{ik}\geq 1\text{ otherwise} \tag{22}$$

Commonly used "gravity variables" (D_n^{ik}) include bilateral distance and indicator variables for common language, colonial heritage etc. These areconstructed in such a way as to impose $\tau_t^{ii} = 1$. In general, the number of gravity variables $J \ll N^2$, where N^2 is the number of bilateral pairs included in the regressions. Ideally, non-resource costs of trade due to policy barriers should also be controlled for (the form of the test would be unaffected). However constructing the required data is beyond the scope of this paper. Given that there is no time variation in the standard set of gravity variables, I allow for time variation in trade costs by estimating a different vector of coefficients γ_t on the gravity variables for each year. Given the set of gravity variables used, symmetry of trade costs is imposed by construction, i.e. $\tau_t^{ik} = \tau_t^{ki}$.

Estimating equations

Taking models (15), (17), (18) and (19), substituting in the expressions for P_t^i and trade costs where appropriate, and taking logs yields the following four estimating equations:

Estimating Equations

	Asset market	Trade costs	Estimating equation
(a)	General	Yes	$w_t^{ik} = \theta_t^i + \phi_t^k + \sum_{n=1}^J \gamma_{nt} d_n^{ik} + \varepsilon_t^{ik}$
(b)	Frictionless	Yes	$w_t^{ik} = \psi^i + \phi_t^k + \beta_c vc_t^i + \beta_h h_t^i + \sum_{n=1}^J \gamma_{nt} d_n^{ik} + \varepsilon_t^{ik}$
(c)	General	No	$w_t^{ik} = \phi_t + \varepsilon_t^{ik}$
(d)	Autarky	Yes	$w_t^{ik} = \theta_t^i + \theta_t^k + \sum_{n=1}^J \gamma_{nt} d_n^{ik} + \varepsilon_t^{ik}$

Here, $w_t^{ik} = \ln\left(IM_t^{ik}/EXP_t^i OUT_t^k\right)$, $vc_t^i = \ln\left(VC_t^i/L_t^i\right)$, $h_t^i = \ln\left(H_t^i/L_t^i\right)$ and $d_n^{ik} = \ln D_n^{ik}$. $\theta_t^i = (\eta - 1)\ln P_t^i$ is an importer-year fixed effect, while $\phi_t^k = (\eta - 1)\ln \Pi_t^k$ is an exporter-year fixed effect, $\psi^i = ((\eta - 1)/(\rho - 1))\ln \lambda^i$ is an importer fixed effect and $\phi_t = (\eta - 1)\ln(\Pi_t)$ is a time fixed effect in the case of zero trade costs. I do not impose restriction (16) on the relationship between θ_t^i and ϕ_t^k, because I do not have access to data on the full universe of countries and it is not necessary to impose the restriction in order to perform the tests.

Since (a) nests in turn (b), (c) and (d), an F-test can be used to test the null of (b), (c) or (d) against the alternative of (a).

4 Data and test results

All data is annual. Population, the current dollar value of GDP, the current dollar value of total imports and total exports, and the current dollar value of private consumption expenditures are taken from the World Bank's *World Development Indicators* (WDI). The employment rate is taken from the *Penn World Tables*, version 6.3 (PWT), a choice based on breadth of coverage rather than data quality. Bilateral merchandise imports in current dollars from 1970 to 2000 are taken from the *NBER-United Nations Trade Data* (NBER-UN) prepared by Feenstra and Lipsey.[8] Bilateral service trade flows are not available for the sample period. Additional data to fill in gaps for Taiwan and the former Western Germany are taken from the national statistical agencies of these countries. Details are available on request.

To be consistent with the model, I need data on gross output to calculate the dependent variable for the test. Gross output is not generally available, but for a given value for σ, a model-consistent estimate is given by

$$OUT_t^i = \frac{1}{1-\sigma} GDP_t^i$$

[8]Bilateral merchandise imports data are available for later years from the IMF's *Direction of Trade Statistics* (DOTS). There are some inconsistencies across the NBER-UN and DOTS trade data, as the latter have been less extensively cleaned, so I do not make use of DOTS in the baseline sample.

I assume a value of $\sigma = 0.5$, based on the fact that the average ratio of gross output to GDP in the *OECD Intersectoral Database* is approximately equal to 2. A country's absorption of its own output (IM_t^{ii}) is then calculated as gross output less total exports. Total expenditure including expenditure on materials (i.e. EXP_t^i) is calculated as gross output less total exports plus total imports. In the absence of data on bilateral service flows, I assume they follow the same pattern as bilateral merchandise flows. IM_t^{ik} is constructed by calculating bilateral merchandise imports as a share of the importer's total merchandise imports (using NBER-UN data), and multiplying this by total imports at the country level from the WDI.

For something over one third of the country pairs in the sample, bilateral imports are recorded as zero. The literature on estimating gravity equations has taken a variety of different approaches to dealing with this issue. As a baseline, I add one to bilateral imports (i.e. construct the dependent variable as $w_t^{ik} = \ln\left(\left(1 + IM_t^{ik}\right)/EXP_t^i OUT_t^k\right)$). In section 4.2, I describe tests of the robustness of the results to alternative approaches to this issue.

For the purpose of estimating the gravity equations, variables that are correlated with trade costs are required. The baseline set of variables consists of bilateral distance in kilometers from largest population center to largest population center, and indicator variables for contiguity, common language and former colonial history. These are taken from the dataset made available by CEPII. A dummy variable indicating common legal origin (British, French, German, Scandinavian or Socialist) is constructed based on the categorization provided by la Porta et al [1999]. The distance variable in the regression is calculated as $\ln\left(1 + dist^{ik}\right)$, where $dist^{ii} = 0$. The indicator variables are normalized such that they equal zero for the case when a country trades with itself. This imposes the normalization that trade costs within a country are zero.

The largest possible sample given the requirement that all of the above variables be available for all years 1970-2000 consists of 88 developed and developing countries. These are listed in the Appendix. Over the sample period, these countries cover between 90% and 94% of world GDP. Within-sample trade (i.e. trade that does not involve a partner not in the sample) accounts for between 72% and 83% of world trade, assuming that the bilateral pattern of service trade is the same as that for merchandise.

4.1 Results

Baseline bilateral tests

The results from estimating the four models described above using the full 88-country sample are reported in the online Appendix. The estimated coefficients on the gravity variables are fairly standard. They differ somewhat between the optimal risk sharing specification and the unrestricted and autarky specifications, but not markedly so, and the implied fitted values of trade costs are very similar. The R^2s are 0.55, 0.52, 0.04 and 0.52 for the unrestricted, optimal risk sharing, zero trade cost and autarky models respectively. Table 1 reports the F-test statistics and p-values for the three hypothesis tests. The null hypothesis of optimal consumption risk sharing conditional on trade costs is rejected at all significance levels in favor of the alternative of some friction in asset markets that impedes risk sharing. The null hypothesis of no trade costs is rejected at all significance levels in favor of the alternative of trade costs. The null hypothesis of financial autarky is rejected at all significance levels in favor of the alternative of some international asset trade or transfers.

Do developed and developing countries face different frictions?

To test the hypothesis that developed countries are closer to optimal risk sharing with each other than are developed and developing countries, I repeat the exercise just described, using only observations on bilateral imports between 22 OECD countries.[9] This amounts to testing whether there is optimal consumption risk sharing between these 22 countries, conditional on estimated trade costs. At the same time, I also implement the other two tests on the OECD subsample. The results from estimating the four models are reported in the online Appendix. The magnitude of the coefficients on the gravity variables in the models with trade costs are different from those in the full sample, and the implied trade costs are substantially smaller. The R^2s of the three models with trade costs are higher when estimated on the OECD subsample than when estimated on the full sample - they range from 0.81 to 0.83 instead of from 0.52 to 0.55.

Table 1 reports the F-test statistics and associated p-values for the OECD sub-sample. The null hypothesis of optimal consumption risk sharing conditional on trade costs *cannot* be rejected for this sample. The null hypothesis of no trade costs is rejected at all levels of significance, despite the fact that the implied trade costs are smaller than in the full sample. The null hypothesis of financial autarky is rejected at all levels of significance. Implementation of the same tests for randomly selected samples of 22 countries indicates that the failure to reject the optimal risk sharing null for developed countries but not the world as a whole is not driven by the difference in sample size. Results from this exercise are reported in the online Appendix.

Are frictions declining in importance over time?

[9]These are the 22 founding OECD economies, less Turkey. The full list is in the online Appendix.

I also estimate the four models separately on the period 1970-1984 and the period 1985-2000. Table 1 reports the F-test statistics and associated p-values for the two sub-samples. In both the earlier period and the later period, the null hypothesis of optimal risk sharing and costly trade is rejected against the alternative of frictions in both goods and asset markets, though the F-test statistic is lower in the later period. Similarly, the null hypothesis of no trade costs is strongly rejected in both periods, and the null hypothesis of financial autarky is rejected in favor of the alternative of some asset trade or transfers. The fact that the F-test statistic falls over time is suggestive of some weakening in asset market frictions.

4.2 Robustness

I examine the robustness of these results along several dimensions. Detailed results are provided in the online Appendix.

I use four alternative approaches to dealing with zeros in the dependent variable. First, I repeat the same tests with an alternative dependent variable, constructed as:

$$w_t^{ik} = \ln\left(\left(\min_j\left(IM_t^{ij}\right) + IM_t^{ik}\right)/EXP_t^i OUT_t^k\right) \tag{23}$$

Second, I estimate using the baseline sample, but dropping observations where there are zeros in the dependent variable. Third, I estimate using data aggregated over five-year intervals to reduce the number of zeros. Finally, I implement a version of the strategy suggested by Helpman, Melitz and Rubinstein [2008], which involves a correction for selection based on a first-stage probit or logit. In all cases, the results are qualitatively unchanged.

Next, I test the robustness of the results to dropping the employment variable. The results are not affected. I also test robustness to using different sets of gravity variables. In particular, since the estimated coefficient on distance differs between the full sample and the OECD sub-sample, I allow for a full set of main effects and interactions with distance for two dummy variables, one indicating that only one country of the bilateral pair is in the OECD, and the other that neither country in the pair is in the OECD. The results are qualitatively unchanged.

I also check the robustness of the results to including Turkey (the excluded founder member of the OECD) in the OECD sample. The p-value rises, and the null of optimal risk sharing in the group including Turkey is (just) rejected at the 5% level, though not at the 10% level. The other results are unchanged. However if in addition Korea is included in the "OECD" group, the null of optimal risk-sharing within the OECD is overwhelmingly rejected in favor of the alternative. I also check whether splitting the sample by de jure financial openness yields similar results to the baseline split. The measure of financial openness I use is that constructed by Chinn and Ito [2008]. I rank countries within years by this measure, and pick groups of countries that are on average most financially open. There is considerable overlap between this sample and the OECD sample. For comparable sample sizes, the null hypothesis of optimal risk sharing is not rejected for the "financially open" sample, though I also cannot reject the null hypothesis of financial autarky for this sample.

4.3 Discussion and interpretation

The most surprising aspect of the test results is that they are considerably more favorable to the null hypothesis of optimal consumption risk sharing between OECD countries than the tests previously used in the literature, so it is worth explaining how and why they are different. There are two key differences between the test I implement and those used in the previous literature. I discuss each in turn.

First, the test is implicitly based on a decomposition of the value of consumption into price and real quantity terms that relies not on measured prices, but on the prices consistent with fitting a structural gravity model of bilateral trade to the data. To understand how these non-standard prices contribute to a different view of consumption risk sharing, I do the following. First, I use standard data on CPIs and nominal exchange rates (taken from *World Development Indicators*) to decompose the value of consumption into price and real quantity terms, and use these to calculate the correlation between changes in log relative real consumption and log relative prices (with the US as numeraire) similar to the past literature. I then perform the decomposition using instead the estimated prices (this requires an assumption about η, since the gravity equations identify $(\eta - 1) \ln P_t^i$, not P_t^i), and calculate the same correlations. I repeat the exercise for a range of values of η. Using measured prices, the correlation between the two variables for the median OECD country is positive and close to zero, similar to what is found by Backus and Smith [1993].[10] For the full 88-country sample, the correlation for the median country is negative and close to zero. This contrasts with the fact that using the estimated prices, the correlation is systematically and strongly negative, both for OECD countries and for the world as a whole. This goes at least some way towards explaining why the results I present above for the OECD are different from the standard literature. Full results, and results from regressing changes in log relative real consumption on log relative prices are reported in the online Appendix.

[10]They work with real exchange rates, the inverse of relative prices.

Of course, while a negative correlation between relative prices and relative real consumption is certainly less inconsistent with optimal risk sharing than a zero or positive correlation, it is by no means conclusive evidence of optimal risk sharing. Even ignoring other potential determinants of marginal utility besides consumption, the sign of the reduced form correlation between relative prices and relative consumption need not be reversed by frictions in asset markets. Indeed, the test I present rejects the null hypothesis of no asset market frictions for the full sample of countries, *despite* the negative correlation between relative prices and relative consumption implied by the prices used in the test.

This brings me to the second key difference between the test I present and the approach of the literature on consumption-real exchange rate correlations. The test I present infers the salience of asset market frictions from the relative ability of constant versus time-varying relative marginal utilities of wealth to explain the variation in bilateral import data. In contrast, the literature on consumption-real exchange rate correlations infers the presence of frictions from the fact that the joint behavior of relative prices and relative real consumption deviates from that predicted by the null hypothesis of optimal risk sharing.

It is possible to nest a version of this approach in mine, as follows. Under the null hypothesis of optimal risk sharing, following equation (21) the coefficient on the log of the value of consumption is equal to $\frac{-\rho}{1-\rho}(\eta-1)$. The estimated value of this coefficient for the OECD sample is negative (equal to -0.31 with standard error 0.6). This is only possible if (a) $\rho < 1$, i.e. the representative agent is less risk averse than log utility, or (b), $\eta < 1$, implying that trade costs are lower for countries that are far away than for countries that are close. Clearly (b) does not make sense, and (a) violates the priors of many macroeconomists. Moreover, reasonable values of η imply values of ρ that, though positive, are very close to zero. It is straightforward to implement a test of the joint null of optimal risk sharing *and* a particular value for $\frac{-\rho}{1-\rho}(\eta-1)$ by subtracting $\frac{-\rho}{1-\rho}(\eta-1)vc_t^i$ from the dependent variable in the optimal risk sharing model. In results reported in the online Appendix, I do this for a variety of values of $\frac{-\rho}{1-\rho}(\eta-1)$. I reject the joint null hypothesis for a range of pairs $\{\eta,\rho\}$ considered standard by the literature. In this sense, my findings are consistent with those of the consumption-real-exchange-rate literature, even though I use different prices.

Obviously, this raises questions about how to interpret the test results presented above. If a more stringent test rejects the null of optimal risk sharing for both OECD countries and the world as a whole, can the baseline results be interpreted as implying OECD countries are at least closer to optimal risk sharing than the rest of the world? A disadvantage of the testing framework is that it is not clear if this is the case. In the next section, I address this issue by explicitly measuring the changes in ex-post welfare from moving to optimal risk sharing, holding trade costs fixed at their historical levels.

Before going to the full counterfactual exercise, it is interesting to perform the following back-of-the-envelope calculation. Assume that the marginal utility of per capita consumption takes the form $(C_t^i/L_t^i)^{-\rho}$, for a particular value of ρ. Given data on real consumption per capita and the consumer price index, the first order condition for aggregate consumption can then be used to back out time-series of λ_t^i/λ_t^j for pairs of countries i and j. I assume $\rho = 2$ and construct $\lambda_t^i/\lambda_t^{US}$ for all the countries in the data set. I then calculate for each country-pair the coefficient of variation of log changes in $\lambda_t^i/\lambda_t^{US}$ over the sample period. The distribution of this measure of volatility for non-OECD countries has significant mass well to the right of the distribution for OECD countries - the median for OECD countries is 0.22, while the median for non-OECD countries is 0.42 (details are reported in the online Appendix). This is consistent with OECD countries being on average closer to optimal risk sharing with the US (the biggest OECD country) than most non-OECD countries are with the US. Obviously this is just suggestive evidence. But with this in mind, I move on to the full counterfactual exercise.

5 Trade and welfare under counterfactual risk sharing

In this section, I calculate the changes in trade and ex-post welfare for OECD and non-OECD countries from moving from historical frictions in asset markets to optimal consumption risk sharing, holding trade costs fixed at their historical levels. I also examine the incremental effect of moving from optimal risk sharing with historical trade costs to optimal risk sharing with zero trade costs. The exercises focuses on consumption risk sharing rather than the full impact of moving to frictionless asset markets.[11] Moreover, I focus on changes in ex-post welfare rather than changes in ex-ante welfare. I restrict the exercises in this way because I do not want to have to take a stand on country wealth in 1970, the historical structure of asset markets, or the underlying process for shocks. My strategy makes use of the fact that estimates of the marginal utility of wealth provide a convenient summary statistic for the historical impact of asset market frictions on consumption risk sharing, without having to characterize those frictions. The exercises have the disadvantage that they demand that the model be taken more literally than in previous section. But relative to testing the null

[11] Gourinchas and Jeanne [2006] address precisely the question of the welfare effects of financial integration in a world where all the potential gains come through the effect on investment, and find that the effects are on the order of a 1% permanent increase in domestic consumption for the typical non-OECD country.

hypothesis of one model against the alternative of another, they have the distinct advantage that they more directly address the question of how far different groups of countries are from optimal risk sharing.

I proceed in two stages. In the first stage I construct estimates of historical trade costs and output prices, and use these to construct estimates of the output historically devoted to consumption. To do this, I first choose functions of trade costs and output prices to match data on bilateral imports, while imposing that the predicted ratio of net exports to GDP for each country exactly matches that in the data. This makes use of the information on the extent of asset trade contained in net exports in a way that the reduced form estimation strategy in the previous section does not. The estimated model does a good job of matching the pattern of bilateral imports and the evolution over time of ratios of trade to GDP. The estimation does not tie down values for the elasticity of substitution η, or for the parameters of the period utility function, $u(\cdot)$. I choose values for these based on the previous literature, and combine them with the estimated functions of trade costs and output prices to calculate trade costs, real output, real consumption and ex-post welfare for both OECD and non-OECD countries. In the process, I also obtain estimates of the historical inverse marginal utility of wealth.

In the second stage, I reallocate the output devoted to real consumption across countries, first, consistent with optimal consumption risk sharing within OECD countries alone, and second, consistent with optimal consumption risk sharing across the world as a whole, while holding trade costs fixed at their calibrated historical levels. The estimates of the historical inverse marginal utility of wealth are used to choose "plausible" points on the Pareto frontier. Finally, I reallocate consumption consistent with optimal risk sharing and zero trade costs. I then calculate a measure of compensating variation for each country under each counterfactual. I also compare trade and net exports as a share of GDP under the fitted model and under the counterfactuals. I now describe the procedure in more detail, followed by the results.

5.1 Estimation strategy

In order to perform an internally consistent counterfactual exercise, it is necessary to assume that one has data on the universe of countries. Having made this assumption, it makes sense to exploit the full structure of the model in estimating trade costs and prices, something I do not do in the previous section. For technical reasons, it is more convenient to estimate output prices rather than consumption prices.[12] The estimating equation is based on:

$$IM_t^{ik} = \frac{\left(\tau_t^{ik}\right)^{1-\eta}\left(Q_t^k\right)^{1-\eta}}{\sum_{j=1}^N \left(\tau_t^{ij}\right)^{1-\eta}\left(Q_t^j\right)^{1-\eta}} EXP_t^i \tag{24}$$

which is obtained by rearranging the first order condition for Z_t^{ik}. I divide across by total expenditure, take logs and include an error term:[13]

$$\ln\left(\frac{IM_t^{ik}}{EXP_t^i}\right) = d^{ik}\gamma_t + \alpha_t^k - \ln\left(\sum_{j=1}^N \exp\left(d^{ij}\gamma_t + \alpha_t^j\right)\right) + \varepsilon_t^{ik} \tag{25}$$

where $\alpha_t^i = (1-\eta)\ln\left(Q_t^i\right)$, $d^{ik}\gamma_t = (1-\eta)\ln\left(\tau_t^{ik}\right)$, and d^{ik} is a vector of (time-invariant) gravity variables, constructed such that $\exp\left(d^{ii}\gamma_t\right) = 1$. This imposes that the estimated Q's and τ's are positive, and that $\tau_t^{ii} = 1$. I then choose the vectors γ_t and α_t to minimize the weighted sum of squared errors:

$$\sum_{i=1}^N \sum_{j=1}^N \omega_t^{ij}\left(\varepsilon_t^{ij}\right)^2 \tag{26}$$

subject to the restriction that the fitted value of the ratio of net exports to GDP for each country is exactly equal to the ratio in the data, i.e.:

$$OUT_t^i = \sum_{j=1}^N \frac{\exp\left(d^{ji}\gamma_t + \alpha_t^i\right)}{\sum_{k=1}^N \exp\left(d^{jk}\gamma_t + \alpha_t^k\right)} EXP_t^j \tag{27}$$

I set the weights ω_t^{ij} equal to $EXP_t^i OUT_t^j$ (the predicted size of IM_t^{ik} in a zero trade cost world). Upweighting large flows relative to small flows ensures that the predicted values of trade as a share of GDP match world aggregates. The procedure is repeated for each period independently.

The data is exactly the data used in the section on testing, reformulated in terms of in-sample shares. The exact transformation is described in the online Appendix. I use a parsimonious vector d^{ij} of gravity variables that includes $\ln\left(1+dist^{ik}\right)$ and main effects and interactions of $\ln\left(1+dist^{ik}\right)$ with indicators for one of the trading partners being in the OECD and the other outside the OECD, and for both trading partners being outside the OECD.[14] Allowing the trade-reducing effect of distance to differ depending on the nature

[12] This guarantees that both output prices and consumption prices are positive.

[13] In the baseline estimation I drop observations where $IM_t^{ik} = 0$. The results are robust to this.

[14] These are set equal to zero when a country trades with itself

of the bilateral pair is necessary to simultaneously match trade-GDP ratios for OECD and non-OECD countries. The estimated parameters on the gravity variables are reported in the online Appendix.

5.2 Model fit

The first panel of Figure 1 is a scatter plot of the fitted values of bilateral imports against the data, for all bilateral pairs and all years in the sample (bilateral imports are expressed as a share of total in-sample world expenditure so that all years can be presented in the same figure). The second panel is a scatter plot of the log of the fitted values against the log of the actual values. These figures illustrate the fact that particularly for large flows (which are important in matching world aggregates) the fit is excellent. The relative fit of large and small flows is controlled by the weighting of the error term in the bilateral trade equations. The cross-sectional correlations between fitted and actual values are systematically high for all types of bilateral pair, except for trade between two (different) non-OECD countries. This is illustrated in the online Appendix. The relatively poor fit for trade between non-OECD countries is consistent with the view that measurement error is a particular problem for flows between developing countries.

I also examine the fit of the model in terms of matching the evolution of trade-GDP ratios. This depends heavily on how well the model matches the share of a country's imports from itself in total expenditure. I focus on five different ratios. First is the ratio of world imports to world GDP. I then look at the ratio of within-OECD imports to OECD GDP, the ratio of trade between OECD and non-OECD countries to OECD GDP and to non-OECD GDP respectively, and finally, the ratio of within-non-OECD imports to non-OECD GDP. The four panels of Figure 2 plot the fitted and actual values of the latter four ratios respectively, with the fitted and actual values for the world ratios also reported in each panel. The model slightly overpredicts within-OECD trade and slightly underpredicts within-non-OECD trade. However the order of magnitude is corrrect, and in terms of time-series evolution, the fitted values closely follow actual values.

By construction, the fitted model exactly matches the ratio of net exports to GDP from the data, so I do not report statistics on fit along this dimension.

5.3 Parameters, trade costs and Pareto weights

To back out an estimate of real consumption, and of the Pareto weights, I require a value for η, the elasticity of substitution, an assumption about the functional form for marginal utility, and values for the parameters of this function. Anderson and van Wincoop [2004] suggest a range of values between 5 and 8 for η, while Eaton and Kortum [2002] suggest a value for the corresponding parameter in their model between 3 and 12. The macro literature has traditionally tended to choose values on the lower end of the spectrum. I choose a benchmark value of 6, and check the robustness of the results to this choice. I assume that marginal utility of per capita consumption takes the form $u_c(\cdot) = (C/L)^{-\rho}$. For ρ, I choose a benchmark value of 2, and again check the robustness of the results to this choice. Then, using the estimates $\hat{\gamma}_t$ and $\hat{\alpha}_t$, and data on population and the share of expenditure devoted to consumption, I recover fitted values of $\hat{\tau}_t^{ij}$, $\hat{Q}_t^i$, $\hat{P}_t^i$, $\hat{C}_t^i$, $\hat{\lambda}_t^i$, $\hat{Y}_t^i$, and $\hat{Y}_t^{iC}$, the real output devoted to consumption purposes. The details of these calculations are described in the online Appendix. The relationship between output in one period and the next (and the time-series behavior of all other variables) is pinned down by the evolution of real GDP in the US, taken from the WDI. This choice affects the welfare calculations, but not predictions on counterfactual trade-to-GDP ratios or net export-to-GDP ratios.

Trade costs

Figure 3 plots median trade costs for trade within OECD countries, trade between OECD and non-OECD countries and trade within non-OECD countries under the assumption that $\eta = 6$. Trade costs are expressed as a percentage of the sales price, so $\tau = 3$ implies a trade cost of 200%. These trade costs are large.[15] However, as Figure 4 illustrates, trade costs on this order of magnitude do a good job of matching trade-GDP ratios.[16] Moreover, while observed trade costs for goods and services that are actually traded are not this high, observed trade costs are not representative of those that apply to output as a whole, as in equilibrium, only the goods and services with the lowest trade costs are actually traded.

Pareto weights

It is also interesting to examine the behavior of the estimated Pareto weights (normalized such that $\sum_{i=1}^{N} \hat{\lambda}_t^i = 1$). For each country, I calculate the coefficient of variation based on the time series of that country's estimated Pareto weight from 1970 to 2000. Figure 4 plots the histogram of the distribution of these coefficients of variation for OECD and non-OECD

[15] Anderson and van Wincoop [2004] summarize the evidence on the size of trade costs.

[16] This statement is of course conditional on a value for η, as the fit of the model depends only on $\hat{\gamma}_t$.

countries separately. The median for OECD countries is 0.26, while the median for non-OECD countries is 0.67. The distribution for non-OECD countries is shifted to the right and is much more spread out than for OECD countries. This greater volatility is consistent with non-OECD countries being further from optimal risk sharing than OECD countries.

Figure 5 illustrates the time-series behavior of the estimated Pareto weights for a selected set of countries. This figure demonstrates that there is substantial co-movement within OECD countries, and to some extent, within groups of non-OECD countries. However, the volatility of the estimated weights is much more substantial for non-OECD countries. Moreover, their evolution reflects episodes such as the 1980s debt crisis and the East Asian financial crisis.

5.4 Counterfactual exercises

The first two counterfactual exercises are designed to measure the distance from optimal risk sharing for both OECD countries and non-OECD countries. This is not exactly the same as measuring the distance from frictionless asset markets. To calculate the decentralized equilibrium that would have obtained if various groups of countries had had access to frictionless asset markets from 1970, it is not sufficient to hold relative Pareto weights constant at the estimated levels for 1970. I would instead have to characterize each country's lifetime budget constraint under frictionless asset markets in 1970. This is beyond the scope of the paper, so I pick a point on the Pareto frontier that seems to be a reasonable benchmark. The baseline point is chosen by averaging the estimated relative Pareto weights over time within the risk-sharing group, I presently examine the robustness of the results to this choice.

First, I impose optimal risk sharing within the OECD only, while keeping trade costs fixed at their estimated historical levels. I do this by imposing that the weights of OECD countries relative to each other are constant over time, equal to the time-series average of their estimated relative weights. The weight of the OECD as a whole relative to that of the rest of the world is held fixed at its estimated level (remember $\sum_{i=1}^{N} \hat{\lambda}_t^i = 1$):

$$\tilde{\lambda}_t^{i,OECD}(1) = \frac{\sum_{t=1970}^{2000} \hat{\lambda}_t^i}{\sum_{j \in OECD} \sum_{t=1970}^{2000} \hat{\lambda}_t^j} \left[\sum_{j \in OECD} \hat{\lambda}_t^j \right] \tag{28}$$

The time-series of Pareto weights for non-OECD countries are also fixed at their estimated levels, i.e. $\tilde{\lambda}_t^{i,NOECD}(1) = \hat{\lambda}_t^{i,NOECD}$.

Second, I impose optimal risk sharing for the full sample of countries. I do this by imposing that the weight of each country relative to each other is held constant over time, with each country's weight calculated as the time-series average of its estimated weight:

$$\tilde{\lambda}^i(2) = \frac{\sum_{t=1970}^{2000} \hat{\lambda}_t^i}{\sum_{j=1}^{N} \sum_{t=1970}^{2000} \hat{\lambda}_t^j} \tag{29}$$

The third counterfactual exercise eliminates trade costs in addition to imposing optimal risk sharing (i.e. $\tilde{\lambda}(2)$) for the full sample of countries. The elimination of trade costs is achieved by setting $\tilde{\tau}_t^{ij} = 1$ for all bilateral pairs in all periods.

Given data on population, L_t, and $\left\{\tilde{\tau}_t, \hat{Y}_t^C, \tilde{\lambda}_t; \eta, \rho\right\}$, I calculate the consumption allocation that solves the problem of a social planner who maximizes the weighted sum of country utilities each period, with weights given by $\tilde{\lambda}(1)$ or $\tilde{\lambda}_t(2)$ where appropriate, subject to the aggregate good and intermediate good resource constraints (which depend on trade costs). The numerical algorithm for finding the solution is based on tatonnement as in Alvarez and Lucas [2007] (See the online Appendix for details).

Since I do not characterize either the baseline asset market, or the process for shocks, I cannot calculate a measure of compensating variation based on ex-ante welfare. Instead, I calculate compensating variation based on averaging ex-post welfare across all sample years:

$$W^i = \frac{1}{T} \sum_{t=1}^{T} \frac{\left(C_t^i / N_t^i\right)^{1-\rho}}{1-\rho} \tag{30}$$

The compensating variation, δ, is then such that:

$$\frac{1}{T} \sum_{t=1}^{T} \frac{\left(\delta \hat{C}_t^i / L_t^i\right)^{1-\rho}}{1-\rho} = \frac{1}{T} \sum_{t=1}^{T} \frac{\left(\tilde{C}_t^i / L_t^i\right)^{1-\rho}}{1-\rho} \tag{31}$$

This measure of welfare weights all periods equally, consistent with how the point on the within-group Pareto frontier is chosen. In the online Appendix, I describe an alternative baseline point on the Pareto frontier and corresponding welfare measure that apply discounting from the perspective of 1970, and which yield results that are very similar to the baseline.

5.5 Results

Table 2 reports summary statistics on the distribution of the compensating variation measure δ for the world as a whole, for the OECD, and for non-OECD countries, for all three counterfactual exercises.

Optimal risk sharing

Given the point I choose on the within-OECD Pareto frontier, optimal risk sharing within the OECD alone does not increase ex-post welfare for the median OECD country.[17] It leads to a reduction in the welfare of the median non-OECD country on the order of 2%. In contrast, optimal risk sharing for the world as a whole leads to a reduction in the welfare of the median OECD country on the order of 1%, while it leads to an increase in welfare for the median non-OECD country of 3%. In the online Appendix I report statistics on the distribution of compensating variation weighted by average population share over the period. The implications are qualitatively similar. These results suggest that the potential gains from relaxing frictions in asset markets between OECD countries may be relatively small. On the other hand, the gains from relaxing frictions in asset markets between OECD and non-OECD countries could be considerably larger. Of course, this is not the same as saying there are no frictions in asset markets between OECD countries; it could be that shocks are much more correlated within OECD countries than across OECD and non-OECD countries. Given the relative size of OECD and non-OECD countries, this would imply small gains to the OECD from greater financial integration, but large gains to non-OECD countries from greater integration with the OECD.

In addition, I examine the effects of optimal risk sharing within these two sets of country groups on trade-to-GDP ratios and net export-GDP ratios. Here, I calculate the value of both GDP and imports by adding consumption-related flows to the real output and real trade flows necessary to maintain real absorption devoted to investment, government consumption and materials at their estimated levels (these latter are not changed from estimated model to counterfactual). These total real flows are valued at the estimated and counterfactual prices as appropriate.

Figure 6 plots the following fitted and counterfactual trade-to-GDP ratios for the optimal risk sharing exercises: OECD trade as a share of OECD GDP, trade between OECD and non-OECD countries as a share of OECD GDP, trade between OECD and non-OECD countries as a share of non-OECD GDP, and trade between non-OECD countries as a share of non-OECD GDP. Optimal risk sharing within OECD countries alone leads on average to a small increase in within-OECD trade and trade between OECD and non-OECD countries. Optimal risk sharing for the world as a whole, on the other hand, leads to similarly small

[17]This is as expected. Even if the point chosen on the Pareto frontier is a Pareto improvement from the perspective of (unobserved) historical ex-ante welfare, there is no guarantee that optimal consumption risk sharing within a group of countries will increase *ex-post* welfare for all countries.

changes in within-OECD trade, but quite substantial increases in trade between OECD and non-OECD countries and within non-OECD countries. Table 3 provides summary statistics on time-series averages of the cross-section distribution of the absolute value of net exports over GDP. Optimal risk sharing within the OECD alone is accompanied by a doubling of the median trade balance for OECD countries from 2% to 4% of GDP in absolute value. Median trade balances for non-OECD countries fall from 6% to 5% of GDP. In contrast, optimal risk sharing in the world as a whole is accompanied by an increase in median trade balances for OECD countries from 2% to 5%, and for non-OECD countries from 6% to 15%. This is consistent with a large portion of consumption smoothing for non-OECD countries being provided by OECD countries.[18]

An additional interesting point to note is that *optimal* consumption risk sharing does not imply *perfect* consumption risk sharing. While the standard deviation of differences between country-level consumption growth rates and the (unweighted) cross-country average is dramatically lower under optimal risk sharing than in the baseline model (2.8 percentage points compared with 14.1 percentage points), non-trivial cross-country dispersion in consumption growth rates persists. More details are provided in the online Appendix.

Relationship to past literature

The literature on the welfare gains to international risk sharing has typically focused only on OECD countries, and has found a very wide range of potential gains, including some that are very large.[19] In the online Appendix, I document that the estimated baseline prices imply considerably more volatile consumption than is implied by data on prices and nominal exchange rates. This is true both for OECD and non-OECD countries. It might appear that the scope for gains from consumption risk sharing in this context is very large relative to what one would calculate based on using real consumption series derived from data on prices. However unlike the previous literature on international risk sharing, I focus on optimal consumption risk sharing conditional on trade costs. Since I estimate that trade costs are large, this substantially limits the optimal degree of risk sharing relative to a world without trade costs. The difference in welfare gains between OECD and non-OECD countries arises because OECD countries are large, so their fluctuations are strongly correlated with aggregate risk, and there is little scope for additional consumption smoothing, while this is not the case for non-OECD countries.

Relative to the literature on the welfare costs of business cycles, the estimated gains

[18]The results on the ratios of trade and net exports to GDP are independent of the values chosen for η and ρ.

[19]See van Wincoop [1999].

to non-OECD countries from optimal risk sharing are big. The choice of constant Pareto weights means that the exercise I perform does not just smooth out short-run fluctuations in consumption growth, but also medium-run deviations of growth from long-run trends, or differences in growth trends over different parts of the 31-year sample. Of course, optimal ex-ante consumption risk sharing may not be a reasonable benchmark, in the sense that international asset markets might be unlikely ever to allow countries to insure against long-run growth differentials. I consider an alternative where the counterfactual Pareto weights are calculated to smooth deviations (of the estimated Pareto weights) from trend, rather than to eliminate trends. The results of this exercise (presented in the online Appendix) are qualitatively similar to those from the baseline exercise: the median OECD country gains from this exercise, but less than the median non-OECD country gains.

Zero trade costs

Eliminating trade costs in addition to moving to optimal risk sharing in the world as a whole leads to estimated increases in welfare that dwarf those from optimal risk sharing alone. This is not surprising, given the estimated size of trade costs. While clearly of first order importance for understanding the failure of *perfect* consumption risk sharing, the policy relevance of this finding is less clear, as the estimated trade costs are very substantial relative to policy barriers to trade (e.g. tariffs).

Robustness

I examine the robustness of the results along the following dimensions. Details of these results are reported in the online Appendix. First, I look at how the results change when different values are chosen for η and ρ. The effect of changing these parameters on the magnitude of welfare changes is not straightforward, because changing each parameter affects the baseline consumption allocation and Pareto weight estimates, the counterfactual point on the Pareto frontier, and the price and welfare effects of consumption reallocation. However the comparative static result that the welfare gains to the median OECD country from optimal risk sharing within the OECD are smaller than those to the median non-OECD country from optimal risk sharing in the world as a whole is unaffected by choosing different parameter values. The gains to all countries from eliminating trade costs are always starkly greater than those from optimal risk sharing alone.

Second, I perform the same three counterfactual exercises at baseline parameter values, but for an alternative baseline point on the Pareto frontier that applies discounting from the perspective of 1970 to calculate weighted averages over time of relative Pareto weights within risk-sharing groups. Corresponding to this, I apply a different welfare measure that discounts utility from the perspective of 1970. The results from this exercise are both qualitatively and quantitatively very similar to those from the baseline exercise.

Finally, I perform the same three counterfactual exercises at baseline parameter values, but holding relative Pareto weights within risk-sharing groups constant at their estimated levels for 1970. Once again, these are not necessarily those that would have obtained had the world shifted to frictionless asset markets in 1970, because the intertemporal budget constraint actually faced by countries in 1970 was different from what they would have faced under frictionless asset markets. The results of this exercise imply that the median OECD country loses substantially from risk sharing among OECD countries, while the median non-OECD country gains substantially from risk sharing in the world as a whole.

6 Conclusion

Why does perfect international consumption risk sharing fail? The previous literature has suggested both trade costs and asset market frictions as important candidate explanations. There has been some debate over the relative importance of these explanations, and the relevance of each for developed and developing countries. I show that a gravity model with an intertemporal dimension provides a natural framework for examining these issues. I use the model both to test for the presence of the two frictions, and to estimate the ex-post gains relative to historical welfare from eliminating them. The results from both exercises suggest that OECD countries are relatively closer to optimal risk sharing than non-OECD countries, and by some measures, are very close indeed to optimal risk sharing conditional on trade costs. In contrast, trade costs are of first order importance in welfare terms, principally because they affect consumption of variety within dates and states, but also because they affect countries' ability to smooth consumption across dates and states.

The results on trade costs are in line with the findings of the previous literature (though an explicit comparison of the welfare losses associated with trade costs and asset market frictions is new). The results on optimal risk sharing contrast with the previous literature, in that they suggest that for OECD countries, the scope for welfare gains from optimal risk sharing conditional on trade costs is small. However it is interesting that this is not the case for the world as a whole, as under the very same parameter values, the results indicate that many developing countries gain significantly from being able to optimally smooth consumption with the rest of the world. There is a substantial body of interesting research devoted to understanding the exact nature of these frictions. By clarifying their relative importance for different groups of countries, this paper hopes to contribute to directing that literature.

References

[1] Alvarez, F. and R. Lucas, [2007], "General equilibrium analysis of the Eaton-Kortum model of international trade," *Journal of Monetary Economics*, 54(6), 1726-1768.

[2] Anderson, J. and E. van Wincoop, [2003], "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review* 93, 170-192.

[3] Anderson, J. and E. van Wincoop, [2004], "Trade Costs," *Journal of Economic Literature* 42, 691-751.

[4] Arkolakis, C., A. Costinot and A. Rodriguez-Clare, [2009], "New Trade Models, Same Old Gains?," NBER Working Papers 15628.

[5] Backus, D., P. Kehoe and F. Kydland, [1992], "International Real Business Cycles," *Journal of Political Economy* 100 (4), 745-775.

[6] Chinn, M. and H. Ito, [2008], "A New Measure of Financial Openness" *Journal of Comparative Policy Analysis*, 10(3), 309–322.

[7] Cole, H. and M. Obstfeld, [1991], "Commodity Trade and International Risk Sharing," *Journal of Monetary Economics* 28, 3-24.

[8] Dekle, R., J. Eaton and S. Kortum, [2008], "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 55(3), 511-540.

[9] Dumas, B., [1992], "Dynamic Equilibrium and the Real Exchange Rate in a Spatially Separated World," *Review of Financial Studies* 5 (2), 153-180.

[10] Dumas, B. and R. Uppal, [2001], "Global Diversification, Growth and Welfare with Imperfectly Integrated Markets for Goods," *Review of Financial Studies* 14 (1), 277-305.

[11] Eaton, J. and S. Kortum, [2002], "Technology, Geography and Trade," *Econometrica* 70 (5), 1741-1779.

[12] Fitzgerald, D., [2008], "Can Trade Costs Explain Why Exchange Rate Volatility Doesn't Feed Into Consumer Prices?" *Journal of Monetary Economics*, 55(3), 606-628.

[13] Gourinchas, P. and O. Jeanne, [2006], "The Elusive Gains from International Financial Integration," *Review of Economic Studies* 73, 1-27.

[14] Heathcote, J. and F. Perri, [2000], "Financial Autarky and International Business Cycles," *Journal of Monetary Economics* 49 (3), 601-627.

[15] Heathcote, J. and F. Perri, [2004], "Financial Globalization and Real Regionalization," *Journal of Economic Theory* 119, 207-243.

[16] Helpman, E., M. Melitz and Y. Rubinstein, [2008], "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, 123(2), 441-487.

[17] Kollmann, R. [1995], "Consumption, Real Exchange Rates and the Structure of International Asset Markets," *Journal of International Money and Finance* 14 (2), 191-211.

[18] Kose, M. and K.-M. Yi, [2006], "Can the Standard International Business Cycle Model Explain the Relation Between Trade and Comovement?" *Journal of International Economics,* 68(2), 267-295.

[19] Krugman, P., [1980], "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950-59.

[20] la Porta, R., F. Lopez-de-Silanes, A. Shleifer and R. Vishny, [1999], "The Quality of Government," *Journal of Law, Economics and Organization* 15 (1), 222-279.

[21] Lewis, K., [1996], "What Can Explain the Apparent Lack of International Consumption Risk Sharing?" *Journal of Political Economy* 104 (2), 267-297.

[22] Mazzenga, E. and M. Ravn, [2004], "International Business Cycles: The Quantitative Role of Trade Costs," *Journal of International Money and Finance* 23 (4), 645-71.

[23] Obstfeld, M. and K. Rogoff, [2000], "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause," in Ben Bernanke and K. Rogoff (eds.), *NBER Macroeconomics Annual 2000* (Cambridge: MIT Press), 339-390.

[24] Ravn, M., [2001], "Consumption Dynamics and Real Exchange Rates," mimeo.

[25] van Wincoop, E., [1999], "How Big Are Potential Welfare Gains From International Risksharing," *Journal of International Economics*, 47, 109-135.

Table 1: Baseline test results

Null	Alternative	F-stat	# rest	d.f.	p-val
Full sample					
Optimal risk sharing	Trade costs, asset mkt friction	6.52	2608	234453	1
No trade costs	Trade costs, asset mkt friction	47.77	5580	234453	1
Financial autarky	Trade costs, asset mkt friction	5.87	2697	234453	1
OECD only					
Optimal risk sharing	Trade costs, asset mkt friction	0.71	628	13485	0
No trade costs	Trade costs, asset mkt friction	39.70	1488	13485	1
Financial autarky	Trade costs, asset mkt friction	2.62	651	13485	1
1970-1984, Full sample					
Optimal risk sharing	Trade costs, asset mkt friction	6.29	1216	113445	1
No trade costs	Trade costs, asset mkt friction	47.68	2700	113445	1
Financial autarky	Trade costs, asset mkt friction	6.37	1305	113445	1
1985-2000, Full sample					
Optimal risk sharing	Trade costs, asset mkt friction	1.64	1303	121008	1
No trade costs	Trade costs, asset mkt friction	69.86	2880	121008	1
Financial autarky	Trade costs, asset mkt friction	5.46	1392	121008	1

Notes: This table reports results from F-tests of null hypothesis against alternative. Tests are based on estimating restricted and unrestricted log-linear gravity models of trade, as described in the text. Baseline sample includes 88 countries as listed in the Appendix, annual data 1970-2000. All bilateral pairs including trade with self are included. Zeros in bilateral trade are replaced by 1 to generate the dependent variable, which is bilateral imports normalized by importer's expenditure and exporter's gross output. Baseline gravity variables include log distance and six indicator variables constructed to normalize trade costs to zero within countries.

Table 2: Summary statistics on compensating variation

	avg	min	p25	p50	p75	max
Optimal risk sharing in OECD only						
All	0.99	0.95	0.98	0.98	0.99	1.16
OECD	1.01	0.97	0.99	1.00	1.03	1.16
non-OECD	0.98	0.95	0.98	0.98	0.99	1.00
Optimal risk sharing in full sample						
All	1.08	0.95	0.99	1.02	1.11	2.13
OECD	1.00	0.96	0.97	0.99	1.02	1.16
non-OECD	1.10	0.95	1.01	1.03	1.12	2.13
Optimal risk sharing for all, no trade costs						
All	2.23	1.55	1.96	2.19	2.38	4.36
OECD	1.86	1.55	1.78	1.86	1.95	2.15
non-OECD	2.36	1.85	2.11	2.28	2.43	4.36

Notes: This table reports summary statistics on the distribution of δ, the measure of compensating variation based on a measure of ex-post welfare that is a simple average of per-period welfare. These distributions are reported for the three counterfactual exercises described in the text relative to the baseline estimated distribution of real consumption, and for the full sample as well as the OECD and non-OECD subsamples. The three counterfactual exercises are first, the imposition of optimal risk sharing between OECD countries, second, the imposition of optimal risk sharing in the world as a whole, and third, optimal risk sharing and zero trade costs in the world as a whole. The point on within-risk-sharing-group Pareto frontier is chosen based on the simple time-series average of the within-risk-sharing group estimated Pareto weights. Pareto weights are held fixed at their estimated levels for countries outside the risk-sharing group, while the weight of the risk-sharing group as a whole varies with respect to the weights of countries outside the group as estimated. The baseline in each case is based on the structural estimation of the nonlinear gravity equation. For both fitted and counterfactual exercises, $\{\eta, \rho\} = \{6, 2\}$.

Table 3: Net exports over GDP: Actual and counterfactual

	avg	min	p50	max
Actual				
All	0.07	0.00	0.05	0.38
OECD	0.03	0.00	0.02	0.11
non-OECD	0.08	0.00	0.06	0.38
Optimal risk sharing, OECD only				
All	0.07	0.00	0.05	0.37
OECD	0.05	0.00	0.04	0.18
non-OECD	0.07	0.00	0.05	0.37
Optimal risk sharing, All				
All	0.20	0.00	0.10	1.77
OECD	0.06	0.00	0.05	0.19
non-OECD	0.24	0.00	0.15	1.77
Optimal risk sharing, no trade cost				
All	0.29	0.01	0.26	1.47
OECD	0.17	0.02	0.16	0.51
non-OECD	0.33	0.02	0.30	1.47

Notes: This table reports the average across time of summary statistics of the cross-section distribution of net exports over GDP. These are reported for the whole sample, and for OECD and non-OECD countries only. The actual values (equal to the predicted values from the structural estimation of the nonlinear gravity equation by construction) are subject to a renormalization that imposes a zero aggregate trade balance for the sample. Statistics are reported for three different counterfactual exercises. The three counterfactual exercises are first, the imposition of optimal risk sharing between OECD countries, second, the imposition of optimal risk sharing in the world as a whole, and third, optimal risk sharing and zero trade costs in the world as a whole. In each chase, $\{\eta, \rho\} = \{6, 2\}$.

Figure 1: Predicted and actual bilateral trade flows

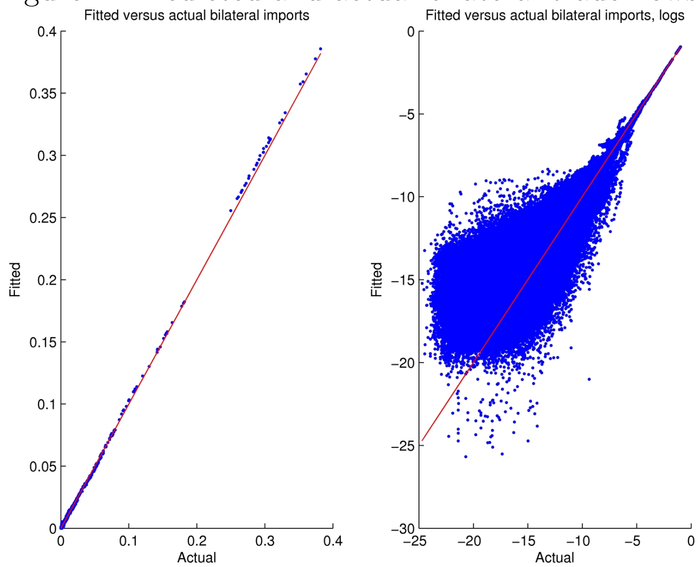

Notes: This figure shows two scatter plots of predicted against actual bilateral trade flows. The predictions are taken from the structural estimation of the nonlinear gravity equation. Actual bilateral imports in this case are the bilateral imports renormalized as in-sample shares, which are used to construct the dependent variable in the estimation. The first figure shows the scatter plot for bilateral imports in levels, where each bilateral flow is expressed as a share of world output. The second figure shows the scatter plot for the same variables, but in logs.

Figure 2: Predicted and actual trade-to-GDP ratios

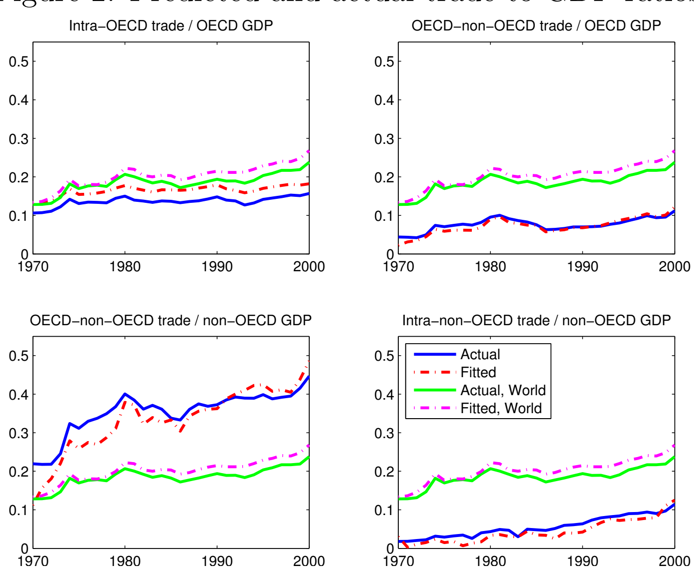

Notes: The panels of this figure plot the evolution over time of predicted and actual trade-to-GDP ratios for trade between different groups of partners, and normalized by the GDP of different groups of countries. The predictions are taken from the structural estimation of the nonlinear gravity equation. By construction, predicted and actual GDP are equal. The first panel shows the ratio of within-OECD trade (excluding a country's imports from itself) to OECD GDP. The second panel shows the ratio of trade between OECD and non-OECD countries to OECD GDP. The third panel shows the ratio of trade between OECD and non-OECD countries to non-OECD GDP. The final panel shows the ratio of within-non-OECD trade (excluding a country's imports from itself) to non-OECD GDP. All panels show the ratio of world trade (excluding a country's imports from itself) to world GDP.

Figure 3: Median trade costs for different types of trade

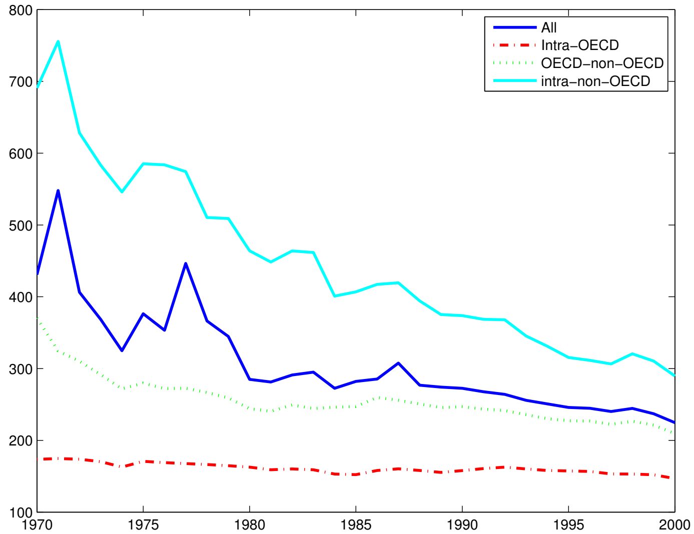

Notes: This figure plots the time-series evolution of the median across trade between specific groups of partners of the fitted values of trade costs based on the structural estimation of the nonlinear gravity equation, and the assumption that $\eta = 6$. The types of bilateral trade are intra-OECD trade, trade between OECD and non-OECD countries, and intra-non-OECD country trade. Also reported is the median trade cost for all trade. Trade costs are expressed as a percentage of the sales price, so $\tau = 3$ implies a trade cost of 200%.

Figure 4: Coefficient of variation of Pareto weights from structural estimation

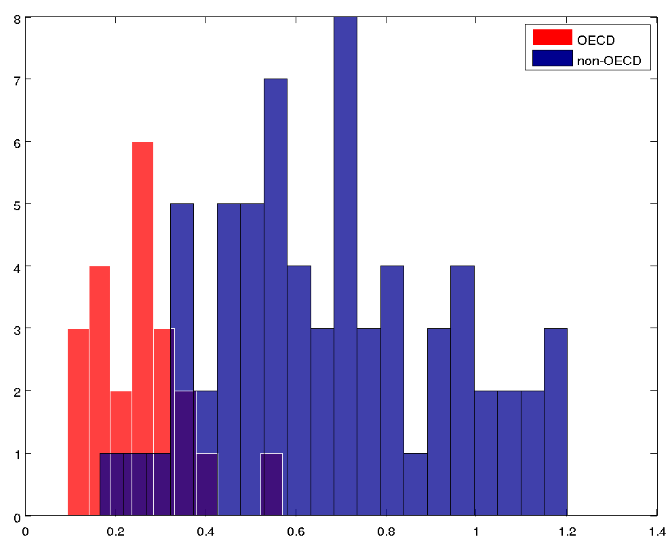

Notes: This figure plots the distribution of the within-country coefficient of variation (standard deviation/mean) of estimated Pareto weights expressed relative to the rest of the world. The Pareto weight estimates are based on the structural estimation of the nonlinear gravity equation, combined with the assumption that $\{\eta, \rho\} = \{6, 2\}$. The distributions for OECD and non-OECD countries are plotted separately.

Figure 5: Normalized Pareto weight estimates for select countries

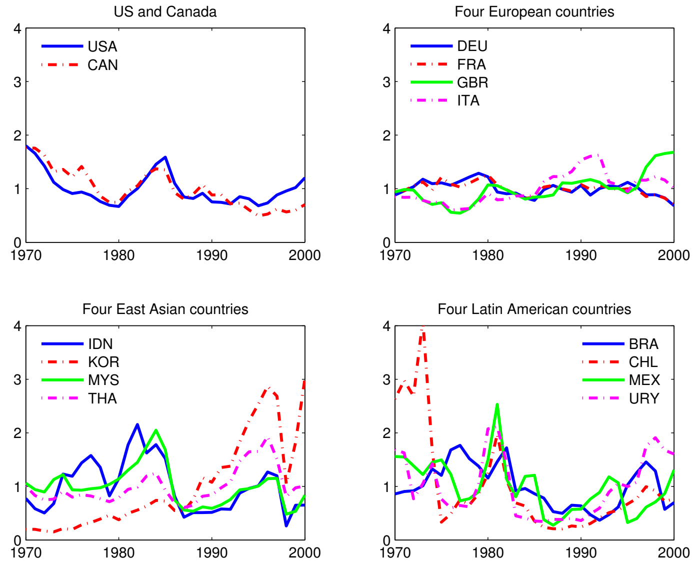

Notes: This figure plots the time-series evolution of estimated Pareto weights for select groups of countries. The Pareto weight estimates are based on the structural estimation of the nonlinear gravity equation, combined with the assumption that $\{\eta, \rho\} = \{6, 2\}$. Weights are normalized by the within-country cross-time mean so different countries can be plotted on the same axes.

Figure 6: Predicted and counterfactual trade-to-GDP ratios

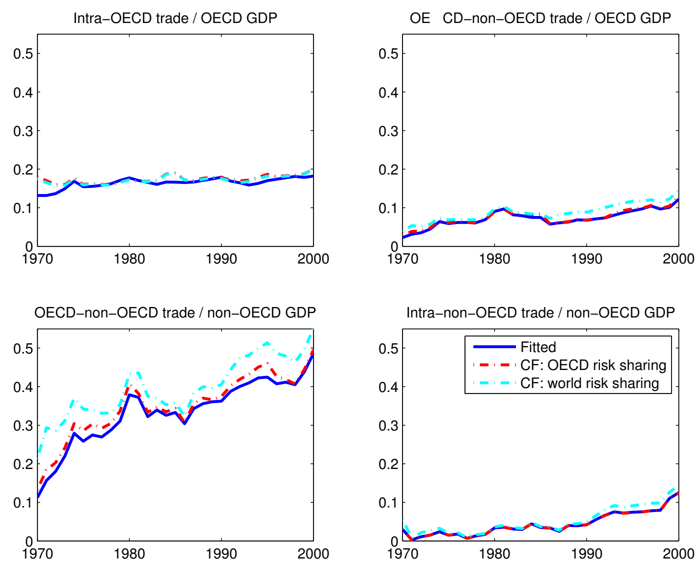

Notes: The panels of this figure plot the evolution over time of predicted and counterfactual trade-to-GDP ratios for trade between different groups of partners, and normalized by the GDP of different groups of countries. The predictions are taken from the structural estimation of the nonlinear gravity equation. By construction, predicted and actual GDP are equal. The two counterfactual exercises for which trade-to-GDP ratios are reported are first, the case of optimal risk sharing for the OECD alone, and second, optimal risk sharing in the world as a whole. The first panel shows the ratio of within-OECD trade (excluding a country's imports from itself) to OECD GDP. The second panel shows the ratio of trade between OECD and non-OECD countries to OECD GDP. The third panel shows the ratio of trade between OECD and non-OECD countries to non-OECD GDP. The final panel shows the ratio of within-non-OECD trade (excluding a country's imports from itself) to non-OECD GDP. All panels show the ratio of world trade (excluding a country's imports from itself) to world GDP.