Trade Costs, Asset Market Frictions and Risk Sharing*

Doireann Fitzgerald†

July 2011

Abstract

I use bilateral import data to test for and quantify the importance of trade costs and asset market frictions in explaining the failure of perfect international consumption risk sharing. I find that while frictions in international asset markets significantly impede optimal consumption risk sharing between developed and developing countries over the period 1970-2000, developed countries are close to optimal risk sharing with each other. Trade costs, in contrast, significantly impede risk sharing for all countries.

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 international consumption risk sharing include costs of trading goods internationally and deviations of international asset markets 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

*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

to perfectly share consumption risk, while 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, the literature finds that frictions in asset markets also play an important role in impeding optimal consumption risk sharing between rich countries (see Backus and Smith [1993], Kollmann [1995] and Ravn [2001]).

I first develop a new test that that makes use of the information contained in bilateral import data to separately identify the role of trade costs and asset market frictions in impeding perfect consumption risk sharing. I implement the test on a set of 88 developed and developing countries over the period 1970-2000. I then explicitly quantify the welfare gains for these countries over this period, first from moving from historical frictions to optimal consumption risk sharing conditional on historical trade costs, and second, from additionally eliminating historical trade costs. The results of both tests and welfare analysis point in the same direction. In contrast with the previous literature, I find that while frictions in international asset markets significantly impede optimal consumption risk sharing between developed and developing countries, developed countries are close to optimal risk sharing with each other. Trade costs, in contrast, significantly impede risk sharing for all countries.

To motivate the tests and to provide a framework for calculating welfare gains, 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 international trade literature (see Arkolakis, Costinot and Rodriguez-Clare [2010]). 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 demand and iceberg costs of trade, this yields predictions about the matrix of bilateral trade which are observationally equivalent to the predictions of the Ricardian model of Eaton and Kortum [2002], and 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 - both within and across states and periods - are limited by the presence of resource costs of trade. I do not take a stand on the structure of asset markets. 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. Hence, a natural metric for the salience of asset market frictions in impeding optimal consumption risk sharing 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 framework for testing for the role of trade costs and asset market frictions in limiting consumption risk sharing. 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 import 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 compares the ability of a restricted gravity model to explain the variation in bilateral trade data with that of an unrestricted gravity model. The restricted model imposes a constant relative inverse marginal utility of wealth for each country; the unrestricted model allows it to vary over time. 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 import data for 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 developed and developing countries conditional on trade costs is also rejected. However the results for developed countries are strikingly different from the previous literature: the test does not reject the null hypothesis of optimal consumption risk sharing within developed countries.

Next, I use the model to provide a more continuous metric of relative distance from optimal consumption risk sharing. More precisely, I structurally estimate the model to benchmark historical trade costs, and the historical impact of asset market frictions on consumption risk sharing. I then solve the problem of a planner who reallocates the output historically devoted to consumption across countries, consistent with optimal risk sharing conditional on historical trade costs. I calculate the implied changes in trade and ex-post welfare relative to the historical benchmark. I do this twice, first imposing optimal risk sharing between developed countries only, and second, imposing optimal risk sharing in the world as a whole. This allows me to see how far developed countries are from optimal risk sharing with each other relative to how far the world as a whole is from optimal risk sharing. I also examine the incremental effect of eliminating trade costs in addition to imposing optimal world risk sharing.

Under optimal consumption risk sharing within developed countries alone, trade within developed countries increases modestly, and the trade balances of developed countries widen. However ex-post welfare in the median developed country is almost unchanged from the benchmark. In contrast, under optimal risk sharing in the world as a whole, there are substantial increases in trade, both between developed and developing countries, and within developing countries. Developing country trade balances widen very significantly. The increase in ex-post welfare in the median developing country is on the order of 3%, while there is a reduction of 1% in welfare in the median developed country. This suggests that over the sample period, developed countries were closer to optimal risk sharing with each other than was the world as a whole. The effect on trade and welfare of eliminating trade costs is enormous for all countries. This is not surprising given that in the context of the model, the size of the trade costs needed to rationalize observed openness are very large.

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 perfect consumption risk sharing null for countries classified as having relatively 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 developed with developing countries. I also base my test for optimal risk sharing on a different metric: the time-series behavior of the estimated inverse marginal utility of wealth, rather than correlations of consumption and real exchange rates.

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 [2010] 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 nonzero 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 the 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 $\pi_t(s^t)$.

Preferences and technology

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

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

where L_t^i is the deterministic population of country i, $C_t^i(s^t)$ is total consumption and $H_t^i(s^t)$ 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:

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

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

The production function for the final good, X is:

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

where $Z_t^{ik}(s^t)$ 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:

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

where $I_t^i(s^t)$ is investment in country i. Capital in country i accumulates according to:

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

where δ is the rate of depreciation.

Resource costs of trade

Intermediate goods trade may be costly: in order for one unit of j's good to arrive in i, $\tau_t^{ij}(s^t)$ units must be shipped, with $\tau_t^{ii}(s^t) = 1$, $\tau_t^{ij}(s^t) \geq 1$ and $\tau_t^{ij}(s^t)\tau_t^{jk}(s^t) \geq \tau_t^{ik}(s^t)$. Trade costs need not be symmetric, i.e. it may be the case that $\tau_t^{ij}(s^t) \neq \tau_t^{ji}(s^t)$. 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 $\tau_t^{ki}\left(s^t\right)Q_t^{i}\left(s^t\right)$. All prices are expressed in terms of a numeraire world currency.

Asset market

Country i enters s^t with a vector of asset holdings $\mathbf{B}_t^i\left(s^{t-1}\right)$ that pays dividends with value given by $\mathbf{D}_t\left(s^t\right)\cdot\mathbf{B}_t^i\left(s^{t-1}\right)$. The vector of asset prices, taken as given by each country, is $\mathbf{R}_t\left(s^t\right)$. The value of asset holdings is given by $\mathbf{R}_t\left(s^t\right)\cdot\mathbf{B}_t^i\left(s^{t-1}\right)$. The country can choose to re-optimize by selling $\mathbf{B}_t^i\left(s^{t-1}\right)$ and purchasing a new vector $\mathbf{B}_{t+1}^i\left(s^t\right)$, with $\mathbf{B}_{t+1}^i\left(s^t\right)\in\mathcal{B}_i\left(s^t,\mathbf{K}_t\left(s^{t-1}\right),\mathbf{B}_t\left(s^{t-1}\right)\right)$. Assets are distinguished by their dividend vectors, while the asset market structure determines $\mathcal{B}_i\left(s^t,\mathbf{K}_t\left(s^{t-1}\right),\mathbf{B}_t\left(s^{t-1}\right)\right)$, 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 takes prices $\mathbf{R}_t\left(s^t\right)$, $\mathbf{D}_t\left(s^t\right)$ and $\mathbf{Q}_t\left(s^t\right)$ as given and chooses $C_t^i\left(s^t\right)$, $H_t^i\left(s^t\right)$, $\mathbf{Z}_t^i\left(s^t\right)$, $M_t^i\left(s^t\right)$, $K_{t+1}^i\left(s^t\right)$ and $\mathbf{B}_{t+1}^i\left(s^t\right)$ to maximize expected utility (1)

[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, see e.g. Alvarez and Jermann [2000].

subject to its aggregate good resource constraint:

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

its budget constraint, expressed in nominal terms:

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

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

and $\mathbf{B}_{t+1}^i\left(s^t\right) \in \mathcal{B}_i\left(s^t, \mathbf{K}_t\left(s^{t-1}\right), \mathbf{B}_t\left(s^{t-1}\right)\right)$. The left hand side of the budget constraint 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. 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).

Some first order conditions

Assume the existence of a competitive equilibrium.[2] Then at any s^t, given $\mathbf{K}^*$ and $\mathbf{B}^*$, the first order conditions with respect to absorption of each intermediate and with respect to aggregate consumption are necessary for competitive equilibrium (i.e. the problem of country i at s^t is convex).[3]

The first order condition with respect to $Z_t^{ik}\left(s^t\right)$ 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 $C_t^i\left(s^t\right)$ 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 $\lambda_t^i\left(s^t\right)$ 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), $u_c(\cdot)$

[2] Alvarez and Lucas [2007] give conditions under which existence can be proven in a model of this type with financial autarky.

[3] I do not make use of the other necessary conditions.

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.:

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

The $\lambda_t^i(s^t)$'s can also be interpreted 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 and frictionless asset markets and optimal risk sharing

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 time-and-state-invariant $\left\{\lambda^1,\ldots,\lambda^N\right\}$ the planner chooses sequences $\mathbf{C}$, $\mathbf{H}$, $\mathbf{Z}$, $\mathbf{M}$, and $\mathbf{K}$ 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 $\mathbf{C}$ and $\mathbf{Z}$ are exactly (11) and (10), with constant λ^i. Any allocation that is the outcome of a planning problem with time-and-state-invariant $\left\{\lambda^1,\ldots,\lambda^N\right\}$ satisfies optimal consumption risk sharing. However the competitive equilibrium allocation under complete and frictionless asset markets is exactly the optimal consumption risk sharing allocation where the $\left\{\lambda^1,\ldots,\lambda^N\right\}$ are such that the allocation satisfies the ex-ante budget constraints.

3 Why does consumption risk sharing fail? A test

Tests for consumption risk sharing make use of the first order condition for aggregate consumption, (11). Dropping state-contingent notation for brevity, the relative marginal utility of consumption between country i and country j can be written:

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

This expression illustrates the fact that the relative marginal utility of consumption can vary over time and across states for two reasons. First, if there are frictions in asset markets, λ_t^i/λ_t^j need not be constant.[4] In particular, relative wealth may respond to current output. 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^j/P_t^i, which in general will comove 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. A disadvantage of this approach is that if traded goods prices differ across countries, the test cannot distinguish between failures of perfect risk sharing due to frictions in goods markets, and failures of perfect 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 the relationship between changes in real exchange rates and relative consumption. Under the null hypothesis of complete and frictionless asset markets, λ_t^i/λ_t^j is constant and the correlation between the relative marginal utility of consumption and real exchange rates is equal to -1. This literature examines the correlation between relative consumption and real exchange rates for developed countries, and finds that it is close to zero or of the opposite sign from that predicted given constant λ^i/λ^j and concave utility. The conclusion drawn from this finding is that there must be frictions in asset markets impeding consumption risk sharing between developed countries. This approach has the disadvantage that it does not provide a natural metric for the relative salience of asset

[4]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 in spite of frictions in asset markets. See e.g. Cole and Obstfeld [1991].

market frictions and trade costs in impeding perfect consumption risk sharing. Moreover, it does not provide a natural way of identifying the relative salience of asset market frictions between different groups of countries and for different periods of time.

I now describe a unified testing framework that can separately identify the role of asset market frictions and trade costs in the failure of perfect consumption risk sharing, and provide a metric for the salience of frictions for different groups of countries and periods. It has the further advantage that it provides a way of measuring whether international asset markets are further from complete markets or from financial autarky. The test is based on the relative ability of restricted and unrestricted gravity equations to match the observed pattern of bilateral trade. It uses the prices consistent with a gravity model of trade (rather than measured prices) to decompose the value of consumption into price and quantity terms. Identification of asset market frictions relies on the first order condition for consumption, but it is the volatility of λ_t^i/λ_t^j over time rather than the correlation of relative prices with relative consumption that is used to identify the presence of frictions in asset markets.

3.1 Tests based on bilateral import flows

The model described in the previous section yields an expression for bilateral imports that takes the gravity form. Combining the first order condition for intermediate goods (10), with the resource constraints for intermediate goods we get:[5]

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

where

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

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 price terms that are known in the gravity literature as "multilateral resistance" terms.[6] This expression is valid for any asset market structure, under the assumed utility function and under the

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

[6]This terminology is due to Anderson and van Wincoop [2003]. It refers to the fact that the price terms capture the fact that conditional on τ_t^{ik}, i will import more from k if i is distant from other potential trading partners (P_t^i is high) or if k is distant from other potential trading partners (Π_t^k is high).

assumption that trade costs take the iceberg form.[7]

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 [2010]). The treatment of gains from inter-temporal trade here is consistent with that literature.

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

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

Since it is not possible to distinguish between optimal risk sharing that is achieved through complete and frictionless asset markets and optimal risk sharing through other means, I say that (17) holds under the null hypothesis of optimal risk sharing.

In the absence of trade costs, the model predicts that 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 relative 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],[8] under financial autarky and

[7] Helpman, Melitz and Rubinstein [2008] develop a generalization of this expression under fixed and per unit costs of trade and firm heterogeneity. This complicates the gravity expression, but does not affect the point that the gravity equation can be used to test for failures of optimal risk sharing.

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

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. As a baseline, I assume that utility is separable in consumption and leisure, and that marginal utility takes the form:

$$u_{ct}^i(\cdot,\cdot) = \left(C_t^i / L_t^i\right)^{-\rho} \tag{20}$$

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:

$$P_t^i = \left(\lambda_t^i\right)^{\frac{1}{1-\rho}} \left(VC_t^i / L_t^i\right)^{\frac{-\rho}{1-\rho}} \tag{21}$$

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 can be constructed 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 $\boldsymbol{\gamma}_t$ on the gravity variables for each year. In the baseline 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	$w_t^{ik} = \psi^i + \phi_t^k + \beta_c vc_t^i + \sum_{n=1}^J \gamma_{nt} d_n^{ik} + \varepsilon_t^{ik}$
(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, $w_t^{ik} = \ln\left(IM_t^{ik}/EXP_t^i OUT_t^k\right)$, $vc_t^i = \ln\left(VC_t^i/L_t^i\right)$ 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.[9]

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

[9]It is not necessary to impose this restriction in order to implement the test, as equations (a) and (b) are identified without it. Imposing the restriction that (16) implies on the relationship between θ_t^i and ϕ_t^k, would require one to have (or assume that one had) data on the full universe of countries, and to use a nonlinear estimation strategy.

expenditures are taken from the World Bank's *World Development Indicators* (WDI). The employment-population rate (for robustness checks) 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.[10] Data on bilateral service trade flows are not available until the last two years of the sample period, and only for a limited set of bilateral pairs.[11] Additional data to fill in gaps for Taiwan and the former Western Germany are taken from the national statistical agencies of these countries.

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 online appendix, which also discusses sample construction issues. Over the sample period, these countries account for between 90% and 94% of world GDP. Within-sample merchandise trade (i.e. trade that does not involve a partner not in the sample) accounts for between 72% and 83% of world merchandise trade.

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

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

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

[10]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 exercises.

[11]Kimura and Lee (2006) find that for the years 1999 and 2000 and for country-pairs for which the OECD makes available data on bilateral service and merchandise trade flows, the gravity equation fits better for bilateral trade in services than it does for bilateral merchandise trade.

different approaches to dealing with this issue. As a baseline, I 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 when a country trades with itself. This restricts the fitted values of trade costs within a country to be equal to zero.

4.1 Results

Baseline bilateral tests

The results from estimating the four models described above using the full 88-country sample are reported in Tables A.1 to A.4 of 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 optimal 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. While it is not possible to nest the hypotheses in such a way as to test the null of optimal risk sharing against the alternative of financial autarky (or vice versa), these two sets of restrictions do a roughly similar job of matching the variation in the bilateral trade data.

These results are imply that though there is some international asset trade, frictions in international asset markets impede optimal consumption risk sharing in the world as a whole.

However the R^2s of the different models suggest that trade costs may be quantitatively a lot more important in impeding consumption risk sharing than frictions in asset markets.

Do developed and developing countries face different frictions?

To examine whether 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 developed countries, all of which were members of the OECD for the majority of the sample period.[12] 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 developed country subsample.

The results from estimating the four models are reported in Tables A.5 to A.8 of 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 developed country 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 developed country sub-sample. The null hypothesis of optimal consumption risk sharing conditional on trade costs *cannot* be rejected for this sample. This contrasts with rejection of the optimal risk sharing null for the full sample. Athough the implied trade costs are smaller than in the full sample, the null hypothesis of no trade costs is rejected at all levels of significance. The null hypothesis of financial autarky is rejected at all levels of significance, and in contrast to the full sample, this model does a weaker job of explaining the data for developed countries than does the optimal risk sharing model.

Implementation of the same tests for randomly selected sub-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 Figure A.1 of the online Appendix.

These results are consistent with developed countries being closer to optimal risk sharing with each other than they are to optimal risk sharing with the world as a whole, either because their need to share idiosyncratic risk through asset trade is lower, or because their ability to share risk with each other is greater. This contrasts with the findings of the literature following Backus and Smith [1993], which infers the failure of optimal risk sharing

[12]The developed country group consists of Australia, Austria, Belgium-Luxembourg, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom, United States.

for developed countries from the comovement of relative consumption and real exchange rates.

Are frictions declining in importance over time?
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 between the earlier and the later period.

4.2 Robustness

I examine the robustness of these results along several dimensions.

I check the robustness of the results to varying the set of countries defined as developed. I repeat the same set of tests for both narrower definitions (e.g. the 6 countries used by Kollmann [1995], the 8 used by Backus and Smith [1993], the 12 used by Ravn [2001]) and broader definitions (e.g. all countries in the 88-country sample that were OECD members at some point during the sample period, all countries that were members of the OECD in 2010). The sample definitions are in the online appendix, and the results are reported in Table A.9. Adding Turkey (a founding member of the OECD but usually classified as an emerging market) does not affect the results. However adding increasing numbers of emerging market economies to the sample leads to rejection of the null hypothesis of optimal risk sharing. 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 baseline developed country 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. These results are reported in Table A.10 of 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 logit. In all cases, the results are qualitatively unchanged. These results are reported in Table A.11 of the online appendix.

I test the robustness of the results to allowing for nonseparability of utility over consumption and leisure by adding the log of the employment rate as an independent variable in the optimal risk sharing model. The results are not affected. Finally, I 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 of 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. They are reported in Table A.12 of the online appendix.

4.3 Discussion and interpretation

Relation to the previous literature

The test results are considerably more favorable to the null hypothesis of optimal consumption risk sharing within developed 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.

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 may contribute to a different view of consumption risk sharing, I do the following. First, I use data on CPIs and nominal exchange rates (taken from the WDI) 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) as the past literature has done. I then perform the decomposition of the value of consumption 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 developed country is positive and close to zero, similar to what is found by Backus and Smith [1993].[13] For the full 88-country sample, the correlation for the median country is negative and close to zero. However, using the estimated prices, the correlation is systematically and strongly negative, both for developed countries and for the world as a whole. This goes some way towards explaining why the results I present above for developed countries 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 tables A.14 and A.15 of the online Appendix.

However, 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, I reject 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 previous literature infers the presence of frictions from the fact that the correlation of relative prices and relative real consumption deviates from that predicted under the null hypothesis of constant relative marginal utilities of wealth. In contrast, I infer 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. One of the advantages of my approach is that my test is plausibly less likely to reject the null under small deviations (in terms of allocations) from optimal consumption risk sharing, and more likely to reject under big deviations.

A modified test

I now describe a modified test that reconciles my baseline results with those of the previous

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

literature. As a reminder, the gravity equation under the optimal risk sharing null is:

$$w_t^{ik} = \psi^i + \phi_t^k + \beta_c vc_t^i + \sum_{n=1}^{J} \gamma_{nt} d_n^{ik} + \varepsilon_t^{ik}$$

where, following (21), $\beta_c = \frac{-\rho}{1-\rho}(\eta - 1)$. The baseline version of the test for optimal risk sharing allows the data to choose the (time- and country-invariant) combination of values of η and ρ for which a constant inverse marginal utility of wealth best fits. The estimated value of β_c for the developed country sample is negative (equal to -0.30 with standard error 0.06). 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. (a) violates the priors of many macroeconomists, while (b) does not make sense in the context of the model I use.[14] It is straightforward to implement a stronger test, of the joint null of optimal risk sharing *and* a particular value for $\frac{-\rho}{1-\rho}(\eta - 1)$. This is done by subtracting $\frac{-\rho}{1-\rho}(\eta - 1)vc_t^i$ from the dependent variable in the optimal risk sharing model. In Table A.13 of the online appendix, I report the results 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. Can the baseline results be interpreted as implying developed countries are at least closer to optimal risk sharing than the rest of the world? I address this issue explicitly in the next section.

5 Trade and welfare under counterfactual risk sharing

I now use the model from section 2 to explicitly quantify the changes in trade and ex-post welfare that result from moving from historical frictions in asset markets to optimal consumption risk sharing conditional on historical trade costs. I also examine the incremental effect of moving from optimal risk sharing with historical trade costs to optimal risk sharing with zero trade costs. These counterfactual exercises demand that the model be taken more literally than in the previous section.

[14] If $\eta < 1$, trade volumes are increasing in trade costs. The assumption that $\eta < 1$ is often made in two-country models where preference asymmetry rather than trade costs are used to match the volume of gross trade flows.

Before going into the details, it is instructive to perform the following back-of-the-envelope calculation that illustrates what drives the results as they pertain to optimal risk sharing. Assume that the marginal utility of per capita consumption takes the form $(C_t^i/L_t^i)^{-\rho}$, so the first order condition for aggregate consumption implies $\lambda_t^i = P_t^i\,(C_t^i/L_t^i)^{\rho}$. Given a value for ρ, data on real consumption per capita and real consumption exchange rates can then be used to estimate λ_t^i/λ_t^j for pairs of countries i and j. I assume $\rho = 2$ and use WDI data on the value of consumption per capita and the CPI to construct estimates of $\Delta \ln\left(\lambda_t^i/\lambda_t^{US}\right)$ for all the countries in the data set. I then calculate for all $i \neq US$ the standard deviation of this variable over the sample period. Figure 1 shows the distribution of this measure of the volatility of the relative marginal utility of wealth for the sample of developed and developing countries used in the previous section. The distribution for developing countries has significant mass well to the right of the distribution for developed countries - the median for developed countries is 0.22, while the median for developing countries is 0.42. The fact that their relative marginal utility of wealth is less volatile is consistent with developed countries being on average closer to optimal risk sharing with the US (the biggest developed country) than are most developing countries. With this in mind, I move on to the full counterfactual exercises.

To keep the exercises tractable, I restrict them along several dimensions. First, I fix the output devoted to consumption in the world as a whole at its historical level, and ignore the potential impact on output available for consumption of relaxing asset market frictions.[15] Second, the optimal risk sharing allocations are obtained as the solution to a planner's problem with time-invariant Pareto weights, subject to the appropriate resource constraints. I choose the time-invariant Pareto weights with some desirable properties, and look at robustness to varying these choices, but do not attempt to solve for the weights that would have obtained if asset market frictions had been eliminated in 1970. Third, I calculate the resulting changes in ex-post welfare rather than ex-ante welfare. Because of these restrictions, I do not have to take a stand on the nature of historical asset market frictions, the joint process for shocks in the world as a whole, or country wealth in 1970.

I proceed in three stages. In the first stage I estimate functions of historical trade costs and output prices for each year. In contrast to the previous section, the estimation strategy makes use of the full structure of the gravity model and of the information on the extent of asset trade that is contained in data on net exports. In the second stage, I choose values for

[15] Gourinchas and Jeanne [2006] address precisely the question of the welfare effects of financial integration in a world where all of 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.

η and ρ. Together with the estimated functions of historical trade costs and output prices, and data on the value of output, this allows me to construct estimates of trade costs and of the real output historically devoted to consumption.[16] US real output is used to benchmark year-on-year variation. The vector of outputs and the matrix of bilateral trade costs are the fundamentals determining the resource constraints of the counterfactual. Using trade costs and output prices together with data on the value of consumption, I obtain estimates of both historical consumption prices and real consumption, which are used to calculate historical inverse marginal utility of wealth and historical ex-post welfare. In the third stage, the estimates of the historical inverse marginal utility of wealth are used to choose the point on the within-risk-sharing-group Pareto frontier for countries in the risk sharing group, and to fix the evolution of relative wealth at the historical benchmark for countries not in the risk sharing group. I then solve the problem of the planner who each period allocates consumption across countries to maximize the weighted sum of country utilities, subject to the appropriate resource constraints. Based on the resulting allocations, I calculate the implied changes in trade and ex-post welfare relative to the historical benchmark. I do this three times, first imposing optimal risk sharing between developed countries only, second, imposing optimal risk sharing in the world as a whole, and third, eliminating trade costs in addition to imposing optimal world risk sharing.

5.1 Estimation strategy

I exploit the full structure of the model in estimating the cross section of trade costs and relative output prices. For technical reasons, it is more convenient to estimate output prices than consumption prices.[17] The estimating equation is based on the following equation of the model, which is obtained by rearranging the first order condition for Z_t^{ik}:

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

[16]I estimate prices because data on price levels (e.g. from the Penn World Tables) are inconsistent with the structure of the model because they imply negative values of τ.

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

To arrive at the estimating equation, I divide across by total expenditure, take logs, approximate trade costs as is standard in the gravity literature, and include an error term:[18]

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

where $\alpha_t^k = (1-\eta)\ln\left(Q_t^k\right)$, $\mathbf{d}^{ik}\cdot\boldsymbol{\gamma}_t = (1-\eta)\ln\left(\tau_t^{ik}\right)$, and $\mathbf{d}^{ik}$ is a vector of (time-invariant) gravity variables, constructed such that $\exp\left(\mathbf{d}^{ii}\cdot\boldsymbol{\gamma}_t\right) = 1$. This imposes that the estimated Q's and τ's are positive, and that $\tau_t^{ii} = 1$. I normalize $Q_t^1 = 1$ by setting $\alpha_t^1 = 0$. For each t, I choose the vectors $\boldsymbol{\gamma}_t$ and $\boldsymbol{\alpha}_t$ 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.:

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

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 used is exactly the data from 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 $\mathbf{d}^{ij}$ 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 developed country group and the other in the developing country group, and for both trading partners being in the developing country group.[19] Allowing the trade-reducing effect of distance to differ depending on the nature of the bilateral pair is necessary to simultaneously match trade-GDP ratios for both developed and developing countries.

[18] In the baseline estimation I drop observations where $IM_t^{ik} = 0$. The results are robust to including these observations along the lines of the testing section.

[19] These are set equal to zero when a country trades with itself such that $\tau_t^{ii} = 1$ is imposed.

5.2 Model fit

The estimated parameters on the gravity variables are reported in Table A.17 of the online Appendix. The first panel of Figure 2 is a scatter plot of the fitted values of bilateral imports against the data used in estimation, 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 crucial in matching world aggregates, the fit of the estimated model 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 pairs of developing countries. This is illustrated in Figure A.2 of the online Appendix. The relatively poor fit for trade between developing countries is consistent with measurement error being a particular problem for this type of flows.

I also examine the fit of the model in terms of matching the evolution of trade-GDP ratios. This is not an explicit target of the objective function, but fit along this dimension is sensitive to 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-developed-country imports to developed country GDP, the ratio of trade between developed and developing countries to developed country GDP and developing country GDP respectively, and finally, the ratio of within-developing-country imports to developing country GDP. The four panels of Figure 3 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. Within-developed-country trade is slightly overpredicted and within-developing-country trade is slightly underpredicted. 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 estimates of prices, quantities and trade costs, 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 lower values. I choose a benchmark value of 6, and check the robustness of the results to this choice. As in the section of the paper on testing, 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.

The relationship between real output in one period and the next (and the time-series behavior of all other variables) is not pinned down by the estimated functions of relative output prices. I use the evolution of measured real GDP in the US, taken from the WDI to benchmark year-to-year variation in real quantities. This choice affects the welfare calculations, but not predictions on counterfactual trade-GDP ratios or net export-GDP ratios. Using the estimates $\hat{\boldsymbol{\gamma}}_t$ and $\hat{\boldsymbol{\alpha}}_t$, the parameters η and ρ, data on the current dollar value of output, on population and the share of expenditure devoted to consumption, and the time-series of US real GDP over the period, 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.

Trade costs

Figure 4 plots median trade costs for trade within developed countries, trade between developed and developing countries and trade within developing 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.[20] However, as Figure 3 illustrates, trade costs on this order of magnitude do a good job of matching observed trade-GDP ratios.[21] Moreover, while measured trade costs for goods and services that are actually traded are not this high, measured trade costs are not necessarily 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. For each country, I calculate the standard deviation of log changes in $\hat{\lambda}_t^i/\hat{\lambda}_t^{US}$ over the sample period.[22] Figure 5 plots the histogram of the distribution of these standard deviations for developed

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

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

[22] The normalization with respect to the US is for comparability with the Pareto weights calculated based on price data whose distribution was reported in 1.

and developing countries separately. The distribution for developing countries has more mass in the right tail than the distribution for developed countries. The median for developed countries is 0.24, while the median for developing countries is 0.29. This greater volatility is consistent with developing countries being further from optimal risk sharing with the US than are developed countries, exactly as was suggested by the results of the simple exercise reported in figure 1.

Figure 6 illustrates the time-series behavior of the estimated Pareto weights (normalized such that $\sum_{i=1}^{N} \hat{\lambda}_t^i = 1$). for a selected set of countries. This figure demonstrates that there is substantial co-movement within developed countries, and to some extent, within groups of developing countries. However, the volatility of the estimated weights is much more substantial for developing countries. Moreover, their evolution reflects episodes such as the 1980s debt crisis and the East Asian financial crisis.

5.4 Counterfactual exercises

All of the counterfactual exercises involve imposing optimal consumption risk sharing within a particular group of countries. The choice of the point on the within-risk-sharing-group Pareto frontier is crucial. Ideally, I would like to pick the point that would have obtained had the risk-sharing-group had access to complete and frictionless asset markets from 1970. However calculating this would require characterizing each country's lifetime budget constraint under frictionless asset markets in 1970, which is beyond the scope of the paper. Instead, I pick a baseline point on the within-risk-sharing-group Pareto frontier that has some desirable properties, and examine robustness of the results to the baseline choice. The baseline point is chosen by averaging the estimated relative Pareto weights over time within the risk-sharing group. Averaging the historical weights means that the exercise examines the welfare impact of smoothing temporary deviations of relative wealth from long-run average wealth, rather than the welfare impact of permanent wealth transfers from one group of countries to another.

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

$$\tilde{\lambda}_t^i(1) = \begin{cases} \frac{\sum_{t=1970}^{2000} \hat{\lambda}_t^i}{\sum_{j \in rich} \sum_{t=1970}^{2000} \hat{\lambda}_t^j} \left[\sum_{j \in rich} \hat{\lambda}_t^j \right] & \text{if } i \in rich \\ \hat{\lambda}_t^i & \text{if } i \in poor \end{cases} \tag{28}$$

Second, I impose optimal risk sharing for the full sample of countries. I do this by fixing the weight of each country relative to each other at 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, $\mathbf{L}_t$, and $\left\{ \tilde{\boldsymbol{\tau}}_t, \hat{\mathbf{Y}}_t^C, \tilde{\boldsymbol{\lambda}}_t; \eta, \rho \right\}$, 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 $\tilde{\boldsymbol{\lambda}}(1)$ or $\tilde{\boldsymbol{\lambda}}_t(2)$ 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:

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

This measure of welfare weights all periods equally, consistent with how the point on the within-group Pareto frontier is chosen. Compensating variation, δ, is then defined such that:

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

5.5 Results

Optimal risk sharing: trade

I first examine the effects of the two optimal risk sharing experiments on predicted trade-GDP ratios and net export-GDP ratios. I calculate the value of both counterfactual GDP and imports by adding counterfactual real consumption-related flows to the historical flows associated with investment, government consumption and materials. These total real flows are then valued at the appropriate counterfactual prices.

Figure 7 plots the same fitted and counterfactual trade-GDP ratios as those reported in figure 3. Optimal risk sharing within developed countries alone is associated with a modest increase in within-developed-country trade and trade between developed and developing countries. Optimal risk sharing for the world as a whole is associated with similarly modest increases in within-developed-country trade, but substantial increases in trade between developed and developing countries and within developing countries.

Table 3 provides summary statistics on time-series averages of the cross-section distribution of the absolute value of the net exports-GDP ratio. Optimal risk sharing within the developed world alone is associated with a doubling of the median trade balance for developed countries from 2% to 4% of GDP in absolute value, while trade balances for developing countries fall. In contrast, optimal risk sharing in the world as a whole is associated with an increase in median trade balances for developed countries from 2% to 5%, and for developing countries from 6% to 15%. This is consistent with a large portion of consumption smoothing for developing countries being provided by developed countries. This is illustrated in figures A.6 and A.7 in the online appendix, which show that in the world risk sharing counterfactual, developing countries as a whole run large trade deficits in the second half of the sample.

Optimal risk sharing: consumption and welfare

Table 2 reports summary statistics on the distribution of δ, the measure of compensating variation, for the world as a whole, for developed countries, and for developing countries, for all three counterfactual exercises. Given the baseline point on the within-developed-country Pareto frontier, optimal risk sharing between developed countries alone is associated with unchanged welfare for the median developed country.[23] It leads to a reduction in the welfare of the median developing country of the order of 2%. In contrast, optimal risk sharing for the world as a whole is associated with an increase in welfare for the median developing country of 3%, while it leads to a reduction in the welfare of the median developed country of the order of 1%. In Table A.19 of the online appendix I report statistics on the

[23]This is not surprising. There is no guarantee that the baseline point on the Pareto frontier is inside the "lens" of gains from trade, where ex-ante welfare is increased for all countries. Moreover, 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.

distribution of compensating variation weighted by average population share over the period. The implications are qualitatively similar to the unweighted distribution. These results suggest that the potential gains from relaxing frictions in asset markets between developed countries may be relatively small. On the other hand, the gains from relaxing frictions in asset markets between developed and developing countries could be considerably larger.

It is important to make the distinction between being close to optimal risk sharing and asset markets being close to frictionless. The results do not necessarily imply that historically, developed countries enjoyed access to asset markets that were close to frictionless. The difference in welfare gains between developed and developing countries is at least partly due to the fact that because developed countries are rich, they are large, and their fluctuations are strongly correlated with aggregate risk, and hence uninsurable. In contrast, developing countries are small and on average less correlated with aggregate risk. Because they are small, the cost to large developed countries of insuring them against their idiosyncratic risk is modest relative to the benefit they enjoy from this insurance.

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. Elimination of the estimated resource costs of trading has a first-order impact on the level of consumption, as well as affecting consumption variety within dates and states, and countries' incentives to smooth consumption across dates and states. 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 several dimensions, focusing particularly on the choice of Pareto weights.

The baseline choice of Pareto weights for risk sharing groups has the benefit of simplicity. However under the baseline weights, some countries in the risk sharing group lose in terms of ex-post welfare relative to the historical benchmark. It is possible to improve on this. For both risk sharing groups considered - developed countries and the world as a whole - it is possible to find vectors of constant relative Pareto weights such that *all* countries within the risk-sharing group are better off in terms of ex-post welfare relative to the historical benchmark, while the time series of weights of the risk-sharing group relative to the rest of the world are fixed at their estimated levels.[24] The baseline result - that developing countries gain more from optimal risk sharing in the world as a whole than developed countries do from optimal risk sharing with each other - is preserved using these alternative sets of Pareto weights. These results are reported in Table A.20 of the online appendix.

Second, the baseline choice of constant Pareto weights means that the exercise I perform does not just smooth out short-run fluctuations in consumption growth around trend, but also medium-run deviations of growth from long-run trends, or differences in growth trends over different parts of the 31-year sample. This 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 Table A.21 of the online appendix) are qualitatively similar to those from the baseline exercise: the median developed country gains in this case, but less than the median developing country does.

Third, I perform the three counterfactual exercises at baseline parameter values, but for a point on the within-risk-sharing-group Pareto frontier that applies discounting from the perspective of 1970 to calculate weighted averages over time of estimated relative Pareto weights within the 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 and are reported in Table A.22 of the online appendix.

Fourth, I perform the three counterfactual exercises, holding relative Pareto weights within risk-sharing groups constant at their estimated levels for 1970. These are not necessarily the Pareto weights that would have obtained had the relevant groups shifted to frictionless asset markets in 1970, because the intertemporal budget constraints actually faced by countries in 1970 are different from those they would have faced under frictionless asset markets. The results of this exercise imply that the median developed country loses substantially from risk sharing among developed countries, while the median developing country gains substantially from risk sharing in the world as a whole (Table A.21 of the

[24]The algorithm for finding this vector is quite straightforward. Starting from the baseline risk sharing vector of Pareto weights, it reduces marginally the weight of the country that gains the most, and uses this to increase the weight of the country that loses the most. The algorithm terminates when the minimum compensating variation equals 1. There is no guarantee that this condition will be fulfilled, but it is in fact fulfilled in both cases considered.

online appendix).

Finally, 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, the Pareto weight estimates, the counterfactual point on the Pareto frontier, and the price and welfare effects of a given consumption reallocation. However the comparative static result - that the welfare gains to the median developed country from optimal risk sharing within developed countries are smaller than those to the median developing country from optimal risk sharing in the world as a whole - is unaffected. The gains to all countries from eliminating trade costs are always starkly greater than those from optimal risk sharing alone. These results are reported in Tables A.23 to A.27 of the online appendix.

Relationship to past literature

The literature on the welfare gains to international risk sharing has typically focused only on developed countries, and finds a wide range of potential gains, including some that are very large.[25] The gains that I find for developed countries are on the low end of the range in this literature. My approach to calculating welfare gains differs from the previous literature in that I allow for the fact that perfect consumption risk sharing can fail due to costs of trading goods (which are estimated to be large) as well as frictions in asset markets. I calculate separately the gains from optimal risk sharing conditional on trade costs and the incremental gains from further eliminating trade costs. In calculating the gains from optimal risk sharing I calculate the real output that can be reallocated across countries using relative prices estimated so as to be consistent with trade costs. In Figures A.8 and A.9 of the online appendix, I show that these estimated prices imply greater historical consumption volatility than one would infer based on measured relative prices. This is true both for developed and developing countries. In the absence of trade costs, this would imply considerable scope for consumption risk sharing. But precisely because I estimate that trade costs are large, the resulting estimates of welfare gains to optimal consumption risk sharing are modest, especially for developed countries. Relative to the literature on the welfare costs of business cycles, the gains I estimate to developing countries from optimal risk sharing are big. As the exercise that makes use of smoothed Pareto weights illustrates, this is not just because the baseline weights smooth out growth trends.

[25]See van Wincoop [1999].

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 welfare gains (relative to the benchmark of historical frictions) from eliminating them. The results from both exercises suggest that developed countries are relatively closer to optimal risk sharing than developing 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 both the level of consumption and consumption 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 developed 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 could 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 U. Jermann [2000], "Efficiency, Equilibrium and Asset Pricing with Risk of Default," Econometrica, 68(4), 775-797.

[2] 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.

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

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

[5] Arkolakis, C., A. Costinot and A. Rodriguez-Clare, [2010], "New Trade Models, Same Old Gains?," *American Economic Review*, forthcoming.

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

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

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

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

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

[11] 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.

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

[13] 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.

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

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

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

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

[18] Kimura, F. and H.-H. Lee (2006), "The Gravity Equation in International Trade in Services," *Review of World Economics* 142 (1), 92-121.

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

[20] 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.

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

[22] 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.

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

[24] 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.

[25] 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.

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

[27] 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	6.64	2609	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
		22 developed countries only			
Optimal risk sharing	Trade costs, asset mkt friction	0.72	629	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	6.40	1217	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	1.67	1304	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, 22 developed countries only				
All	0.99	0.95	0.98	0.98	0.99	1.16
Developed	1.01	0.97	0.99	1.00	1.03	1.16
Developing	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
Developed	1.00	0.96	0.97	0.99	1.02	1.16
Developing	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
Developed	1.86	1.55	1.78	1.86	1.95	2.15
Developing	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 developed and developing country subsamples. The three counterfactual exercises are first, the imposition of optimal risk sharing between developed 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
Developed	0.03	0.00	0.02	0.11
Developing	0.08	0.00	0.06	0.38
Optimal risk sharing, developed only				
All	0.07	0.00	0.05	0.37
Developed	0.05	0.00	0.04	0.18
Developing	0.07	0.00	0.05	0.37
Optimal risk sharing, All				
All	0.20	0.00	0.10	1.77
Developed	0.06	0.00	0.05	0.19
Developing	0.24	0.00	0.15	1.77
Optimal risk sharing, no trade cost				
All	0.29	0.01	0.26	1.47
Developed	0.17	0.02	0.16	0.51
Developing	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 the absolute value of net exports over GDP. These are reported for the whole sample, and for developed and developing 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 developed 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 case, $\{\eta, \rho\} = \{6, 2\}$.

Figure 1: Volatility of log changes in simple relative Pareto weight estimates

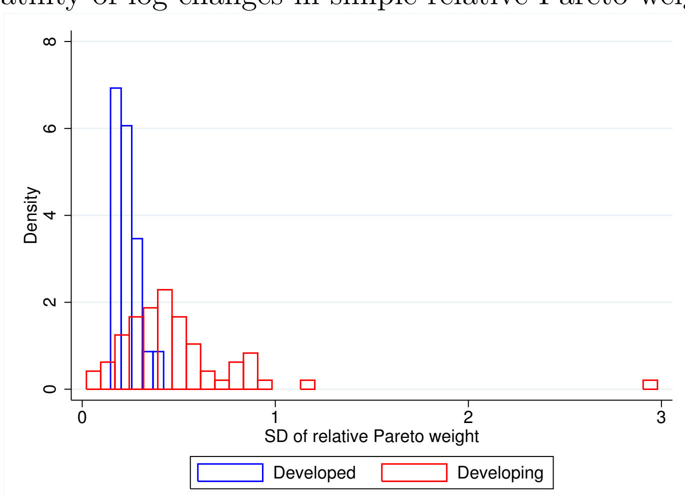

Notes: This figure plots the histogram of the standard deviation of log changes in estimates of Pareto weights relative to the US, for developed and developing countries separately. The estimates are constructed based on the first order condition $\lambda_t^i \left(C_t^i/N_t^i\right)^{-\rho} = P_t^i$, using data on CPIs and nominal exchange rates to recover real consumption from the value of consumption, and assuming $\rho = 2$.

Figure 2: Predicted and actual bilateral trade flows

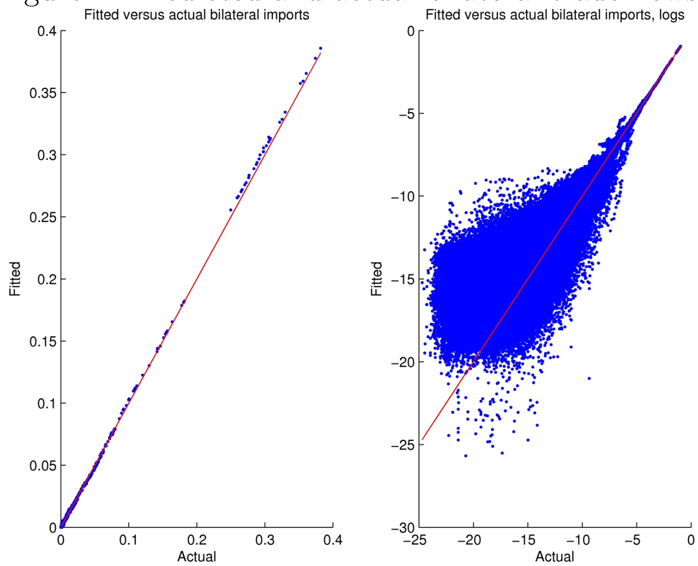

Notes: This figure shows two scatter plots of predicted against actual bilateral trade flows. The predictions are taken from the structural estimation of the nonlinear gravity equation. Actual bilateral imports in this case are the bilateral imports renormalized as in-sample shares, which are used to 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 3: Predicted and actual trade-to-GDP ratios

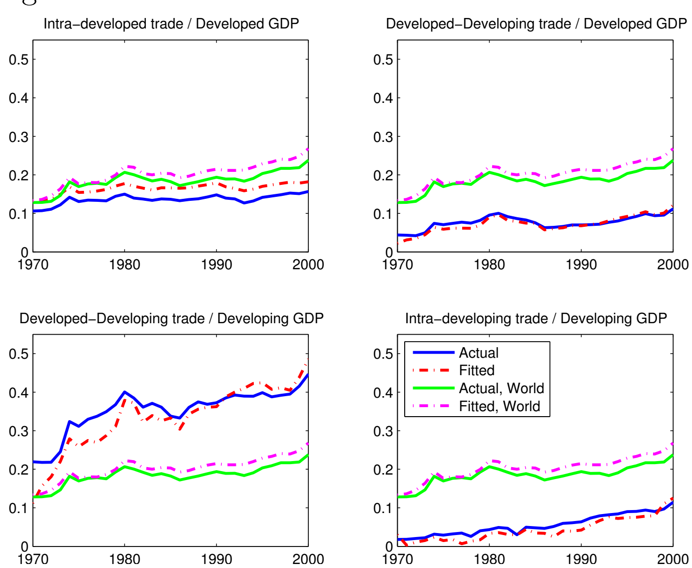

Notes: The panels of this figure plot the evolution over time of predicted and actual trade-to-GDP ratios for trade between different groups of partners, and normalized by the GDP of different groups of countries. The predictions are taken from the structural estimation of the nonlinear gravity equation.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

Figure 4: Median trade costs for different types of trade

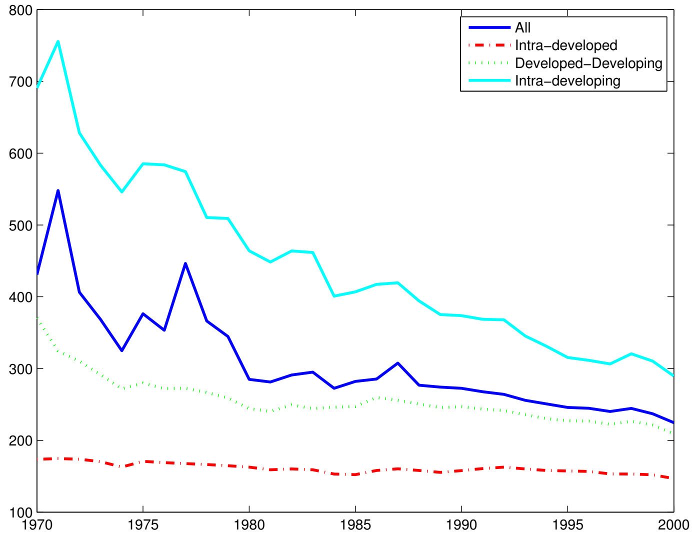

Notes: This figure plots the time-series evolution of the median across trade between specific groups of partners of the fitted values of trade costs based on the structural estimation of the nonlinear gravity equation, and the assumption that $\eta = 6$. The types of bilateral trade are intra-developed country trade, trade between developed and developing countries, and intra-developing-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 5: Volatility of log changes in structural relative Pareto weight estimates

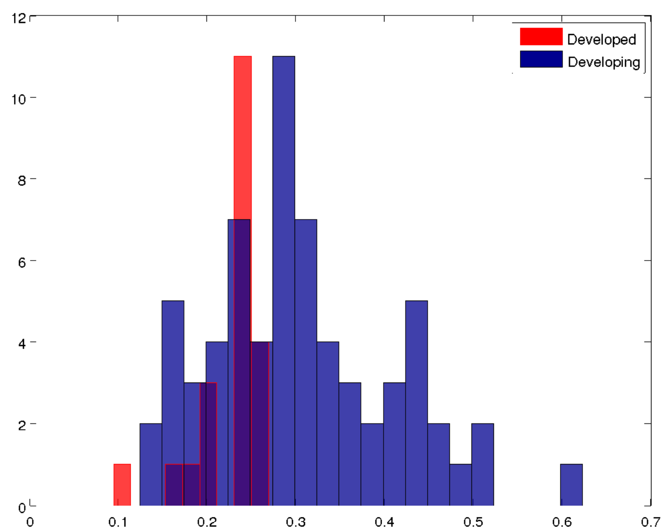

Notes: This figure plots the histogram of the standard deviation of log changes in estimates of Pareto weights relative to the US, for developed and developing countries separately. The Pareto weight estimates are based on the structural estimation of the nonlinear gravity equation, combined with the assumption that $\{\eta, \rho\} = \{6, 2\}$.

Figure 6: Normalized Pareto weight estimates for select countries

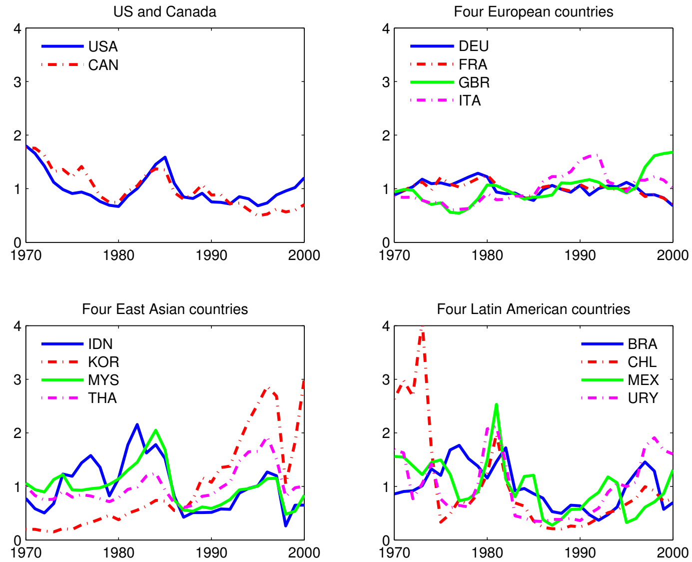

Notes: This figure plots the time-series evolution of estimated Pareto weights for select groups of countries. The Pareto weight estimates are based on the structural estimation of the nonlinear gravity equation, combined with the assumption that $\{\eta, \rho\} = \{6, 2\}$. Weights are normalized by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

Figure 7: Predicted and counterfactual trade-to-GDP ratios

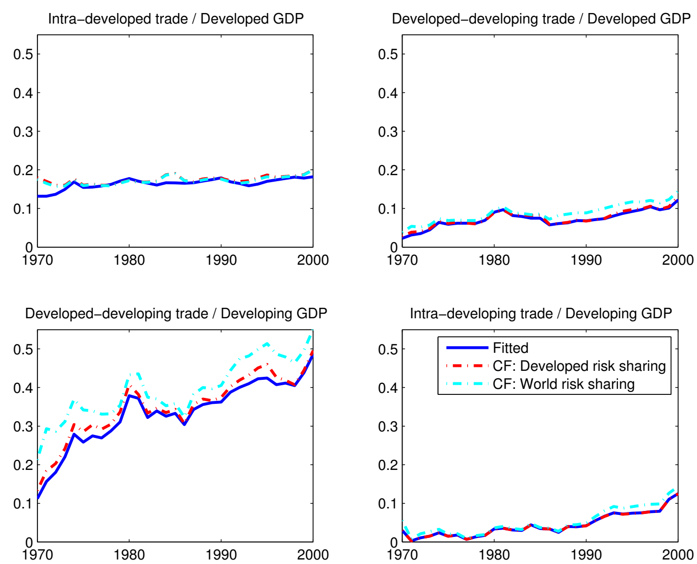

Notes: The panels of this figure plot the evolution over time of predicted and counterfactual trade-to-GDP ratios for trade between different groups of partners, and normalized by the GDP of different groups of countries. The predictions are taken from the structural estimation of the nonlinear gravity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