How Good is International Risk Sharing?

Stepping outside the Shadow of the Welfare Theorems*

Mark Aguiar
maguiar@princeton.edu

Oleg Itskhoki
itskhoki@fas.harvard.edu

Dmitry Mukhin
d.mukhin@lse.ac.uk

March 20, 2025

preliminary — download updated draft here:
https://itskhoki.com/papers/IRS.pdf

Abstract

We revisit whether global output is (Pareto) efficiently distributed across countries over time. The efficient allocation of goods across regions requires that the relative marginal utilities of consumption across countries comoves with the relative costs of the country-specific consumption bundles. Standard approaches to evaluating this property exploit the Welfare Theorems and equate the observed real exchange rate with the social relative costs of consumption. Given the large literature documenting the disconnect between exchange rates, relative prices faced by consumers in a given market, and relative quantities consumed, we develop a methodology that measures relative costs that is robust to this disconnect. We find that relative consumption growth across regions is significantly more correlated with our computed shadow prices than it is with observed real exchange rates, suggesting an allocation closer to efficient risk sharing. Moreover, we provide a decentralization that matches observed prices and quantities, enabling us to rationalize the better implied risk sharing with the failure of the standard correlations. The decentralization involves a combination of segmented foreign exchange markets and pricing-to-market behavior in goods markets. The model implies that consumption allocations are insulated from excess fluctuations in the exchange rate via the optimal equilibrium pricing behavior of exporters.

*We thank Manuel Amador, Andy Atkeson, David Baqaee, Ariel Burstein, Mick Devereux, Alessandro Dovis, Charles Engel, Doireann Fitzgerald, Pierre-Olivier Gourinchas, Xitong Hui, Fabrizio Perri, Rosen Valchev, Steve Wu and Haonan Zhou for helpful comments, and Andrei Voronin for outstanding research assistance. Dmitry Mukhin acknowledges the support from the Economic and Social Research Council (ESRC) [grant number ES/Y001540/1].

1 Introduction

In this paper, we revisit whether global output is (Pareto) efficiently distributed across countries over time. In a world in which countries consume heterogenous baskets of goods, the core question becomes whether the relative shadow costs of delivering a unit of country A's consumption composite versus country B's varies systematically with the relative (Pareto weighted) marginal utilities.[1] The canonical study by Backus and Smith (1993) started a literature that reaches a pessimistic conclusion. Importantly, Backus and Smith (1993) leverages the welfare theorems to equate the unobserved relative shadow price of consumption baskets with the observed real exchange rate. A vast literature, summarized below, concludes that pricing frictions are pervasive both in exchange rate markets and in markets for goods. With this in mind, we revisit this question by developing a methodology that does not rely on the welfare theorems and is robust to both types of frictions.

Our launching point is a standard social planning problem in which a given sequence of output is divided across countries according to fixed Pareto weights. The environment builds on Armington (1969), in which each country produces a specific variety of goods. Our exercise takes world output as given, and hence we abstract from whether production of individual varieties is optimal; our focus is whether the observed production is efficiently distributed across countries. Each country admits a representative consumer that has utility over a country-specific composite of goods. The relative shadow cost of providing consumption to a particular country depends on whether that country's composite is biased toward the goods that are in abundant or scarce supply. From the observed sequences of real output and consumption, we construct the relative shadow price (Marginal Rate of Transformation or MRT) of consumption between two regions, which is the planner's "real exchange rate." Crucially, this implied real exchange rate is constructed using only quantities, and does not rely on observed exchange rates or prices. With the planner's real exchange rate in hand, we can then re-evaluate whether the sequence of relative marginal utilities across regions is correlated with the relative cost of the two consumption commodities.

Our approach allows for two types of distortions, or "wedges" between the data and the planning allocation. The first is the inter-temporal wedge, which corresponds to a violation of risk sharing. Specifically, we measure the extent to which the Pareto weight implied by a country's observed relative consumption is stable over time. To the extent it varies, there is scope for an inter-temporal trade between regions that generates a Pareto improvement. The second distortion is a static distortion that represents a wedge between the marginal rates of substitution (MRS)

[1] With identical preferences, the relative cost is one, and marginal utilities should move in lock step. This motivates the early real business cycle literature's (as in (Backus et al., 1992; Baxter and Crucini, 1995; Ohanian et al., 2018)) focus on consumption correlations.

for different varieties across regions. Absent time-variation in trade costs (a point we consider in an extension), the relative cost between different varieties is the same (up to a constant of proportionality) across regions, and hence a static optimality condition equates the MRS between domestic and foreign goods in each region. Our methodology provides measures of both wedges, and we can access the extent of risk sharing both conditional and unconditional on the static distortion.

We provide a decentralization of the planning problem and map the two distortions into particular market frictions. The static distortion arises due to pricing-to-market frictions. We show that the familiar "alphabet soup" of pricing frictions considered in the literature – producer currency pricing (PCP), local currency pricing (LCP), and dominant currency pricing (DCP) – along with assumptions about market structure generate different implications for the static wedge. The inter-temporal wedge maps into frictions in the foreign exchange market, including segmented markets that decouple the relative price of currencies (and consumption baskets) from the representative agent's stochastic discount factor.

The are several core assumptions we make in mapping the data into the shadow costs. First, we assume each country's consumption commodity is a constant-elasticity composite of world varieties and coincides (in growth rates) with the real consumption index computed by national statistical agencies. This assumption is standard in much of the literature. Second, we need to parameterize preference weights. We do so by using trade shares in a base year, and then hold the preference weights constant over time. For the base year, we are assuming that prices at which goods are traded are the same as those faced by consumers in each region. Third, we need to take a stand on the static distortion in the base year. We assume that this distortion is captured by deviations in measured purchasing power parity across regions, as measured by the Penn World Tables.

TO BE COMPLETED

Related Literature We build on two literatures. The first is that testing international risk sharing. In addition to Backus and Smith (1993), other papers in this literature include Brandt et al. (2006); Bai and Zhang (2010, 2012); Heathcote and Perri (2014). See Heathcote and Perri (2014) for a summary.

Perhaps the closest antecedent to our study is the very important paper by Fitzgerald (2012). That paper and ours share a kindred spirit; namely, using model-based prices rather than observed prices as the appropriate counterparts to the relative costs of consumption in evaluating risk sharing. However, there are key differences, as well. Fitzgerald (2012) uses bilateral trade shares and a gravity model of trade costs to compute implied prices. In particular, all countries have identical preferences and trade in competitive markets, and hence the price of consumption

bundles varies across countries solely due to time-varying trade costs. These trade costs, in turn, are linear functions, with time-varying coefficients, of standard gravity variables. Fitzgerald estimates prices for each country's Armington variety, along with a set of trade-cost coefficients, that best fits the bilateral trade patterns, subject to the vector of global market clearing conditions. Thus, if a country is importing a disproportionate number of goods, its variety must be relatively expensive domestically, which then enters into all other countries' price indices with the addition of the estimated bilateral trade costs. By solving this simultaneous system for each country's traded variety price and its consumption price index, Fitzgerald obtains the implied real exchange rate between all pairs of countries.

In contrast, we exploit home-bias in consumption bundles and use a different set of empirical moments. In particular, under home-bias, a country that has a disproportionately large level of domestic output in a period will correspondingly have a lower relative shadow price for their consumption basket, as its preferred input is in abundant supply. Thus we rely on the relative growth rates of consumption and output across countries, and do not use the time series of trade flows (other than in a base year to calibrate the home-bias parameters). These are distinct approaches to the same objective – estimating the relative cost of consumption bundles in each countries – and hence reach a different answer. We compare our estimates to the implied price series obtain using Fitzgerald's methodology in Section xx. Despite this difference, both studies find much better risk sharing than that inferred using observed real exchange rates. This reflects the disconnect between observed real exchange rates and the model-implied shadow relative costs, which brings us to the second difference with Fitzgerald's work. We provide a decentralization of the constrained planning economy that rationalizes the gap between implied planning prices and observed prices. In particular, we show how wedges in exchange rate markets due to segmentation and in goods markets due to pricing-to-market behavior partially offset each other in equilibrium, leading to a disagreement between true risk sharing based on shadow social costs and that implied by Backus-Smith type exercises using observed real exchange rates.

The other literature we build on relates to our decentralization. In particular, the large number of papers on incomplete exchange rate pass through, pricing to market, and exchange rate volatility. Segmentation in exchange rate markets is the focus of Alvarez et al. (2002, 2009); Jeanne and Rose (2002); Gabaix and Maggiori (2015); Itskhoki and Mukhin (2021). The disconnect between local currency prices and exchange rate movements is now vast. See the Handbook chapters Burstein and Gopinath (2012) and Gopinath and Itskhoki (2021). TO BE COMPLETED

2 Environment and Planning Problem

2.1 Environment

The world consists of two regions, Home and Foreign, each of which produces a specialized good. We abstract from production and take output in each country as an exogenous stochastic process. In particular, let the state of nature at time t be denoted $s_t \in S$, where S is a discrete set, and let histories be denoted $s^t = \{s_0, s_1, \ldots, s_t\} \in \mathcal{S}^{\sqcup}$. Let $\pi_t(s^t)$ denote the probability of history s^t and $\mathbb{E}$ be the associated expectation operator. We denote output in Home (of the Home good) at time t as $Y_t(s^t)$, and Foreign output as $Y_t^*(s^t)$. In the text, we present the environment as if the two goods are freely traded, but allow for constant iceberg transportation costs in the empirical implementation. In what follows, we suppress the history notation and time subscripts whenever convenient.

Let C_H and C_F denote Home's consumption of the Home and Foreign good, respectively, while C_H^* and C_F^* denote Foreign's consumption of the same goods. Preferences at Home are given by

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t U(C_t),$$

where

$$U(C) = \frac{C^{1-\sigma}}{1-\sigma},$$

with $\sigma > 0$ and C denotes a constant-elasticity consumption aggregator

$$C = C(C_H, C_F) \equiv \left((1-\gamma)^{\frac{1}{\theta}} C_H^{\frac{\theta-1}{\theta}} + \gamma^{\frac{1}{\theta}} C_F^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}. \tag{1}$$

The parameter $\gamma \in [0,1]$ captures the extent of home bias in consumption, with $\gamma = 0$ representing complete home bias and $\gamma = 1/2$ denoting equal weights. Foreign consumer share the same β and σ as Home, but enjoy utility over a different composite good:

$$C^* = C^*(C_H^*, C_F^*) \equiv \left(\gamma^{*\frac{1}{\theta}} C_H^{*\frac{\theta-1}{\theta}} + (1-\gamma^*)^{\frac{1}{\theta}} C_F^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}. \tag{2}$$

Note that γ^* captures the weight on Foreign's imported good, and thus $\gamma^* = 0$ represents complete home bias for Foreign. For the analysis below, we use the following terminology:

Definition 1. *The environment features* Home Bias *in preferences if* $\gamma < 1/2$ *and* $\gamma^* < 1/2$.

The aggregate resource constraints on each good require

$$\begin{aligned} C_{H,t} + C^*_{H,t} &= Y_t \\ C_{F,t} + C^*_{F,t} &= Y^*_t. \end{aligned} \tag{3}$$

2.2 Efficient Allocations

Given preferences and output, we can characterize Pareto efficient allocations as sequences that maximize a weighted sum of welfare:

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \left[\omega U(C(C_{H,t}, C_{F,t})) + U(C^*(C^*_{H,t}, C^*_{F,t})) \right] \tag{4}$$

subject to the sequence of resource conditions (3). We can trace out the Pareto frontier by varying the weight $\omega \geq 0$.

As there is nothing that links resources across periods, the set of efficient allocations can be obtained by a sequence of static planning problems:

$$\max_{\{C_{H,t}, C_{F,t}, C^*_{H,t}, C^*_{F,t}\}} \omega U(C(C_{H,t}, C_{F,t})) + U(C^*(C^*_{H,t}, C^*_{F,t})) \tag{P}$$

$$\text{subject to}$$

$$\begin{aligned} C_{H,t} + C^*_{H,t} &= Y_t \\ C_{F,t} + C^*_{F,t} &= Y^*_t. \end{aligned}$$

The necessary and sufficient conditions for an optimum are (dropping t subscripts):

$$\omega U'(C) \frac{\partial C}{\partial C_k} = U'(C^*) \frac{\partial C^*}{\partial C^*_k} \quad k = H, F, \tag{5}$$

plus the two resource conditions. For each weight ω, these four equations pin down a unique efficient allocation period-by-period. Note that taking the ratio of (5) for the two goods implies:

$$\frac{C_H}{C_F} = \left(\frac{1-\gamma}{\gamma} \right) \left(\frac{1-\gamma^*}{\gamma^*} \right) \frac{C^*_H}{C^*_F}, \tag{6}$$

which implies the marginal rates of substitution across the two goods are equalized. This is optimal given that, in the absence of transportation costs, the relative cost to the planner is identical across regions (we will consider trade costs below).

We shall make use of two geometric representations of efficient and in-efficient allocations. In Figure 1, we depict in panel (a) the Edgeworth box associated with a given state (Y, Y^*). The

horizontal axis measures C_H, which ranges from $C_H = 0$ on the left to $C_H = Y$ on the right. From the resource condition, $C_H^* = Y - C_H$ can be read from right to left. Similarly, the vertical axis measures C_F. Efficient points use all resources and satisfy (5), which implies the indifference curves between Home and Foreign are tangent, a geometric representation of (6). For example, point A in the figure depicts a particular efficient allocation associated with a single weight ω. As we vary ω, we trace out the set of efficient allocations, depicted by the solid line connecting the two opposite corners.

A second geometric representation of the same problem is the "Consumption Possibility Frontier" (CPF). This is the production possibility frontier for the technologies that take the two traded inputs and produces the two final consumption composites that enter the respective utility functions. In particular, consider the following problem:

$$\begin{aligned} \mathbf{C}(C^*; Y, Y^*) \equiv \max_{\{C_H, C_F, C_H^*, C_F^*\}} & \; \mathcal{C}(C_H, C_F) \\ & \text{subject to} \\ & C_H + C_H^* = Y \\ & C_F + C_F^* = Y^* \\ & \mathcal{C}^*(C_H^*, C_F^*) \geq C^*. \end{aligned} \tag{7}$$

That is, the problem solves for the maximum consumption enjoyed by Home conditional on aggregate resources and the requirement that the planner provide a minimum consumption C^* delivered to Foreign. This problem is equivalent to (P), with each C^* mapping into a Pareto weight ω.

Letting $\tilde{\mathcal{Q}}^{FB}$ denote the multiplier on the final constraint, necessary conditions are

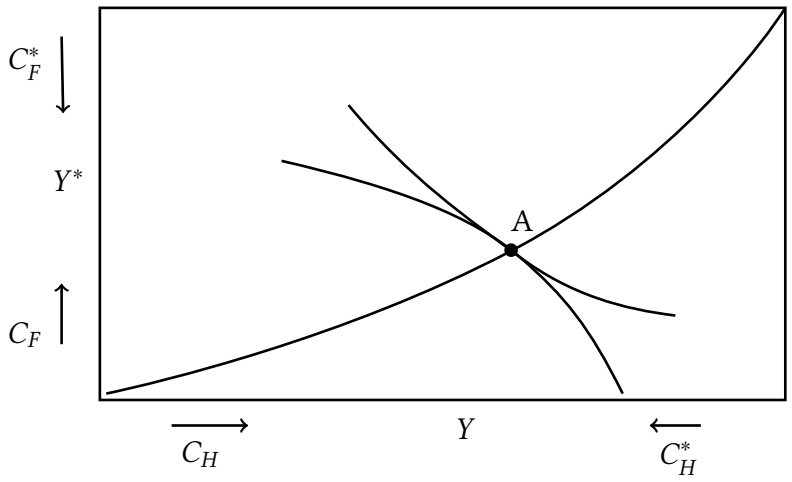
$$\frac{\partial \mathcal{C}}{\partial C_k} = \tilde{\mathcal{Q}}^{FB} \frac{\partial \mathcal{C}^*}{\partial C_k}. \tag{8}$$

Here, $\tilde{\mathcal{Q}}^{FB} = -\partial \mathbf{C} / \partial C^*$ is the shadow price of Foreign consumption in terms of Home; that is, $\tilde{\mathcal{Q}}^{FB}$ is the amount of C that is given up in order to increase C^* at the margin. Hence, $\tilde{\mathcal{Q}}^{FB}$, which is a function of the states $\{C^*, Y, Y^*\}$, represents the Marginal Rate of Transformation (MRT) between Foreign and Home consumption at an efficient allocation.

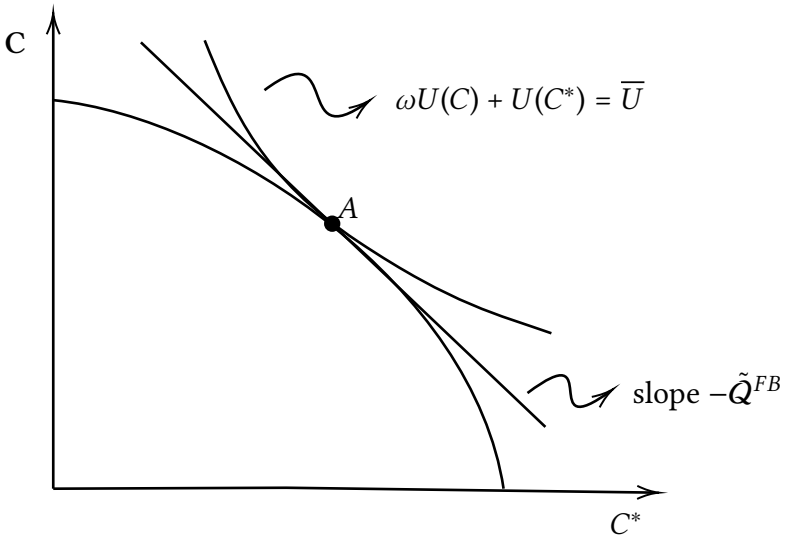
Note that the Planning allocation at any point in time can now be written as

$$\max_{C^*} \omega U(\mathbf{C}(C^*; Y, Y^*)) + U(C^*). \tag{9}$$

Figure 1: First Best Allocation



(a) Edgeworth Box



(b) Consumption Possibility Frontier

The first-order conditon is $\omega U'(\mathbf{C})\tilde{Q}^{FB} = U'(\mathbf{C}^*)$, or

$$\frac{\omega U'(C^*)}{U'(C)} = \tilde{Q}^{FB}. \tag{10}$$

The equalities state that the marginal rate of substitution (MRS) between Foreign and Home consumption equals the shadow price (MRT) of Foreign consumption relative to Home.

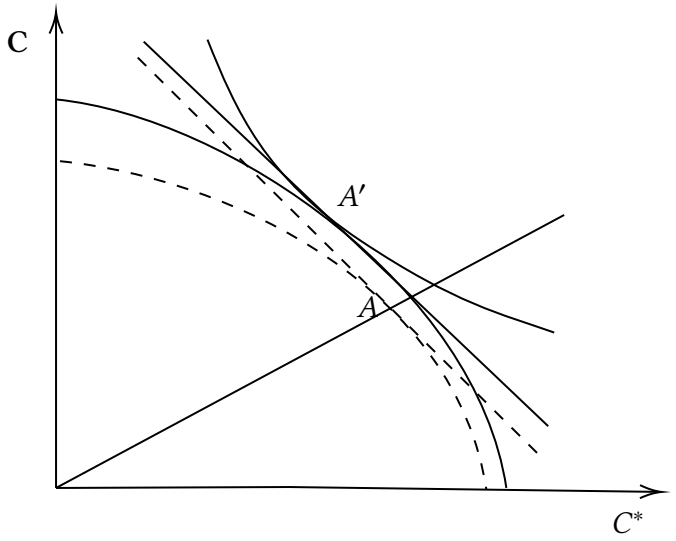
The intuition for equation (9) is that the Planner prefers to tilt consumption towards periods in which a regions's consumption composite is relatively cheap. That is, when $\tilde{Q}^{FB}$ is relatively small, it is a period in which Home's composite is relatively cheap to construct from $\{Y, Y^*\}$, and vice versa when Q^{FB} is relatively large. If the composites were identical (that is, $\gamma = 1 - \gamma^*$), then $\tilde{Q}^{FB} = 1$ in every period. In that case, the ratio of marginal utilities would be constant.

In Panel (b) of Figure 1 we plot the Consumption Possibility Frontier (CPF). Point A depicts an efficient allocation for a given ω. The slope of the frontier represents the MRT, while the slope of the depicted indifference curve for the planner represents the MRS. The figure is depicted for the case in which the input intensities of the two consumption baskets differ (i.e., $\gamma \neq 1 - \gamma^*$), which implies the frontier is strictly concave, a property demonstrated and discussed in Samuelson (1949). The shadow cost of providing Foreign with more consumption is increasing the greater is Foreign's share of total resources.

The intuition for the concavity of the frontier is well known (and provided by Samuelson), but worth reviewing given what will follow. One feasible allocation is to set $C_H/C_F = C_H^*/C_F^* = Y/Y^*$, so that each region consumes the goods in the same ratio (which is the same as the ratio of the endowments). With this rule, the split between C and C^* is achieved by varying the (common) fraction of each endowment given to Home and Foreign. This will trace out a linear frontier with slope minus one. If the regions prefer the goods in different intensities, a strictly superior sharing rule will be to allow them to consume in different proportions; namely, those given by (6). This pushes the frontier to the northeast of a line, generating a CPF that is concave to the origin.

It will be useful to consider how the first-best allocation is altered by changes in the relative endowment. In particular, consider an increase in Y, holding Y^* constant. Under home bias, this makes the good that Home uses intensively relative abundant and the good the Foreign uses intensively relatively scarce. This raises the relative cost of producing C^* in terms of the Home consumption commodity. In particular, we show in the appendix that at a fixed ratio C/C^*, $\tilde{Q}^{FB}$ is increasing in Y and decreasing in Y^*, all else equal. We depict this case in Figure 2, where the original allocation is at point A along the dashed consumption possibility frontier. The ray from the origin through point A depict allocations that keep the ratio of C/C^* constant. The new CPF is the solid line and is steeper than A at the intersection of the ray. As preferences do not change with endowments, to achieve a tangency with the planner's indifference curve along the new CPF,

Figure 2: Response to a Change in Y/Y^*



Note: This figure depicts a change in Y/Y^*. The original CPF is the dashed concave line intersecting point A, which denotes the original allocation. The dashed line tangent to A represents the MRT. The new CPF is the solid line that is concave to the origin. As depicted, Y has increased, allowing for greater consumption of both composites. The new optimal allocation is A'.

the first-best allocation (for given Pareto weight ω) increases the ratio of Home consumption to Foreign. The new tangency is point A'. This comparative static is summarized in the following lemma:

Lemma 1. *For a given Pareto weight ω, under home bias the first-best C/C^* is increasing in Y/Y^*.*

Thus, first-best risk-sharing under home bias implies Home's relative consumption is positively correlated with its relative endowment.

2.3 Static and Dynamic Distortions

If the Welfare Theorems hold, then $\tilde{Q}^{FB}$ equals the competitive equilibrium's real exchange rate. This implication was used by Backus and Smith (1993) to test efficient risk sharing by relating observed consumption to the empirical real exchange rate, which we denote Q. One of our goals is to explore to what extent the failure of their test is due to the failure of the Welfare Theorems.

We do so by considering distortions along two dimensions. The first is a "static" wedge that distorts the relative consumption of Home and Foreign goods within a period. The second is

a "dynamic" wedge that distorts the inter-temporal allocation of consumption. In Section 3 we introduce a decentralized economy in which these distortions will arise from particular frictions in goods and asset markets.

The Static Wedge: In an efficient allocation, condition (6) states that the marginal rate of substitution between Home and Foreign goods in each country should be equated. The static wedge breaks this equality. Namely, letting δ denote the static wedge, we have

$$\begin{aligned} \frac{\partial C/\partial C_H}{\partial C/\partial C_F} &= (1+\delta)\frac{\partial C^*/\partial C_H^*}{\partial C^*/\partial C_F^*} \\ &\Rightarrow \\ \frac{C_H}{C_F} &= (1+\delta)^{\theta}\left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{C_H^*}{C_F^*}. \end{aligned} \tag{11}$$

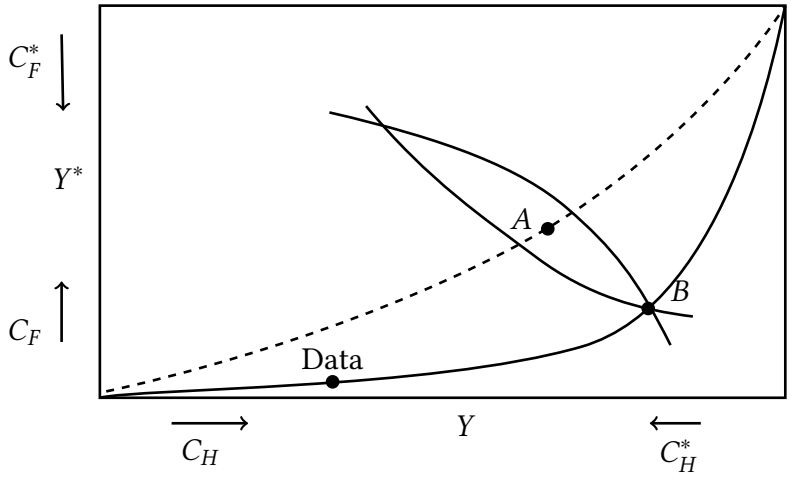
If $\delta \neq 0$, there are static gains from trade. In particular, Home is willing to sell $C/\partial C_H$ units of Home good in exchange for a unit of the Foreign good. Conversely, Foreign is willing to buy Home goods at a relative price $\partial C/\partial C_F$. If $\delta > 0$, there are strict gains from such a trade (and vice versa for $\delta < 0$), implying the presence of a feasible Pareto improvement. Thus δ represents an unrealized gain from a static exchange of Home for Foreign goods. In Section 3, we show how various familiar models of Pricing-to-Market, with and without nominal rigidities, map into δ in a decentralized equilibrium. In particular, δ will be generated by different relative prices faced by consumers in each region. As noted above, in our empirical work, we do allow for trade costs, which will also generate a wedge between C_H/C_F and C_H^*/C_F^*. However, we will treat such costs as stable over time and interpret time variation in δ as reflecting pricing-to-market frictions.

In Figure 3 Panel (a), we show how the static wedge distorts the allocation within the Edgeworth box of Figure 1. Recall that point A is on the locus of efficient allocations, which is now depicted by a dotted line. Point B is distorted so that Home consumes disproportionately more Home good relative to Foreign. At B, Home's MRS between Home and Foreign goods differs from Foreign's by a factor $1+\delta$. As we vary the Pareto weight ω, we can trace out a distorted contract curve, along which the MRS distortion is held constant at $1+\delta$. This is the solid locus in the Edgeworth box.

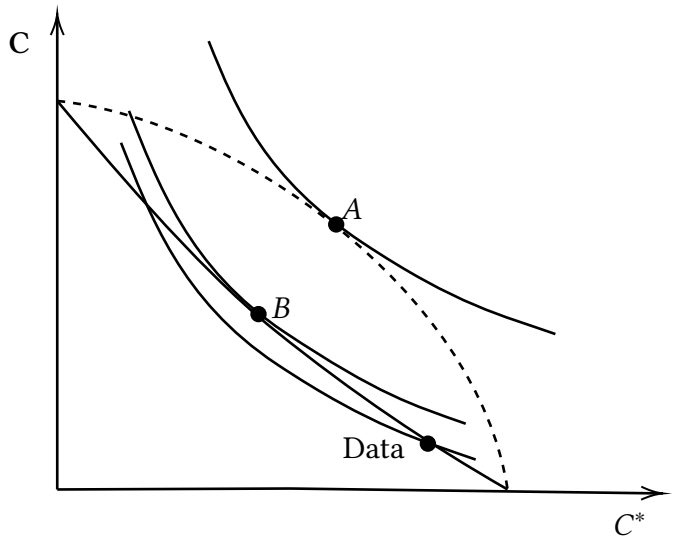
Panel (b) of Figure 3 depicts how the Consumption Possibility Frontier is distorted by the static wedge. The dashed line is the first-best frontier depicted in Figure 1. The solid line is the distorted frontier. This locus solves Problem (7) with the additional constraint (11).[2] We can

[2]While useful as a conceptual link to the first best, note that the additional constraint associated with the static wedge reduces the constraint set of (7) to a single allocation when the C^* constraint is evaluated as an equality. This is clear in the Edgeworth box. The promised C^* to Foreign pins down the indifference curve of Foreign, and then the δ requirement pins down the single point on that indifference curve that is consistent with aggregate resources.

Figure 3: Two Distortions



(a) Static Distortion



(b) Distorted CPF

extend our notation to write $\mathbf{C}(C^*; Y, Y^*, \delta)$ to denote the maximum consumption of Home given the aggregate resources and the static wedge, and similarly $\tilde{Q}(C^*; Y, Y^*, \delta) = \partial \mathbf{C}/\partial C^*$ to denote the MRT at the frontier. Note that $\tilde{Q}^{FB} = \tilde{Q}(C^*; Y, Y^*, \delta = 0)$.

At the end points of the Consumption Possibility Frontiers, the two frontiers coincide.[3] This reflects that in both cases, one country's consumption goes to zero and the other goes to consuming both endowments. Away from these points, as long as the static wedge is non-zero, the second-best frontier lies interior to the first best. Moreover, it need not be concave. Recall that the intuition for concavity of the first-best CPF exploited the notion that regions use the two goods with different intensities. With the static distortion, the different intensities are not optimized, and may be worse than (interior to) the equal ratio benchmark that would deliver a linear frontier.

We can consider allocations that solve Problem (P) with the additional constraint (11). The second best allocation obtained in this manner maximizes the weighted sum of utility subject to the static market distortion. This allocation is depicted as point B in panel (b) of Figure 3. At B, we have

$$\frac{U'(C^*)}{\omega U'(C)} = \tilde{Q}, \tag{12}$$

where $\tilde{Q}(C^*; Y, Y^*, \delta)$ is the multiplier on the C constraint in the distorted version of Problem (7). $\tilde{Q}$ represents the distorted Marginal Rate of Transformation and is the slope of the interior frontier in Panel (b) of Figure 3. Note that for general δ, the distorted frontier need not be concave, and hence $\tilde{Q}$ can be a non-monotonic function of C^*. Condition (12) is a necessary condition for an efficient allocation that is constrained by the static wedge, something that will motivate our empirical implementation.

The Dynamic Wedge: The second wedge we consider is the inter-temporal distortion, which is the deviation from condition (12):

$$1 + \lambda_t \equiv \frac{U'(C_t^*)}{\omega U'(C_t)} \frac{1}{\tilde{Q}_t}. \tag{13}$$

Comparing this expression with (5), the inter-temporal wedge reflects an implicit shift in the Pareto weight that rationalizes the observed allocation. This implies that if $\lambda_t \neq \lambda_{t_0}$, there is an inter-temporal reallocation between (C_{t_0}, C_t) and $(C^*_{t_0}, C^*_t)$ that represents a Pareto improvement.[4]

Thus, the constraint set is essentially a singleton.

[3] To see this, as we approach the horizontal intercept, we can achieve (in the limit) the efficient allocation $C \to C(Y, Y^*)$ and $C^* \to C^*(0, 0) = 0$, by letting C^*_H and C^*_F both approach zero along a path such that their ratio continues to satisfy (11) when $C_H/C_F \to Y/Y^*$. The same logic applies at the vertical intercept.

[4] In particular, if $\lambda_t > \lambda_{t_0}$, consider increasing C^*_t at the margin by $\Delta C^* > 0$ and decreasing $C^*_{t_0}$ by $-\beta^{t-t_0} U'(C^*_t)/U'(C^*_{t_0}) * \Delta C^*$, such that Foreign's welfare is unchanged to the first order. Feasibility implies that

A useful summary of the above can be provided by the following "distorted" planning problem:

$$\max_{\{C_{H,t},C_{F,t},C^*_{H,t},C^*_{F,t}\}} \omega U(C(C_{H,t},C_{F,t})) + U(C^*(C^*_{H,t},C^*_{F,t})) \tag{P'}$$

subject to

$$C_{H,t} + C^*_{H,t} = Y_t$$
$$C_{F,t} + C^*_{F,t} = Y^*_t$$
$$\frac{C_{H,t}}{C_{F,t}} = (1+\delta)^\theta \left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{C^*_{H,t}}{C^*_{F,t}}$$
$$\omega U(C(C_{H,t},C_{F,t})) = \overline{U}_t.$$

The final two constraints reflect our two distortions. The penultimate constraint is the static distortion. The last constraint is the dynamic distortion, and forces the planner to provide a particular level of income to Home independent of ω. The link to the earlier discussion is as follows: if we let λ_t denote the multiplier on the last constraint, then, after combining terms, the Lagrangian contains $\omega\lambda_t U$, rather than ωU as the objective for splitting the pie between Home and Foreign, which represents the dynamic distortion to risk sharing.

2.4 Identification

With these conceptual notions in hand, we can turn to identification of δ, $\tilde{\lambda}$, and $\tilde{Q}$ from observation on real consumption and output in two regions. We postpone to Section 5 how we aggregate the world into two regions as well as other measurement challenges, like the presence of investment and government expenditures. We also postpone how we calibrate the preference parameters σ, γ and γ^*. Because real consumption and output are indices, we will eventually work in growth rates, but for the current discussion we assume we observe $\{C_t, C^*_t, Y_t, Y^*_t\}$ for some period $t = 0, 1, \ldots, T$.

In a given period, we would like to recover $\{C_H, C_F, C^*_H, C^*_F\}$ from observables. We have four equations – the two resource conditions (3), as well as the two aggregations of the Home and Foreign goods into the respective consumption composites given in equations (1) and (2). These are four equations in four unknowns. But these are nonlinear equations, so it is not clear yet how many potential solutions there are, if any. The Edgeworth box is a useful tool to address this question. Given observable $\{Y, Y^*\}$ we have the dimensions of the box, and given $\{C, C^*\}$ we have two indifference curves. If these curves are tangent, then $\delta = 0$ and there is a unique

$\Delta C_t = -\tilde{Q}_t \Delta C^*$ and $\Delta C_{t_0} = Q_{t_0}\beta^{t-t_0}U'(C^*_t)/U'(C^*_{t_0}) * \Delta C^*$. Using the definition of λ_t, the net increase in Home's welfare is proportional to $\lambda_t - \lambda_{t_0} > 0$.

decomposition into $\{C_H, C_F, C_H^*, C_F^*\}$. If they intersect twice, which is the most they can intersect, there are two possible decompositions, one with $\delta < 0$ and one with $\delta > 0$. We assume the static distortion enhances home bias, which is associated with $\delta > 0$. Given this identifying assumption, we select the decomposition that lies "below" the efficient contract curve. Finally, if the two indifference curves never intersect, then we have violated feasibility. This is evidence of a mis-specification of preferences and/or a mis-measurement of the observed aggregates. In Section 5, we discuss if and when such pathologies arise.

In Figure 3 Panel (a), we depict a point labelled "Data." This is the solution to the four equations mentioned in the previous paragraph, and pins down the static wedge, δ. Relative to point B, the Data has relatively more consumption to Foreign and less to Home, which we discuss next.

In Panel (b) of Figure 3, we also indicate the observed data $\{C, C^*\}$. As we use the same data points in computing the empirical δ, the distorted CPF contains this point by construction. Relative to B, the Data is biased toward Foreign consumption. This reflects the dynamic wedge.

In particular, given $\{C, C^*, Y, Y^*\}$ plus our measured δ, we can turn to the version of problem (7) with the addition of constraint (11) to infer $\tilde{Q}$. Recall that $\tilde{Q}$ is the multiplier on the $C^*(C_H, C_F) \geq C^*$ constraint. Given the allocation, this can be read directly off the Lagrangian's first-order conditions.[5]

Repeating this each period, we obtain a sequence $\{\tilde{Q}_t\}_{t=0}^T$. From (13), and the sequences for real consumption in both regions, we can obtain the associated path of the dynamics wedge (with the normalization $\lambda_0 = 0$ for some reference year $t = 0$). In Figure 3 panel (b), the indifference curves are drawn for the Pareto weight that matches the data at $t = 0$ assuming $\lambda_0 = 0$. In the Figure, the difference in the slope of the indifference curve and the slope of the distorted frontier at the point "Data" represents that period's dynamic wedge λ.

3 Decentralized Equilibrium

In this section, we address two issues regarding a decentralized equilibrium and how it relates to the planning problem. The first is how to map the observed empirical allocation to a sequence of prices and wedges that support the allocation as an equilibrium. This is essentially an accounting exercise, similar to the measurement of the wedges between planning allocation and the observed allocation. The second task is to understand what type of environments endogenously generate the price and wedge behavior observed in the data. In particular, what types of market and pricing frictions discussed in the literature best fit the observed allocation.

[5] More specifically, using the resource conditions to substitute $C_H^* = Y - C_H$ and $C_F^* = Y^* - C_F$, we are left with two choices, $\{C_H, C_F\}$ and two constraints – the $C \geq C$ constraint and the static wedge constraint (11). Given an allocation, the two first-order conditions give two linear equations in the two multipliers, which can be solved for the multiplier of interest, $\tilde{Q}$.

3.1 Household Behavior

Consider the representative Home household's problem. At history s^t, it faces prices $P_{H,t}(s^t)$ and $P_{F,t}(s^t)$ for the Home and Foreign goods, respectively. The ideal price index for the Home composite is therefore (suppressing the history indexing)

$$P = \left((1-\gamma)P_H^{1-\theta} + \gamma P_F^{1-\theta}\right)^{1/(1-\theta)}.$$

For the initial exposition, we write the problem as if the Home household has access to a full set of Arrow-Debreu securities. While the complete-markets notation is more convenient for the exposition, we show below how to map the wedges we recover into more realistic asset structures, such as incomplete and segmented asset markets. Let $\Lambda_t(s^t)$ be the price of an Arrow-Debreau security in Home that pays one unit of the Home's currency in history s^t. For some initial asset position w, the representative Home agent's problem is:

$$\begin{aligned}
&\max_{\{C_t(s^t)\}} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t(s^t) U(C_t(s^t)) \\
&\text{subject to: } w \geq \sum_{t} \sum_{s^t} \Lambda_t(s^t) \left(P_t(s^t)C_t(s^t) - P_{H,t}(s^t)Y_t(s^t) - T_t(s^t)\right),
\end{aligned}$$

where $T_t(s^t)$ is an income stream to be discussed below.

The first-order conditions for the Home representative agent imply (again, suppressing the history notation when convenient):

$$\frac{C_H}{C} = (1-\gamma)\left(\frac{P_H}{P}\right)^{-\theta}, \quad \frac{C_F}{C} = \gamma\left(\frac{P_F}{P}\right)^{-\theta}, \tag{14}$$

and

$$\beta^t \pi_t(s^t) U'(C_t(s^t)) = \xi \Lambda_t(s^t) P_t(s^t), \tag{15}$$

where ξ is the Lagrange multiplier on the budget constraint.

Given a sequence of observed allocations $\{C_{H,t}, C_{F.t}, C_t\}$, we can infer the sequence of prices $\{P_{H,t}, P_{F,t}, P_t\}$ up to a constant of normalization (for example, we can set $P_{H,t}(s^t) = 1$ at every history). Note that we do this only for the realized history associated with the oberved data. Similarly, we can infer the sequence of implied Arrow-Debreu prices $\Lambda_t(s^t)$ up to a constant of normalization (associated with ξ and similar to ω in the planning problem). Again, these prices are restricted to the realized history.

The Foreign agent faces prices $\{P^*, P_H^*, P_F^*\}$ for the Foreign composite, Home good, and Foreign

good, respectively, with

$$P^* = \left(\gamma^* P_H^{*1-\theta} + (1-\gamma^*)P_F^{*1-\theta}\right)^{1/(1-\theta)}.$$

These prices are denominated in units of Foreign's currency. Let $\mathcal{E}$ denote the nominal exchange rate expressed as Foreign currency per unit of Home currency. At this stage, we can abstract from nominal exchange rates and assume $\mathcal{E} = 1$ in all histories, implying that all prices in both Home and Foreign are expressed in a common numeraire. We suppress $\mathcal{E}$ in the notation for now, and shall re-introduce a nominal exchange rate when we discuss pricing protocols in Section 3.4.

Given some initial wealth w^*, transfers $T_t^*(s^t)$, and prices of Arrow-Debreau securities $\Lambda_t^*(s^t)$ in Foreign (that pay out in the same numeraire as Λ), the Foreign household's problem is:

$$\max_{\{C_t^*(s^t)\}} \sum_{t=0}^{\infty} \sum_{s^t} \pi_t(s^t) U(C_t^*(s^t))$$

subject to:

$$w^* \geq \sum_t \sum_{s^t} \Lambda_t^*(s^t) \left(P_t^*(s^t)C_t^*(s^t) - P_{F,t}^*(s^t)Y_t^*(s^t) - T_t^*(s^t)\right).$$

The first-order conditions for the Foreign representative agent imply:

$$\frac{C_H^*}{C^*} = \gamma^* \left(\frac{P_H^*}{P^*}\right)^{-\theta}, \quad \frac{C_F^*}{C^*} = (1-\gamma^*)\left(\frac{P_F^*}{P^*}\right)^{-\theta}, \tag{16}$$

and

$$\beta^t \pi_t(s^t) U'(C_t^*(s^t)) = \xi^* \Lambda_t^*(s^t) P_t(s^t), \tag{17}$$

where ξ^* is the Lagrange multiplier on Foreign's budget constraint. Mirroring Home, we can infer the sequence of relative prices $\{P_{H,t}^*, P_{F,t}^*, P_t^*\}$ and the sequence of Arrow-Debreau prices $\Lambda_t^*(s^t)$ up to constants of normalization for goods and security prices.

Given the observed real exchange rate, Q, we can pin down the ratio of Home to Foreign price levels. In particular, we have

$$Q \equiv \frac{P^*}{P} = \frac{P_H^*}{P_H} \left(\frac{\gamma^* + (1-\gamma^*)\left(\frac{P_F^*}{P_H^*}\right)^{1-\theta}}{(1-\gamma) + \gamma\left(\frac{P_F}{P_H}\right)^{1-\theta}} \right)^{1/(1-\theta)}, \tag{18}$$

which allows us to measure P_H^*/P_H given the relative prices in each country obtained from the allocations.

3.2 Implied Wedges

The ratio P_H^*/P_H is the deviation in the Law of One Price (LOP) for the Home good, and represents the relative price of the same good in the two locations. Given this, and the relative prices within each market, we can also compute the LOP deviation for Foreign: P_F/P_F^*. We define the wedges μ_H and μ_F as the premium (if greater than zero, and discount otherwise) paid by the Home agent relative to a consumer in Foreign for the Home and Foreign good, respectively:

$$P_H = (1 + \mu_H)P_H^* \quad P_F = (1 + \mu_F)P_F^*. \tag{19}$$

From (15) and (17), we have

$$\frac{\xi}{\xi^*}\frac{U'(C_t^*(s^t))}{U'(C_t(s^t))} = \frac{\Lambda_t^*(s^t)}{\Lambda_t(s^t)}\frac{P_t^*(s^t)}{P_t(s^t)} = \frac{\Lambda_t^*(s^t)}{\Lambda_t(s^t)}Q_t(s^t).$$

We define the "Backus-Smith" wedge as:

$$1 + \psi_t(s^t) \equiv \frac{\Lambda_t^*(s^t)}{\Lambda_t(s^t)} = \frac{\xi}{\xi^*}\frac{U'(C_t^*(s^t))}{U'(C_t(s^t))}\frac{1}{Q_t(s^t)}. \tag{20}$$

This is the wedge in private risk-sharing. The first expression in (20) defines ψ as the deviation across the two regions in the period-0 price of a dollar delivered in a certain history. The final term uses the households' first order conditions to relate ψ the wedge that was measured by Backus and Smith (1993); namely, the relative marginal utility of consumption in Foreign versus Home compared to the relative equilibrium price of the two composite goods (scaled by a constant ratio reflecting initial wealth).

The ratio ξ^*/ξ reflects the relative wealth of Home to Foreign. This is the equilibrium counterpart to the Planner's Pareto weight on Home, ω. As these will play no role in our measurement or counter-factual exercises, it is without loss for our purposes to normalize $\omega = \xi^*/\xi$ in the expressions that follow to simplify notation.

The following lemma shows how the equilibrium wedges relate to the planner's distortions discussed in the previous section:

Lemma 2. *(i) The static wedge defined in* (11) *is related to deviations in the Law of One Price in* (19) *by:*

$$1 + \delta = \frac{1 + \mu_F}{1 + \mu_H}.$$

(ii) The dynamic wedge defined in (13) is related to the Backus-Smith wedge of (20) by

$$\frac{1+\lambda_t}{1+\psi_t} = \frac{Q_t}{\tilde{Q}_t}.$$

Proof. By definition,

$$\frac{P_H}{P_F} = \frac{P_H^*}{P_F^*}\frac{1+\mu_F}{1+\mu_H} \Rightarrow \frac{C_H}{C_F} = \left[\frac{1+\mu_F}{1+\mu_H}\right]^{\theta}\left(\frac{1-\gamma}{\gamma}\frac{1-\gamma^*}{\gamma^*}\right)\frac{C_H^*}{C_F^*},$$

where the right-hand side expression is obtained using (14) and (15). Comparing this expression with (11), we have the first result. The second result is obtained by taking the ratio of (13) to (20) and using our assumed normalization $\omega = \xi^*/\xi$. □

Note that the single distortion in the Planner's problem, δ, is mapped into two distortions in the goods market. This is a consequence that we use an additional observable, the real exchange rate Q. Fixing P_H/P_F and P_H^*/P_F^*, doubling both $1+\mu_H$ and $1+\mu_F$ is consistent with the same δ, but from (18) we see that this would also halve the observed real exchange rate. However, for allocations rather than prices, δ is sufficient to characterize the goods market distortion.

The second result in the lemma establishes that the Planner's dynamic wedge and the Backus-Smith wedge are equivalent to the difference between the technological marginal rate of transformation between the two consumption composites and the equilbrium real exchange rate. If the Welfare Theorems hold, these are the same. To the extent the Welfare Theorems fail in practice, the two wedges will be distinct.

From Part (i) of Lemma 2, we see that the static distortion disappears if the deviation from the law of one price is equal across the two goods. This is the case when households in both Home and Foreign face the same relative prices (although the levels of each price may differ). In this case, we have

Lemma 3. *If $\mu_H = \mu_F = \mu$, then $\delta = 0$ and*

$$\frac{1+\lambda}{1+\psi} = \frac{Q}{\tilde{Q}} = 1+\mu.$$

Proof. TBA □

This lemma states that when relative prices are equated across regions, the difference between the planner's dynamic wedge and the Backus-Smith wedge (and hence, the difference between the MRT and the real exchange rate) is equal to the deviations from the law of one price. We immediately have the following corollary:

Corollary 1. *If $\mu_H = \mu_F = \mu$, we have:* (a) *when LOP holds ($\mu = 0$), then $\lambda = \psi$;* (b) *when the Backus-Smith condition holds ($\psi = 0$), then $\lambda = \mu$;* (c) *and when $1 + \mu = \frac{1}{1+\psi}$, then $\lambda = 0$.*

These are useful special cases, as they nest many familiar models, which we discuss below. They also illustrate that there is no one-to-one mapping between ψ and λ. First, we can have $\lambda = \psi$ when LOP deviations are absent. Second, even when the Backus-Smith condition holds perfectly ($\psi = 0$), we can have $\lambda = \mu \neq 0$. Third, for any Backus-Smith wedge ψ, we can have $\lambda = 0$ if μ offsets ψ.

3.3 Financial Markets

In this subsection, we briefly discuss alternative interpretations of ψ. The most straightforward interpretation is that there are complete markets subject to region-specific distortionary taxes (which are lump-sum rebated). Thus, Home and Foreign households face different after-tax prices on the Arrow-Debreu securities, driving a wedge between $\Lambda_t(s^t)$ and $\Lambda^*_t(s^t)$. A similar interpretation is obtained if we introduce segmented markets with intermediation frictions. For example, intermediaries charging mark-ups on asset trades between Home and Foreign can rationalize the different prices faced by Home and Foreign for the same security. A polar opposite case is financial autarky but fiscal transfers between the regions in the spirit of Farhi and Werning (2017). In this case, the wedge ψ will reflect the extent to wich such transfers are not consistent with the Backus-Smith condition for private risk sharing. Finally, ψ can be interpreted as the difference in marginal utilities that arises from incomplete markets. For example, if only a single risk-free bond paying out in the common numeraire (dollars) was traded, the associated Euler Equations in the two regions would imply for each period t:

$$\mathbb{E}_t\left[\frac{U'(C_{t+1})}{U'(C_t)}\frac{P_t}{P_{t+1}}\right] = E_t\left[\frac{U'(C^*_{t+1})}{U'(C^*_t)}\frac{P^*_t}{P^*_{t+1}}\right],$$

where we have cancelled β^t and the interest rate from both sides. In this environment, the sequence $\{\psi_t\}$ represents the net differences between realized and expected marginal utilities across the two regions.

Regardless of the interpretation, we can define the income accounted for by financial market frictions as

$$\begin{aligned}\Pi_X &\equiv \sum_t \sum_{s^t} \left(\Lambda^*_t(s^t) - \Lambda_t(s^t)\right)\left(P_t(s^t)C_t(s^t) - P_{H,t}(s^t)Y_t(s^t) - T_t(s^t)\right)\\ &= \sum_t \sum_{s^t} \psi(s^t)\Lambda_t(s^t)\left(P_t(s^t)C_t(s^t) - P_{H,t}(s^t)Y_t(s^t) - T_t(s^t)\right).\end{aligned}$$

That is, Π_X is the extra (or less) Home pays to finance its current account relative to what Foreign pays evaluated at a common set of Arrow-Debreu prices. This implies that the value of Home's current account flows evaluated at Foreign's prices is:

$$\begin{aligned}\overline{w} \equiv w + \Pi_X &= \sum_t \sum_{s^t} (1 + \psi(s^t))\Lambda_t(s^t) \left(P_t(s^t)C_t(s^t) - P_{H,t}(s^t)Y_t(s^t) - T_t(s^t)\right) \\ &= \sum_t \sum_{s^t} \Lambda_t^*(s^t) \left(P_t(s^t)C_t(s^t) - P_{H,t}(s^t)Y_t(s^t) - T_t(s^t)\right). \end{aligned} \tag{21}$$

Hence, Π_X is the difference in the value of net exports evaluated at Foreign's versus Home's stochastic discount factor.

The market clearing condition for financial markets can be written

$$\overline{w} + w^* = w + w^* + \Pi_X = 0.$$

That is, the present value of global current accounts, evaluated at common Arrow-Debreu prices (in this case Foreign's), must sum to zero.

3.4 Pricing Protocols

In this subsection, we discuss alternative models of deviations from the law of one price. To do so, we re-introduce the nominal exchange. Recall that Home consumers face prices $\{P_H, P_F\}$ denominated in Home's currency and Foreign consumers face prices $\{P_H^*, P_F^*\}$ in their currency. With $\mathcal{E}$ denoting the nominal exchange rate in terms of Foreign currency per unit of Home currency, we have $\mathcal{E}P_H = (1 + \mu_H)P_H^*$ and $\mathcal{E}P_F = (1 + \mu_F)P_H^*$.

A simple benchmark is a producer sells its product for a common price in its own currency across all destinations. This is consistent with a constant mark-up across destinations in a flexible price model and with what is known as "producer currency pricing" (PCP) in the Open Economy New Keynesian literate (see, for example,). Under PCP, a producer sets the price of their product in their own currency regardless of the destination market. The price that consumers face for the good in the destination market will then be proportional to the nominal exchange rate between the producer's and consumer's currencies. That is, $P_H^* = P_H/\mathcal{E}$ and $P_F = \mathcal{E}P_F^*$. Regardless of the nominal exchange rate realization, the price that consumers face in different markets will then be identical once converted into a common numeraire. This implies that under PCP, $\mu_H = \mu_F = 0$.

An alternative benchmark is that producers set a destination-specific price for each market. In particular, producers fix a price in each market in units of the consumers' currency, and this price is "sticky" in the sense it does not adjust to the realized value of the nominal exchange. This is referred to as "local currency pricing" (LCP) in the Open Economy New Keynesian literate (see, for

example,). Under this pricing protocol, the law of one price deviation will be determined by the nominal exchange rate. In paricular, suppose that Home sets prices P_H in the Home market and sets the Foreign-currency price P_H^* in the Foreign market. Consistent with much of the literature, suppose $P_H^* = \mathbb{E}[\mathcal{E}]P_H$, so that the expected Home-currency price for the Home good in Foreign equals the price in Home, where expectation is taken with respect to the information set at the time of price setting. Similarly, Foreign sets P_F^* in its own market and set's the Home-currency price P_F in the Home market, with $P_F = P_F^*/\mathbb{E}[\mathcal{E}]$. All such prices are fixed relative to nominal exchange rate realizations. Under this protocol, we have $1 + \mu_H = 1 + \mu_F = \frac{\mathcal{E}}{\mathbb{E}\mathcal{E}}$. Note that the relative prices in each market are the same: $P_H/P_F = P_H^*/P_F^*$.

Collecting these results and using Lemmas 2 and 3, we have the following pricing protocol taxonomy:

Proposition 1. *Under PCP, we have $\delta = 0$ and $\lambda = \psi$. Under LCP, we have $\delta = 0$ and $\frac{1+\lambda}{1+\psi} = \frac{Q}{\bar{Q}} = \frac{\mathcal{E}}{\mathbb{E}[\mathcal{E}]}$.*

Similar to the accounting done for the financial market wedges, we define (suppressing the state notation and setting $\mathcal{E} = 1$):

$$\Pi_H \equiv (P_H^* - P_H)C_H^* = -\mu^* P_H^* C_H^*$$
$$\Pi_F \equiv (P_F - P_F^*)C_F = \mu P_F^* C_F,$$

as the additional profit or loss associated with the goods-market LOP deviations. Market clearing for Home goods requires:

$$\begin{aligned} P_H(C_H + C_H^*) = P_H Y \quad &\Rightarrow \\ P_H C_H + P_H^* C_H^* + \mu_H C_H^* = P_H Y \quad &\Rightarrow \\ P_H C_H + P_H^* C_H^* = P_H Y + \Pi_H, \end{aligned}$$

where the second line uses $P_H = (1 + \mu_H)P_H^*$ and the last line follows from the definition of Π_H. Similarly, market clearing for Foreign goods requires:

$$P_F C_F + P_F^* C_F^* = P_F^* Y^* + \Pi_F.$$

In a given period, Home spends PC and earns $P_H Y + T$. The difference is its current account deficit. From Balance of Payments accounting, this must equal the negative of Foreign's current

account deficit. That is:

$$PC - P_H Y - T = -\left(P_F^* C_F^* - P_F^* Y^* - T^*\right).$$

By definition, $PC = P_H C_H + P_F C_F$ and similarly for Foreign. Substituting this, and the fact the resource conditions for Home and Foreign goods, we have:

$$T + T^* = \Pi_H + \Pi_F.$$

Hence, the transfers are how the income generated by LOP deviations is distributed across the two regions.[6]

3.5 Equilibrium Concept

Our notion of a decentralized equilibrium takes the sequence of realized wedges $\{\mu_H, \mu_F, \psi\}$ as primitives. That is, while motivated by the examples of market frictions discussed above, we abstract from the exact micro-foundations of these wedges. This is not relevant for measurement, but is a restriction imposed when computing counter-factual allocations. We view this as a useful approach to compute how each type of wedge affects the allocations and welfare. It should be stressed, however, that we are not claiming that, for example, removing the Backus-Smith wedge in practice would leave pricing wedges unchanged. Indeed, the empmirical results suggest that these wedges are tightly inter-related. Nevertheless, using such counter-factuals provide an insight into the sources and magnitudes of deviations from the first best.

Given this assumption, we can also abstract from the nominal exchange rate and implicitly set $\mathcal{E} = 1$ at all histories. We also abstract from how the implicit income generated from these wages are distributed and treat the income sharing rules $\{T, T^*\}$ as primitives. This is without loss, as the transfers are essentially components of the initial relative wealths. Conditional on the wedges and the process for endowments, we can define equilibrium in the usual way:

Definition 2. *Conditional on sequences of endowments* $\{Y_t(s^t), Y_t^*(s^t)\}$ *and wedges* $\{\mu_H(s^t), \mu_F(s^t), \psi_t(s^t)\}$ *for* $t \geq 0, s^t \in S^t$, *an equilibrium consists of allocations* $\{C_{H,t}(s^t), C_{F,t}(s^t), C_t(s^t)\}$ *and* $\{C_{H,t}^*(s^t), C_{F,t}^*(s^t), C_t^*(s^t)\}$, *as well as prices* $\{P_{H,t}(s^t), P_{F,t}(s^t), P_t(s^t), \Lambda(s^t)\}$, $\{P_{H,t}^*(s^t), P_{F,t}^*(s^t), P_t^*(s^t), \Lambda^*(s^t)\}$, *and* $Q_t(s^t)$ *such that:*

(i) Households optimize given prices and initial wealth positions $\{\overline{w}, s^*\}$ *and an income*

[6] Note that without loss we can introduce claims on these income flows and embed them in the initial wealth positions. That is, what matters for the household decision is the sum of initial wealth plus the present value of the transfers.

transfer rule $\{T_t(s^t), T_t^*(s^t)\}$;

(ii) Financial and goods markets clear;

(iii) Prices and wedges are consistent with definitions given by (19) *and* (20).

We can frame our measurement exercise as observing an allocation generated by such an equilibrium and our task is to identify and recover the underlying wedges, both for the equilibrium and the associated wedges for the Planner's problem. With these realized wedges in hand, we can consider alternative sequences of wedges to explore counter-factual equilibria.

4 Welfare

4.1 Social welfare costs (λ)

We consider the welfare of the Home representative agent along a given a given history s^t, $t \geq 0$. In particular, let $C_t = C_t(s^t)$ and similarly for other variable. We can use our definition of the consumption possibility frontier to define

$$\mathcal{U}(\{C_t^*, \delta_t\}) \equiv \sum_{t=0}^{\infty} \beta^t U(\boldsymbol{C}(C_t^*, \delta_t)).$$

Recall that from the definition of λ_t, we have

$$(1 + \lambda_t)\omega U'(\boldsymbol{C}(C_t^*, \delta_t))\boldsymbol{Q}_t = U'(C_t^*),$$

and, given a static wedge δ_t and Pareto weight ω, there is a one-to-one mapping between λ_t and C_t^*.[7] With a slight abuse of notation, we therefore will also express Home welfare equivalently as a function of our two wedges $\mathcal{U}(\{\lambda_t, \delta_t\})$.

One natural counter-factual question is how much can Home welfare increase without reducing Foreign's utility, conditoinal on a sequence of static wedges $\{\delta_t\}$. That is,

$$\max_{\{C_t^*\}} \mathcal{U}(\{C_t^*, \delta_t\}) \quad s.t. \quad \sum_{t=0}^{\infty} \beta^t U(C_t^*) \geq \overline{U}^*,$$

where $\overline{U}^*$ is the utility of Foreign along the realized empirical path. We shall compute the gap between the solution to this problem and the observed Home utility (in consumption equivalent units) for various countries in Section 6. Here we provide some insight into the nature of the welfare gains.

To do so, consider a second order approximation of $\mathcal{U}(\{\lambda_t, \delta_t\})$ around the first best ($\lambda_t = 0, \delta_t = 0, \forall t$). In the appendix we show:

Proposition 2. *Home's welfare can be expressed to a second-order approximation as:*

$$\mathcal{U}(\{\lambda_t, \delta_t\}) = \mathcal{U}(\{0, 0\}) + \sum_{t=0}^{\infty} \beta^t \left(A_t \lambda_t^2 + B_t \delta_t^2\right),$$

where A_t and B_t are functions of the first-best allocation.

[7] This assumes that the consumption possibility frontier – and hence $\boldsymbol{Q}_t$ – is well behaved, ensuring there is a unique C_t^* that satisfies this equation.

The fact that the first-order terms are absent is due to the Envelope Theorem and the fact that the point of approximation is the first-best allocation. The notable omission is the "cross derivative" terms $\lambda_t \delta_t$. The intuition for this also has an envelope logic.

At the first best, static efficiency states that if the planner were to increase Foreign consumption by one unit, both Home and Foreign are indifferent regarding how to construct the Foreign composite from the two underlying goods. That is, they are indifferent regarding static exchanges of C_H for C_F at the margin. The static wedge δ forces them to choose a particular ratio, but this has a second order cost. Hence, small movements in δ do not affect the costs of shifting consumption from Home to Foreign. In other words, $\partial \boldsymbol{Q}_t / \partial \delta_t = 0$ at $\delta_t = 0$. Hence, there is no second-order consequence of δ on the costs of the inter-temporal distortion λ.

The squared terms λ_t^2 and δ_t^2 have a Harberger triangle intuition. Recall that λ_t is like a tax on inter-temporal trades and δ_t is a tax on static trades. The welfare costs of these wedges in the neighborhood of zero is equal to the square of the tax rate times a term reflecting elasticities, which is captured by the terms A_t and B_t. The absence of a cross-product term implies that the interaction of these tax-wedges are of third-order in terms of welfare.

4.2 Private welfare costs (ψ)

From the perspective of a private agent in Home or Foreign, the presence of ψ (or, equivalently, a difference in the stochastic discount factors across regions) represents a potential welfare gain from better risk sharing. We can compute a partial equilibrium measure of this gain, whereby partial equilibrium we mean that we hold prices and exchange rates fixed and consider the perspective of an atomistic Home agent.

In particular, consider holding $\overline{w}$ fixed in (21), but allowing a single Home agent to trade at Foreign asset prices $\Lambda^*(s^t)$. That is, we set $\psi(s^t) = 0$ for all s^t for this agent. This is equivalent to choosing a consumption sequence holding constant the second line of (21). We can than ask what are the perceived gains to an individual Home agent from engaging in these trades relative to the observed allocation over a particular history s^t.

The welfare gain in consumption equivalent units, denoted τ, is defined by

$$(1+\tau)^{1-\sigma} = \frac{\sum_t \beta^t U(\tilde{C}_t)}{\sum_t \beta^t U(C_t)},$$

where $\tilde{C}_t = \tilde{C}(s^t)$ represents the counter-factual consumption sequence for this Home agent. In the appendix, we prove the following:

Proposition 3. *Let $\{\hat{C}(s^t)\}$ denote the original allocation relative to a reference year and $\hat{\psi}$ the distortion relative to a reference year. Define $\zeta_t \equiv \frac{\beta^t U(\hat{C}_t)}{\sum_k \beta^k U(\hat{C}_k)}$. The welfare gain for an individual Home agent from removing the wedge $\{\psi(s^t)\}$ is given by*

$$1 + \tau = \frac{\sum_t \zeta_t \hat{\psi}_t}{\left(\sum_t \zeta_t \hat{\psi}_t^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}} \geq 1.$$

Note that since $\zeta_t \in (0, 1)$ and sums to one, Jensen's inequality implies the final inequality, which is strict if $\hat{\psi}_t$ varies over time. This variation is why there are perceived private gains from additional inter-temporal trades. Note that one advantage of this proposition is that it involves only the original allocation and the measured wedges $\{\psi_t\}$, and hence is amenable to measurement.

5 Measurement

5.1 Data

One of the key advantages of our approach is its minimal data requirement: to measure the quality of risk sharing, we only need the time series for aggregate real quantities and a few cross-sectional moments to calibrate parameters of the model. To this end, we use the World Development Indicators (WDI) — the World Bank's premier compilation of cross-country comparable data — and focus on a balanced panel of 62 countries from 2000-2019. The period is chosen to maximize the sample size and excludes the high volatility of 2020 pandemic. We use annual data to focus on risk sharing at the medium frequency, the one at which the Backus-Smith puzzle is most pronounced (Corsetti et al., 2012).

Using the WDI data, we construct three aggregate series for the following quantities: GDP, consumption, and absorption, which we denote Y, C, and A, respectively. Absorption is defined as total local spending, computed as GDP minus net exports, which in turn equals the sum of consumption, investment, and government spending. For each series $X \in \{Y, C, A\}$, we compute real growth rates between a base year and period t. We select 2019 as our base year, and let $\widehat{X}_{i,t} \equiv X_{i,t}/\overline{X}_i$ denote the growth in X for country i between 2019 and year t, where $\overline{X}$ denotes base-year quantities. For Y and C, we use the respective constant local currency indices provided by WDI to compute real growth rates. For absorption, we deflate the nominal series by the GDP deflator.

For each country i, we take Foreign to be the rest of the world (ROW). To compute ROW growth for series X, we first construct a global growth rate as follows. In the base year 2019, we compute country i's share in the global economy for each series. In particular, letting $X_{i,t}$ denote

the current (2019) dollar value of X for country i, define its share (weight) by $wt_{i,X} \equiv X_{i,t}/\sum_{j\in J} X_{j,t}$, where J is our sample of 62 countries. The global growth in X is then $G_{X,t} = \sum_{j\in J} wt_{j,X}\widehat{X}_{j,t}$. For country i, the growth in the ROW is computed as

$$\widehat{X}^*_t \equiv \frac{(G_{X,t} - wt_{i,X}\widehat{X}_{i,t})}{1 - wt_{i,X}}.$$

Finally, the real exchange rate for country i is computed as the weighted average real exchange rate against all other countries, where weights are each country's share of ROW GDP in the base year. That is

$$Q_{i,t} \equiv \sum_{j\neq i} wt_{Y,j} Q_{i,j,t}/(1 - wt_{Y,i}),$$

where $Q_{i,j,t}$ is the real exchange between i and j in year t relative to 2019.[8]

5.2 Calibration and the Measurement of Inputs

To recover the quantities of Home and Foreign goods consumed in each country we follow the procedure discussed in Section 2. We assume that absorption in Home and Foreign is produced using the CES aggregators C and C^* defined in (1) and (2), respectively. In particular, $A_t = C\left(A_{H,t}, A_{F,t}\right)$, where $A_{H,t}$ and $A_{F,t}$ is the amount of the Home and Foreign goods, respectively, used to produce A_t. Similarly, $A^*_t = C^*\left(A^*_{H,t}, A^*_{F,t}\right)$.

Given our homothetic functional form assumptions, it is straightforward to define the environment in growth rates, as in the "exact hat algebra" of Dekle et al. (2007). Letting $t = 0$ denote the base year, we have for the Home country

$$\widehat{A}_t \equiv \frac{A_t}{\bar{A}} = \frac{C\left(A_{H,t}, A_{F,t}\right)}{\bar{A}} = C\left(\frac{A_{H,t}}{\bar{A}}, \frac{A_{F,t}}{\bar{A}}\right) = C\left(\frac{\bar{A}_H}{\bar{A}}\widehat{A}_{H,t}, \frac{\bar{A}_F}{\bar{A}}\widehat{A}_{F,t}\right).$$

where $t = 0$ is our base year and $A_{H,t}$ and $A_{F,t}$ is the amount of the Home and Foreign goods, respectively, used to produce A_t. We can rewrite the final term in more compact form by defining an aggregator in terms of growth rates. Defining $\bar{\gamma}$ as

$$\bar{\gamma} \equiv \gamma^{\frac{1}{\theta}}\left(\frac{\bar{A}_F}{\bar{A}}\right)^{\frac{\theta-1}{\theta}},$$

[8] The mean correlation of the resulting series with the real effective exchange rate from the WDI is 0.84 and the relative standard deviation is about 1.4.

we have

$$\widehat{A} = \left((1-\bar{\gamma})\widehat{A_H}^{\frac{\theta-1}{\theta}} + \bar{\gamma}\widehat{A_F}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}$$
$$\widehat{A^*} = \left(\bar{\gamma}^*(\widehat{A_H^*})^{\frac{\theta-1}{\theta}} + (1-\bar{\gamma}^*)(\widehat{A_F^*})^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}.$$

Similarly, we can scale the resource constraints:

$$\hat{Y}_t \equiv \frac{Y_t}{\bar{Y}} = \frac{A_{H,t} + A^*_{H,t}}{\bar{Y}} = (1-\bar{\chi})\hat{A}_H + \bar{\chi}\hat{A}^*_H \tag{22}$$
$$\hat{Y}^*_t \equiv \frac{Y^*_t}{\bar{Y}^*} = \frac{A_{F,t} + A^*_{F,t}}{\bar{Y}^*} = \bar{\chi}^*\hat{A}_F + (1-\bar{\chi}^*)\hat{A}^*_F, \tag{23}$$

where $\bar{\chi} \equiv \bar{A}^*_H/\bar{Y}$ is the share of Home's production exported in the base year, and $\bar{\chi}^* \equiv \bar{A}_F/\bar{Y}^*$ is the share of Foreign's output exported in the base year.

We parameterize the model as follows. We use conventional values for the inter- and intra-temporal elasticities: $\sigma = 2$ and $\theta = 4$ (see e.g. Chari et al., 2002; Backus et al., 1994). To calibrate $\overline{\chi}$, we take the ratio of exports in the base year and divide by a measure of GDP adjusted for purchasing power parity (PPP). The PPP adjustment reflects that relative prices are not equalized across countries, consistent with the static distortion δ introduced in Section 2. In particular, we compute non-exported output as the dollar value of GDP minus the dollar value of exports, all divided by the WDI's PPP adjustment factor. Hence, a country with a relatively low price for nontradables (i.e., a PPP adjustment less than one), is scaled up. With this adjustment, we compute

$$\bar{\chi} = \frac{Exp}{Exp + (GDP - Exp)/PPP},$$

where Exp is the value of base-year exports, GDP is the value of base-year GDP, both expressed in dollars, and PPP is the PPP adjustment factor. Similarly, $\bar{\chi}^*$ is computed as the weighted average of all other country's $\bar{\chi}$. The constant $\bar{\gamma}$ is the elasticity of absorption with respect to the foreign input, which we assume is related to its spending share in the base year. In particular, we set $\bar{\gamma}$ equal to the ratio of the dollar value of imports to the dollar value of domestic absorption (which, recall, is output minus net exports). For $\overline{\gamma}^*$, we again take the weighted sums across all other countries, using relative GDP as weights. Note that we only need the level of variables to compute the constants reflecting base-year shares. All other years are expressed in terms of growth rates.

With these parameters in hand, we then use the system of four equations in four unknowns, now in growth rates, outlined in Section 2. This provides a series for $\widehat{A}_{H,t}, \widehat{A}_{F,t}, \widehat{A}^*_{H,t}$ and $\widehat{A}^*_{F,t}$. We can then compute the inputs for consumption (absorption net of investment and government

spending) using the relative growth of consumption to absorption:

$$\widehat{C}_{k,t} = \widehat{A}_{k,t} \times \frac{\widehat{C}_t}{\widehat{A}_t}, \quad k = H, F, \tag{24}$$

and similarly for $\widehat{C}^*_{k,t}$. That is, the inputs of Home and Foreign goods for consumption are a scaled version of the inputs for absorption, with the scale factor representing how much consumption has grown relative to total absorption. In the counter-factuals, we take $\widehat{C}_{H,t} + \widehat{C}^*_{H,t}$ as the total Home goods (relative to the base year) available for use in consumption in the two regions. This assumes that the growth in investment and government expenditure are primitives and are not changed in the counter-factual.

5.3 Identification of Wedges

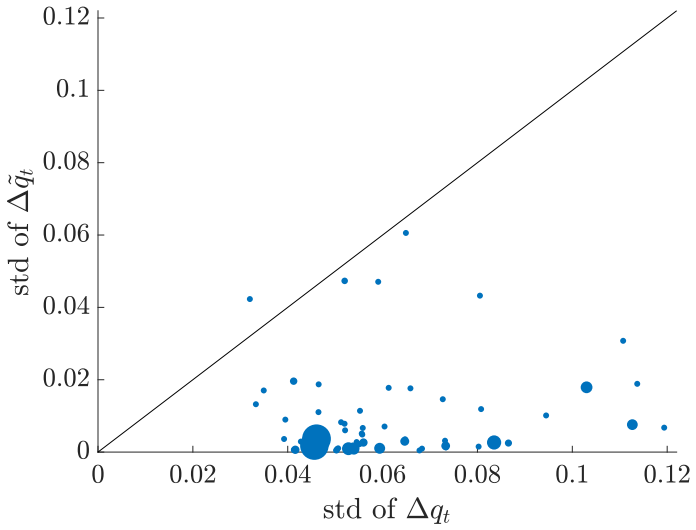
To back out the wedges, the system of optimality conditions from the constrained planner's problem can also be rewritten in terms of base-year trade shares, growth rates of real quantities, shadow prices of home and foreign goods ν, ν^* and two distortions η, λ:

$$\begin{aligned}
(1-\bar{\gamma})\left[\lambda\left(\frac{\bar{C}}{\bar{C}^*}\right)^{1-\sigma}\right]\hat{C}^{1-\sigma}\left(\frac{\hat{C}_H}{\hat{C}}\right)^{\frac{\theta-1}{\theta}} &= (1-\bar{\chi})\hat{C}_H(\bar{Y}\nu) - \eta, \\
\bar{\gamma}\left[\lambda\left(\frac{\bar{C}}{\bar{C}^*}\right)^{1-\sigma}\right]\hat{C}^{1-\sigma}\left(\frac{\hat{C}_F}{\hat{C}}\right)^{\frac{\theta-1}{\theta}} &= \bar{\chi}^*\hat{C}_F(\bar{Y}^*\nu^*) + \eta, \\
\bar{\gamma}^*\hat{C}^{*1-\sigma}\left(\frac{\hat{C}^*_H}{\hat{C}^*}\right)^{\frac{\theta-1}{\theta}} &= \bar{\chi}\hat{C}^*_H(\bar{Y}\nu) + \eta, \\
(1-\bar{\gamma}^*)\hat{C}^{*1-\sigma}\left(\frac{\hat{C}^*_F}{\hat{C}^*}\right)^{\frac{\theta-1}{\theta}} &= (1-\bar{\chi}^*)\hat{C}^*_F(\bar{Y}^*\nu^*) - \eta.
\end{aligned}$$

Notice that Lagrange multipliers and Pareto weights enter in level, not changes, and are not separately identified from base-year quantities $\bar{Y}, \bar{Y}^*, \bar{C}, \bar{C}^*$. Because the values of the latter are the same in all periods, however, the changes relative to the base year reveal time variation in shadow prices $\hat{\nu}, \hat{\nu}^*$ and the risk-sharing distortion $\hat{\lambda}$. With these estimates at hand, one can evaluate the MRT and the quality of risk sharing based on

$$\left[\frac{1}{\lambda}\left(\frac{\bar{C}}{\bar{C}^*}\right)^{\sigma-1}\right]\left(\frac{\hat{C}}{\hat{C}^*}\right)^{\sigma} = \frac{\bar{\chi}\frac{\hat{C}^*_H}{\hat{C}^*}(\bar{Y}\nu) + (1-\bar{\chi}^*)\frac{\hat{C}^*_F}{\hat{C}^*}(\bar{Y}^*\nu^*)}{(1-\bar{\chi})\frac{\hat{C}_H}{\hat{C}}(\bar{Y}\nu) + \bar{\chi}^*\frac{\hat{C}_F}{\hat{C}}(\bar{Y}^*\nu^*)}.$$

Figure 4: Relative Volatility of MRT and RER



The system also sheds new light on our calibration strategy. Applied to the base year when changes in all variables are equal one, we get a system with four trade shares $\bar{\gamma}, \bar{\gamma}^*, \bar{\chi}, \bar{\chi}^*$, two shadow prices ν, ν^* and two distortions proportionate to η and λ. The reason one needs four moments to calibrate the model despite only two unknown demand shifters γ and γ^* is that there are also the two wedges that parametrize the deviations from consumption possibility frontier. Intuitively, import shares capture the degree of home bias in consumption, the net export is informative about Pareto weight λ, and the ratio of real exports to GDP provides information about distortion in relative prices across countries η.[9]

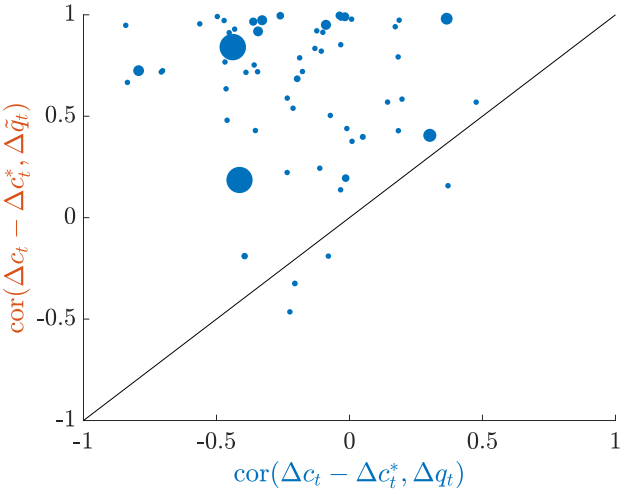
6 Empirical Results

We begin with a comparison of the observed real exchange rate Q_t with our planner's MRT $\tilde{Q}_t$. Letting Δq_t denote the log of the first difference of Q_t, and similarly for $\Delta\tilde{q}_t$ and $\tilde{Q}_t$, Figure 4 provides a scatter plot of the relative volatities. Each point is a country, where the standard deviation is taken over the sample period. The main takeaway from the figure is that the implied *MRT* is roughly five to ten times smaller. Viewed through the model, this suggests that at annual frequencies the observed real exchange rate has significant volatility unrelated to the social marginal rate of transformation between Home and Foreign consumption composites. The standard deviation of $\tilde{Q}_t$ is on the order of one percent, which is in line with the volatility of many macroeconomic aggregates at annual frequencies.

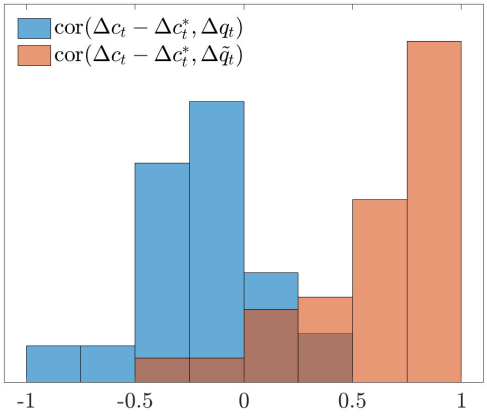
[9] Adjusting for PPP differences is important to ensure that export shares (for quantities) contain additional information relative to import shares (in values) and is equivalent to $\eta \neq 0$ in the base period.

Figure 5: Consumption Differences and the Real Exchange Rate

(a) Relative Correlations



(b) Histogram of Correlations



The lower volatility of $\tilde{Q}_t$ speaks to one dimension of exchange rate disconnect. Another dimension is the lack of correlation between the exchange rate and macro variables. In Figure 5 we depict the correlation of the respective notions of the exchange rate with the difference in log consumption growth between Home and Foreign. Panel (a) is a scatter plot, with each point again representing a country (with Foreign a composite of the rest of the world). Panel (b) is the histogram of correlations, depicting the respective distributions across countries. We see that the risk sharing condition holds much better with our planner's MRT than with the observed real exchange rate, at least at annual frequencies.

Figure 5 speaks to the fact that the optimal risk sharing condition becomes more evident in the data once we compute the planner's marginal rate of transformation. To show this more directly, in Figure 6 we depict the volatilities of the log change in the inter-temporal wedges, λ_t and ψ_t in Panel (a) and the volatilities of the level of the respective wedges in Panel (b). In both, the volatility of λ_t is almost uniformly lower than its counterpart ψ_t. However, the difference is more modest at longer horizons.

To give a sense of the underlying time series, Figure 7 depicts the path of the three wedges, λ_t, ψ_t, and δ_t, for a selection of countries. Panel (a) is in annual log first differences and Panel (b) is in levels.

We now turn to how the model can implement the better risk sharing in a decentralized equilibrium. The crucial component is that the implied relative prices faced by consumers in each region are not highly volatile and are insulated from exchange rate movements. In Figure 8 we plot the model implied relative prices faced by consumers in the two regions (P_F/P_H and P_F^*/P_H^*) as well as the deviation from the law of one price for the Home good (P_H^*/P_H). The relative prices

Figure 6: The Volatilities of Inter-Temporal Wedges

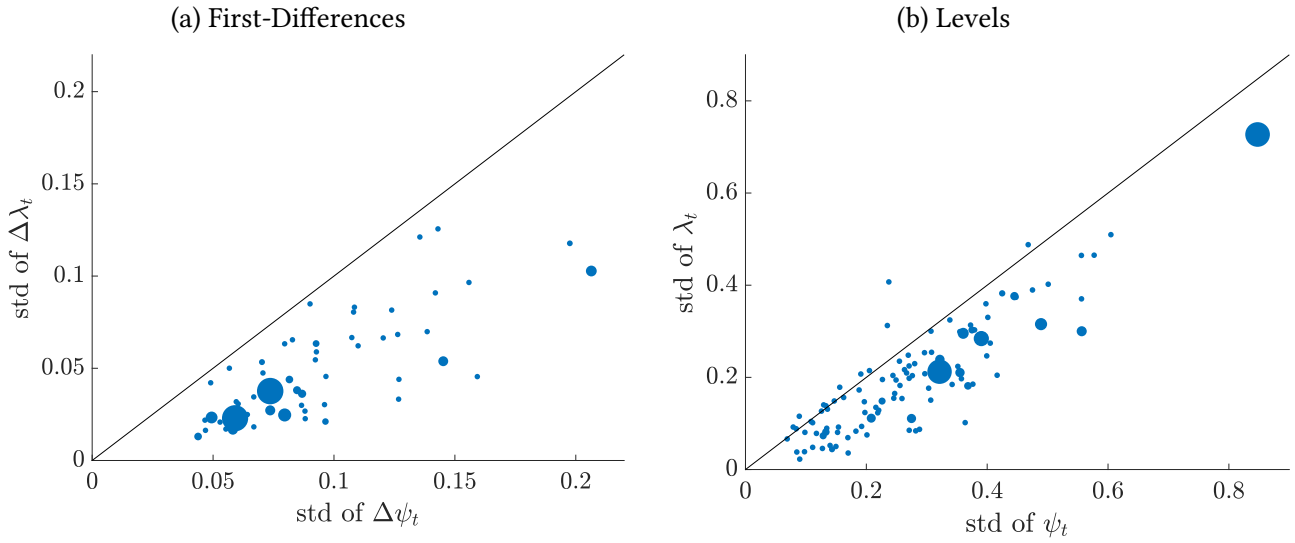
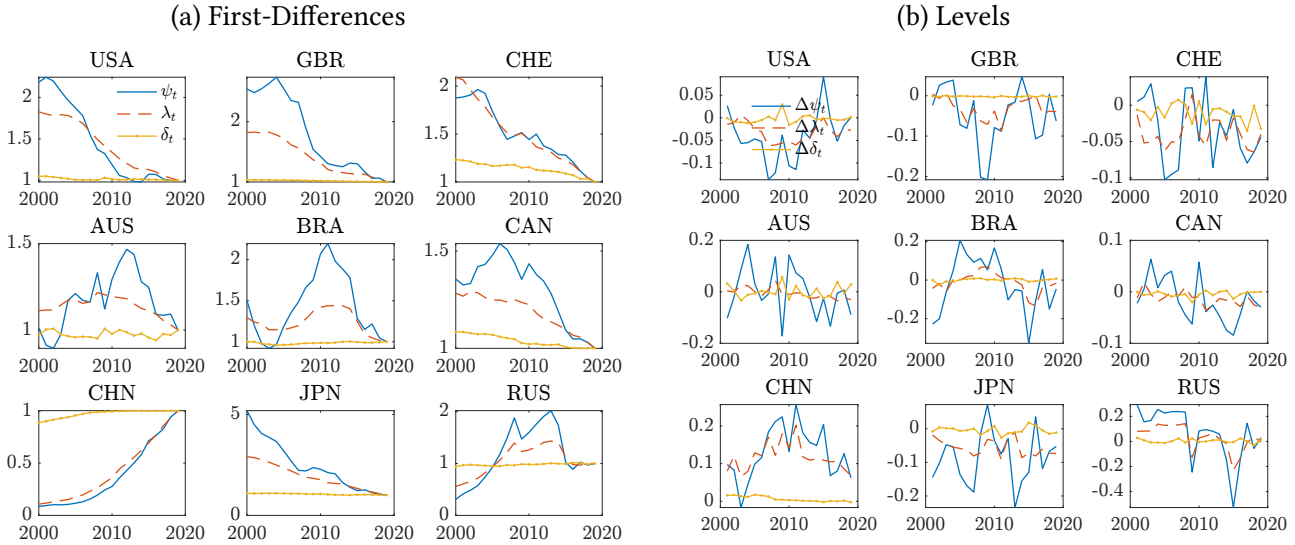


Figure 7: Time Series of Wedges



faced by consumers in the repsective regions are notably more stable over time than the real exchange rate and the deviation from the law of one price. In fact, the LOP deviations track the exchange rate quite closely. Hence, the volatility of the observed exchange rate are not reflected in the marginal rate of substitution of consumers in either region. Consequently, the decentralization implies that relative markups across borders are moving to offset any volatility to the exchange rate. This is broadly consistent with the LCP pricing protocol.

With these results in hand, we can return to the welfare calculations implied by the wedges. In

Figure 8: Welfare Gains

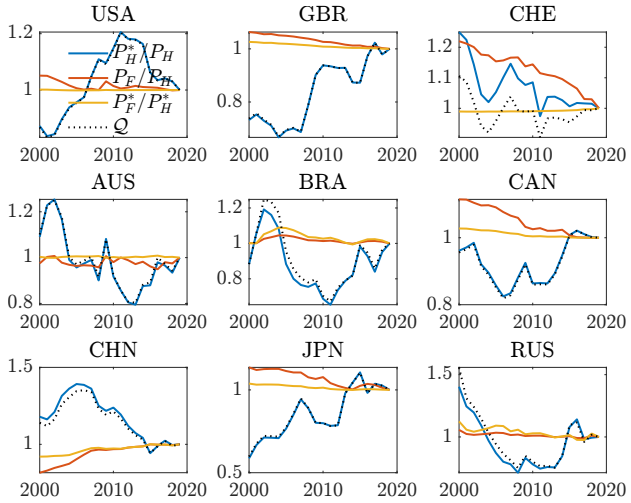


Figure 9: Relative Volatility of MRT and RER

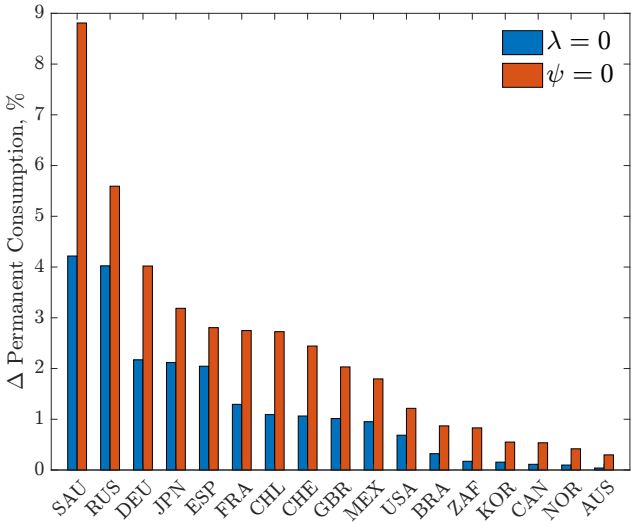


Figure 9 we plot the welfare gain for Home enjoyed by shutting down the planner's wedge versus the private sector wedge. The implied welfare gains are much larger when measured with respect to ψ_t than for λ_t. That is, private agents perceive a much larger potential gain to undistorted intertemporal trades than does the planner. This reflects that the private agents do not internalize that some of these perceived gains are simply movements in relative markups.

7 Conclusion

In this paper we proposed an alternative measure of the relative costs of consumption across regions. To do so, we considered a planner optimally allocating output toward consumption in two regions over time. A key construct is the planner's marignal rate of transformation between Home and Foreign consumption composites, which represents a shadow real exchange rate. This shadow exchange rate is an alternative to the observed real exchange in testing the well known Backus-Smith condition for optimal risk sharing. We show that risk sharing distortions are significantly smaller when the planner's relative costs are used to price relative consumption. We also showed how such an allocation can be decentralized. The primary distortion in the decentralized equilibrium is a shock to the real exchange rate that is not reflected in macroeconomic aggregates. Rather, this shock is offset by movements in the relative markups implied by the deviations from the law of one price. In particular, consumers in each region face stable implied prices while the price of the same good sold in two different locations inherits the volatility of the exchange rate. This pattern is broadly consistent with the local currency pricing protocol. The net result of the empirical exercise is that observed consumption allocations are much closer to efficient than would be inferred using observed movements in exchange rates.

References

Alvarez, Fernando, Andrew Atkeson, and Patrick J. Kehoe, "Money, Interest Rates, and Exchange Rates with Endogenously Segmented Markets," *Journal of Political Economy*, Feb 2002, *110* (1), 73–112.

_ , _ , and _ , "Time-Varying Risk, Interest Rates, and Exchange Rates in General Equilibrium," *Review of Economic Studies*, 2009, *76* (3), 851–878.

Armington, Paul S., "A Theory of Demand for Products Distinguished by Place of Production (Une theorie de la demande de produits differencies d'apres leur origine) (Una teoria de la demanda de productos distinguiendolos segun el lugar de produccion)," *Staff Papers - International Monetary Fund*, March 1969, *16* (1), 159.

Backus, David K. and Gregor W. Smith, "Consumption and real exchange rates in dynamic economies with non-traded goods," *Journal of International Economics*, November 1993, *35* (3–4), 297–316.

Backus, David K, Patrick J Kehoe, and Finn E Kydland, "International Real Business Cycles," *Journal of Political Economy*, August 1992, *100* (4), 745–75.

_ , _ , and _ , "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?," *American Economic Review*, March 1994, *84* (1), 84–103.

Bai, Yan and Jing Zhang, "Solving the Feldstein-Horioka Puzzle With Financial Frictions," *Econometrica*, 2010, *78* (2), 603–632.

_ and _ , "Financial Integration and International Risk Sharing," *Journal of International Economics*, January 2012, *86* (1), 17–32.

Baxter, Marianne and Mario J Crucini, "Business Cycles and the Asset Structure of Foreign Trade," *International Economic Review*, 1995, *36* (4), 821–854.

Brandt, Michael W., John H. Cochrane, and Pedro Santa-Clara, "International risk sharing is better than you think, or exchange rates are too smooth," *Journal of Monetary Economics*, May 2006, *53* (4), 671–698.

Burstein, Ariel T. and Gita Gopinath, "International Prices and Exchange Rates," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. IV 2012.

Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan, "Can Sticky Price Models Generate Volatile and Persistent Exchange Rates?," *Review of Economic Studies*, 2002, *69* (3), 533–63.

Corsetti, Giancarlo, Luca Dedola, and Francesca Viani, "The international risk sharing puzzle is at business cycle and lower frequency," *Canadian Journal of Economics/Revue canadienne d'économique*, 2012, *45* (2), 448–471.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced trade," *American Economic Review*, 2007, *97* (2), 351–355.

Farhi, Emmanuel and Iván Werning, "Fiscal Unions," *American Economic Review*, December 2017, *107* (12), 3788–3834.

Fitzgerald, Doireann, "Trade costs, asset market frictions, and risk sharing," *American Economic Review*, 2012, *102* (6), 2700–2733.

Gabaix, Xavier and Matteo Maggiori, "International Liquidity and Exchange Rate Dynamics," *The Quarterly Journal of Economics*, 2015, *130* (3), 1369–1420.

Gopinath, Gita and Oleg Itskhoki, "The Existence and Implications of Dominant Currencies," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. V 2021. forthcoming.

Heathcote, Jonathan and Fabrizio Perri, "Assessing international efficiency," in "Handbook of International Economics," Vol. 4, Elsevier, 2014, pp. 523–584.

Itskhoki, Oleg and Dmitry Mukhin, "Exchange Rate Disconnect in General Equilibrium," *Journal of Political Economy*, 2021, *129* (8), 2183–2232.

Jeanne, Olivier and Andrew K. Rose, "Noise Trading and Exchange Rate Regimes," *The Quarterly Journal of Economics*, 2002, *117* (2), 537–569.

Ohanian, Lee E, Paulina Restrepo-Echavarria, and Mark LJ Wright, "Bad investments and missed opportunities? postwar capital flows to asia and latin america," *American Economic Review*, 2018, *108* (12), 3541–3582.

Samuelson, Paul A., "International Factor-Price Equalisation Once Again," *The Economic Journal*, June 1949, *59* (234), 181.

Appendix

A Data Appendix

B Calibration

In this appendix, we outline the calibration and measurement procedures. We take the perspective of a single country ("Home") facing the Rest of the World ("Foreign"). Unless otherwise stated, all variables are measured in current US dollars.

For notation, we use the following subscripts: H denotes Home, F denotes the rest of the world (Foreign). Let $\mathcal{F}$ denote the set of countries in Foreign. When indexing individual countries, the index i will be reserved for the Home country.

We shall calibrate certain ratios for a "base year" in order to use growth relative to the base year in the remaining periods. We use a "bar" to denote base-year values. We will need to make some assumptions about the data in the base year, which we will outline here:

(i) Let $\tau_{ij} \geq 1$ denote the iceberg transportation cost of Home's exports to country j. We assume that transport costs are constant over time and that $\tau_{ij} = \tau$ for all $j \in \mathcal{F}$.

(ii) Let $\bar{P}^{PPP}$ denote Home's PPP-adjusted GDP price deflator in the base year t_0. That is, Home's PPP-adjusted GDP in the base year can be expressed as $\bar{P}^{PPP}\bar{Y}$. We assume that the price of Home's exports to country j is the same as the price of Home's GDP in the base year. That is, $\bar{P}_{ij} = \bar{P}^{PPP}$ for all $j \in \mathcal{F}$, where $\bar{P}_{ij}$ is the price of Home's exports to j in t_0 (net of transportation costs).

(iii) Let $\tau_{jk} \geq 1$ denote the iceberg transportation cost of country j's exports to country k. We assume that transport costs are constant over time and that $\tau_{jk} = \tau^*$ for all $j \in \mathcal{F}$ and $j \neq k$.

(iv) Let $\bar{P}_{ji}$ denote the price of an import from country $j \in \mathcal{F}$ to Home (country i) in the base year. Let $\bar{P}_j^{PPP}$ denote country j's PPP-adjusted GDP price deflator in the base year. We assume that $P_{ji} = \tau_{ji}\bar{P}_j^{PPP} = \tau^*\bar{P}_j^{PPP}$.

Assumptions (i) and (iii) are consistent with treating Foreign as a single aggregate of the rest-of-the-world. Assumptions (ii) and (iv) are essentially states that Home exports to the rest-of-the-world at a price level equivalent to the PPP adjusted price of GDP. This assumes exports are priced the same regardless of destination, an assumption we relax in all other years than the base year. The assumption in the base year is necessitated given that we do not observe destination-specific or source-specific pricing in Home's exports and imports. Note that export prices will differ than domestic prices to the extent that the PPP adjustment differs from one.

B.1 Home Country Parameters

Let ABS_t denote Home's "absorption", which is the sum of consumption, investment, and government expenditure, expressed in period t US dollars (we will omit H or i subscripts for brevity when the context is clear). Let real absorption be generated by same aggregator as consumption, C, defined in equation (1). Specifically,

$$A_t \equiv C\left(A_{H,t}, A_{F,t}\right) = \left((1-\gamma)^{\frac{1}{\theta}} A_{H,t}^{\frac{\theta-1}{\theta}} + \gamma^{\frac{1}{\theta}} A_{F,t}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}},$$

where A_H and A_F are the Home and Foreign components of absorption, respectively. For a base year $t = t_0$, let a "bar" denote the base-year value of a variable and a "hat" denote growth relative to the base year. Dividing through by base-year absorption $\bar{A}$, we obtain

$$\hat{A}_t \equiv \frac{A_t}{\bar{A}} = \hat{C}(\hat{A}_{H,t}, \hat{A}_{F,t}) \equiv \left((1-\bar{\gamma})\hat{A}_{H,t}^{\frac{\theta-1}{\theta}} + \bar{\gamma}\hat{A}_{F,t}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}, \tag{25}$$

where, from the definition of C, we have

$$\bar{\gamma} \equiv \gamma^{\frac{1}{\theta}} \left(\frac{\bar{A}_F}{\bar{A}}\right)^{\frac{\theta-1}{\theta}} \text{ and } 1-\bar{\gamma} = \gamma^{\frac{1}{\theta}} \left(\frac{\bar{A}_H}{\bar{A}}\right)^{\frac{\theta-1}{\theta}}.$$

The resource constraint on Home goods implies

$$Y_t = A_{H,t} + \tau A_{H,t}^*,$$

where Y_t is Home's real GDP, $A_{H,t}^*$ is the rest of the world's absorption of Home goods (Home's real exports net of transport costs), and τ is the cost of exporting from Home to the Foreign (see Assumption (i) above). The real growth rate of Home's GDP as

$$\hat{Y}_t = \frac{\bar{A}_H}{\bar{Y}}\hat{A}_{H,t} + \frac{\tau\bar{A}_H^*}{\bar{Y}}\hat{A}_{H,t}^* = (1-\bar{\alpha})\hat{A}_{H,t} + \bar{\alpha}\hat{A}_{H,t}^*, \tag{26}$$

where the second equality defines $\bar{\alpha} \equiv \tau\bar{A}_H^*/\bar{Y}$.

Calibrating $\bar{\alpha}$: The parameter $\bar{\alpha}$ is the ratio of the quantity of Home exports (inclusive of transport costs) divided by real GDP. However, we do not directly observe the quantity of exports or output in the base year. We observe nominal exports in US dollars, $\overline{EXP}$, which we assume are measured including transportation costs. From Assumption (ii) above, exports are priced at $\overline{P}^{PPP}$. Letting $\overline{GDP}^{PPP} = \bar{P}^{PPP}\bar{Y}$ denote Home's PPP-adjusted GDP in the base year. We then have,

$$\bar{\alpha} = \frac{\bar{P}^{PPP}\tau\bar{A}_H^*}{\bar{P}^{PPP}\bar{Y}} = \frac{\overline{EXP}}{\overline{GDP}^{PPP}},$$

where $\overline{EXP}$ is Home's exports in the base year (inclusive of transport costs) and $\overline{GDP}^{PPP}$ is Home's PPP-adjusted GDP in the base year, both in dollars and both obtained from the WDI database.

Calibrating $\bar{\gamma}$: If Home agents faced a base-year price $\bar{P}_F$ for Foreign goods and a price $\bar{P}$ for the composite commodity, cost-minimization implies

$$\gamma^{\frac{1}{\theta}} \left(\frac{\bar{A}_F}{\bar{A}} \right)^{\frac{-1}{\theta}} = \frac{\bar{P}_F}{\bar{P}}.$$

Multiplying through by $\bar{A}_F/\bar{A}$, we obtain

$$\bar{\gamma} = \frac{\bar{P}_F \bar{A}_F}{\bar{P}\bar{A}}$$

as the expenditure share of foreign imports out of total expenditure. The challenge is that imports are measured by statistical agencies "at the dock" of the respective importing country, but not necessarily at the prices faced by consumers when choosing between Home and Foreign goods at the retail level. Nevertheless, for the base year only, we assume that the price of imports at the dock is the price for Foreign goods faced by consumers at the retail level. Under this relative price assumption, we have

$$\bar{\gamma} = \frac{\overline{IMP}}{\overline{ABS}},$$

where $\overline{IMP}$ is base-year nominal imports in US dollars, and $\overline{ABS}$ is the sum of consumption, investment, and government expenditure, also expressed in base-year US dollars.

From Absorption to Consumption: We have assumed that the same aggregator is used for all components of absorption. Assuming that the relative price of Home and Foreign inputs are the same across absorption components, then the expenditure shares on Home versus Foreign will be the same for consumption and absorption. That is,

$$\frac{P_{H,t}C_{H,t}}{P_t C_t} = \frac{P_{H,t}A_{H,t}}{P_t A_t} \quad \Rightarrow C_{H,t} = \frac{C_t}{A_t} A_{H,t} \quad \Rightarrow \hat{C}_{H,t} = \frac{\hat{C}_t}{\hat{A}_t} \hat{A}_{H,t} \tag{27}$$

$$\frac{P_{F,t}C_{F,t}}{P_t C_t} = \frac{P_{F,t}A_{F,t}}{P_t A_t} \quad \Rightarrow C_{F,t} = \frac{C_t}{A_t} A_{F,t} \quad \Rightarrow \hat{C}_{F,t} = \frac{\hat{C}_t}{\hat{A}_t} \hat{A}_{F,t} \tag{28}$$

B.2 Foreign Parameters

For $j \in \mathcal{F}$, let $\hat{Y}_{j,t}$ denote real GDP growth in country j relative to the base year. We define world GDP growth as

$$\hat{Y}_t^* \equiv \sum_{j \in \mathcal{F}} \omega_j^Y \hat{Y}_{j,t},$$

where the weights ω_j^Y are defined as country j's share of PPP-adjusted world GDP in the base year:

$$\omega_j^Y \equiv \frac{\overline{GDP}_j^{PPP}}{\sum_{k\in\mathcal{F}} \overline{GDP}_k^{PPP}}.$$

Let $A_{jk,t}$ denote country j's exports to country k in period t and τ_{jk} the associated trade cost, and $A_{jj,t}$ denote domestic sales within j (and $\tau_{jj} = 1$). Letting i denote the Home country, the country-j resource constraint is

$$Y_{j,t} = \tau_{ji} A_{ji,t} + \sum_{k\in\mathcal{F}} \tau_{jk} A_{jk,t}.$$

Let $\bar{\alpha}_{jk}^* \equiv \tau_{jk}\bar{A}_{jk}/\bar{Y}_j$ denote country j's exports to country k, inclusive of trade costs, as a share of GDP in the base year. We then have

$$\hat{Y}_{j,t} = \bar{\alpha}_{ji}^* \hat{A}_{ji,t} + \sum_{k\in\mathcal{F}} \bar{\alpha}_{jk}^* \hat{A}_{jk,t}.$$

Taking the weighted sum across all foreign countries:

$$\hat{Y}_t^* = \sum_{j\in\mathcal{F}} \omega_j^Y \hat{Y}_{j,t} = \sum_{j\in\mathcal{F}} \omega_j^Y \bar{\alpha}_{ji}^* \hat{A}_{ji,t} + \sum_{j\in\mathcal{F}} \omega_j^Y \sum_{k\in\mathcal{F}} \bar{\alpha}_{jk}^* \hat{A}_{jk,t}.$$

Define the growth in Home's absorption of Foreign goods as

$$\hat{A}_{F,t} \equiv \frac{1}{\bar{\alpha}^*} \sum_{j\in\mathcal{F}} \omega_j^Y \bar{\alpha}_{ji}^* \hat{A}_{ji,t}, \tag{29}$$

where $\bar{\alpha}^*$ is a constant defined below. Similarly,

$$\hat{A}_{F,t}^* \equiv \frac{1}{1-\bar{\alpha}^*} \sum_{j\in\mathcal{F}} \omega_j^Y \sum_{k\in\mathcal{F}} \bar{\alpha}_{jk}^* \hat{A}_{jk,t}.$$

We then have

$$\hat{Y}_t^* = \bar{\alpha}^* \hat{A}_{F,t} + (1-\bar{\alpha}^*)\hat{A}_{F,t}^*.$$

We assume that Home (country i) consumes imports from the various foreign countries by aggregating into a composite A_F according to a Cobb-Douglas aggregator (or a first-order approximation of a more general aggregator function):

$$\hat{A}_{F,t} = \frac{1}{\bar{\gamma}} \sum_{j\in\mathcal{F}} \gamma_{ji} \hat{A}_{ji,t}, \tag{30}$$

where γ_{ji} is the expenditure share of country j's exports in Home's total absorption, which, using base-year values, is defined as:

$$\gamma_{ji} = \frac{\bar{P}_{ji}\bar{A}_{ji}}{\bar{P}\bar{A}},$$

where $\bar{P}_{ji}$ is the price of country j's exports to Home in the base year and $\bar{P}\bar{A}$ is Home's total (nominal)

absorption in the base year.

Calibratin $\bar{\alpha}^*$: As expression (29) equals (30) at every t, we need that weights on country j in the respective expressions must be equal. That is,

$$\begin{aligned}\bar{\alpha}^* &= \frac{\omega_j^Y \alpha_{ji}^*}{\gamma_{ji}} = \frac{\overline{GDP}_j^{PPP}}{\sum_{k\in\mathcal{F}} \overline{GDP}_k^{PPP}} * \frac{\tau_{ji}\bar{A}_{ji}}{\bar{Y}_j} * \frac{\sum_{k\in\mathcal{F}} \bar{P}_{ki}\bar{A}_{ki}}{\bar{P}_{ji}\bar{A}_{ji}} \\ &= \frac{\overline{IMP}}{\sum_{k\in\mathcal{F}} \overline{GDP}_k^{PPP}} * \frac{\tau_{ji}\overline{GDP}_j^{PPP}}{\bar{P}_{ji}\bar{Y}_j},\end{aligned}$$

where $\overline{IMP} = \sum_{k\in\mathcal{F}} \bar{P}_{ki}\bar{A}_{ki}$ is the nominal US dollar value of Home's imports in the base year. From Assumption (iv) above, $\bar{P}_j^{PPP} = \bar{P}_{ji}/\tau_{ji}$, and hence the last term is one. This implies:

$$\bar{\alpha}^* = \frac{\overline{IMP}}{\sum_{k\in\mathcal{F}} \overline{GDP}_k^{PPP}}. \tag{31}$$

Growth in Foreign absorption, $\hat{A}^*$ is a weighted sum of the growth rates observed in the individual foreign countries:

$$\hat{A}_t^* \equiv \sum_{j\in\mathcal{F}} \omega_j^A \hat{A}_{j,t},$$

with $\{\omega_j^A\}$ defined analogously to ω_j^Y:

$$\omega_j^A \equiv \frac{\overline{ABS}_j}{\sum_{k\in\mathcal{F}} \overline{ABS}_k},$$

and where $\overline{ABS}_j = \bar{P}_j\bar{A}_j$ is total nominal (in US dollars) absorption of country j in the base year. Define country j's expenditure share on country k's exports in the base year as

$$\gamma_{kj} \equiv \frac{\bar{P}_{kj}\bar{A}_{kj}}{\overline{ABS}_j}.$$

To a first-order approximation, we can write

$$\hat{A}_{j,t} = \gamma_{ij}\hat{A}_{ij} + \sum_{k\in\mathcal{F}} \gamma_{kj}\hat{A}_{kj,t}.$$

Combining these equations, we have

$$\hat{A}_t^* = \sum_{j\in\mathcal{F}} \omega_j^A \gamma_{ij}\hat{A}_{ij,t} + \sum_{j\in\mathcal{F}} \omega_j^A \sum_{k\in\mathcal{F}} \gamma_{kj}\hat{A}_{kj,t}. \tag{32}$$

Let $\hat{A}^*_H$ be the growth rate of Foreign absorption of Home goods. In particular, define

$$\hat{A}^*_{H,t} = \sum_{j\in\mathcal{F}} \alpha_{ij}\hat{A}_{ij,t}, \tag{33}$$

where

$$\alpha_{ij} \equiv \frac{\bar{A}_{ij}}{\sum_{k\in\mathcal{F}}\bar{A}_{ik}},$$

is the share of Home's exports that go to country j in the base year.[10] In order to write (32) as

$$\hat{A}^*_t = \bar{\gamma}^*\hat{A}^*_{H,t} + (1-\bar{\gamma}^*)\hat{A}^*_{F,t}$$

we search for weights such that

$$\hat{A}^*_{H,t} = \frac{1}{\bar{\gamma}^*}\sum_{j\in\mathcal{F}} \omega^A_j\gamma_{ij}\hat{A}_{ij,t}, \tag{34}$$

where $\bar{\gamma}^*$ is a constant defined below.

Calibrating $\bar{\gamma}^*$ and ω^A_j: Comparing (33) to (34), we require

$$\omega^A_j = \bar{\gamma}^*\frac{\alpha_{ij}}{\gamma_{ij}} = \bar{\gamma}^*\frac{\bar{A}_{ij}}{\sum_{k\in\mathcal{F}}\bar{A}_{ik}} * \frac{\bar{P}_j\bar{A}_j}{\bar{P}_{ij}\bar{A}_{ij}} = \bar{\gamma}^*\frac{\bar{P}_j\bar{A}_j}{\bar{P}_{ij}\sum_{k\in\mathcal{F}}\bar{A}_{ik}} = \bar{\gamma}^*\frac{\overline{ABS}_j}{\sum_{k\in\mathcal{F}}\bar{P}_{ik}\bar{A}_{ik}},$$

where the last equality uses Assumption (ii) above that states $\bar{P}_{ij} = \bar{P}_{ik} = \bar{P}^{PPP}$ for all $j,k \in \mathcal{F}$. We then have

$$\bar{\gamma}^* = \omega^A_j\frac{\overline{EXP}}{\overline{ABS}_j} = \frac{\overline{EXP}}{\sum_{j\in\mathcal{F}}\overline{ABS}_j},$$

where $\overline{EXP} = \sum_{k\in\mathcal{F}}\bar{P}_{ik}\bar{A}_{ik}$ is Home's dollar value of exports in the base year.

B.3 Measurement

We measure the parameters $\bar{\alpha}$, $\bar{\alpha}^*$, $\bar{\gamma}$, and $\bar{\gamma}^*$ using the World Bank's World Development Indicators (WDI) database. We use the following variables:

[10] Here, we can adjust for transport costs, as well, but under Assumption (iii), $\tau_{ik} = \tau_{ij} = \tau$ for every $jk \in \mathcal{F}$, and the transport costs cancel from α_{ij}.

For each year t, we have the following four relationships:

$$\hat{A}_t = \left((1-\bar{\gamma})\hat{A}_{H,t}^{\frac{\theta-1}{\theta}} + \bar{\gamma}A_{F,t}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}},$$

$$\hat{A}_t^* = \left((1-\bar{\gamma}^*)\hat{A}_{H,t}^{*\frac{\theta-1}{\theta}} + \bar{\gamma}^*\hat{A}_{F,t}^{*\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}},$$

$$\hat{Y}_t = (1-\bar{\alpha})\hat{A}_{H,t} + \bar{\alpha}\hat{A}_{H,t}^*,$$

$$\hat{Y}_t^* = \bar{\alpha}^*\hat{A}_{F,t} + (1-\bar{\alpha}^*)\hat{A}_{F,t}^*.$$

For each year t, these four equations determine the four unobserved quantities $\{\hat{A}_{H,t}, \hat{A}_{F,t}, \hat{A}_{H,t}^*, \hat{A}_{F,t}^*\}$. A unique solution corresponds to a tangency in the Edgeworth box diagram of Figure **??**. Two solutions correspond to a wedge $\delta_t \neq 0$, and we pick the solution for which the implied $\delta_t > 0$ (that is, home bias is exacerbated by the wedge). In a handful of cases (less than xx percent), there is no solution. This represents a mis-specification or measurement error, as the allocation is not feasible given the resources and preferences. We drop country-years for which that happens. We go from absorption to consumption using (27).

In each country, we take investment and government expenditure as given. The planning problem and counter-factuals involve reallocating the share of output that is devoted to consumption. Specifically, define

$$Y_{C,t} \equiv C_{H,t} + \tau C_{H,t}^* \text{ and } Y_{C,t}^* \equiv \tau^* C_{F,t} + C_{F,t}^*$$

where $\tau^* C_{F,t}$ is the consumption in Home of Foreign goods inclusive of transport costs. Dividing $Y_{C,t}$ by the base year value, we have

$$\hat{Y}_{C,t} \equiv \frac{Y_{C,t}}{\bar{Y}_C} = \frac{C_{H,t} + \tau C_{H,t}^*}{\tau^*\bar{C}_H + \bar{C}_H^*} = \frac{\frac{\bar{C}_H}{\bar{Y}}\hat{C}_{H,t} + \frac{\tau\bar{C}_H}{\bar{Y}}C_{H,t}^*}{\frac{\tau^*\bar{C}_H}{\bar{Y}} + \frac{\bar{C}_H^*}{\bar{Y}}},$$

where the last equality is obtained by dividing the numerator and denominator by $\bar{Y}$. Using the fact that $C_H = \frac{C}{A}A_H$, we have

$$\frac{\bar{C}_H}{\bar{Y}} = \frac{\bar{C}}{\bar{A}}\frac{\bar{A}_H}{\bar{Y}} = (1-\bar{\alpha})\frac{\bar{C}}{\bar{A}},$$

where the last equality uses the definition of $\bar{\alpha}$, and

$$\frac{\tau\bar{C}_H^*}{\bar{Y}} = \frac{\bar{C}^*}{\bar{A}^*}\frac{\tau\bar{A}_H}{\bar{Y}} = \bar{\alpha}\frac{\bar{C}}{\bar{A}}.$$

We then have

$$\hat{Y}_{C,t} = \frac{(1-\bar{\alpha})\frac{\bar{C}}{\bar{A}}\hat{C}_{H,t} + \bar{\alpha}\frac{\bar{C}^*}{\bar{A}^*}\hat{C}^*_{H,t}}{(1-\bar{\alpha})\frac{\bar{C}}{\bar{A}} + \bar{\alpha}\frac{\bar{C}^*}{\bar{A}^*}}.$$

Defining

$$\bar{\alpha}_C \equiv \frac{\bar{\alpha}\bar{C}^*/\bar{A}^*}{(1-\bar{\alpha})\bar{C}/\bar{A} + \bar{\alpha}\bar{C}^*/\bar{A}^*},$$

we have

$$\hat{Y}_{C,t} = (1-\bar{\alpha}_C)\hat{C}_{H,t} + \bar{\alpha}_C\hat{C}^*_{H,t}. \tag{35}$$

A similar derivation holds for $\hat{Y}^*_{C,t}$. In particular,

$$\bar{\alpha}^*_C \equiv \frac{\bar{\alpha}^*\bar{C}/\bar{A}}{(1-\bar{\alpha}^*)\bar{C}^*/\bar{A}^* + \bar{\alpha}^*\bar{C}/\bar{A}},$$

and

$$\hat{Y}^*_{C,t} = \bar{\alpha}^*_C\hat{C}_{H,t} + (1-\bar{\alpha}^*_C)\hat{C}^*_{H,t}. \tag{36}$$

Computing Planning Problem Wedges

We begin with computing the (distorted) Consumption Possibility Frontier of Figure 3 in growth rates. Recall the problem in levels (dropping the t subscript):

$$\begin{aligned}
\boldsymbol{C}(C^*; Y_C, Y^*_C, \delta) = \max_{\{C_H, C_F, C^*_H, C^*_F\}} & \; C(C_H, C_F) \\
\text{s.t. } & C_H + C^*_H \leq Y_C \\
& C_F + C^*_F \leq Y^*_C \\
& C^*(C^*_H, C^*_F) \geq C^* \\
& \frac{C_H}{C_F} = (1+\delta)^\theta \left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{C^*_H}{C^*_F}.
\end{aligned}$$

Now consider a normalized problem. In particular, we take $\{\bar{C}_H, \bar{C}_F, \bar{C}^*_H, \bar{C}^*_F\}$ and the associated $\bar{C}$ and $\bar{C}^*$ as arbitrary but fixed constants, and consider choosing an allocation *relative* to those benchmarks. Using

the above definitions in regard to normalization, we can rewrite this problem as

$$\begin{aligned}
\hat{\boldsymbol{C}}(\hat{C}^*; \hat{Y}_C, \hat{Y}^*_C, \widehat{1+\delta}) = \max_{\{\hat{C}_H, \hat{C}_F, \hat{C}^*_H, \hat{C}^*_F\}} \; & \hat{C}(\hat{C}_H, \hat{C}_F) && (\hat{P}) \\
\text{s.t. } & (1-\bar{\alpha}_C)\hat{C}_H + \bar{\alpha}_C \hat{C}^*_H \leq \hat{Y}_C && \nu \\
& \bar{\alpha}^*_C \hat{C}_F + (1-\bar{\alpha}^*_C)\hat{C}^*_F \leq \hat{Y}^*_C && \nu^* \\
& \hat{C}^*(\hat{C}^*_H, \hat{C}^*_F) \geq \hat{C}^* && \tilde{\boldsymbol{Q}}' \\
& \frac{\hat{C}_H}{\hat{C}_F} = (\widehat{1+\delta})^\theta \frac{\hat{C}^*_H}{\hat{C}^*_F}, && \zeta
\end{aligned}$$

where we indicate the Lagrange multipliers on the constraints on the right-hand margin and we define

$$(\widehat{1+\delta})^\theta \equiv (1+\delta)^\theta * \left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{\bar{C}_F}{\bar{C}_H}\frac{\bar{C}^*_H}{\bar{C}^*_F} = \frac{(1+\delta)^\theta}{(1+\bar{\delta})^\theta}.$$

Given an allocation $\{\hat{C}_H, \hat{C}_F, \hat{C}^*_H, \hat{C}^*_F\}$, we can compute $\widehat{1+\delta}$ directly from the final constraint and solve for the multipliers $\{\nu, \nu^*, \tilde{Q}', \zeta\}$ using the first-order conditions for Problem $\hat{P}$. Note that

$$\tilde{\boldsymbol{Q}}' = -\frac{\partial \hat{\boldsymbol{C}}}{\partial \hat{C}^*},$$

and represents the shadow cost of marginal *growth* in Foreign consumption in terms of Home consumption growth. If $\tilde{\boldsymbol{Q}} = -\partial \boldsymbol{C}/\partial C^*$ is the relative price in levels, we have

$$\tilde{\boldsymbol{Q}}' = -\frac{\partial \hat{\boldsymbol{C}}}{\partial \hat{C}^*} = -\frac{\bar{C}^*}{\bar{C}}\frac{\partial \boldsymbol{C}}{\partial C^*} = -\frac{\bar{C}^*}{\bar{C}}\tilde{\boldsymbol{Q}}.$$

Note that we can compute $\tilde{Q}'$ in any year from the above problem, but cannot compute $\tilde{\boldsymbol{Q}}$ without knowing the level of reference consumptions in each regions. However, we can use $\tilde{Q}'$ in any two periods to compute the ratio of $\tilde{\boldsymbol{Q}}$ across time.

Turning to the preferences side, note that the Planner's objective can be rewritten in growth rates – again, *given* fixed reference constants $\bar{C}$ and $\bar{C}^*$ – as follows:

$$\begin{aligned}
\sum_{t=0}^{T} \beta^t \left[\omega U(C_t) + U(C^*_t)\right] &= U(\bar{C}^*) \sum_{t=0}^{T} \beta^t \left[\omega \frac{U(\bar{C})}{U(\bar{C}^*)}\frac{U(C_t)}{U(\bar{C})} + \frac{U(C^*_t)}{U(\bar{C}^*)}\right] \\
&= (1-\sigma) U(\bar{C}^*) \sum_{t=0}^{T} \beta^t \left[\omega' U(\hat{C}_t) + U(\hat{C}^*_t)\right],
\end{aligned}$$

where the last line uses the functional form of U and defines $\omega' \equiv \omega \frac{U(\bar{C})}{U(\bar{C}^*)}$. Using the function $\hat{\boldsymbol{C}}$ defined

above, we can write the Planner's problem (up to a positive multiplicative constant) as:

$$\max_{\{\hat{C}_t^*\}} \sum_{t=0}^{T} \beta^t \left[\omega' U(\hat{\boldsymbol{C}}(\hat{C}_t^*)) + U(\hat{C}_t^*)\right].$$

The first order conditions are:

$$\omega' U'(\hat{C}_t)\tilde{\boldsymbol{Q}}_t' = U'(\hat{C}_t^*) \quad \forall t.$$

We define the Planner's risk sharing wedge in period t, denoted λ_t, by

$$\lambda_t \equiv \frac{U'(\hat{C}_t^*)}{\omega' U'(\hat{C}_t)\tilde{\boldsymbol{Q}}'} - 1.$$