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

August 26, 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 equilibrium pricing behavior of exporters.

*We thank Manuel Amador, Andy Atkeson, David Baqaee, Ariel Burstein, Luca Dedola, 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 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

[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 focus on consumption correlations (see e.g. Backus et al., 1992; Baxter and Crucini, 1995; Ohanian et al., 2018).

(MRS) 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 allows us to assess the welfare losses from imperfect 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.

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 (2013). 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. 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 surveys in Burstein and Gopinath (2012), Itskhoki (2021), 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 (or rest of the world, ROW), 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 $s^t = \{s_0, s_1, \ldots, s_t\} \in S^{t+1}$. 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 what follows, we suppress the history notation and time subscripts whenever convenient. In this section, we present the environment as if the two goods are freely traded, but allow for constant iceberg transportation costs in our empirical implementation.

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

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t U(C_t), \qquad \text{where} \qquad U(C) = \frac{C^{1-\sigma}}{1-\sigma},$$

with constant relative risk aversion $\sigma > 0$. C is 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 $\gamma \to 0$ representing complete home bias and $\gamma = 1/2$ denoting equal weights.

Foreign consumer have the same β and σ as Home, but enjoy utility over a different composite:

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

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 assume that the environment features *home bias* in preferences defined as $\gamma < 1 - \gamma^*$. The symmetric case with $\gamma, \gamma^* < 1/2$ is sufficient, but not necessary under this definition, which applies to general asymmetric regions.

The aggregate resource constraints for each good require

$$\begin{aligned} C_{Ht} + C_{Ht}^* &= Y_t, \\ C_{Ft} + C_{Ft}^* &= Y_t^*. \end{aligned} \tag{3}$$

This determines the set of feasible allocation of individual goods, while (1) and (2) quantify the associated aggregate consumption levels in the two regions of the world.

2.2 Efficient and distorted allocations

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

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

subject to the sequence of resource constraints (3). The entire Pareto frontier can be traced by varying the weight $\omega \geq 0$.

We can further split the planning problem into two sub-problems. The first concerns choosing

an optimal static allocation within a period, including Home consumption C, given a level of Foreign consumption C^*. The second stage chooses the optimal dynamic sequence of C_t and C_t^*.

Static Efficiency: Consider the following planning problem:

$$\mathbf{C}(C^*; Y, Y^*) \equiv \max_{\{C_H, C_F, C_H^*, C_F^*\}} C(C_H, C_F) \tag{5}$$
$$\text{subject to (3) and } \; C^*(C_H^*, C_F^*) \geq C^*.$$

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. As we vary C^*, we trace out the Consumption Possibility Frontier (CPF). This is the production possibility frontier for the technologies that take the two traded inputs and produce the two final consumption composites that enter the respective utility functions.

Letting $\tilde{Q}$ denote the multiplier on the final C^* constraint, the necessary conditions are:

$$\frac{\partial C}{\partial C_k} = \tilde{Q} \frac{\partial C^*}{\partial C_k} \qquad \text{for } \; k = H, F. \tag{6}$$

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

By taking the ratio of (6) for the two goods $k = H, F$, we can eliminate $\tilde{Q}$ and obtain the static optimality condition:

$$\frac{C_H}{C_F} = \left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{C_H^*}{C_F^*}, \tag{7}$$

which implies that the marginal rates of substitution across the two goods are equalized for Home and Foreign. Note that (7) features home bias rather than trade costs, which are absent under our assumptions in this section (but will be included in the empirical implementation).

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 constraint, $C_H^* = Y - C_H$ can be read from right to left. Similarly, the vertical axis measures C_F. Efficient points use all resources and satisfy (7), which implies the indifference curves between Home and Foreign are tangent, a geometric representation of (7). For example, point A in the figure depicts a particular efficient static allocation associated with a given weight ω. As we vary C^*, we trace out the set of efficient allocations, depicted by the dashed line connecting the two opposite corners, the *contract curve*.

Panel (b) of Figure 1 depicts the Consumption Possibility Frontiers (CPF). Problem (5) defines the undistorted CPF, which is the outer-most dashed frontier. This frontier is strictly concave to the origin when preferences differ due to home bias.[2] Concavity implies that the shadow cost $\tilde{Q}$ of providing Foreign with more consumption is increasing the greater is the Foreign's share of total resources.

If the Welfare Theorems hold, then $\tilde{Q}$ 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 consider possible distortions from efficiency 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. These wedges allow us to rationalize observed allocations as the outcome of a constrained planning problem, where the wedges represent additional constraints on the planner's choice set. In Section 3, we introduce a decentralized economy in which these distortions arise from particular frictions in goods and asset markets.

The Static Wedge: The *static wedge* δ breaks the equality of MRS in equation (7) such that:

$$\frac{\partial C^*/\partial C_H^*}{\partial C^*/\partial C_F^*} = (1+\delta)\frac{\partial C/\partial C_H}{\partial C/\partial C_F} \qquad \Leftrightarrow \qquad \frac{C_H}{C_F} = (1+\delta)^{\theta}\left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{C_H^*}{C_F^*}. \tag{8}$$

If $\delta \neq 0$, there are static gains from trade as Home and Foreign have different marginal rates of substitution between the Home and Foreign good. Thus δ represents an unrealized gain from a static exchange of Home for Foreign goods, where $\delta > 0$ results in an additional frictional home bias. In particular, δ is generated by different relative prices faced by consumers in each region. In Section 3, we show how various familiar models map into δ in a decentralized equilibrium.[3]

The point labelled "Data" in Figure 1a depicts a potential empirical deviation from static efficiency. As depicted, Home's MRS between Home and Foreign goods differs from Foreign's by a factor $1+\delta$ with $\delta > 0$, resulting in Home's consumption being frictionally biased towards the Home good, according to (8). As we vary the Pareto weight ω, we can trace out a *distorted*

[2]The concavity and its intuition is discussed in Samuelson (1949). In brief, 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 proportion and equal to the ratio of the endowments. 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, pushing the frontier to the northeast of a line and generating a CPF that is concave to the origin.

[3]As noted above, our empirical work allows for trade costs which also generate a wedge between C_H/C_F and C_H^*/C_F^*. However, we treat such costs as stable over time in the medium term and interpret year-to-year time variation in δ as reflecting the consequences of pricing-to-market frictions.

contract curve, along which the MRS distortion is held constant at $1+\delta$. This is the solid locus in the Edgeworth box.[4] In Section 3 we discuss why a fixed δ arises as a natural constraint in the second-best risk sharing problem in a wide range of standard models.

Formally, consider an extended planning problem that solves Problem (5) with the additional constraint (8):[5]

$$\mathbf{C}(C^*; Y, Y^*, \delta) \equiv \max_{\{C_H, C_F, C_H^*, C_F^*\}} C(C_H, C_F) \tag{9}$$
$$\text{subject to (3), (8) and } \; C^*(C_H^*, C_F^*) \geq C^*,$$

where we extend the notation for the $\mathbf{C}(C^*; Y, Y^*, \delta)$ to include the static wedge δ. In words, now $\mathbf{C}$ denotes the maximum consumption of Home given consumption of Foreign C^*, the aggregate resources (Y, Y^*), and the static wedge δ. Similarly, $\tilde{Q} = \tilde{Q}(C^*; Y, Y^*, \delta)$ is the Lagrange multiplier on the C^* constraint in this problem, which by virtue of optimality conditions satisfies $\tilde{Q} = -\partial \mathbf{C}/\partial C^*$ and equals the MRT along the distorted frontier $\mathbf{C}(C^*; Y, Y^*, \delta)$. The earlier solution for the first-best problem is nested as a special case with $\delta = 0$.

Figure 1b depicts with the inner solid line the *distorted* Consumption Possibility Frontier frontier $C = \mathbf{C}(C^*; Y, Y^*, \delta)$ given the static wedge δ. At the end points of the CPF, the distorted solid frontier coincides with the undistorted dashed frontier.[6] 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 distorted frontier lies interior to the first best. Moreover, it need not be concave: with the static distortion, the regional consumption intensities are not optimized, and may be inferior to the equal ratio benchmark that would deliver a linear frontier.

The Dynamic Wedge: With $C = \mathbf{C}(C^*; Y, Y^*, \delta)$ obtained from the distorted static problem (9) for a given static wedge δ, we can write the inter-temporal Planning problem as:

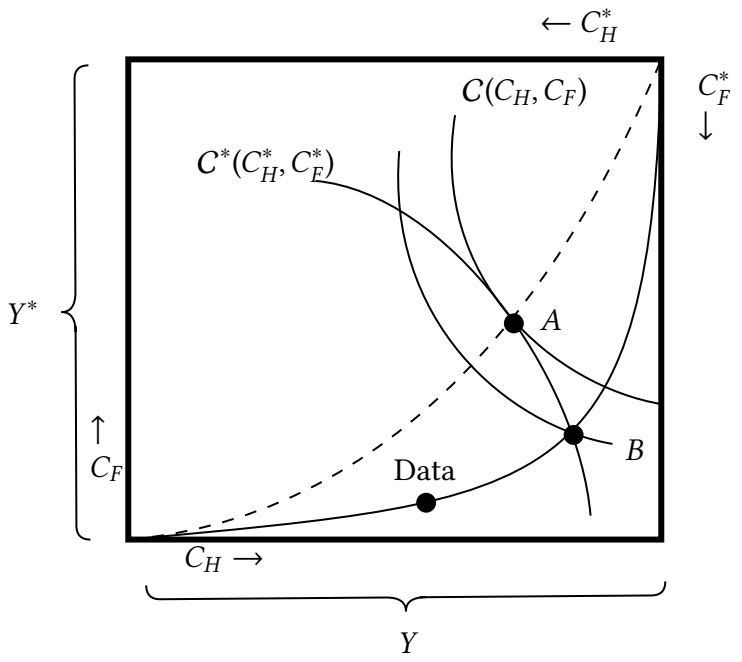
$$\max_{\{C_t^*\}} \mathbb{E} \sum_{t \geq 0} \beta^t \left[\omega U(\mathbf{C}(C_t^*; Y_t, Y_t^*, \delta_t)) + U(C_t^*) \right]. \tag{10}$$

[4] Aguiar et al. (2025) and Itskhoki and Mukhin (2025) use such an Edgeworth box to study the impact of tariffs which, just like pricing to market frictions, result in a static wedge $\delta > 0$ and a distorted contract curve.

[5] While useful as a conceptual link to the planner's problem, note that the additional constraint (8) in problem (9) associated with the static wedge δ reduces the constraint set to a *singleton*, when the C^* constraint is evaluated as an equality. This is clear in the Edgeworth box in Figure 1a. 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 where it intersects the δ-distorted contract curve (while aggregate resources Y and Y^* pin the size of the Edgeworth box).

[6] 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 (8) when $C_H/C_F \to Y/Y^*$. The same logic applies at the vertical intercept.

(a) Edgeworth Box



(b) Consumption Possibility Frontier

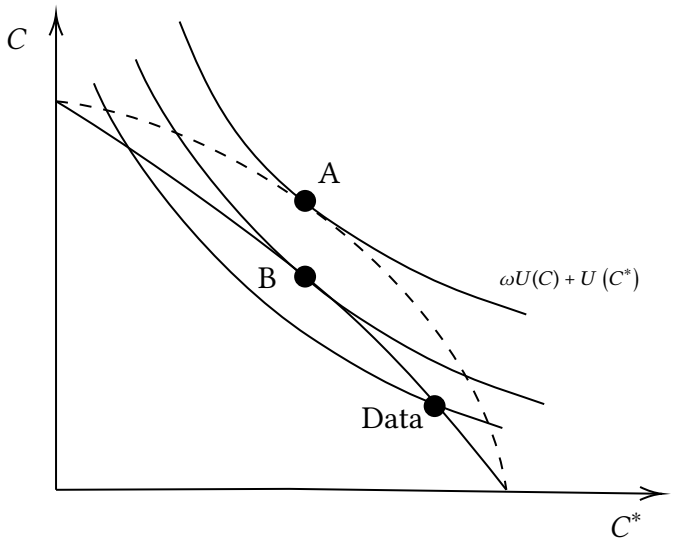


Figure 1: Geometric Represenation of Constrained Efficient Allocations

Note: Panel (a) contains the Edgeworth box for a given (Y, Y^*). Point A represents an efficient allocation. Point B lies along the distroted contract curve defined by a constant static wedge δ, which determines the difference in slopes between the two curves. "Data" represents a particular observed allocation, which has the same δ as point B but a relative shift in consumption toward Foreign. Panel (b) contains the Consumption Possibility Frontier (CPF) for a given (Y, Y^*). The lines convex to the origin represent the planner's indifference curves. The dashed frontier is the first-best CPF, for which $\delta = 0$. The inner, solid frontier represents the distorted CPF given a non-zero static wedge δ. As in Panel (a), point A represents a first-best allocation. Point B represents the second best risk-sharing allocation, at which the planner's indifference curve is tangent to the distorted CPF. As in panel (a), the point "Data" represents a potential observed allocation. The difference in slope between the planner's indifference curve and the distorted CPF is governed by the dynamic wedge λ.

The first-order condition for a particular history s^t (suppressed in the notation) is $\omega U'(\mathbf{C})\cdot\partial\mathbf{C}/\partial C^* + U'(C^*) = 0$, or, equivalently, using the property that $\tilde{Q} = -\partial\mathbf{C}/\partial C^*$ we have:

$$\frac{U'(C^*)}{\omega U'(C)} = \tilde{Q}. \tag{11}$$

This equality states that the marginal rate of substitution (MRS) between Foreign and Home consumption equals the shadow price (MRT) of Foreign consumption relative to Home. Note that we can use the optimality condition (11) for different values of $\omega \geq 0$ to trace out the δ-distorted contract curve depicted with a solid line in the Edgeworth Box in Figure 1a.

The intuition for equation (11) is that the Planner prefers to tilt consumption towards the region for which the consumption composite is relatively cheap in a given state. That is, when $\tilde{Q}$ is relatively small, it is a period in which Home's composite is relatively cheap to construct from (Y, Y^*), given static distortion δ, and vice versa when $\tilde{Q}$ is relatively large. If the composites were identical (that is, when $\gamma = 1-\gamma^*$), then $\tilde{Q} = 1$ in every period for $\delta = 0$. In that case, the ratio of marginal utilities should be constant in response to changes in output Y and Y^*.

The dynamic or inter-temporal wedge is the deviation from condition (11). We define:

$$1 + \lambda_t \equiv \frac{U'(C_t^*)}{\omega U'(C_t)}\frac{1}{\tilde{Q}_t}. \tag{12}$$

The inter-temporal wedge reflects an implicit shift in the Pareto weight that rationalizes the observed allocation. This implies that if $\lambda_t \neq \lambda_0$, there is an inter-temporal reallocation between (C_0, C_t) and (C_0^*, C_t^*) that represents a Pareto improvement.[7]

In Figure 1 Panel (b) we consider the optimal choice of $\{C, C^*\}$ given a Pareto weight ω. Point A represents the first-best allocation, when $\lambda = \delta = 0$. The second-best allocation maximizes the weighted sum of utility subject to the static market distortion. The resulting allocation is depicted as point B in both panels of Figure 1. The tangency at B in panel (b) implies that $\frac{U'(C^*)}{\omega U'(C)} = \tilde{Q}$ where $\tilde{Q} \equiv \tilde{Q}(C^*; Y, Y^*, \delta)$ is the distorted MRT and is the slope of the interior consumption frontier. Note that for general δ, the distorted frontier need not be concave, and hence $\tilde{Q}$ can be a non-monotonic function of C^*. The point labelled "Data" reflects both a static wedge $\delta > 0$ and a dynamic wedge $\lambda < 0$. The static wedge represents the inward shift of the CPF. The dynamic wedge represents the lack of tangency between the weighted sum of utilities and the CPF. The difference in slopes between the CPF and the indifference curve is λ.

[7] In particular, if $\lambda_t > \lambda_0$, consider increasing C_t^* at the margin by $\Delta > 0$ and decreasing C_0^* by $-\beta^t U'(C_t^*)/U'(C_0^*)\Delta$, such that Foreign welfare is unchanged. Feasibility implies that $\Delta C_t = -\tilde{Q}_t\Delta$ and $\Delta C_0 = \tilde{Q}_0\beta^t U'(C_t^*)/U'(C_0^*)\Delta$. Using the definition of λ_t, the net increase in Home welfare is proportional to $(\lambda_t - \lambda_0)\Delta > 0$. By convention, we adopt $t = 0$ as the baseline period and normalize $\lambda_0 = 0$, as its value is not separately identified from the planner's weight ω.

2.3 Identification

We now turn to identification of δ, λ, and $\tilde{Q}$ from observable data on real consumption and output in the 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 the discussion of how we calibrate the preference parameters σ, θ, γ, γ^* and ω. 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 in Figure 1 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 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 $\delta > 0$. For concreteness, we focus here on the solution that enhances home bias of the regions, that is $\delta > 0$ which corresponds to the area below the efficient contract curve.[8] Finally, if the two indifference curves never intersect, then we have violated feasibility (i.e., the empirical allocation is outside the planner's CPF). This is evidence of a mis-specification of preferences and/or a mis-measurement of the observed aggregates. In Section 5, we discuss when such pathologies arise in the data.

In Figure 1, in both panels (a) and (b), we depict a point labelled "Data." In panel (a), this is the solution $\{C_H, C_F, C_H^*, C_F^*\}$ to the four equations mentioned above, and it pins down the static wedge $\delta > 0$. In panel (b), we also indicate the observed data on aggregate consumption $\{C, C^*\}$. Given δ, we can turn to the planner's sub-problem (9) of delivering maximum C given C^* in order to infer $\tilde{Q}$. Relative to B, the Data is biased toward Foreign consumption. In the figure, the gap between the slope of the indifference curve and the slope of the distorted frontier at the "Data" point represents that period's dynamic wedge λ_t, which is negative in the case depicted.

Over time, the Data point moves around inside both the Edgeworth box and the first-best CPF, resulting in the time series variation in the implied wedges δ and λ. So long as the resulting Data point is inside the first-best CPF, there exists a unique solution for (λ, δ) with $\delta > 0$.[9]

[8] In our empirical estimation in Section 5, we bring additional data on trade flows to identify the value of δ_t in the baseline year, and assume the static distortion is biased in the same direction for a country over time, which in practice results in an enhancement of home bias in all cases but one.

[9] Strictly speaking, the dimensions of the Edgeworth box and location of the CPF move around as well with variation in output (Y, Y^*), which we take into account when recovering the wedges.

The Cole-Obstfeld case The identification procedure described above and implemented in the data in Section 5 is general and applies for any calibrated model parameters. However, the system of equations used to invert the data is highly non-linear and, in general, does not feature a tractable closed-form solution. For illustration, we focus here on the Cole-Obstfeld case with $\sigma = \theta = 1$ which admits a more tractable inversion for recovering wedges and planner's MRT from observable data on consumption and output. We have the following characterization (with the complete non-linear solution provided in Appendix A):

Lemma 1. *Suppose that $\sigma = \theta = 1$:*

(i) For fixed λ and δ, the following ratios are constant as Y and Y^* vary:

$$\frac{C}{Y^{1-\gamma}(Y^*)^{\gamma}}, \qquad \frac{C^*}{Y^{\gamma^*}(Y^*)^{1-\gamma^*}}, \qquad \frac{\tilde{Q}}{(Y/Y^*)^{1-\gamma-\gamma^*}};$$

(ii) For fixed δ, Y, and Y^*, an increase in λ implies an increase in C and decrease in C^*;

(iii) For fixed λ, Y, and Y^*, an increase in $|\delta|$ implies a decrease in both C and C^*.

This lemma shows that C and C^* expand with both outputs, yet exhibiting home bias, and the planner's MRT depends on the ratio of outputs. Furthermore, when the two wedges remain constant, i.e., $\Delta\lambda = \Delta\delta = 0$, the following relationships hold in log changes:

$$\Delta \log(C/C^*) = \Delta \log \tilde{Q} = (1 - \gamma - \gamma^*)\Delta \log(Y/Y^*).$$

Despite this relationship under constant wedges, an implication of the lemma is that — in the presence of home bias — there is no simple relationship between consumption and output correlations.[10] As a result, the inference about distortions must rely on the full structure of the model combined with the full covariance matrix of consumption and output in both region.

In the Cole-Obstfeld case, consumptions can be scaled with the Cobb-Douglas mixtures of the two outputs, as shown in the lemma. Then the wedges can be inferred from the joint movements in these ratios, depending on whether they move in concert or in opposite directions. A reduction in both consumption ratios — a movement towards the origin in Figure 1b — implies an increase in δ, while a movement in opposite directions — along a distorted CPF — implies a change in λ.

[10] Assuming symmetry, $\gamma = \gamma^*$ and $\text{var}(\Delta \log Y) = \text{var}(\Delta \log Y^*)$, the lemma implies that:

$$\text{cov}(\Delta \log C, \Delta \log C^*) = 2\gamma(1-\gamma)\text{var}(\Delta \log Y) + [1 - 2\gamma(1-\gamma)]\text{cov}(\Delta \log Y, \Delta \log Y^*),$$

such that consumption correlation is generally larger than output correlation, but this is not the case in general outside this symmetric case or when $\sigma\theta \neq 1$ (see discussion in Heathcote and Perri, 2014).

2.4 Welfare loss

The two distortions (λ, δ) introduced above allow us to rationalize any observed allocation of output to consumption across regions as an outcome of a constrained planner's problem. In particular, in a given period t and state s^t, a fictitious planner solves:[11]

$$\mathbb{W}(\bar{U}, \delta) \equiv \max_{\{C_H, C_F, C_H^*, C_F^*\}} \omega U(C(C_H, C_F)) + U(C^*(C_H^*, C_F^*)) \tag{13}$$
$$\text{subject to (3), (8) and } \; U(C(C_H, C_F)) = \overline{U}$$

for a given $\overline{U} \in [U(C(0,0)), U(C(Y, Y^*))]$. Recall that constraint (3) is the resource condition and constraint (8) imposes a static wedge δ on the allocation and this constraint is slack when $\delta = 0$.

The final constraint in (13) imposes a dynamic distortion by forcing the planner to provide a particular level of utility to Home independent of parameter ω. Denoting with $\omega\lambda$ the Lagrange multiplier on this constraints, we can show that its value is increasing in the level of the required Home utility $\overline{U}$, in any given period and state and for a given static wedge δ. Note that the Lagrangian associated with (13) contains $\omega(1 + \lambda)U$, with ωU from the objective and $\omega\lambda U$ from the constraint. The latter part reflects the dynamic distortion to risk sharing arising from the final constraint on the planner. In particular, the first-order optimality conditions for this problem result in (12) which we use to identify the dynamic wedge λ.

Finally, observe that the constraint set in (13) — just like in (9) — is a singleton, and therefore this planning problem is merely a medium to define the static and dynamic wedges. An additional consequence of this representation is a characterization of the welfare loss associated with the two distortions. Given the monotonicity of the Lagrange multiplier λ in the constraint $\overline{U}$ for a given static wedge δ, we denote the inverse mapping from (λ, δ) to $\overline{U}$ as $\bar{U}(\lambda, \delta)$. Then we define the within-state welfare loss from the two distortions as:

$$\mathbb{L}(\lambda, \delta) = -[\mathbb{W}(\bar{U}(\lambda, \delta), \delta) - \mathbb{W}(\bar{U}(0,0), 0)], \tag{14}$$

where $\bar{U}(0,0)$ corresponds to the utility U allocated to Home by the unconstrained planner. Irrespective of the structural data-generating process, any observed allocation (C, C^*, Y, Y^*) can be mapped into the pair of distortions (λ, δ) with the associated welfare loss $\mathbb{L}(\lambda, \delta)$ given by (14). An important property of our static-dynamic wedge decomposition is that the two distortions are orthogonal in the following formal sense (see Appendix A):

[11] Formally, start with a sequence problem to maximize the intertemporal objective (4) subject to a sequence of constraints for each time period t and state s^t and given a sequence of distortions $\delta_t(s^t)$ and $\overline{U}_t(s^t)$. As all constraints are within period, this problem separates into a sequence of state-by-state allocation problems given by (13) and linked by a common planner's weight ω which captures the risk-sharing objective across time and states.

Proposition 1. *Irrespective of the value of the static distortion δ, the welfare loss from the dynamic distortion is minimized when $\lambda = 0$. Formally, for any δ,* $\arg\min_\lambda \mathbb{L}(\lambda, \delta) = 0$.

This is an unusual property in the general theory of the second best, where the presence of a distortion on one margin calls for optimal distortions of other margins, as for example is the case with general markup wedges.[12] This is not the case under our (λ, δ) representation, justifying our focus on λ as a measure of the risk-sharing distortion independent from the nature and extent of static inefficiencies in the economy captured by δ. More precisely, the intertemporal reallocation of consumption due to δ is entirely captured by $\tilde{Q}$. Given this distorted MRT, the planner has no additional incentive to deviate from (11), and chooses $\lambda = 0$. In other words, the identification of $\lambda \neq 0$ necessarily implies the presence of a distortion to intertemporal risk sharing.

A corollary of Proposition 1 is that a second-order expansion of the welfare loss around the efficient allocation characterized by $\lambda = \delta = 0$ can be represented by a quadratic form in λ^2 and δ^2 without a cross-product term. That is, the second-order welfare loss is additively separable in the static and dynamic wedges. If this were not the case, the value of λ that minimizes the welfare loss for some small $\delta \neq 0$ must be different than zero, contradicting Proposition 1.

We provide a general derivation of the second order approximation to the welfare loss function in Appendix A, and display here the simplified expression for the symmetric case with $\omega = 1$ and $\gamma = \gamma^*$. We have:

$$\mathbb{L}(\lambda, \delta) = \theta\gamma(1-\gamma)\left[\kappa\lambda^2 + \frac{1}{4}\delta^2\right], \tag{15}$$

where $\kappa \equiv \frac{1}{1+4\gamma(1-\gamma)(\sigma\theta-1)}$, and the loss is expressed in units of home flow marginal utility. For relevant parameter values, welfare loss from distortions to static and dynamic international exchange is increasing in the value of openness γ and in the elasticity of substitution θ which makes the given wedges more consequential for the resulting allocations. Formula (15) describes the *Harberger triangles* associated with the dynamic and static wedges, respectively.[13]

Importantly, the welfare loss characterized in (15) is the property of the observed allocation rather than its representation as the planner's solution. In other words, it applies as a characterization of welfare losses in any decentralization that is consistent with the data on consumption and output. Therefore, the abstraction of the constrained planner's problem (13) and the two associated wedges (λ, δ) is useful for mapping the observed data into a quantitative measure of

[12] Note that the (λ, δ) representation is not unique, and the observed allocation as well as the associated welfare loss can be rationalized with alternative sets of wedges. For example, another natural way to quantify distortions are the product-region-specific markup wedges following Baqaee and Farhi (2020). We discuss the mapping between the two sets of wedges in the Appendix. See also Bhandari et al. (2025), Dávila and Schaab (2025) and Baqaee and Burstein (2025) for the discussion of welfare in frictional economies with heterogenous agents.

[13] In Appendix A, we show that the welfare loss in (15) from the dynamic wedge λ can be represented as $\frac{1}{2}\,\mathrm{d}\log C\,\mathrm{d}\lambda$, where $\mathrm{d}\log C = 2\gamma(1-\gamma)\kappa\theta\,\mathrm{d}\lambda$ when expanded around $\lambda = 0$.

departure from efficiency. In particular, the planner's dynamic wedge λ offers a measure of welfare losses from inefficient international risk sharing, given the arbitrary distortions in the static allocation captured by δ. Importantly, there is no such property for the Backus-Smith wedge calculated using the observed real exchange rate, as we discuss in Section 3.

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.

3.1 Household Behavior

The households' problem in the decentralized equilibrium is standard. We therefore relegate details to the appendix and flag key aspects in the text. The representative Home household faces a sequence of prices $\{P_{H,t}(s^t), P_F H, t(s^t)\}$ for the Home and Foreign goods in Home's currency. With the CES consumption aggregator (1), the ideal price index for the Home composite is therefore (suppressing time and history indexing when convenient):

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

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 (in Section 3.3) 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 wealth w, the representative Home agent's problem is:

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

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):

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

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 the budget constraint which is in a one-to-one (monotonically decreasing) mapping with the initial wealth w.

The Foreign agent's problem is symmetric. It faces prices $\{P^*_{H,t}(s^t), P^*_{F,t}(s^t)\}$ for the Home and Foreign good, with an associated CES ideal price index $P^* = \left(\gamma^* P_H^{*1-\theta} + (1-\gamma^*)P_F^{*1-\theta}\right)^{1/(1-\theta)}$. Without loss, these prices are expressed in Home's currency, as well.[14] The Foreign agent faces prices of Arrow-Debreu securities $\Lambda^*_t(s^t)$. Given some initial wealth w^* and transfers $T^*_t(s^t)$, the Foreign household's problem is:

$$\begin{aligned} &\max_{\{C^*_t(s^t)\}} \sum_{t=0}^{\infty} \sum_{s^t} \pi_t(s^t) U(C^*_t(s^t)) \\ &\text{subject to:} \quad w^* \geq \sum_t \sum_{s^t} \Lambda^*_t(s^t) \left(P^*_t(s^t) C^*_t(s^t) - P^*_{F,t}(s^t) Y^*_t(s^t) - T^*_t(s^t) \right). \end{aligned}$$

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

$$\frac{C^*_H}{C^*} = \gamma^*\left(\frac{P^*_H}{P^*}\right)^{-\theta}, \qquad \frac{C^*_F}{C^*} = (1-\gamma^*)\left(\frac{P^*_F}{P^*}\right)^{-\theta} \tag{18}$$

and

$$\beta^t \pi_t(s^t) U'(C^*_t(s^t)) = \xi^* \Lambda^*_t(s^t) \mathcal{E}(s^t) P^*_t(s^t), \tag{19}$$

where ξ^* is the Lagrange multiplier on Foreign's budget constraint decreasing in the initial wealth w^*.

Measuring Implied Prices: Given measured allocations in each region, we can use the first-order conditions to infer prices faced by Home and Foreign agents. In particular, given a sequence

[14] We can denote the Foreign currency price sequence by $\{\tilde{P}^*_{H,t}(s^t), \tilde{P}^*_{F,t}(s^t)\}$. Letting $\mathcal{E}_t(s^t)$ denote the exchange rate, expressed as Foreign currency per unit of Home currency, we have $P_{k,t}(s^t) = \tilde{P}_{k,t}(s^t)/\mathcal{E}_t(s^t)$, for $k = H, F$. Since only relative prices will matter for the household's problem, the nominal exchange rate is not relevant for the allocation.

of observed allocations $\{C_{H,t}, C_{F,t}, C_t\}$, we can use (16) to infer the sequence of relative 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). Similarly, from (17) 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). Similarly, from (18) and (19) we can infer $\{P^*_{H,t}, P^*_{F,t}, P^*_t\}$ up to a normalization (for example $P^*_{F,t} = 1$ for all t and s^t).

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

$$\mathcal{Q} \equiv \frac{P^*}{P} = \frac{P^*_H}{P_H}\left(\frac{\gamma^* + (1-\gamma^*)\left(P^*_F/P^*_H\right)^{1-\theta}}{(1-\gamma) + \gamma\left(P_F/P_H\right)^{1-\theta}}\right)^{1/(1-\theta)}. \tag{20}$$

The first term, P^*_H/P_H is the deviation from the Law of One Price (LOP) for the Home good. This reflects that consumers in the two countries may pay a different amount for the same good. The term in parentheses varies if the *relative* prices in each country move independently from each other.

3.2 Implied Wedges

The ratio P^*_H/P_H is the deviation from 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:

$$P_H = (1+\mu_H)P^*_H \qquad \text{and} \qquad P_F = (1+\mu_F)P^*_F. \tag{21}$$

Using (17), (19) and (20), we define the "Backus-Smith" wedge as:

$$1 + \psi_t(s^t) \equiv \frac{\Lambda^*_t(s^t)}{\Lambda_t(s^t)} = \frac{\xi}{\xi^*}\frac{U'(C^*_t(s^t))}{U'(C_t(s^t))}\frac{1}{\mathcal{Q}_t(s^t)}. \tag{22}$$

This is the wedge in private risk-sharing. The first expression in (22) 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 (17) and (19) and the definition of the real exchange rate (20) to relate ψ to 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.

The constant 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 $\omega \equiv \xi^*/\xi$ 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:

Proposition 2. *(i) The static wedge defined in* (8) *is related to deviations in the Law of One Price in* (21) *by:*

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

(ii) The dynamic wedge defined in (12) *is related to the Backus-Smith wedge of* (22) *by*

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

Note that the pair of LOP deviations (μ_H, μ_F) in the decentralized allocation map into a single static distortion in the Planner's problem, δ. For relative consumption allocations, it is sufficient to know the wedge between relative prices in the two markets, that is:

$$1 + \delta = \frac{P_H/P_F}{P_H^*/P_F^*} = \frac{1 + \mu_F}{1 + \mu_H},$$

as follows from combining the respective static first-order conditions for Home and Foreign and the definitions (21) and comparing to (8). Proportional changes in both $1 + \mu_H$ and $1 + \mu_F$ do not affect the static wedge δ, yet they can be identified in the decentralized allocation from the real exchange rate Q, which is an additional observable. For example, doubling both domestic prices for given values of the relative prices P_H/P_F and P_H^*/P_F^* halves the real exchange rate, as can be seen from (20), without changing the static wedge δ and hence the relative consumption allocation in the goods market.

The second result in the proposition establishes that the gap between the Planner's dynamic wedge λ and the Backus-Smith wedge ψ is equivalent to the gap between the observable equilibrium real exchange rate Q and the unobservable technological marginal rate of transformation between the two consumption composites $\tilde{Q}$. Our next results show that when the Welfare Theorems hold, the two variables are the same, $\tilde{Q} = Q$. Furthermore, to the extent the Welfare Theorems fail, the private and social measures of MRT, Q and $\tilde{Q}$, can be distinct, and so can be the two dynamic wedges.

We observe from Part (i) of Proposition 2 that the static distortion disappears if the deviations from the law of one price are 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 prices may differ. This case allows for a simple characterization of the gap between Q and $\tilde{Q}$ and, hence, between λ and δ:

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

$$\frac{1+\lambda}{1+\psi} = \frac{Q}{\tilde{Q}} = 1+\mu. \tag{23}$$

This Proposition emphasizes a direct link between the planner's dynamics wedge λ, the decentralized risk-sharing (Backus-Smith) wedge ψ, and the difference in the level of prices across regions μ. Importantly, this simple relationship (23) holds only when the LOP of one price deviations are equal across goods, $\mu_H = \mu_F$, and therefore there is no planner's static distortion, $\delta = 0$, according to Proposition 2. In other words, this works when all prices are uniformly higher in one of the regions without distorting the relative prices across regions.[15] Under these circumstances, the proposition shows that the real exchange rate is a shifted version of the planner's MRT by the markup-induced difference in the price levels: $Q = (1+\mu)\tilde{Q}$. By consequence, this drives a gap between the planner's and the decentralized risk-sharing wedges. Intuitively, μ is the non-technological contributor to the price level difference across regions, which affects the private risk-sharing calculations, but not for the planner.

Proposition 3 has the following immediate 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. In this case, the Welfare theorems hold for the goods market implying $Q = \tilde{Q}$, i.e., no disagreement between decentralized aggregate prices and the planner's marginal rate of transformation. By consequence, the only source for dynamic inefficiency λ are the risk-sharing frictions ψ in the decentralized allocation.

Second, even when the private risk-sharing (Backus-Smith) condition holds perfectly, $\psi = 0$, we can have $\lambda = \mu \neq 0$. In this case, frictional price level differences create a gap between the

[15] The prove of the proposition shows how, in the absence of the static wedge, the algebra of the planner's problem is similar to the algebra of the decentralized allocation. This ceases to be the case when $\mu_H \neq \mu_F$ and $\delta \neq 0$, resulting in a substantially more complicated mapping from (ψ, μ_H, μ_F) to λ. A fortiori, the presence of $\mu_H \neq \mu_F$ leads to a gap between Q and $\tilde{Q}$ and, hence, between λ and ψ.

planner's MRT $\tilde{Q}$ and the real exchange rate Q, making privately-optimal risk sharing inefficient from the planner's perspective. This leaves dynamic welfare gains on the table even when all households are on their undistorted Euler equations with $\psi = 0$.

Finally, for any Backus-Smith wedge ψ, we can have $\lambda = 0$ if μ offsets ψ. Specifically, by Proposition 3, we have that:

$$(1 + \mu)(1 + \psi) = 1 \qquad \Rightarrow \qquad \lambda = 0,$$

even though in general $Q \neq \tilde{Q}$ in this case. In other words, frictional differences in the price levels can act to offset frictions in risk sharing, and in this limiting case perfectly so. While abstractly this last case may appear knife-edge, we argue below that this case may offer a better point of approximation to the real world than the more intuitive cases with either $\mu = 0$ or $\psi = 0$. This is intriguing because this requires correlated movements in the goods market and asset market frictions. In order to investigate this possibility further, we introduce microfoundations for both frictions.

3.3 Equilibrium Wedges

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:

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

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.

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. Suppose Foreign consumers face prices $\{P_H^*, P_F^*\}$, now denominated in units of Foreign currency, and let $\mathcal{E}$ denote 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 $\mathcal{E}P_F = (1+\mu_F)P_F^*$.

Producer Currency Pricing: 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 literature (see, for example, Obstfeld and Rogoff (1995)). 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, $\mu_H = \mu_F = 0$.

Local Currency Pricing: 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, Devereux and Engel (2002)). Under this pricing protocol, the law of one price deviation will be determined by the nominal exchange rate. In particular, suppose that Home sets prices P_H in the Home market and sets the Foreign-currency price P_H^* in the Foreign market. Suppose $P_H^* = \overline{\mathcal{E}}P_H$, where, for example, the literature typicall assumes $\overline{\mathcal{E}}$ is the expected exchange rate conditional on information 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 $P_F = P_F^*/\overline{\mathcal{E}}$. All such prices are fixed relative to nominal exchange rate realizations. Note that the relative prices in each market are the same: $P_H/P_F = P_H^*/P_F^*$. That is, under LCP, we have $1+\mu_H = 1+\mu_F = \mathcal{E}/\overline{\mathcal{E}}$, the relative price wedges move in lock-step and, from Proposition 3 there is no static wedge.

Using the equilibrium characterization and Propositions 2 and 3, we have the following pricing protocol taxonomy:

Proposition 4. *Under PCP, we have $\delta = 0$ and $\lambda = \psi$. Under LCP, we have $\delta = 0$ and*

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

In models where LOP holds, including under PCP, the Backus-Smith wedge ψ equals the planner's risk-sharing wedge λ. Under LCP, LOP of price violations move in concert, maintaining undistorted relative prices across markets $\delta = 0$, yet the planner's wedge and the Backus-Smith wedge differ. To the extent the Backus-Smith wedge ψ is the main driver of the exchange rate $\mathcal{E}$, the planner's risk-sharing wedge λ may be significantly less volatile than ψ. Indeed, the variation in the exchange rate drives correlated LOP deviations μ, which offset the effect of ψ on the planner's risk-sharing wedge, according to Proposition 3. In contrast, exchange rate fluctuations driven by other shocks result in LOP deviations under LCP which translate into the planner's risk-sharing wedge λ even when $\psi = 0$. Models of pricing-to-market with endogenous markups as in Atkeson and Burstein (2008) lie in between the two limiting cases of PCP and LCP. Alternative pricing protocols such as dominant-currency pricing (DCP; Gopinath et al., 2020) distort relative prices across countries generally resulting in variation in both δ and λ. In our empirical analysis, we study the joint statistical properties of the data-inferred wedges δ, λ, ψ and μ, and discuss the structural models that are most likely to give rise to such wedges.

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^*.[16] 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 5. *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.

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, $\partial Q_t / \partial \delta_t = 0$ 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,

[16] This assumes that the consumption possibility frontier – and hence Q_t – is well behaved, ensuring there is a unique C_t^* that satisfies this equation.

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 (**??**), 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 (**??**). 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 6. *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 104 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 104 countries. The global growth in X is then $G_{X,t} = \sum_{j\in J} wt_{j,Y}\widehat{X}_{j,t}$. For country i, the growth in the ROW is computed as

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

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

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

where $Q_{i,j,t}$ is the real exchange between i and j in year t relative to 2019.[17] Appendix B provides

[17] The mean correlation of the resulting series with the real effective exchange rate from the WDI is 0.84 and the

further details.

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

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

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:

$$\hat{Y}_t \equiv \frac{Y_t}{\bar{Y}} = \frac{A_{H,t} + A_{H,t}^*}{\bar{Y}} = (1-\bar{\chi})\hat{A}_H + \bar{\chi}\hat{A}_H^* \tag{24}$$
$$\hat{Y}_t^* \equiv \frac{Y_t^*}{\bar{Y}^*} = \frac{A_{F,t} + A_{F,t}^*}{\bar{Y}^*} = \bar{\chi}^*\hat{A}_F + (1-\bar{\chi}^*)\hat{A}_F^*, \tag{25}$$

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.

relative standard deviation is about 1.4.

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:

$$\widehat{C}_{k,t} = \widehat{A}_{k,t} \times \frac{\widehat{C}_t}{\widehat{A}_t}, \quad k = H, F, \tag{26}$$

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 v, v^* and two distortions η, λ:

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

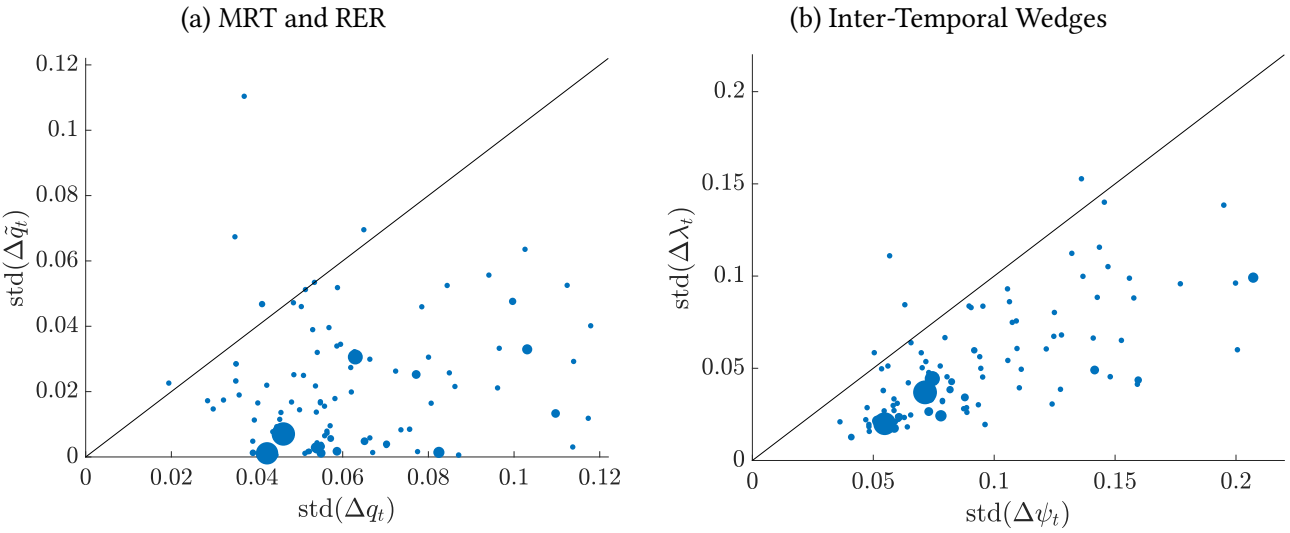
Notice that Lagrange multipliers and Pareto weights enter in level, not changes, and are not separately identified from base-year quantities $\bar{Y}, \bar{Y}^*, \bar{C}, \bar{C}^*$. Because the values of the latter are the same in all periods, however, the changes relative to the base year reveal time variation in shadow prices $\hat{v}, \hat{v}^*$ 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

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

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 v, v^* 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.[18]

[18] 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 2: Relative Volatilities



6 Empirical Results

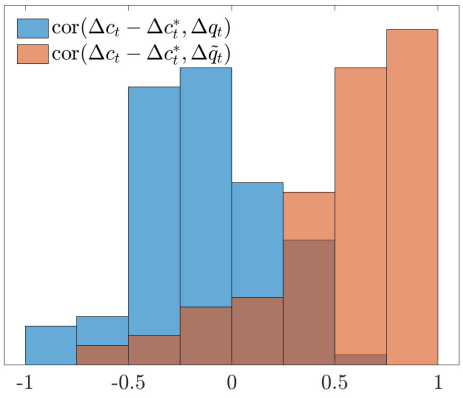
We begin with a comparison of the observed real exchange rate Q_t with our planner's *MRT* $\tilde{Q}_t$. Letting Δq_t denote the log of the first difference of Q_t, and similarly for $\Delta\tilde{q}_t$ and $\tilde{Q}_t$, Figure 2a provides a scatter plot of the relative volatilities. Each point is 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 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.

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 3, 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 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. Panel (b) is a scatter plot, with each point again representing a country (with Foreign a composite of the rest of the world), showing that the correlation with the planner's *MRT* is almost uniformly higher than the Backus-Smith correlation.

Figure 3 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,

Figure 3: Consumption Differences and the Real Exchange Rate

(a) Histogram of Correlations



(b) Relative Correlations

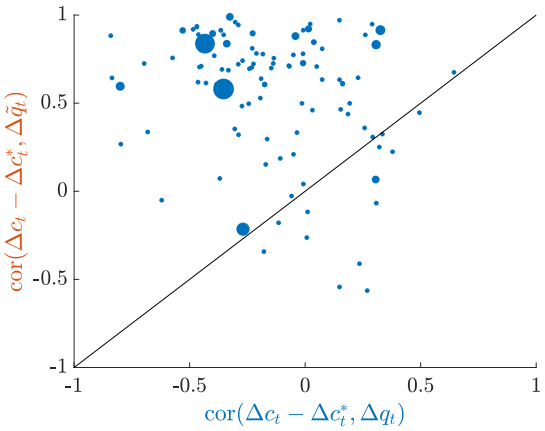
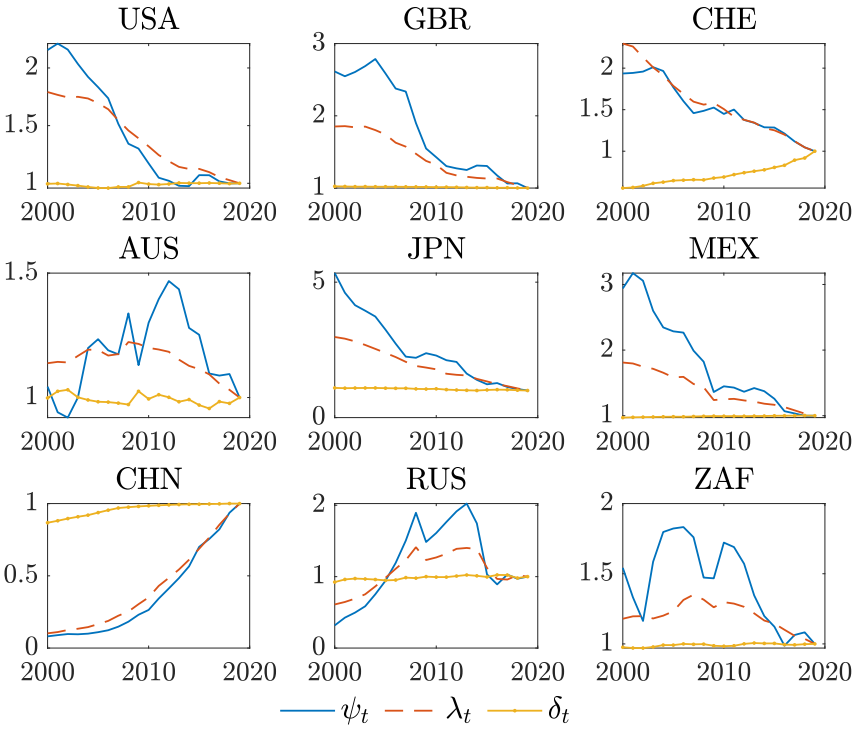


Figure 4 depicts the path of the three wedges, λ_t, ψ_t, and δ_t, for a selection of countries. Panel (a) is in levels and Panel (b) is in annual log first differences. Three facts about wedges stand out from the figure. First, the static wedge δ_t is surprisingly stable in time for most countries. This is despite large deviations from efficient allocation in the base year $\bar{\delta} \neq 0$. Second, the inter-temporal wedge λ_t is mostly determined by relative consumption at lower frequencies and therefore, has a clear trend in most countries. This trend is also usually in line with the low-frequency dynamics of the Backus-Smith wedge ψ_t and suggests poor risk sharing of permanent productivity shocks (Aguiar and Gopinath, 2007). Finally, at higher frequencies, the inter-temporal wedge is much less volatile than the Backus-Smith wedge. This can clearly be seen from Figure 2b that depicts the volatilities of the log changes in the two wedges and shows an almost uniform ranking of the two. This observations suggests a much better risk sharing at business cycle frequencies.

With these results in hand, we can do the welfare calculations implied by the wedges. In Figure 5, we plot the welfare gain for Home enjoyed by shutting down the planner's wedge versus the private sector wedge. As explained above, a constant Pareto weight for each country is calibrated to keep the welfare of the rest of the world unchanged, i.e. all welfare gains from better risk sharing are allocated to a given country. Panel (a) considers the case of shutting down any volatility in the two wedges. For most countries, the implied welfare gains are larger when measured with respect to ψ_t than for λ_t. That is, private agents perceive a larger potential gain to undistorted inter-temporal 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. Panel (b) shows welfare gains from shutting down the cyclical component of each wedge, while keeping the same trend that can be attributed to cross-country differences in discount factor β

Figure 4: Time Series of Wedges

(a) Levels



(b) First-Differences

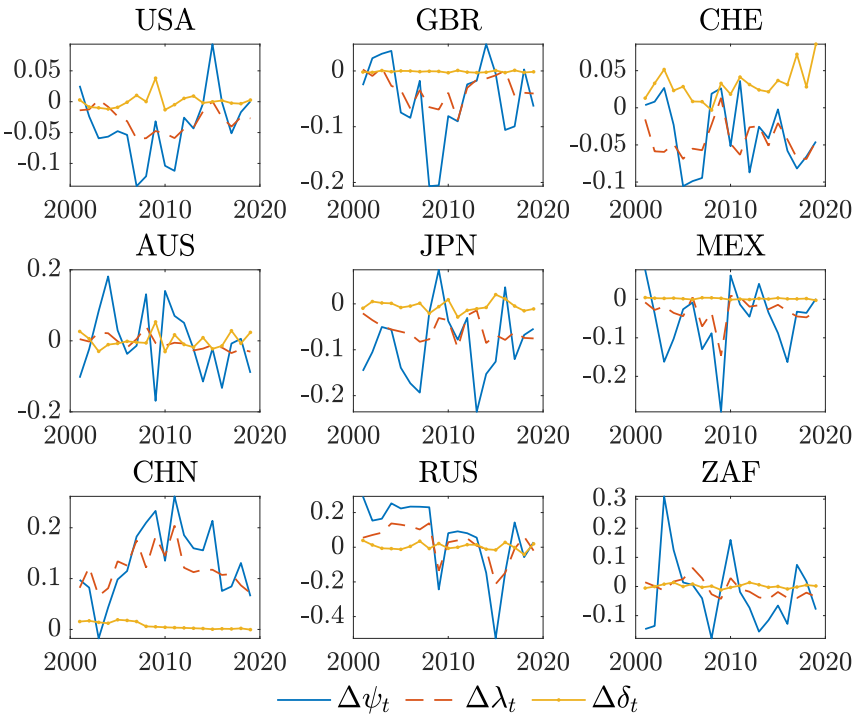
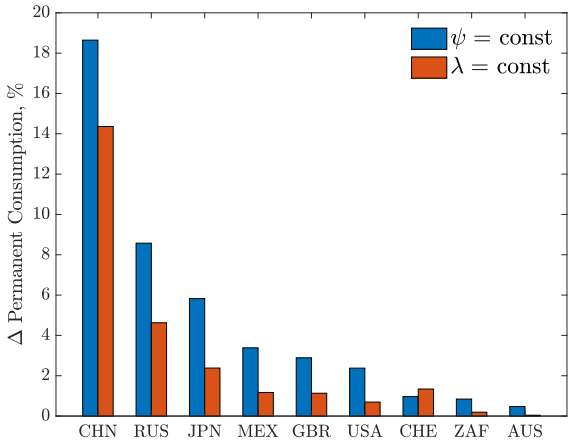
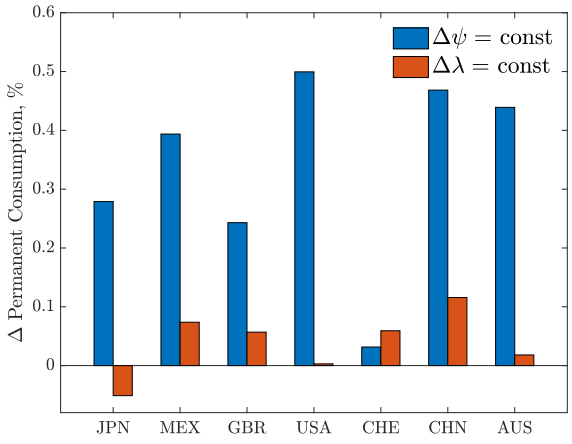


Figure 5: Welfare Gains

(a) Total



(b) High Frequency



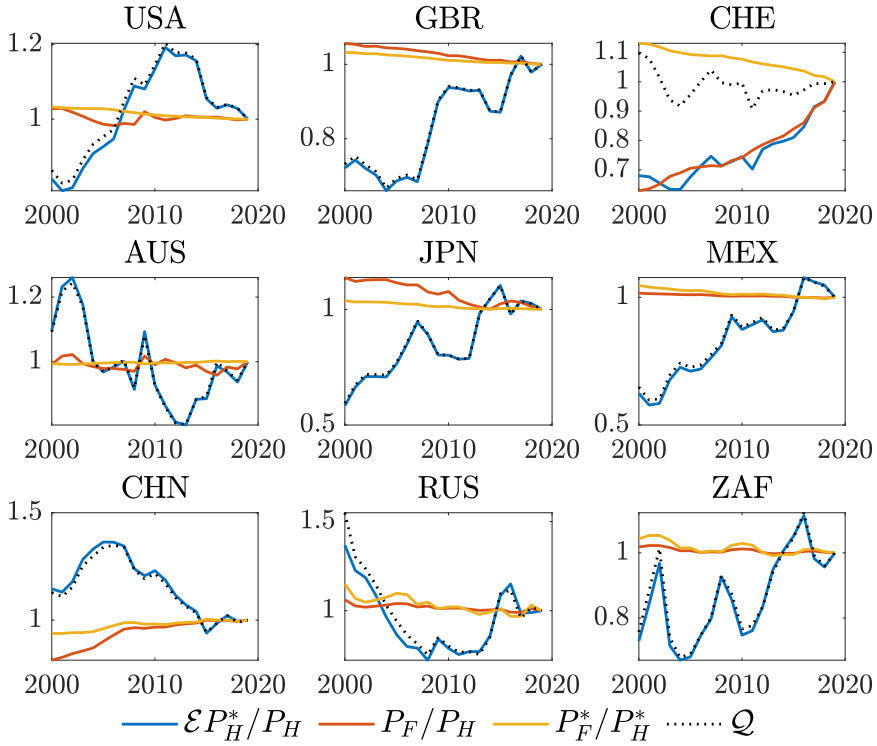
and/or permanent shocks.

Lastly, we 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 6 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, and the real exchange rate Q. The relative prices faced by consumers in the respective 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.

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

Figure 6: Implied Prices



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

Aguiar, Mark and Gita Gopinath, "Emerging Market Business Cycles: The Cycle Is the Trend," *Journal of Political Economy*, 2007, *115*, 69–102.

_ , Manuel Amador, and Doireann Fitzgerald, "Tariff Wars and Net Foreign Assets," 2025. working paper.

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.

Atkeson, Andrew and Ariel T. Burstein, "Trade Costs, Pricing-to-Market, and International Relative Prices," *American Economic Review*, December 2008, *98* (5), 1998–2031.

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.

Baqaee, David R. and Ariel T. Burstein, "Aggregate Efficiency with Heterogeneous Agents," 2025. working paper.

Baqaee, David Rezza and Emmanuel Farhi, "Productivity and Misallocation in General Equilibrium*," *The Quarterly Journal of Economics*, 09 2020, *135* (1), 105–163.

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

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas Sargent, "Efficiency, Insurance, and Redistribution Effects of Government Policies," 2025. Working Paper.

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.

Dávila, Eduardo and Andreas Schaab, "Welfare Assessments with Heterogeneous Individuals," *Journal of Political Economy*, 2025. forthcoming.

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

Devereux, Michael B. and Charles Engel, "Exchange rate pass-through, exchange rate volatility, and exchange rate disconnect," *Journal of Monetary Economics*, July 2002, *49* (5), 913–940.

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.

__ **, Emine Boz, Camila Casas, Federico J. Díez, Pierre-Olivier Gourinchas, and Mikkel Plagborg-Møller**, "Dominant Currency Paradigm," *American Economic Review*, 2020, *110* (3), 677–719.

Heathcote, Jonathan and Fabrizio Perri, "The International Diversification Puzzle Is Not as Bad as You Think," *Journal of Political Economy*, 2013, *121* (6), 1108–159.

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

Itskhoki, Oleg, "The Story of the Real Exchange Rate," *Annual Review of Economics*, 2021, *13* (1), 423–455.

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

__ **and** __ , "The Optimal Macro Tariff," 2025. NBER Working Paper No. 33839.

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

Obstfeld, Maurice and Kenneth Rogoff, "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 1995, *103*, 624–60.

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 Model Appendix

A.1 Characterization for the Planner's Problem in Section 2

Consider the constrained planner's problem (13) using the consumption aggregators defined in (1) and (2). Denote with μ and μ^* the Largrange multipliers on these conditions, with ν and ν^* the Lagrange multipliers on the resource constraints in (3) and, for convenience, with $\eta C_H/C_F$ the Lagrange multiplier on the δ-constraint (8) and with $\omega\lambda$ the Lagrange multiplier on the final constraint $U(\mathcal{C}(C_H, C_F)) = \overline{U}$. The optimality conditions for (C_H, C_F, C_H^*, C_F^*) are given by:

$$\begin{aligned}
\mu C(1-\gamma)^{\frac{1}{\theta}}\left(\frac{C_H}{C}\right)^{\frac{\theta-1}{\theta}} &= C_H\nu - \eta,\\
\mu C\gamma^{\frac{1}{\theta}}\left(\frac{C_F}{C}\right)^{\frac{\theta-1}{\theta}} &= C_F\nu^* + \eta,\\
\mu^* C^*\gamma^{*\frac{1}{\theta}}\left(\frac{C_H^*}{C^*}\right)^{\frac{\theta-1}{\theta}} &= C_H^*\nu + \eta,\\
\mu^* C^*(1-\gamma^*)^{\frac{1}{\theta}}\left(\frac{C_F^*}{C^*}\right)^{\frac{\theta-1}{\theta}} &= C_F^*\nu^* - \eta,
\end{aligned}$$

where we use the relationship in (8) to simplify the optimality conditions for C_H^* and C_F^*. The optimality conditions for C and C^* are:

$$\mu = \omega(1+\lambda)U'(C) \qquad \text{and} \qquad \mu^* = U'(C^*),$$

where we used the fact that $U'(C) = C^{-\sigma}$ under CRRA utility. These six conditions, together with six conditions in (1)–(3), (8) and $U(C) = \bar{U}$, allow to solve for $(C_H, C_F, C_H^*, C_F^*, C, C^*)$ and $(\mu, \mu^*, \nu, \nu^*, \eta, \lambda)$ for a given exogenous vector $(Y, Y^*, \delta, \bar{U})$.

To characterize the planner's MRT $\tilde{Q}$, we take the ratio of the optimality conditions for C and C^*:

$$\frac{U'(C^*)}{\omega U'(C)} = (1+\lambda)\frac{\mu^*}{\mu} \qquad \Leftrightarrow \qquad \tilde{Q} = \mu^*/\mu,$$

where this implication follows from the definition of the wedge λ in (12).[19] Indeed, μ^*/μ is the ratio of Lagrange multipliers on C^* and C respectively, which also corresponds to the slope of the distorted CPF

[19] Note that λ is a Lagrange multiplier which ensures that $U(\mathcal{C}(C_H, C_F)) = \overline{U}$ holds. Given $\delta \in (-1, \infty)$ and $\overline{U} \in [U(\mathcal{C}(0,0)), U(\mathcal{C}(Y, Y^*))]$, the solution to the planner's problem (13) exists and is unique (the constraint set is a singleton), including a unique value of λ that satisfies the first-order optimality condition (12). Under CES consumption aggregators, the optimum to the planner's problem is internal, and therefore this first-order optimality condition is also sufficient to recover the unique value of λ that corresponds to $(Y, Y^*, \delta, \overline{U})$.

defined by (9), $\tilde{Q} = -\partial C/\partial C^*$. Adding up the optimality conditions for C_H and C_F and for C_H^* and C_F^*, we obtain:

$$\mu C = C_H \nu + C_F \nu^* \qquad \text{and} \qquad \mu^* C^* = C_H^* \nu + C_F^* \nu^*,$$

where we used the expressions for consumption aggregators in (1) and (2), but the same property holds for any homothetic aggregators. We therefore have the following expression for the planner's MRT:

$$\tilde{Q} = \frac{\mu^*}{\mu} = \frac{\frac{C_H^*}{C^*}\nu + \frac{C_F^*}{C^*}\nu^*}{\frac{C_H}{C}\nu + \frac{C_F}{C}\nu^*}.$$

Taking the ratio of the optimality conditions for C_F and C_H and, respectively, C_F^* and C_H^* we obtain:

$$\left(\frac{\gamma}{1-\gamma}\frac{C_H}{C_F}\right)^{\frac{1}{\theta}} = \frac{\nu^* + \eta/C_F}{\nu - \eta/C_H} \qquad \text{and} \qquad \left(\frac{1-\gamma^*}{\gamma^*}\frac{C_H^*}{C_F^*}\right)^{\frac{1}{\theta}} = \frac{\nu^* - \eta/C_F^*}{\nu + \eta/C_H^*}.$$

Taking further the ratio of these two expressions and using the constraint (8) yields:

$$\frac{\nu + \eta/C_H^*}{\nu^* - \eta/C_F^*}\frac{\nu^* + \eta/C_F}{\nu - \eta/C_H} = \left[\frac{\gamma}{1-\gamma}\frac{\gamma^*}{1-\gamma^*}\frac{C_H/C_F}{C_H^*/C_F^*}\right]^{1/\theta} = 1 + \delta,$$

which is a relationship between the wedge δ and the Lagrange multiplier η such that $\eta > 0$ if and only if $\delta > 0$, and vice versa (which also implies that $\eta = 0$ if and only if $\delta = 0$).

Lastly, we aggregate the optimality conditions for C_H and C_F and, respectively, C_H^* and C_F^* and use the resource constraint (3) to obtain:

$$\begin{aligned}
\mu C(1-\gamma)^{\frac{1}{\theta}}\left(\frac{C_H}{C}\right)^{\frac{\theta-1}{\theta}} + \mu^* C^* \gamma^{*\frac{1}{\theta}}\left(\frac{C_H^*}{C^*}\right)^{\frac{\theta-1}{\theta}} &= \nu Y, \\
\mu C \gamma^{\frac{1}{\theta}}\left(\frac{C_F}{C}\right)^{\frac{\theta-1}{\theta}} + \mu^* C^*(1-\gamma^*)^{\frac{1}{\theta}}\left(\frac{C_F^*}{C^*}\right)^{\frac{\theta-1}{\theta}} &= \nu^* Y^*,
\end{aligned}$$

and hence $\mu C + \mu^* C^* = \nu Y + \nu Y^*$. This provides an implicit characterization of the planner's allocation and the associated Lagrange multipliers in the general case.

A.2 The Cole-Obstfeld case. Proof of Lemma 1

The optimality conditions above simplify significantly in the Cole-Obstfeld case when $\sigma = \theta = 1$, that is, when the consumptions aggregators are Cobb-Douglas and the utility is logarithmic. When $\sigma = 1$, the optimality conditions for C and C^* become:

$$\mu C = \omega(1+\lambda) \qquad \text{and} \qquad \mu^* C^* = 1.$$

When $\theta = 1$, the optimality conditions for (C_H, C_F, C_H^*, C_F^*) simplify to:

$$(1-\gamma)\mu C = C_H \nu - \eta, \qquad \gamma^* \mu^* C^* = C_H^* \nu + \eta,$$
$$\gamma \mu C = C_F \nu^* + \eta, \qquad (1-\gamma^*)\mu^* C^* = C_F^* \nu^* - \eta.$$

and imply, together with resource constraints, that:

$$(1-\gamma)\mu C + \gamma^* \mu^* C^* = \nu Y \qquad \text{and} \qquad \gamma \mu C + (1-\gamma^*)\mu^* C^* = \nu^* Y^*.$$

Combing together, we can solve for:

$$\nu = \frac{(1-\gamma)\omega(1+\lambda) + \gamma^*}{Y} \qquad \text{and} \qquad \nu^* = \frac{\gamma\omega(1+\lambda) + (1-\gamma^*)}{Y^*}.$$

We can then characterize the planner's allocation as follows:

$$C_H = \frac{(1-\gamma)\omega(1+\lambda) + \eta}{(1-\gamma)\omega(1+\lambda) + \gamma^*} Y, \qquad C_H^* = \frac{\gamma^* - \eta}{(1-\gamma)\omega(1+\lambda) + \gamma^*} Y,$$
$$C_F = \frac{\gamma\omega(1+\lambda) - \eta}{\gamma\omega(1+\lambda) + (1-\gamma^*)} Y^*, \qquad C_F^* = \frac{(1-\gamma^*) + \eta}{\gamma\omega(1+\lambda) + (1-\gamma^*)} Y^*.$$

Substituting this solution into the constraint (8), we have:

$$\frac{1 + \eta/[(1-\gamma)\omega(1+\lambda)]}{1 - \eta/[\gamma\omega(1+\lambda)]} \frac{1 + \eta/(1-\gamma^*)}{1 - \eta/\gamma^*} = 1 + \delta,$$

which provides a strictly monotonic relationship between η and δ, for any given value of λ, such that $\eta = 0$ if and only if $\delta = 0$. Furthermore, if δ and λ are constant, then η is constant even if (Y, Y^*) are changing.

Finally, we use the Cobb-Douglas utility aggregators, e.g., $C = \left(\frac{C_H}{1-\gamma}\right)^{1-\gamma}\left(\frac{C_F}{\gamma}\right)^{\gamma}$, to solve for aggregate consumption:

$$C = \omega(1+\lambda)\left(\frac{1 + \eta/[(1-\gamma)\omega(1+\lambda)]}{(1-\gamma)\omega(1+\lambda) + \gamma^*}\right)^{1-\gamma}\left(\frac{1 - \eta/[\gamma\omega(1+\lambda)]}{\gamma\omega(1+\lambda) + (1-\gamma^*)}\right)^{\gamma} Y^{1-\gamma} Y^{*\gamma},$$
$$C^* = \left(\frac{1 - \eta/\gamma^*}{(1-\gamma)\omega(1+\lambda) + \gamma^*}\right)^{\gamma^*}\left(\frac{1 + \eta/(1-\gamma^*)}{\gamma\omega(1+\lambda) + (1-\gamma^*)}\right)^{1-\gamma^*} Y^{\gamma^*} Y^{*1-\gamma^*},$$

and the planner's MRT:

$$\tilde{Q} = \frac{1}{\omega(1+\lambda)}\frac{C}{C^*} = \underbrace{\frac{(1 + \eta/[(1-\gamma)\omega(1+\lambda)])^{1-\gamma}(1 - \eta/[\gamma\omega(1+\lambda)])^{\gamma}}{(1 - \eta/\gamma^*)^{\gamma^*}(1 + \eta/(1-\gamma^*))^{1-\gamma^*}}}_{\text{effect of static distortion } \delta}\left(\underbrace{\frac{\gamma\omega(1+\lambda) + (1-\gamma^*)}{(1-\gamma)\omega(1+\lambda) + \gamma^*}\frac{Y}{Y^*}}_{=\nu^*/\nu}\right)^{1-\gamma-\gamma^*}$$

When $\delta = 0$, and hence $\eta = 0$, the first terms equals to one for any value of $\omega(1+\lambda)$, and therefore the undistorted MRT is given by $\tilde{Q} = \left(\frac{\nu^*}{\nu}\right)^{1-\gamma-\gamma^*} = \left(\frac{\gamma\omega(1+\lambda)+(1-\gamma^*)}{(1-\gamma)\omega(1+\lambda)+\gamma^*}\frac{Y}{Y^*}\right)^{1-\gamma-\gamma^*}$. When δ and λ are unchanged, so

is η, and therefore the proportional changes in C, C^* and $\tilde{Q}$ are given by:

$$C \propto Y^{1-\gamma} Y^{*\gamma}, \qquad C^* \propto Y^{\gamma^*} Y^{*1-\gamma^*}, \qquad \tilde{Q} \propto \left(\frac{Y}{Y^*}\right)^{1-\gamma-\gamma^*},$$

or in other words:

$$\Delta \log(C/C^*) = \Delta \log \tilde{Q} = (1-\gamma-\gamma^*)\Delta \log(Y/Y^*).$$

Furthermore, by full differentiation, one can establish that $\partial C/\partial\lambda > 0$, $\partial C^*/\partial\lambda < 0$ for a given (Y, Y^*, δ), and $\partial C/\partial\delta < 0$ and $\partial C^*/\partial\delta < 0$ when $\delta > 0$, and vice versa, for a given (Y, Y^*, λ). Note that full differentiation takes into account the effects of a changing λ on η for a given δ. This completes the proof of Lemma 1. ■

A.3 Second-order Approximation to the Welfare Loss

Point of approximation We characterize a second order expansion of the welfare loss function $\mathbb{L}(\lambda, \delta)$ defined in (14) around an efficient point with $\delta = \lambda = 0$. We denote by $\bar{Y}$ and $\bar{Y}^*$ the values of Y and Y^* in the *baseline* allocation, and we denote with bars the other variables corresponding to this allocation. We also define the following constants measuring openness of the economies under the baseline allocation:

$$\bar{\gamma} \equiv \gamma^{\frac{1}{\theta}} \left(\frac{\bar{C}_F}{\bar{C}}\right)^{\frac{\theta-1}{\theta}}, \qquad \bar{\gamma}^* \equiv \gamma^{*\frac{1}{\theta}} \left(\frac{\bar{C}^*_H}{\bar{C}^*}\right)^{\frac{\theta-1}{\theta}}, \qquad \bar{\alpha} \equiv \frac{\bar{C}^*_H}{\bar{Y}}, \qquad \bar{\alpha}^* \equiv \frac{\bar{C}_F}{\bar{Y}^*}.$$

Note that in the Cobb-Douglas case with $\theta = 1$ we simply have $\bar{\gamma} = \gamma$ and $\bar{\gamma}^* = \gamma^*$. Also note from the definition of the consumption aggregator (1) we have that $1-\bar{\gamma} = (1-\gamma)^{\frac{1}{\theta}} \left(\frac{\bar{C}_H}{\bar{C}}\right)^{\frac{\theta-1}{\theta}}$ and from the resource constraint (3) we have $1-\bar{\alpha} = \bar{C}_H/\bar{Y}$. Symmetric equations hold for $1-\bar{\gamma}^*$ and $1-\bar{\alpha}^*$. Rewriting the planner's optimality conditions for $(C_H, C_F, C^*_H, C^*_F, C, C^*)$ from Appendix A.1 for the baseline undistorted allocation, we have:[20]

$$(1-\bar{\gamma})\omega\bar{C}^{1-\sigma} = (1-\bar{\alpha})\bar{\nu}\bar{Y}, \qquad \bar{\gamma}^*\bar{C}^{*1-\sigma} = \bar{\alpha}\bar{\nu}\bar{Y},$$
$$\bar{\gamma}\omega\bar{C}^{1-\sigma} = \bar{\alpha}^*\bar{\nu}^*\bar{Y}^*, \qquad (1-\bar{\gamma}^*)\bar{C}^{*1-\sigma} = (1-\bar{\alpha}^*)\bar{\nu}^*\bar{Y}^*.$$

Taking the ratios to first eliminate $\omega\bar{C}^{1-\sigma}$ and $\bar{C}^{*1-\sigma}$ and then another ratio to eliminate $(\bar{\nu}\bar{Y})/(\bar{\nu}^*\bar{Y}^*)$, we have:

$$\frac{\bar{\gamma}}{1-\bar{\gamma}}\frac{\bar{\gamma}^*}{1-\bar{\gamma}^*} = \frac{\bar{\alpha}}{1-\bar{\alpha}}\frac{\bar{\alpha}^*}{1-\bar{\alpha}^*},$$

a condition that holds for any allocation with $\delta = 0$, but does not hold when $\delta \neq 0$.

[20] Note that $\eta = 0$ when $\delta = 0$. Therefore, the optimality condition for C_H simplifies to $\bar{\mu}\bar{C}(1-\gamma)^{\frac{1}{\theta}} \left(\frac{\bar{C}_H}{\bar{C}}\right)^{\frac{\theta-1}{\theta}} = \bar{C}_H\bar{\nu}$, while $\bar{\mu}\bar{C} = \omega$ when evaluated at the baseline allocation with $\lambda = \delta = 0$. Similar expressions hold for $\bar{C}_F$, $\bar{C}^*_H$ and $\bar{C}^*_F$, and result in the four conditions in the text after we apply definitions of $\bar{\gamma}$, $\bar{\alpha}$, etc.

First-order expansion for the observed allocation The observed allocation is (C, C^*, Y, Y^*), where $Y = \bar{Y}$ and $Y^* = \bar{Y}^*$ by construction, while C and C^* may be different from $\bar{C}$ and $\bar{C}^*$ indicating the presence of wedges λ and/or δ. We now characterize the first order deviations in the consumption allocation $(C_H, C_F, C_H^*, C_F^*, C, C^*)$ from its baseline level in response to $(\lambda, \delta) \neq 0$, and denote the corresponding proportional (or log) deviations with small letters, $(c_H, c_F, c_H^*, c_F^*, c, c^*)$, such that, e.g., $c \equiv (C - \bar{C})/\bar{C} \approx \log(C/\bar{C})$.

From the resource constraint, we have:

$$(1 - \bar{\alpha})c_H + \bar{\alpha}c_H^* = 0 \qquad \text{and} \qquad \bar{\alpha}^* c_F + (1 - \bar{\alpha}^*)c_F^* = 0.$$

From the consumption aggregators, we have:

$$c = (1 - \bar{\gamma})c_H + \bar{\gamma}c_F \qquad \text{and} \qquad c^* = \bar{\gamma}^* c_H^* + (1 - \bar{\gamma}^*)c_F^*.$$

Next we expand the planner's optimality conditions for (C_H, C_F, C_H^*, C_F^*) from Appendix A.1:[21]

$$\begin{aligned}
-\frac{1}{\theta}c_H + \left(\frac{1}{\theta} - \sigma\right)c + \lambda - \hat{\nu} &= -\frac{1}{(1 - \bar{\gamma})\omega\bar{C}^{1-\sigma}}\eta, \\
-\frac{1}{\theta}c_F + \left(\frac{1}{\theta} - \sigma\right)c + \lambda - \hat{\nu}^* &= \frac{1}{\bar{\gamma}\omega\bar{C}^{1-\sigma}}\eta, \\
-\frac{1}{\theta}c_H^* + \left(\frac{1}{\theta} - \sigma\right)c^* - \hat{\nu} &= \frac{1}{\bar{\gamma}^*\bar{C}^{*1-\sigma}}\eta, \\
-\frac{1}{\theta}c_F^* + \left(\frac{1}{\theta} - \sigma\right)c^* - \hat{\nu}^* &= -\frac{1}{(1 - \bar{\gamma}^*)\bar{C}^{*1-\sigma}}\eta,
\end{aligned}$$

where $\hat{\nu} \equiv (\nu - \bar{\nu})/\bar{\nu}$ and η as before (since $\bar{\eta} = 0$). We take differences of these equations to solve out $\hat{\nu}$ and $\hat{\nu}^*$, and use the resource constraints to solve out c_H^* and c_F. This yields:

$$\begin{aligned}
-\frac{1}{\bar{\alpha}\theta}c_H + \left(\frac{1}{\theta} - \sigma\right)(c - c^*) + \lambda &= -\frac{1}{\bar{\alpha}(1 - \bar{\gamma})\omega\bar{C}^{1-\sigma}}\eta, \\
\frac{1}{\bar{\alpha}^*\theta}c_F^* + \left(\frac{1}{\theta} - \sigma\right)(c - c^*) + \lambda &= \frac{1}{(1 - \bar{\alpha}^*)\bar{\gamma}\omega\bar{C}^{1-\sigma}}\eta,
\end{aligned}$$

where we used the relationships $(1 - \bar{\gamma})\omega\bar{C}^{1-\sigma} = \frac{1-\bar{\alpha}}{\bar{\alpha}}\bar{\gamma}^*\bar{C}^{*1-\sigma}$ and $\bar{\gamma}\omega\bar{C}^{1-\sigma} = \frac{\bar{\alpha}^*}{1-\bar{\alpha}^*}(1 - \bar{\gamma}^*)\bar{C}^{*1-\sigma}$ that hold for the baseline allocation.

Next using the δ-constraint (8) on the planner and the characterization above, we derive the approxi-

[21] For example, the optimality conditions for C_H can be rewritten without approximation (i.e., the "exact hat algebra") as:

$$\omega(1 - \bar{\gamma})\bar{C}^{1-\sigma}(1 - \lambda)(C/\bar{C})^{1-\sigma}\left(\frac{C_H/\bar{C}_H}{C/\bar{C}}\right)^{\frac{\theta-1}{\theta}} = \bar{\nu}\bar{C}_H(C_H/\bar{C}_H)(\nu/\bar{\nu}) - \eta,$$

and note that $\bar{\nu}\bar{C}_H = \omega(1 - \bar{\gamma})\bar{C}^{1-\sigma}$ for the baseline allocation with $\eta = \lambda = 0$. Dividing by $\omega(1 - \bar{\gamma})\bar{C}^{1-\sigma}$ and taking a first-order Taylor expansion results in the first line in the text.

mate relationship between δ and η:

$$\delta = \frac{1}{\theta}\left[c_H - c_F - c_H^* + c_F^*\right] = \frac{1}{\bar{\alpha}\theta}c_H + \frac{1}{\bar{\alpha}^*\theta}c_F^* = \frac{(1-\bar{\gamma})\bar{\alpha} + (1-\bar{\alpha}^*)\bar{\gamma}}{(1-\bar{\alpha}^*)\bar{\gamma}}\frac{1}{\bar{\alpha}(1-\bar{\gamma})\omega\bar{C}^{1-\sigma}}\eta.$$

Lastly, we use consumption aggregators to solve for:

$$\begin{aligned} c - c^* &= (1-\bar{\gamma})c_H + \bar{\gamma}c_F - \bar{\gamma}^* c_H^* - (1-\bar{\gamma}^*)c_F^* \\ &= \left[(1-\bar{\gamma})\bar{\alpha} + \bar{\gamma}^*(1-\bar{\alpha})\right]\frac{c_H}{\bar{\alpha}} - \left[(1-\bar{\gamma}^*)\bar{\alpha}^* + \bar{\gamma}(1-\bar{\alpha}^*)\right]\frac{c_F^*}{\bar{\alpha}^*} = \frac{\phi}{1+(\sigma\theta-1)\phi}\theta\lambda, \end{aligned}$$

where $\phi \equiv (1-\bar{\gamma})\bar{\alpha} + \bar{\gamma}^*(1-\bar{\alpha}) + (1-\bar{\gamma}^*)\bar{\alpha}^* + \bar{\gamma}(1-\bar{\alpha}^*)$, and the last equality substitutes in the solutions for c_H and c_F^* and simplifies using the fact that $\frac{\bar{\gamma}}{1-\bar{\gamma}}\frac{\bar{\gamma}^*}{1-\bar{\gamma}^*} = \frac{\bar{\alpha}}{1-\bar{\alpha}}\frac{\bar{\alpha}^*}{1-\bar{\alpha}^*}$.

Using these derivations, we can now express the full consumption allocation as a function of wedges (λ, δ), in proportional deviations around the efficient baseline:

$$\begin{aligned} \frac{1}{\bar{\alpha}}c_H = -\frac{1}{1-\bar{\alpha}}c_H^* &= \kappa\theta\lambda + \chi\theta\delta, \\ \frac{1}{\bar{\alpha}^*}c_F^* = -\frac{1}{1-\bar{\alpha}^*}c_F &= -\kappa\theta\lambda + (1-\chi)\theta\delta, \end{aligned}$$

where $\chi \equiv \frac{(1-\bar{\alpha}^*)\bar{\gamma}}{(1-\bar{\gamma})\bar{\alpha}+(1-\bar{\alpha}^*)\bar{\gamma}}$ and $\kappa \equiv \frac{1}{1+(\sigma\theta-1)\phi}$, and recall that $\phi = (1-\bar{\gamma})\bar{\alpha} + \bar{\gamma}^*(1-\bar{\alpha}) + (1-\bar{\gamma}^*)\bar{\alpha}^* + \bar{\gamma}(1-\bar{\alpha}^*)$. Furthermore, this implies that

$$c = \left[(1-\bar{\gamma})\bar{\alpha} + (1-\bar{\alpha}^*)\bar{\gamma}\right]\kappa\theta\lambda \qquad \text{and} \qquad c^* = -\left[(1-\bar{\alpha})\bar{\gamma}^* + (1-\bar{\gamma}^*)\bar{\alpha}^*\right]\kappa\theta\lambda,$$

and indeed there is no first-order effect of δ on aggregate consumption around $\delta = 0$.

Second-order expansion for the welfare loss Consider next the welfare function maximized by the constrained planner in problem (13), which we can write as:

$$\mathbb{W} = \omega U(C) + U(C^*) = \omega U(\mathcal{C}(C_H, C_F)) + U(\mathcal{C}^*(C_H^*, C_F^*)),$$

where $(C_H, C_F, C_H^*, C_F^*, C, C^*)$ is the observed allocation corresponding to wedges (λ, δ), and for a given $Y = \bar{Y}$ and $Y^* = \bar{Y}^*$. The first-best welfare is $\bar{\mathbb{W}} = \omega U(\bar{C}) + U(\bar{C}^*)$, evaluated at the efficient baseline allocation. We now characterize the second order expansion in (λ, δ) to the welfare loss defined as:

$$\mathbb{L}(\lambda, \delta) = \bar{\mathbb{W}} - \mathbb{W}.$$

Note that the first-order terms are absent by the optimality of the undistorted allocation, which can be verified directly by differentiating $\omega U(\mathcal{C}(C_H, Y^* - C_F^*)) + U(\mathcal{C}^*(Y - C_H, C_F^*))$ with respect to C_H and C_F^* and evaluating the resulting expressions at $(\bar{C}_H, \bar{C}_F^*)$. Note that we used the resource constraint (3) to substitute out C_H^* and C_F, resulting in an unconstrained optimization problem for an undistorted planner.

Taking the second-order expansion for $\mathbb{W}(C_H, C_F^*) = \omega U(C(C_H, Y^* - C_F^*)) + U(C^*(Y - C_H, C_F^*))$ in C_H and C_F^* around $\bar{C}_H$ and $\bar{C}_F^*$ yields the following result:

$$\mathbb{L}(\lambda, \delta) = -\frac{\bar{\mathbb{W}}_{HH}\bar{C}_H^2}{2}c_H^2 - \frac{\bar{\mathbb{W}}_{FF}\bar{C}_F^{*2}}{2}c_F^{*2} - \bar{\mathbb{W}}_{HF}\bar{C}_H\bar{C}_F^* c_H c_F^*,$$

where $\bar{\mathbb{W}}_{HF} \equiv \frac{\partial^2 \mathbb{W}(C_H, C_F^*)}{\partial C_H \partial C_F^*}\big|_{\bar{C}_H, \bar{C}_F^*}$ and similarly for $\bar{\mathbb{W}}_{HH}$ and $\bar{\mathbb{W}}_{FF}$. Note that $\frac{\partial \mathbb{W}}{\partial C_H} = \omega U' \frac{\partial \mathcal{C}}{\partial C_H} - U^{*\prime}\frac{\partial \mathcal{C}^*}{\partial C_H^*}$ and $\frac{\partial \mathbb{W}}{\partial C_F^*} = -\omega U' \frac{\partial \mathcal{C}}{\partial C_F} + U^{*\prime}\frac{\partial \mathcal{C}^*}{\partial C_F^*}$. Therefore, we have:

$$\begin{aligned}
\bar{\mathbb{W}}_{HH} &= \omega \bar{U}''(\bar{\mathcal{C}}_H)^2 + \omega \bar{U}' \bar{\mathcal{C}}_{HH} + \bar{U}^{*\prime\prime}(\bar{\mathcal{C}}_H^*)^2 + \bar{U}^{*\prime}\bar{\mathcal{C}}_{HH}^*, \\
\bar{\mathbb{W}}_{FF} &= \omega \bar{U}''(\bar{\mathcal{C}}_F)^2 + \omega \bar{U}' \bar{\mathcal{C}}_{FF} + \bar{U}^{*\prime\prime}(\bar{\mathcal{C}}_F^*)^2 + \bar{U}^{*\prime}\bar{\mathcal{C}}_{FF}^*, \\
\bar{\mathbb{W}}_{HF} &= -\omega \bar{U}''\bar{\mathcal{C}}_H\bar{\mathcal{C}}_F - \omega \bar{U}'\bar{\mathcal{C}}_{HF} - \bar{U}^{*\prime\prime}\bar{\mathcal{C}}_H^*\bar{\mathcal{C}}_F^* - \bar{U}^{*\prime}\bar{\mathcal{C}}_{HF}^*,
\end{aligned}$$

where $\bar{\mathcal{C}}_H = \frac{\partial \mathcal{C}(C_H, C_F)}{\partial C_H}\big|_{\bar{C}_H, \bar{C}_F = \bar{Y}^* - \bar{C}_F^*}$, $\bar{\mathcal{C}}^*{}_H = \frac{\partial \mathcal{C}^*(C_H^*, C_F^*)}{\partial C_H^*}\big|_{\bar{C}_H^* = \bar{Y} - \bar{C}_H, \bar{C}_F}$, $\bar{\mathcal{C}}_{HF} = \frac{\partial^2 \mathcal{C}(C_H, C_F)}{\partial C_H \partial C_F}\big|_{\bar{C}_H, \bar{C}_F = \bar{Y}^* - \bar{C}_F^*}$, and similarly for other derivatives. We have:

$$\bar{\mathcal{C}}_H = (1-\gamma)^{\frac{1}{\theta}}\left(\frac{\bar{C}_H}{\bar{C}}\right)^{-\frac{1}{\theta}} = (1-\bar{\gamma})\frac{\bar{C}}{\bar{C}_H}, \qquad \bar{\mathcal{C}}_F = \bar{\gamma}\frac{\bar{C}}{\bar{C}_F},$$

$$\bar{\mathcal{C}}_{HH} = -\frac{1}{\theta}\bar{\gamma}(1-\bar{\gamma})\frac{\bar{C}}{\bar{C}_H^2}, \qquad \bar{\mathcal{C}}_{FF} = -\frac{1}{\theta}\bar{\gamma}(1-\bar{\gamma})\frac{\bar{C}}{\bar{C}_F^2}, \qquad \bar{\mathcal{C}}_{HF} = \frac{1}{\theta}\bar{\gamma}(1-\bar{\gamma})\frac{\bar{C}}{\bar{C}_H\bar{C}_F},$$

and similarly for $\mathcal{C}^*$. We also have $\bar{U}''\bar{C} = -\sigma\bar{U}' = \bar{C}^{-\sigma}$, and same for $\bar{U}^*$. Finally, $\bar{C}_F/\bar{C}_F^* = \bar{\alpha}^*/(1-\bar{\alpha}^*)$ and $\bar{C}_H^*/\bar{C}_H = \bar{\alpha}/(1-\bar{\alpha})$ from the definition of $\bar{\alpha}$. Using this, we can solve for:

$$\begin{aligned}
\frac{1}{\omega\bar{C}^{1-\sigma}}\bar{\mathbb{W}}_{HH}\bar{C}_H^2 &= -(1-\bar{\gamma})\left[\left(\sigma(1-\bar{\gamma}) + \frac{1}{\theta}\bar{\gamma}\right) + \frac{1-\bar{\alpha}}{\bar{\alpha}}\left(\sigma\bar{\gamma}^* + \frac{1}{\theta}(1-\bar{\gamma}^*)\right)\right], \\
\frac{1}{\omega\bar{C}^{1-\sigma}}\bar{\mathbb{W}}_{FF}\bar{C}_F^{*2} &= -\bar{\gamma}\frac{1-\bar{\alpha}^*}{\bar{\alpha}^*}\left[\frac{1-\bar{\alpha}^*}{\bar{\alpha}^*}\left(\sigma\bar{\gamma} + \frac{1}{\theta}(1-\bar{\gamma})\right) + \left(\sigma(1-\bar{\gamma}^*) + \frac{1}{\theta}\bar{\gamma}^*\right)\right], \\
\frac{1}{\omega\bar{C}^{1-\sigma}}\bar{\mathbb{W}}_{HF}\bar{C}_H\bar{C}_F^* &= -(1-\bar{\gamma})\left[\bar{\gamma}\frac{1-\bar{\alpha}^*}{\bar{\alpha}^*} + (1-\bar{\gamma}^*)\right]\left(\frac{1}{\theta} - \sigma\right),
\end{aligned}$$

where we the fact that $\frac{\bar{C}^{*1-\sigma}}{\omega\bar{C}^{1-\sigma}} = \frac{1-\bar{\gamma}}{\bar{\gamma}^*}\frac{\bar{\alpha}}{1-\bar{\alpha}} = \frac{\bar{\gamma}}{1-\bar{\gamma}^*}\frac{1-\bar{\alpha}^*}{\bar{\alpha}^*}$.

Using this characterization, we can express the second-order of the welfare loss functions as follows:

$$\begin{aligned}
\frac{1}{\omega\bar{C}^{1-\sigma}}\mathbb{L}(\lambda, \delta) =& \frac{\theta}{2}\bar{\alpha}(1-\bar{\gamma})\left[1 + (\bar{\alpha}(1-\bar{\gamma}) + \bar{\gamma}^*(1-\bar{\alpha}))(\sigma\theta - 1)\right]\left(\frac{c_H}{\bar{\alpha}\theta}\right)^2 \\
&+ \frac{\theta}{2}\bar{\gamma}(1-\bar{\alpha}^*)\left[1 + (\bar{\gamma}(1-\bar{\alpha}^*) + \bar{\alpha}^*(1-\bar{\gamma}^*))(\sigma\theta - 1)\right]\left(\frac{c_F^*}{\bar{\alpha}^*\theta}\right)^2 \\
&- \theta\bar{\alpha}(1-\bar{\gamma})\left[\bar{\gamma}(1-\bar{\alpha}^*) + \bar{\alpha}^*(1-\bar{\gamma}^*)\right](\sigma\theta - 1)\left(\frac{c_H}{\bar{\alpha}\theta}\right)\left(\frac{c_F^*}{\bar{\alpha}^*\theta}\right),
\end{aligned}$$

where we solved earlier for $\frac{c_H}{\bar{\alpha}\theta} = \kappa\lambda + \chi\delta$ and $\frac{c_F^*}{\bar{\alpha}^*\theta} = -\kappa\lambda + (1-\chi)\delta$ with $\chi \equiv \frac{(1-\bar{\alpha}^*)\bar{\gamma}}{(1-\bar{\gamma})\bar{\alpha} + (1-\bar{\alpha}^*)\bar{\gamma}}$.

Lemma A1. *The second-order expansions to the welfare loss function $\mathbb{L}(\lambda,\delta)$ is separable quadratic in λ and δ, i.e., takes the form $\mathbb{L}(\lambda,\delta)=\frac{1}{2}[\bar{A}\lambda^2+\bar{B}\delta^2]$ for some $\bar{A},\bar{B}>0$ and features no interaction term $\lambda\delta$. In particular, $\bar{A}=\omega U'(\bar{C})\cdot\frac{\partial \mathcal{C}}{\partial\lambda}\big|_{\bar{C}}=\omega\bar{C}^{1-\sigma}[\bar{\alpha}(1-\bar{\gamma})+\bar{\gamma}(1-\bar{\alpha}^*)]\kappa\theta$.*

Proof: Given the expansion for $\mathbb{L}(\lambda,\delta)$ above, for there to be no interaction term between λ and δ, it is sufficient that the following two conditions hold:

$$\bar{\alpha}(1-\bar{\gamma})\chi=\bar{\gamma}(1-\bar{\alpha}^*)(1-\chi),$$
$$[\bar{\alpha}(1-\bar{\gamma})+\bar{\gamma}^*(1-\bar{\alpha})]\chi=[\bar{\gamma}(1-\bar{\alpha}^*)+\bar{\alpha}^*(1-\bar{\gamma}^*)]\frac{\bar{\gamma}(1-\bar{\alpha}^*)}{\bar{\alpha}(1-\bar{\gamma})}(1-\chi)+[\bar{\gamma}(1-\bar{\alpha}^*)+\bar{\alpha}^*(1-\bar{\gamma}^*)]\,(1-2\chi).$$

The first condition holds immediately by the definition of χ which implies $\frac{\chi}{1-\chi}=\frac{\bar{\gamma}(1-\bar{\alpha}^*)}{\bar{\alpha}(1-\bar{\gamma})}$. Using this, we can simplify the second condition as follows:

$$[\bar{\alpha}(1-\bar{\gamma})+\bar{\gamma}^*(1-\bar{\alpha})]\bar{\gamma}(1-\bar{\alpha}^*)=[\bar{\gamma}(1-\bar{\alpha}^*)+\bar{\alpha}^*(1-\bar{\gamma}^*)]\,\bar{\alpha}(1-\bar{\gamma}).$$

This condition always holds by the property that $\bar{\gamma}^*(1-\bar{\alpha})\bar{\gamma}(1-\bar{\alpha}^*)=\bar{\alpha}^*(1-\bar{\gamma}^*)\bar{\alpha}(1-\bar{\gamma})$.

Finally, we characterize the value of $\bar{A}$ in the expression for $\mathbb{L}(\lambda,\delta)$. By direct algebraic calculation:

$$\mathbb{L}(\lambda,\delta)=\frac{1}{2}\omega\bar{C}^{1-\sigma}[\bar{\alpha}(1-\bar{\gamma})+\bar{\gamma}(1-\bar{\alpha}^*)]\kappa\theta\lambda^2+\frac{1}{2}\bar{B}\delta^2.$$

Note that $\bar{C}^{1-\sigma}=U'(\bar{C})\bar{C}$. Furthermore, recall the first-order solution for aggregate consumption, $c=[(1-\bar{\gamma})\bar{\alpha}+(1-\bar{\alpha}^*)\bar{\gamma}]\,\kappa\theta\lambda$. Therefore, $\frac{1}{\bar{C}}\frac{\partial \mathcal{C}}{\partial\lambda}\big|_{\bar{C}}=\frac{\partial\log C}{\partial\lambda}\big|_{\bar{C}}=[(1-\bar{\gamma})\bar{\alpha}+(1-\bar{\alpha}^*)\bar{\gamma}]\,\kappa\theta$. ■

Two special cases While the general expression for $\mathbb{L}(\lambda,\delta)$ characterized above is burdened by $(\bar{\gamma},\bar{\gamma}^*,\bar{\alpha},\bar{\alpha}^*)$, there are two special case which simplify the expression considerably. First, consider the generalized Cole-Obstfeld case such that $\sigma\theta=1$. We have:

$$\begin{aligned}\frac{1}{\omega\bar{C}^{1-\sigma}}\mathbb{L}(\lambda,\delta)=&\frac{\theta}{2}\left[\bar{\alpha}(1-\bar{\gamma})\left(\frac{c_H}{\bar{\alpha}\theta}\right)^2+\bar{\gamma}(1-\bar{\alpha}^*)\left(\frac{c_F^*}{\bar{\alpha}^*\theta}\right)^2\right]\\=&\frac{\theta}{2}[\bar{\alpha}(1-\bar{\gamma})+\bar{\gamma}(1-\bar{\alpha}^*)]\left[\lambda^2+\chi(1-\chi)\delta^2\right],\end{aligned}$$

where we used $\frac{c_H}{\bar{\alpha}\theta}=\lambda+\chi\delta$, $\frac{c_F^*}{\bar{\alpha}^*\theta}=-\lambda+(1-\chi)\delta$ as $\kappa=1$ in this case, and we still have $\chi\equiv\frac{(1-\bar{\alpha}^*)\bar{\gamma}}{(1-\bar{\gamma})\bar{\alpha}+(1-\bar{\alpha}^*)\bar{\gamma}}$.

Second, the general expression simplifies considerably in the symmetric case such that $\bar{\gamma}=\bar{\gamma}^*=\bar{\alpha}=\bar{\alpha}^*$. A sufficient requirement for such symmetry is $\omega=1$, $\gamma=\gamma^*$ and $\bar{Y}=\bar{Y}^*$. We have:

$$\begin{aligned}\frac{1}{\omega\bar{C}^{1-\sigma}}\mathbb{L}(\lambda,\delta)=&\frac{\theta}{2}\bar{\gamma}(1-\bar{\gamma})\left[\left(\frac{c_H}{\bar{\alpha}\theta}\right)^2+\left(\frac{c_F^*}{\bar{\alpha}^*\theta}\right)^2+2\bar{\gamma}(1-\bar{\gamma})(\sigma\theta-1)\left(\frac{c_H}{\bar{\alpha}\theta}-\frac{c_F^*}{\bar{\alpha}^*\theta}\right)^2\right]\\=&\theta\bar{\gamma}(1-\bar{\gamma})\left[\kappa\lambda^2+\frac{1}{4}\delta^2\right],\end{aligned}$$

where we used $\frac{c_H}{\bar{\alpha}\theta} = \kappa\lambda + \chi\delta$ and $\frac{c_F^*}{\bar{\alpha}^*\theta} = -\kappa\lambda + (1-\chi)\delta$ with $\kappa \equiv \frac{1}{1+(\sigma\theta-1)\phi}$, $\phi = 4\bar{\gamma}(1-\bar{\gamma})$ and $\chi = 1-\chi = 1/2$.

Harberger Triangle for λ Lemma A1 provides the structure and the closed-form expression for the welfare loss from the risk-sharing friction λ. This expression has the following structure:

$$\text{Welfare loss from } \lambda = \frac{1}{2}\omega U'(\bar{C})\bar{C} \cdot \mathrm{d}\log C \cdot \mathrm{d}\lambda,$$

where $\mathrm{d}\lambda = \lambda$ since $\bar{\lambda} = 0$, and $\mathrm{d}\log C$ is the proportional change in consumption due to the λ wedge, which from our first-order solution is $\mathrm{d}\log C = c = [(1-\bar{\gamma})\bar{\alpha} + (1-\bar{\alpha}^*)\bar{\gamma}]\kappa\theta \cdot \mathrm{d}\lambda$. This is a Harberger triangle, where $\mathrm{d}\lambda = \lambda$ is the wedge, or the size of the distortion, and $\mathrm{d}\log C$ is the impact of the distortion on the allocation, while $\omega U'(\bar{C})\bar{C}$ is the welfare impact of this change in the allocation. The same calculation can be done using foreign consumption C^* and foreign utility $U(C^*)$.

A.4 Product Markup Wedges

A.5 Proof of Proposition 1

Consider the optimization problem (13). First, note that for any value of $\bar{U}$ in the final constraint, the δ constraint (8) is slack for $\delta = 0$. To see this, compare the optimality conditions in the relaxed problem without the constraint (8) and in the full problem with the constraint (8). The two sets of conditions coincides if and only if $\delta = 0$. Therefore, the objective in the full problem is strictly greater for $\delta = 0$ than for any $\delta \neq 0$, and thus $\arg\max_\delta \mathbb{W}(\bar{U}, \delta) = 0$ for any value of $\bar{U}$.

Next, for a given value of δ in the constraint (8) consider the optimization problem (13) with and without the final constraint. The value of the unconstrained optimization is weakly greater and corresponds to $\bar{U}$ such that the constraint is slack with a corresponding Lagrange multiplier $\lambda = 0$. Furthermore, any other value of the Lagrange multiplier, $\lambda \neq 0$, implies that the constraint is tight, and hence the values of the objective $\mathbb{W} = \omega U + U^*$ is strictly lower. Therefore, for any value of δ, the minimum loss from the dynamic distortions corresponds to λ, that is, $\arg\min_\lambda \mathbb{L}(\lambda, \delta) = 0$ for any given value of δ.

The proof relies on the existence of the mapping $\bar{U}(\lambda, \delta)$ to define the loss function $\mathbb{L}(\lambda, \delta)$ in equation (14). Consider an alternative optimization problem:

$$\mathbb{V}(\lambda, \delta) \equiv \max_{\{C_H, C_F, C_H^*, C_F^*\}} \omega(1+\lambda)U(\mathcal{C}(C_H, C_F)) + U(\mathcal{C}^*(C_H^*, C_F^*)) \qquad \text{subject to (3) and (8).}$$

An increase in λ translates into a weak increase in $U(\mathcal{C}(C_H, C_F))$ resulting from the solution to this problem. Furthermore, the resulting allocation corresponds to the solution to the original constrained problem (13) given the value of $\bar{U}$ equal to $U(\mathcal{C}(C_H, C_F))$. Therefore, $\bar{U}$ is weakly increasing in λ for a given value of δ, and we denote this mapping with $\bar{U}(\lambda, \delta)$. All this relationships are continuous and monotonically increasing when the utility possibility frontier is concave. Importantly, for any δ and any λ, including $\lambda = 0$, there exists a corresponding $\bar{U}$, which allows us to define the loss function $\mathbb{L}(\lambda, \delta)$ used in the proof

of the proposition. ■

A.6 Proof of Proposition 2

By definition,

$$\frac{P_H}{P_F} = \frac{P_H^*}{P_F^*}\frac{1+\mu_F}{1+\mu_H} \qquad \Rightarrow \qquad \frac{C_H}{C_F} = \left[\frac{1+\mu_F}{1+\mu_H}\right]^{\theta}\left(\frac{1-\gamma}{\gamma}\frac{1-\gamma^*}{\gamma^*}\right)\frac{C_H^*}{C_F^*}, \tag{27}$$

where the right-hand side expression is obtained using (16) and (18). Comparing this expression with (8), we have the first result. The second result is obtained by taking the ratio of (12) to (22) and using our assumed normalization $\omega = \xi^*/\xi$. ■

A.7 Proof of Proposition 3

Combine (16) and (18) with (21) after imposing $\mu_H = \mu_F = \mu$ to obtain:

$$\frac{C_H}{C_F}\frac{\gamma}{1-\gamma} = \left(\frac{P_F}{P_H}\right)^{\theta} = \left(\frac{P_F^*}{P_H^*}\right)^{\theta} = \frac{C_H^*}{C_F^*}\frac{1-\gamma^*}{\gamma^*}.$$

This allocation corresponds to $\delta = 0$ in the constraint (8) on the Planner's problem (13), which implies that the constraint is slack: the optimality conditions in the Planner's problem with and without the constraint (8) coincide when $\delta = 0$. As the resource constraints (3) in the planner's problem coincide with the market clearing conditions, $\mu_H = \mu_F = \mu$ and $\delta = 0$ imply that the relative prices in the decentralize allocation must reflect the relative Lagrange multipliers ν and ν^* on the resource constraints (3) in the planner's problem (13), which we denote $\tilde{\mathcal{S}} \equiv \nu^*/\nu = P_F/P_H = P_F^*/P_H^*$. This is directly verified from the optimality conditions to the Planner's allocation when $\delta = 0$. Finally, we substitute $\tilde{\mathcal{S}}$ into the expression for the planner's MRT (from the planner's optimality conditions[22]) and the definition of the real exchange rate (20):

$$\tilde{Q} = \left(\frac{\gamma^* + (1-\gamma^*)\tilde{\mathcal{S}}^{1-\theta}}{(1-\gamma) + \gamma\tilde{\mathcal{S}}^{1-\theta}}\right)^{\frac{1}{1-\theta}} \qquad \text{and} \qquad Q = (1+\mu)\left(\frac{\gamma^* + (1-\gamma^*)\tilde{\mathcal{S}}^{1-\theta}}{(1-\gamma) + \gamma\tilde{\mathcal{S}}^{1-\theta}}\right)^{\frac{1}{1-\theta}},$$

[22] Write out the definitions (1) and (2) as constraints $C = C(C_H, C_F)$ and $C^* = C^*(C_H^*, C_F^*)$ on the planner's problem (13) with associated Lagrange multipliers μ and μ^*. The first order conditions with respect to C and C^* yield $\omega(1+\lambda)U'(C) = \mu$ and $U'(C^*) = \mu^*$, which by the definition of the λ wedge (12) implies that $\tilde{Q} = \mu^*/\mu$. The first order conditions with respect to C_H and C_F after dropping the static wedge constraint (8) are $\mu\partial\mathcal{C}/\partial C_H = \nu$ and $\mu\partial\mathcal{C}/\partial C_F = \nu^*$, and by the homotheticity of $\mathcal{C}$ aggregator we have $\mu C = \nu C_H + \nu^* C_F$. Using consumption aggregator (1), the optimality also implies that $C_H/C = (1-\gamma)(\mu/\nu)^{\theta}$ and $C_F/C = \gamma(\mu/\nu^*)^{\theta}$. Similarly, for foreign consumption allocation we have that $\mu^* C^* = \nu C_H^* + \nu^* C_F^*$ and $C_H^*/C^* = \gamma^*(\mu^*/\nu)^{\theta}$ and $C_F^*/C^* = (1-\gamma^*)(\mu^*/\nu^*)^{\theta}$. Combining together these insights, we have that

$$\frac{\mu^*}{\mu} = \frac{C_H^*/C^* + (\nu^*/\nu)\cdot C_F^*/C^*}{C_H/C + (\nu^*/\nu)\cdot C_F/C} = \frac{\gamma^*(\mu^*/\nu)^{\theta} + (\nu^*/\nu)(1-\gamma^*)(\mu^*/\nu^*)^{\theta}}{(1-\gamma)(\mu/\nu)^{\theta} + (\nu^*/\nu)\gamma(\mu/\nu^*)^{\theta}} = \left(\frac{\mu^*}{\mu}\right)^{\theta}\cdot\frac{\gamma^* + (1-\gamma^*)(\nu^*/\nu)^{1-\theta}}{(1-\gamma) + \gamma(\nu^*/\nu)^{1-\theta}},$$

resulting in the expression used in the proof after expressing $\tilde{Q} = \mu^*/\mu$ as a function of $\tilde{\mathcal{S}} \equiv \nu^*/\nu$.

where we also used the fact that the LOP deviation $P_H^*/P_H = 1+\mu$ in the second expression. We immediately see that $Q/\tilde{Q} = (1+\mu)$, and the result follows by Proposition 2. ■

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

$$\hat{A}_t \equiv \frac{A_t}{\bar{A}} = \hat{C}(\hat{A}_{H,t}, \hat{A}_{F,t}) \equiv \left((1-\bar{\gamma})\hat{A}_{H,t}^{\frac{\theta-1}{\theta}} + \bar{\gamma}\hat{A}_{F,t}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}, \tag{28}$$

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

$$\hat{Y}_t = \frac{\bar{A}_H}{\bar{Y}}\hat{A}_{H,t} + \frac{\tau\bar{A}^*_H}{\bar{Y}}\hat{A}^*_{H,t} = (1-\bar{\alpha})\hat{A}_{H,t} + \bar{\alpha}\hat{A}^*_{H,t}, \tag{29}$$

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,

$$\frac{P_{H,t}C_{H,t}}{P_t C_t} = \frac{P_{H,t}A_{H,t}}{P_t A_t} \quad \Rightarrow C_{H,t} = \frac{C_t}{A_t} A_{H,t} \quad \Rightarrow \hat{C}_{H,t} = \frac{\hat{C}_t}{\hat{A}_t} \hat{A}_{H,t} \tag{30}$$

$$\frac{P_{F,t}C_{F,t}}{P_t C_t} = \frac{P_{F,t}A_{F,t}}{P_t A_t} \quad \Rightarrow C_{F,t} = \frac{C_t}{A_t} A_{F,t} \quad \Rightarrow \hat{C}_{F,t} = \frac{\hat{C}_t}{\hat{A}_t} \hat{A}_{F,t} \tag{31}$$

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

$$\hat{A}_{F,t} \equiv \frac{1}{\bar{\alpha}^*} \sum_{j\in\mathcal{F}} \omega_j^Y \bar{\alpha}_{ji}^* \hat{A}_{ji,t}, \tag{32}$$

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):

$$\hat{A}_{F,t} = \frac{1}{\bar{\gamma}} \sum_{j\in\mathcal{F}} \gamma_{ji} \hat{A}_{ji,t}, \tag{33}$$

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 (32) equals (33) 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:

$$\bar{\alpha}^* = \frac{\overline{IMP}}{\sum_{k\in\mathcal{F}} \overline{GDP}_k^{PPP}}. \tag{34}$$

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.[23] 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

$$\hat{A}_t^* = \sum_{j\in\mathcal{F}} \omega_j^A \gamma_{ij}\hat{A}_{ij,t} + \sum_{j\in\mathcal{F}} \omega_j^A \sum_{k\in\mathcal{F}} \gamma_{kj}\hat{A}_{kj,t}. \tag{35}$$

[23] In practice, we use ω_j^Y instead of ω_j^A to aggregate countries' absorptions into the RoW series. The two weights are highly correlated, but output-based one allows us to get a larger sample of countries with solution for $\hat{A}_{H,t}, \hat{A}_{F,t}, \hat{A}_{H,t}^*, \hat{A}_{F,t}^*$ in every period.

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

$$\hat{A}^*_{H,t} = \sum_{j\in\mathcal{F}} \alpha_{ij}\hat{A}_{ij,t}, \tag{36}$$

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.[24] In order to write (35) as

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

we search for weights such that

$$\hat{A}^*_{H,t} = \frac{1}{\bar{\gamma}^*}\sum_{j\in\mathcal{F}} \omega^A_j\gamma_{ij}\hat{A}_{ij,t}, \tag{37}$$

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

Calibrating $\bar{\gamma}^*$ and ω^A_j: Comparing (36) to (37), 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:

[24] 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 (30).

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

$$\hat{Y}_{C,t} = (1-\bar{\alpha}_C)\hat{C}_{H,t} + \bar{\alpha}_C\hat{C}^*_{H,t}. \tag{38}$$

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

$$\hat{Y}^*_{C,t} = \bar{\alpha}^*_C\hat{C}_{H,t} + (1-\bar{\alpha}^*_C)\hat{C}^*_{H,t}. \tag{39}$$

Computing Planning Problem Wedges

We begin with computing the (distorted) Consumption Possibility Frontier of Figure 1 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

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