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

December 10, 2025

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, Diego Perez, 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. That is, has the recent wave of globalization that witnessed large movements of goods and assets across borders realize the inter-temporal and static gains from trade. 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 marginal utilities.[1] The canonical study by Backus and Smith (1993) started a literature that reaches a pessimistic conclusion. Importantly, that paper 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 simple planning problem involving two countries, Home and Foreign, with distinct preferences over two traded inputs. The environment builds on Armington (1969), in which each country produces distinct goods. Our exercise takes world output as given, and hence we focus on consumption efficiency given available output rather than production efficiency. We first ask what is the most efficient manner of delivering consumption to Home given global output and subject to delivering a certain level of consumption to Foreign. The shadow cost (i.e., the Lagrange multiplier) of this last constraint is the planner's price of Foreign consumption in terms of Home consumption. This is our conceptual notion of the social marginal rate of transformation (MRT). For given output, as we vary the promised consumption to Foreign, we trace out the Consumption Possibility Frontier (CPF) between Home and Foreign. At any point on this frontier, the slope of the CPF corresponds to the planner's MRT.

We show how to map observed sequences of real output and consumption to the MRT of consumption between two regions, which is the planner's "real exchange rate." 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. 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 MRT 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. That is, whether the planner's optimality condition equating the relative movement in consump-

[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).

tion to the change in the MRT is satisfied. To the extent it is violated, there is scope for an inter-temporal trade between regions that generates a Pareto improvement. We refer to violations of the planner's dynamic optimality as the "dynamic" wedge. Time variation in the dynamic wedge is equivalent to observing a time-varying Pareto weight implied by the observed path of relative consumption and the MRT.

We introduce a decentralized equilibrium to map the wedge into distortions in asset and goods markets. For private agents solving their household maximization problem, the relevant price of consumption is the retail price and the relevant exchange rate is the market exchange rate. Hence, the private risk-sharing wedge is that measured by Backus and Smith (1993) using observed real exchange rates, henceforth referred to as the "Backus-Smith" (BS) wedge. The BS wedge is the focus of a large literature on exchange rate puzzles. In the decentralization, the BS wedge corresponds to differences in the (common currency) stochastic discount factors between Home and Foreign agents. Frictions like incomplete markets, capital controls and taxation of capital flows, or barriers to asset trades (segmented or intermediated markets) give rise to a BS wedge. However, the planner can "see through" the prices faced by consumers to the true social costs. In particular, the difference between the planner's dynamic wedge and the BS wedge equals the difference between the observed real exchange rate and the MRT. The empirical question becomes whether this latter gap behaves in such a way as to "undo" the risk-sharing frictions implied by the Backus-Smith wedge.

Another important friction in our decentralization is the extent to which the same good sells for different prices in the two regions. If this "law of one price" (LOP) deviation exceeds physical trade costs, then the planner disagrees with consumers regarding the marginal cost of moving the good across borders. If the LOP deviations are common across goods, we show that the planner's dynamic wedge is the product of the LOP deviation and the BS wedge. This relationship yields a number of insights. First, if there are no LOP deviations, then the planner's dynamic wedge equals the BS wedge. Alternatively, if the LOP deviation and the BS wedge are perfectly negatively correlated with each other, then there is no planning dynamic wedge even in the presence of large BS wedges. In this case, the distortions in the foreign currency market are completely undone by LOP deviations in the goods market. Finally, if there is no BS wedge, then the planner's dynamic wedge is equal to the LOP deviation. This shows that perfect risk sharing can be consistent with either the presence or absence of a BS wedge, depending on the nature of LOP deviations.

The fact that the price of the same good differs across two markets is familiar from the "alphabet soup" of pricing-to-market frictions considered in the literature: pricing to market (PTM), producer currency pricing (PCP), local currency pricing (LCP), and dominant currency pricing (DCP). These pricing protocols have different implications for the comovement of LOP deviations and the observed real exchange rate, and hence the distortions to the planner's dynamic optimality.

In particular, PCP implies zero LOP deviations and hence equality between the BS and planner's dynamic wedges. In this case, the methodology of Backus and Smith (1993) accurately measures Pareto efficient risk sharing. On the other hand, LCP implies that LOP deviations are perfectly correlated with nominal exchange rate movements. If nominal exchange rates are "disconnected" from fundamentals (i.e., output), then the BS wedge will be completely undone by the LOP deviations, at least in the short run, insulating efficient risk sharing from the disconnect.

Another implication of the planning problem is that the static gains from trade are fully realized: that is, the marginal utility of any good is equalized across regions (up to trade costs). We define the "static" wedge as differences in the marginal rates of substitution (MRS) for different goods across regions. Absent variation in trade costs, the relative shadow price of the goods move in tandem across regions, and hence a static optimality condition equates the MRS between domestic and foreign goods in each region. The planner's static wedge is then defined by the gap between the two regions' MRS, which can be inferred from the observed allocations. The presence of the static wedge implies that the gains from trade are not fully exploited within a period. In the decentralization, the static wedge reflects the extent that consumers in Home face different relative prices than those in Foreign. This gap is equal to the ratio of the two LOP deviations: that is, the deviation in the price of the Home good across the two markets divided by the deviation in the price of the Foreign good. Hence, our static wedge corresponds to a double log difference, or a *relative* relative price.

Interestingly, the welfare consequences of the two distortions — static and dynamic — are orthogonal in the following sense: a second-order approximation to the welfare loss around the efficient allocation is additively separable in the (squared) dynamic and static wedges. Furthermore, we show that at any level of the static distortion, it is always welfare improving to completely eliminate fully or partially the dynamic wedge. This is a surprising result from the perspective of the theory of the second best, emphasizing the orthogonality property of our decomposition of distortions into the dynamic and static wedges.

Our methodology allows us to measure the wedges between the planner's allocation and the observed equilibrium using standard data from the World Development Indicators (WDI). We do so for a 20-year period from 2000 to 2019.[2] For each country, we construct a two-region environment by aggregating the rest of the world using relative PPP-adjusted GDP. The parameterization of preferences rests on several key assumptions. First, we assume commonly used functional forms: specifically, the aggregator combining domestically produced goods and foreign goods has a constant elasticity of substitution and flow utility over the composite has a constant elasticity

[2]We truncate at 2000 in order to include a wide set of countries, particularly China; it also reflects our assumption that trade costs and preference for imports are stable within the sample. The truncation at 2019 is to omit the global pandemic.

of inter-temporal substitution. We assume these two elasticities are common across the two regions and set them to standard values. A key preference parameter is the extent of home bias in consumption, which we calibrate in a base year (which we take to be 2019). With the calibrated preference parameters in hand, the measurement exercise for the remaining years uses only the series for real output and consumption. The dynamic wedges are measured as changes relative to 2019, allowing us to difference out the unobserved Pareto weights on each country.

We organize our results by frequency, contrasting longer trends with annual first-differences. At annual frequencies, the average MRT has a standard deviation that is one third that of the observed real exchange rate. In particular, the MRT has a volatility similar in size to that of relative consumption and output growth, while the observed real exchange rate suffers from well documented "excess" volatility. Remarkably, the average correlation of relative consumption at annual frequencies with the MRT measure is 0.56 versus −0.14 for the real exchange rate (see Figure 4 in Section 5). Furthermore, the standard deviation of the planner's dynamic wedge is almost half that of the Backus-Smith measure (0.06 vs 0.10).

At longer horizons, however, there are large deviations from risk-sharing (or inter-temporal) efficiency. The mean absolute value of the log change in the dynamic wedge over the full sample is 50 log points. This reflects that the magnitude of average growth in relative consumption over the twenty year sample is 26 log points while the change in the MRT is only 7 log points. The large gap in consumption growth across regions reflects that we have in our sample economies like China, that outpaced the rest of the world, and Japan, that lagged behind. On the other side, the stability of the MRT over twenty years reflects that our implied shadow cost of relative consumption is quite stable even in the presence of large movements in differential output growth. More broadly (and intuitively), the results suggest that insurance is better than that implied by high frequency movements in the real exchange rate, but longer run trend growth differences are not well insured across countries, leading to the low-frequency comovement of domestic consumption with domestic output.

To assess the implications for welfare, we compute several counter-factuals. Specifically, we compute the home country's welfare gain (while holding constant welfare in the rest of the world) from eliminating the country's dynamic wedge, i.e., the inter-temporal misallocation between this country and the rest of the world. The median gain in consumption equivalents is 1.1 percent. An equivalent exercise that eliminates the Backus-Smith wedge for individual risk sharing implies a welfare gain of 1.4 percent. The small size of this difference reflects in part that long-run risk is not well shared empirically. To isolate the contribution to welfare from better risk sharing at higher frequencies, we also compute the welfare gain from eliminating each wedge's deviation from a log-linear trend. Here, the median gain from eliminating the planning wedge is only 0.3 percent, versus 0.7 percent for the BS wedge.

Finally, we can map the observed allocation into implied prices using our parameterized preferences and the equilibrium conditions from the decentralization. The results indicate that implied relative prices within each region are relatively stable, while LOP deviations track the observed real exchange rate closely. The implication is that consumers in each region are insulated from exchange rate volatility by offsetting movements in cross-border LOP deviations. This is consistent with the empirical facts documented by Engel (1999) and the literature on LOP deviations that followed, and echoes the risk-sharing properties we derive for the local currency pricing (LCP) protocol when exchange rates are driven by financial-market shocks. In brief, much of the exchange rate's excess volatility at high frequency is absorbed by LOP movements (and implicitly the markups of importers and exporters) and not passed onto consumers, while longer run differences in trend growth are not well insured and are not strongly reflected in movements in either the observed real exchange rate or the planner's MRT.

Related Literature Our paper speaks to work on international (mis)allocation and exchange-rate-based tests of risk sharing. On the consumption side, a large literature beginning with Backus and Smith (1993) documents weak comovement between relative consumption and the real exchange rate.[3] Closest to us, Fitzgerald (2012) estimates model-implied price indices and bilateral real exchange rates from a gravity framework, finding more favorable risk sharing than suggested by observed real exchange rates. We similarly replace observed prices with model-implied shadow costs, but we rely on real macroeconomic quantities rather than bilateral trade values, and we provide a decentralization that rationalizes why observed prices need not track social costs.

Our decentralization connects to asset-pricing and portfolio approaches to international risk sharing (e.g., Baxter and Jermann, 1997; Cole and Obstfeld, 1991; Heathcote and Perri, 2013; Coeurdacier and Rey, 2013; Farhi and Werning, 2017; Lewis and Liu, 2022). The primary distinction between our work and these previous studies is that we construct the planner's real exchange rate (MRT) to assess the quality of risk sharing. Our analysis of wedges between a planner's allocation and the empirical allocation relates to a large literature on "wedge accounting" (e.g., Chari et al., 2007; Hsieh and Klenow, 2009; Capelle and Pellegrino, 2023; Kleinman et al., 2023).

On the exchange-rate side, we relate to work on segmented financial markets and exchange-rate determination (e.g., Alvarez et al., 2002; Jeanne and Rose, 2002; Kollmann, 2005; Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021, 2025a) and to the extensive literature on goods-market frictions, pricing-to-market, and incomplete pass-through (e.g., Rogoff, 1996; Engel, 1999; Devereux and Engel, 2003; Atkeson and Burstein, 2008; Amiti et al., 2019; Burstein and Gopinath,

[3] See also Kollmann (1995); Backus et al. (1992); Lewis (1996); Aguiar and Gopinath (2007); Corsetti et al. (2008); Bai and Zhang (2010, 2012); Heathcote and Perri (2013); Ohanian et al. (2018); Corsetti et al. (2023) and the survey in Heathcote and Perri (2014).

2012; Itskhoki, 2021; Gopinath and Itskhoki, 2021). We show that frictions in asset and goods markets can offset each other in equilibrium, so that observed real exchange rates provide a poor proxy for the social MRT relevant for assessing international risk sharing.

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. Throughout the analysis, we assume that the environment features *home bias* in preferences, which is defined as follows:

Assumption 1 (Home Bias). *Preferences satisfy:* $\gamma + \gamma^* < 1$.

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:

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

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, $\tilde{Q} = -\partial \mathrm{C}/\partial C^*$ 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 under Assumption 1, condition (7) implies that $C_H/C_F > C_H^*/C_F^*$.

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.[4] 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.

The following result provides an additional insight about the slope of the efficient CPF, $\tilde{Q}$:

Proposition 1. *Suppose allocation* $\{C_H, C_F, C_H^*, C_F^*\}$ *solves Problem* (5) *and let* $\tilde{Q}$ *denote the multiplier on the final constraint. Then:* (i) $\tilde{Q}$ *is strictly increasing in* C_H/C_F *and* C_H^*/C_F^**; and* (ii) *holding constant* C/C^* *(i.e., along a ray from the origin in Figure 1b), an increase in* Y/Y^* *implies an increase in* $\tilde{Q}$.

This proposition gives a sense of how home bias influences the marginal rate of transformation. As Y/Y^* increases, Home's good becomes more abundant. In response, at each point along the CPF, each region's allocation becomes more intensive in the Home good. As Home's con-

[4] 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.

sumption aggregator uses the Home good intensively, it becomes relatively cheaper to deliver C and more expensive to deliver C^*, raising the shadow cost of the latter. In the decentralization of Section 3, we map ratios C_H/C_F and C_H^*/C_F^* to relative prices and leverage this result to show how pricing frictions affect the implied MRT.

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 of pricing map into δ in a decentralized equilibrium.[5]

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* contract curve, along which the MRS distortion is held constant at $1+\delta$. This is the solid locus in the Edgeworth box.[6] In Section 3, we discuss why a constant δ arises as a natural constraint in the second-best risk sharing problem in a wide range of standard models.

[5]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 frictions.

[6]Aguiar et al. (2025) and Itskhoki and Mukhin (2025b) 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.

Formally, consider an extended planning problem that solves Problem (5) with the additional constraint (8):[7]

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

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.[8] This reflects that in both cases, one country's consumption goes to zero and the other country consumes 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:[9]

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

[7]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). To ensure $\tilde{Q}$ remains well defined, we deal with this technicality by replacing the static wedge constraint with an inequality in the formal analysis and proofs. See Appendix A for details.

[8]To see this, as we approach the horizontal intercept, we can achieve (in the limit) the efficient allocation $C \to \mathcal{C}(Y,Y^*)$ and $C^* \to \mathcal{C}^*(0,0) = 0$, 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.

[9]Note that we treat δ as a constraint on the distorted planning problem. As we discuss in Section 3, this assumes that shifts in aggregate consumption do not affect the *double*-difference in prices across borders or, equivalently, the frictions in exploiting the gains from trade. Many models of pricing frictions, including PTM, PCP and LCP, respect this property.

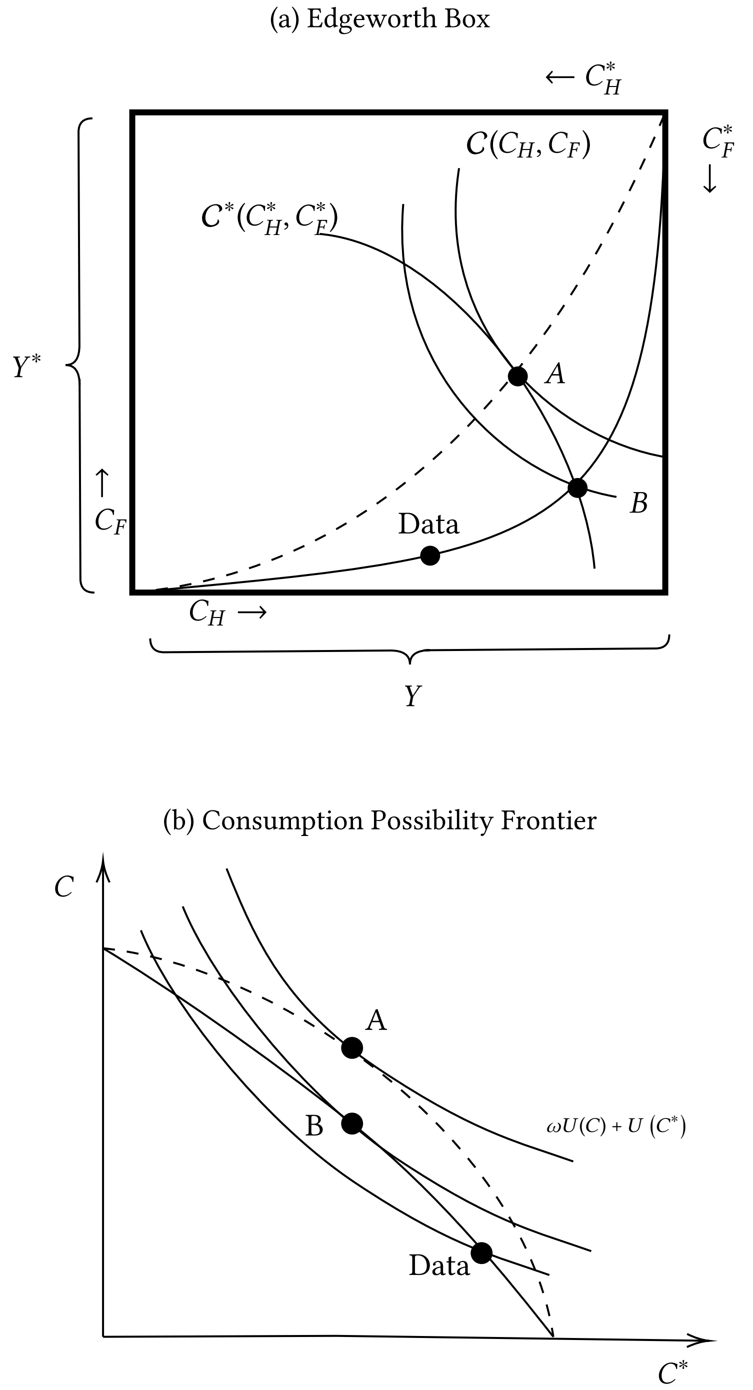

Figure 1: Geometric Representation of Constrained Efficient Allocations

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 $\delta > 0$, 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, $\lambda < 0$. 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, which we suppress in the notation, is $\omega U'(\mathrm{C}) \cdot \partial \mathrm{C}/\partial C^* + U'(C^*) = 0$. Equivalently, using the property that $\tilde{Q} = -\partial \mathrm{C}/\partial C^*$, 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.[10]

In Figure 1b 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 relative slopes between the CPF and the indifference curve determines λ.

[10] 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 sign of the net increase in Home welfare is determined by $(\lambda_t - \lambda_0)\Delta > 0$. 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 4 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 in (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. Finally, if the two indifference curves do not intersect inside the Edgeworth box, then we have violated feasibility, that is, 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.[11]

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. 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. So long as the resulting Data point is inside the first-best CPF, there exists a unique solution for (λ, δ) with $\delta > 0$.

[11] In our empirical estimation in Section 4, we discuss when such pathologies arise in the data, and we also bring additional data on trade flows to identify the value of δ_t in the baseline year.

The Cole-Obstfeld Case The identification procedure described above and implemented in the data in Section 4 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 and Obstfeld (1991) 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 $\sigma = \theta = 1$. Then there exist strictly positive functions of δ and λ, g and g^*, such that:*

$$C = g(\lambda, \delta)\, Y^{1-\gamma} Y^{*\gamma} \qquad \text{and} \qquad C^* = g^*(\lambda, \delta)\, Y^{\gamma^*} Y^{*1-\gamma^*},$$

and therefore $\tilde{Q} = \frac{1}{\omega(1+\lambda)} \frac{g(\lambda,\delta)}{g^*(\lambda,\delta)} \, (Y/Y^*)^{1-\gamma-\gamma^*}$. Moreover, g is strictly increasing in λ; g^* is strictly decreasing in λ; and g and g^* are strictly decreasing in $|\delta|$ for $\delta \neq 0$.

The derivative of C and C^* with respect to δ reflects that as δ moves away from zero in either direction, the static distortion intensifies and less consumption can be generated from the same output. The relationship of the respective consumptions with respect to λ reflects that λ shifts consumption away from Foreign and toward Home. The lemma states that for given wedges, C and C^* expand with both outputs, while 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.[12] 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.

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

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

and depending on parameters the consumption correlation may be smaller, equal to, or larger than the output correlation (see Heathcote and Perri, 2014).

2.4 Welfare Losses

The decomposition of the distance from the first-best allocation into static and dynamic wedges provides an intuitive "two-stage budgeting" interpretation: the planner would like to ensure the static gains from trade are exploited, and given whatever static distortion remains, she will equate the relative marginal utilities to the marginal rate of transformation. To see this dichotomy more explicitly, we can shed some light on how this "separability" arises.

For a given Y and Y^*, define

$$\tilde{C}^*(\lambda, \delta) \equiv \arg\max_{C^*} \left[\omega(1+\lambda)U(\mathbf{C}(C^*; Y, Y^*, \delta)) + U(C^*)\right].$$

The first-order condition for this problem is

$$U'(C^*) = \omega(1+\lambda)U'(\mathbf{C}) \cdot \tilde{\mathcal{Q}},$$

where we used the definition of $\tilde{\mathcal{Q}}(C^*; Y, Y^*, \delta) \equiv -\partial\mathbf{C}/\partial C^*$. Note that $\tilde{C}^*(0,0)$ is the first-best allocation, and $\tilde{C}^*(0,\delta)$ is the second-best allocation that maximizes the planner's objective given the static wedge. We define a period's welfare deviation from the first-best allocation as:

$$\mathbb{L}(\lambda, \delta) \equiv \left[\omega U(\mathbf{C}(\tilde{C}^*(0,0))) + U(\tilde{C}^*(0,0))\right] - \left[\omega U(\mathbf{C}(\tilde{C}^*(\lambda,\delta))) + U(\tilde{C}^*(\lambda,\delta))\right], \tag{13}$$

where we have suppressed the additional arguments in the $\mathbf{C}$ functions (keeping in mind that the static wedge is zero in the first bracketed term and δ in the second). By construction, $\mathbb{L}(\lambda, \delta) \geq 0$ and $\mathbb{L}(0,0) = 0$. Also, we have for any δ

$$\begin{aligned} \frac{\partial \mathbb{L}(\lambda,\delta)}{\partial \lambda} &= -\left[\omega U'(\mathbf{C})\frac{\partial \mathbf{C}}{\partial C^*} + U'(\tilde{C}^*)\right]\frac{\partial \tilde{C}^*(\lambda,\delta)}{\partial \lambda} \\ &= \lambda \omega U'(\mathbf{C})\frac{\partial \mathbf{C}}{\partial C^*}\frac{\partial \tilde{C}^*(\lambda,\delta)}{\partial \lambda} \gtreqless 0 \iff \lambda \gtreqless 0, \end{aligned}$$

where the second equality uses the first-order conditon from the argmax problem in the definition of $\tilde{C}^*$ and the last statement uses the fact that $\partial\mathbf{C}/\partial C^* < 0$ and $\partial\tilde{C}^*/\partial\lambda < 0$.[13] Hence, $\mathbb{L}(\cdot, \delta)$ has a global minimum at $\lambda = 0$ for every δ:

Proposition 2. *For any value of the static distortion δ, the welfare loss from the dynamic distortion is minimized when $\lambda = 0$:* $\arg\min_\lambda \mathbb{L}(\lambda,\delta) = 0, \forall\delta$.

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

[13] The second-order condition for the $\tilde{C}^*$ problem implies that $\tilde{C}^*$ is strictly decreasing in λ given δ.

with general markup wedges.[14] This is not the case under our (λ, δ) representation, motivating 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, a constrained planner has no additional incentive to deviate from (11), and would choose $\lambda = 0$. In other words, $\lambda \neq 0$ necessarily implies the presence of a distortion to intertemporal risk sharing.[15]

A corollary of Proposition 2 is that a second-order expansion of the welfare loss around the efficient allocation, $\lambda = \delta = 0$, 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 2.

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

$$\mathbb{L}(\lambda, \delta) = \theta\gamma(1-\gamma)\Big[\kappa\lambda^2 + \frac{1}{4}\delta^2\Big] \qquad \text{with} \quad \kappa \equiv \frac{1}{1 + 4\gamma(1-\gamma)(\sigma\theta - 1)}, \tag{14}$$

where the loss is expressed in units of home flow marginal utility. For standard parameter values, $\sigma\theta \geq 1$, 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. The squared terms λ^2 and δ^2 have a Harberger triangle intuition. Recall that λ is like a tax on inter-temporal trades and δ 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 the tax elasticity of the consumption allocation (see appendix).

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

[14] There are other decompositions that can be used to characterized the distance to the first-best. For example, another way to quantify distortions are the product-region-specific markup wedges following Baqaee and Farhi (2020). 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.

[15] This separation rests on two assumptions. The first is homotheticity of preferences, so the MRS between Home and Foreign goods is independent of aggregate consumption. The second is the assumption that δ does not vary with aggregate consumption, which is equivalent to assuming that frictions to bilateral trade (i.e., trade impediments and differential relative price distortions, as we discuss below) are not sensitive to the level of total consumption.

allocation. The second task is to understand what type of environments endogenously generate the price and wedge behavior observed in the data. In particular, how the market and pricing frictions discussed in the literature can be used to interpret the observed allocations we present in Section 5.

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 $\{P_{Ht}(s^t), P_{Ft}(s^t)\}$ 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 convenient for the exposition, we discuss 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-Debreu security in Home that pays one unit of the Home's currency in history s^t. The representative Home agent has an initial wealth w and an endowment stream $\{Y_t(s^t)\}$, and solves the following problem:[16]

$$\max_{\{C_t(s^t)\}} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t(s^t) U(C_t(s^t)) \tag{15}$$

$$\text{subject to} \quad w \geq \sum_{t} \sum_{s^t} \Lambda_t(s^t) \left(P_t(s^t)C_t(s^t) - P_{Ht}(s^t)Y_t(s^t)\right).$$

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

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

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 strictly monotonically decreasing in the initial wealth w.

[16] The initial wealth consists of any claims on Foreign residents as well as claims to rebated tax revenues and/or profits arising from the wedges in foreign exchange markets and goods markets discussed below.

The Foreign agent's problem is symmetric. It faces prices $\{P^*_{Ht}(s^t), P^*_{Ft}(s^t)\}$ for the Home and Foreign good in Foreign currency, with an associated CES ideal price index:

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

The nominal exchange rate as $\mathcal{E}$ is the amount of Home currency needed to purchase one unit of Foreign currency. In turn, the real exchange rate $\mathcal{Q}$ is the ratio of Foreign to Home price levels:

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

which depends on local relative prices, P_F/P_H and P^*_F/P^*_H, as well as on the deviation from the Law of One Price (LOP) for the Home good, $\mathcal{E}P^*_H/P_H$, as consumers in the two countries may pay a different amount for the same good.[17]

The Foreign agent faces Arrow-Debreu prices $\Lambda^*_t(s^t)$ expressed as the period $t = 0$ Foreign currency price of one unit of Foreign currency delivered at t in history s^t. Note that Foreign and Home residents may face different prices for identical securities and goods even when converted to the same currency. Given an initial wealth w^* and a stream of endowments $\{Y^*_t(s^t)\}$, the Foreign household solves:

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

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

$$\frac{C^*_H}{C^*} = \gamma^*\left(\frac{P^*_H}{P^*}\right)^{-\theta}, \qquad \frac{C^*_F}{C^*} = (1-\gamma^*)\left(\frac{P^*_F}{P^*}\right)^{-\theta} \tag{19}$$

and

$$\beta^t \pi_t(s^t) U'(C^*_t(s^t)) = \xi^* \Lambda^*_t(s^t) P^*_t(s^t), \tag{20}$$

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

[17] Note that $\mathcal{Q}$ in (18) can be equivalently rewritten in terms of the local relative prices and the LOP deviation for the Foreign good, $\mathcal{E}P^*_F/P_F$. Two other aggregate relative prices, the terms of trade $P_F/(\mathcal{E}P^*_H)$ and $\mathcal{E}P^*_F/P_H$, can also be recovered from the three relative prices that we focus on, but they are not needed for our empirical implementation.

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 of observed allocations $\{C_{Ht}, C_{Ft}, C_t\}$, we can use (16) to infer the sequence of relative prices $\{P_{Ht}/P_t, P_{Ft}/P_t\}$. Similarly, from (17) we can infer the sequence of implied Arrow-Debreu prices, or more precisely $\xi\Lambda_t(s^t)P_t(s^t)/\pi_t(s^t)$. Similarly, from (19) and (20) we can infer $\{P^*_{Ht}/P^*_t, P^*_{Ft}/P^*_t\}$ and $\xi^*\Lambda^*_t(s^t)P^*_t(s^t)/\pi_t(s^t)$. In addition to the observed quantities, we observe the real exchange rate $\{\mathcal{Q}_t\}$ defined in (18), which together with the local relative prices allows us to construct LOP deviations for both goods, P^*_H/P_H and P^*_F/P_F.

3.2 Decentralized Wedges

Given the relative prices of goods and prices of states identified above, we can recover the implied allocation-relevant wedges in the decentralized equilibrium. First, we define the two law-of-one-price deviations for the Home and Foreign goods as:

$$1 + \mu_H \equiv \frac{\mathcal{E}P^*_H}{P_H} \qquad \text{and} \qquad 1 + \mu_F \equiv \frac{\mathcal{E}P^*_F}{P_F}. \tag{21}$$

In other words, the wedges μ_H and μ_F are the premiums (if greater than zero, and discount otherwise) paid by the Foreign consumer relative to the Home consumer for the Home and Foreign goods, respectively.

Second, we define the Backus-Smith wedge (normalizing $\psi_0 = 0$):

$$1 + \psi_t(s^t) \equiv \frac{\Lambda^*_t(s^t)}{\Lambda_t(s^t)}\frac{\mathcal{E}_0(s_0)}{\mathcal{E}_t(s^t)} = \frac{1}{\bar{\omega}}\frac{U'(C^*_t(s^t))}{U'(C_t(s^t))}\frac{1}{\mathcal{Q}_t(s^t)}, \tag{22}$$

where the constant $\bar{\omega} \equiv \frac{1}{\mathcal{E}_0}\frac{\xi^*}{\xi}$ is decreasing in the relative Foreign wealth, $w^*\mathcal{E}_0/w$, and it is the equilibrium counterpart to the Planner's Pareto weight on Home ω. In what follows, we normalize $\omega = \bar{\omega}$ as this has no effect on our measurement and counterfactuals in the following sections. The wedge ψ measures deviations from efficient *private* risk sharing. The first expression in (22) defines it as the deviation across the two regions in the period-0 price of the Home currency delivered in a certain future state s^t at date t, hence the exchange rate conversion for Λ^*. The second equality uses the households' first order conditions (17) and (20) and the definition of the real exchange rate (18) to relate ψ to the wedge that we measure in the data following Backus and Smith (1993): namely, the relative marginal utility of consumption in Foreign versus Home compared to the relative equilibrium price of the two consumption aggregates $\mathcal{Q}$.

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

Proposition 3. (i) *The static wedge δ defined in (8) is related to deviations from the Law of One Price in (21) by:*

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

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

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

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. This 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 foreign prices P_F^* and P_H^* does not change either local relative price P_F^*/P_H^* or P_F/P_H, but doubles the real exchange rate, as can be seen from (18). This, however, does not affect 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 — i.e., the goods market relative prices are undistorted — the two variables are the same, $\tilde{Q}=Q$. To the extent the Welfare Theorems fail, the private and social measures of MRT, Q and $\tilde{Q}$, can be distinct, and hence so can the two dynamic wedges, ψ and λ.

We observe from Part (i) of Proposition 3 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 ψ:

Proposition 4. *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 dynamic wedge λ, the decentralized risk-sharing (Backus-Smith) wedge ψ, and the difference in the level of prices across regions μ. The condition that $\mu_H = \mu_F$ implies that all prices may be uniformly higher in one region, but the relative prices are undistorted, and, in particular, $Q = (1+\mu)\tilde{Q}$. The mapping from LOP deviations to $Q/\tilde{Q}$ becomes analytically intractable when $\delta \neq 0$, however, the general principle that LOP deviations μ_H and μ_F may insulate λ from fluctuations in ψ still applies.[18]

Proposition 4 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 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 4, we have that:

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

as in this case $Q \neq \tilde{Q}$, and, in particular, $Q/\tilde{Q} = 1 + \mu = 1/(1+\psi)$. In other words, frictions in cross-border pricing exactly offset movements in the real exchange rate, resulting in an efficient inter-temporal allocation of consumption despite deviations from perfect private risk sharing. In the next subsection, we show that this implication holds under a popular pricing protocol. Furthermore, it will be a useful benchmark to interpret the empirical results of Section 5.

[18] The proof of Proposition 4 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 complex mapping from (ψ, μ_H, μ_F) to λ, with no closed-form characterization even in the Cole-Obstfeld case. 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.

3.3 Equilibrium Wedges

We complete our description of the decentralized economy with a discussion of alternative models of deviations from perfect risk sharing in the financial market and from the law of one price in the goods market.

Financial Markets We begin by briefly discussing alternative interpretations of the private risk-sharing wedge ψ. The most straightforward interpretation is that there are complete markets subject to region-specific distortionary taxes (e.g., capital controls), which are rebated lump-sum. Thus, Home and Foreign households face different after-tax prices on the Arrow-Debreu securities, resulting in a wedge $\psi_t(s^t)$ between $\Lambda_t(s^t)$ and $\Lambda_t^*(s^t)$, as in Itskhoki and Mukhin (2025c). A similar interpretation is obtained if we introduce segmented markets with intermediation frictions. For example, intermediaries charging premia or mark-ups on asset trades can rationalize the different prices faced by Home and Foreign households for the same security.

More generally, ψ can be interpreted as the difference in marginal utilities that arises from exogenous or endogenous market incompleteness. For example, if only a single risk-free bond paying out in the common numeraire (e.g., in the home currency) is traded, the associated Euler equations in the two regions imply for each period t:

$$\mathbb{E}_t\left[\frac{U'(C_{t+1})}{U'(C_t)}\frac{P_t}{P_{t+1}}\right] = \mathbb{E}_t\left[\frac{U'(C^*_{t+1})}{U'(C^*_t)}\frac{P^*_t\mathcal{E}_t}{P^*_{t+1}\mathcal{E}_{t+1}}\right], \tag{24}$$

where we have cancelled the common discount factor β 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 (see, e.g., Backus et al., 2001; Lustig and Verdelhan, 2019). In incomplete intermediated markets, ψ emerges in response to shifts in demand in the currency and various other asset markets (see, e.g., Gabaix and Maggiori, 2015).[19]

It is important to distinguish the planner's optimal risk sharing condition from that of private agents in partial equilibrium. From the perspective of individual agents, the observed real exchange rate constructed from retail prices is the relevant measure of the relative price of consumption baskets. Hence, the presence of ψ implies that individual Home and Foreign agents would like to engage in additional inter-temporal trade and tilt their consumption profiles to eliminate ψ. However, the planner recognizes that retail prices may be distorted by cross-border markups, and therefore such trades may not bring the allocation closer to the first best.

Pricing Protocols We discuss three familiar pricing protocols with sticky and flexible prices which differ in their implications for the endogenous LOP deviations μ_H and μ_F. Under flexible

[19] Chernov et al. (2024) generalize condition (24) to an arbitrary international asset market structure that nests cases of both (partially) integrated and (frictionally) intermediated asset trading between Home and Foreign households.

price setting, output Y and Y^* can be still viewed as exogenously given endowments, with prices adjusting to ensure market clear. Under sticky prices, output is demand determined. To explore this alternative, we describe a sticky-price extension to our environment.

Specifically, we assume the monetary authority controls nominal aggregate demand. In particular, let $M = PC$ and $M^* = P^*C^*$ denote nominal demand in Home and Foreign, respectively. A useful microfoundation is a cash-in-advance constraint, but one could also think of a nominal interest rate as the instrument of demand management as in standard New Keynesian models. The nominal exchange rate varies with relative demand, but distorted by the presence of the financial wedge ψ. For example, if $\sigma = 1$, condition (22) immediately implies the following relationship:

$$\mathcal{E} = \frac{1}{\bar{\omega}(1+\psi)} \frac{M}{M^*}. \tag{25}$$

The timing within a period is as follows: (i) retailers set prices in each region (described below); (ii) the monetary authorities decide on the nominal money supplies and the foreign exchange distortion/tax ψ is realized; and (iii) production occurs and markets clear. We provide full details of the environment in Appendix B, while Appendix A contains proofs of the propositions.

Producer Currency Pricing To complete the description of the sticky-price environment, we need to specify price-setting behavior. A simple pricing benchmark is when a producer sells its product for a common price across all destinations. This is consistent with a marginal-cost or constant mark-up pricing across destinations in a flexible-price model. In the sticky-price Open Economy New Keynesian literature, it is also known as *producer currency pricing* (PCP) (e.g. Obstfeld and Rogoff, 1995). Under PCP, a producer sets the price of their product in their own currency regardless of the destination market. As a result, the price that consumers face for the good in the destination market is proportional to the nominal exchange rate between the producer and consumer currencies. That is, $P_H^* = P_H/\mathcal{E}$ and $P_F = P_F^*\mathcal{E}$. 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. By conditions (21), this implies that under PCP, $\mu_H = \mu_F = 0$. Furthermore, consumers in both countries face the same relative price that moves one-for-one with the nominal exchange rate, $P_F/P_H = P_F^*/P_H^* = \mathcal{E}P_F^*/P_H$, given the preset producer prices P_F^* and P_H.[20]

[20] For all parts of Proposition 5, we are holding $\bar{\omega}$ constant: that is, we are comparing different values of M/M^* or ψ holding constant the present value of resources in the household budget sets. Given preset prices (P_H, P_F^*), the realized values of (M, M^*, ψ) determine the equilibrium outcomes $(C, C^*, Y, Y^*, \mathcal{Q}, \mathcal{E})$. The resulting output (Y, Y^*) is, in general, inefficient unless $\psi = 0$ and M and M^* are chosen to relax the preset-price constraints, an open-economy version of the divine coincidence (Galí and Monacelli, 2005; Itskhoki and Mukhin, 2023). Our focus, however, is on the efficiency of consumption allocations – which is achieved when $\lambda = \delta = 0$ – given the *observed* output, irrespective of whether output is exogenously given, produced, or demand-determined.

Proposition 5 (a). *Under PCP, the equilibrium features $\mu_H = \mu_F = 0$, $\delta = 0$, and $\lambda = \psi$. Moreover, $\tilde{\mathcal{Q}} = \mathcal{Q}$ and both are increasing in M/M^* and decreasing in ψ.*

PCP is a case in which the planner's MRT equals the observed real exchange rate. The same relative prices in each market ensure that $\delta = 0$, and hence the equilibrium lies on the first-best CPF. Moreover, the LOP deviations are zero, and so private agents and the planner agree on the relative price of the two consumption bundles. This equates the planner's dynamic wedge λ with the private risk-sharing wedge ψ. Given that output is demand determined, Y/Y^* positively co-moves with M/M^*. From Proposition 1, this implies that $\tilde{\mathcal{Q}}$, and hence $\mathcal{Q}$, both increase in M/M^* and decrease in ψ.

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 respective market's currency, and this price is "sticky" in the sense that it does not adjust to the realized value of the nominal exchange rate. This is referred to as *local currency pricing* (LCP) in the Open Economy New Keynesian literature (e.g. Devereux and Engel, 2002). Under this pricing protocol, LOP deviations move one-for-one with the nominal exchange rate. Specifically, suppose that a Home-good producer sets a price of P_H in the Home market and a Foreign-currency price of P_H^* in the Foreign market. Let $\bar{\mathcal{E}} \equiv P_H/P_H^*$ be the ratio of the two, and hence the LOP deviation is $\mathcal{E}P_H^*/P_H = \mathcal{E}/\bar{\mathcal{E}}$. In the literature on LCP, it is typically assumed that the Foreign-good producer sets prices analogously, in particular, $P_F/P_F^* = \bar{\mathcal{E}}$. Thus, a natural benchmark is that $\bar{\mathcal{E}} = \mathbb{E}\mathcal{E}$, where $\mathbb{E}\mathcal{E}$ is the expected exchange rate conditional on information at the time of price setting, making the two LOP deviations zero in expectation. However, LCP is also consistent with producers targeting a higher price in one market versus the other by a uniform proportion. From (21), the realized LOP deviations are the same for both goods, and proportional to the nominal exchange rate: $1 + \mu_H = 1 + \mu_F = \mathcal{E}/\bar{\mathcal{E}}$. Note that the realized exchange rate determines the LOP deviations, but does not change the relative prices faced by consumers in each market, $P_F/P_H = P_F^*/P_H^*$, which are pre-determined under LCP.

Proposition 5 (b). *Under LCP, the equilibrium features $\mu_H = \mu_F \equiv \mu$ such that $1 + \mu \propto \mathcal{Q}$, $\delta = 0$, $\tilde{\mathcal{Q}}$ is constant within the period, $\mathcal{Q} \propto \frac{1}{1+\psi}(M/M^*)^\sigma$, $1 + \lambda = (1 + \mu)(1 + \psi) \propto (M/M^*)^\sigma$, that is λ does not vary with ψ given M/M^*.*

Under LCP, relative prices in each market are pre-set. Hence, relative consumption of Home and Foreign goods in each market is also pre-determined ex ante. By part (i) of Proposition 1, this implies that $\tilde{\mathcal{Q}}$ is constant. In particular, movements in *aggregate* consumption driven by monetary policy are matched by changes in output leaving the MRT constant. In turn, financial shocks ψ are fully absorbed by the nominal exchange rate, and hence $\mathcal{Q}$ and the LOP deviations μ, leaving allocations (as well as the planner's wedge λ) unchanged.[21] The LOP deviations drive a wedge between equilibrium relative prices and the planner's shadow prices, and hence $\tilde{\mathcal{Q}}$ does not equal $\mathcal{Q}$ in general. However, as M/M^* varies, relative consumptions are affected and these changes map directly into λ. This echoes the tradeoff between fixed versus floating exchange rate regimes analyzed in Devereux and Engel (2003).

Flexible Pricing to Market Lastly, we consider a flexible-price model with *pricing to market* (PTM) à la Atkeson and Burstein (2008). In this model, the optimal markup differs across markets for each product, as it depends on the average competitor price in each destination. Home bias (and trade costs) result in different consumer price levels, and the real exchange rate $\mathcal{Q}$ provides an aggregated metric for these differences. Following Amiti et al. (2019), we approximate the strategic price-setting behavior of Atkeson-Burstein and assume that in setting prices for a destination market, firms put weight ϕ on prices of their competitors. In particular, we show in the appendix that $\mathcal{E}P_H^*/P_H = \mathcal{E}P_F^*/P_F = (\mathcal{E}P^*/P)^{\phi}$. Parameter $\phi \in [0,1)$ is the weight on the price index in the target market, and captures the extent of strategic complementarities in price setting. Similarly, $1-\phi$ captures the weight on the common marginal cost, which is converted into the target market's currency using $\mathcal{E}$. The equilibrium LOP deviations are therefore given by $1+\mu_H = 1+\mu_F = \mathcal{Q}^{\phi}$. LOP holds if and only if $\phi = 0$; otherwise, both Home and Foreign LOP deviations track the real exchange rate with an elasticity $\phi > 0$.

Proposition 5 (c)**.** *Under flexible-price PTM with $\phi \in [0,1)$, the equilibrium features $\mu_H = \mu_F \equiv \mu$ such that $1+\mu = \mathcal{Q}^{\phi}$, $\delta = 0$, $\tilde{\mathcal{Q}} = \mathcal{Q}^{1-\phi}$, $1+\lambda = (1+\psi)\mathcal{Q}^{\phi}$, and $\mathcal{Q}$ is strictly increasing in Y/Y^* and decreasing in ψ.*

We see that PTM shares many features of PCP and LCP, depending on the value of ϕ. In particular, at $\phi = 0$, firms set prices based on domestic costs alone. In this case, the planner and private agents agree on the risk-sharing distortions, and all shocks to ψ map into λ, similar to PCP. On the other hand, as $\phi \to 1$, all weight is put on the destination market, and we approach LCP. In this case, financial wedges ψ are increasingly absorbed by LOP deviations μ with vanishing affects on $\tilde{\mathcal{Q}}$ and λ. For interior values of ϕ, the equilibrium shares features of both extremes.

Equilibrium Wedges We can leverage the fact that PTM provides a bridge between PCP and LCP to compactly collect the above results:

[21] In fact, given pre-set prices, if M/M^* is chosen such that $(C/C^*)^{\sigma}$ equals the constant $\omega\tilde{\mathcal{Q}}$, then $\lambda = 0$ regardless of ψ.

Corollary 2. *Under PCP, LCP, and PTM, the equilibrium features $\mu_H = \mu_F \equiv \mu$ and $\delta = 0$, and for each protocol there exists a constant $\bar{\mathcal{Q}}$ such that:*

$$\frac{1+\lambda}{1+\psi} = \frac{\mathcal{Q}}{\tilde{\mathcal{Q}}} = 1 + \mu = \left(\frac{\mathcal{Q}}{\bar{\mathcal{Q}}}\right)^{\phi},$$

with $\phi = 0$ for PCP, $\phi = 1$ for LCP, and $\phi \in [0, 1)$ for PTM. Under all protocols, $\mathcal{Q}$ and, hence, μ are strictly decreasing in ψ.

As we already saw in Corollary 1, in models where LOP holds, including under PCP and flexible prices with $\phi = 0$, the planner's risk-sharing wedge λ equals the Backus-Smith wedge ψ. Under LCP and PTM with $\phi > 0$, LOP violations move in concert and together with the real exchange rate. This keeps the relative prices in sync across the two markets, $P_F/P_H = P_F^*/P_H^*$, resulting in no static wedge, $\delta = 0$.[22] At the same time, the planner's dynamic wedge λ differs from ψ, as the real exchange rate $\mathcal{Q}$ differs from the planner's MRT $\tilde{\mathcal{Q}}$ under LOP deviations. Furthermore, the LOP deviations μ comove with $\mathcal{Q}^{\phi}$, endogenously offsetting the effects of financial wedge ψ on both the planner's MRT $\tilde{\mathcal{Q}}$ and the planner's dynamic wedge λ, rendering both of them more stable relative to their respective decentralized equilibrium benchmarks $\mathcal{Q}$ and ψ.

This last observation about LCP and PTM models is essential for understanding the patterns in the data. Consider models in which the main driver of the exchange rate is variation in the private risk-sharing wedge ψ, as in Itskhoki and Mukhin (2021). For example, under LCP with prices pre-set such that expected LOP deviations are zero, both nominal and real exchange rate innovations are given by:

$$\frac{\mathcal{Q}}{\mathbb{E}\mathcal{Q}} = \frac{\mathcal{E}}{\mathbb{E}\mathcal{E}} = \frac{\frac{1}{1+\psi}(M/M^*)^{\sigma}}{\mathbb{E}\left[\frac{1}{1+\psi}(M/M^*)^{\sigma}\right]}.$$

Therefore, the Home exchange rates appreciate ($\mathcal{Q}$ and $\mathcal{E}$ decrease) with $1 + \psi$ and depreciate ($\mathcal{Q}$ and $\mathcal{E}$ increase) with $(M/M^*)^{\sigma}$. Holding M/M^* constant, variation in $1+\psi$ is then fully offset by endogenous variation in LOP deviations, $1 + \mu \propto \mathcal{Q}$, resulting in fully stable planner's MRT $\tilde{\mathcal{Q}}$ and dynamic wedge λ. By contrast, when exchange rate fluctuations are driven by other shocks, like money or productivity (endowments), the resulting LOP deviations μ translate into the planner's risk-sharing wedge λ even when $\psi = 0$.

[22] All pricing protocols in Proposition 5 (a)-(c) imply $\delta = 0$. This is useful for analytical tractability and building intuition, but is not imposed in our empirical analysis. In particular, we do not impose any pricing-protocol restrictions in Section 5. Thus, we do not rule out alternative pricing protocols such as dominant-currency pricing (DCP; Gopinath et al., 2020), which may distort relative prices across countries. In particular, we study the joint statistical properties of the data-inferred wedges $(\mu_H, \mu_F, \psi, \delta, \lambda)$, and discuss the structural models that are consistent with measured wedges.

4 Measurement

In this section, we summarize our approach to calibration and mesasurement. Full details are contained in Appendix D.

4.1 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 to 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 $\widehat{X}_{it} \equiv X_{it}/\overline{X}_i$ 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 constant-local-price indices provided by WDI to compute real growth rates. For absorption, we deflate the nominal series by the GDP deflators.

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, we construct the GDP weight of country i in the base year (2019) as $w_i^Y \equiv \bar{Y}_i / \sum_{j \in J} \bar{Y}_j$, where J is our sample of 104 countries. The global growth in any variable X is then $G_t^X \equiv \sum_{j \in J} w_j^Y \widehat{X}_{jt}$. Furthermore, for country i, the growth in the ROW is computed as:

$$\widehat{X}_{it}^* \equiv \frac{1}{1 - w_i^Y} (G_t^X - w_i^Y \widehat{X}_{it}).$$

We drop the index i below where it causes no confusion as our analysis is country-by-country with the ROW correspondingly defined in each case.

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:

$$Q_{it} \equiv \sum_{j \neq i} \frac{w_j^Y}{1 - w_i^Y} Q_{ijt},$$

where Q_{ijt} is the real exchange between i and j in year t relative to 2019.[23] Appendix D provides further details.

4.2 Calibration and Measurement

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. Specifically, $A_t = C(A_{Ht}, A_{Ft})$, where A_{Ht} and A_{Ft} is the amount of the Home and Foreign goods, respectively, used to produce A_t. Similarly, $A_t^* = C^*(A_{Ht}^*, A_{Ft}^*)$.

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). We have for the Home country:

$$\widehat{A}_t \equiv \frac{A_t}{\bar{A}} = \frac{C(A_{Ht}, A_{Ft})}{\bar{A}} = C\Big(\frac{A_{Ht}}{\bar{A}}, \frac{A_{Ft}}{\bar{A}}\Big) = C\Big(\frac{\bar{A}_H}{\bar{A}}\widehat{A}_{Ht}, \frac{\bar{A}_F}{\bar{A}}\widehat{A}_{Ft}\Big),$$

where $\bar{A}$ corresponds to the base-year values of the variables. We can rewrite the final term in more compact form by re-defining the aggregator (1) and (2) in terms of growth rates:

$$\widehat{A}_t = \left((1-\bar{\gamma})\widehat{A_{Ht}^{\frac{\theta-1}{\theta}}} + \bar{\gamma}\widehat{A_{Ft}^{\frac{\theta-1}{\theta}}}\right)^{\frac{\theta}{\theta-1}}, \tag{26}$$

$$\widehat{A}_t^* = \left(\bar{\gamma}^*\widehat{A^*{}_{Ht}^{\frac{\theta-1}{\theta}}} + (1-\bar{\gamma}^*)\widehat{A^*{}_{Ft}^{\frac{\theta-1}{\theta}}}\right)^{\frac{\theta}{\theta-1}}, \tag{27}$$

where we define $\bar{\gamma} \equiv \gamma^{\frac{1}{\theta}} \left(\bar{A}_F/\bar{A}\right)^{\frac{\theta-1}{\theta}}$ and $\bar{\gamma}^* \equiv \gamma^{*\frac{1}{\theta}} \left(\bar{A}_H^*/\bar{A}^*\right)^{\frac{\theta-1}{\theta}}$.

Similarly, we can scale the resource constraints (3) as follows:

$$\widehat{Y}_t \equiv \frac{Y_t}{\bar{Y}} = \frac{A_{Ht} + A_{Ht}^*}{\bar{Y}} = (1-\bar{\alpha})\widehat{A}_{Ht} + \bar{\alpha}\widehat{A}_{Ht}^* \tag{28}$$

$$\widehat{Y}_t^* \equiv \frac{Y_t^*}{\bar{Y}^*} = \frac{A_{Ft} + A_{Ft}^*}{\bar{Y}^*} = \bar{\alpha}^*\widehat{A}_{Ft} + (1-\bar{\alpha}^*)\widehat{A}_{Ft}^*, \tag{29}$$

where $\bar{\alpha} \equiv \bar{A}_H^*/\bar{Y}$ and $\bar{\alpha}^* \equiv \bar{A}_F/\bar{Y}^*$ are the shares of, respectively, the Home's and Foreign's production exported in the base year.

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

We 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 $\overline{\alpha}$, 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 domestic prices of local goods (i.e., a PPP adjustment factor less than one), is scaled up. With this adjustment, we compute $\bar{\alpha} = Exp/GDP_{PPP}$, where Exp is the value of base-year exports and GDP_{PPP} is the value of base-year GDP in dollars divided by the PPP adjustment factor. Similarly, $\bar{\alpha}^*$ is computed as the imports of the Home country in USD divdided by the sum of the rest-of-world's PPP-adjusted GDPs.

The constant $\bar{\gamma}$ is the elasticity of absorption with respect to the foreign input, which we measure in the spirit of the Shephard's lemma as the absorption shares of imports 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 take Home's exports divided by the rest-of-world's (aggregate) absorption, all expressed in base year USD. Note that we only need the levels of variables in the base year to compute these constants reflecting base-year shares. All other years are expressed in terms of growth rates relative to the base year.

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 $\widehat{A}_{Ht}, \widehat{A}_{Ft}, \widehat{A}^*_{Ht}$ and $\widehat{A}^*_{Ft}$. We can then compute the inputs for consumption (absorption net of investment and government spending) using the relative growth of consumption to absorption:

$$\widehat{C}_{kt} = \widehat{A}_{kt} \cdot \frac{\widehat{C}_t}{\widehat{A}_t} \quad \text{for } k = H, F, \tag{30}$$

and similarly for $\widehat{C}^*_{kt}$. 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.

Given $\{C_{kt}, C^*_{kt}\}$, $k \in \{H, F\}$, we construct the share of Home and Foreign GDP devoted to consumption. These quantities are held fixed in total but allowed to be re-allocated across countries in our counter-factuals. This assumes that the growth in investment and government expenditure are primitives and are not changed by the planner.

4.3 Identification of Wedges

With the allocations in hand, it is straightforwd to compute the wedges. The static wedge relative to the base year, δ_t, is computed directly from the ratio of $\widehat{C}_{Ht}/\widehat{C}_{Ft}$ to $\widehat{C}^*_{Ht}/\widehat{C}^*_{Ft}$ (see (62) in Appendix D). Given our measures of $\widehat{C}^*$, δ, and the GDP of each region devoted to consumption, we have the constraints on the distorted CPF problem (9), which can be rewritted in growth rates (see problem ($\hat{P}$) in Appendix D). By substituting the observed allocation into the first-order conditions of this problem, we can solve for the multiplier $\tilde{Q}_t$, which is the MRT, and hence for the planner's dynamic wedge relative to the base year λ_t.

5 Empirical Results

In this section we present the results of our measurement exercise. As will become clear, the results differ depending on the frequency of observation. In light of this, we divide each set of results into long differences — that is, the 20 year trend between the start of our sample and the end — and annual growth rates (log differences).

5.1 Real Exchange Rates and Marginal Rates of Transformation

Table 1 contains summary statistics of the model generated series and the empirical moments for consumption and the real exchange rate. The first column is the absolute value of the log difference between the first and last periods, averaged across the n countries in our sample.[24] The second column is the average standard deviation of annual log changes. The first column captures the longer run "trend" component and the second column captures the higher frequency variation.

We begin with a comparison of the observed real exchange rate Q with our planner's MRT $\tilde{Q}$. The real exchange rate is the GDP-weighted average real exchange of each country with respect to all other countries. The MRT is computed by solving a calibrated planning problem for each country, as described in the preceding sections. For both the long differences and the annual changes, the implied MRT is about 2.5 times more stable than the observed real exchange rate.

To get a better sense of the variation across countries contained in our sample, Figure 2 displays a scatter plot of the MRT against the real exchange rate, where each point represents a country in our sample. Panel (a) contains the long differences used in column 1 of Table 1 and panel (b) is the standard deviation of annual log differences. The size of each point reflects their relative PPP-adjusted GDP in the benchmark year. In panel (a), we see that the distribution of

[24] Recall that the last period's value (the benchmark year) is always one, and hence the difference is effectively the magnitude of the log of the first period value.

Table 1: Summary Statistics

	Mean Absolute Value of Long Difference $\frac{1}{n}\sum_i \lvert\log(x_{iT}/x_{i0})\rvert$	Mean Standard Deviation of Annual Change $\frac{1}{n}\sum_i \mathrm{std}(\Delta \log x_{it})$
Relative consumption C/C^*	0.26	0.03
Real exchange rate Q	0.17	0.06
Marginal rate of transformation $\tilde{Q}$	0.07	0.02
Private risk-sharing wedge ψ	0.59	0.10
Planner's dynamic wedge λ	0.48	0.06
Planner's static wedge δ	0.07	0.02

Note: The first column is the absolute value of the log difference between the first and last periods, averaged across the countries in our sample. The second column is the standard deviation of annual log changes, averaged across countries. Note that C/C^* and Q are directly measured in the data, the remaining moments are inferred using the the model. Moments for the wedges are for *gross* wedges, i.e., $1+\psi$, $1+\lambda$ and $1+\delta$.

long differences for the planner's implied MRT is much more concentrated than is the case for the real exchange rate.This is a visualization of the fact reported in Table 1 that the mean absolute change for the real exchange rate is 2.5 times that of the MRT.

In panel (b), we see the different behavior at annual frequencies. As reported in Table 1, the standard deviation of the real exchange rate in annual changes is also roughly 2.5 times that of the MRT. However, Figure 2 indicates that there is significant dispersion in the relative volatilities across countries. Viewed through the model, this suggests that at both 20-year and 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 $\tilde{Q}$ at annual frequencies is on the order of 2.5 percent, which is in line with the volatility of many macroeconomic aggregates, in contrast to a 6 percent volatility of the real exchange rate.

5.2 Risk Sharing Wedges

The lower volatility of $\tilde{Q}$ compared to Q begs the question of whether the MRT is less "disconnected" from relative consumption. Recall that λ is a measure of the wedge between relative consumptions and the MRT (compared to the benchmark year), while ψ is the wedge between relative consumptions and the observed real exchange rate. From Table 1, we see that on average the initial $1+\lambda$ is roughly 48 log points different from the benchmark year (which is normalized to 1), while the gap in $1+\psi$ is approximately 10 log points greater than this. This indicates a modest improvement in the measured risk sharing wedge when the real exchange rate is replaced with the MRT. We shall map this into a welfare measure below.

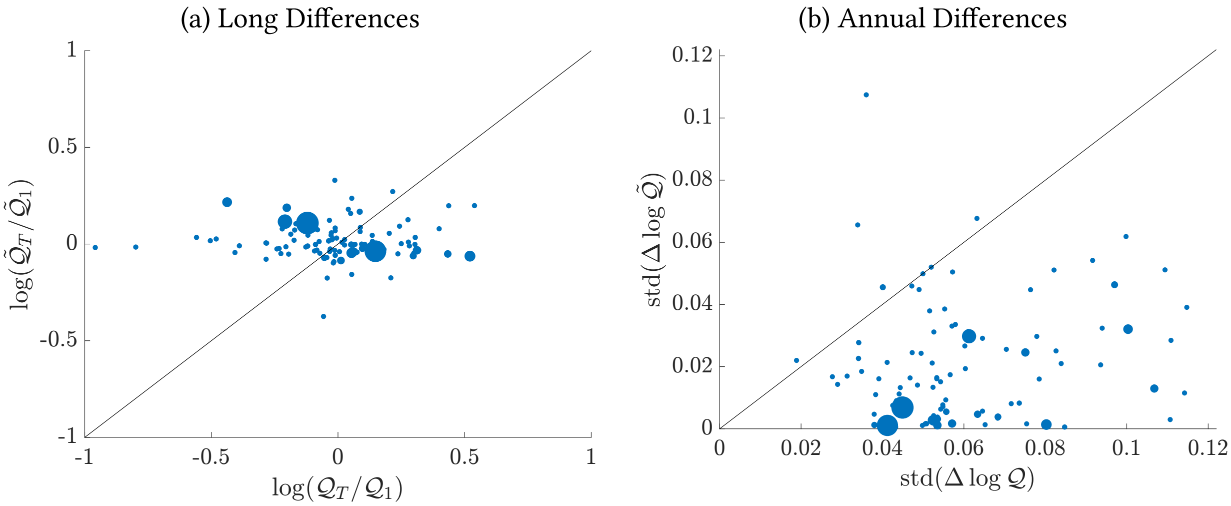

Figure 2: Real Exchange Rates and Marginal Rates of Transformation

Note: Both panels depict a scatter plot in which each circle represents a country with the size corresponding to the PPP-adjusted GDP. Panel (a) plots the log difference between the last and first years of the sample for $\tilde{Q}$ (vertical axis) against the log difference for the real exchange rate (horizontal axis). Panel (b) plots the standard deviation over the sample period for the log annual change in $\tilde{Q}$ against the same moment for Q. The real exchange rate is the GDP-weighted exchange rate across the other countries in the sample.

Note that the magnitudes of the long differences are quite big. The first row of the table indicates that the average difference in the (cumulative) growth in consumption between each country and the rest of the world is roughly 26 log points. Multiplying by $\sigma = 2$ to obtain the marginal rate of substitution gives a difference of 52 log points. The change in the MRT is an order of magnitude smaller (7 log points). That is, the long-run differences in consumption growth are much larger than the change in the relative cost of the two consumption composites. Correspondingly, a planner would like to substantially reallocate consumption across countries at longer horizons. Put another way, we do not see significant risk sharing of the long-run growth differences whether we use MRT or the real exchange rate as our measure of relative costs. This is consistent with the model of Aguiar and Gopinath (2007).

However, the story is markedly different at higher frequencies. In particular, Column 2 of Table 1 indicates that the difference between λ and ψ at annual frequencies is much more pronounced. The average standard deviation of $\Delta \ln(1 + \psi)$ is almost double that of $\Delta \ln(1 + \lambda)$. This is also reflected in the "excess" volatility of the real exchange rate at higher frequencies relative to MRT.

Figure 3 depicts the scatter plots of long-run differences in λ and ψ in panel (a), and the distribution of the standard deviations of annual log changes in panel (b). As with the summary statistics, long changes in λ and ψ line up well, while the high frequency differences in their behavior are substantial.

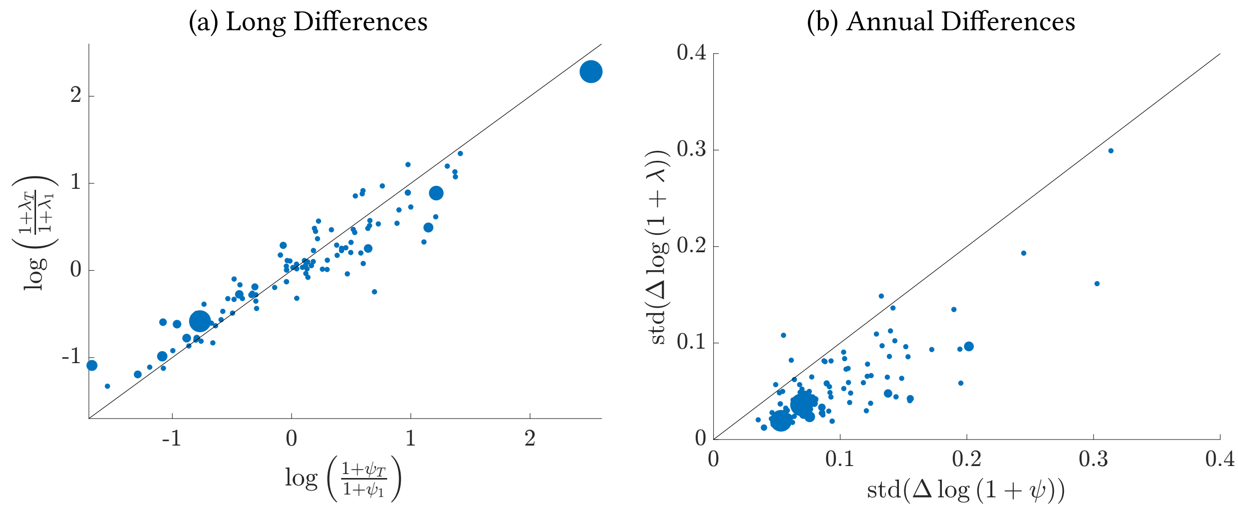

Figure 3: Risk Sharing Wedges

Note: Both panels depict a scatter plot in which each circle represents a country observation, the size of which corresponds to relative PPP-adjusted GDP. Panel (a) plots the log difference between the last and first years of the sample for the planner's dynamic wedge $1 + \lambda$ (vertical axis) against the log difference for the financial wedge $1 + \psi$ (horizontal axis). Panel (b) plots the standard deviation over the sample period for the log annual change in $1 + \lambda$ against the same moment for $1 + \psi$.

The main difference between our approach to risk sharing and the traditional approach is visualized in Figure 4. In this figure, we depict the correlation of the respective notions of the exchange rate changes, $\Delta \log Q$ and $\Delta \log \tilde{Q}$, respectively, with the difference in log consumption growth between Home and ROW. Panel (a) is the histogram of correlations, depicting the respective distributions across countries. We see that the risk sharing condition holds much better with the 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, showing that the correlation with the planner's MRT is almost uniformly higher than the Backus-Smith correlation. Across our sample of countries, the average correlation of relative annual consumption growth with the log change in MRT is 0.56, while the correlation with the observed real exchange rate is -0.14.

5.3 The Static Wedge

Table 1 further indicates that on average the static wedge δ is relatively stable, both at high frequencies and in long differences. Note that this is distinct from the level of the static wedge being small. Differently from the dynamic wedge, for which we can only compute changes, we can compute the level of the static wedge in the benchmark year using our calibrated parameters

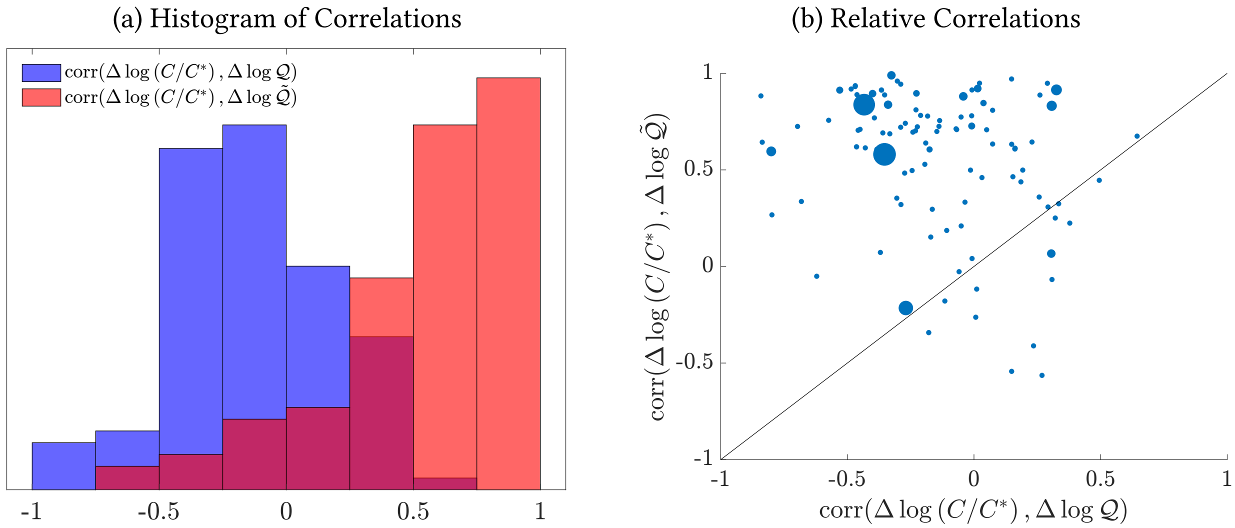

Figure 4: Consumption Differences and the Real Exchange Rate

Note: Panel (a) is a histogram of correlations in the cross section of countries. The blue (left-most) bars represent the distribution of the correlation between annual log change in the real exchange rate and the differential growth in consumption between each country and the rest of the world; the red (right-most) bars plot the same for the log changes in the planner's marginal rate of transformation. Panel (b) is a scatter plot of the same data, with the vertical axis representing the correlation of relative consumption growth with the log change in the planner's MRT and the horizontal axis the correlation between differential consumption growth and the log change in the real exchange rate, with each dot representing a country.

(see Appendix D.4 for details).[25] The wedge is a simple function of the expenditure shares on imports in absorption ($\bar{\gamma}$ and $\bar{\gamma}^*$) and the share of output exported to the other country ($\bar{\alpha}$ and $\bar{\alpha}^*$), all of which are calibrated in the baseline year. Computed this way, the level of the mean baseline static wedge is 4.1 across the sample of countries, with a standard deviation of 3.3, reflecting excess home bias relative to the first-best. Table 1 indicates that while the level of the static wedge may be significant, it is relatively stable over time.[26]

5.4 Time Path of Wedges

Figure 4 speaks to the fact that the efficiency of international risk sharing at annual frequencies becomes more noticeable in the data once we compute the planner's marginal rate of transformation. To show this more directly, Figure 5 depicts the path of the three wedges — λ, ψ, and δ — for a selection of countries. Panel (a) is in levels and panel (b) is in annual log changes.

Three facts about wedges stand out from the figure. First, as noted above, the static wedge δ

[25] We allow for constant trade costs, which represent real resource costs and are not planning problem wedges. Hence the baseline static wedge is independent of the presence of trade costs. See Appendix D.4 for the derivation.

[26] While useful for reference, keep in mind that we do not use the baseline static wedge in any of our results, as we can compute the change in the wedge without knowing its level.

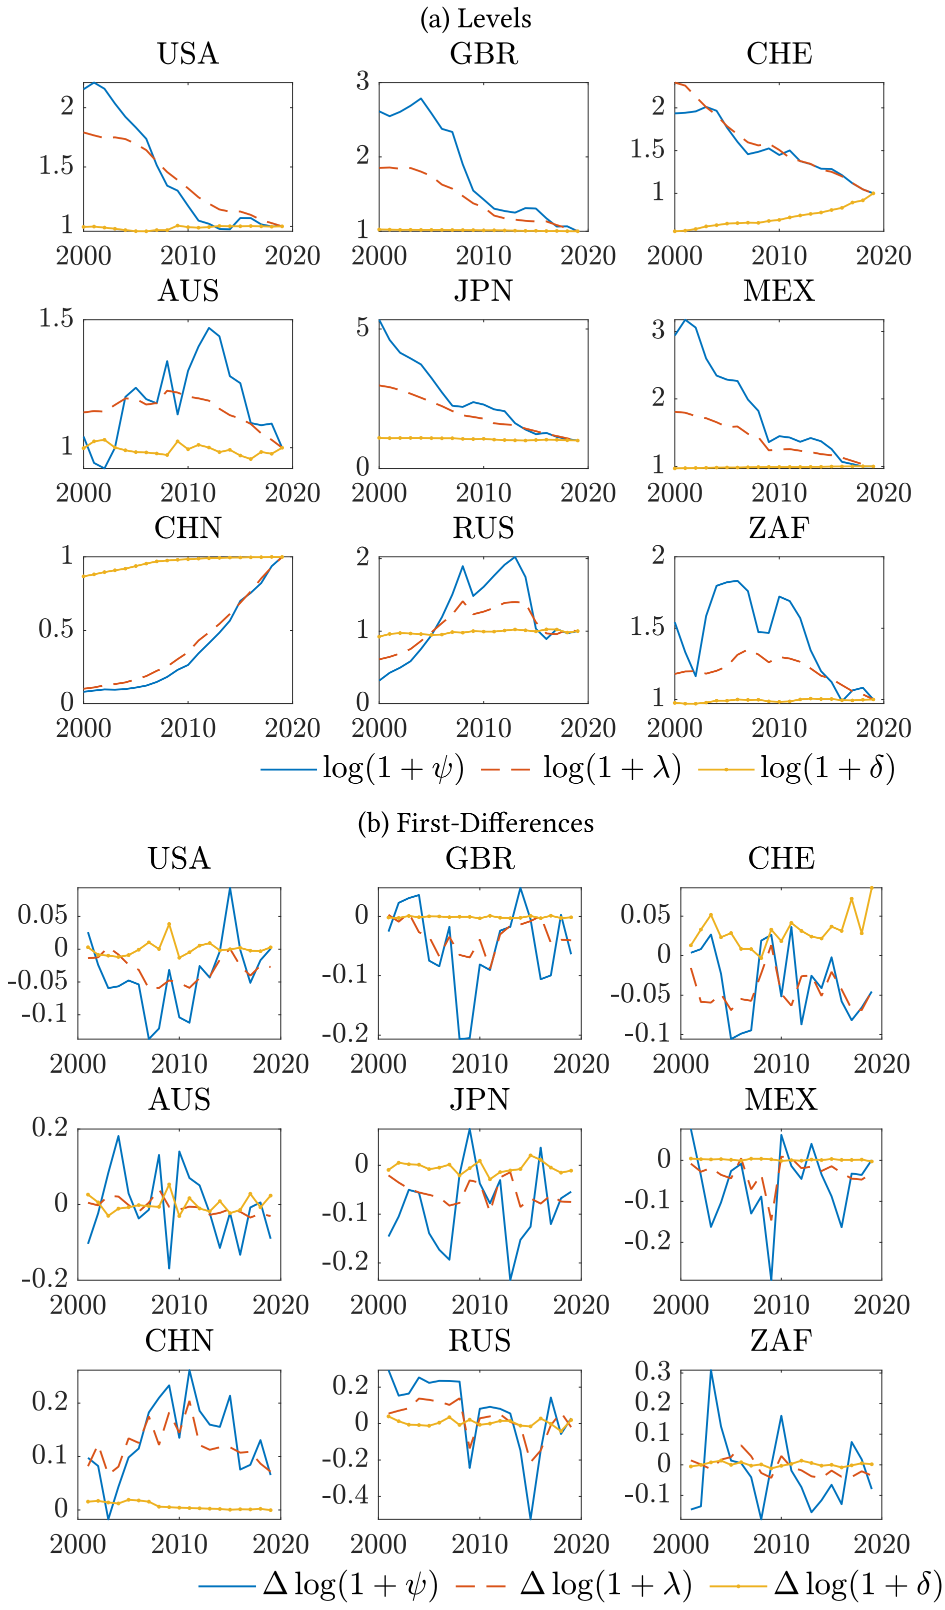

Figure 5: Time Series of Wedges

Note: Each frame represents a time series within a country. The solid blue line corresponds to the private risk-sharing wedge ψ; the dashed red line corresponds to the planner's dynamic wedge λ; and the yellow line with the diamonds corresponds to the static wedge δ. Panel (a) is in log levels and panel (b) is annual log changes.

is surprisingly stable over time for most countries. This is despite large deviations from efficient allocation in the base year (i.e., the benchmark year δ is not equal to zero). Second, the planner's dynamic wedge λ 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 private risk-sharing (Backus-Smith) wedge ψ, and suggests poor risk sharing of permanent productivity shocks. Finally, at higher frequencies, the planner's dynamic wedge is much less volatile than the Backus-Smith wedge. This can be clearly seen in panel (b) of Figure 3 that depicts the volatilities of the log changes in the two wedges and shows their almost uniform ranking. These observations suggests a much more efficient risk sharing at business cycle frequencies.

5.5 Welfare Implications

With these results in hand, we can do the welfare calculations implied by the wedges. For the planner's counter-factual, we solve a Pareto problem that holds Foreign utility constant while maximizing Home's utility. This will involve inter-temporal reallocation of consumption between Home and Foreign such that the resulting dynamic wedge is zero in all periods.

For the financial wedge, we compute the partial equilibrium gains for an individual Home agent that can share risk with an individual Foreign agent. The partial equilibrium nature of the exercise is reflected in the fact that we assume the path of the observed real exchange rate is unchanged by these trades. The gains can be computed using the approach outlined in Section 2.

We can also compute the welfare gain from eliminating the higher frequency volatility in the two wedges. To do this, for each wedge and each country we assume the counter-factual "tax" on the inter-temporal reallocation of consumption decays at a constant geometric rate (i.e., declines log-linearly in time), eliminating the higher frequency fluctuations.[27] While the welfare gain must be positive for eliminating the wedge at all frequencies, it is not necessarily the case for only eliminating the high frequency variation.

Table 2 reports the implied welfare gains for Home across our sample of countries. The first column computes the welfare gains from eliminating the respective wedges completely and the second column assumes the counter-factual wedge moves log-linearly with time. The top panel is the planner's dynamic wedge λ and the bottom panel is the partial equilibrium gain from eliminating the financial friction ψ on private trades.

Throughout the distribution, the implied welfare gains are typically larger when measured with respect to ψ than for λ. 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

[27] Specifically, the counter factual planner's wedge is $\lambda' = (1+\lambda_1)^{1-\frac{t-1}{T-1}} - 1$, given that $\lambda_T = 0$, and similarly for ψ. The counter-factual wedges equal the original wedges at $t = 1$ and $t = T$, but decline log-linearly in-between.

Table 2: Welfare Implications

	Zero Wedge	Linear Trend
Planner's dynamic wedge, λ		
Mean	1.88	0.69
10th percentile	0.15	−0.01
50th percentile	1.13	0.31
90th percentile	4.66	1.98
Private risk-sharing wedge, ψ		
Mean	2.23	1.08
10th percentile	0.28	0.09
50th percentile	1.44	0.73
90th percentile	5.02	2.58

Note: Counterfactuals welfare gains at Home, holding Foreign welfare constant. All numbers represent the welfare-equivalent percentage gain in permanent consumption. The top panel concerns removing the planner's dynamic wedge and the bottom panel concerns the private risk-sharing (Backus-Smith) wedge. The column denoted "Zero Wedge" removes the respective wedge, while the column denoted "Linear Trend" removes fluctuations around a log-linear time trend. Each panel reports the welfare gain at different percentiles of the distribution of countries.

that some of these perceived gains embedded in ψ are offset by the distorted goods market prices reflected in relative markups μ_H and μ_F. Moreover, and as expected from the preceding results, the welfare difference between the two wedges is much more pronounced at higher frequencies. This is consistent with the results presented above that show better risk sharing at annual frequencies when using the MRT instead of the real exchange rate as the relative cost of consumption.

In Figure 6, we plot the welfare gains for a selection of countries. Panel (a) considers the case of shutting down all volatility in λ or ψ. Panel (b) shows welfare gains from shutting down the cyclical component of each wedge, while maintaining a linear trend in the wedge. Consistent with Table 2, for most countries the welfare gain is larger from the private perspective than the planning perspective, particularly for shutting down the higher frequency variation.

The gains from eliminating the planner's risk-sharing wedge λ are modest, but non-trivial. For a median country, eliminating high-frequency fluctuations in λ are equivalent to a 0.3% increase in real consumption; for context, this is almost an order of magnitude larger than the Lucas (1987) gains from eliminating the business cycle. Eliminating the entire wedge λ has a more substantial effect, equivalent to a 1.1% increase in real consumption for a median country (see Table 2). Figure 6 depicts two outlier countries which much larger gains — China and Russia. The reason is that these countries experience unusual growth trajectories in our sample relative to the rest of the world, and with limited international risk sharing their consumption tracks closely their output.

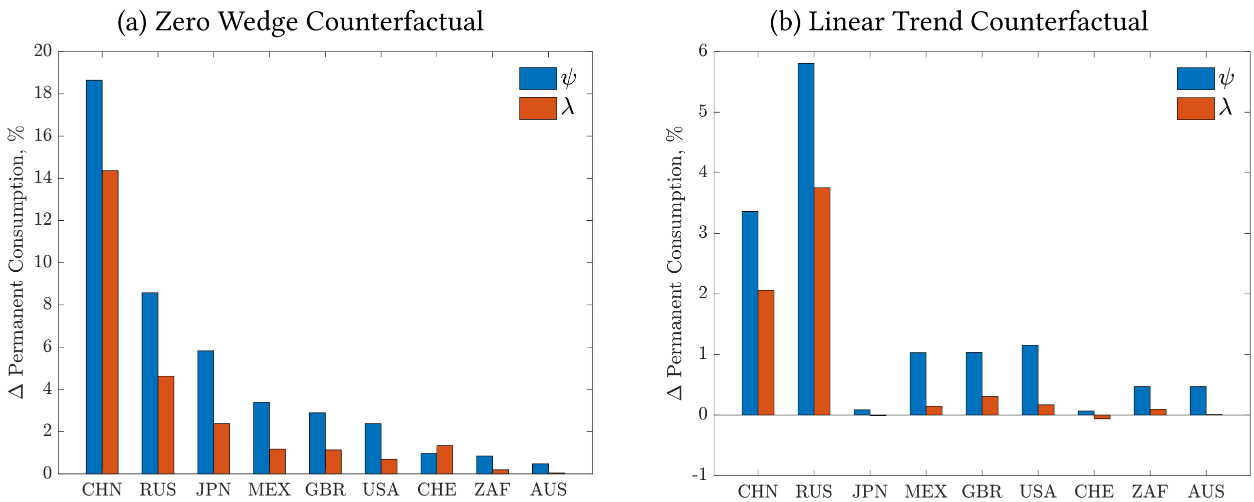

Figure 6: Welfare Gains for a Selection of Countries

Note: Welfare gains from counterfactual exercises for selected countries. Vertical axes refer to the welfare-equivalent percentage gain in permanent consumption. Panel (a) reports gains from completely removing the private risk-sharing (Backus-Smith) wedge ψ (left-most, blue bars) and the planner's dynamic wedge λ (right-most, red bars). Panel (b) repeats the exercise when removing deviations from a log-linear trend in each wedge.

5.6 Decentralization

In this subsection, we explore how the superior risk sharing — evident from the less volatile planner's wedge λ relative to the Backus-Smith wedge ψ — is achieved in a decentralized equilibrium. As described in Section 3, we can use the allocations and the household's first-order conditions to compute implicit prices faced by consumers in Home and Foreign. As also discussed in Section 3, the covariance of the implied cross-border markups (μ_H, μ_F) with the financial wedge ψ determines the variation in the planner's dynamic wedge λ.

In Table 3 we report summary statistics of the implied prices and cross-border markups. The first block of three columns is for log levels relative to the baseline and the second block is for annual first differences. For each block, the first column reports standard deviations computed for each country and then averaged across our sample of countries. The second and third columns are the (average) correlation with the real exchange rate $\mathcal{Q}$ and the private risk-sharing wedge ψ, respectively.

The first two rows are the implied relative prices faced locally by Home and Foreign consumers, respectively. The variation in these relative prices is quite modest, and the prices are largely uncorrelated with the exchange rate and the risk-sharing wedge. The third and fourth columns are the deviations from the Law of One Price for the Home and Foreign goods, respectively. These deviations are quite volatile and almost perfectly correlated with the real exchange

Table 3: Implied Prices

	Levels			Differences		
	StD	Corr Q	Corr ψ	StD	Corr Q	Corr ψ
P_F/P_H	0.02	-0.11	0.10	0.01	0.07	-0.15
P_F^*/P_H^*	0.04	-0.12	0.14	0.03	0.06	-0.14
$1+\mu_H = \mathcal{E}P_H^*/P_H$	0.12	0.90	-0.59	0.07	0.89	-0.66
$1+\mu_F = \mathcal{E}P_F^*/P_F$	0.11	0.99	-0.61	0.06	0.99	-0.73
$1+\delta = (1+\mu_H)/(1+\mu_F)$	0.03	-0.03	0.11	0.02	-0.05	0.13

Note: This table reports moments of implied relative prices. Prices are constructed from the underlying allocations using the procedure outlined in the text. The first three columns report moments using log levels relative to the baseline year and the last three columns using annual log differences. "StD" reports the standard deviation; "Corr Q" reports the correlation with the real exchange rate; and "Corr ψ" reports the correlation with the private risk-sharing ("Backus-Smith") wedge. The first two rows report the ratio of the Foreign good price to Home good price in Home and Foreign (ROW) markets, respectively. The third and fourth rows report the deviation from the law of one price for Home and Foreign goods, respectively. The final row is the ratio of the LOP deviations, which equals the planner's static wedge $1 + \delta$.

rate.[28] They are also negatively correlated with the risk-sharing wedge, although less than perfectly. The final row is the relative LOP deviation, which from Proposition 3 is equivalent to the planner's static wedge. This wedge is not only relatively stable but also uncorrelated with Q or ψ.

In Figure 7, 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, $\mathcal{E}P_H^*/P_H$, and the real exchange rate Q for a selection of countries. 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, LOP deviations track the exchange rate quite closely.

The main takeaway from the table and figures is that the lion's share of variation in the real exchange rate — and its key proximate wedge, the financial shock ψ — is absorbed by LOP deviations, that is, by the cross-border markups. The consumers in each country are largely shielded from these fluctuations and face stable relative prices. Consequently, the decentralization implies that relative markups across borders are moving to offset volatility in the exchange rate. To the extent that fluctuations in Q are driven by something other than the true relative costs, it is consistent with allocational efficiency that these shocks are not reflected in domestic relative prices, but rather absorbed in LOP deviations.[29]

[28] Note that these results are based entirely on the real quantity data and the inversion of the calibrated model. Nonetheless, they are consistent with the empirical evidence measuring LOP deviations directly from the price data (see, e.g., Gopinath et al., 2011; Fitzgerald and Haller, 2013).

[29] Keep in mind that the LOP deviations imply changes in profits for exporters and importers, and our representative-agent model implicitly assumes that these profits or losses are rebated lump-sum to consumers. Similarly, the financial wedge ψ potentially implies profits from intermediation of cross-border asset trades. See Section 3 for details of the decentralization. Fixed costs of entry into exporting and moderate Sharpe ratios of carry trades can rationalize the existence of such profits in equilibrium.

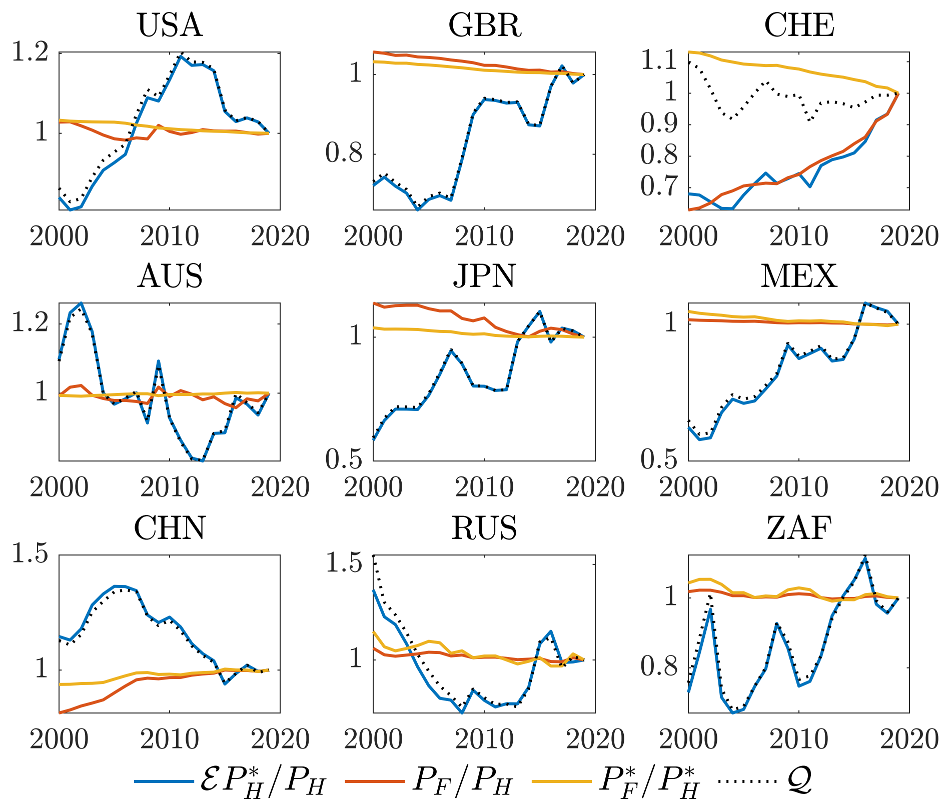

Figure 7: Implied Prices

Note: Each frame displays model-implied time series for relative prices for a given country and its respective ROW. The solid blue line corresponds to the deviation from the law of one price for the Home good; the solid red and yellow lines depict the ratio of the Foreign to Home good prices at Home and in Foreign (ROW). The dashed black line is the observed real exchange rate. All series are normalized to one in the baseline year (2019).

This is broadly consistent with the LCP pricing protocol. The main distinction is that LCP generates $\delta = 0$ in all periods, while our measurement indicates that δ is relatively stable over time at a non-zero level. From Proposition 5(b) and Corollary 2, we have that LCP implies that λ is relatively stable if the LOP deviations μ_H and μ_F are negatively correlated with ψ. Table 3 suggests this is substantially true in the data, and provides an intuition based on decentralized pricing behavior about why the planner's dynamic wedge λ is less volatile than the private risk-sharing wedge ψ.

6 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 marginal 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 at annual frequencies when the planner's relative costs are used to price relative consumption. However, at longer horizons, there are still large unexploited gains from inter-temporal trade.

We also show how the observed allocation is decentralized. The primary distortion in the decentralized equilibrium is a financial 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 pricing to market and local-currency price setting. The net result of the empirical exercise is that observed consumption allocations are much closer to efficient at annual frequencies 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.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings, "International Shocks, Variable Markups and Domestic Prices," *Review of Economic Studies*, November 2019, *6* (86), 2356–2402.

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.

Backus, David K., Silverio Foresi, and Chris I. Telmer, "Affine Term Structure Models and the Forward Premium Anomaly," *The Journal of Finance*, February 2001, *56* (1), 279–304.

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.

— **and Urban J Jermann**, "The International Diversification Puzzle Is Worse Than You Think," *The American Economic Review*, 1997, *87* (1), 170–180.

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

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.

Capelle, Damien and Bruno Pellegrino, "Unbalanced Financial Globalization," Technical Report 10642, CESifo – Center for Economic Studies and ifo Institute 2023.

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.

— **,** — **, and** — , "Business Cycle Accounting," *Econometrica*, 05 2007, *75* (3), 781–836.

Chernov, Mikhail, Valentin Haddad, and Oleg Itskhoki, "What do Financial Markets say about the Exchange Rate?," 2024. NBER Working Paper No. 32436.

Coeurdacier, Nicolas and Helene Rey, "Home bias in open economy financial macroeconomics," *Journal of Economic Literature*, 2013, *51* (1), 63–115.

Cole, Harold L and Maurice Obstfeld, "Commodity trade and international risk sharing: How much do financial markets matter?," *Journal of monetary economics*, 1991, *28* (1), 3–24.

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.

— **,** — **, and Sylvain Leduc**, "International Risk Sharing and the Transmission of Productivity Shocks," *The Review of Economic Studies*, 2008, *75* (2), 443–473.

— **, Lucio D'Aguanno, Aydan Dogan, Simon Lloyd, and Rana Sajedi**, "Global Value Chains and International Risk Sharing," *Robert Schuman Centre for Advanced Studies Research Paper*, 2023, (2023_61).

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.

— **and** — , "Monetary Policy in the Open Economy Revisited: Price Setting and Exchange Rate Flexibility," *Review of Economic Studies*, 2003, *70*, 765–84.

Engel, Charles, "Accounting for U.S. Real Exchange Rate Changes," *Journal of Political Economy*, June 1999, *107* (3), 507–538.

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.

__ and Stefanie Haller, "Pricing-to-Market: Evidence From Plant-Level Prices," *Review of Economic Studies*, February 2013, *81* (2), 761–786.

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

Galí, Jordi and Tommaso Monacelli, "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, 2005, *72* (3), 707–734.

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.

__ , Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li, "International Prices, Costs, and Markup Differences," *American Economic Review*, October 2011, *101* (6), 2450–86.

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.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, November 2009, *124* (4), 1403–1448.

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 __ , "Optimal Exchange Rate Policy," 2023. NBER Working Paper No. 31933.

__ and __ , "Mussa Puzzle Redux," *Econometrica*, January 2025, *93* (1), 1–39.

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

__ and __ , "What Drives the Exchange Rate?," *IMF Economic Review*, March 2025, *73*, 86–117.

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

Jr., Robert E. Lucas, "Interest rates and currency prices in a two-country world," *Journal of Monetary Economics*, 1982, *10* (3), 335–359.

—, *Models of business cycles*, Oxford: Basil Blackwell, 1987.

Kehoe, Patrick J. and Virgiliu Midrigan, "Sticky Prices and Sectoral Real Exchange Rates," 2008. Working Paper, NYU.

Kleinman, Benny, Ernest Liu, Stephen J. Redding, and Motohiro Yogo, "Neoclassical Growth in an Interdependent World," Technical Report 31951, National Bureau of Economic Research 2023.

Kollmann, Robert, "Consumption, real exchange rates and the structure of international asset markets," *Journal of International Money and Finance*, April 1995, *14* (2), 191–211.

—, "Macroeconomic effects of nominal exchange rate regimes: new insights into the role of price dynamics," *Journal of International Money and Finance*, March 2005, *24* (2), 275–292.

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

— **and Edith X. Liu**, "How Can Asset Prices Value Exchange Rate Wedges?," Technical Report 30422, National Bureau of Economic Research 2022.

Lustig, Hanno N. and Adrien Verdelhan, "Does Incomplete Spanning in International Financial Markets Help to Explain Exchange Rates?," *American Economic Review*, June 2019, *109* (6), 2208–44.

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.

Rogoff, Kenneth, "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 1996, *34*, 647–668.

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

Appendix

A Proofs

A.1 Proofs for Section 2

We first characterize the solution to the distorted Consumer Possibilty Frontier, Problem (9), which nests the undistorted frontier by setting $\delta = 0$. Letting $\{\nu, \nu^*, \tilde{Q}, \eta\}$ denote the multipliers on the constraints, we write the Lagrangian for the problem as:

$$\begin{aligned} \mathscr{L} = \Big\{ & C(C_H, C_F) + \nu \left(Y - C_H - C_H^*\right) + \nu^* \left(Y^* - C_F - C_F^*\right) \\ & + \tilde{Q}(C^*(C_H^*, C_F^*) - C^*) + \eta \left(\frac{C_H}{C_F} - (1+\delta)^\theta \Gamma^{-1} \frac{C_H^*}{C_F^*} \right) \Big\}, \end{aligned} \tag{31}$$

where

$$\Gamma \equiv \left(\frac{\gamma}{1-\gamma}\right)\left(\frac{\gamma^*}{1-\gamma^*}\right) < 1$$

by Assumption 1. To avoid a singleton constraint set, we can rewrite equation (8) as an inequality such that the planner is constrained from lowering C_H/C_F if it is distorted above its efficient level ($\delta \geq 0$). In this case, $\eta \geq 0$ in the above Lagrangian. The reverse case of $\delta < 0$ is handled in a similar way, with the multiplier $\eta < 0$. The first-order conditions are

$$\partial C / \partial C_H = \nu - \eta / C_F \tag{32a}$$

$$\partial C / \partial C_F = \nu^* + \eta C_H / C_F^2 \tag{32b}$$

$$\tilde{Q} \partial C^* / \partial C_H^* = \nu + \eta (1+\delta)^\theta \Gamma^{-1} / C_F^* \tag{32c}$$

$$\tilde{Q} \partial C^* / \partial C_F^* = \nu^* - \eta (1+\delta)^\theta \Gamma^{-1} C_H^* / C_F^{*2}. \tag{32d}$$

If $\delta = 0$, then there is no static wedge and we can drop the final constraint in (31) by setting $\eta = 0$. In the $\delta = 0$ case, the ratio of (32a) to (32b) and similarly with (32c) and (32d) implies:

$$\frac{\gamma}{1-\gamma} \frac{C_H}{C_F} = \left(\frac{\nu}{\nu^*}\right)^{-\theta} = \frac{1-\gamma^*}{\gamma^*} \frac{C_H^*}{C_F^*} \quad \Rightarrow \quad \frac{C_H}{C_F} = \Gamma^{-1} \frac{C_H^*}{C_F^*}, \tag{33}$$

which yields equation (7) in the text.

A.1.1 Proof of Proposition 1

Proof. Part (i): The premise is $\{C_H, C_F, C_H^*, C_F^*\}$, which solves (32a)-(32d) with $\eta = 0$. The ratio of (32b) to (32d) yields with some rearranging:

$$\begin{aligned}\tilde{Q}^\theta &= \left(\frac{\gamma}{1-\gamma^*}\right)\left(\frac{C_F^*}{C_F}\right)\frac{C(C_H, C_F)}{C^*(C_H^*, C_F^*)} \\ &= \left(\frac{\gamma}{1-\gamma^*}\right)\frac{C(C_H/C_F, 1)}{C^*(C_H^*/C_F^*, 1)} = \left(\frac{\gamma}{1-\gamma^*}\right)\frac{C(C_H/C_F, 1)}{C^*(\Gamma C_H/C_F, 1)},\end{aligned}$$

where the last equality uses the static efficiency condition (7). Note that this relationship is independent of Y and Y^*. Taking the derivative of the log of the right-hand side with respect to $x \equiv C_H/C_F$ yields a derivative that is strictly positive if $\Gamma < 1$. As $C_H/C_F = \Gamma C_H^*/C_F^*$ at an efficient allocation, $\tilde{Q}$ is strictly increasing in that ratio, as well.

Part (ii): Let C_0/C_0^* be the reference ratio of C/C^*. Where this ray intersects the original CPF, denote the slope of the CPF as $-\tilde{Q}_0$. Suppose Y/Y^* increases, generating a new CPF. We first characterize the ratio C_1/C_1^* at which the new CPF has slope $-\tilde{Q}_0$. Take the ratio of the resource conditions:

$$\frac{Y}{Y^*} = \frac{C_H + C_H^*}{C_F + C_F^*} = \frac{C_H}{C_F}\left(\frac{1 + C_H^*/C_H}{1 + C_F^*/C_F}\right) = \frac{C_H}{C_F}\left(\frac{1 + C_H^*/C_H}{1 + \Gamma^{-1}C_H^*/C_H}\right),$$

where the last equality uses (33). As $\Gamma < 1$, the last term in parentheses is strictly increasing in C_H/C_H^*. From Part (i), if $\tilde{Q} = \tilde{Q}_0$ is constant, then so is C_H/C_F (for any Y/Y^*). Hence, an increase in Y/Y^* for a given $\tilde{Q}_0$ is associated with an increase in C_H/C_H^*, which from static efficiency also implies an increase in C_F/C_F^*, and hence implies an increase in C/C^*. Thus, an increase in Y/Y^* implies there is a ray $C_1/C_1^* > C_0/C_0^*$ at which the slope of the new CPF is $-\tilde{Q}_0$. Given concavity of the CPF (i.e., the magnitude of the slope, $\tilde{Q}$, is decreasing in C/C^*), that implies that at the original C_0/C_0^*, the slope of the new CPF is larger in absolute value. □

A.1.2 Proof of Lemma 1

Proof. In the Cole-Obstfeld case, we have

$$\begin{aligned}C(C_H, C_F) &= C_H^{1-\gamma}C_F^{\gamma} \\ C^*(C_H^*, C_F^*) &= C_H^{*\gamma^*}C_F^{*1-\gamma^*} \\ U(\cdot) &= \log(\cdot)\end{aligned}$$

Let $\{C_H, C_F, C_H^*, C_F^*\}$ solve Problem (9) with associated $\tilde{Q}$, C^*, and $C = \boldsymbol{C}(C^*; Y, Y^*, \delta)$. By definition (12), we have

$$1 + \lambda \equiv \frac{C}{\omega C^*}\frac{1}{\tilde{Q}}.$$

Define

$$\tilde{\eta} \equiv \frac{\eta C_H}{C_F C} = (1+\delta)\Gamma^{-1}\frac{C_H^*}{C_F^*}\frac{\eta}{C},$$

where the second equality is the static wedge constraint. Using the fact that $\tilde{Q}C^* = C/(\omega(1+\lambda))$, we can rewrite equations (32a)-(32b) more compactly as

$$1-\gamma = \nu C_H/C - \tilde{\eta} \tag{34a}$$

$$\gamma = \nu^* C_F/C + \tilde{\eta} \tag{34b}$$

$$\frac{\gamma^*}{\omega(1+\lambda)} = \nu C_H^*/C + \tilde{\eta} \tag{34c}$$

$$\frac{1-\gamma^*}{\omega(1+\lambda)} = \nu^* C_F^*/C - \tilde{\eta}. \tag{34d}$$

We first establish that $\tilde{\eta}$ is a function only of δ and λ. To do so, move $\tilde{\eta}$ to the left-hand side of the above and take the ratios of the first two and second two conditions, respectively:

$$\frac{1-\gamma+\tilde{\eta}}{\gamma-\tilde{\eta}} = \frac{\nu}{\nu^*}\frac{C_H}{C_F}$$

$$\frac{\gamma^*-\tilde{\eta}\omega(1+\lambda)}{1-\gamma^*+\tilde{\eta}\omega(1+\lambda)} = \frac{\nu}{\nu^*}\frac{C_H^*}{C_F^*} = \frac{\nu}{\nu^*}\frac{\Gamma C_H}{C_F(1+\delta)^\theta}.$$

The terms on the right are well defined and strictly positive. This restricts $\tilde{\eta} \in (\underline{\eta}, \overline{\eta})$, where

$$\underline{\eta} \equiv \max\left\{\gamma-1, \frac{\gamma^*-1}{\omega(1+\lambda)}\right\} \text{ and } \overline{\eta} \equiv \min\left\{\gamma, \frac{\gamma^*}{\omega(1+\lambda)}\right\}.$$

Rearranging and equating the above first-order conditions, we have:

$$(1+\delta)\left(\frac{\gamma^*-\tilde{\eta}\omega(1+\lambda)}{1-\gamma^*+\tilde{\eta}\omega(1+\lambda)}\right) = \Gamma\left(\frac{1-\gamma+\tilde{\eta}}{\gamma-\tilde{\eta}}\right).$$

The left-hand side is strictly decreasing in $\tilde{\eta}$ and the right-hand side is strictly increasing. For $\delta \geq 0$, at $\tilde{\eta} = 0$, the left-hand side is weakly greater than the right, and equal if $\delta = 0$. There is a unique solution for $\tilde{\eta} \in [0, \overline{\eta})$. For $\delta < 0$, there is a unique solution for $\tilde{\eta} \in (\underline{\eta}, 0)$. Let us define $H(\lambda, \delta)$ as the mapping between the wedges and $\tilde{\eta}$. Differentiating and imposing the relevant bounds on parameters yields:

$$\frac{\partial H}{\partial \delta} > 0 \qquad \text{and} \qquad \text{sign}\left\{\frac{\partial H}{\partial \lambda}\right\} = -\text{sign}\{\delta\}.$$

We now turn to the main proposition. Summing (34a) and (34c) and imposing the resource constraint yields:

$$1-\gamma+\frac{\gamma^*}{\omega(1+\lambda)} = \nu Y/C.$$

Similarly, adding (34b) and (34d) yields

$$\gamma+\frac{1-\gamma^*}{\omega(1+\lambda)} = \nu^* Y^*/C.$$

We can use these to eliminate ν/C and ν^*/C from (34a) and (34b), respectively, to obtain:

$$C_H = (1-\gamma+\tilde{\eta})C/\nu = \frac{(1-\gamma+\tilde{\eta})Y}{1-\gamma+\frac{\gamma^*}{\omega(1+\lambda)}}$$

$$C_F = (\gamma-\tilde{\eta})C/\nu^* = \frac{(\gamma-\tilde{\eta})Y^*}{\gamma+\frac{1-\gamma^*}{\omega(1+\lambda)}}.$$

Using the Cobb-Douglas aggregator, we obtain

$$C = C_H^{1-\gamma} C_F^{\gamma} = g(\lambda, \delta) Y^{1-\gamma} Y^{*\gamma},$$

where

$$g(\lambda, \delta) \equiv \left(\frac{1 - \gamma + H(\lambda, \delta)}{1 - \gamma + \frac{\gamma^*}{\omega(1+\lambda)}} \right)^{1-\gamma} \left(\frac{\gamma - H(\lambda, \delta)}{\gamma + \frac{1-\gamma^*}{\omega(1+\lambda)}} \right)^{\gamma} > 0.$$

Differentiating and imposing the bounds on parameters, we have

$$\frac{\partial g}{\partial \lambda} > 0 \qquad \text{and} \qquad \text{sign}\left\{ \frac{\partial g}{\partial \delta} \right\} = -\text{sign}\{\delta\}.$$

We follow the same steps with conditions (34c) and (32d) to obtain

$$C_H^* = (\gamma^*/(\omega(1+\lambda)) - \tilde{\eta})\, C/\nu = \frac{(\gamma^* - \tilde{\eta}(\omega(1+\lambda)))\, Y}{(1-\gamma)(\omega(1+\lambda)) + \gamma^*}$$
$$C_F^* = \left(\frac{1-\gamma^*}{\omega(1+\lambda)} + \tilde{\eta} \right) C/\nu^* = \frac{(1 - \gamma^* + \tilde{\eta}\omega(1+\lambda)) Y^*}{\gamma\omega(1+\lambda) + 1 - \gamma^*}.$$

This yields

$$C^* = C_H^{*\gamma^*} C_F^{*1-\gamma^*} = g^*(\lambda, \delta) Y^{\gamma^*} Y^{*1-\gamma^*},$$

where

$$g^*(\lambda, \delta) \equiv \left(\frac{(\gamma^* - H(\lambda, \delta)(\omega(1+\lambda)))}{(1-\gamma)(\omega(1+\lambda)) + \gamma^*} \right)^{\gamma^*} \left(\frac{(1 - \gamma^* + H(\lambda, \delta)\omega(1+\lambda))}{\gamma\omega(1+\lambda) + 1 - \gamma^*} \right)^{1-\gamma^*} > 0.$$

Differentiating and imposing the bounds on parameters, we have :

$$\frac{\partial g^*}{\partial \lambda} < 0 \qquad \text{and} \qquad \text{sign}\left\{ \frac{\partial g^*}{\partial \delta} \right\} = -\text{sign}\{\delta\}.$$

Using these expressions and the definition of λ, we have

$$\tilde{Q} = \frac{C}{C^*\omega(1+\lambda)} = h(\lambda, \delta) \left(\frac{Y}{Y^*} \right)^{1-\gamma-\gamma^*}, \qquad \text{where} \qquad h(\lambda, \delta) \equiv \frac{1}{\omega(1+\lambda)} \frac{g(\lambda, \delta)}{g^*(\lambda, \delta)}.$$

□

A.2 Proofs for Section 3.2

A.2.1 Proof of Proposition 3

Proof. By the definitions in equation (21),

$$\frac{P_H}{P_F} = \frac{P_H^*}{P_F^*} \frac{1+\mu_F}{1+\mu_H} \qquad \Rightarrow \qquad \frac{C_H}{C_F} = \left[\frac{1+\mu_F}{1+\mu_H} \right]^{\theta} \left(\frac{1-\gamma}{\gamma} \frac{1-\gamma^*}{\gamma^*} \right) \frac{C_H^*}{C_F^*}, \tag{35}$$

where the right-hand side expression is obtained using (16) and (19). Comparing this expression with (8), we have the first result. For the second result, divide the definitions in (12) and (22) to obtain

$$\frac{1+\lambda}{1+\psi} = \frac{\frac{U'(C^*)}{\omega U'(C)\tilde{Q}}}{\frac{U'(C^*)}{\bar{\omega} U'(C) Q}} = \frac{\bar{\omega}}{\omega}\frac{Q}{\tilde{Q}}.$$

The result follows from our assumed normalization $\omega = \bar{\omega}$. □

A.2.2 Proof of Proposition 4 and Corollary 1

Proof. From Proposition 3, when $\mu_H = \mu_F = \mu$, we have $\delta = 0$. From the proof of Proposition 1, we have

$$\tilde{Q} = \left[\left(\frac{\gamma}{1-\gamma^*}\right)\frac{C(C_H/C_F, 1)}{C^*(C_H^*/C_F^*, 1)}\right]^{\frac{1}{\theta}} = \left(\frac{(1-\gamma)\,(P_F/P_H)^{\theta-1} + \gamma}{\gamma^*\left(P_F^*/P_H^*\right)^{\theta-1} + 1 - \gamma^*}\right)^{\frac{1}{\theta-1}}.$$

where the second line uses the household's first-order conditions (16) and (19). When $\mu_F = \mu_H$, we have $P_H/P_F = P_H^*/P_F^*$. This implies that the above is equivalent to

$$\tilde{Q} = \left(\frac{(1-\gamma) + \gamma\,(P_F/P_H)^{1-\theta}}{\gamma^* + (1-\gamma^*)\left(P_F^*/P_H^*\right)^{1-\theta}}\right)^{\frac{1}{\theta-1}} = Q\frac{P_H}{\mathcal{E}P_H^*} = Q/(1+\mu),$$

where the middle equality follows from definition of Q in (18) and the last equality is by the definition of $1+\mu = 1+\mu_H$. This proves the second equality in equation (23). The first equality is Proposition 3 Part (ii).
Part (a) of Corollary 1 follows by substitution $\mu = 0$ in the above to obtain $Q = \tilde{Q}$, and then using this fact in Part (ii) of Proposition 3. Part (b) follows by setting $\psi = 0$ in equation (23). Part (c) is a direct implication of equation (23). □

A.3 Proofs for Section 3.3

The proofs in this subsection use the equilibrium conditions discussed in Appendix B. For convenience, the key equilibrium conditions are:

1. Nominal expenditure M and M^* is controlled by monetary authority at Home and in Foreign, respectively, based on (43) and its counterpart in Foreign, which we assume holds with equality:

$$PC = M \qquad \text{and} \qquad P^*C^* = M^*.$$

2. Market clearing and household optimality imply (equation (46)):

$$Y = (1-\gamma)\left(\frac{P_H}{P}\right)^{-\theta} C + \gamma^*\left(\frac{P_H^*}{P^*}\right)^{-\theta} C^*,$$
$$Y^* = \gamma\left(\frac{P_F}{P}\right)^{-\theta} C + (1-\gamma^*)\left(\frac{P_F^*}{P^*}\right)^{-\theta} C^*,$$

with price indexes $P = \mathcal{P}(P_H, P_F)$ and $P^* = \mathcal{P}^*(P_H^*, P_F^*)$ defined in Section 3.

3. The private risk-sharing (Backus-Smith) condition (22) combined with (43) and the definition of the real exchange rate in (18) is (equation (47)):

$$\mathcal{E} = \frac{1}{\bar{\omega}(1+\psi)} \frac{M}{M^*} \left(\frac{C}{C^*}\right)^{\sigma-1} = \frac{1}{\bar{\omega}(1+\psi)} \left(\frac{M}{M^*}\right)^{\sigma} \left(\frac{P^*}{P}\right)^{\sigma-1}.$$

4. Price setting:

   (a) Under PCP: P_H and P_F^* are set prior to the period and $P_H^* = P_H/\mathcal{E}$ and $P_F = P_F^*\mathcal{E}$;

   (b) Under LCP: All local-currency prices (P_H, P_H^*, P_F, P_F^*) are pre-set; and

   (c) Under PTM: Prices are flexible and set according to (48).

A.3.1 Proof of Proposition 5

Proof. **Part (a):** Under PCP, by definition $P_H = \mathcal{E}P_H^*$ and $P_F = \mathcal{E}P_F^*$ and P_H and P_F^* are set ex ante. From (21), we immediately have $\mu_H = \mu_F = 0$. From Proposition (4), this implies $\delta = 0$, $\tilde{Q} = Q$ and $\lambda = \psi$. From the definition of Q in equation (18), and using the fact that PCP implies $\mathcal{E}P_H^* = P_H$,we have

$$Q = \left(\frac{\gamma^* + (1-\gamma^*)\left(P_F^*/P_H^*\right)^{1-\theta}}{(1-\gamma) + \gamma\left(P_F/P_H\right)^{1-\theta}}\right)^{1/(1-\theta)} = \left(\frac{\gamma^* + (1-\gamma^*)\left(\mathcal{E}P_F^*/P_H\right)^{1-\theta}}{(1-\gamma) + \gamma\left(\mathcal{E}P_F^*/P_H\right)^{1-\theta}}\right)^{1/(1-\theta)}, \tag{36}$$

where the second equality imposes $P_H = \mathcal{E}P_H^*$ and $P_F = \mathcal{E}P_F^*$. Under PCP, P_H and P_F^* are fixed ex ante, and hence this relates the realized Q to the realized $\mathcal{E}$. A useful property of this mapping is that $Q > 0$ at $\mathcal{E} = 0$ and

$$\frac{\mathcal{E}}{Q}\frac{dQ}{d\mathcal{E}} \in (0,1).$$

From the definition of ψ in (22), we have

$$Q = \frac{1}{\bar{\omega}(1+\psi)}\left(\frac{C}{C}\right)^{\sigma} = \frac{1}{\bar{\omega}(1+\psi)}\left(\frac{M}{M^*}\frac{P^*}{P}\right)^{\sigma} = \frac{1}{\bar{\omega}(1+\psi)}\left(\frac{M}{M^*}\frac{Q}{\mathcal{E}}\right)^{\sigma},$$

where the second equality uses $M = PC$ and $M^* = P^*C^*$ and the final equality uses $Q = \mathcal{E}P^*/P$. Rearranging, we have

$$Q = (\bar{\omega}(1+\psi))^{\frac{1}{\sigma-1}}\left(\frac{M}{M^*}\right)^{\frac{\sigma}{1-\sigma}}\mathcal{E}^{\frac{\sigma}{\sigma-1}}. \tag{37}$$

Equations (36) and (37) are two equations to determine $\{Q, \mathcal{E}\}$ given M/M^* and ψ. Recall that (36) is positively sloped with an elasticity between zero and one and has a positive intercept at $\mathcal{E} = 0$. We now consider three cases for the size of $\sigma > 0$ to establish existence, uniqueness, and comparative statics:

(i) Suppose $\sigma \in (0,1)$: Then (37) is negatively sloped with $Q \to \infty$ as $\mathcal{E} \to 0$ and $Q \to 0$ as $\mathcal{E} \to \infty$. Hence, there is a unique solution. Moreover, an increase in M/M^* or a decrease in ψ increases Q in (37) for each $\mathcal{E}$. This implies that both Q and $\mathcal{E}$ are increasing in M/M^* and decreasing in ψ;

(ii) Suppose $\sigma = 1$: Then $\mathcal{E}$ is pinned down by (25) and (36) determines Q. From (25), $\mathcal{E}$ is increasing in M/M^* and decreasing in ψ. From (36), Q inherits these comparative statics.

(iii) Suppose $\sigma > 1$: Then (37) lies below (36) at $\mathcal{E} = 0$, but has an elasticity strictly greater than one, which is greater than that of (37). This also implies a unique solution to the two equations. Moreover, an increase in M/M^* or a decrease in ψ implies an increase in $\mathcal{E}$ in (37) for a given $\mathcal{Q}$. Together with (36), this implies that $\mathcal{Q}$ is increasing in M/M^* and decreasing in ψ.

This completes the proof of Part (a).

Part (b): Under LCP, all prices are set before the period such that $P_H = \bar{\mathcal{E}}P_H^*$ and $P_F = \bar{\mathcal{E}}P_F^*$ for some constant $\bar{\mathcal{E}}$. This implies that P and P^* are also pre-set and $\mathcal{Q} \propto \mathcal{E}$ within a period. From (21), the price-setting protocol also implies that $\mu_H = \mu_F = \mu = \mathcal{E}/\bar{\mathcal{E}} - 1$ and from Proposition 4 we have $\delta = 0$ and $\mathcal{Q} = (1+\mu)\tilde{\mathcal{Q}}$. Using the definition (18), the fact that $\mathcal{Q} = (1+\mu)\tilde{\mathcal{Q}}$ also implies

$$\tilde{\mathcal{Q}} = \frac{\mathcal{Q}}{1+\mu} = \left(\frac{\gamma^* + (1-\gamma^*)\left(P_F^*/P_H^*\right)^{1-\theta}}{(1-\gamma) + \gamma\left(P_F/P_H\right)^{1-\theta}} \right)^{1/(1-\theta)}.$$

As $P_F/P_H = P_F^*/P_H^*$ are set ex ante, this implies that $\tilde{\mathcal{Q}}$ does not vary within the period and $\mathcal{Q}$ is proportional to $1+\mu$. Equation (37) holds for any pricing protocol, which implies

$$\mathcal{Q} = (\bar{\omega}(1+\psi))^{-1} \left(\frac{M}{M^*}\right)^{\sigma} \left(\frac{\mathcal{Q}}{\mathcal{E}}\right)^{\sigma}. \tag{38}$$

As $\mathcal{Q} \propto \mathcal{E}$, this implies that $\mathcal{Q} \propto (1+\psi)^{-1}(M/M^*)^{\sigma}$. Finally, $\mathcal{Q} = (1+\mu)\tilde{\mathcal{Q}}$ plus Proposition 4 implies $1+\lambda = (1+\mu)(1+\psi) = \mathcal{Q}(1+\psi)/\tilde{\mathcal{Q}}$, where the latter follows from $\mathcal{Q} = (1+\mu)\tilde{\mathcal{Q}}$. As $\tilde{\mathcal{Q}}$ does not vary within the period, $1+\lambda \propto \mathcal{Q}(1+\psi) \propto (M/M^*)^{\sigma}$, where the final term follows from (38).

Part (c): The statement that $\mu_H = \mu_F = \mu = \mathcal{Q}^{\phi} - 1$ follows from (49). The fact that $\delta = 0$ follows from $\mu_H = \mu_F$ and Proposition 3. Proposition 4 gives $\tilde{\mathcal{Q}} = \mathcal{Q}/(1+\mu) = \mathcal{Q}^{1-\phi}$, where the last equality uses $\mu = \mathcal{Q}^{\phi} - 1$. This plus Proposition 4 gives $1+\lambda = (1+\psi)\mathcal{Q}/\tilde{\mathcal{Q}} = (1+\psi)\mathcal{Q}^{\phi}$. Finally, note that by (18), $1+\mu = \mathcal{E}P_H^*/P_H = \mathcal{Q}^{\phi}$, and letting $\mathcal{S} \equiv P_F/P_H = P_F^*/P_H^*$, we have

$$\mathcal{Q}^{1-\phi} = \left(\frac{\gamma^* + (1-\gamma^*)\mathcal{S}^{1-\theta}}{1-\gamma+\gamma\mathcal{S}^{1-\theta}} \right)^{\frac{1}{1-\theta}} = \frac{\mathcal{P}^*(1,\mathcal{S})}{\mathcal{P}(1,\mathcal{S})}, \tag{39}$$

where the second equality uses the definitions of $\mathcal{P}^*$ and $\mathcal{P}$ as the CES ideal price indices for Home and Foreign, respectively. From (46), we have

$$\frac{Y}{Y^*} = \frac{(1-\gamma)\mathcal{P}(1,\mathcal{S})^{\theta}\frac{C}{C^*} + \gamma^*\mathcal{P}^*(1,\mathcal{S})^{\theta}}{\gamma\mathcal{P}(\mathcal{S}^{-1},1)^{\theta}\frac{C}{C^*} + (1-\gamma^*)\mathcal{P}^*(\mathcal{S}^{-1},1)^{\theta}}, \tag{40}$$

where we have used $P_H/P = 1/\mathcal{P}(1,\mathcal{S})$, $P_F/P = 1/\mathcal{P}(\mathcal{S}^{-1},1)$ and similarly for $P_H^*/P^* = 1/\mathcal{P}^*(1,\mathcal{S})$ and $P_F^*/P^* = \mathcal{P}^*(\mathcal{S}^{-1},1)$. Note that

$$\frac{\mathcal{P}^*(1,\mathcal{S})}{\mathcal{P}(1,\mathcal{S})} = \frac{\mathcal{S}\mathcal{P}^*(\mathcal{S}^{-1},1)}{\mathcal{S}\mathcal{P}(\mathcal{S}^{-1},1)} = \frac{\mathcal{P}^*(\mathcal{S}^{-1},1)}{\mathcal{P}(\mathcal{S}^{-1},1)} = \mathcal{Q}^{1-\phi},$$

where the first two equalities use the constant-returns-to-scale feature of the price aggregators and the last equality follows from (39). Differentiating this expression, we have

$$d\log\mathcal{S} = \frac{1-\phi}{A} d\log\mathcal{Q}, \tag{41}$$

where

$$A \equiv \frac{(1-\gamma-\gamma^*)\mathcal{S}^{1-\theta}}{[1-\gamma+\gamma\mathcal{S}^{1-\theta}][\gamma^*+(1-\gamma^*)\mathcal{S}^{1-\theta}]} \in (0,1).$$

Moreover, from (22), we have

$$\frac{C}{C^*} = (\bar{\omega}(1+\psi))^{\frac{1}{\sigma}} \mathcal{Q}^{\frac{1}{\sigma}}.$$

Substituting in these expressions into (40) and rearranging, we have

$$\frac{Y}{Y^*} = \mathcal{S}^{\theta} \frac{(1-\gamma)\,(\bar{\omega}(1+\psi))^{\frac{1}{\sigma}}\,\mathcal{Q}^{\frac{1}{\sigma}+(\phi-1)\theta}+\gamma^*}{\gamma\,(\bar{\omega}(1+\psi))^{\frac{1}{\sigma}}\,\mathcal{Q}^{\frac{1}{\sigma}+(\phi-1)\theta}+1-\gamma^*}, \tag{42}$$

where we have used $\mathcal{P}(1,\mathcal{S}) = \mathcal{S}\mathcal{P}(\mathcal{S}^{-1},1)$. Differentiating, we obtain

$$\begin{aligned} d\log(Y/Y^*) &= \theta d\log\mathcal{S} - B\left[\left(\theta(1-\phi)-\frac{1}{\sigma}\right)d\log\mathcal{Q} - \frac{1}{\sigma}d\log(1+\psi)\right] \\ &= (\theta(1-\phi)(A^{-1}-B)+B/\sigma)d\log\mathcal{Q} + (B/\sigma)d\log(1+\psi), \end{aligned}$$

where the second equality uses (41) and $B \in (0,1)$ is:

$$B \equiv \frac{(1-\gamma-\gamma^*)\mathcal{Q}^{\theta(1-\phi)-1/\sigma}\,[\bar{\omega}(1+\psi)]^{-1/\sigma}}{[(1-\gamma)+\gamma^*\mathcal{Q}^{\theta(1-\phi)-1/\sigma}\,[\bar{\omega}(1+\psi)]^{-1/\sigma}\,][\gamma+(1-\gamma^*)\mathcal{Q}^{\theta(1-\phi)-1/\sigma}\,[\bar{\omega}(1+\psi)]^{-1/\sigma}\,]}.$$

As $A \in (0,1)$, $B \in (0,1)$ and $\phi \in [0,1)$, the term multiplying $d\log\mathcal{Q}$ is strictly positive. Hence, $\mathcal{Q}$ is strictly increasing in Y/Y^* and decreasing in ψ. □

B Monetary Equilibrium

In this appendix, we introduce a monetary economy with endogenous production. This augments Section 3.3 in the text.

B.1 Households

Flow utility for the Home representative agent becomes $U(C) - v(n)$, where n represents hours worked and v represents (dis)utility over labor. In the monetary model, households trade in the asset market at the start of the period, and then use cash in the latter half for transactions. This follows the timing of Lucas (1982) and Kehoe and Midrigan (2008). Specifically, after trading in the asset market at time t in history s^t, the Home representative agent has the following budget constraint:

$$\begin{aligned}
&M_t(s^t) + \sum_{s_{t+1}} \Lambda_{t+1}(s_{t+1}|s^t) B_{t+1}(s^{t+1}) = \\
&R_{t-1}(s^{t-1}) W_{t-1}(s^{t-1}) n_{t-1}(s^{t-1}) + B_t(s^t) + \left[M_{t-1}(s^{t-1}) - P_{t-1}(s^{t-1}) C_{t-1}(s^{t-1})\right] + T_t(s^t) + \Pi_t(s^t),
\end{aligned}$$

where the left-hand side consists of purchases in the period t asset markets – namely, cash plus next-period's state-contingent assets – and the right-hand side is the cash brought into the period (last-period's wages and unspent cash) plus transfers and dividends. $\Lambda_{t+1}(s_{t+1}|s^t) \equiv \Lambda_{t+1}((s_{t+1}, s^t))/\Lambda_t(s^t)$ are the one-period Arrow prices and $R_t(s^t) = 1/\sum_{s_{t+1}} \Lambda_{t+1}(s_{t+1}|s_t)$ is the period t interest rate. Following Kehoe and Midrigan (2008), we assume that worker's receive the previous period's wages with interest, financed by implicit government transfers. The transfers also include any cash injections or reductions made by the government.

The cash-in-advanced constraint is

$$P_t(s^t) C_t(s^t) \leq M_t(s^t). \tag{43}$$

We assume this binds at every history. We can then write the household's period-0 constraint as:

$$w \geq \sum_t \sum_{s^t} \Lambda_t(s^t) \left(P_t(s^t) C_t(s^t) - W_t(s^t) n_t(s^t)\right),$$

where w contains the period-0 value of transfers, dividends, and any cash held at the start of time. This is equivalent to the constraint in (15) with labor income replacing the endowment. The first-order conditions (16) and (17) continue to characterize the household's optimal choices. The Foreign household's problem is symmetric.

Our methodology can remain agnostic about how labor markets clear. If wages are flexible, households choose $n_t(s^t)$ and $n_t^*(s^t)$ given wages. Due to the fact that wages are paid with interest, this yields an undistorted static labor-consumption tradeoff in each region:

$$U'(C_t(s^t))W_t(s^t)/P_t(s^t) = v'(n_t(s^t)) \text{ and } U'(C_t^*(s^t))W_t^*(s^t)/P_t^*(s^t) = v'(n_t^*(s^t)).$$

Alternatively, we can assume that households are not on their labor supply curves and take labor income as beyond their control. In what follows, we consider the case of flexible wages for concreteness.

B.2 Technology

Production occurs according to

$$Y_t(s^t) = Z_t(s^t)n_t(s^t)$$
$$Y_t^*(s^t) = Z_t^*(s^t)n_t^*(s^t),$$

where Z and Z^* represent labor-productivity shocks. Under flexible prices and wages, cost minimization implies for interior labor demand that

$$P_{H,t}(s^t)Z_t(s^t) = W_t(s^t) \text{ and } P_{F,t}^*(s^t)Z_t^*(s^t) = W_t^*(s^t). \tag{44}$$

If output is demand determined, then $n_t(s^t)$ and $n_t^*(s^*)$ are such that production equals demand.

B.3 Asset Markets

As noted in the text, there are many alternative (and for our purposes equivalent) interpretations of the Backus-Smith wedges, $\{\psi_t(s^t)\}$, which we take as primitives. For concreteness, suppose that a Home resident must pay $(1 + \psi_t(s^t))\mathcal{E}_t(s^t)$ units of Home currency to purchase a unit of Foreign currency, while the Foreign resident receives only $\mathcal{E}_t(s^t)$. The amount $\psi_t(s^t)\mathcal{E}_t(s^t)$ is either profit to financial intermediaries or a tax payment. In either case, it is rebated back to the households and the present value of these transfers and/or dividends are part of the initial wealth positions w and w^*.

Recall that Foreign resident's can purchase a claim to a unit of foreign currency delivered in history s^t at a period-zero price in foreign currency of $\Lambda_t^*(s^t)$, which is equivalent to a period-zero price in Home currency of $\mathcal{E}_0(s_0)\Lambda_t^*(s^t)$. A Home resident is willing to pay $(1+\psi_t(s^t))\mathcal{E}_t(s^t)\Lambda_t(s^t)$ in period-0 units of Home currency for a claim on a unit of Foreign currency in history s^t. Allowing the residents' to trade in period zero such claims on currency in t, s^t, arbitrage implies

$$\mathcal{E}_0(s_0)\Lambda_t^*(s^t) = (1 + \psi_t(s^t))\mathcal{E}_t(s^t)\Lambda_t(s^t), \tag{45}$$

which is the definition in equation (22).

B.4 Pricing

To close the model, we need to state how prices are set. As noted in the text, for our purposes we do not model how prices are set beyond the constraints of the respective pricing protocol. That is, we remain agnostic about the market structure of the goods market and the desired level of markups. We proceed by discussing our three pricing protocols.

Producer Currency Pricing (PCP) Under PCP, we take the sequence of domestic prices as primitives $\{P_{H,t}(s^t), P^*_{F,t}(s^t)\}$. Given these prices and the sequence of exchange rates, export prices are $P^*_{H,t}(s^t) = P_{H,t}(s^t)/\mathcal{E}(s^t)$ and $P_{F,t} = \mathcal{E}(s^t)P^*_{Ft}(s^t)$. We define a PCP equilibrium as follows:

Definition 1. *Given technology $\{Z_t(s^t), Z^*_t(s^t)\}$, monetary policies $\{M_t(s^t), M^*_t(s^t)\}$, asset market taxes/wedges $\{\psi_t(s^t)\}$, domestic price sequences $\{P_{H,t}(s^t), P^*_{F,t}(s^t)\}$, and initial assets w and w^*, a PCP equilibrium consists of allocations $\{C_{H,t}(s^t), C_{F,t}(s^t), C^*_{H,t}(s^t), C^*_{F,t}(s^t)\}$ with associated aggregates $\{C_t(s^t), C^*_t(s^t)\}$, labor $\{n_t(s^t), n^*_t(s^t)\}$, nominal exchange rates $\{\mathcal{E}_t(s^t)\}$, nominal wages $\{W_t(s^t), W^*_t(s^t)\}$, and Arrow-Debreu prices $\{\Lambda_t(s^t), \Lambda^*_t(s^t)\}$ such that:*

(i) Households in each region optimize;

(ii) Goods markets clear;

(iii) Asset prices are consistent with $\{\psi_t(s^t)\}$ (equation (45));

(iv) Nominal expenditure equals the money supply ("cash-in-advance"): $M_t(s^t) = P_t(s^t)C_t(s^t)$ and $M^*_t(s^t) = P^*_t C^*_t(s^t)$, $\forall t, s^t$;

(v) Export prices are given by $P^*_{H,t}(s^t) = P_{H,t}(s^t)/\mathcal{E}(s^t)$ and $P_{F,t} = \mathcal{E}(s^t)P^*_{Ft}(s^t)$.

To clarify the restrictions imposed by equilibrium, note that household optimality conditions (16)-(20), market clearing, and cash-in-advance imply

$$\begin{aligned} Y &= Zn = (1-\gamma)\left(\frac{P_H}{P}\right)^{-\theta} C + \gamma^*\left(\frac{P^*_H}{P^*}\right)^{-\theta} C^* = (1-\gamma)\left(\frac{P_H}{P}\right)^{-\theta}\frac{M}{P} + \gamma^*\left(\frac{P^*_H}{P^*}\right)^{-\theta}\frac{M^*}{P^*}, \\ Y^* &= Z^*n^* = \gamma\left(\frac{P_F}{P}\right)^{-\theta} C + (1-\gamma^*)\left(\frac{P^*_F}{P^*}\right)^{-\theta} C^* = \gamma\left(\frac{P_F}{P}\right)^{-\theta}\frac{M}{P} + (1-\gamma^*)\left(\frac{P^*_F}{P^*}\right)^{-\theta}\frac{M^*}{P^*}, \end{aligned} \tag{46}$$

where, letting $\mathcal{P}(P_H, P_F)$ and $\mathcal{P}^*(P_H, P_F)$ denote the ideal price indices for Home and Foreign, respectively, we have

$$\begin{aligned} P &= \mathcal{P}(P_H, \mathcal{E}P^*_F) \\ P^* &= \mathcal{P}^*(P_H/\mathcal{E}, P^*_F). \end{aligned}$$

The Backus-Smith condition (22) implies

$$\left(\frac{C}{C^*}\right)^{\sigma} = \bar{\omega}(1+\psi)\mathcal{Q} = \bar{\omega}(1+\psi)\frac{\mathcal{E}P^*}{P}.$$

Imposing the cash-in-advance conditions and rearranging yields:

$$\mathcal{E} = \frac{1}{\bar{\omega}(1+\psi)}\left(\frac{M}{M^*}\right)^{\sigma}\left(\frac{P^*}{P}\right)^{\sigma-1}. \tag{47}$$

Equations (46) and (47) determine $\mathcal{E}$, n, and n^* (or, equivalently, Y and Y^*) given $\{P_H, P_F^*\}$ and $\{M, M^*\}$, conditional on the PCP condition for export prices and the price aggregators. Real consumption allocations can then be recovered from prices and the cash-in-advance conditions. Arrow-Debreu prices are given by (17) and (20). If wages are flexible, then W and W^* are given from the households' static optimization condition. Firm profits are $(P_H Z - W)n$ and $(P_F^* Z^* - W^*)n^*$.

Local Currency Pricing (LCP) Under LCP, are local currency prices are pre-set.

Definition 2. *Given technology $\{Z_t(s^t), Z_t^*(s^t)\}$, monetary policies $\{M_t(s^t), M_t^*(s^t)\}$, asset market taxes/wedges $\{\psi_t(s^t)\}$, local-currency price sequences $\{P_{H,t}(s^t), P_{H,t}^*(s^t), P_{F,t}, P_{F,t}^*\}$, and initial assets w and w^*, a LCP equilibrium consists of allocations $\{C_{H,t}(s^t), C_{F,t}(s^t), C_{H,t}^*(s^t), C_{F,t}^*(s^t)\}$ with associated aggregates $\{C_t(s^t), C_t^*(s^t)\}$, labor $\{n_t(s^t), n_t^*(s^t)\}$, nominal exchange rates $\{\mathcal{E}_t(s^t)\}$, nominal wages $\{W_t(s^t), W_t^*(s^t)\}$, and Arrow-Debreu prices $\{\Lambda_t(s^t), \Lambda_t^*(s^t)\}$ such that:*

(i) Households in each region optimize;

(ii) Goods markets clear;

(iii) Asset prices are consistent with $\{\psi_t(s^t)\}$ (equation (45)); and

(iv) Nominal expenditure equals the money supply ("cash-in-advance"): $M_t(s^t) = P_t(s^t)C_t(s^t)$ and $M_t^*(s^t) = P_t^* C_t^*(s^t), \forall t, s^t$.

Under LCP, equation (46) pins down n and n^*, given monetary policies, as

$$P = \mathcal{P}(P_H, P_F^*)$$
$$P^* = \mathcal{P}^*(P_H, P_F^*).$$

Note that $\mathcal{E}$ does not play a role in output determination. The nominal exchange rate is then given by (47). The rest of the equilibrium objects are straightforward.

Pricing to Market (PTM) Under PTM, we assume prices are set according to

$$\begin{aligned} P_H &= (\tilde{W})^{1-\phi} P^{\phi}; \qquad P_H^* = (\tilde{W}/\mathcal{E})^{1-\phi}(P^*)^{\phi}, \\ P_F &= (\tilde{W}^*\mathcal{E})^{1-\phi} P^{\phi}; \quad P_F^* = (\tilde{W}^*)^{1-\phi}(P^*)^{\phi}, \end{aligned} \tag{48}$$

for $\phi \in [0,1)$, where $\tilde{W} \equiv W/Z$ and $\tilde{W}^* \equiv W^*/Z$ represent the respective marginal costs of production and

$$P = \mathcal{P}(P_H, P_F^*)$$
$$P^* = \mathcal{P}^*(P_H^*, P_F^*).$$

This implies

$$\frac{\mathcal{E}P_H^*}{P_H} = \frac{\mathcal{E}P_F^*}{P_F} = Q^{\phi}. \tag{49}$$

Definition 3. *Given technology $\{Z_t(s^t), Z_t^*(s^t)\}$, monetary policies $\{M_t(s^t), M_t^*(s^t)\}$, asset market taxes/wedges $\{\psi_t(s^t)\}$, and initial assets w and w^*, a* <u>*PTM equilibrium*</u> *consists of allocations $\{C_{H,t}(s^t), C_{F,t}(s^t), C_{H,t}^*(s^t), C_{F,t}^*(s^t)\}$ with associated aggregates $\{C_t(s^t), C_t^*(s^t)\}$, labor $\{n_t(s^t), n_t^*(s^t)\}$, nominal exchange rates $\{\mathcal{E}_t(s^t)\}$, nominal wages $\{W_t(s^t), W_t^*(s^t)\}$, Arrow-Debreu prices $\{\Lambda_t(s^t), \Lambda_t^*(s^t)\}$, and local-currency prices $\{P_{H,t}(s^t), P_{H,t}^*(s^t), P_{F,t}(s^t), P_{F,t}^*(s^t)\}$ such that:*

(i) Households in each region optimize;

(ii) Goods markets clear;

(iii) Asset prices are consistent with $\{\psi_t(s^t)\}$ (equation (45)); and

(iv) Nominal expenditure equals the money supply ("cash-in-advance"): $M_t(s^t) = P_t(s^t)C_t(s^t)$ and $M_t^*(s^t) = P_t^*C_t^*(s^t), \forall t, s^t$; and

(v) The PTM condition (48) is satisfied.

For each period, there are 13 equilibrium objects: 6 real allocations $\{C_H, C_F, C_H^*, C_F^*, n, n^*\}$ plus 7 nominal prices $\{P_H, P_F, P_H^*, P_F^*, W, W^*, \mathcal{E}\}$. These are simultaneously determined by the four household first-order conditions (16) – (20), the two static optimality conditions for leisure, the two market clearing conditions (46), the Backus-Smith condition (47), and the four PTM equations (48).

C Second-Order Approximation to Welfare

This appendix provides the details behind the second order approximation of welfare discussed in the last paragraph of Section 2.

Point of approximation We characterize a second order expansion of the welfare loss function $\mathbb{L}(\lambda, \delta)$ defined in (13) 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 for the baseline undistorted allocation, we have:[30]

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

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

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 $\delta \neq 0$.

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:

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

[30]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.

Next we expand the planner's optimality conditions for (C_H, C_F, C_H^*, C_F^*) from Appendix A:[31]

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

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 approximate 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),

[31] 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.

in proportional deviations around the efficient baseline:

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

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 (10), which we can write as:

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

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 $\omega U(C(C_H, Y^* - C_F^*)) + U(C^*(Y - C_H, C_F^*))$ 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 $\frac{\partial \mathbb{W}}{\partial C_H} = \omega U' \frac{\partial C}{\partial C_H} - U^{*\prime}\frac{\partial C^*}{\partial C_H^*}$ and $\frac{\partial \mathbb{W}}{\partial C_F^*} = -\omega U' \frac{\partial C}{\partial C_F} + U^{*\prime}\frac{\partial C^*}{\partial C_F^*}$. Therefore, we have:

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

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

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

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

and similarly for 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, $\bar{A} = \omega U'(\bar{C}) \cdot \frac{\partial C}{\partial \lambda}\big|_{\bar{C}} = \omega\bar{C}^{1-\sigma}[\bar{\alpha}(1-\bar{\gamma}) + \bar{\gamma}(1-\bar{\alpha}^*)]\kappa\theta$.*

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:

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

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, $\frac{1}{C}\frac{\partial C}{\partial\lambda}\big|_{\bar{C}} = \frac{\partial\log C}{\partial\lambda}\big|_{\bar{C}} = [(1-\bar{\gamma})\bar{\alpha}+(1-\bar{\alpha}^*)\bar{\gamma}]\,\kappa\theta$. □

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 $U(C^*)$.

D 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. The code to replicate all tables and figures can be found at the github repository: https://github.com/markaguiar/IRS-Replication-Package.git.

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 $j \neq k$.[32]

(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 trade prices are given by the price level used for PPP adjusted GDP; that is, Home exports to the rest-of-the-world at a price level equivalent to the PPP adjusted price of GDP, and imports at Foreign's PPP adjusted price plus a transport cost. This assumes Home exports (imports) are priced the same regardless of destination (origin), 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.

D.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

[32] Note that shipments from $j \in \mathcal{F}$ to Home carry the cost τ^*, as well.

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

$$\hat{A}_t \equiv \frac{A_t}{\bar{A}} = \hat{C}(\hat{A}_{H,t}, \hat{A}_{F,t}) \equiv \left((1-\bar{\gamma})\hat{A}_{H,t}^{\frac{\theta-1}{\theta}} + \bar{\gamma}\hat{A}_{F,t}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}, \tag{50}$$

where, from the definition of C, we have

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

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

$$\hat{Y}_t = \frac{\bar{A}_H}{\bar{Y}}\hat{A}_{H,t} + \frac{\tau\bar{A}^*_H}{\bar{Y}}\hat{A}^*_{H,t} = (1-\bar{\alpha})\hat{A}_{H,t} + \bar{\alpha}\hat{A}^*_{H,t}, \tag{51}$$

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,

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

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,

$$\frac{P_{H,t}C_{H,t}}{P_t C_t} = \frac{P_{H,t}A_{H,t}}{P_t A_t} \quad \Rightarrow C_{H,t} = \frac{C_t}{A_t} A_{H,t} \quad \Rightarrow \hat{C}_{H,t} = \frac{\hat{C}_t}{\hat{A}_t}\hat{A}_{H,t} \tag{52}$$

$$\frac{P_{F,t}C_{F,t}}{P_t C_t} = \frac{P_{F,t}A_{F,t}}{P_t A_t} \quad \Rightarrow C_{F,t} = \frac{C_t}{A_t} A_{F,t} \quad \Rightarrow \hat{C}_{F,t} = \frac{\hat{C}_t}{\hat{A}_t}\hat{A}_{F,t} \tag{53}$$

D.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

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

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

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

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:

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

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

$$\hat{A}_{F,t} \equiv \frac{1}{\bar{\alpha}^*} \sum_{j \in \mathcal{F}} \omega^Y_j \bar{\alpha}^*_{ji} \hat{A}_{ji,t}, \tag{54}$$

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

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

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):

$$\hat{A}_{F,t} = \frac{1}{\bar{\gamma}} \sum_{j \in \mathcal{F}} \gamma_{ji} \hat{A}_{ji,t}, \tag{55}$$

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 (54) equals (55) at every t, we need that weights on country j in the respective expressions must be equal. That is,

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

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:

$$\bar{\alpha}^* = \frac{\overline{IMP}}{\sum_{k\in\mathcal{F}} \overline{GDP}_k^{PPP}}. \tag{56}$$

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.[33] 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

$$\hat{A}_t^* = \sum_{j\in\mathcal{F}} \omega_j^A \gamma_{ij}\hat{A}_{ij,t} + \sum_{j\in\mathcal{F}} \omega_j^A \sum_{k\in\mathcal{F}} \gamma_{kj}\hat{A}_{kj,t}. \tag{57}$$

[33]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

$$\hat{A}^*_{H,t} = \sum_{j \in \mathcal{F}} \alpha_{ij} \hat{A}_{ij,t}, \tag{58}$$

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.[34] In order to write (57) as

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

we search for weights such that

$$\hat{A}^*_{H,t} = \frac{1}{\bar{\gamma}^*} \sum_{j \in \mathcal{F}} \omega^A_j \gamma_{ij} \hat{A}_{ij,t}, \tag{59}$$

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

Calibrating $\bar{\gamma}^*$ and ω^A_j: Comparing (58) to (59), 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.

D.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:

[34] 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:

$$
\begin{aligned}
\hat{A}_t &= \left((1-\bar{\gamma})\hat{A}_{H,t}^{\frac{\theta-1}{\theta}} + \bar{\gamma} A_{F,t}^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}, \\
\hat{A}_t^* &= \left((1-\bar{\gamma}^*)\hat{A}_{H,t}^{*\frac{\theta-1}{\theta}} + \bar{\gamma}^* \hat{A}_{F,t}^{*\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}, \\
\hat{Y}_t &= (1-\bar{\alpha})\hat{A}_{H,t} + \bar{\alpha}\hat{A}_{H,t}^*, \\
\hat{Y}_t^* &= \bar{\alpha}^*\hat{A}_{F,t} + (1-\bar{\alpha}^*)\hat{A}_{F,t}^*.
\end{aligned}
$$

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 1 Panel (a). 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 (52).

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

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

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

$$\hat{Y}_{C,t} = (1-\bar{\alpha}_C)\hat{C}_{H,t} + \bar{\alpha}_C\hat{C}^*_{H,t}. \tag{60}$$

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

$$\hat{Y}^*_{C,t} = \bar{\alpha}^*_C\hat{C}_{F,t} + (1-\bar{\alpha}^*_C)\hat{C}^*_{F,t}. \tag{61}$$

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):

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

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

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

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

$$
(\widehat{1+\delta})^\theta \equiv (1+\delta)^\theta * \left(\frac{1-\gamma}{\gamma}\right)\left(\frac{1-\gamma^*}{\gamma^*}\right)\frac{\bar{C}_F}{\bar{C}_H}\frac{\bar{C}_H^*}{\bar{C}_F^*} = \frac{(1+\delta)^\theta}{(1+\bar{\delta})^\theta}. \tag{62}
$$

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

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

D.4 Baseline Static Wedge

We can use the two-country planning problem to map the baseline static wedge into our parameters. In the first-best planning problem with trade costs, the static optimality condition becomes:

$$\tau^* \left(\frac{1-\gamma}{\gamma}\right)^{\frac{1}{\theta}} \left(\frac{A_H}{A_F}\right)^{\frac{-1}{\theta}} = \frac{1}{\tau}\left(\frac{1-\gamma^*}{\gamma^*}\right)^{\frac{1}{\theta}} \left(\frac{A_H^*}{A_F^*}\right)^{\frac{-1}{\theta}}.$$

We can define the baseline static wedge $\bar{\delta}$ such that:

$$\begin{aligned}
(1+\bar{\delta})\tau^* \left(\frac{1-\gamma}{\gamma}\right)^{\frac{1}{\theta}} \left(\frac{\bar{A}_H}{\bar{A}_F}\right)^{\frac{-1}{\theta}} &= \frac{1}{\tau}\left(\frac{\gamma^*}{1-\gamma^*}\right)^{\frac{1}{\theta}} \left(\frac{\bar{A}_H^*}{\bar{A}_F^*}\right)^{\frac{-1}{\theta}} \Rightarrow \\
(1+\bar{\delta})\frac{\tau^*\bar{A}_F}{\bar{A}_H} \left(\frac{1-\gamma}{\gamma}\right)^{\frac{1}{\theta}} \left(\frac{\bar{A}_H}{\bar{A}_F}\right)^{1-\frac{1}{\theta}} &= \frac{\bar{A}_F^*}{\tau\bar{A}_H^*}\left(\frac{\gamma^*}{1-\gamma^*}\right)^{\frac{1}{\theta}} \left(\frac{\bar{A}_H^*}{\bar{A}_F^*}\right)^{1-\frac{1}{\theta}} \Rightarrow \\
(1+\bar{\delta})\frac{\tau^*\bar{A}_F}{\bar{A}_H}\left(\frac{1-\bar{\gamma}}{\bar{\gamma}}\right) &= \frac{\bar{A}_F^*}{\tau\bar{A}_H^*}\left(\frac{\bar{\gamma}^*}{1-\bar{\gamma}^*}\right) \Rightarrow \\
(1+\bar{\delta})\frac{\tau^*\bar{A}_F}{\bar{A}_F^*}\left(\frac{1-\bar{\gamma}}{\bar{\gamma}}\right) &= \frac{\bar{A}_H}{\tau\bar{A}_H^*}\left(\frac{\bar{\gamma}^*}{1-\bar{\gamma}^*}\right) \Rightarrow \\
(1+\bar{\delta})\left(\frac{\bar{\alpha}^*}{1-\bar{\alpha}^*}\right)\left(\frac{1-\bar{\gamma}}{\bar{\gamma}}\right) &= \left(\frac{1-\bar{\alpha}}{\bar{\alpha}}\right)\left(\frac{\bar{\gamma}^*}{1-\bar{\gamma}^*}\right) \Rightarrow \\
1+\bar{\delta} &= \left(\frac{1-\bar{\alpha}}{\bar{\alpha}}\right)\left(\frac{1-\bar{\alpha}^*}{\bar{\alpha}^*}\right)\left(\frac{\bar{\gamma}}{1-\bar{\gamma}}\right)\left(\frac{\bar{\gamma}^*}{1-\bar{\gamma}^*}\right),
\end{aligned}$$

where the third line uses the definitions of $\bar{\gamma}$ and $\bar{\gamma}^*$ as the expenditure shares on imports in absorption, and the fifth line uses the definitions of $\bar{\alpha}$ and $\bar{\alpha}^*$ as the share of exports in output. Given our calibrated parameters, we can measure the implied static wedge in the baseline year. Note that we do not use this in order to compute the growth in the static wedge, which can be calculated without knowledge of the base year's value.