

November 2013

Assessing International Efficiency¹

Jonathan Heathcote

Federal Reserve Bank of Minneapolis and CEPR
heathcote@minneapolisfed.org

Fabrizio Perri

Federal Reserve Bank of Minneapolis, NBER, and CEPR
fabri.perri@gmail.com

Abstract

This paper is structured in three parts. The first part outlines the methodological steps, involving both theoretical and empirical work, for assessing whether an observed allocation of resources across countries is efficient. The second part applies the methodology to the long-run allocation of capital and consumption in a large cross section of countries. We find that countries that grow faster in the long run also tend to save more both domestically and internationally. These facts suggest that either the long-run allocation of resources across countries is inefficient, or that there is a systematic relation between fast growth and preference for delayed consumption. The third part applies the methodology to the allocation of resources across developed countries at the business cycle frequency. Here we discuss how evidence on international quantity comovement, exchange rates, asset prices, and international portfolio holdings can be used to assess efficiency. Overall, quantities and portfolios appear consistent with efficiency, while evidence from prices is difficult to interpret using standard models. The welfare costs associated with an inefficient allocation of resources over the business cycle can be significant if shocks to relative country permanent income are large. In those cases partial financial liberalization can lower welfare.

KEYWORDS: International risk sharing, Long-run risk, Long-run growth, International business cycles, Real exchange rate

JEL CLASSIFICATION CODES: F21, F32, F36, F41, F43, F44

¹Draft prepared for the Handbook of International Economics, Volume 4. We thank the editors Gita Gopinath and Ken Rogoff for their suggestions and their patience, Karen Lewis and Maury Obstfeld for insightful discussions, Alberto Polo for outstanding research assistance, and seminar participants at the Handbook Conference in Cambridge for great comments. Perri thanks the European Research Council for financial support under grant RESOCONBUCY. The data sets and computer codes used in the paper are available on the authors' websites. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

Is the observed allocation of resources across residents in different countries Pareto efficient? Or is it possible for a single government or an international organization to devise a mechanism (for example, a tax/subsidy system or the introduction of a new asset) so as to achieve a different allocation of resources that improves the welfare of residents in all countries? If observed allocations are inefficient, how large are the potential welfare gains from improving efficiency?

These questions cannot be answered by using theory alone, as our interest is in the efficiency of allocations we observe in the data, in a given set of countries and in a given time frame. At the same time, they cannot be answered with data alone, since the same data are in principle consistent with either efficiency or inefficiency depending on the underlying model of preferences, technologies and frictions.

Researchers have attempted to answer these questions in two popular strands of literatures in international macroeconomics. The first is the international consumption risk sharing literature (see, for example, the seminal work of Cole and Obstfeld, 1991) that deals with the allocation of consumption across countries and states of the world, taking as given the distribution of output. The second is the literature on the efficient distribution of productive assets across countries (see, for example, the work on capital by Lucas, 1990, and the work on labor by Hamilton and Whalley, 1984). The issue in this strand is whether world output and welfare can be increased by reallocating factors of production across countries.

The main objective of this chapter is to provide a simple but integrated methodological framework that lays down precisely the issues involved in combining data and theory to assess international efficiency along both of these dimensions.

The first part of the chapter (Section 2) describes in a general form the methodological steps that are needed to assess the efficiency of a given allocation of resources, and highlights the potential problems associated with each of these steps. The second part of the chapter (Sections 3 and 4) discusses two applications of the general methodology. These applications are closely related to various influential articles in the international macro literature, and our discussion of these applications within a single framework will highlight new connections and complementarities among these papers. Section 3 analyzes the long-run allocation of consumption and investment in a large cross section of countries. Section 4 deals with the allocation of consumption and investment in developed countries over the business cycle. Section 5 concludes, attempting an answer to the efficiency questions posed at the beginning and pointing to future interesting research directions.

The main limitation of our survey is that we follow the traditional approach in international macro and assume an efficient distribution of resources within a country (i.e., the existence of a representative agent/firm within a country). We will not discuss recent and interesting research (e.g., Kocherlakota and Pistaferri, 2007, and Mendoza, Quadrini, and Rios-Rull, 2009) that studies

the international allocation of resources in a world where the intranational distribution of resources is not efficient.

2 A Methodology for Assessing International Efficiency

In this section we outline the general methodological steps that are needed to assess the efficiency of a given allocation.

2.1 Specifying Preferences, Technologies and Frictions

The first necessary step in assessing whether various features of the data (e.g., the international comovement of consumption, capital flows between countries) are consistent with efficiency is to specify a model economy, i.e., preferences, technologies, and frictions. The model economy can then be used to generate theoretical counterparts to the empirical variables of interest.

2.1.1 Preferences

This step is essential as preferences ultimately determine the value of transferring resources across countries. Absent restrictions on preferences, it is impossible to determine whether allocations are efficient and to quantify the welfare costs of any inefficiencies. To see this, consider the following example. Suppose that during a global recession, we observe country A reducing consumption by more than country B. In some models – with symmetric preferences – this observation would be interpreted as a lack of consumption risk sharing and hence inefficiency. In alternative models – with asymmetric preferences – this same allocation can be efficient. For example, if country A is more risk tolerant than country B, then it is efficient for country A to take on a bigger share of global risk and hence reduce consumption more in a global recession (for a model of this type, see Gourinchas, Rey, and Govillot, 2010).

As another example, consider a model in which the efficient consumption allocation is the one that equalizes consumption growth rates across different countries, and suppose that a researcher is interested in assessing the gains of moving from the observed consumption allocations (in which growth rates are not equalized) to the efficient allocation. Different assumptions about preferences can make the gains from risk sharing arbitrarily large (for example, if agents are extremely risk averse) or arbitrarily small (if preferences are close to linear).

As is well known (see Stigler and Becker, 1977), the preference problem is endemic in economics; here we just want to stress that it is of first order importance in international efficiency problems. Ideally, researchers should justify assumptions about preferences, preference heterogeneity, and/or preference shocks using observables (e.g., asset price data, long-run trends, trade flows).

2.1.2 Technologies

The specification of technology concerns the primitive (i.e., taken as given by the researcher) distribution of resources across countries/agents, time, and states of the world in the economy. Examples include the endowments of goods, labor, capital, total factor productivity, or productive opportunities. As with preferences, the distribution of resources should be pinned down by observables, but, unlike preferences, the connection between model and observables is usually more direct. Consider, for example, the issue of specifying a process of endowments of consumption goods in each country in the classic international consumption risk-sharing problem. In this case, a researcher can identify these endowments simply by constructing time series of the production of tradable consumption goods in each country, using national accounts data.

In many international business cycles studies, the primitive resource that is assumed to differ across countries is total factor productivity (TFP). A researcher can construct time series for TFP across countries using data from national accounts plus assumptions on the production functions.

An important remark is that observables are sometimes not sufficient to distinguish whether differences in resources among countries are due to *ex post* risk or *ex ante* heterogeneity, but this distinction has important implications for efficiency. Consider, for example, two poor countries and assume at some point that we observe one of the countries extracting a lot of oil. If the presence of oil was not known to residents in the two countries at the beginning of time, it would be (*ex ante*) efficient for them to share this resource risk, and efficient allocations would all involve substantial transfers from the lucky country to the unlucky. If, on the other hand, this difference was known to the residents at date zero, allocations that do not involve any transfer can also be efficient.

2.1.3 Frictions

Frictions are constraints that all market allocations (efficient or not) have to satisfy because of some physical or technological features of the environment. A classical example of a friction in international macro is limited tradability of goods: it is often assumed that a fraction of resources in a given country cannot be shipped to other countries. To see why frictions matter for assessing efficiency, consider the extreme example (borrowed from Brandt, Cochrane, and Santa Clara, 2006) of Earth and Mars. Suppose both planets face income risk, but shipping any goods between them is impossible. In this case, the resulting market allocation is that in each planet consumption is equal to income. This allocation of resources is efficient because no other allocation satisfies the physical no-trade constraint (the friction) and yields higher welfare to residents. A very influential paper by Obstfeld and Rogoff (2001) argues that many aspects of international macro data that suggest inefficiency no longer do so once they are analyzed within a model that features limited tradability.

Another friction often introduced in international macro models is the assumption of a limited

enforcement technology on international contracts. In particular, it is assumed that if countries default on international obligations, the harshest punishment that can be imposed on them is exclusion from future trade (autarky). This friction implies that any market allocation has to yield, in each date and in each state, expected welfare to any country at least as high as the expected welfare under autarky (see, for example, Kehoe and Perri, 2002). This typically rules out allocations that involve large intertemporal transfers between countries, which reduces the set of allocations that can be achieved by a world planner/policymaker. If there is no Pareto-improving reallocation of resources that preserves incentives to repay debts or report truthfully in environments with enforcement or information frictions, allocations are labeled “constrained efficient.”

2.2 Efficient Allocations and Market Allocations

Once the fundamentals of the economy are described, a researcher can first characterize (analytically or numerically) the set of efficient allocations, which is usually done by solving a (constrained) planning problem. These efficient allocations are a natural baseline to compare with data.

However, which features of efficient allocations are hallmarks of international efficiency will not always be obvious. For this reason a useful step is to compute the theoretical predictions of alternative models in which there is no world social planner, and in which allocations are determined in a competitive equilibrium in which agents trade an exogenously limited set of assets internationally. Examples of commonly assumed market structures are autarky (no markets across countries), financial autarky (no intertemporal markets between countries), limited asset trade between countries (e.g., a single noncontingent bond), and complete markets within and between countries. Before turning to the data, an instructive approach will be to compare and contrast the predictions of alternative market structures alongside the constrained efficient baseline, in order to learn which features of the data are more or less sensitive to the scope for international asset trade, and – relatedly – which moments offer the sharpest tests of international efficiency. It will also be useful to learn when and whether trade in a limited set of assets can perfectly decentralize constrained-efficient allocations.

2.3 Comparing Models and Data

This step involves the comparison of several model allocations with data, to get a sense of which setup can better account for the data. Obviously, there are many dimensions along which one can perform this comparison. Many authors have focused on the international correlations of quantities such as GDP, consumption, and investment at a business cycle frequency (see, for example Baxter and Crucini, 1995), since in some models efficient and inefficient allocations yield very different correlations of these quantities. Another commonly used statistic involves comovement between consumption ratios and real exchange rates (Backus and Smith, 1993). More recently some

authors have also suggested using asset prices (Brandt, Cochrane, and Santa Clara, 2006), portfolios (Heathcote and Perri, 2013), or capital flows (Gourinchas and Jeanne, 2011) as additional pieces of evidence against which researchers should benchmark models. Other authors have used seminatural experiments, such as financial liberalizations, to assess whether responses to these observed changes suggested efficient or inefficient allocations of resources across countries (e.g., Kose et al., 2009).

Ideally, one should use as much relevant empirical evidence as possible to discriminate between different models, because bringing in more data gives the researcher more confidence in evaluating whether an observed allocation is efficient. However, when considering any particular dimension of the data, it is important that at least one theoretical allocation (efficient or inefficient) comes close to replicating the empirical moments of interest. If none of the models on the table can account for certain features of the data, then the combination of those models and those moments is not useful for learning about efficiency. An example of this issue, which we will discuss in detail, is that it is difficult to use moments involving the real exchange rate to assess efficiency in the context of models that cannot replicate basic properties of real exchange rate dynamics.

2.4 Assessing Welfare Gains and Designing Policy Interventions

Once we have established that a model and an associated market structure offers a reasonable account of several dimensions of relevant data, we can use the model to assess efficiency and answer additional questions. The first is simply to ask, in case the allocations resulting from the model that best fits the data are not efficient, how big are the welfare gains of moving from the observed allocation to an allocation within the set of efficient allocations. This is a number in which many researchers have been interested (see, for example, Cole and Obstfeld, 1991, or Gourinchas and Jeanne, 2006) and a number that, unfortunately, differs widely across different studies. A second question is why, within the context of the model, efficiency is not achieved, and whether instruments are available to a policymaker that could improve welfare while still respecting the frictions in the environment.

We will now proceed to illustrate all of these steps in concrete applications.

3 Assessing Long-Run Efficiency

We now follow the steps described above to assess the efficiency of long-run allocations of consumption and capital across a large cross section of countries, specifically the set of countries in the Penn World Tables, which have continuous data for the period 1960-2010.

3.1 Preferences, Technologies, and Frictions

We will think about each country in the data as being small relative to a fictional “world economy.” The role of the world economy is to pin down the world interest rate. There is one tradable good used for consumption and investment (later we will discuss introducing a nontradable sector). A representative agent in each small country i has preferences

$$\sum_{t=0}^{\infty} \beta^t u(C_{it}, \phi_{it}),$$

where

$$u(C_{it}, \phi_{it}) = \phi_{it} \frac{C_{it}^{1-\sigma}}{1-\sigma}$$

and ϕ_{it} is a country and date-specific preference shifter.

The production technology in country i is

$$Y_{it} = F(K_{it}, A_{it}) = K_{it}^{\alpha} A_{it}^{1-\alpha},$$

where K_{it} and A_{it} denote, respectively, capital and labor productivity (hours worked are assumed constant and normalized to one). At date zero, per-capita capital and productivity in each country i are assumed identical to those in the world economy: $K_{i0} = K_0$ and $A_{i0} = A_0$.

The representative agent in the world economy has a similar utility function, absent the preference shifters. World productivity grows at a constant rate $A_{t+1}/A_t = \gamma$. Thus, the world economy features a balanced growth path along which output, consumption and investment all grow at rate γ . The constant gross interest rate along this balanced growth path is given by $R = \gamma^{\sigma}/\beta$.

The risk each small country i faces is growth rate risk. Country i will experience a country-specific growth rate for labor productivity, $A_{i,t+1}/A_{i,t} = \gamma_i$ for all $t \geq 0$. We consider two alternative models for how information about γ_i is revealed. In the first model, which we label “perfect foresight,” we assume that γ_i is revealed at date 0, and from that date onward agents are perfectly informed about productivity at each future date. Thus, for example, this model presumes that in 1960 everyone knew that Korea would subsequently grow quickly while Argentina would grow slowly.

In the second model, which we label “repeated surprises,” we make the opposite assumption and assume that at each date t , agents assign probability 1 to the event $A_{i,\tau+1}/A_{i,\tau} = \gamma$ for all $\tau \geq t$. Subsequently, they are repeatedly surprised to observe realized growth $A_{i,\tau+1}/A_{i,\tau} = \gamma_i$.

3.2 Efficient Allocations

Allocations $\{C_{it}, K_{i,t+1}\}$ in country i are efficient if they solve the following two planner problems:

1. The time path for consumption $\{C_{it}\}$ solves

$$\max_{\{C_{it}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(C_{it}, \phi_{it})$$

subject to

$$\sum_{t=0}^{\infty} \frac{C_{it}}{R^t} \leq B_{i0}$$

for some present value of consumption $B_{i0} > 0$ allocated to country i .

2. The time path for capital $\{K_{i,t+1}\}$ solves the following series of problems:

$$\max_{K_{i,t+1}} \{E_t [F(K_{i,t+1}, A_{i,t+1})] + (1 - \delta)K_{i,t+1} - RK_{i,t+1}\} \quad \forall t,$$

where the expectation is over possible values for $A_{i,t+1}$ and is conditional on the sequence $\{A_{i\tau}\}_{\tau=0}^t$. Under the information structures described above, agents (and the planner) assign probability 1 to the value $A_{i,t+1} = \gamma_i A_{it}$ in the perfect foresight model, and assign probability 1 to $A_{i,t+1} = \gamma A_{it}$ in the repeated surprises model.

Allocations for consumption that solve the first problem satisfy consumption efficiency. The first-order conditions with respect to C_{it} and $C_{i,t+1}$ imply

$$\frac{\phi_{it}}{\beta \phi_{i,t+1}} \left(\frac{C_{it}}{C_{i,t+1}} \right)^{-\sigma} = R \quad \forall t. \quad (1)$$

Thus, the intertemporal marginal rate of substitution of consumption in each country i is equated to the world gross return to capital. Different choices for B_{i0} correspond to different levels for country i 's consumption, each of which corresponds to a different point on the Pareto frontier.

Allocations for capital that solve the second problem satisfy production efficiency. The first-order condition with respect to $K_{i,t+1}$ is

$$E_t \left[\alpha \left(\frac{K_{i,t+1}}{A_{i,t+1}} \right)^{\alpha-1} \right] + (1 - \delta) = R \quad \forall t \quad (2)$$

Thus, the expected marginal product of capital is equal to that in the world economy for all $t \geq 1$.

Note that consumption efficiency (eq. 1) is a difficult equation to test empirically, absent knowledge of the preference shifters ϕ_{it} . Production efficiency (eq. 2) is in principle easier to test because it does not involve preferences.²

²However, note that since capital must be put in place one period in advance, if the realized value for $A_{i,t+1}$ differs from the expected value, the realized marginal product of capital will differ from the world interest rate. Still, given those expectations, the allocation of capital is efficient ex ante. There is no expectation sign in the consumption efficiency condition, because consumption can be instantaneously reallocated across countries.

Suppose we assume common preferences across countries (i.e., $\phi_{it} = 1$ for all i and for all t). This will be our baseline assumption. Then the consumption efficiency condition 1 simplifies to

$$\frac{1}{\beta} \left(\frac{C_{i,t+1}}{C_{it}} \right)^\sigma = R \quad \forall t,$$

which implies that all countries share the same consumption growth rate.

3.3 Market Allocations

Now that we have characterized efficient allocations, we will consider decentralized competitive equilibria under alternative explicit market structures, to investigate when and where deviations from efficiency arise.

Financial Autarky Here we assume no asset trade between countries. The absence of asset trade means that each country's net exports must be zero at each date, because there is no way to import the tradable good in return for a contractual promise to export the tradable good at a future date. The resource constraint is

$$C_{it} + K_{i,t+1} = F(K_{it}, A_{it}) + (1 - \delta)K_{it} \quad \forall t.$$

Under financial autarky, we can envision allocations in each country being determined by a country-specific planner who maximizes expected lifetime utility subject to the resource constraint. The first-order condition for capital accumulation is

$$\phi_{it} C_{it}^{-\sigma} = \beta E_t \left[\phi_{i,t+1} C_{i,t+1}^{-\sigma} \left(1 + \alpha \left(\frac{K_{i,t+1}}{A_{i,t+1}} \right)^{\alpha-1} - \delta \right) \right].$$

The previous equation indicates that agents will choose to equate the expected marginal rate of substitution to the expected marginal rate of transformation. However, absent international asset trade, the marginal rate of substitution will not in general be equalized across countries. In contrast, equilibrium consumption growth rates will be country specific and mirror country-specific productivity growth rates. Given differential consumption growth, countries will optimally choose country-specific marginal products of capital. This teaches us something useful about the two efficiency conditions described above, namely, that efficiency requires that both hold jointly. In the financial autarky economy, when missing asset markets lead to a deviation from consumption efficiency, it is not optimal to equate the marginal product of capital across countries, and so the production efficiency condition is not satisfied either.

Bond Economy Under this market structure, agents in country i can borrow and lend from the world economy by trading an international one period bond whose price is R^{-1} . We assume

that residents in country i hold all the domestic capital and finance all domestic investment. At each date, country i faces a budget constraint of the form

$$C_{it} + K_{i,t+1} + \frac{B_{i,t+1}}{R} = F(K_{it}, A_{it}) + (1 - \delta)K_{it} + B_{it} \quad \forall t,$$

where $B_{i0} = 0$.³

Whether efficiency is achieved in the bond economy model depends on the model for expectations. Given perfect foresight, trade in a bond delivers efficiency. To understand this result, consider the capital and bond accumulation choices for the representative agent in country i . Given perfect foresight, the two corresponding intertemporal first-order conditions deliver the two hallmark conditions for efficiency described above: the intertemporal marginal rate of substitution is equalized across countries, and the marginal product of capital is equalized across countries. Note that asset trade is crucial to delivering this outcome. In particular, bond trade equates the marginal rate of substitution across countries, since the bond offers countries a common rate at which to exchange current for future consumption. Then arbitrage within each country leads to investment choices that equate the expected (country-specific) marginal product of capital to the (common global) interest rate.

Things are slightly different in the repeated surprise version of the bond economy. The expected marginal rate of substitution is again equalized across countries and equal to the world interest rate. Arbitrage again equates the expected marginal product of capital across countries to the world interest rate, thereby achieving productive efficiency. However, consumption efficiency is not achieved. Although the *expected* marginal rate of substitution is equated, ex post fast-growing countries will enjoy faster consumption growth than slow-growing countries. Given common preferences, this is inefficient.

Complete Markets In this economy, people trade a full set of state-contingent claims at each date.⁴

In the complete markets model, allocations are always efficient. Each country invests to equate the expected return to capital to the world interest rate. Trade in contingent claims ensures that the marginal utility of consumption in each country i grows at the same rate as in the world economy. Absent preference shifters (i.e., assuming $\phi_{it} = 1$), the level of consumption is equal to that in the world economy at each date, $C_{it} = C_t$.

³An alternative would be to assume that agents hold all their wealth in the international bond and that foreigners own all domestic capital. Given perfect foresight about productivity growth, these two alternative assumptions on portfolios would be identical, since returns will be equalized across countries. In the repeated surprises model, expected returns will be equated, but the assumption about who owns domestic capital will have a minor impact on ex post returns.

⁴In the perfect foresight version of the model, an alternative way to complete markets is to assume that agents initially trade Arrow securities contingent on the realization of the vector $\{\gamma_i\}$. After the vector $\{\gamma_i\}$ is realized, the securities pay out. Then, and in every subsequent period, the market structure corresponds to that in the bond economy model, where the starting bond position B_{i0} is the payoff from initial trade in Arrow securities.

Although allocations in the complete markets and bond economy models are both efficient under perfect foresight, the two market structures pick out different points on the Pareto frontier. With complete markets, insurance in the initial period translates into growth-rate-specific initial transfers that equate the present value of consumption across countries. With only a bond, in contrast, the present value of each country’s consumption reflects the present value of country-specific net output, corresponding to the planner’s problem defining efficient allocations with $B_{i0} = 0$. In this case (assuming common preferences), countries with faster productivity growth will enjoy higher consumption at each date.

3.4 Comparing models and data

We start the section by first describing some general features of the data we are going to use to assess efficiency.

3.4.1 Data

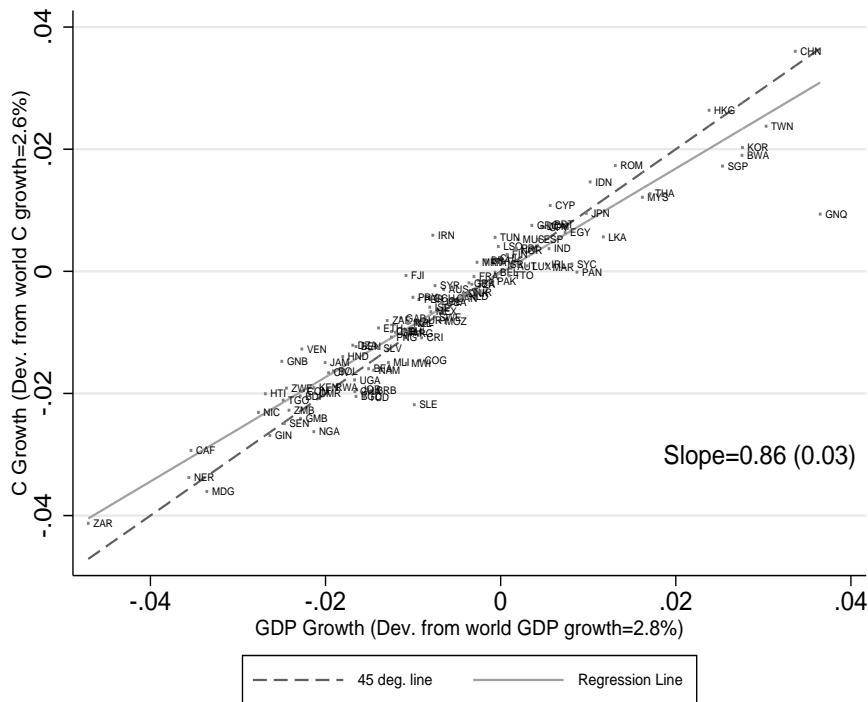
In Figures 1, 2, and 3 we describe some details of the growth experiences of all the 112 countries that have continuous data in the Penn World Tables over the period 1960–2010. First consider in Figure 1 the plot of consumption growth against output growth. Growth rates show dramatic variation, ranging from some African countries, where output grew as many as 4 percentage points per year slower than the world average, to China, where output grew 4 percentage points faster. In terms of corresponding growth in consumption, countries almost line up along the 45 degree line. However, the least squares regression line suggests that faster output growth does not translate quite one-for-one into faster consumption growth: if one country’s output grows 1 percentage point per year faster than another’s, the faster-growing country on average enjoys a 0.86 percentage point faster growth rate for consumption.

Next, Figure 2 plots the relationship between output growth on the one hand and the growth of investment on the other. Notice that once again, countries tend to line up along the 45 degree line. However, the least squares regression line suggests that faster output growth translates more than one-for-one into faster investment growth. If one country’s output grows 1 percentage point per year faster than another’s, the faster growing country experiences a 1.07 percentage point faster growth rate for investment.

Finally, in Figure 3 we plot the relationship between output growth and the end of sample net foreign asset (NFA) position.⁵ Here we note that there is not much evidence of a systematic relationship between the two variables: the set of country points form something of a cloud. To the extent that there is a relationship, it is positive, as Gourinchas and Jeanne (2011) originally

⁵Net foreign asset position data are from Lane and Milesi-Ferretti (2007) and refer to the end of 2007. The number of countries represented in Figure 3 is slightly smaller than the number in Figures 1 and 2 (108 vs. 112), since NFA data are not available for all countries in the Penn World Tables.

Figure 1: Long-run GDP and consumption growth

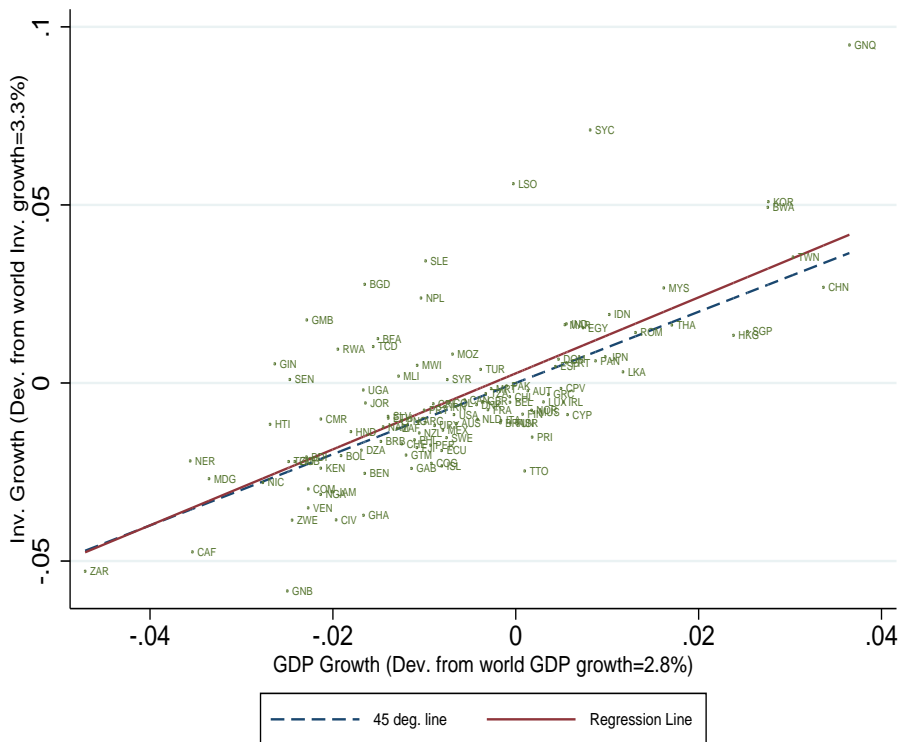


emphasized for the set of developing countries. Faster-growing countries (like Singapore or China) tend to have accumulated positive NFA positions, whereas slow-growing countries (like Niger or Nicaragua) have accumulated negative positions. Alfaro et al. (2011) argue that once one strips official sovereign flows out of international capital flows, fast growers on average are net recipients of private international capital flows, though the relationship remains noisy. Finally, note that most countries' absolute NFA positions are smaller than 50 percent of their GDP.

3.4.2 The Perfect Foresight Model: Predictions

We now quantitatively compare the predictions of the three market structures with the data. We set the preference parameters β and σ to relatively standard values of 0.97 and 2. We set the technology parameters α and δ to 0.36 and 0.06. We set the growth rate of labor productivity in the world economy γ so that aggregate consumption at each date is equal to average consumption across a set of bond economies, where the distribution of country productivity growth rates in the set corresponds to the distribution of output growth in our Penn World Table sample. The implied growth rate for world productivity is 2.46% per year, so $\gamma = 1.0246$. Our choices for γ , σ , and β

Figure 2: Long-run GDP and investment growth



translate into a constant world interest rate of 8.2%.⁶

We then consider a range of constant country-specific growth rates from 4% slower to 4% faster than the world growth rate, which covers the range of country experiences in our data. Thus, $\gamma_i \in [1.0246 - 0.04, 1.0246 + 0.04]$.⁷

Figure 4 plots the model predictions for each market structure, assuming perfect foresight about country-specific productivity growth. Panels *A* and *B* plot average consumption and investment

⁶This interest rate is high relative to most empirical estimates. It is high because this is a model with growth. Setting $\beta = 1$ would reduce R to 5.0%, but would be less appealing from the standpoint of the welfare calculations presented later. An alternative approach would be to use non-time-separable preferences à la Epstein and Zin (1989).

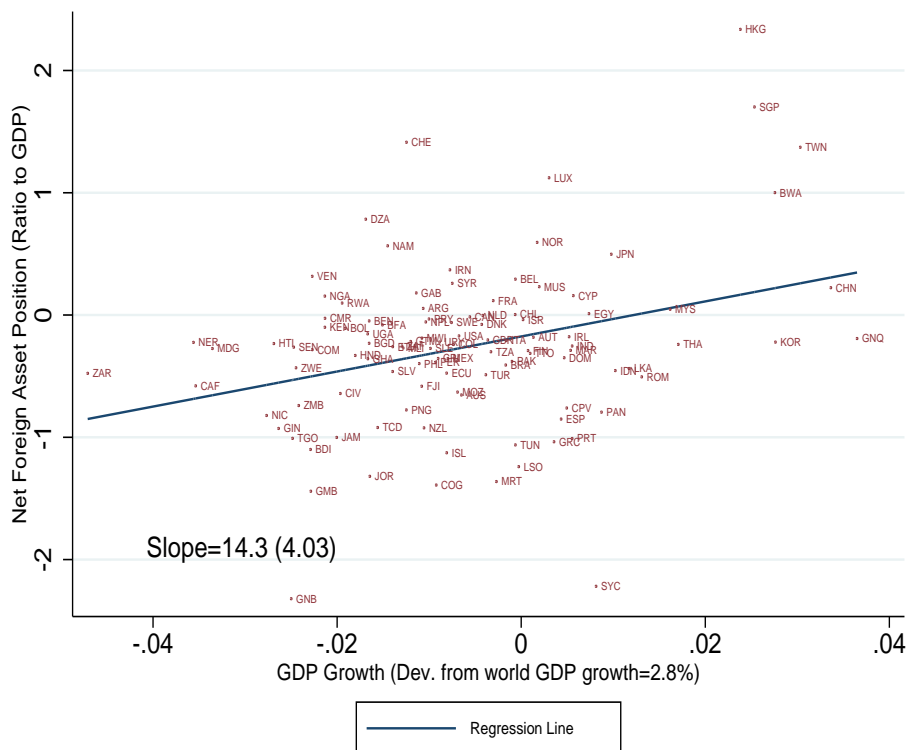
⁷Characterizing equilibria for these economies is fairly straightforward. Given a fixed and exogenous world interest rate, the bond economy model is analytically tractable. At each date, consumption is set such that the expected present value of current and future consumption equals the expected present value of current and future labor earnings plus the gross return on initial wealth. For example, date 0 consumption in the perfect foresight version of the bond model is given by

$$\frac{C_{i0}}{1 - \frac{\gamma}{R}} = \frac{(1 - \alpha)K_{i0}^\alpha A_{i0}^{1-\alpha}}{1 - \frac{\gamma_i}{R}} + RK_{i0},$$

where the left-hand side defines the present value of consumption (which grows at rate γ), and the right-hand side captures the expected present value of lifetime earnings (which grows at rate γ_i), plus the gross return on initial wealth.

Allocations in the autarky model must be characterized numerically. We guess an initial value for consumption and then use the intertemporal first-order condition for investment alongside the resource constraint to iterate forward and verify convergence to the balanced growth path.

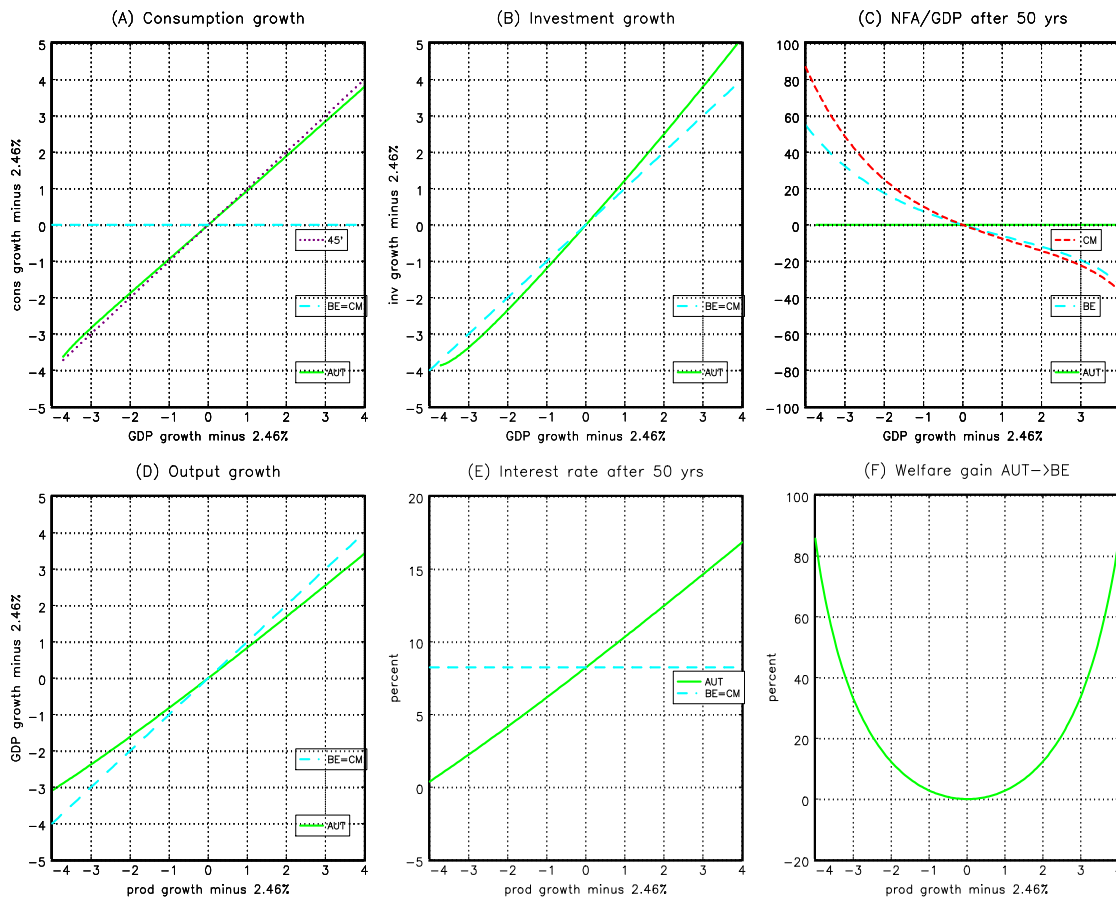
Figure 3: Long-run GDP growth and Net Foreign Asset Positions



growth against average output growth. Panel *C* plots the ratio of net foreign assets to output in year 50 – a value of one means that holdings of the international bond are equal to output. Panel *D* plots average annual output growth over a 50-year period relative to annual labor productivity growth γ_i . Panel *E* plots the net interest rate in year 50. Finally, Panel *F* shows the welfare gain of being able to trade the international bond relative to autarky, measured as the constant percentage increase in autarky consumption required to deliver equal lifetime utility to that achieved in the bond economy.

First consider autarky. Here, because net exports are zero at each date, the net foreign asset position remains constant at zero. Faster productivity growth translates into faster output growth (Panel *D*), and because there is no scope for international borrowing and lending, faster output growth translates into faster consumption growth (Panel *A*). Indeed, if countries have time-invariant preferences, then in the limit as $t \rightarrow \infty$, each country i will converge to a country-specific balanced growth path, in which capital, output, and consumption will all grow at the country-specific growth rate for labor productivity γ_i . Differentials in consumption growth translate into interest rate differentials, with faster-growing countries having higher interest rates (Panel *E*). These interest rate differentials are very large: at a 5% growth rate, the balanced growth path interest rate is $1.05^2/0.97 = 13.7\%$ whereas at a 0% growth rate, the interest rate is $1/0.97 = 3.1\%$.

Figure 4: Long-run growth with perfect foresight



A couple of details are worth noting about the autarky economy. First, for slow-growing countries, because the domestic balanced growth path interest rate is lower than for the world economy, capital must grow more rapidly than productivity during transition. Thus, in slow-growing countries, output growth exceeds productivity growth (Panel *D*). Second, in slow-growing countries, consumption growth tends to exceed output growth, whereas investment grows more slowly than output (Panels *A* and *B*). Again, this is because slow-growing countries divert a relatively large share of output to investment rather than consumption early in the transition.

Now look at the bond economy and the complete markets economy. The first thing to note is that the implications of these two economies are quite similar. First, consumption efficiency (eq. 1) implies that country consumption growth is divorced from country output growth. Production efficiency (eq. 2) implies that interest rates and the marginal product of capital are equated across countries (Panel *E*), and the growth rates of country output and investment are therefore identical to the growth rate of country productivity (Panels *B* and *D*). With common preferences across countries, as in this example, consumption growth is equated across countries (Panel *A*).

Although the paths for capital and output in the complete markets and bond economy models are identical, the levels of country-specific consumption paths, as well as the dynamics for net exports and net foreign asset positions, differ slightly across the two market structures. In both cases, however, fast-growing countries have accumulated large negative net foreign asset positions after 50 years of fast growth. For example, in the bond economy model, a country growing consistently 1 percentage point faster than the world economy has a negative net foreign asset position exceeding 600% of GDP. The logic is simply that a country that knows it will grow fast and thus has high permanent income relative to current income at date 0 sets initial consumption equal to permanent income and finances the gap between permanent and current income by borrowing from abroad. Net foreign asset positions are even larger in the complete markets economy, since countries that draw fast growth rates must make large initial transfers, and thus begin the transition with large negative net foreign asset positions.

3.4.3 The Perfect Foresight Model: Comparing with Data

Which market structure predicts outcomes that most closely approximate the historical experiences of actual economies as described above?

At first glance, the implications of the bond and complete markets models appear grossly counterfactual. First, consumption growth closely tracks output growth in our sample of countries. Countries that have enjoyed relatively fast economic growth (like Korea) now enjoy higher consumption levels than countries that have not (like Argentina). This stands in stark contrast to the complete markets and bond economy models, in which – given perfect foresight – consumption should grow at the same rate in all countries.

However, it is important to note that ours is a model in which all output is tradable. Suppose a fraction of output actually comprises nontradable goods. Nontradable consumption will – by definition – track nontradable output. If countries with faster growth in aggregate output also enjoy faster growth in nontradable output, then aggregate consumption will tend to track aggregate output, even if asset markets are complete. The quantitative predictions of such a model for growth in aggregate output and consumption will depend on the details of how nontradables are introduced. With separable preferences over tradable and nontradable consumption, as well as an endowment process for nontradable output, all the pictures plotted in Figure 4 still apply, except that now they should be interpreted as applying to the tradable sector only. How would the predictions for aggregate consumption and output change? To develop one concrete example, suppose that preferences over tradables c_t^T and nontradables c_t^N take the form

$$u(c_t^T, c_t^N) = \alpha \log c_t^T + (1 - \alpha) \log c_t^N.$$

Suppose in addition that the nontradable endowment grows at rate γ_i , the growth rate of labor

productivity in the tradable sector. In such a model, aggregate consumption in country i will grow at gross rate $\gamma^\alpha \gamma_i^{(1-\alpha)}$. The growth rate of aggregate output will vary over time, but at the date when tradable consumption equals tradable output, output will be growing at rate γ_i . Thus, the larger is nontradables share in consumption $(1 - \alpha)$, the closer will be the growth rate of consumption to the growth rate of output. In the data, the slope of a regression of consumption growth on output growth is 0.86 over the period 1960-2010, which the model can replicate given a nontradable share of $(1 - \alpha) = 0.86$. This exceeds all reasonable estimates of the fraction of output that is nontradable, indicating that while introducing nontradables can account for some of the comovement between output and consumption in our cross section of countries, it cannot explain all of it: comovement remains between growth in output of tradables and growth in consumption of tradables.

A second problem with the complete markets and bond economy models is that they predict enormous net foreign asset positions, with fast-growing countries accumulating large negative net foreign asset positions. Introducing a nontradable sector in the model would imply smaller net foreign asset positions. For example, suppose for a country with a particular growth rate γ_i , tradable output was 50% of total output after 50 years. Then the net foreign asset position relative to total output would be half as large as the one suggested by Panel *C* of Figure 4. However, such positions would still be much larger than those observed in the data. Moreover, the systematic theoretical link between faster growth and more negative NFA positions is absent in the data, where the correlation between past growth and the current NFA position is positive.

That seems to leave the financial autarky model as the most plausible baseline market structure. Indeed, in some respects the autarky model offers a reasonable account of the nature of growth across fast- versus slow-growing countries. As noted above, the autarky model replicates the fact that consumption (investment) growth in relatively slow-growing countries tends to be faster (slower) than output growth. However, the financial autarky model faces some challenges of its own. In particular, that model implies very large differences in the marginal product of capital across countries, whereas in practice the marginal product of capital appears to be roughly equalized (see Caselli and Feyrer, 2007). Relatedly, the model is also inconsistent with Kaldor's (1957) economic growth facts, since the capital-to-output ratio rises over time in slow-growing countries and falls in fast-growing countries.

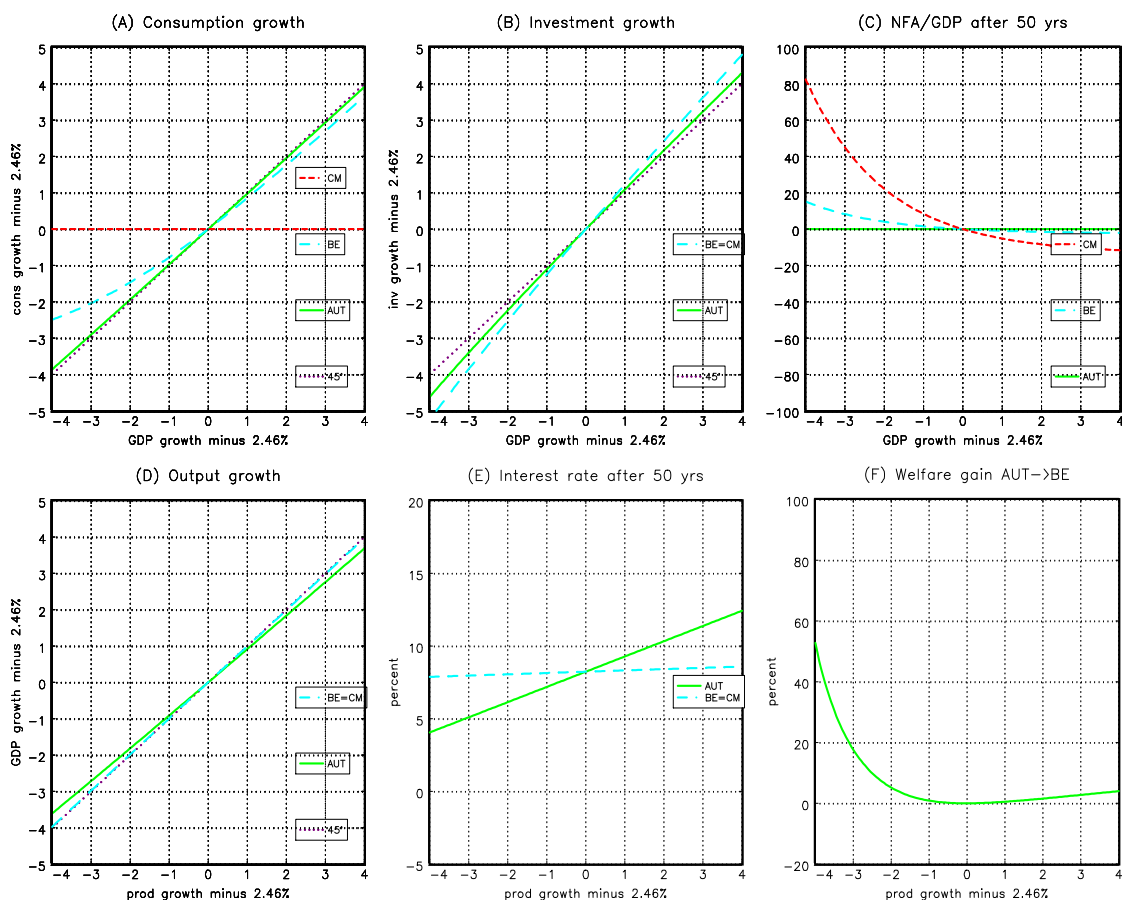
3.4.4 Alternative Model 1: Repeated Surprises

A key challenge to the models presented thus far is that it is difficult to reconcile cross-country differences in expected consumption growth rates with common-across-countries returns to capital. The autarky model generates differences in long-term expected growth rates but also implies country-specific returns, whereas conversely asset trade ensures a common world interest but seems

to dictate equalization of expected growth rates.

One way to reconcile the two facts is to postulate that fast-growing countries never expected such rapid growth, and that slow-growing countries never expected to stagnate relative to the rest of the world. If all countries expect identical future productivity growth, then trade in a bond will equate expected consumption growth rates (and expected returns to capital) but will not equate realized consumption growth if some countries consistently enjoy faster productivity growth than others. Figure 5 describes the growth dynamics under the “repeated surprises” scenario, in which agents always expect country productivity to grow at a 2.46% rate.⁸

Figure 5: Long-run growth with repeated surprises



Here, the plots for complete markets and autarky look very similar to those for the perfect foresight model. Under complete markets, consumption is again equalized across countries. Because agents underestimate growth in fast-growing countries, and installed capital cannot be instantly

⁸This model is slightly more complicated to solve. In autarky, at each date t , we solve for consumption such that given expected future productivity growth at rate γ , the economy converges to the world economy balanced growth path. This value for consumption determines the capital-to-output ratio in $t + 1$. When $\gamma_i \neq \gamma$, this ratio is not the one expected at date t , and thus a new transition path must be computed to determine consumption at $t + 1$.

reallocated, the complete markets model now delivers small differences in returns to capital across countries, with slightly higher returns in fast-growing countries. The autarky model now generates smaller interest rate differentials across countries relative to the perfect foresight model. The logic is that fast-growing countries now (mistakenly) expect slower consumption growth, and thus a lower interest rate leaves them indifferent on the margin between consuming and investing.

The economy that looks most different relative to the perfect foresight specification is the bond economy. In the repeated surprise version of this economy, consumption growth broadly follows output growth. Thus, the allocation of consumption is no longer efficient. However, consumption growth exceeds output growth for slow-growing countries, whereas consumption growth is weaker than output growth for fast-growing countries, so some consumption smoothing is achieved. Slow-growing countries still accumulate large positive NFA positions relative to output, whereas fast-growing countries accumulate large deficits. The logic for these patterns is that at each date during transition, a slow-growing country sees current income turn out lower than expected and revises downward expected permanent income. Relative to actual income, wealth is higher than expected. But the slow-growing country does not want to reduce savings, because the representative agent is a permanent income consumer. Rather, the slow-growing country invests its excess wealth abroad, and the NFA position rises. Consumption growth for the slow-growing country is faster than output growth because as the NFA position rises, an ever-increasing share of consumption comes from interest income out of saving, and thus the consumption-to-output ratio rises.

Overall, with the repeated-surprise model for expectations, the bond economy offers a more reasonable account of the data. The key strength of the model is that it can deliver an equilibrium outcome in which realized long-run consumption growth rates differ across countries, while returns to capital are roughly equalized. The one dimension along which the model remains most at odds with the data is the dynamics of capital flows. For example, a country that grows (unexpectedly) 1 percentage point faster than the rest of the world for 50 years should end up with a negative net foreign asset position approaching 100% of GDP, whereas a country that grows 1 percentage point slower should end up with a positive position of around 150% of GDP. These numbers are very large relative to the actual variation in the NFA position across countries (see Figure 3). As in the perfect foresight version of the model, introducing a nontradable sector would imply smaller NFA positions.

An alternative or complementary explanation for the relatively small net foreign asset positions observed in the data is that countries differ with respect to preferences, and that preferences vary systematically with productivity growth such that fast-growing countries also tend to be more patient. We will explore this possibility in the next section.

3.4.5 Alternative Model 2: Preference Variation

Suppose that in the data we observe cross-country equality in marginal products of capital (as argued by Caselli and Feyrer, 2007) but cross-country variation in consumption growth. In the context of a model with asset trade, this can only be explained by country variation in preferences. Moreover, if one assumes that countries can trade a bond freely, then the first-order condition for bonds can be used to identify preference shocks from data on consumption and interest rates. For example, if $\frac{\phi_{i,t+1}}{\phi_{it}} = \frac{\beta_i}{\beta}$, so that countries differ with respect to their rates of time preference, then the constant expected consumption growth rate for country i will be given by

$$\frac{E[C_{i,t+1}]}{C_{it}} = (\beta_i R)^{\frac{1}{\sigma}}.$$

The identification of preference shocks is important since it affects the calculation of the welfare gains from alternative market structures and the assessment of whether allocations are efficient. To see this, consider the allocation for capital and consumption along the equilibrium path of the bond economy model under the baseline calibration in which $\phi_{it} = 1$ for all t . Denote this allocation $\{K_{it}^{BE}, C_{it}^{BE}\}$. One can then construct an alternative time-varying path $\{\tilde{\phi}_{it}\}$ such that given $\{\tilde{\phi}_{it}\}$, the equilibrium in the bond economy model features exactly the same path $\{K_{it}^{BE}\}$ but where the path for consumption $\{\tilde{C}_{it}\}$ is different and such that there is no bond trade. In particular, this path $\{\tilde{\phi}_{it}\}$ can be reverse engineered from the intertemporal first-order condition, by setting consumption equal to the difference between domestic output and domestic investment at each date t and computing the value for $\tilde{\phi}_{i,t+1}$ such that the first-order condition is satisfied:

$$\tilde{\phi}_{i,t+1} = \frac{\tilde{\phi}_{it} \tilde{C}_{it}^{-\sigma}}{\beta R \tilde{\phi}_{i,t+1} E_t [\tilde{C}_{i,t+1}^{-\sigma}]}$$

where

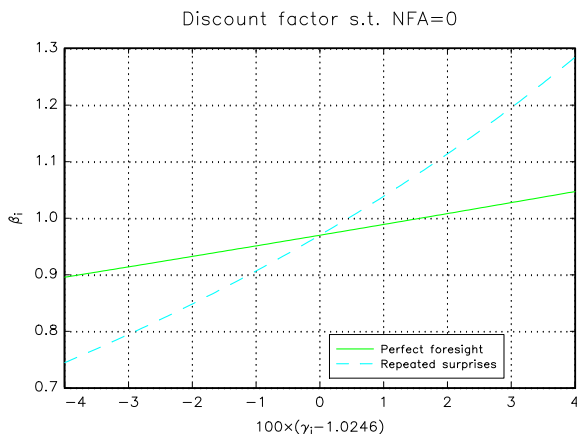
$$\tilde{C}_{it} = (K_{it}^{BE})^\alpha A_{it}^{1-\alpha} - K_{i,t+1}^{BE} + (1 - \delta)K_{it}^{BE}$$

and $\tilde{\phi}_{i0} = 1$. Note that because the time path for capital is identical to that in the original bond economy model, expected returns are equal to the world interest rate at each date. At this common world interest rate, given preferences described by $\{\tilde{\phi}_{it}\}$, agents have no incentives to trade bonds, and thus allocations with a bond market are identical to those under financial autarky.⁹ It follows that if preferences were truly described by $\{\tilde{\phi}_{it}\}$, then the welfare gains of moving from financial autarky to a bond economy market structure would be zero. Moreover, given perfect foresight, allocations under financial autarky would be efficient.

⁹Bond economy allocations are identical to those under financial autarky given the preference path $\{\tilde{\phi}_{it}\}$. They are not identical to allocations under financial autarky assuming $\phi_{it} = 1$ for all t .

What do the paths $\{\tilde{\phi}_{it}\}$ look like for the countries with the sorts of growth experiences in our Penn World Tables sample? It turns out that under both models for expectations, the required growth rate for $\tilde{\phi}_{it}$ is constant after the initial period. Thus, we can express preference differences across countries in terms of differences in the rate of time preference. Given perfect foresight, the mapping from γ_i to β_i such that the bond economy and autarky allocations coincide is defined by $\frac{\beta_i}{\gamma_i^\sigma} = \frac{\beta}{\gamma^\sigma} = R$. The mapping is more involved in the repeated surprises model. We plot both mappings in Figure 6.

Figure 6: Discount factors such that autarkic allocations are efficient



Note that much larger cross-country variation in the rate of time preference is required to generate an absence of asset trade in the repeated surprises model. The logic is that in that model, in each period a slow-growing country finds itself with too much capital relative to productivity. To be willing to immediately consume all this excess capital (rather than invest in the bond), the slow-growing country must be very impatient.

From a positive perspective, the models with preference heterogeneity seem to offer a reasonable approximation to the experiences of fast- versus slow-growing economies. Consumption growth tends to track output growth, and fast-growing countries are not net foreign savers. At the same time, interest rates are equated across countries. However, before concluding that cross-country heterogeneity in preferences accounts for cross-country heterogeneity in savings rates, it would be nice to see some independent evidence (besides the differences in savings rates) that preferences really do differ across countries.

In this light, one leading candidate explanation for observed “global imbalances” (see the chapter by Gourinchas and Rey 2013 in this volume) is that countries in which idiosyncratic household-level risk is larger and/or less well insured have a stronger precautionary motive to save. In practice, a stronger precautionary motive will deliver very similar aggregate predictions to a higher discount

factor. In particular, following international capital market liberalization, those countries with a stronger precautionary motive will reduce current consumption and lend to the rest of the world (see Mendoza, Quadrini, and Rios-Rull, 2009).

Although differences in idiosyncratic risk are an appealing potential justification for differences in patience it is not obvious why greater idiosyncratic risk should be systematically connected to faster growth. Some recent research develops models in which, because of domestic financial frictions, reforms which fuel growth simultaneously generate an increase in saving with the rest of the world. A common feature of those models is that the determinant of external saving is not the marginal rate of substitution (MRS) of the representative agent (i.e., aggregate consumption growth), but rather the MRS of a subset of agents in the country which has a stronger desire to save. This can be either because the savers face entrepreneurial risk (Sandri, 2010), because the savers are not the ones benefitting from faster growth (Song et al., 2011), or because fast growth implies rising wages and thus declining income for employers (Buera and Shin, 2010). It remains an open quantitative question whether these mechanisms can generate strong enough incentives for external saving to dominate the standard inter-temporal motives that link relatively fast growth to capital inflows in the simple models we have worked through in this section. An alternative way to connect patience and growth that does not rely on financial frictions is to posit that patience drives growth, because more patient countries are more likely to devote current resources towards investments that are conducive to long-run growth (for more on this see, Doepke and Zilibotti, 2011).

3.4.6 Comparing All Models

Table 1 summarizes various features of the data for our sample of countries and for the various models we have considered.

Column (1) simply reports the least squares slope coefficients plotted in Figures 1, 2 and 3. Columns (2) and (3) present results from splitting the sample into two equal length subperiods, one from 1960 to 1985, and a second from 1986 to 2010. Because of greater international financial integration in the second subperiod, one might expect to find more evidence of international risk sharing. Indeed it appears that country consumption growth has become less tightly linked to country output growth over time, suggesting movement toward a more efficient cross-country allocation of consumption. Comparing the empirical consumption regression coefficients with the ones for various models reported in columns (4)-(10), one interpretation of this reduced sensitivity is that autarky was a reasonable approximation to the world economy in the first half of the sample period, whereas in the second part the bond economy under the repeated surprises model for expectations looks like a reasonable candidate model. However, note that the allocation of consumption is not efficient under either of these models.

Table 1: Long-run growth patterns in the data and in theory

	DATA			MODELS						
				Perfect Foresight			Repeated Surprises			Alt.Prefs Foresight
Dep. Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	1960-2010	60-85	86-10	AUT	BE	CM	AUT	BE	CM	BE=AUT
(1) Cons. Growth	0.86 (0.03)	0.88 (0.036)	0.70 (0.038)	0.94 (0.00)	0.00*	0.00*	0.97 (0.00)	0.78 (0.01)	0.00*	1.00* (0.00)
(2) Inv. Growth	1.07 (0.11)	1.04 (0.13)	1.39 (0.12)	1.21 (0.00)	1.00*	1.00*	1.11 (0.00)	1.29* (0.00)	1.25* (0.00)	1.00* (0.00)
(3) Δ NFA/GDP	14.3 (4.03)	7.75 (2.81)	3.82 (2.79)	0.00	-865* (24.9)	-1196* (48.0)	0.00	-174 (10.2)	-947* (55.0)	0.00*

Notes: * next to a coefficient denotes that the allocation for the dependent variable is efficient. Each coefficient corresponds to the OLS coefficient on average annual growth in GDP per capita (standard errors are in parentheses). The dependent variables are: (1) growth in consumption pc, (2) growth in investment pc, and (3) the sample period change in the ratio of net foreign assets to GDP. Each regression includes a constant. We assume that $NFA/GDP_{1960}=0$. The sample of countries for regressions (1) and (2) is the set of 112 countries in the Penn World Tables with data over the entire 1960-2010 period. The sample for the regressions in row (3) is the set of 108 countries with NFA position data in the IMF International Financial Statistics.

In addition, capital flows in the wrong direction relative to all the theoretical models in columns (5), (6), (8), and (9). The models predict that fast-growing countries should be using international financial markets to fund high investment rates and (in the foresight models) to increase consumption in line with high expected future income. Instead, fast-growing countries have on average been exporting savings and increasing their net foreign asset positions over time. Comparing row (3) columns (2) and (3) indicates no evidence of a change in the direction of capital flows over time.

Since none of the models offers a compelling positive theory of capital flows, it is difficult to interpret what the capital flow evidence has to say about risk sharing. In addition, although it seems that we can reject consumption efficiency, recall that our simple models abstract from nontradables. Conceivably, if nontradables are very important, one might be able to reconcile the second-period consumption regression coefficient with efficiency. To bring some more evidence to bear on these issues, we now assess how the three indicators of efficiency developed above vary with two popular measures of openness. The first is the *de jure* index of capital market openness developed by Chinn and Ito (2008). The second is simple trade openness. In Table 2 we report regression coefficients analogous to those in Table 1 for four subgroups of our 1986-2010 sample: countries with Chinn-Ito index values above the sample median (Capital Open) and below the sample median (Capital Closed), and countries with trade (imports plus exports) over GDP above

Table 2: Efficiency and openness

Sample: 1986-2010				
	(1) Capital Open	(2) Capital Closed	(3) Trade Open	(4) Trade Closed
Dependent variable:				
(1) Cons. Growth	0.80 (0.06)	0.65 (0.05)	0.61 (0.06)	0.83 (0.05)
(2) Inv. Growth	0.95 (0.24)	1.55 (0.15)	1.56 (0.17)	1.16 (0.18)
(3) Δ NFA/GDP	14.8 (6.8)	0.68 (2.8)	6.17 (4.20)	-2.50 (3.61)

Note: Each number is the OLS coefficient on average annual growth in GDP per capita (standard errors in parentheses).

the sample median (Trade Open) and below the sample median (Trade Closed).

Surprisingly, countries with more open capital markets seem to enjoy less risk sharing according to all three measures. In particular, within the set of financially open countries, consumption and investment track output quite closely, and capital flows out of rather than into relatively fast-growing countries. One interpretation of this evidence is that actual international capital flows typically work against risk sharing and that financially closed economies achieve more efficient allocations than financially open ones (we will develop a simple model with this feature in Section 4.4.2). An alternative interpretation is that the model in which countries differ only with respect to productivity growth is the wrong model for understanding international capital flows. In the spirit of the example economy with preference heterogeneity and/or differential precautionary saving motives described above, perhaps fast-growing countries are effectively so much more patient than slow-growing countries that, given the chance, they would choose to lend to their slow-growing neighbors rather than borrow from them.

Greater trade openness translates into better risk sharing, according to the consumption risk-sharing indicator. One possible interpretation of this finding is that nontradables account for a larger fraction of output in countries that trade relatively little, leading to a stronger tendency for consumption growth to track income growth. If a country trades relatively little because it faces relatively high transportation costs, relatively strong comovement between output and consumption is efficient.

3.5 Welfare

In our numerical example, we can compare welfare across alternative market structures. Panel *F* in Figures 4 and 5 shows the permanent percentage increase in consumption under autarky that delivers equal welfare to the bond economy, conditional on the country growth rate being γ_i . Given perfect foresight about productivity growth, expected welfare *ex ante* and realized welfare *ex post* are the same. In the repeated surprise model, the two welfare measures are different. We focus on realized welfare.

In the perfect foresight economy, welfare gains are U-shaped and approximately symmetrical. If $\gamma_i = \gamma$, so that country productivity growth is identical to that in the world economy, then there are no welfare gains from being able to trade a bond: autarkic allocations are efficient. As the absolute difference between γ_i and γ increases, the welfare gains from bond trade increase and become very large. For countries growing 4 percentage points faster or slower than the world economy, the welfare gains reach 100% of consumption. One might wonder whether welfare gains would be larger or smaller in an endowment economy version of the model. In fact, in an endowment economy, one can solve for the welfare gain in closed form.¹⁰ Quantitatively, welfare gains turn out to be similar in the endowment and production economy versions of the model.

Now look at the corresponding plot for the repeated surprise economy (Panel *F*, Figure 5). Welfare gains from being able to trade a bond are smaller here. The reason is that fast-growing countries do not anticipate fast growth and therefore do not borrow so much early in transition. However, even in this case the welfare gains from trade in a bond are nontrivial. The reason is that bond trade gives access to foreign capital to fast-growing (high autarkic interest rate) countries that are net foreign borrowers, and gives higher returns on saving to slow-growing (low autarkic interest rate) countries that are net foreign lenders.

Suppose we were to extend the model to include a nontradable sector as discussed above. How would that change the welfare results? The answer, assuming tradables and nontradables enter preferences separably, is that the welfare gains of introducing bond trade will be exactly those plotted, except that now they should be interpreted as measuring the percentage increase in tradable consumption required to leave the agent indifferent between the two market structures (given the same nontradable consumption in both cases).

We now estimate the average welfare gain of moving from autarky to free bond trade in our panel of countries in the Penn World Tables. More precisely, we ask what is the permanent percentage increase in consumption under autarky that would leave an individual indifferent between being allocated a nationality at random in 1960, knowing that each country will remain permanently in autarky, versus experiencing the same lottery given free bond trade between countries at the world

¹⁰The expression for the welfare gain is $100 \times \left(\left(\frac{\gamma^\sigma - \beta\gamma}{\gamma^\sigma - \beta\gamma_i} \right) \left(\frac{1 - \beta\gamma_i^{1-\sigma}}{1 - \beta\gamma^{1-\sigma}} \right)^{\frac{1}{1-\sigma}} - 1 \right)$, where γ_i and γ now denote the growth rates of the endowment in country i and in the world economy.

Table 3: Welfare gains of moving from market structure A to B

Structure A	Structure B	Gain (% of cons.) Perfect foresight	Gain (% of cons.) Repeated surprises
<i>AUT</i>	<i>BOND</i>	9.6	4.8
<i>AUT</i>	<i>CM</i>	16.0	16.8
<i>BOND</i>	<i>AUT</i>	-8.8	-4.6
<i>BOND</i>	<i>CM</i>	5.9	11.4

risk-free rate?¹¹ We also compute the analogous welfare gain associated with moving from autarky to complete markets. We then redo both experiments assuming that the lottery is between the bond economy on the one hand versus autarky or complete markets on the other.¹² Table 3 shows the results, under both models for expectations.

The welfare gain from being able to trade a bond is equivalent to a 9.6% increase in consumption in financial autarky, assuming perfect foresight about future productivity growth. Given access to a bond, the additional gain from being able to insure against growth rate risk is worth 5.9% of consumption. Recall that bond trade in our example offers no insurance, in the sense that bond markets open only after γ_i is drawn, after which point each country's destiny is known. Rather, bond trade just allows countries to allocate capital and consumption more efficiently across countries. It is perhaps surprising that the welfare gains from achieving efficient intertemporal allocations within countries exceed the additional potential gains from being able to insure *ex ante* against the draw for γ_i and equating the level of consumption across countries. Note also that these welfare gains are extremely large relative to Lucas' (1987) estimates of the welfare costs of business cycles (0.008%). The reason they are so much larger is simply that here we are evaluating the cost of lack of insurance against different long run growth outcomes as opposed to the cost of lack of insurance against transitory business cycle fluctuations. In other words, we are looking at big shocks, whereas Lucas focused on small shocks.

How do these welfare numbers change when differences in growth rates come as a surprise? Now, as expected, the gains from bond trade are smaller, because trade in a bond is no longer sufficient to deliver an efficient intertemporal allocation of consumption. The counterpart to that

¹¹Note that this thought experiment abstracts from inequality in initial conditions in 1960, since we assume all countries share the same initial capital and productivity, irrespective of market structure. By assuming that agents assign the same probability to drawing China's growth rate as Iceland's, the calculation also abstracts from heterogeneity in country size.

¹²For the purposes of these calculations, we take the distribution of γ_i to be the set of annual output per capita growth rates in our Penn World Tables sample over the period 1960 to 2010, since country output and country productivity grow at the same rate in the bond economy given perfect foresight. Since aggregate consumption grows at the same rate in the (perfect foresight) bond and complete markets models, the welfare gains associated with moving from a bond economy to complete markets come solely from equalizing the distribution of consumption.

result is that the gains from being able to explicitly insure against growth rate risk are larger.

The welfare gains of moving to complete markets are of theoretical interest, but it is hard to imagine what sorts of markets or institutions might provide insurance against long-run growth rate risk. Suppose we take autarky under the repeated surprise model for expectations as the closest theoretical approximation to the actual global economy over the past 50 years. We would then conclude that an expected welfare gain worth 4.8% of consumption would be an upper bound for the potential welfare gains from countries having access to a globally integrated bond market over this period. It is an upper bound for three reasons. First, as discussed above, this gain should more properly be viewed as being expressed as a percentage of tradable consumption rather than total consumption. Second, realizing these gains would have required especially fast-growing countries to accumulate very large negative net foreign asset positions – positions that would presumably be difficult to sustain in practice given difficulties in enforcing repayment of international debts. Third, these welfare calculations were computed assuming no preference asymmetries across countries. In the example economy described in section 3.4.5 preferences differ systematically across countries in such a way that allocations and welfare are identical under the autarkic and bond economy market structures.

3.6 Summary

Our overall assessment of the evidence in Tables 2 and 3 and in Figures 1 through 5 is that the long-run allocations of consumption across countries are inefficient. With sufficient creativity, one can conjure up cross-country variation in preferences and technologies such that observed allocations are efficient, but in our view those models require implausibly large nontradable sectors, and an implausible pattern of covariation between growth rates and rates of time preference across countries. The fraction of output devoted to domestic consumption varies little across countries of very different income levels, which is another way of saying that countries that have experienced fast output growth have experienced similarly rapid consumption growth. The simplest and most compelling explanation for this fact in our view is that there is limited consumption insurance against long-run risk. Put differently, the long-run welfare of a country's citizens is much more tightly linked to the performance of their home country than to that of the world economy.

On the other hand, productive efficiency is harder to reject. Caselli and Feyrer (2007) have argued that marginal products of capital are roughly equated across countries, which suggests a high degree of international capital mobility. One way to reconcile small cross-country differences in returns on the one hand with large cross-country differences in consumption growth rates on the other is to postulate that differences in expected consumption growth rates across countries are small, even though differences in realized growth rates are large. The repeated surprises model for expectations we outlined has that feature, and that model offers a reasonable account of the data

in those dimensions. However, recall that although capital is allocated efficiently in that model, consumption is not.

Finally, we note once again that explaining the observed dynamics of capital flows remains an open challenge. All the models with asset trade we have considered predict a strong negative correlation between long-run growth and the net foreign asset position, while in practice cross-country NFA variation is modest (relative to the theory) and is positively (though not closely) related to output growth. As more satisfactory positive theories of global imbalances develop, we expect the question of long-run efficiency to be revisited.

4 Assessing Efficiency in International Business Cycles

The seminal contribution of Backus, Kehoe, and Kydland (1992) has started a very active research line that has tried to assess whether the international allocation of resources across developed economies over the business cycle is efficient. In this section, we use the methodological framework described above to organize and describe the main contributions of this literature, to summarize its main findings so far, and to suggest future research directions. To stay in close contact with the literature, the theoretical framework we use in this example is the two-country, two-good international business cycle model developed in Backus, Kehoe, and Kydland (1994). The key difference between this model and the one used in the previous section is that in this model countries produce different goods that are imperfect substitutes. In response to country-specific shocks, this imperfect substitutability will give rise to changes in relative prices. It is plausible that domestic and foreign goods are less substitutable in the short run than in the long run, when production processes and supply chains can be adjusted in response to changes in international relative prices (see, for example, Ruhl, 2008).

4.1 Preferences, Technologies and Frictions

The economy is composed of two countries, indexed $i = 1$ and $i = 2$, each populated by mass one of identical, infinitely lived households. In each period t , the economy experiences one event $s_t \in S$. We denote by s^t the history of events up to and including date t . The probability at date 0 of any particular history s^t is given by $\pi(s^t)$.

Each household derives utility from consumption, $c_i(s^t)$, and disutility from labor supply, $n_i(s^t)$. Preferences are given by

$$\sum_{t=0}^{\infty} \beta^t \sum_{s^t} \pi(s^t) U(c_i(s^t), n_i(s^t)), \quad (3)$$

where the parameter β captures the rate of time preference and the period utility function is $U(c_i, n_i) = \left(c_i^\mu (1 - n_i)^{1-\mu} \right)^{1-\gamma} / (1 - \gamma)$.

Capital in place $k_i(s^{t-1})$ (chosen in the previous period) and labor are combined to produce two country-specific intermediate goods. These are the only tradable goods in the economy. The intermediate good produced in country 1 is labeled a , and the good produced in country 2 is labeled b . The intermediate goods production functions are Cobb-Douglas:

$$F_i(z_i(s^t), k_i(s^{t-1}), n_i(s^t)) = \exp(z_i(s^t)) (k_i(s^{t-1}))^\theta (n_i(s^t))^{1-\theta}, \quad (4)$$

where $z_i(s^t)$ is an exogenous productivity shock that follows a symmetric autoregressive process:

$$\begin{aligned} \begin{bmatrix} z_1(s^t) \\ z_2(s^t) \end{bmatrix} &= \begin{pmatrix} \rho & \psi \\ \psi & \rho \end{pmatrix} \begin{bmatrix} z_1(s^{t-1}) \\ z_2(s^{t-1}) \end{bmatrix} + \begin{bmatrix} \varepsilon_1(s^t) \\ \varepsilon_2(s^t) \end{bmatrix} \\ \begin{bmatrix} \varepsilon_1(s^t) \\ \varepsilon_2(s^t) \end{bmatrix} &\sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma_\varepsilon^2 \begin{pmatrix} 1 & Corr_{\varepsilon_1, \varepsilon_2} \\ Corr_{\varepsilon_1, \varepsilon_2} & 1 \end{pmatrix}\right). \end{aligned}$$

Within each country, the intermediate goods a and b are combined to produce country-specific nontradable final consumption and investment goods according to the following constant returns to scale technology:

$$G_i(a_i(s^t), b_i(s^t)) = \begin{cases} \left[\omega a_i(s^t)^{\frac{\sigma-1}{\sigma}} + (1-\omega) b_i(s^t)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, & i = 1 \\ \left[(1-\omega) a_i(s^t)^{\frac{\sigma-1}{\sigma}} + \omega b_i(s^t)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, & i = 2, \end{cases} \quad (5)$$

where $a_i(s^t)$ and $b_i(s^t)$ denote the quantities of intermediate goods a and b used in country i as inputs, σ is the elasticity of substitution between domestic and foreign-produced inputs, and ω determines the extent to which there is a home or foreign bias in the composition of domestically produced final goods. This bias allows the model to replicate empirical measures for the volume of trade relative to GDP.

Investment augments the capital stock in the standard way:

$$k_i(s^t) = (1 - \delta)k_i(s^{t-1}) + x_i(s^t), \quad (6)$$

where δ is the depreciation rate and $x_i(s^t)$ is the amount of the final good devoted to investment in country i .

The resource constraints for this economy are

$$a_1(s^t) + a_2(s^t) = F(z_1(s^t), k_1(s^{t-1}), n_1(s^t)) \quad (7)$$

$$b_1(s^t) + b_2(s^t) = F(z_2(s^t), k_2(s^{t-1}), n_2(s^t)) \quad (8)$$

and

$$c_i(s^t) + x_i(s^t) = G_i(a_i(s^t), b_i(s^t)), \quad i = 1, 2. \quad (9)$$

We will consider two alternative measures of output in this economy. The first, following Backus et al. (1994), is the physical quantity of intermediate goods produced, which we denote $y_i(s^t) = F_i(z_i(s^t), k_i(s^{t-1}), n_i(s^t))$. The second values intermediate goods output in country i in units of i 's final consumption-investment good. We denote this alternative $y_1^c(s^t) = G_{a1}(s^t)y_1(s^t)$ and $y_2^c(s^t) = G_{b2}(s^t)y_2(s^t)$, where G_{ai} and G_{bi} denote the marginal products of intermediates a and b in country i 's final goods production. The value of country 1's net exports, in units of the domestic final good, is $nx_1(s^t) = G_{a1}(s^t)a_2(s^t) - G_{b1}(s^t)b_1(s^t)$. Note that when each component of GDP is measured in units of the final good, the national income accounting identity is preserved: $y_i^c(s^t) = c_i(s^t) + x_i(s^t) + nx_i(s^t)$.

4.2 Efficient Allocations

Efficient allocations in this framework are easily computed using a planning problem that maximizes a weighted average of the welfare of the representative agents in the two countries. Let κ denote the relative weight on country 1. The key first-order conditions defining efficiency are

$$\begin{aligned} \kappa U_{c1}(s^t)G_{a1}(s^t) &= (1 - \kappa)U_{c2}(s^t)G_{a2}(s^t) \\ \kappa U_{c1}(s^t)G_{b1}(s^t) &= (1 - \kappa)U_{c2}(s^t)G_{b2}(s^t) \end{aligned} \quad (10)$$

$$\begin{aligned} U_{c1}(s^t)G_{a1}(s^t)F_{n1}(s^t) &= -U_{n1}(s^t) \\ U_{c2}(s^t)G_{b2}(s^t)F_{n2}(s^t) &= -U_{n2}(s^t) \end{aligned} \quad (11)$$

$$U_{c1}(s^t) = \beta E_{s^t} [U_{c1}(s^{t+1}) [G_{a1}(s^{t+1})F_{k1}(s^{t+1}) + (1 - \delta)]] \quad (12)$$

$$U_{c2}(s^t) = \beta E_{s^t} [U_{c2}(s^{t+1}) [G_{b2}(s^{t+1})F_{k2}(s^{t+1}) + (1 - \delta)]] , \quad (13)$$

where U_{ci} , U_{ni} , F_{ki} , and F_{ni} denote, respectively, marginal utilities from consumption and hours and the marginal products of capital and labor in intermediate goods production.

The first pair of equations (eq. 10) defines an efficient division of tradable goods across countries. This is the generalization of the consumption-risk-sharing condition from the one-good model to a two-good world. It is efficient to divide good a such that the marginal value to the planner from putting an additional unit in either country is the same. One interpretation of these conditions is that the planner equates the marginal rate at which it is willing to substitute domestic for foreign consumption, $(1 - \kappa)U_{c2}(s^t)/\kappa U_{c1}(s^t)$, to the marginal rate at which it is able to transform domestic into foreign consumption by reallocating intermediates across countries, $G_{a1}(s^t)/G_{a2}(s^t) =$

$G_{b1}(s^t)/G_{b2}(s^t)$.

Since the relative weight κ is constant, the intertemporal marginal rate of substitution for intermediate goods is equated state-by-state across countries:

$$\frac{U_{c1}(s^{t+1})G_{a1}(s^{t+1})}{U_{c1}(s^t)G_{a1}(s^t)} = \frac{U_{c2}(s^{t+1})G_{a2}(s^{t+1})}{U_{c2}(s^t)G_{a2}(s^t)} = Q(s^t, s^{t+1}) \quad \forall s^t, s^{t+1}.$$

The second two pairs of efficiency conditions define efficient allocations of labor and capital. These conditions are standard. Note that the second intertemporal efficiency condition 13 can alternatively be written as

$$\frac{G_{b1}(s^t)}{G_{b2}(s^t)}U_{c1}(s^t) = \beta E_{s^t} \left[\frac{G_{b1}(s^{t+1})}{G_{b2}(s^{t+1})}U_{c1}(s^{t+1}) [G_{b2}(s^{t+1})F_{k2}(s^{t+1}) + (1 - \delta)] \right].$$

Comparing this equation to eq. 12 reveals an analogue to the result from the one-good model that efficiency dictates equating the marginal product of capital across countries (eq. 2). In this economy, capital in country 1 and capital in country 2 are different goods. Productive efficiency here requires equating the expected returns to investing in the two countries, given that the “prices” of capital in country 2 today and tomorrow relative to corresponding capital in country 1 are $G_{b1}(s^t)/G_{b2}(s^t)$ and $G_{b1}(s^{t+1})/G_{b2}(s^{t+1})$, respectively.

4.2.1 Efficiency in Endowment Economy Example

Before exploring this model’s predictions for productive efficiency, it will be useful to consider a simpler version of the model (Pakko, 1997) in which labor supply and capital are both fixed and equal to one, so that

$$y_i(s^t) = z_i(s^t) \tag{14}$$

and $U(c_i) = c_i^{(1-\gamma)}/(1 - \gamma)$.

In this simpler problem, the planner’s only choice is how to allocate intermediate endowments across countries, and an efficient consumption allocation is defined by eqs. 7, 10, and 14. Although this model is very simple, it clarifies the characterization of efficient consumption allocations in two-good models. The richer business cycle model will endogenize output, but the consumption efficiency condition (eq. 10) is the same, and thus much of the intuition that can be gleaned from the simpler model will carry over.

To warm up, consider a one-good model, or equivalently the special case of the model described in which $\omega = 1/2$ and $\sigma \rightarrow \infty$. In a one-good model, the marginal rate of transformation between domestic and foreign consumption is one, and thus efficiency simply dictates equating planner-weighted marginal utilities of consumption. Each country i receives a fixed fraction of the world endowment: the fraction for country 1 is $\kappa^{\frac{1}{\gamma}} / \left(\kappa^{\frac{1}{\gamma}} + (1 - \kappa)^{\frac{1}{\gamma}} \right)$. Thus, consumption comoves per-

fectly across countries, while the correlation of output is just dictated by the correlation of the exogenous productivity shocks. If country 1 has a relatively favorable productivity shock, it should increase exports to country 2. Thus, net exports should be procyclical.

These stark predictions are the starting point for a large fraction of the empirical work on international consumption risk sharing. As we will report below, in the data movements in consumption are typically less strongly positively correlated across countries than output, contrary to efficiency in the one-good model. In addition, net exports are typically counter- rather than procyclical, again in apparent contradiction to efficiency.

The characterization of efficient consumption sharing is less stark in the two-good model with σ finite. Suppose country 1 enjoys a positive productivity shock and therefore produces more of good a . The first difference with respect to the one-good model is that the productivity (shadow price) of good a will now fall relative to good b . Thus, output in country 1 valued in terms of the final good will increase by less than the increase in the endowment, and output in country 2 will rise, even though productivity there did not move. Thus the cross-country correlation between y_1^c and y_2^c will tend to be larger than the correlation between z_1 and z_2 . The second difference relative to the one-good model is that with two goods the planner faces a trade-off in deciding where to allocate the extra good a that is produced. On the one hand, the incentive to equalize consumption will push the planner toward exporting a good chunk of it abroad. On the other hand, imperfect substitutability between intermediate goods in producing the final consumption good, coupled with the bias in preferences toward the locally produced intermediate ($\omega > 0.5$), will push the planner toward devoting more of the extra good a to country 1.

We now show that for certain combinations of parameter values, the business cycle properties of efficient allocations in the two-good model differ sharply from those familiar from one-good models. In particular, efficient allocations can feature countercyclical net exports and a cross-country output correlation exceeding the cross-country consumption. A critical locus dividing the parameter space into regions in which business cycle properties differ qualitatively between the two- and one-good models is

$$\tilde{\sigma}(s, \gamma) = \frac{1}{\gamma} - \frac{(1 - \gamma)}{2s\gamma},$$

where $s = (1 - \omega)^{-\sigma} / ((1 - \omega)^{-\sigma} + \omega^{-\sigma})$ is the steady-state fraction of the domestic intermediate allocated to producing the domestic final good (in the case $\sigma = 1$, $s = \omega$).

PROPOSITION 1: Efficient allocations have the following properties if and only if $\sigma < \tilde{\sigma}(s, \gamma)$:

1. Pass-through from relative output to relative consumption is larger than one.
2. Net exports are countercyclical.
3. The cross-country output correlation exceeds the cross-country consumption correlation.

Proof: See the Appendix. The characterization in Proposition 1 is based on a log-linearization of eqs. 7, 10, and 14. We then solve in closed form for the solutions to this log-linear system. Note that Proposition 1 states results for output measured in units of the final good (y_i^c). Pakko (1997) explored the same model but measured output in units of the intermediate good (y_i). The analogous condition for properties 1 and 3 when output is defined that way is $\sigma < (2s - 1) / 2s\gamma$.

The three properties listed in the proposition are obviously closely interrelated, and they all run counter to the conventional wisdom about efficiency derived from one-good models. The intuition centers on the trade-off sketched above between minimizing fluctuations in consumption mix within a country versus minimizing fluctuations in total consumption across countries. Consider a concrete example, with country 1 producing apples and country 2 producing bananas. Suppose country 1 has a particularly good harvest. If apples and bananas are poor substitutes, and if residents in country 1 have a preference bias toward apples, concentrating fruit consumption in country 1 will be efficient. In that case, the relative value of country 1's consumption will increase even more than the value of their output (property 1), net exports will fall (property 2), and consumption will end up comoving less strongly than output across countries (property 3).

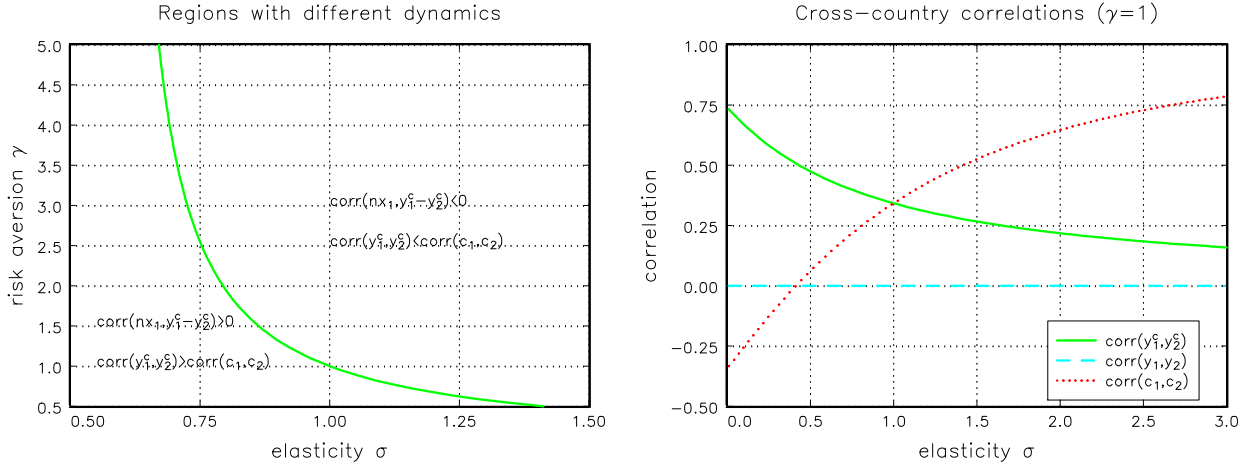
Panel A of Figure 7 plots the locus $\tilde{\sigma}(s, \gamma)$ below for $s = 0.85$, corresponding to an import share of 15%. The locus goes through the point $\sigma = \gamma = 1$, indicating that with (i) a unitary elasticity of substitution between domestic and foreign goods and (ii) logarithmic utility over consumption, net exports are always exactly zero, and the cross-country consumption correlation is identical to the cross-country output correlation. Higher risk aversion strengthens the planner's incentive to equalize consumption across countries, and thus a stronger incentive to maintain the steady-state mix of goods in consumption (a lower value for σ) is required to prevent the planner from wanting the more productive country to run a trade surplus.

A large part of the literature on international risk sharing investigates whether country-specific output growth helps predict country-specific consumption growth (see, for example, Lewis, 1996). If it does, that is taken as evidence against efficiency. However, property 1 indicates that even a large positive relationship between the two is not necessarily indicative of inefficiency.

With respect to property 2, Cole and Obstfeld (1991) were the first to emphasize that at certain parameter configurations, allocations are efficient absent any intertemporal borrowing and lending. We will revisit their paper when discussing evidence on efficiency from international portfolios.

Panel B of Figure 7 plots consumption and output correlations in this model as functions of σ , for $s = 0.85$ $\gamma = 1$ and uncorrelated endowment shocks. As $\sigma \rightarrow \infty$, so that the model collapses to a one-good model, the consumption correlation tends to 1, while the output correlation (in units of the final good) tends to 0.

Figure 7: Business cycle dynamics in a two-good endowment economy



4.3 Market Allocations

We now turn to the version of the economy with production. Here the literature has explored a variety of alternative market structures. Our baseline will be an economy where a full set of Arrow securities is traded internationally (complete markets). We will also consider economies where a limited set of assets is traded internationally: only stocks (as in Heathcote and Perri, 2013), only a bond (Arvanitis and Mikkola, 1996), or no assets at all (Heathcote and Perri, 2002).

In all market economies households rent labor to competitive intermediate goods-producing firms at wage $w_i(s^t)$ (measured in units of the final good). They also trade intermediate goods at prices $q_i^a(s^t)$, $q_i^b(s^t)$. Final goods-producing firms purchase the intermediate inputs and produce the final consumption/investment good, solving

$$\max_{a_i(s^t), b_i(s^t)} \left\{ G_i(a_i(s^t), b_i(s^t)) - q_i^a(s^t)a_i(s^t) - q_i^b(s^t)b_i(s^t) \right\}. \quad (15)$$

Intermediate goods-producing firms hold capital and make investment decisions. The intermediate goods firm's maximization problem in country i is to choose $k_i(s^t)$, $n_i(s^t)$ for all s^t and for all $t \geq 0$ to maximize

$$\sum_{t=0}^{\infty} \sum_{s^t} Q_i(s^t) d_i(s^t)$$

taking as given $k_i(s^{-1})$, where $Q_i(s^t)$ is the price the firm uses to value dividends at s^t relative to

consumption at date 0, and dividends (in units of the final good) are given by

$$d_1(s^t) = q_1^a(s^t)y_1(s^t) - w_1(s^t)n_1(s^t) - x_1(s^t) \quad (16)$$

$$d_2(s^t) = q_2^b(s^t)y_2(s^t) - w_2(s^t)n_2(s^t) - x_2(s^t). \quad (17)$$

The state-contingent consumption prices $Q_i(s^t)$ play a role in intermediate goods firms' state-contingent decisions regarding how to divide earnings between investment and dividend payments. We assume that firms use the discount factor of the representative local household to price the marginal cost of forgoing current dividends in favor of extra investment:

$$Q_i(s^t) = \frac{\pi(s^t)\beta^t U_{ci}(s^t)}{U_{ci}(s^0)}. \quad (18)$$

We now describe how the representative households' budget constraints differ across the different market structures.

Complete markets

Without loss of generality, we can assume that a complete set of Arrow securities is denominated in units of good a . Let $B_i(s^t, s_{t+1})$ be the quantity of the security purchased by households in country i after history s^t that pays one unit of good a in period $t+1$ if and only if the state of the economy is s_{t+1} . Let $Q(s^t, s_{t+1})$ be the price in units of good a of this security. The budget constraint for the representative household in country i is

$$c_i(s^t) + q_i^a(s^t) \sum_{s_{t+1}} Q(s^t, s_{t+1}) B_i(s^t, s_{t+1}) = w_i(s^t)n_i(s^t) + d_i(s^t) + q_i^a(s^t) B_i(s^{t-1}, s_t). \quad (19)$$

Stock economy

In this economy, agents trade internationally equity of the intermediate goods-producing firms. Let $P^a(s^t)$ and $P^b(s^t)$ denote the price of shares in the representative firms in countries 1 and 2, in units of those countries' respective consumptions. Let $\lambda_i^a(s^t)$ and $\lambda_i^b(s^t)$ denote the shares of country 1 and 2 stocks purchased by agents in country i . The budget constraint for the representative household in country 1 (country 2 is analogous) is

$$\begin{aligned} & c_1(s^t) + P^a(s^t)\lambda_1^a(s^t) + e(s^t)P^b(s^t)\lambda_1^b(s^t) \\ & = w_1(s^t)n_1(s^t) + \lambda_1^a(s^{t-1}) [P^a(s^t) + d_1(s^t)] + \lambda_1^b(s^{t-1})e(s^t) [P^b(s^t) + d_2(s^t)], \end{aligned} \quad (20)$$

where $e(s^t)$ is the real exchange rate.

Bond economy

In this model, only a single noncontingent bond is traded. Let $B_i(s^t)$ denote the quantity and

$P(s^t)$ the price (in units of good a) of bonds bought by households in country i after history s^t . The bond pays one unit of good a in period $t + 1$ irrespective of the state in $t + 1$. The budget constraint for the representative household in country i is

$$c_i(s^t) + q_i^a(s^t)P(s^t)B_i(s^t) = w_i(s^t)n_i(s^t) + d_i(s^t) + q_i^a(s^t)B_i(s^{t-1}). \quad (21)$$

Financial autarky

In the financial autarky model, no assets are traded internationally; hence, the budget constraint for the representative household in country i is given by

$$c_i(s^t) = w_i(s^t)n_i(s^t) + d_i(s^t). \quad (22)$$

Households' problems

Households choose $c_i(s^t) \geq 0$, $n_i(s^t) \in [0, 1]$ and asset purchases (if assets are traded) for all s^t and for all $t \geq 0$ to maximize 3 subject to the appropriate sequence of budget constraints given by eq. 19, 20, 21, or 22, taking as given initial productivity shocks, initial capital stocks and, if assets are traded internationally, the initial distribution of wealth.

Definition of equilibrium

An equilibrium is a set of prices for all s^t and for all $t \geq 0$ such that when households solve their problems taking these prices as given, all markets clear. The goods market-clearing conditions are 7 and 9. The asset market conditions for the complete markets, stock, and bond economies are, respectively,

$$B_1(s^t, s_{t+1}) + B_2(s^t, s_{t+1}) = 0, \quad \forall s_{t+1} \in S. \quad (23)$$

$$\begin{aligned} \lambda_1^a(s^t) + \lambda_2^a(s^t) &= 1 \\ \lambda_1^b(s^t) + \lambda_2^b(s^t) &= 1. \end{aligned}$$

$$B_1(s^t) + B_2(s^t) = 0 \quad (24)$$

4.4 Comparing Models and Data

We now compare allocations (efficient and market) in the setup described above with the data in order to assess international efficiency over the business cycle. In the following four subsections, we explore comparisons for four different sorts of observables that have been used in the literature: (1) standard macroeconomic quantities, (2) the real exchange rate, (3) international portfolio

diversification, and (4) asset prices.

4.4.1 Assessing Efficiency Using Quantities

Our first comparison is based on international comovement of macro quantities and international flows of resources (i.e., net exports) at business cycle frequencies. In order to obtain predictions for business cycle dynamics in the various setups described above, we first calibrate the parameters of the model and the productivity process. We then numerically solve the models using standard linearization techniques. Finally, we simulate the models to compute statistics that can be compared with data. The first row of Table 4 reports our main statistics for quantities: international correlations of macro aggregates between the United States and the G6 (an aggregate of Canada, France, Germany, Italy, Japan and the UK), the standard deviation of GDP (as a measure of business cycle risk), and the standard deviation and cyclical of net exports (as a fraction of GDP).¹³ Following most of the literature since Backus et al. (1994), we focus on output measured in units of intermediate goods, $y_i(s^t)$.

The key features of the data can be summarized as follows:

1. Output, investment, and employment co-move positively and strongly across countries.
2. The cross-country correlation of consumption is positive but smaller than the correlation of output.
3. Net exports are not very volatile (their standard deviation is about one-third that of GDP) and are strongly countercyclical.

These features have been documented in many business cycle studies, and they are typical of many developed countries in different periods of time (see, for example, Backus and Kehoe, 1992). The first question we ask is whether a reasonably calibrated version of the model described above can generate efficient allocations with these features. The first to address this question were Backus, Kehoe, and Kydland (1992 and 1994, hereafter BKK). They concluded that features (1) and (2) of the data appear inconsistent with efficiency, naming this inconsistency the “quantity anomaly.” Here we first revisit the anomaly using the same calibration as in BKK, reported in Table 5.¹⁴ The second row of Table 4 reports the statistics for efficient allocations in the complete markets BKK economy. The cross-country correlation of consumption (0.94) exceeds the corresponding correlation of output (0.55), contrary to the data. This discrepancy was also present in the endowment version of

¹³We focus here on correlations and standard deviations at business cycle frequencies. Statistics at other frequencies might also be informative about risk sharing. See, for example, Baxter 2011, Pakko, 2004, and Rabanal and Rubio-Ramirez, 2010.

¹⁴The only difference between our calibration and BKK’s is that we set the standard deviation and correlation of the productivity innovations to match the standard deviation of US GDP and the correlation between US GDP and G6 GDP in our sample.

the model with relatively high elasticities of substitution between goods ($\sigma = 1.5$ in the BKK calibration).¹⁵

The version of the economy with production introduces additional implications for investment and employment. As long as goods are not too complementary, it is efficient to increase labor supply and to invest more in the country where productivity is relatively high, that is, to “make hay where the sun shines.” This implies that efficient allocations feature low cross-country correlations of investment (-0.07) and employment (-0.05), again at odds with the data. One dimension along which the production economy does much better than the endowment model is the dynamics of net exports, which are now countercyclical in line with data. Efficient allocations predict countercyclical net exports because of the incentive to invest in good times, which makes domestic absorption more procyclical (relative to the endowment case) and hence net exports more countercyclical. Overall, though, a comparison of lines (1) and (2) in Table 4 seems to point strongly against efficiency. Before we consider alternative market structures, however, we will show that alternative reasonable parameterizations make the efficient complete markets model allocation consistent with the three quantity facts described earlier.

Three important elements of the BKK calibration account for the large differences between the quantity dynamics observed in the data versus those predicted by the model.

The first element is the estimation of the productivity process, which includes a positive spillover term ψ (see Table 5). Because a positive productivity shock in one country signals high future productivity in the other, it is efficient for both countries to increase current consumption. Thus, consumption comovement ends up exceeding output comovement. Empirical work (see, for example, Baxter and Crucini, 1995, or Heathcote and Perri, 2004) has shown that precisely estimating spillovers is difficult, and that estimates are sensitive to the details of whether and how productivity series are detrended prior to estimating the transmission coefficients ρ and ψ .

The second element is the nonseparability between consumption and leisure in the utility function used by BKK. This feature makes it difficult to resolve the quantity anomaly, because it ties together comovement in labor and consumption. In their calibration, if both countries work more in response to a shock, the marginal utility of consumption in both countries will rise, which is a force toward equalizing consumption.

The third element is the relatively high value for the elasticity of substitution between foreign and domestic goods assumed by BKK which, as discussed in subsection 4.2.1, implies strong cross-country comovement in consumption as a feature of efficient allocations.

In lines (3), (4), and (5) we modify the BKK parameterization by changing these elements one at a time. In line (3) we consider a process for productivity with no spillovers ($\psi = 0$), as estimated by Heathcote and Perri (2004). In line (4) we consider log-separable preferences ($\gamma = 1$). In line (5) we consider a low elasticity of substitution ($\sigma = 0.6$). Comparing lines (3), (4), and (5) with

¹⁵Defining output in units of the consumption good slightly increases the model output correlation, to 0.60.

the data (1), we see that each change in the parameterization moves the model closer to the data, but none of the changes alone can solve the quantity anomaly. In line (6) we introduce all three changes simultaneously. Comparing lines (1) and (6) shows that a reasonable calibration of the complete markets model generates fluctuations in quantities that are very similar to those observed in the data.¹⁶

The next question is whether these fluctuations could also be replicated by a (nonefficient) market economy with limited asset trade. In particular, Baxter and Crucini (1995) argue that a bond economy (as described in subsection 4.3) can generate international comovement of consumption that is lower than the correlation of GDP. In line (7) of Table 4, we reproduce the Baxter and Crucini (BC) finding, using a special case of our general setup. In particular, we consider the BC process for productivity (unit roots with no spillovers), and, since theirs is a one-good setup, we set elasticity of substitution between goods to a high value ($\sigma = 5$). Line (7) shows that this model does indeed generate a consumption correlation lower than the output correlation. The economic mechanism through which the model delivers this is completely different from the one discussed in the two-good model. It does not hinge on imperfect substitutability of goods but rather on changes in the international interest rate.

Suppose country 1 experiences a positive productivity shock. Its demand for international loans increases strongly for two reasons. The first is that because markets are incomplete, risks are not shared and residents in country 1 are the sole beneficiaries of the increase in productivity. The second is that shocks are permanent, so residents in country 1 want to increase both consumption and investment. The strong demand for funds on the international market causes an increase in the world interest rate, which in turn induces residents in country 2 to supply more labor, to save more in international bonds and less in domestic capital, and to consume less. The result is an international consumption correlation that is below the output correlation, but also counterfactually negative international comovement in investment (-0.39 in the model versus 0.51 in the data).

To summarize, many authors have interpreted the quantity anomaly as evidence against international efficiency and have considered variants of the one-good international business cycle model in which observed international comovement can be explained by frictions that preclude efficient allocations across countries. Here we have shown that international comovement of quantities is perhaps better captured by assuming complete financial markets in the context of a two-good model with low elasticity of substitution.¹⁷ Thus, the international comovement of quantities is not necessarily inconsistent with international efficiency.¹⁸

¹⁶In this calibration, the cross-country correlation of output measured in units of consumption is 0.63.

¹⁷See also Viani, 2011, for a conclusion in this spirit.

¹⁸Fitzgerald (2012) uses a different model and a different quantity-based moment to assess efficiency, but comes to a similar conclusion. Her framework is a multicountry model, in which each country produces a country-specific intermediate good. All countries have identical symmetric preferences over a composite of all country varieties, but country-pair-specific transportation costs generate differences in final consumption prices P_i across countries. Fitzgerald notes that the quantity of imports into country i from country k (relative to the size of the two economies)

Table 4: Assessing efficiency using quantities

	International correlations				Domestic statistics		
	(y_1, y_2)	(c_1, c_2)	(x_1, x_2)	(n_1, n_2)	% sd y	% sd $\frac{nx}{y}$	corr($\frac{nx}{y}, y$)
1. Data	0.55	0.31	0.51	0.57	1.54	0.44	-0.51
Complete markets models							
2. BKK (see Table 5)	0.55	0.93	-0.07	-0.01	1.54	0.23	-0.43
3. No spillovers: $\rho = 0.91, \psi = 0$	0.55	0.71	0.35	0.56	1.54	0.19	-0.40
4. Separable utility: $\gamma = 1$	0.55	0.94	0.02	0.15	1.54	0.23	-0.43
5. Low elasticity: $\sigma = 0.6$	0.55	0.88	-0.08	0.10	1.54	0.28	-0.47
6. All: $\rho = 0.91, \psi = 0, \gamma = 1, \sigma = 0.6$	0.55	0.35	0.39	0.71	1.54	0.47	-0.46
Bond economy model							
7. BC: $\rho = 1, \psi = 0, \sigma = 5$	0.55	0.29	-0.39	0.92	1.54	0.82	-0.39

Notes: All data are from the OECD Quarterly National Accounts (GDP and components) and Main Economic Indicators (employment). The sample for the data statistics is 1960.1–2012.2. The variable y denotes real GDP, c denotes real consumption (both private and public), n denotes civilian employment, x denotes real gross fixed capital formation, nx/y denotes net exports over GDP (all nominal). All variables except net exports are in logs. All variables are HP filtered with a smoothing parameter of 1600. Statistics from the model are produced by simulating the model for the same numbers of periods as the the data and taking averages over 20 simulations. In lines 2 through 7 the standard deviation and correlation of shock innovations are calibrated to replicate the standard deviation of output and the international correlation of GDP.

Table 5: Baseline parameter values (from Backus et al., 1994)

Preferences	
Discount factor	$\beta = 0.99$
Weight on consumption	$\mu = 0.34$
Curvature	$\gamma = 2$
Technology	
Capital's share	$\theta = 0.36$
Depreciation rate	$\delta = 0.025$
Elasticity of substitution	$\sigma = 1.5$
Import share	$1 - s = 0.15$
Productivity process	
Persistence and spillover	$\rho = 0.906$ $\psi = 0.088$
Variance and correlation	$\sigma_\varepsilon^2 = 0.0097$ $\eta = 0.65$

4.4.2 Assessing Efficiency Using Real Exchange Rates

In the setup described above a direct implication of efficiency is that the ratio of marginal utilities across countries should be proportional to the rate at which foreign consumption can be transformed into domestic consumption by reallocating intermediate goods (see eq. 10). In our decentralized economies the marginal products of intermediate goods are their respective prices (relative to final goods), so that the ratio of the marginal products across countries ($G_{a1}(s^t)/G_{a2}(s^t) = G_{b1}(s^t)/G_{b2}(s^t)$) is the price of foreign consumption relative to domestic consumption, i.e., the real exchange rate, hereafter denoted $e(s^t)$. Thus, efficiency implies perfect comovement between the ratio of marginal utility in country 2 to marginal utility in country 1 and the real exchange rate. Intuitively, an increase in marginal utility in, say, country 2 relative to country 1 is compatible with efficiency only if resources consumed by country 2 become more expensive relative to those consumed by country 1. If one assumes that the utility function is separable between consumption and leisure, with exponent $(1 - \gamma)$ on the consumption component, this implies a perfect linear relationship (and hence a correlation of 1) between the ratio of domestic to foreign log consumption and the log real exchange rate:

$$\log \frac{1 - \kappa}{\kappa} + \gamma \log \left(\frac{c_1(s^t)}{c_2(s^t)} \right) = \log e(s^t). \quad (25)$$

In an influential paper, Backus and Smith (1993) show that for various pairs of developed countries, this correlation is actually close to zero or even negative. Moreover, equation (25) implies a relationship between the volatilities of relative log consumption and the log real exchange rate: with logarithmic preferences ($\gamma = 1$) the two should be equally volatile. In the data, however, the real exchange rate is typically much more volatile than relative consumption. In the first row of Table 6, we report the standard deviations of the real exchange rate and relative consumption for the United States versus the G6, as well as the correlation between the two variables. In the second row, we report the predictions for these variables in the BKK model assuming complete markets (we use the parameterization of the model in line 6 of Table 4 that resolves the quantity anomaly). Comparing lines 1 and 2 suggests a sharp rejection of efficiency, because efficient model allocations feature a real exchange rate that is both not very volatile and at the same time perfectly correlated with relative consumption.

Can alternative market structures with limited scope for international asset trade account for the observed features of the real exchange rate? If they can, that would be evidence against efficiency. If they cannot, we would instead conclude that the BKK framework does not offer a

should be systematically linked to P_i , to a “multilateral resistance” term for country k (capturing its distance from other potential trading partners) and to trade costs between i and k . With perfect risk sharing, P_i can be mapped directly into consumption for country i . She tests this specification for the import equation against an alternative in which P_i is treated as a time dummy, and finds that for developed countries the risk-sharing specification is not rejected.

satisfactory theory of real exchange rates and that its implications for exchange rates should not be used to assess efficiency. Before we address this question (in the next subsection), a useful step is to first relate the exchange rate e (price of foreign consumption relative to domestic) to the terms of trade p (price of foreign intermediate b relative to domestic intermediate a). Since consumption in both countries is a mix of domestic and foreign intermediates for which the law of one price holds, movements in the real exchange rate are mechanically related to changes in the terms of trade. Indeed, a log-linear approximation gives

$$\hat{e} = (2s - 1)\hat{p}, \tag{26}$$

where \hat{e} and \hat{p} denote log deviations from the steady state. Note that when both countries consume the same bundle of intermediate goods ($s = 0.5$), the real exchange rate is constant. If each country consumes only its own intermediate goods ($s = 1$), the real exchange rate and the terms of trade are the same variable. Also, as long as the trade share is less than 50%, ($s > 1/2$), the model predicts that the real exchange rate and the terms of trade should move together.

Are real exchange rates informative about efficiency? The literature to date disagrees on the answer to this question. Chari, Kehoe, and McGrattan (2002) have argued that when allocations are not efficient, the BKK setup can generate a more volatile exchange rate but cannot account for the low correlation between the exchange rate and relative consumption (even considering variants with nominal rigidities). To see why incomplete markets can generate more volatility, consider a positive productivity shock in country 1 that increases the supply of good a and thus pushes up the terms of trade and (by 26) depreciates the real exchange rate. As discussed in Section 4.2.1, when goods are imperfect substitutes, efficiency calls for consumption of country 1, which is intensive in good a , to go up. This increase in “demand” for good a mitigates its fall in price because of higher relative productivity, reducing the size of the real exchange rate depreciation. When markets are incomplete, the increase in consumption in country 1 is smaller, the demand effect is weaker, and the exchange rate depreciates more, implying more volatility. Notice, though, that the Backus-Smith puzzle is not really solved: lines 2, 3, and 4 of Table 6 show that going from efficient to incomplete markets (inefficient) allocations increases exchange rate volatility, but the real exchange rate remains strongly positively correlated with relative consumption.

In contrast, Corsetti, Dedola, and Leduc (2008) have argued that inefficiency can, in two different parameterizations of this setup, solve the Backus-Smith puzzle. The first case is a very low elasticity of substitution between intermediates. In this case, in response to a positive productivity shock in country 1, the relative price of good a must actually increase, implying a real exchange rate appreciation. This perverse dynamic arises because with a very low elasticity, the only way to generate additional demand for good a is to increase its relative price and thereby the income

and purchasing power of residents of country 1, who are the main customers for good a . This is the case considered in lines 5 and 6 of Table 6, where we set the elasticity of substitution to 0.38. Notice that in this case the inefficient allocations in the bond economy generate both a fairly volatile real exchange rate and a negative correlation between the real exchange rate and relative consumption.¹⁹

The second case considered by Corsetti et al. (2008) is one of a very high elasticity of substitution and very persistent shocks. The intuition here is similar to the one discussed in the previous section, describing the Baxter and Crucini (1995) model. When markets are incomplete, a very persistent shock in country 1 makes residents want to increase both consumption and investment. Demand for good a is so strong that the exchange rate appreciates.²⁰ In lines 7 and 8 of Table 6 we contrast efficient and inefficient allocations in this case. Note that although the bond economy generates exchange rates that are negatively correlated with consumption, the exchange rate remains much less volatile than in the data.

We conclude that within this framework it is possible to account for real exchange rate fluctuations, but doing so requires a very low – and arguably implausible – elasticity of substitution between imported and domestically produced goods. It is also important to note that this setup (along with many variants, including versions with nontradable goods) has the feature that real exchange rates are driven by fundamentals. Unfortunately, however, a solid link between exchange rates and fundamentals has not yet been established, for two related reasons.

First, at least since Mussa (1986), it has been well known that real exchange rates are much more volatile between pairs of countries with flexible exchange rates as compared with pairs of countries with fixed nominal exchange rates or a common currency. This suggests that a satisfactory theory of real exchange rate dynamics requires a theory of nominal exchange rates, and that even if the BKK model can be parameterized to generate realistic volatility, the fact that it is purely a real model suggests that it might not do so for the right reasons. From the perspective of using real exchange moments to assess efficiency, Hadzi-Vaskov (2008), Hess and Shin (2010), and Devereux and Hnatkovska (2011), report a particularly telling fact: the correlation between exchange rates and relative consumption is negative only for country pairs with flexible exchange rates, whereas for countries or regions sharing a fixed exchange rate, the correlation is mostly positive. Once again, it appears that a theory of nominal exchange rates is needed to make inferences about risk sharing. However, a satisfactory theory of nominal exchanges rates remains work in progress. Engel (2013) concludes his chapter in this handbook as follows: “Although this survey has suggested many different models, it is questionable that the models allow us to explain, even after the fact, the movements in major currency rates.”

¹⁹Although this setup can account for the Backus-Smith puzzle, it can do so only for a very narrow range of elasticities of substitution.

²⁰Another way to obtain this mechanism would be to introduce trend shocks as in Aguiar and Gopinath, 2007.

Table 6: Assessing efficiency using real exchange rates

	% sd e	% sd $\frac{c_1}{c_2}$	corr($\frac{c_1}{c_2}, e$)
1. Data	6.39	0.97	-0.21
Baseline parameters: $\rho = 0.91, \psi = 0, \gamma = 1, \sigma = 0.6$			
2. Efficient allocations	0.47	0.47	1
3. Bond Economy	0.73	0.36	0.99
4. Financial Autarky	3.15	0.02	0.79
Very low elasticity: $\rho = 0.91, \psi = 0, \gamma = 1, \sigma = 0.38$			
5. Efficient allocations	0.54	0.54	1
6. Bond Economy	2.88	0.15	-0.17
High elasticity and pers. shocks: $\rho = 1, \psi = 0, \gamma = 1, \sigma = 5$			
7. Efficient allocations	0.14	0.14	1
8. Bond Economy	0.23	1.28	-0.69

Notes: Real exchange rate data and relative consumption refer to US v/s G6. Real exchange rate between the US and the G6 is computed as the ratio between a weighted average of consumption deflators (all converted into US dollars) in the G6 countries and the US consumption deflator. Weights are proportional to GDP over the sample. Consumption and consumption deflators are from the OECD Quarterly National Accounts while nominal exchange rates are from the IMF International Financial Statistics. The sample for the data statistics is 1960.1–2012.2. In each parameterization, the standard deviation and correlation of innovations of productivity shocks are set so that the model reproduces the standard deviation of GDP in the US and international correlation of GDP between the US and the G6.

A second challenge when trying to connect the exchange rate to fundamentals is that empirical evidence suggests that a large share of real exchange rate movements in the data reflect changes in the relative price of traded goods (Engel, 1999) and that in accounting for changes in the relative price of traded goods, deviations from the law of one price and pricing to market play an important role (see Burstein and Gopinath, 2013, in this volume). The efficient ratio of foreign to domestic consumption should respond to changes in the real exchange rate that reflect changes in fundamental preferences or technologies, but not to changes that simply reflect changes in cross-country price differentials for the same goods. Berka et al. (2012) sketch a model in which a large share of real exchange rate movements reflect nominal shocks that move the nominal exchange rate and which are not offset by changes in relative pricing thanks to assumptions of infrequent price adjustment and local currency pricing. Thus, the real exchange rate moves in line with the nominal rate. These real exchange rate movements lead to fluctuations in real allocations that are inefficient.

This discussion highlights that evidence on efficiency from exchange rates (or asset price data) is indirect. Allocations are efficient if they satisfy a planner's problem and the planner's problem does not have any prices in it. Exchange rate data are informative about international efficiency only to the extent that exchange rates are informative about preferences or technologies (fundamentals). To make this point as sharply as possible, in the next subsection we compare and contrast two simple alternative exchange rate models. In one of these models, exchange rate movements are informative about changes in the underlying technology, and efficient allocations feature a systematic relationship between the exchange rate and relative consumption. In the second model, exchange rate movements are disconnected from fundamentals, in the spirit of Berka et al., and the efficient allocation of consumption is in turn disconnected from the exchange rate.

As work on decomposing the fundamental sources of exchange rate movements progresses, it should become easier to disentangle the fundamental-driven component of exchange rate movements from exchange rate movements that do not reflect changes in preferences and technologies. In the meantime, we put little weight on moments involving the real exchange rate, such as the Backus-Smith correlation, when assessing international efficiency.

Exchange rates and efficiency: a simple example Consider the following static (repeated) two-country, two-good economy. The representative agent in country 1 receives one unit of good a , while the representative agent in country 2 receives one unit of good b . Goods a and b are freely tradable. Agents in country 1 derive utility from consuming good a , while those in country 2 derive utility from consuming good b . The period utility function is $u_i(c_i) = c_i^{(1-\gamma)}/(1-\gamma)$.

The only source of uncertainty is the real exchange rate e , defined as the price of a unit of good b in units of good a . We will consider two different theories for the exchange rate.

In the first model, a linear technology can convert 1 unit of good b into e units of good a , or 1 unit of good a into $1/e$ units of good b . Thus, in this model e is the stochastic production price of

b in units of a .

In the second model, the technology is constant: goods a and b can always be transformed one into the other at a rate of one-for-one. Exchange rate fluctuations are driven by non-fundamental stochastic tariffs. Country 1 imposes a stochastic import tax / subsidy $\tau^{im} = 1 - e$ on imports of good b and an export tax / subsidy of $\tau^{ex} = 1 - \frac{1}{e}$ on exports of good a . Thus, to receive 1 unit of b in 1 (which can be transformed into 1 unit of a and consumed) requires buying $(1 - \tau^{im})^{-1} = e^{-1}$ units of good b in 2, while receiving 1 unit of a in 2 requires buying $(1 - \tau^{ex})^{-1} = e$ units of a in 1. From the perspective of agents, who care only about after-tax prices, in both countries the price of b relative to a is e , just as in the first model. Revenues from these taxes are rebated lump-sum to residents of country 2.

Let $e \in E = \{e_1, e_2, \dots, e_N\}$ and let $\pi(e_j)$ denote the probability of drawing e_j . We assume the same support and probability distribution for e across both models for the exchange rate.

We consider two alternative financial market structures. The first is financial autarky, under which agents can only barter good a in exchange for good b . Note, though, that there will be no barter in equilibrium, since after the exchange rate e is realized, there are no gains from trade.

The second market structure is complete markets. We assume that each period, before the shock e is realized, agents can trade state-contingent contracts that deliver a unit of good a in country 1 in a particular state e_j . Let $b_i(e_j)$ denote the number of units of the contract purchased by the representative agent in country i contingent on the exchange rate being e_j , and let $q(e_j)$ be the corresponding price. Market clearing requires $b_1(e_j) + b_2(e_j) = 0$ for all $e_j \in E$.

We will compare allocations and welfare in autarky and under complete markets, and contrast the welfare gains from financial integration under the two alternative models for the exchange rate.

Under financial autarky, allocations are simply given by $c_i(e_j) = 1$, for $i = 1, 2$ and $\forall e_j \in E$. Since relative consumption is constant in autarky, there is zero correlation between relative consumption and the real exchange rate, and thus the autarkic version of this model will replicate the low correlation between the real exchange rate and relative consumption that Backus and Smith document in the data.

Under complete markets, the problems agents face are formally similar across the two alternative exchange rate models. In both cases, the respective problems for agents in countries 1 and 2 are

$$\begin{aligned} & \max_{\{b_1(e_j)\}} \sum_{e_j \in E} \pi(e_j) u(c_1(e_j)) \\ & \text{s.t.} \quad \sum_{e_j \in E} q(e_j) b_1(e_j) = 0 \\ & \quad c_1(e_j) = b_1(e_j) + 1 \end{aligned}$$

and

$$\begin{aligned} & \max_{\{b_2(e_j)\}} \sum_{e_j \in E} \pi(e_j) u(c_2(e_j)) \\ & \text{s.t.} \quad \sum_{e_j \in E} q(e_j) b_2(e_j) = 0 \\ & \quad c_2(e_j) = \frac{b_2(e_j)}{e_j} + 1 + T(e_j) \end{aligned}$$

Note that if agents in country 2 have bought $b_2(e_j)$ units of a claim to good a in state e_j and that state is realized, they will be able to exchange the payoff into an additional $b_2(e_j)/e_j$ units of good b , which is the good they consume. $T(e_j)$ denotes the lump-sum rebate of tax revenue. This term is zero in the first exchange rate model.

By taking first-order conditions to these problems, it is straightforward to show that the equilibrium ratio of consumption across countries has the familiar form from complete markets models:

$$\left(\frac{c_1(e_j)}{c_2(e_j)}\right)^\gamma = e_j \quad \forall e_j \in E. \quad (27)$$

Thus, agents use financial markets to divide aggregate world resources in the way that seems optimal given the price e .

Note that this relationship holds under both of the alternative exchange rate theories. The key difference between the two theories is as follows. In the first model, the exchange rate e is truly the marginal rate of transformation between a and b , so the price e is sending the correct signal about the relative costs of delivering consumption to the two representative agents. In the second model, in contrast, the true technology for transforming one good into the other never changes, so the price e is sending a false signal about the relative costs of producing consumption. In this model, the efficient allocation is characterized by a constant (e -invariant) consumption ratio across countries.

This difference between the two exchange rate models shows up in their respective world resource constraints. In the first model, this constraint (expressed in units of good a) is

$$c_1(e_j) + e_j c_2(e_j) = 1 + e_j,$$

whereas in the second model it is

$$c_1(e_j) + c_2(e_j) = 2.$$

Combining the consumption-sharing rule (27) with these two resource constraints allows us to solve for equilibrium consumption in the two models.

In the first,

$$c_1(e) = e^{\frac{1}{\gamma}} \frac{1+e}{e^{\frac{1}{\gamma}} + e}, \quad c_2(e) = \frac{1+e}{e^{\frac{1}{\gamma}} + e}.$$

In the second,

$$c_1(e) = e^{\frac{1}{\gamma}} \frac{2}{e^{\frac{1}{\gamma}} + 1}, \quad c_2(e) = \frac{2}{e^{\frac{1}{\gamma}} + 1}.$$

We now compare welfare across market structures. For this purpose, we suppose $E = \{1/\chi, \chi\}$ and that $\pi(e_1) = \pi(e_2) = 0.5$, so that the mean log exchange rate is zero and the standard deviation is $\ln \chi$. The standard deviation of the real exchange rate reported in Table 6 is 0.0639. Thus we

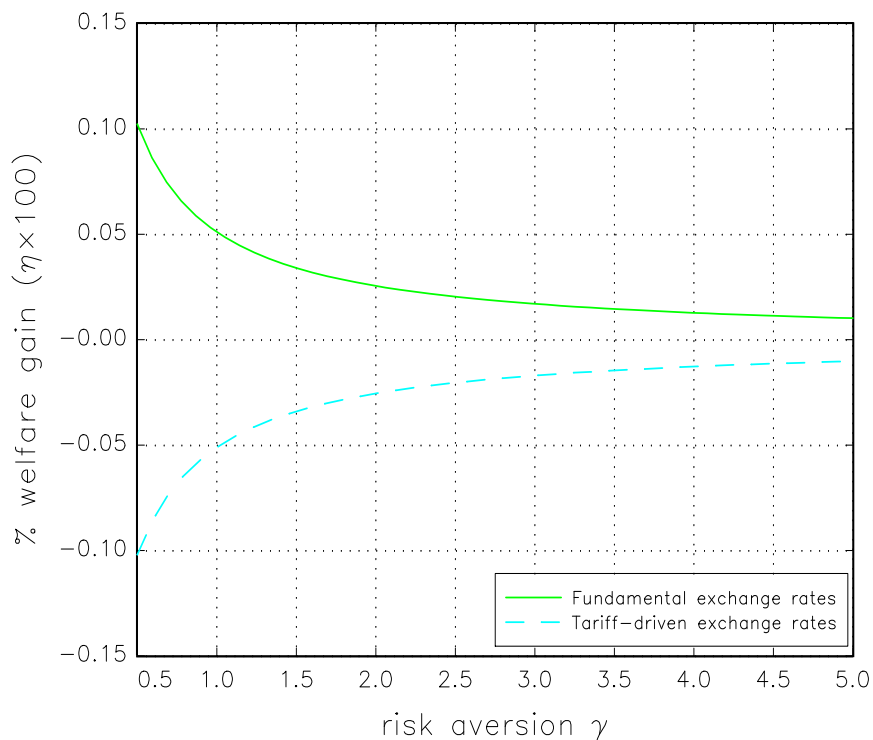
set $\chi = \exp(0.0639) = 1.066$.

Expected utility under autarky is simply $U^{AUT} = (1 - \gamma)^{-1}$. Expected utility under complete markets is $U_i^{CM} = 0.5u(c_i(1/\chi)) + 0.5u(c_i(\chi))$. In both exchange rate models, $U_1^{CM} = U_2^{CM}$. We define the welfare gain of moving from autarky to complete markets as the value for η that solves

$$(1 + \eta)^{1-\gamma} U^{AUT} = U^{CM}.$$

Figure 8 plots η for a range of values for risk aversion γ .

Figure 8: Welfare gains from financial integration



Focus first on the fundamentals-driven model for the exchange rate. For this model, the welfare gains of moving from autarky to an efficient allocation are small but nontrivial. For $\gamma = 1$ (log preferences) the gain is 0.051% of consumption, which is much larger than Lucas' (1987) estimate of the welfare cost of business cycles of 0.008%. One reason gains are larger is that the real exchange rate is much more volatile than consumption at business cycle frequencies. Welfare gains are declining in risk aversion. The logic is that the less tolerant are agents of fluctuations in consumption, the smaller are the gains to diverting resources to the country in which consumption is cheapest. In the limit as $\gamma \rightarrow \infty$, $\eta \rightarrow 0$. As $\gamma \rightarrow 0$, $\eta \rightarrow (\chi - 1)/2 = 3.3\%$.

Now look at the second tariff-driven model for the exchange rate. In this model, autarky is

efficient, and introducing complete markets reduces welfare. In fact, the welfare plot is the mirror image of that for the fundamentals-driven exchange rate model!

The key message to take away from this example is that identifying the source of exchange rate variation is critical for interpreting exchange-rate-based measures of efficiency. When the exchange rate is driven by fundamental technology shocks, autarky is inefficient, and there are welfare gains from introducing asset trade. The Backus-Smith correlation is a useful diagnostic for this inefficiency. In autarky the correlation is zero, and in complete markets (when allocations are efficient) it is one.

When the exchange rate is driven by random tariffs and disconnected from preferences and technologies, everything is reversed. Autarky is efficient, and there are welfare losses from introducing asset trade. The Backus-Smith correlation is a completely misleading diagnostic tool. In autarky (when allocations are efficient) the correlation is zero, whereas in complete markets (when allocations are inefficient), it is one.

A second message from the example is that in a distorted economy, financial liberalization can be welfare reducing. In particular, in our second example distortionary tariffs lead to an inefficient allocation of resources when international financial markets are complete. One remedy is to eliminate the tariffs. An alternative is to prevent agents from acting on the resulting distorted price signals by ruling out international asset trade.

4.4.3 Assessing Efficiency Using Asset Prices

Taking first differences of the consumption efficiency condition (eq. 10) gives

$$\log e(s^{t+1}) - \log e(s^t) = \log m_2(s^{t+1}) - \log m_1(s^{t+1}), \quad (28)$$

where the left-hand side is the log change in the real exchange rate, and the right-hand side is the difference between the growth rate of the marginal utility of consumption in country 2 and the corresponding growth rate in country 1. In finance, these growth rates are called stochastic discount factors (SDFs) because they define the appropriate way to discount payoffs when pricing assets. In particular, any asset j traded in country i with payoffs $x^j(s^{t+1})$ in state s^{t+1} has price $p^j(s^t)$ given by

$$p^j(s^t) = E_{s^t} [m_i(s^{t+1})x^j(s^{t+1})].$$

As noted before, growth rates of marginal utility (SDFs) are not observed directly. The standard macroeconomic approach that we have followed up to now has been to make assumptions on preferences (time separability and the absence of any preference shocks) such that the stochastic discount factor is proportional to consumption growth, which can be measured directly. However, these assumptions are quite restrictive.

An alternative approach is to note that although the SDF is not directly observable, we can learn something about its statistical properties by looking at asset prices. The fact that there is a substantial excess return to stocks over bonds indicates that there must be a large negative covariance between $m_i(s^{t+1})$ and stock returns, which in turn requires that $m_i(s^{t+1})$ must be very volatile.

Brandt, Cochrane, and Santa-Clara (2006) start with eq. 28. They then note that the variance of the left-hand side (the variance of real exchange rate changes) must equal to the variance of the right-hand side. The variance of the right-hand side is the sum of the variances of the two SDFs, minus twice the covariance between them. They plug in empirical values for the variance of exchange rate changes and (high) variances for SDFs consistent with equity premium evidence and conclude that the covariance between SDFs must be positive and large. They conclude that “international risk sharing is better than you think.”

At this point, it is important to emphasize that the Brandt et al. (2006) notion of risk sharing is quite different from our notion of efficiency. We have defined allocations to be efficient if no Pareto-improving reallocation of resources is possible, taking as given physical transportation frictions and taking as given differences in preferences across countries. Brandt et al. define perfect risk-sharing as being achieved when all restrictions on goods trade as well as all restrictions on asset trade are removed. They write, “Risk sharing requires frictionless goods markets. The container ship is a risk sharing innovation as important as 24 hour trading.” In their view, real exchange rate volatility itself is a direct measure of (lack of) risk sharing. If all costs of trading goods and assets could be eliminated, SDFs would comove perfectly across countries, and the real exchange rate would not move at all. Because the real exchange rate already moves so little (relative to SDFs) Brandt et al. conclude that risk sharing in this broad sense is already very good.

Colacito and Croce (2011) take inspiration from the Brandt et al. paper. They set themselves the task of trying to construct a model in which SDFs comove strongly, even though measured consumption growth does not. Their answer is that such a scenario is perfectly possible in a world with non-time-separable preferences (à la Epstein and Zin, 1989) and long run risks (à la Bansal and Yaron, 2004). The idea is that the cross-country correlation of SDFs is driven primarily by highly correlated long-run risks, whereas the cross-country correlation of consumption growth (in the short run) is driven by weakly correlated transitory shocks. In the Colacito and Croce environment, there is full home bias in preferences, so agents in each country only want to consume their local endowment. Thus, equilibrium allocations are efficient (by our definition) for any asset market structure, including autarky. This highlights an undesirable feature of the Brandt et al. (2006) measure of risk sharing: in the context of the Colacito and Croce model, nothing can be learned about risk sharing by examining the covariance between SDFs across countries or the volatility of the real exchange rate.

Colacito and Croce (2010) explore the welfare gains generated by moving from financial autarky

to complete markets in a similar environment to their (2011) paper. One important extension relative to their previous paper is that they do not impose perfect home bias in preferences, so the welfare gains from increasing financial integration are not zero by construction. They find that the potential welfare gains in moving from autarkic to efficient allocations are potentially very large when (i) there is intermediate home bias in preferences, (ii) risk aversion is high, and (iii) there is substantial long-run risk that is not strongly correlated across countries. Lewis and Liu (2012) conduct a similar exercise, using a different strategy for identifying the critical persistent risk correlation that does not rely on assuming complete financial markets but only exploits standard asset pricing equations.

These papers connect the assessment of insurance against long-run risk (Section 3) and the assessment of insurance against business cycle risk (Section 4). In particular, both papers conclude that little is at stake when it comes to insuring transitory business cycle shocks, whereas the potential welfare costs from inefficient allocations are much larger when shocks are very persistent. However, it remains an open question how well these long-run risks are actually insured. Colacito and Croce (2010) and Lewis and Liu (2012) argue that long-run shocks are either highly correlated or well insured across countries. Is this conclusion consistent with the substantial cross-country variation in long-run consumption growth described in Section 3? Nakamura et al. (2012) present some new empirical evidence that can potentially be used to address this question.

4.4.4 Assessing Efficiency Using International Portfolios

The most obvious mechanism via which agents can hedge country-specific risk is by holding foreign assets that increase agents' exposure to shocks in other countries and appropriately reduce exposure to domestic shocks. Within the context of a specific model, one can ask what portfolios (if any) deliver an efficient cross-country allocation of consumption and capital. One can then compare those portfolios to the ones observed in the data. If one has confidence in the model, then the distance between observed and efficient portfolios can be used to gauge efficiency.

A very simple model is a symmetric two-country world in which two trees (one in each country) produce stochastic dividends of apples in each period (see Lucas, 1982). Agents in the two countries enjoy apples equally. An efficient allocation involves agents in each country consuming fixed fractions of the world apple endowment in each period. Now consider a decentralized environment. Absent any opportunities for international asset trade, consumption in each country would equal the country-specific dividend. This allocation is inefficient. If stock markets are introduced that allow agents to freely trade shares in trees, then in equilibrium the representative agent in each country will choose to hold half the shares in each tree. This equilibrium is efficient because each representative agent receives and consumes half of the world apple endowment in each period.

It is well known that in practice, portfolios tend to be heavily biased toward domestic assets.

Thus, for example, Americans mostly hold stocks in US companies and US government or corporate bonds. Relative to the simple model outlined above, these portfolio choices would appear to be inefficient. But this setup is obviously too simple. One important respect in which it is too simple is that it assumes that asset income is agents' only source of income. In practice, most of household income comes from labor income, which is inherently almost impossible to diversify (short of working in multiple countries). Baxter and Jermann (1997) emphasize that introducing nondiversifiable labor income increases the gap between portfolios predicted by theory versus those observed in the data. In the context of our apple tree example, suppose that domestic agents receive half of their domestic tree's apple endowment, as compensation for the work of picking the apples. Stock owners have the rights to the remaining half. Now equilibrium portfolios will be 100% foreign biased: domestic agents will buy all the shares in the foreign tree and vice versa. With those portfolios, agents will again end up consuming half the world endowment of apples in each period. Thus, introducing labor income seems to make the international diversification puzzle worse, as Baxter and Jermann emphasized.

The simple model we started with is unrealistic along another dimension: it is a one-good model. Consider the simple two-good endowment economy described in subsection 4.2.1. Suppose that the tree in the domestic country produces apples and the foreign tree produces bananas, and that fruit can be freely traded between countries. Now if the domestic tree has a particularly good year and produces lots of apples, the world relative price of apples will fall and the price of bananas will rise. This relative price movement provides automatic insurance against country-specific shocks. In fact, as Cole and Obstfeld (1991) showed, given a unitary elasticity of substitution in preferences between fruit, the terms of trade will move one-for-one with the relative fruit endowment. In that case, irrespective of the portfolio mix between shares in domestic and foreign trees, income will automatically be equated across countries. Thus, if domestic and foreign agents have identical preferences, any portfolio will deliver the same efficient allocation. In that special case, even complete home bias is efficient.

Heathcote and Perri (2013) explore equilibrium portfolio choice in the stock economy model described in subsection (4.3) above. That is a natural environment for studying portfolio diversification, since it is the same model that has been widely used to study other dimensions of risk sharing, notably cross-country correlations of consumption, output, and investment, as well as comovement between relative consumption and the real exchange rate. The model features both labor income, as emphasized by Baxter and Jermann (1997), and the relative price effects emphasized by Cole and Obstfeld (1991). Recall that agents trade shares in representative domestic and foreign firms, which pay dividends equal to revenue less wage payments and investment spending. It turns out that the portfolios implied by a calibrated version of that model feature domestic-foreign compositions that are quantitatively similar to those observed in the data. Moreover, given (i) logarithmic utility over the consumption composite ($\gamma = 1$) and (ii) a unitary elasticity of substitution between

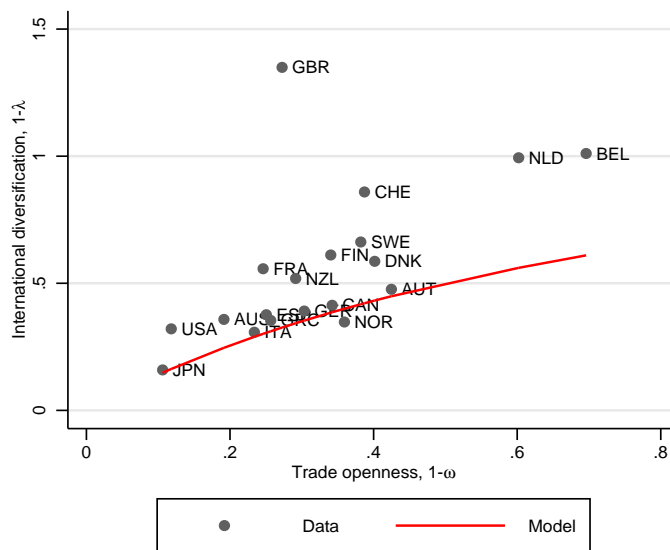
domestic and foreign intermediates in producing the final good (so that $G_1(a_1, b_1) = a_1^\omega b_1^{1-\omega}$), it is possible to characterize the equilibrium fraction of wealth invested in foreign stocks in closed form. Heathcote and Perri (Proposition 1) show that this share is constant and given by

$$1 - \lambda = \left(\frac{1 - \theta}{1 - \omega} + 2\theta \right)^{-1}. \quad (29)$$

They also show that allocations in this equilibrium are efficient. Note that diversification is increasing in the steady-state trade share in the model, $1 - \omega$. In addition, assuming $1 - \omega < 0.5$, diversification is decreasing in labor's share $1 - \theta$, contrary to the Baxter and Jermann result.

How do the (efficient) portfolio predictions of this model compare with the data? To answer that question, Heathcote and Perri assemble data on diversification ($1 - \lambda$) and trade shares ($1 - \omega$) for OECD economies over the period 1990 to 2007 (they assume a common-across-countries share for capital income $\theta = 0.36$). Figure 9 reproduces Figure 2 from that paper and plots diversification against the trade share as predicted by equation 29 (the solid line) alongside an analogous scatter plot for their sample of countries. Most country points are close to the line, and thus observed portfolios are close to the ones that in theory are consistent with efficiency.²¹

Figure 9: Portfolio diversification in the Heathcote-Perri model and in the data



How is it possible that home-biased portfolios allow for perfect pooling of country-specific risk, even though agents are already heavily exposed to domestic risk, thanks to nondiversifiable labor income? Perfect risk sharing in this environment requires that the ratio of marginal utilities of

²¹Great Britain is an exception, but high observed diversification there reflects the country's special position as an international financial center.

consumption across countries should be equal to the real exchange rate (which is the marginal rate of transformation between domestic and foreign consumption). Given logarithmic utility over consumption, this implies that the value of consumption should be equated across countries. A positive domestic shock raises the relative present value of domestic labor earnings. To equate the relative value of permanent income across countries, it is clear that the return on domestic agents' portfolios must decline relative to the return on foreign agents' portfolios. In the model, the relative return to domestic stocks falls following a positive domestic shock because the real exchange rate depreciates, reducing the relative value of domestic capital. Thus, a portfolio biased towards domestic assets offers a hedge against nondiversifiable labor income risk.

For risk aversion above one, efficient relative consumption becomes less sensitive to the real exchange rate than in the logarithmic example described above. In decentralized environments, the covariance between relative equity returns and the real exchange rate then becomes an additional driver of portfolio choice, since agents have an incentive to bias portfolios toward assets that offer relatively high returns in states in which the real exchange rate appreciates and the relative price of domestic consumption is high (see Van Wincoop and Warnock, 2010). Coeurdacier (2009) explores how introducing trade costs in goods affects the covariation between relative equity returns and the real exchange rate, and thus how trade costs impact portfolio composition. He finds that the theoretical pattern of covariation and thus the direction of portfolio bias depends on the level of trade costs: moderate costs imply a foreign bias, whereas high costs imply home bias.

The literature on international portfolio choice continues to grow. One lesson from this literature to date is that the extent of diversification predicted by theory is quite sensitive to model details, and new insights continue to emerge as richer models are built featuring more general preferences, alternative sources of risk, and more refined asset market structures. Engel and Matsumoto (2009) explore diversification in an environment with nominal frictions and find that the level of diversification that delivers efficient allocations (taking the frictions as given) depends on the degree of price stickiness. Coeurdacier and Gourinchas (2011) explore diversification in bonds and equities separately and find that bonds are well suited to hedging exchange rate risk, whereas equities are good for hedging nontradable (e.g., labor income) risk. Berriel and Bhattarai (2013) argue that domestic nominal bonds are a natural hedge for domestic agents against price level risk: if policy generates a surprise increase in the price level, the real return on domestic bonds will go down at the same time that the real value of debt and thus the present value of future taxes on domestic agents is reduced. Coeurdacier and Rey (2013) offer a much more comprehensive survey of the rapidly growing literature on this topic than we have provided here.

4.5 Welfare and Policies

One lesson that could be drawn from the analysis in this section is that the international allocation of resources across the business cycle is not necessarily inconsistent with efficiency. If that is the case it is interesting to ask what would be the welfare costs of shutting down the international financial markets that generate efficient allocations in our models. In Table 7 we report the welfare changes (in percentages of lifetime consumption) of going from complete markets to restricted asset market structures.²² In particular we report the value of η that solves the following equation

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi(s^t) U((1 + \eta)c_i^{CM}(s^t), n_i^{CM}(s^t)) = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi(s^t) U(c_i^{MKT}(s^t), n_i^{MKT}(s^t))$$

where initial conditions $k_i(s_0), z_i(s_0), i = 1, 2$ are set to their nonstochastic steady state values, and where a *CM* superscript denotes complete markets allocations, while a *MKT* superscript denotes a given market allocation. In particular, we focus on the market structures where only a bond is traded and on financial autarky.²³ Line 1 of the table reports these numbers for the parameters used in row 6 of Table 4 (i.e., those which deliver international correlations close to the data). At business cycle frequencies the efficiency gains associated with complete markets are not worth much in welfare terms. For example, going from complete markets to financial autarky is equivalent to a loss of lifetime consumption of around 0.02%. In line 2 we report welfare numbers for the economy described in Section 4.2.1 in which the two intermediate goods are received as endowments rather than being produced.²⁴

In this case the value of international financial markets is even lower, and the numbers are in line with the similar exercise performed by Cole and Obstfeld (1991). The reason is that in the endowment case financial markets cannot affect productive efficiency and only serve to keep marginal utilities aligned across countries (consumption efficiency). In the production economy, in contrast, financial markets also equalize the expected returns to investment across countries, thereby raising world productivity and allowing agents to enjoy higher average levels of consumption.

Are there plausible parameter values that deliver larger estimates for the value of international financial markets? In line 3 we consider a case in which the standard deviation of innovations to productivity is twice as large as the one needed to replicate the volatility of US GDP. This value is consistent with the volatility of business cycles observed in emerging economies (see Neumeyer and Perri, 2005). Now the welfare benefits of financial markets are larger (about 4 times as large

²²We compute welfare using second order approximations of the model since first order approximations yield inaccurate welfare results (see Kim and Kim, 2003).

²³We do not report the welfare losses of going to a stock economy, as in this simple setup an economy where two stocks are traded always yields welfare extremely close to complete markets (welfare is identical in the case discussed in Section 4.4.4).

²⁴The process for endowments is again chosen to match the standard deviation and international correlation of GDP in the data.

Table 7: The value of international financial markets ($100 \times \eta$)

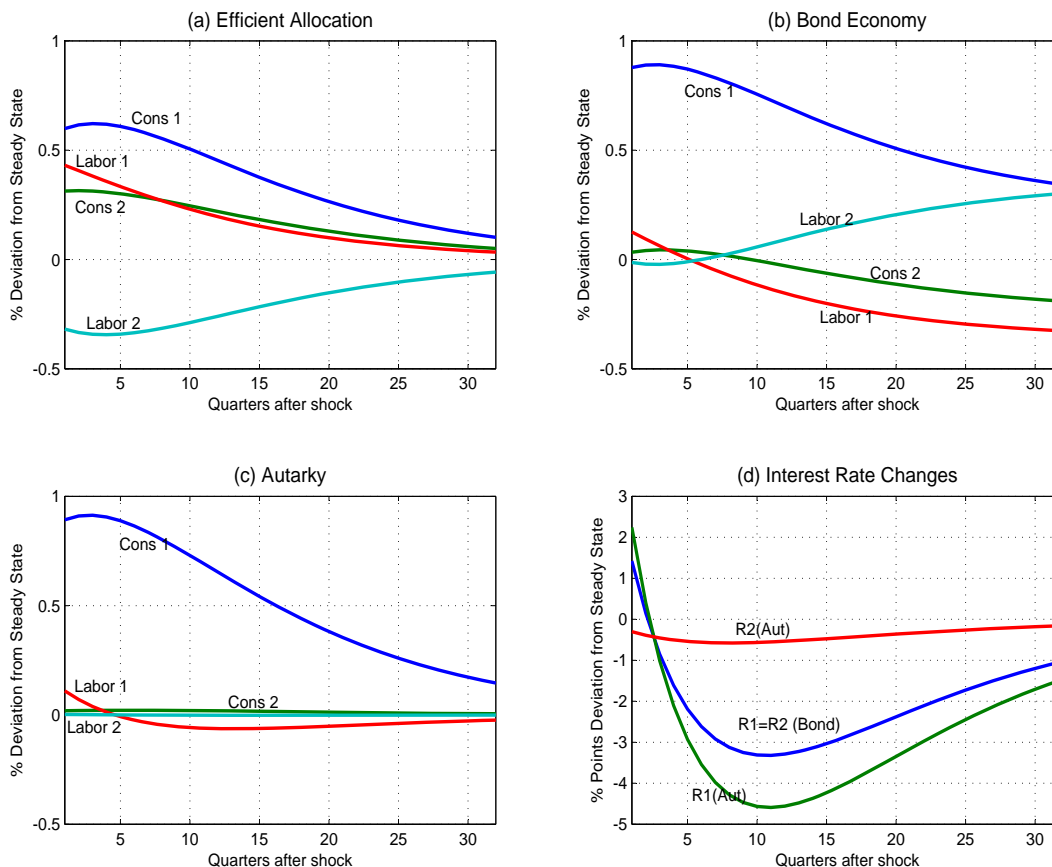
	Bond	Fin. Aut.
	(% of cons.)	
1. Baseline $\rho = 0.91, \psi = 0, \gamma = 1, \sigma = 0.6$ $\sigma_\varepsilon = 0.0082, \text{corr}(\varepsilon_1, \varepsilon_2) = 0.435$	-0.013	-0.021
2. Endowment	-0.003	-0.009
3. High Volatility, $\sigma_\varepsilon = 0.016$	-0.051	-0.086
4. Unit Elasticity, $\sigma = 1$	-0.001	-0.003
5. High Elasticity, $\sigma = 5$	-0.004	-0.007
Persistent shocks		
6. $\beta = 0.8, \rho = 0.91$	-0.132	-0.015
7. $\beta = 0.8, \rho = 0.91, \sigma = 5$	-0.054	-0.012
8. $\beta = 0.99, \rho = 0.999$	-0.409	-0.378

as in the baseline case). Line 4 and 5 consider two different values for the elasticity of substitution, σ . Note that the welfare value of financial markets is non-monotone in this parameter. When the elasticity is close to one, the value of financial markets is minimal because relative price movements provide a lot of insurance against idiosyncratic productivity shocks (recall that autarky is efficient given $\sigma = 1$ in the endowment version of the model).

In all the cases considered so far, welfare in the bond economy is higher than in financial autarky, suggesting that financial liberalization is always welfare-improving. The last three cases considered in Table 7 paint a rather different picture. In these cases parameters are chosen so that shocks generate large differences between the permanent incomes of the two countries. In order to generate such differences productivity shocks have to be persistent relative to individual discounting: ρ must be large relative to β . In lines 6, 7, and 8 we consider three parameterizations with this property, the first two with a low discount factor and baseline persistence, the last with the baseline discount factor and higher persistence. Two aspects of the results are remarkable. The first is that the welfare value of financial markets can be an order of magnitude larger than in the cases previously considered. The second is that the welfare ordering is reversed: welfare is higher under financial autarky than in the bond economy.

It is not surprising that the costs of eliminating all financial markets are larger the more persistent are relative shocks, because more persistent shocks mean larger gains from mutual insurance.

Figure 10: Comparing allocations across alternative market structures



Bonds (being noncontingent) are not able to deliver this insurance. Welfare can actually be higher in financial autarky than in the bond economy because of the presence of pecuniary externalities, i.e., general equilibrium price effects.²⁵ As an illustration of these effects consider the response to a productivity shock in country 1 in the high elasticity parameterization (line 7). The symmetric efficient allocation involves consumption increasing almost identically in both countries (consumption efficiency) while investment (not plotted) and labor rise by more in country 1 than in country 2 (productive efficiency). These responses are displayed in Panel A of Figure 10.

In the incomplete markets models (Panels B and C) allocations are not efficient. Country 1 is the prime beneficiary of its higher productivity (because insurance markets do not exist) and thus residents of that country permanently increase consumption sharply relative to residents of country 2. The shock to relative wealth also shows up in labor supply: relative to the efficient allocation, agents in country 1 work too little, while those in country 2 work too much.

Now consider the differences between the financial autarky and bond economy models. In

²⁵For a discussion of these effects in this class of models see also Corsetti et al., 2012.

financial autarky, country 1's desire to save (given its temporarily high income) translates into a large fall in the equilibrium interest rate (Panel D), and an associated sharply declining consumption profile during most of transition. When countries can trade a bond, country 1 is able to achieve a smoother consumption profile by saving in country 2, and as savings flow abroad there is less downward pressure on the domestic interest rate. Thus, relative to financial autarky, the equilibrium interest rate declines by less in country 1, and declines by more in country 2. Higher interest rates are beneficial for country 1, which is a net saver. In country 2 it is in each atomistic agent's best interests to borrow from country 1. But in aggregate, their borrowing raises the interest rate, which hurts all borrowers in country 2. Thus, introducing trade in a bond changes interest rates in a way that effectively amplifies the differential impact of the initial shock across countries, reduces effective insurance, and moves equilibrium allocations further away from the efficient ones.²⁶

This finding is intriguing from a policy point of view because it suggests that it might be desirable for two countries that can only trade a noncontingent bond to close international financial markets altogether. Interestingly, the model suggests that this situation is more likely in exactly those cases where welfare differs significantly across market structures. Thus the welfare effects of going from a bond economy to financial autarky might be both positive and economically significant (in the example considered in line 6, the gain exceeds 0.1 percent of consumption).

5 Conclusions

The conclusions we take from this paper are simple. First, over the long run allocations appear inefficient. In particular, there is little evidence that consumption responds efficiently to persistent cross-country differentials in output growth. This is important, because the potential welfare gains from achieving more efficient allocations in the long run are large. In contrast, it is difficult to reject the hypothesis that allocations respond efficiently to business cycle frequency fluctuations. Patterns of cross-country comovement in macro aggregates are consistent with efficiency, as are observed levels of portfolio diversification. It is difficult to reconcile observed exchange rate dynamics with efficiency, but it is also difficult to reconcile them with alternative decentralized asset market structures that explicitly limit international risk-sharing. Thus we view evidence on efficiency from prices as inconclusive, pending better theories of nominal exchange rate dynamics. Future work should also focus on connecting evidence on risk-sharing from asset prices to evidence from quantities. To date, the macroeconomic literature has largely neglected asset price evidence, while the finance literature has typically treated the process for consumption as exogenous.

One important area for future research that we have largely neglected in this survey is the interaction between risk sharing against idiosyncratic shocks within a country versus risk sharing

²⁶In general the introduction of an asset also changes the terms of trade, which can be important. Here we isolate the role of general equilibrium interest rate movements by considering a high elasticity parametrization in which the terms of trade moves little.

against country-level shocks between countries. In environments with idiosyncratic risk the extent of openness to international financial flows can interact with information or enforcement frictions that preclude a first best allocation of resources within a country (see, for example, Broer, 2009, Broner and Ventura, 2009, Martin and Taddei 2012, and Mendoza et al. 2009).

Given our assessment that actual international allocations are inefficient, at least in the long run, an obvious next step is to consider whether specific policy interventions might increase efficiency. In many instances, if country-specific risks are not pooled because of frictions in international financial markets, working to remove these frictions should increase efficiency. However, we have discussed two examples in which financial liberalization can be welfare reducing. First, in Section 4.5 we discussed an example in which complete markets guarantees efficiency, but partial liberalization – introducing international bond trade relative to financial autarky – reduces welfare. Second, in Section 4.4.2 we discussed an example in which the source of inefficiency is non-fundamental driven fluctuations in the exchange rate. In that context, closing international financial markets is a way to insulate the economy from these otherwise distortionary shocks.

References

- [1] Aguiar, M., and G. Gopinath, 2007, Emerging Market Business Cycles: The Cycle Is the Trend, *Journal of Political Economy*, 115(1), 69-102.
- [2] Alfaro, L., S. Kalemli-Ozcan, and V. Volosovych (2011), Sovereigns, Upstream Capital Flows and Global Imbalances, NBER Working Paper No. 17396.
- [3] Arvanitis, A.V., and A. Mikkola, 1996, Asset-Market Structure and International Trade Dynamics, *American Economic Review*, 86(2), 67-70.
- [4] Backus, D.K., and P.J. Kehoe, 1992, International Evidence on the Historical Properties of Business Cycles, *American Economic Review*, 82(4), 864-888.
- [5] Backus, D.K., P.J. Kehoe, and F.E. Kydland, 1992, International Real Business Cycles, *Journal of Political Economy*, 100(4), 745-775.
- [6] Backus, D.K., P.J. Kehoe, and F.E. Kydland, 1994, Dynamics of the Trade Balance and the Terms of Trade: The J-Curve? *American Economic Review*, 84(1), 84-103.
- [7] Backus, D.K., and G.W. Smith, 1993, Consumption and Real Exchange Rates in Dynamic Economies with Non-traded Goods, *Journal of International Economics*, 35(3-4), 297-316.
- [8] Bansal, R., and A. Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance*, 59(4), 1481-1509.
- [9] Baxter, M., 2012, International Risk-Sharing in the Short Run and in the Long Run, *Canadian Journal of Economics*, 45(2), 376-393.
- [10] Baxter, M., and M.J. Crucini, 1995, Business Cycles and the Asset Structure of Foreign Trade, *International Economic Review*, 36(4), 821-854.
- [11] Baxter, M., and U. Jermann, 1997, The International Diversification Puzzle Is Worse Than You Think, *American Economic Review*, 87(1), 170-180.
- [12] Berka, M., M.B. Devereux, and C. Engel, 2012, Real Exchange Rate Adjustment in and out of the Eurozone, *American Economic Review*, 102(3), 179-185.
- [13] Berriel, T.C., and S. Bhattarai, 2013, Hedging against the Government: A Solution to the Home Asset Bias Puzzle, *American Economic Journal: Macroeconomics*, 5(1), 102-134.
- [14] Brandt, M.W., J.H. Cochrane, and P. Santa-Clara, 2006, International Risk Sharing Is Better Than You Think, or Exchange Rates Are Too Smooth, *Journal of Monetary Economics*, 53(4), 671-698.

- [15] Broer, T., 2010, Domestic or Global Imbalances? Rising Inequality and the Fall in the US Current Account, Working Paper, Institute for International Economic Studies, Stockholm.
- [16] Broner, F.A., and J. Ventura, 2010, Rethinking the Effects of Financial Liberalization, NBER Working Paper No. 16640.
- [17] Buera, P.J., and Y. Shin, 2011, Productivity Growth and Capital Flows: The Dynamics of Reforms, working paper, University of California at Los Angeles.
- [18] Burstein, A., and G. Gopinath, (2013), International Prices and Exchange Rates, in G. Gopinath, E. Helpman, and K. Rogoff, eds., Handbook of International Economics, vol. 4, Amsterdam: Elsevier.
- [19] Caselli, F., and J. Feyrer, 2007, The Marginal Product of Capital, Quarterly Journal of Economics, 122(2), 535-568.
- [20] Chari, V.V., P.J. Kehoe, and E.R. McGrattan, 2002, Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?, Review of Economic Studies, 69(3), 533-563.
- [21] Chinn, M. D., and H. Ito, 2008, A New Measure of Financial Openness, Journal of Comparative Policy Analysis, 10(3), 309-322.
- [22] Coeurdacier, N., 2009, Do Trade Costs in Goods Markets Lead to Home Bias in Equities? Journal of International Economics, 77(1), 86-100.
- [23] Coeurdacier, N., and P.-O. Gourinchas, 2011, When Bonds Matter: Home Bias in Goods and Assets, NBER Working Paper No. 17560.
- [24] Coeurdacier, N., and H. Rey, 2013, Home Bias in Open Economy Financial Macroeconomics, Journal of Economic Literature, 51(1), 63-115.
- [25] Colacito, R., and M.M. Croce, 2010, The Short and Long Run Benefits of Financial Integration, American Economic Review, 100(2), 527-531.
- [26] Colacito, R., and M.M. Croce, 2011, Risks for the Long Run and the Real Exchange Rate, Journal of Political Economy, 119(1), 153-181.
- [27] Cole, H.L., and M. Obstfeld, 1991, Commodity Trade and International Risk Sharing: How Much Do Financial Markets Matter? Journal of Monetary Economics, 28(1), 3-24.
- [28] Corsetti, G., L. Dedola, and S. Leduc, 2008, International Risk Sharing and the Transmission of Productivity Shocks, Review of Economic Studies, 75(2), 443-473.

- [29] Corsetti, G., L. Dedola, and S. Leduc, 2012, Demand Imbalances, Exchange-Rate Misalignment and Monetary Policy, Working Paper, European Central Bank.
- [30] Doepke M., and F. Zilibotti, 2006, Occupational Choice and the Spirit of Capitalism, *Quarterly Journal of Economics*, 123(2), 747-793.
- [31] Devereux, M.B., and V. Hnatkovska, 2011, Consumption Risk-Sharing and the Real Exchange Rate: Why Does the Nominal Exchange Rate Make Such a Difference? NBER Working Paper No. 17288.
- [32] Engel, C., 1999, Accounting for U.S. Real Exchange Rate Changes, *Journal of Political Economy*, 107(3), 507-538.
- [33] Engel, C., and A. Matsumoto, 2009, The International Diversification Puzzle When Goods Prices Are Sticky: It's Really about Exchange-Rate Hedging, Not Equity Portfolios, *American Economic Journal: Macroeconomics*, 1(2), 155-188.
- [34] Engel, C., 2013, Exchange Rates and Interest Parity, in G. Gopinath, E. Helpman and K. Rogoff, eds., *Handbook of International Economics*, vol. 4, Amsterdam: Elsevier.
- [35] Epstein, L.G., and S.E. Zin, 1989, Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica*, 57(4), 937-969.
- [36] Fitzgerald, D., 2012, Trade Costs, Asset Market Frictions, and Risk Sharing, *American Economic Review*, 102(6), 2700-2733.
- [37] Gourinchas, P.-O., and O. Jeanne, 2006, The Elusive Gains from International Financial Integration, *Review of Economic Studies*, 73(3), 715-741.
- [38] Gourinchas, P.-O., and O. Jeanne, 2013, Capital Flows to Developing Countries: The Allocation Puzzle, *Review of Economic Studies* 80 (4), 1484-1515
- [39] Gourinchas, P.-O., and H. Rey, 2013, External Adjustment, Global Imbalances, Valuation Effects, in G. Gopinath, E. Helpman and K. Rogoff, eds., *Handbook of International Economics*, vol. 4, Amsterdam: Elsevier.
- [40] Gourinchas, P.-O., H. Rey, and N. Govillot, 2010, Exorbitant Privilege and Exorbitant Duty, IMES Discussion Paper Series 10-E-20, Bank of Japan.
- [41] Hadzi-Vaskov, M., 2008, Does the Nominal Exchange Rate Explain the Backus-Smith Puzzle? Evidence from the Eurozone, Discussion Paper 07-32, Utrecht School of Economics.

- [42] Hamilton, B., and J. Whalley, 1984, Efficiency and Distributional Implications of Global Restrictions on Labor Mobility: Calculations and Policy Implications, *Journal of Development Economics*, 14(1-2), 61-75.
- [43] Heathcote, J., and F. Perri, 2002, Financial Autarky and International Business Cycles, *Journal of Monetary Economics*, 49(3), 601-627.
- [44] Heathcote, J., and F. Perri, 2004, Financial Globalization and Real Regionalization, *Journal of Economic Theory*, 119(1), 207-243.
- [45] Heathcote, J., and F. Perri, 2013, The international Diversification Puzzle Is Not As Bad As You Think, *Journal of Political Economy*, Forthcoming
- [46] Hess, G., and K. Shin, 2010, Understanding the Backus-Smith Puzzle: It's the (Nominal) Exchange Rate Stupid, *Journal of International Money and Finance*, 29(1), 169-180.
- [47] Kaldor, N., 1957, A Model of Economic Growth, *Economic Journal*, 67(268), 591-624.
- [48] Kehoe, P.J., and F. Perri, 2002, International Business Cycles with Endogenous Incomplete Markets, *Econometrica*, 70(3), 907-928.
- [49] Kim, J., and S. H. Kim, 2003, Spurious Welfare Reversals in International Business Cycle Models, *Journal of International Economics*, 60(2), 471-500.
- [50] Kocherlakota, N.R., and L. Pistaferri, 2007, Household Heterogeneity and Real Exchange Rates, *Economic Journal*, 117(519), C1-C25.
- [51] Kose, A., E. Prasad, K. Rogoff, and S.J. Wei, 2009, Financial Globalization: A Reappraisal, *IMF Staff Papers*, 59, 8-62.
- [52] Lane, P.R., and G.M. Milesi-Ferretti, 2007, The External Wealth of Nations Mark II: Revised and Extended Estimates of Foreign Assets and Liabilities, 1970-2004, *Journal of International Economics*, 73(2), 223-250.
- [53] Lewis, K.K., 1996, What Can Explain the Apparent Lack of International Consumption Risk Sharing? *Journal of Political Economy*, 104(2), 267-297.
- [54] Lewis, K.K., and E.X. Liu, 2012, International Consumption Risk Is Shared After All: An Asset Return View, NBER Working Paper No. 17872.
- [55] Lucas, R.E., 1982, Interest rates and currency prices in a two-country world, *Journal of Monetary Economics*, 10, 335-359.
- [56] Lucas, R.E., 1987, *Models of Business Cycles*, Oxford: Basil Blackwell.

- [57] Lucas, R.E., 1990, Why Doesn't Capital Flow from Rich to Poor Countries? *American Economic Review*, 80(2), 92-96.
- [58] Martin, A. and F. Taddei, 2012. International capital flows and credit market imperfections: a tale of two frictions, *Journal of International Economics* 89 (2), 441–452.
- [59] Mendoza, E.G., V. Quadrini, and J.V. Rios-Rull, 2009, Financial Integration, Financial Development, and Global Imbalances, *Journal of Political Economy*, 117(3), 371-416.
- [60] Mussa, M., 1986, Nominal Exchange Rate Regimes and the Behavior of Real Exchange Rates: Evidence and Implications, *Carnegie-Rochester Conference Series on Public Policy*, 25(1), 117-213.
- [61] Nakamura, E., D. Sergeyev, and J. Steinsson, 2012, Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence, NBER Working Paper No. 18128.
- [62] Neumeayer P.A., and F. Perri, 2005, Business Cycles in Emerging Economies: The Role of Interest Rates, *Journal of Monetary Economics*, 52, 345-380.
- [63] Obstfeld, M., and K. Rogoff, 2001, The Six Major Puzzles in International Macroeconomics: Is There a Common Cause? in *NBER Macroeconomics Annual 2000*, Volume 15, B. Bernanke and K. Rogoff (eds), MIT Press, 339-412.
- [64] Pakko, M.R., 1997, International Risk Sharing and Low Cross-Country Consumption Correlations: Are They Really Inconsistent? *Review of International Economics*, 5(3), 386-400.
- [65] Pakko, M.R., 2004, A Spectral Analysis of the Cross-Country Consumption Correlation Puzzle, *Economics Letters*, 84(3), pages 341-347.
- [66] Rabanal P., and J.F. Rubio-Ramirez, 2010, Can International Macroeconomic Models Explain Low-Frequency Movements of Real Exchange Rates? Working Paper, Duke University.
- [67] Ruhl, K.J., 2008, The International Elasticity Puzzle, Working Paper 08-30, New York University, Leonard N. Stern School of Business, Department of Economics.
- [68] Sandri D., 2010, Growth and Capital Flows with Risky Entrepreneurship, IMF Working Paper 10/37.
- [69] Song, Z., K. Storesletten, and F. Zilibotti, 2011, Growing Like China, *American Economic Review*, 101(1), 196-233.
- [70] Stigler, G.J., and G.S. Becker, 1977, De Gustibus Non Est Disputandum, *American Economic Review*, 67(2), 76-90.

[71] Van Wincoop, E., and F. Warnock, 2010, Can Trade Costs in Goods Explain Home Bias in Assets? *Journal of International Money and Finance*, 29(6), 1108-1123.

[72] Viani, F., 2011, International Financial Flows, Real Exchange Rates and Cross-Border Insurance, Banco de Espana Working Paper 1038, Banco de Espana.

6 Appendix

PROOF OF PROPOSITION 1

For any variable x , let $\hat{x}(s^t) = \log x(s^t) - \log \bar{x}$, where \bar{x} is the nonstochastic steady-state value for x .

Efficient allocations are defined by values for (a_1, a_2, b_1, b_2) that satisfy the resource constraints and the consumption efficiency conditions.

The log-linearized versions of these equations are

$$\begin{aligned} s\hat{a}_1 + (1-s)\hat{a}_2 &= \hat{z}_1 \\ (1-s)\hat{b}_1 + s\hat{b}_2 &= \hat{z}_2 \\ \left(-\gamma + \frac{1}{\sigma}\right) \left(s\hat{a}_1 + (1-s)\hat{b}_1\right) - \frac{1}{\sigma}\hat{a}_1 &= \left(-\gamma + \frac{1}{\sigma}\right) \left(s\hat{b}_2 + (1-s)\hat{a}_2\right) - \frac{1}{\sigma}\hat{a}_2 \\ \left(-\gamma + \frac{1}{\sigma}\right) \left(s\hat{a}_1 + (1-s)\hat{b}_1\right) - \frac{1}{\sigma}\hat{b}_1 &= \left(-\gamma + \frac{1}{\sigma}\right) \left(s\hat{b}_2 + (1-s)\hat{a}_2\right) - \frac{1}{\sigma}\hat{b}_2. \end{aligned}$$

The solutions are

$$\begin{aligned} \hat{a}_1 - \hat{z}_1 &= (1-s)\Omega(\hat{z}_1 - \hat{z}_2) \\ \hat{a}_2 - \hat{z}_1 &= -s\Omega(\hat{z}_1 - \hat{z}_2) \\ \hat{b}_2 - \hat{z}_2 &= -(1-s)\Omega(\hat{z}_1 - \hat{z}_2) \\ \hat{b}_1 - \hat{z}_2 &= s\Omega(\hat{z}_1 - \hat{z}_2), \end{aligned}$$

where

$$\Omega = \frac{(1-\sigma\gamma)(1-2s)}{4s(1-\sigma\gamma)(1-s)-1}.$$

Consumption and output (measured in units of the final consumption good) deviations are given

by

$$\begin{aligned}\hat{c}_1 &= s\hat{a}_1 + (1-s)\hat{b}_1 \\ \hat{c}_2 &= (1-s)\hat{a}_2 + s\hat{b}_2 \\ \hat{y}_1^c &= \frac{1}{\sigma}\hat{c}_1 - \frac{1}{\sigma}\hat{a}_1 + \hat{z}_1 \\ \hat{y}_2^c &= \frac{1}{\sigma}\hat{c}_2 - \frac{1}{\sigma}\hat{b}_2 + \hat{z}_2.\end{aligned}$$

1. The pass-through coefficient from changes in relative output to relative consumption is given by

$$(\hat{c}_1 - \hat{c}_2) = \frac{(2s-1)}{2\gamma(s-1) + 4s(\sigma\gamma-1)(1-s) + 1} (\hat{y}_1^c - \hat{y}_2^c).$$

2. Net exports are given by

$$(\hat{y}_1^c - \hat{c}_1) = \frac{2s(1-\sigma\gamma) - (1-\gamma)}{2\gamma + 4s(1-\sigma\gamma) - \frac{1}{1-s}} (\hat{y}_1^c - \hat{y}_2^c).$$

3. The cross-country consumption and output correlations are

$$\begin{aligned}\text{corr}(\hat{c}_1, \hat{c}_2) &= \frac{K_c + (1-K_c)\rho}{(1-K_c) + K_c\rho} \\ \text{corr}(\hat{y}_1^c, \hat{y}_2^c) &= \frac{K_y + (1-K_y)\rho}{(1-K_y) + K_y\rho},\end{aligned}$$

where ρ is the correlation of productivity shocks, and if and only if

$$\begin{aligned}K_c &= 2s(1-s) \left(\frac{2(1-\sigma\gamma)(1-s) - 1}{4s(1-\sigma\gamma)(1-s) - 1} \right) \left(1 - 2s \frac{(1-\sigma\gamma)(1-2s)}{4s(1-\sigma\gamma)(1-s) - 1} \right) \\ K_y &= 2 \frac{\gamma(1-s)}{4s(\sigma\gamma-1)(1-s) + 1} \left(1 - \frac{\gamma(1-s)}{4s(\sigma\gamma-1)(1-s) + 1} \right).\end{aligned}$$

Given these expressions, it is straightforward to verify that if and only if $\sigma < \tilde{\sigma}(s, \gamma)$,

$$\begin{aligned}\frac{\hat{c}_1 - \hat{c}_2}{\hat{y}_1^c - \hat{y}_2^c} &> 1 \\ \text{corr}(\hat{y}_1^c - \hat{c}_1, \hat{y}_1^c - \hat{y}_2^c) &< 0 \\ \text{corr}(\hat{y}_1^c, \hat{y}_2^c) &> \text{corr}(\hat{c}_1, \hat{c}_2).\end{aligned}$$