International Trade and Intertemporal Substitution

Fernando Leibovici

New York University

Michael E. Waugh

New York University

First Version: July 2010

This Version: Febuary 2012

ABSTRACT

This paper studies a dynamic extension of international trade models and its ability to explain the behavior of imports at business-cycle frequencies. Our premise is that because international trade is time-intensive, variation in the rate at which agents are willing to substitute across time affects how trade volumes respond to changes in income and prices. We formalize this idea and find that our calibrated model is quantitatively consistent with cyclical properties of U.S. imports. We then show that given the behavior of output and relative prices during the 2008-2009 crisis, our model goes a long way toward explaining the collapse in U.S. imports.

Email: fernando.leibovici@nyu.edu, mwaugh@stern.nyu.edu. Fernando Leibovici gratefully acknowledges the financial assistance provided by the Social Sciences and Humanities Research Council of Canada (SSHRC).

1. Introduction

This paper studies a dynamic extension of international trade models and its ability to explain the behavior of trade volumes at business-cycle frequencies where static models fail. Static trade models have difficulty explaining the response of trade volumes to changes in income and prices.[1] For example, the empirical elasticity of imports to measures of output is well over one; standard models predict a unitary income elasticity. The empirical elasticity of import volumes to measures of relative prices is well below one; calibrations of standard models use values well above one. Accounting exercises using static trade models to measure deviations between predicted and observed fluctuations in imports find these deviations to be pro-cyclical.[2]

We show how introducing dynamics into the import decision can help trade models rationalize these features of the data. The extensions we focus on are a time-to-ship friction to import goods and finite intertemporal elasticity of substitution. The time-to-ship friction makes the importing decision dynamic because resources today must be sacrificed for the delivery of goods tomorrow. With a finite intertemporal elasticity of substitution, the rate at which agents are willing to substitute across time—intertemporal marginal rate of substitution—depends on the trade-off between consumption today versus expectations of consumption tomorrow. Our insight is that variation in the intertemporal marginal rate of substitution with changes in income and relative prices will break the unitary income elasticity, bias the estimated price elasticity relative to static trade models, and show up as a time-varying trade friction.

We formalize these ideas by building a pure exchange, small open-economy model. In the model, there are a large number of countries. An agent within a country receives a stochastic endowment of its own nationally differentiated good. Agents have time-separable preferences of the constant relative risk aversion class over an aggregate consumption good. The aggregate consumption good is a composite of the nationally differentiated goods, and the aggregator of the goods is of the constant elasticity of substitution class. Agents take prices as given, and we model the evolution of relative prices as following a stochastic process. The only friction that agents face is that they must commit resources today for the delivery of imported goods in subsequent periods.

We designed our model to answer a specific quantitative question: Given a stochastic process describing income and prices, in U.S. data, how do the time-to-ship friction and finite intertemporal elasticity of substitution shape the decision to import? To answer this question, we esti-

[1] Examples of models of this type are those of Krugman (1980), Eaton and Kortum (2002), Anderson and van Wincoop (2003), and Melitz (2003). This also includes international real business-cycle models as summarized in Backus, Kehoe, and Kydland (1995).

[2] Houthakker and Magee (1969) is the seminal reference about the high income elasticity for the U.S. and other countries at long-run frequencies. Jacks, Meissner, and Novy (2009) and Levchenko, Lewis, and Tesar (2010a), building on the insights of Chari, Kehoe, and McGrattan (2007) emphasize the third fact's role in accounting for the decline in trade during the 2008-2010 recession.

mate the stochastic process for endowments and prices from U.S. data. We use standard values to calibrate the intertemporal elasticity of substitution and the elasticity of substitution across goods. We then simulate the model and estimate implied income and price elasticities using the simulated data.

The main result is that our model can quantitatively deliver the high income elasticity and low price elasticity, consistent with U.S. time series data. We find that the income elasticity in our model is well over one, and the price elasticity is well below the underlying elasticity of substitution for plausible values of the intertemporal elasticity of substitution. For example, with an intertemporal elasticity of substitution of 0.20 and an elasticity of substitution of 1.5, the income and price elasticities are 1.88 and -0.48; in the data, they are 2.00 and -0.48. Performing the same exercise, but removing the time-to-ship friction or shutting down variation in the intertemporal marginal rate of substitution, we find the estimated income elasticity is effectively one and the price elasticity is the same as our calibrated elasticity of substitution.

Our results have much to say about the large drop in trade during the 2008-2009 crisis.[3] Given that our model that can account well for features U.S. imports at business-cycle frequencies, we ask what are our model's implications for the 2008-2009 crisis? To do so, we apply the Kalman smoother to our model, using U.S. data on income and prices to generate a predicted series of imports. We show that for plausible values of the intertemporal elasticity of substitution, our model explains a majority of the collapse in U.S. import data (see Figure 7). In contrast, the model without the time-to-ship friction or without variation in intertemporal marginal rate of substitution performs poorly. Our theory is simple: a shock unexpectedly reduced output, and agents became unwilling to substitute (on the margin) across time periods. Because international trade is time-intensive, imports declined more than absorption resulting, in a "trade collapse."

The distinguishing feature of our theory relative to the trade collapse literature, is that it does not rely on specifics about the 2008-2009 crisis but applies broadly across time periods. Trade finance, input-output linkages, vertical specialization, compositional features, and inventories played some role in generating a decline in trade during the 2008-2009 crisis. However, over the past 40 years, imports responded about two times as much as income. It remains to be shown that these mechanisms can account for these cyclical fluctuations beyond the 2008-2009 crisis. Our theory meets this challenge by accounting well for the cyclical fluctuations of U.S. imports over the last 40 years (see Figure 6(a)).

Trade elasticities play critical roles in formulating predictions and recommendations for policy

[3]Much attention has focused on this episode. See, for example, the papers in Baldwin (2010); Alessandria, Kaboski, and Midrigan (2010b); Amiti and Weinstein (2011); Bems, Johnson, and Yi (2010); Chor and Manova (2010); Eaton, Kortum, Neiman, and Romalis (2010); Jacks, Meissner, and Novy (2009); Levchenko, Lewis, and Tesar (2010a).

makers with regard to issues such as the effects of a one-percent depreciation of the U.S. dollar (see, e.g., the discussion in Hooper, Johnson, and Marquez (2000)). This has generated a large literature on the estimation of income and price elasticities for imports (see, e.g., Marquez (2002) for a discussion of this research). The key questions surrounding the empirical estimates of trade elasticities concern their stability (and, hence, their usefulness in making predictions) and the disconnect with the predictions of standard trade models.

We contribute to this literature by providing answers to these open questions. First, we rationalize the disconnect between empirical trade elasticities and standard trade models by introducing dynamics into the import decision. Moreover, while our model does not have constant price and income elasticities, it retains the parsimony and performance of statistical models. These two features—theoretical consistency and statistical performance—suggest that our model can contribute to answering certain forecasting and policy questions.

2. Cyclical Features of International Trade Volumes

In this section, we outline some key features of cyclical fluctuations in imports, income, prices and their co-movement in U.S. time series data. The data features we describe are not entirely new; for example see Houthakker and Magee (1969) on the income elasticity of trade at long-run frequencies; Ruhl (2008) on the low price elasticity; and Jacks, Meissner, and Novy (2009) and Levchenko, Lewis, and Tesar (2010a) on the wedge analysis. However, summarizing all three features of the data is important because we will argue one underlying mechanism can rationalize all three features of the data.

Before proceeding, it is worthwhile to outline our language conventions. First, while we use the term "elasticity," the estimates we discuss are best thought of as simply summarizing the statistical properties of how imports, income, and prices behave in the time series. Second, although the measure of output/income we focus on is absorption, we will use the terms income, output, and absorption synonymously throughout.[4]

To summarize the statistical properties of imports, income, and prices in U.S. time series data, we use a log-linear relationship relating imports to prices and income. The rationale for using this relationship comes from standard models of international trade based on CES preferences

[4] Absorption is gross domestic product plus imports minus exports. Because static trade models typically impose balanced trade, absorption corresponds with income. Hence, we will use absorption and income synonymously.

or production functions.[5] These models have the following demand function for imports:

$$\log M_t = -\theta \log\left(\frac{p_{m,t}}{P_t}\right) + \log Abs_t + \omega_t. \tag{1}$$

This equation relates real imports M, real absorption Abs, the price of imports p_m, and the absorption price index P, in a log-linear way. The parameter θ controls the price elasticity of imports, ω_t is a "wedge," which we will describe in more detail below.

We use data and the structure of equation (1) to perform two exercises to summarize meaningful features of the data. The first exercise runs the regression

$$\log M_t = \alpha \log\left(\frac{p_{m,t}}{P_t}\right) + \beta \log Abs_t + \epsilon_t. \tag{2}$$

Relative to equation (1), the coefficient α measures the empirical price elasticity, and β measures the empirical income elasticity in the context of standard trade models.

The second exercise imposes the theoretical restrictions implied by equation (1), a unit income elasticity and an assumed value for the price elasticity, and it then infers the wedge ω by comparing predicted imports versus actual imports. Specifically, the wedge is computed as

$$\omega_t = \log M_t - \left(-\theta \log\left(\frac{p_{m,t}}{P_t}\right) + \log Abs_t\right). \tag{3}$$

This exercise is similar to that of Jacks, Meissner, and Novy (2009), and Levchenko, Lewis, and Tesar (2010a). Following the arguments of Chari, Kehoe, and McGrattan (2007), this exercise is meaningful because *systematic* deviations between theory and data shed light on mechanisms through which underlying primitives operate. Specifically, if ω_t varies systematically with the business cycle, then this suggests that: (i) that there are economic forces that are not reflected in equation (1); and (ii) that any new mechanism posited to explain these deviations should operate through the wedge. We set $\theta = 1.5$, which is a standard international business-cycle calibration of this parameter. Using larger θs as in trade calibrations result in larger wedges.

2.1. Measurement Issues

There are several issues in constructing data for use in the regression in (2) and wedge analysis in (3). They are: (i) the appropriate definitions of imports and absorption; and (ii) how to construct the appropriate real measures and the associated price indices. Because these are

[5] We think of standard models as those that generate log-linear import demand functions also known as gravity equations. Examples of standard models are those of Krugman (1980), Anderson and van Wincoop (2003), Eaton and Kortum (2002), and Melitz (2003) or international business-cycle models such as Backus, Kehoe, and Kydland (1995).

important issues, we will spend several paragraphs here describing the construction of our data series.

We focus our analysis on imports and absorption of goods, excluding oil. The National Income and Product Accounts (NIPA) reports measures of imports and exports of goods and GDP coming from goods sales. Appendix A provides the details of the exact data series that we use.

The focus on goods GDP helps address compositional issues of the sort emphasized by Eaton, Kortum, Neiman, and Romalis (2010). They argue that because manufacturing output declined more than aggregate GDP during the 2008-2009 recession and that most trade is in manufactures, these facts account for a large fraction of the drop in trade. To address these compositional issues, we focus on an absorption measure where most trade occurs (goods only component of GDP).[6] To address compositional issues within goods (i.e., durable vs. non-durable) emphasized by Boileau (1999) and Engel and Wang (2009), we perform the same analysis later in this section using only durable or non-durable goods.

Constructing real measures of these activities and the associated price indices is not as straightforward as it might seem. Real values in the U.S. NIPA accounts are chain-type indexes using an "ideal" chain index advocated by Fisher (1922). While these indexes have desirable properties, they are not additive across categories (see Ehemann, Katz, and Moulton (2002) and Whelan (2002) for detailed discussions). For our purposes, the implication is that one cannot compute real absorption simply by adding real goods GDP to real imports and subtracting real exports. An (approximate) solution to this problem is to use a "Fisher of Fishers" approach suggested by Diewert (1978). The basic idea is to take the real values and the associated price indexes for the categories of interest and then compute Fisher indexes of these measures—hence the "Fisher of Fishers" name.

Using this approach, we construct data series for real absorption of goods, real imports of goods, and the associated price indexes starting in the second quarter of 1967 and ending in the third quarter of 2011. To deal with trends, we HP-filtered the logarithm of these data with smoothing parameter 1600. Results using log-first-differences yielded no significant differences.

2.2. High Income Elasticity, Low Price Elasticity, Pro-Cyclical Wedges

Table 1 presents the results of the estimated income and price elasticities of imports using ordinary least squares to estimate the regression in (2). Figure 2(a) plots the results from the wedge exercise. Below we outline three observations.

[6] Eaton, Kortum, Neiman, and Romalis (2010) find that the decline in manufacturing accounts for 80 percent of the drop in trade (exports + imports) relative to GDP which implies an elasticity of 1.25. However, comparisons are very difficult since they focus on total trade rather than on imports, relative to contemporaneous world trade, and on four quarter log changes rather than on quarterly deviations from trend.

Table 1: Empirical Price and Income Elasticities

Data	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$	R^2	# Obsv.
Goods GDP	-0.48 (0.11)	2.00 (0.14)	0.65	178

Note: Data are in logs and HP filtered over the time period from Q2 1967 to Q3 2011. Heteroskedastic robust standard errors are in parenthesis.

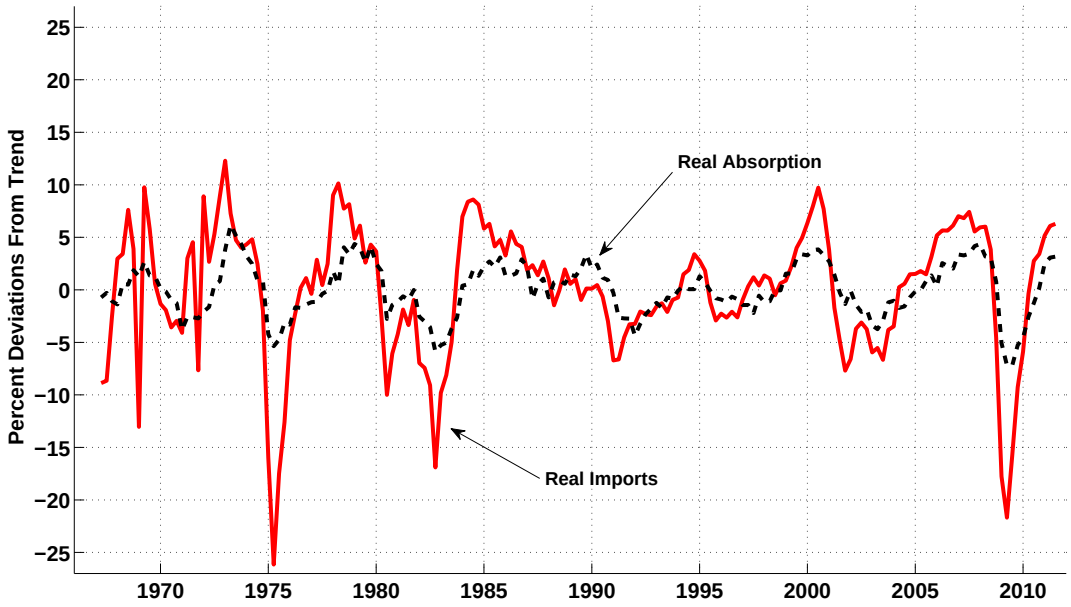
O.1. Income elasticity > 1. The estimated income elasticity of imports is nearly two—i.e., a one-percent increase in absorption results in a two-percent increase in imports. Figure 1(a) illustrates this finding by plotting the percent deviations from trend of real absorption and imports. Consistent with the findings in Table 1, absorption correlates strongly with imports, yet it is less than half as volatile.

This feature of the data is interesting because it contrast with what standard trade models imply. These models imply that a one-percent increase in absorption results in a one-percent increase in imports—i.e., these models have a unit income elasticity. Thus, attempts at modeling trade volumes at business-cycle frequencies must confront the discrepancy between standard trade models and the data.

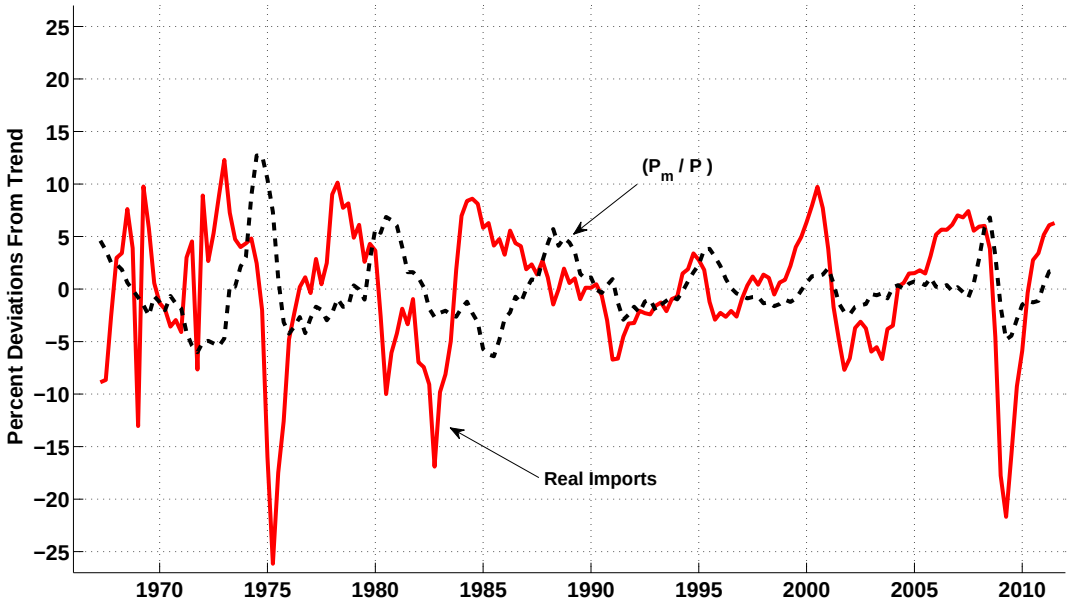
As mentioned, the fact that the estimated income elasticity of demand for U.S. imports exceeds unity is a well-known fact, dating back to Houthakker and Magee (1969). Marquez (2002) examines this feature of the data from a modern perspective and finds that it is intractable to alternative econometric specifications, different frequencies, and commodity disaggregation. The key feature we wish to emphasize here is that it is found at short-run or business-cycle frequencies. The distinction regarding the time-horizon is important because prominent explanations of the high income elasticity of import demand are based on expanding product variety (see, e.g., Krugman (1989) and Feenstra (1994)) and are best though of as medium-/long-run explanations.[7]

O.2. Low price elasticity. The second point is that the estimated import price elasticity is -0.47. Figure 1(b) illustrates this finding. It plots the percent deviations from trend of relative prices $\frac{p_{m,t}}{P_t}$ and import data. Notice that prices and imports weakly correlate with each other negatively and, in some instances, even move in the same direction. Thus, the low price elasticity in Table 1 is not a surprise.

[7]Quantitatively, Feenstra (1994) shows that the expanding product variety explanation can account for only part of the high income elasticity. Ruhl (2008) shows that this margin is not quantitatively important at business-cycle frequencies. Kehoe and Ruhl (2009) measure changes in the extensive margin and find that it plays little role outside of significant structural transformations or trade liberalization.



(a) Real Absorption and Real Imports



(b) Relative Prices and Real Imports

Figure 1: Absorption, Relative Prices, and Import Data

There is nothing inherently wrong with a low price elasticity. Modelers have a choice over this parameter. However, this estimate is low relative to typical calibrations/estimations of static trade models or international business-cycle models. Estimates using static trade models and changes in trade flows during trade liberalizations typically suggest substantially higher values. Lower values typically come from imposing a unitary income elasticity and using time series variation in prices and trade flows relative to absorption. Ruhl (2008) provides an extensive discussion of the conflicting estimates of this elasticity.

O.3. Pro-cyclical wedges. The final point is that the wedges inferred using equation (3) are pro-cyclical and explain much of the variation in imports.

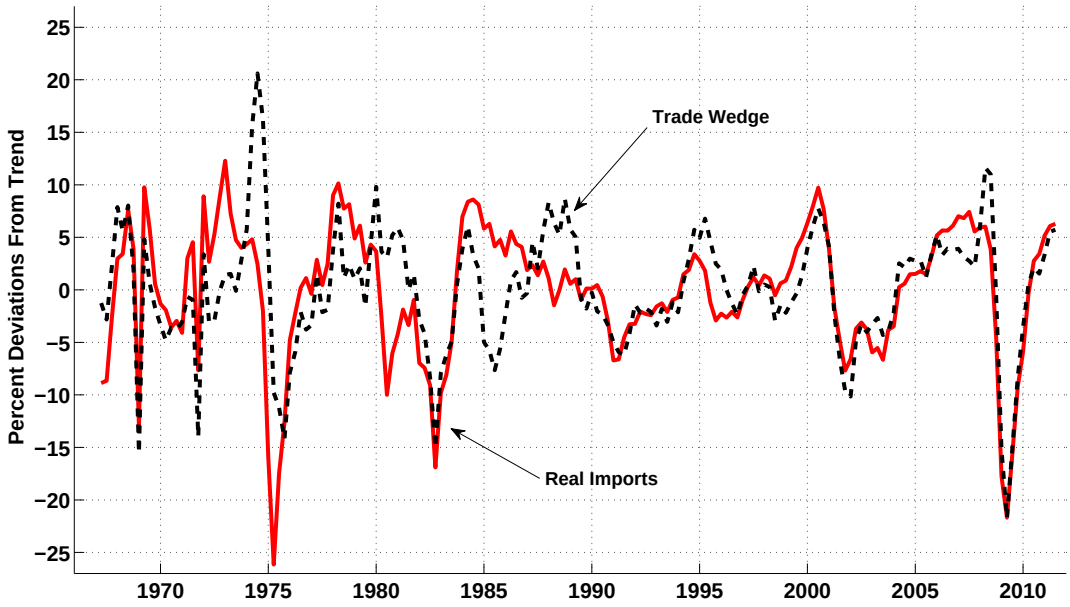
Figure 2(a) simply plots the wedge versus import data. For most of the time period, the wedge tracks imports very closely. Confirming this, a regression of imports on the wedge yields a slope coefficient of 0.70 and an R^2 of 0.44. This suggests that systematic variation in the wedge is quantitatively important to understanding variation in imports. Given the discussion above, this observation suggests that there are economic forces that are operative in the data but not in models consistent with equation (1).

Systematic variation in the trade wedge is not distinct from observations **O.1** and **O.2**. Standard trade models basically have stronger substitution effects relative to income effects—i.e., imports should be more responsive to a one-percent change in prices relative to a one-percent change in income. The data observations **O.1** and **O.2** suggest the complete opposite pattern—imports are less responsive to prices relative to income. Thus, the wedge analysis based on a model that puts more weight on relative price changes versus changes in income is bound to find systematic variation in the trade wedge.

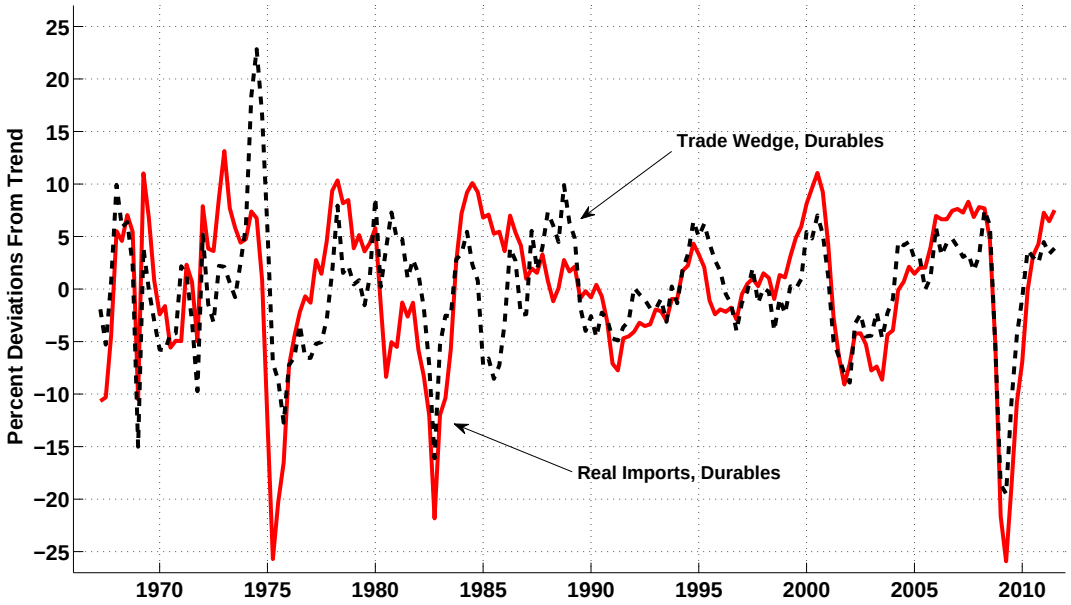
The 2008-2009 crisis illustrates this point well. During this period, absorption decreased and imports decreased even more—this reflects the high income elasticity. When the income elasticity is constrained to be one in the wedge analysis, the wedge must then decrease to rationalize the drop in imports. Furthermore, relative prices decreased and imports did not increase as predicted by the standard model—this reflects the low price elasticity. When the price elasticity is constrained to take on a standard value, this implies that the wedge must decrease even more. Thus, the fact that imports over-respond to income and under-respond to prices manifests itself as a pro-cyclical wedge when events like the 2008-2010 recession are analyzed in the context of a standard trade model.

2.3. Robustness—Durables and Inventories

Recent papers on the collapse in trade during the 2008-2009 crisis have raised two issues: the distinction between durables and non-durables and inventories. Here, we show that observations **O.1-3** are robust to restricting attention to durables or non-durable goods, and accounting



(a) Wedges and Real Imports



(b) Wedges and Real Imports, Durables Only

Figure 2: Wedges and Import Data

for the behavior of inventories.

Durables vs. Non-Durables. One concern is that observations **O.1-3** are just picking up compositional effects of the sort described by Boileau (1999) and Engel and Wang (2009). The argument is that a larger fraction of imports is classified as durables than in, say, absorption of total goods. This idea, combined with the fact that consumption of durables is more volatile than that of nondurables, suggests that an income elasticity larger than unity or pro-cyclical trade wedges may arise because of the compositional difference. Therefore, the durables composition explanation suggests that if we focused on only durables or non-durables, then observations **O.1-3** would disappear.

We addressed this argument by restricting the import, absorption, and price data to include only durable goods or only non-durables and by re-estimating equation (2). Appendix A provides the details of the exact data series that we use.

Table 2 presents our results. It shows that the income elasticity of imports is well above one (albeit mitigated), and the import price elasticity is similar to that found in Table 1. When the data are restricted to only non-durables, one finds a very high income elasticity and low price elasticity.[8]

The durable trade wedge still accounts for a lot of the variation in imports. Figure 2(b) illustrates this by plotting the trade wedge for durable goods only. Similar to the results discussed in **O.3**, a regression of imports on the wedge yields a slope coefficient of 0.75 and an R^2 of 0.40. The reasoning is the same as discussed above: though the data suggests that imports of durables are more responsive to income than to changes in relative prices, standard models predict the opposite pattern. Thus, while the income elasticity is mitigated, the relative weighting of the income and substitution effects conflict with what the data suggest.

Further evidence against the durables explanation is the following: we should observe that the income elasticity of imports should be higher when share of durables in imports is large relative to the share of durables in total—i.e., when there is a large compositional difference between imports and absorption. Thus, we should observe a positive correlation between the income elasticity and the share of durables in imports relative to the share of durables in total.

To explore this implication, we ran the regression in (2) for all goods on a 40-quarter moving window—i.e., for 1967q2-1977q1, 1967q3-1977q2, etc. Figure 3 plots the estimated income elasticities on the y-axis. The x-axis reports the average share of durables in imports relative to the share of durables in total for the same 40-quarter window.

Figure 3 displays a strong negative correlation (-0.84 and statistically different from zero) be-

[8]Some care must be taken with this observation because available measures of non-durable imports includes petroleum products, unlike other data series.

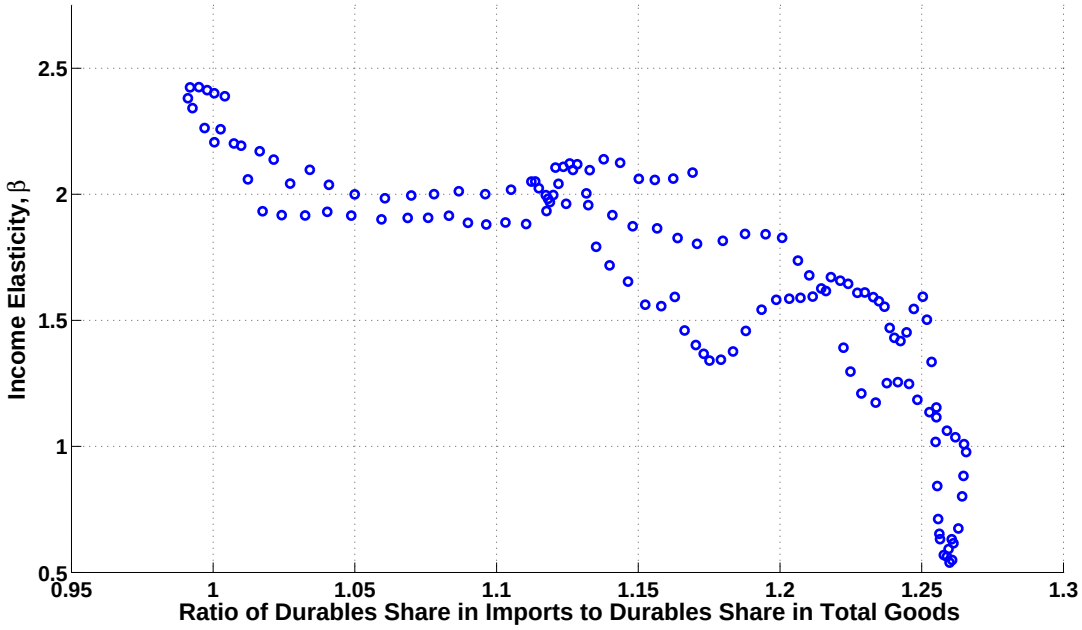


Figure 3: Income Elasticity vs. Share of Durables in Imports

tween the income elasticity and the share of durables in imports.[9] Notice that when the share of durables in imports was the same as the share of durables in total—i.e., when there is no compositional difference—the income elasticity is at its peak. In contrast, when the share of durables in imports was larger than the share of durables in total—i.e., when there is a compositional difference—the income elasticity is near at its lowest point. Figure 3 appears to be in contrast with the positive correlation that the durables explanation implies.

Inventories. Another concern is that **O.1-3** arises because we abstract from changes in inventories.[10] Alessandria, Kaboski, and Midrigan (2010b) make this argument while studying the decline in trade flows during the 2008-2009 crisis; Alessandria, Kaboski, and Midrigan (2010a) argue that inventory considerations are important for understanding the dynamics of devaluations. An implication of Alessandria, Kaboski, and Midrigan's (2010b) model is that the regression equation in (2) should be augmented with the change in imported inventories (Alessandria, Kaboski, and Midrigan (2011) provide this derivation).

We followed this argument by augmenting the regression in (2) by including data on the real change in private inventories as an additional explanatory variable. Separate information on

[9] Engel and Wang (2009) emphasize the volatility of imports and the role of durables; the correlation between the volatility of imports of goods and the share of durables in trade is -0.92, and the correlation between the relative volatility of imports to absorption and the share of durables is -0.81. Both are statistically different from zero.

[10] Feenstra (1994) suggests this, as well, and uses real personal consumption to instrument for the fact that changes in inventories are not controlled for. We did this and found that the estimated price and income elasticities are -0.52 and 2.34.

Table 2: Empirical Price and Income Elasticities — Durables and Inventories

Data Series / Approach	Price Elasticity	Income Elasticity	Inventory Elasticity	R^2	# Obsv.
Durable Goods	-0.49 (0.12)	1.37 (0.10)	—	0.64	178
Non-Durable Goods	-0.17 (0.04)	2.97 (0.28)	—	0.50	178
Goods & Δ Inventories	-0.45 (0.10)	1.66 (0.11)	0.22 (0.02)	0.74	178

Note: Data are in logs and HP filtered over the time period from Q2 1967 to Q3 2011. Heteroskedastic robust standard errors are in parentheses.

changes in inventories of imported goods is unavailable. The third row in Table 2 reports the results. After controlling for changes in inventories, the income elasticity is 1.66, relative to 2.00 without controlling for inventories. Including inventories also improves the fit of the regression from an R^2 of 0.63 without inventories to 0.74 with inventories. These results suggest that inventory adjustments are a partial, but not a complete explantation, of the high income elasticities observed at cyclical frequencies.

In the next sections, we argue that systematic variation in intertemporal substitution can rationalize the features, **O.1-3**, of the data described above. Intertemporal substitution matters because we model the decision to import as a dynamic decision. With a finite intertemporal elasticity of substitution, intertemporal substitution depends on the trade-off between consumption today versus expectations of consumption tomorrow. Changes in endowments and prices affect the trade-off between consumption today and consumption tomorrow, and, thus, this will break the unitary income elasticity, bias the estimated price elasticity, and show up as a time-varying trade friction consistent the features **O.1-3** seen in the data.

3. Model

The world economy consists of a large number of countries of two types, home and foreign. In each country, there is an infinitely-lived representative consumer who has time-separable preferences over a period utility function. Period utility is of the constant relative risk aversion class and defined over a composite consumption good Q_t defined below. Expected discounted

future utility is given by:

$$\mathrm{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \frac{Q_t^{1-\gamma}}{1-\gamma} \right\}, \tag{4}$$

where $\gamma > 0$ and $\beta \in (0,1)$. The parameter γ equals the coefficient of risk aversion, and $\frac{1}{\gamma}$ equals the elasticity intertemporal of substitution over the composite consumption good Q_t. The parameter β is the subjective discount factor. E_0 is the mathematical expectation operator conditional on information at date zero.

The composite consumption good Q_t is a CES aggregate over two goods, x and y:

$$Q_t = \left(x_t^{\rho} + y_t^{\rho}\right)^{1/\rho}, \tag{5}$$

where $\rho \in (0,1)$. The parameter ρ controls the elasticity of substitution across goods, with this elasticity defined as $\theta = \frac{1}{1-\rho}$.

Every period, the consumer in a typical home country receives a stochastic endowment of good x, while the consumer in a typical foreign country receives a stochastic endowment of good y. These endowments are idiosyncratic to each particular country. These endowments can be either consumed domestically or used to acquire foreign goods through a centralized goods market. Goods cannot be re-exported.

International trade is the exchange of one type of good for another type of good. International trade is subject to two technological constraints. First, agents in each country face iceberg trade costs, $\tau > 1$, to move goods across borders. This implies that for every τ units of a good that are shipped, only one unit arrives at a destination.

Second, international purchases are subject to time-to-ship.[11] We model this such that if the home country purchases one unit of good y at date t, the good arrives (and is only available for consumption) at date $t+1$.[12] Along with this assumption, we assume that goods must be paid for before they are delivered.[13] Underlying this international trading structure is an assumed enforcement technology that allows countries to coordinate the dynamic exchange of goods across international borders.

[11] Empirically, Hummels (2007), Djankov, Freund, and Pham (2010),Alessandria, Kaboski, and Midrigan (2010a) and Hummels and Schaur (2012) document the time-intensive nature of shipping goods across borders.

[12] In section 4.3 we loosen the strict one period time-to-ship assumption.

[13] In reality, a variety of payment arrangements are used. The payment structure in the model is known as "cash in advance." See Capela (2008) for a description of how international transactions take place. At the aggregate level, it is hard to find evidence about the composition of payment structures of international trade. Antras and Foley (2011) provide evidence from a large U.S.-based exporter that cash in advance accounts for a large share of international transactions.

We abstract from trade in international financial assets.[14] This does not imply that the typical agent is in financial autarky. Because of the timing assumption on international trade, trade is equivalent to purchasing a one-period non-state contingent bond that pays out in units of the foreign good in the next period. Thus, international trade acts implicitly like a financial asset in this model. With this insight and noting that each country faces idiosyncratic risk in endowments, this model is closely related to one-good incomplete market models of Huggett (1993) and Aiyagari (1994), and specifically to Castro's (2005) reinterpretation of these models to countries rather than to individuals.

With these assumptions, the consumer in a typical home country faces the following budget constraint:

$$p_{xt}x_t + \tau p_{yt}y_{t+1} \;\leq\; p_{xt}z_t, \tag{6}$$

where p_{yt} and p_{xt} are the international prices that the typical consumer faces. The term z_t is the idiosyncratic endowment of the x good that the consumer in the typical the home country receives. The term x_t is the amount of the home good purchased for consumption today; y_{t+1} is the amount of the foreign good purchased today for consumption in the next period.

Inspecting the budget constraint (6), one realizes that the value of orders of imports at date t (i.e., imports delivered at $t+1$) must equal the value of exports shipped at date t. Thus, our model has a "dynamic trade balance" condition holding for every pair of contiguous time periods:

$$\underbrace{\tau p_{yt}y_{t+1}}_{\text{Imports}_{t+1}} = \underbrace{p_{xt}(z_t - x_t)}_{\text{Exports}_t}. \tag{7}$$

This observation is interesting because trade is not balanced period by period as in static trade models, and, thus, our model will have a non-zero current account. In the model, the current account equals the difference in orders of imports and arrivals of imports. Consistent with the idea of international trade being similar to asset trade, a positive current account acts as an increase in savings since consumption of the home good is forgone today for consumption of the foreign good tomorrow.

Given the description of the model, we will now focus on the problem of a consumer located in

[14] Evidence on the proper international asset structure to allow agents to trade is limited, and what is available points towards financial autarky; see, for example, Heathcote and Perri (2002).

a particulary home country. The consumer faces the following dynamic programming problem:

$$V(S,y) \;=\; \max_{x,y'} \left\{ \frac{(x^\rho + y^\rho)^{(1-\gamma)/\rho}}{1-\gamma} \;+\; \beta \mathrm{E}[\; V(S',y') \mid S \,] \right\}, \tag{8}$$

$$\text{subject to} \quad p_x x + \tau p_y y' \leq p_x z,$$

$$S' = g(S), \quad \text{and} \quad S = \{z, p_y, p_x\}.$$

There are four state variables in the problem. The first three state variables are the consumer's idiosyncratic endowment realization z and aggregate prices p_y and p_x which are summarized in the vector S. These evolve according to a law of motion $g(S)$ which we will discuss more below. The consumer takes the aggregate prices p_{yt} and p_{xt} as given and forms expectations of future prices based upon knowledge of $g(S)$. The other state variable is last-period's orders of imports arriving this period for consumption. Given the state variables, the consumer chooses the quantity of good x to be consumed this period and chooses the quantity of good y' to order internationally for consumption in the following period.

We model the law of motion $g(S)$ for a country's idiosyncratic endowment z_t in units of good x and the prices p_{yt} and p_{xt} as following a stationary VAR(1) process:

$$\log S_t = A \log S_{t-1} + \nu_t, \tag{9}$$

where $S_t = \{z_t, p_{yt}, p_{xt}, \}$ and the innovations ν_t are jointly normally distributed with mean zero and variance covariance matrix Σ. This parameterizes the law of motion $g(S)$ in (8).

The idea here is to directly model endowments and prices as following a stochastic process and estimate it from data rather than solving for the recursive competitive (general) equilibrium as in Huggett (1993), Aiyagari (1994), or Castro (2005). The motivation for this assumption is our desire to answer the following quantitative question in a simple and straightforward way: If a typical agent faces an endowment and price process like the one we observe in the data, then what are the implications for imports? Directly specifying a stochastic process over prices and estimating it from data allows us to answer this question in a straightforward manner.

Because we are not interested in world-economy outcomes or performing counterfactuals, this approach allows us to sidestep a more involved alternative. The aggregate state variables p_y and p_x depend on the aggregate supply of the idiosyncratic endowment realizations of good x and good y across all countries. Variation in prices over time are a result of changes in the relative aggregate supply of each good. The recursive competitive (general) equilibrium is a value

function, policy function, and a perceived law of motion that satisfies (8) and the perceived law of motion corresponds with the equilibrium law of motion, subject to market clearing conditions. Thus, to solve for the general equilibrium we would be looking for a fixed point mapping a perceived law of motion into the actual law of motion.

3.1. Qualitative Features

In this section, we present some qualitative features of the model and the intuition behind the mechanism that can deliver a higher-than-unity income elasticity, a price elasticity lower than θ, and terms that look like wedges co-varying with fluctuations in endowments and prices.

Dynamic Import Demand. The key relationship in our model is the dynamic demand function for imports. After solving the representative consumer's problem, the demand for imports in the home country can be written as:

$$\frac{p_{xt}z_t}{P_t}\left[\frac{\tau}{E_t(\tilde{m}_{t+1})}\frac{p_{yt}}{P_t}\right]^{-\theta} = y_{t+1} \quad \text{and} \quad \tilde{m}_{t+1} = \beta\left(\frac{Q_{t+1}}{Q_t}\right)^{-\gamma+\frac{1}{\theta}} \tag{10}$$

$$P_t = \left[p_{xt}^{1-\theta} + \left(\frac{\tau p_{yt}}{E_t(\tilde{m}_{t+1})}\right)^{1-\theta}\right]^{\frac{1}{1-\theta}}.$$

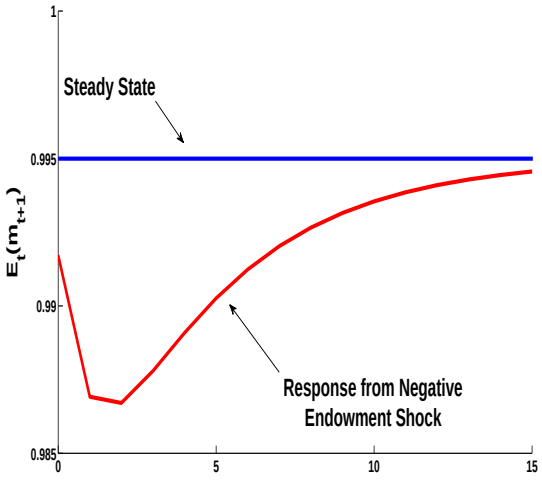
There are several points to note about equation (10). First, the timing, endowments and prices at date t affect imports consumed/delivered at date $t+1$. This is in contrast to the contemporaneous effects of endowments and prices on the import decision in equation (1) that comes from a standard static trade model. This is a direct result of the time-to-ship assumption.

Second, the term $E_t(\tilde{m}_{t+1})$, enters equation (10). The term $\tilde{m}_{t+1}$ is an implicit function of the standard intertemporal marginal rate of substitution (IMRS) in theories of asset prices and consumption with CRRA preferences.[15] This term induces the demand for imports depends on the agent's willingness to substitute consumption today (i.e. spending less on home goods today) for consumption tomorrow (i.e. imports arrive tomorrow); the term $E_t(\tilde{m}_{t+1})$ reflects this valuation.

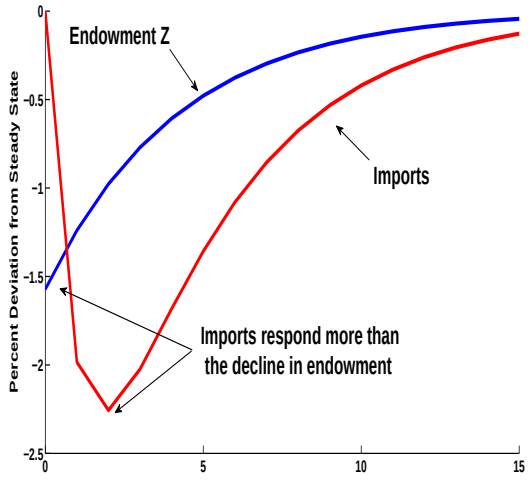
An important difference between $\tilde{m}_{t+1}$ and the IMRS is that $\frac{1}{\theta}$ shows up along with the parameter γ.[16] The reason is that because x is not a perfect substitute for y; thus, the elasticity of substitution across goods shows up with intertemporal elasticity. The special case when x and y are perfect substitutes (implying $\theta = \infty$) illustrates this point, as only the parameter γ shows

[15] To keep terminology and notation clear, the IMRS is $m_{t+1} = \beta\left(Q_{t+1}/Q_t\right)^{-\gamma}$—i.e., the marginal rate of substitution of aggregate consumption. The term $\tilde{m}_{t+1}$ is an implicit function of m_{t+1}.

[16] A more subtle point is that time-to-ship works similarly to habits as decisions yesterday impact utility today (see Campbell and Cochrane (1999)). A key difference is that these decisions are internalized, unlike external habit.



(a) Response in $E_t(\tilde{m}_{t+1})$ to a shock to z_t



(b) Response in Imports to a shock to z_t

Figure 4: Imports Respond More Than Endowments.

up in $\tilde{m}_{t+1}$. In contrast, when the elasticity of substitution equals the intertemporal elasticity, $\theta = \frac{1}{\gamma}$, variation in consumption across time periods plays no role and $\tilde{m}_{t+1} = \beta$.

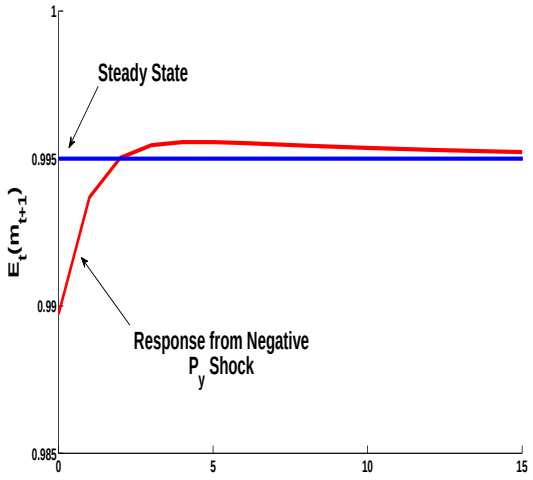
Finally, note that the IMRS will change with the endowments and prices. Thus, systematic variation in the IMRS leads to systematic variation in $E_t(\tilde{m}_{t+1})$, and the model's income elasticity will generally differ from one and the price elasticity will generally differ from θ. Moreover, because $E_t(\tilde{m}_{t+1})$ shows up in the same places as the trade friction, and, thus, variation in the IMRS has the ability to look like a time-varying trade wedge.

Below, we describe how this mechanism can deliver an income elasticity greater than unity and an artificially low price elasticity. We focus on the relevant case of the parameter space ($\gamma - \frac{1}{\theta} > 0$) and with independent, AR(1) processes describing the evolution of endowments and prices.

Income Elasticity Greater than Unity. Our economy can generate an income elasticity of imports that is greater than one. To gain intuition regarding this result, assume that the economy is initially in steady-state and then hit by a negative shock to endowment z_t at date t.

As a result of the shock, two forces work to affect imports. First, a negative shock to z_t directly lowers imports tomorrow since expenditures decreased. This is the standard force in static trade models. Second, a negative shock decreases the agent's willingness to substitute across time, pushing $E_t(\tilde{m}_{t+1})$ down, further lowering imports. This is where dynamics bite. Because of the decline in the endowment *and* the decline in $E_t(\tilde{m}_{t+1})$, imports drop more than proportionally to the drop in the endowment.

The dynamic force deserves more explanation. Because the endowment is scarce at date t,



(a) Change in $E_t(\tilde{m}_{t+1})$ to a shock to p_{yt}



(b) Change in Imports to a shock to p_{yt}

Figure 5: Imports Respond Less Than Prices.

marginal utility today is high. Moreover, because z_t is below its steady-state value, the agent's rationally expect the endowment to be relatively higher tomorrow—i.e., $E_t(z_{t+1}) > z_t$. A relatively higher level of endowment means more consumption tomorrow, decreasing expected marginal utility tomorrow. Together, these forces imply that $E_t(\tilde{m}_{t+1})$ declines with a negative endowment shock. Because of how $E_t(\tilde{m}_{t+1})$ enters equation (10), a decrease in $E(\tilde{m}_{t+1})$ lowers imports more than the drop in the endowment does.

Figures 4(a) and 4(b) trace these arguments out. Figure 4(a) plots the decline in $E_t(\tilde{m}_{t+1})$, when the z_t is perturbed below its non-stochastic steady state and the economy is allowed to transit back. Figure 4(b) plots the corresponding response of imports and endowments. The figure shows that imports decline by more than the endowment shock. The difference is purely because of the change in agent's willingness to substitute across time in response to the endowment shock. In contrast, in a static model, imports would move by the exact same amount as the move in endowments.

Low Price Elasticity. Our economy has a price elasticity that is artificially below the elasticity of substitution θ, as well. To gain the intuition regarding this result, again assume that $\gamma - \frac{1}{\theta} > 0$ and that the economy is initially in steady state and then hit by a negative price shock to $p_{y,t}$ at date t.

The reason is that there are two forces working in different directions. First, a lower $p_{y,t}$ directly increases imports tomorrow by θ. This is the standard force in static trade models. Second, a negative price shock decreases $E_t(\tilde{m}_{t+1})$ pushing imports in the *opposite* direction of the first force. Because these two forces work in opposing directions, imports will increase by less than θ.

Why does a lower price of imports decrease $E_t(\tilde{m}_{t+1})$? This shock makes agents "wealthier" in the future because imports delivered tomorrow cost less. This leads to an increase in expected aggregate consumption and lower expected marginal utility in the next period, lowering $E_t(\tilde{m}_{t+1})$. Because of how $E_t(\tilde{m}_{t+1})$ enters equation (10), a decrease in $E(\tilde{m}_{t+1})$ is a force to decrease imports.

Figure 5(a) and 5(b) illustrate these arguments. Figure 5(a) plots the decline in the stochastic discount factor when the $p_{y,t}$ is perturbed below its non-stochastic steady state and the economy is allowed to transit back. Figure 5(b) plots the corresponding response of imports and prices. Here, the response in prices is the response of the term $(\frac{p_{yt}}{P_t})^{-\theta}\frac{1}{P_t}$ from (10). This is the object of interest because in the static model, imports should respond one for one with this term. In contrast, Figure 5(b) shows that imports rise less than the response of the price term. Again, the difference between imports and the price term arises because of the change in the agent's willingness to substitute across time in response to the price shock.

The previous discussions considered orthogonal shocks to endowments and prices and their effects. If there is non-zero covariance between endowment and prices shocks, then their dynamic effects on the intertemporal marginal rate of substitution could offset or reinforce each other. For example, if states of the world with low endowments are also likely to have high relative prices of imports, then both shocks would have offsetting effects on the intertemporal marginal rate of substitution, nullifying its effect on imports.

This observation is important for two reasons. First, it influenced our choice to model endowments and prices as a stochastic process and estimate it from the data. This allows us to isolate the role of intertemporal substitution with a stochastic endowment and price process that is estimated directly from the data. Second, our model's ability to replicate observations **O.1** and **O.2** is not predetermined. Because there exist covariance structures that can influence the results, the data provides important discipline on our models ability to capture cyclical features of import data.

4. Quantitative Analysis

In this section, we study the quantitative properties of our model. The quantitative question motivating our analysis is: How the time-to-ship friction and finite intertemporal elasticity of substitution shape the decision to import, given a stochastic process describing output and prices estimated from U.S. data? To answer this question we calibrate our model using standard parameter values and estimate an endowment and price process from U.S. data. We then perform three exercises.

First, we simulate time paths of imports, absorption, and prices from our model. We then use

equation (2) as a lens through which to examine co-movement in the simulated data. We then compare these results to data and to a model without the time-to-ship friction and a model without variation in the IMRS playing a roll.

Second, we introduce a more flexible time-to-ship technology. We recognize that there is heterogeneity in the speed of the international delivery of goods, (see, e.g., Hummels (2007) and Hummels and Schaur (2012)). This flexible time-to-ship technology allows us to study the quantitative implications of our model beyond the one-period time-to-ship technology.

Finally, we analyze our model's predictions for the collapse in trade during the 2008-2009 crisis and the time series of U.S. imports between 1967-2011. To do so, we apply the Kalman smoother to our model using U.S. data on absorption and prices to generate a predicted series of imports. Again, we then compare these results to the data and to the predictions from a model without the time-to-ship friction.

4.1. Calibration

We take a time period to represent one quarter, which implies that it takes one quarter for international goods to arrive. This assumption is stark, but in line with previous calibrations—see, e.g., Alessandria, Kaboski, and Midrigan (2010b). They motivate their choice based on evidence from Djankov, Freund, and Pham (2010), who show that the extra time to ship a good internationally is, on average, between 1.5 to two months. Amiti and Weinstein (2011) provide a nice discussion of this evidence and argue that trade finance leads to further time impediments. In Section 4.3, we loosen this assumption.

We explore different values of the parameter $\gamma = \{\frac{1}{\theta}, 2, 5, 10\}$, which controls the agent's intertemporal elasticity of substitution over aggregate consumption. As discussed, the case when $\gamma = \frac{1}{\theta}$ is interesting because this is when $\tilde{m}_{t+1}$ is constant and equals the subjective discount rate. We set the discount factor β equal to 0.995. This value implies a steady-state real annual interest rate of two percent which is consistent with the data on ex-post real returns on near-risk-free assets (see e.g., Tallarini (2000)).

The baseline value we use for the elasticity of substitution is 1.5, which is the standard value used in calibrations of international real business-cycle models, (see Backus, Kehoe, and Kydland (1995)). Note that in the dynamic model, the estimated price elasticity will differ from this value because of the arguments made in the previous section. Only in the static model without time-to-ship or the dynamic model without the IMRS playing a roll will this value correspond with the price elasticity. Appendix C reports the results with an elasticity of substitution equal to 4 which is at the low end of estimates from cross-sectional data or trade liberalization episodes (see e.g., the discussion in Simonovska and Waugh (2010)).

Finally, the trade friction is calibrated so that imports are 15 percent of absorption. The top

Table 3: Summary of Calibration

Preferences	
Time Period	A quarter; will loosen later
Trade Friction	Target Imports 15% of Absorption
Discount factor	$\beta = 0.995$
$1/$ IES	$\gamma = \{\frac{1}{\theta}, 2, 5, 10\}$
Elasticity of Substitution	$\theta = 1.50$
Estimated parameters for $\{z, p_y, p_x\}$ process	
Transition Matrix	$A = \begin{pmatrix} 0.80^{***} & -0.15^{***} & 0.05 \\ 0.17^{***} & 0.94^{***} & -0.25^{*} \\ -0.03 & 0.05^{***} & 0.75^{***} \end{pmatrix}$
Std. dev. of innovations	$\sigma_z = 0.015 \quad \sigma_{py} = 0.012 \quad \sigma_{px} = 0.004$
Corr. of innovations	$\text{corr}(z, p_y) = 0.055, \quad \text{corr}(z, p_x) = -0.25, \quad \text{corr}(p_y, p_x) = 0.26$

Note: Three stars indicate statistical significance at the 1 percent level; one star at the 10 percent level.

panel of Table 3 summarizes our calibration of preferences.

We estimate the stochastic process for endowments and prices (equation (9)) using HP filtered quarterly U.S. data on real absorption, absorption price index, and the import price index. Again, all data series include only the goods component of GDP and non-petroleum imports of goods. The rationale for choosing these series is that one can show that z corresponds to absorption (appropriately modified as discussed below) in the data and p_x with the associated price index.[17] Finally, import prices (appropriately modified as discussed below) inform us about p_y. Appendix A provides the details regarding the construction of these series, and Appendix B details the mapping between these measures in the data and the model.

An issue in the estimation of the stochastic process for endowments and prices is that we must modify standard data series to correspond with the timing in the model. Specifically, we need the variables in the data to reflect the timing at which they would be observed by the agent in the model, not the timing at which they are observed by U.S. statistical agencies. These agencies collect import data and the prices on arrival at the border. Yet, in the model, these prices are observed a quarter before. Thus, we adjust the data variables accordingly to make

[17] A second approach to measuring p_x uses the price index for domestic consumption. We found that the quantitative results were similar across this approach and the baseline.

them consistent with the timing of our model. Appendix B provides the details.

The bottom panel of Table 3 summarizes the estimated relationship between z_t, p_{yt}, p_{xt}.

Finally, in the simulations, we measure output from our model to conform with National Income and Product Accounts (NIPA). In all our simulations, we collect data from our model and compute quarterly chain-type quantity and price indexes for absorption.

4.2. Results

Table 4 presents the results. The first row replicates the empirical income and price elasticities seen in Table 1. The second through fourth rows report the results from our model for different values of γ. These are the cases when variation in the IMRS impacts the decision to import.

Table 4 shows that our model can deliver an income elasticity greater than unity. As the intertemporal elasticity of substitution increases, the income elasticities are 1.3, 1.9, and 2.7. In the case of $\gamma = 5$, the income elasticity is near the order of magnitude seen in the data. These results show that variation in the IMRS can be quantitatively strong enough to account for the high income elasticity of imports in the U.S. data outlined in observation **O.1**.

Table 4 also shows that the price elasticity lies below the true elasticity of substitution θ. Depending on the intertemporal elasticity of substitution, our model delivers a price elasticity of -1.16, -0.48, 0.45. All are meaningfully below the calibrated elasticity of substitution of -1.5. In fact, with $\gamma = 10$, changes in the IMRS more than offset changes in prices such that the price elasticity becomes positive at 0.45. Again, these these results show that variation in the IMRS is strong enough to rationalize the low price elasticity outlined in observation **O.2**—even though the true elasticity of substitution is 1.5.

These results suggests that the observations **O.1** and **O.2** are not independent features. In our model, systematic variation in the IMRS *simultaneously* generates a high income elasticity and a low price elasticity. Quantitatively, the success is quite striking: the same parameterization that generates an income elasticity that is close to the data is also the same parameterization generating an price elasticity close to the data. A priori, there was no reason to expect this outcome.

The second to the last row of Table 4 shows the results from the model retain the time-to-ship technology, but with intertemporal elasticity of substitution equal to the elasticity of substitution across goods. This is the case when variation in the IMRS does not affect the import decision. In this parameterization, the price elasticity corresponds with the elasticity of substitution θ and an income elasticity of unity, exactly like the static no time-to-ship model. This shows that the timing difference between the dynamic and static models is not driving the results. Instead, it is the the systematic variation in the IMRS that drives our findings.

Table 4: Elasticities from Data and Model, $\theta = 1.5$

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.48	2.00
Model, $\gamma = 2$	-1.16 [-1.19,-1.12]	1.31 [1.29, 1.33]
Model, $\gamma = 5$	-0.48 [-0.52,-0.43]	1.88 [1.84, 1.92]
Model, $\gamma = 10$	0.45 [0.35, 0.56]	2.74 [2.63, 2.86]
Model, $\gamma = \frac{1}{\theta}$	-1.51 [-1.55,-1.47]	1.03 [1.00, 1.05]
Model, no time-to-ship	-1.50 [-1.55,-1.44]	0.93 [0.85, 0.99]

Note: Results are averages from 250 simulations, with each simulation being 178 periods long; values in brackets report 95-percent confidence intervals. Section 2 describes the data.

The final row shows the results when time-to-ship is turned off. This is a standard trade model with an import demand equation corresponding with equation (1). Here, the income and price elasticities correspond with what the static model predicts—a price elasticity corresponding with the elasticity of substitution θ and an income elasticity of unity.

4.3. Intermediate Time-to-Ship

The one-period shipping technology may be a stark assumption, given that our model is calibrated to a quarterly frequency. Clearly, there is heterogeneity in the speed of the international shipping of goods and we would like to understand how accounting for the speed of delivery affects our results. A simple way to loosen our one-period lag assumption is to posit the following law of motion for how the consumption of imports relates to orders:

$$c(y)_t = \varphi y_{t+1} + (1-\varphi) y_t, \quad \text{and} \quad Q_t = \left(x_t^\rho + c(y)_t^\rho\right)^{1/\rho}, \tag{11}$$

where $c(y)_t$ is consumption of imports. Here, imports consumed at date t equals a fraction φ of orders made today plus $1-\varphi$ of last period's orders that did not arrive immediately. The law of motion in (11) reflects the idea that some orders may arrive immediately—e.g., the shipping of goods by airplane or from nearby trading partners—while other orders arrive with a delay—e.g., goods shipped by sea and from trading partners far away.

We solved the problem in (8), with the modifications in (11). We then performed the same exercise described above. In the results, we focused on the case with $\gamma = 10$ and varied the

Table 5: Intermediate Time-to-Ship, Elasticities: Data and Model ($\gamma = 10$)

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.48	2.00
Model, $\varphi = 0$	0.45 [0.35, 0.56]	2.74 [2.63, 2.86]
Model, $\varphi = 0.25$	-0.13 [-0.19,-0.06]	2.27 [2.20, 2.35]
Model, $\varphi = 0.50$	-0.70 [-0.78,-0.62]	1.80 [1.73, 1.87]
Model, $\varphi = 0.75$	-1.18 [-1.26,-1.10]	1.32 [1.24, 1.40]

Note: Results are averages from 250 simulations, with each simulation being 178 periods long; values in brackets report 95-percent confidence intervals. Section 2 describes the data.

time-to-ship parameter φ.

Table 5 present the results. The second row reports the results with $\varphi = 0$, the one-period time-to-ship model. The third through fifth rows report the results as we increase the fraction of orders that arrive within the same period. Not surprisingly, as φ increases, the importance of time-to-ship diminishes, and the price and income elasticities begin to converge towards those implied by the static model. Similar to Table 4, the same parameterization that generates an income elasticity that is close to the one in the data is also the same parameterization generating an price elasticity close to the one in the data.

Overall, Table 5 illustrates the effect of intertemporal motives on trade. If one made the argument that 50 percent of international deliveries to the U.S. arrive within a quarter, then with an intertemporal elasticity of substitution of 0.10, the model is well within the ballpark of the data. Or using Hummels's (2007) data, which say that about 30 percent of U.S. imports arrive by air, then our model with intertemporal elasticity of substitutions of between 0.25 and 0.10 can easily account for cyclical properties of imports in the U.S.

The results in Table 5 also suggest that the response of trade flows to shocks should depend on distance and mode of shipment. For example, international transactions shipped by air or those destinations with relatively short shipping times should be less responsive to shocks. In contrast, transactions shipped by sea, with long shipping times, should be very responsive to shocks. This suggests that one should see the volatility of trade flows increase with distance and shipping modes with higher time requirements (i.e., ocean shipping). There is evidence supporting this implication. Levchenko, Lewis, and Tesar (2010b) find that sectors with longer shipping times or higher shares of imports shipped by sea experienced larger falls in trade

Table 6: Measures of Fit, Data and Model

Statistic	Data	Baseline Model	Static Model
Root Mean Square Error, 1967-2011	—	0.037	0.057
Imports, % Deviation From Trend Q2 2009	−21.68	−12.91	−0.33

relative to sectors with shorter shipping times or imports shipped predominantly by air in the recent crisis. Amiti and Weinstein (2011) find that firms that export predominantly by air respond less to financial sector shocks than those that export predominantly by ship. Both of these results are consistent with our mechanism.

4.4. U.S. Imports 1967-2011 and the Trade Collapse

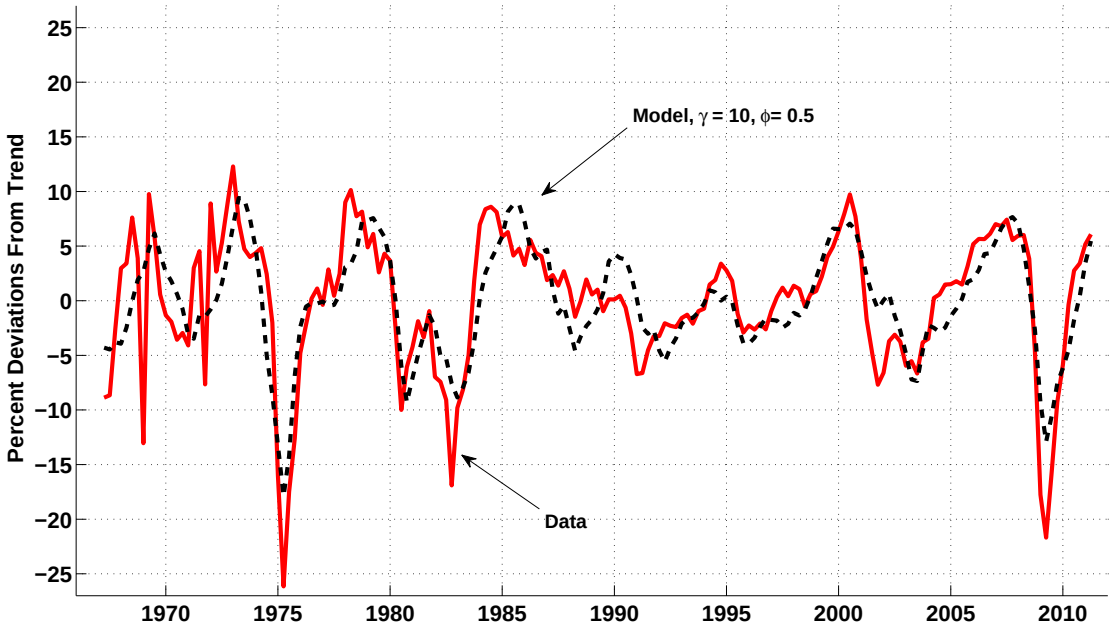
Given that our model that can account well for features U.S. imports at business-cycle frequencies, what are our model's implications for the 2008-2009 crisis?

To answer this question, we first applied the Kalman smoother to our model using U.S. data on absorption and prices over the entire 1967-Q3 2011 time period and generated a predicted series of imports. We set the parameters to $\gamma = 10$, $\theta = 1.5$, and $\varphi = 0.50$. Regarding the the intermediate time-to-ship parameter, as discussed above, we think that there are reasonable arguments that a fraction of trade is not that time-intensive and that 50 percent is a conservative value. The intertemporal elasticity of substitution was picked to be consistent with the nature of our exercise—i.e., given that the model with $\gamma = 10$ can account well for cyclical features U.S. imports over the entire time period, then what are our model's implications for the 2008-2009 crisis?

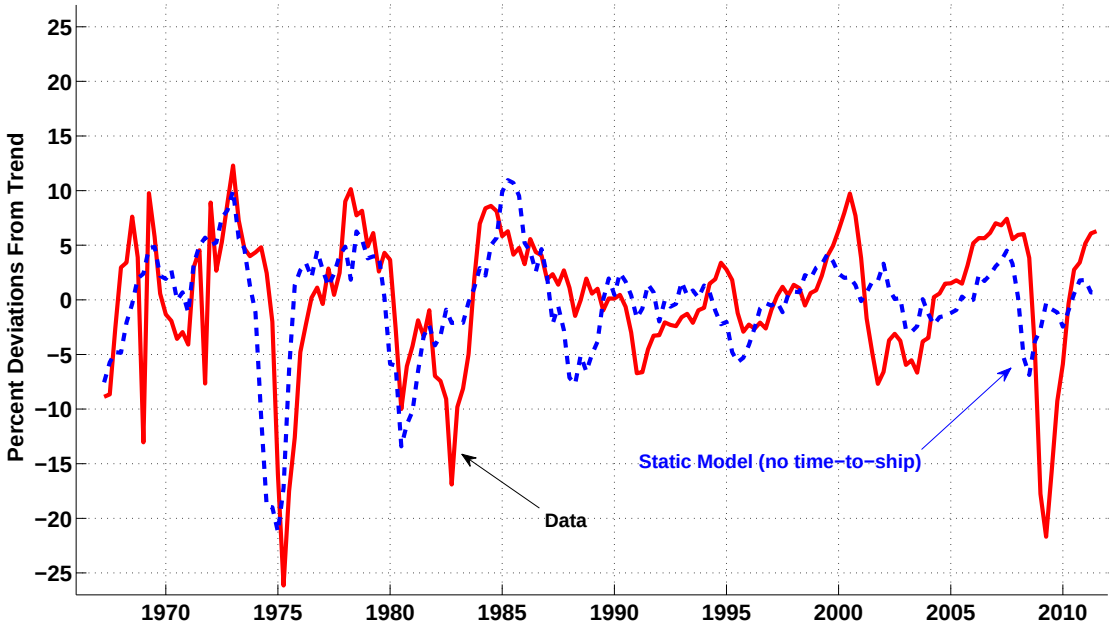
Figure 6(a) illustrates the results. It plots import data and predictions from our model for the time period 1967- Q3 2011. Our model performs very well. Our model tracks the data quite closely by capturing both the overall magnitude of fluctuations and the timing of peaks and troughs.[18] In contrast to these outcomes, Figure 6(b) presents the results from the static model. Unlike our model, the static model has severe problems regarding the timing and magnitudes in certain instances.

The first row in Table 6 provides a measure of fit. It reports the root mean squared error between

[18]There are two drops in imports that are completely unaccounted for by our model—specifically, Q1 1969 and Q4 1971—but, there is an explanation. In Q1 1969 and parts of Q3 and Q4 1971, the U.S. states suffered major shutdowns of U.S. ports due to dock worker strikes; see, e.g., Isard (1975). Thus, these events appear to account for this discrepancy between the model and the data.



(a) Model Prediction vs. Data



(b) No Time-to-Ship Model Prediction vs. Data

Figure 6: Model Predictions vs. Import Data, 1967-2011

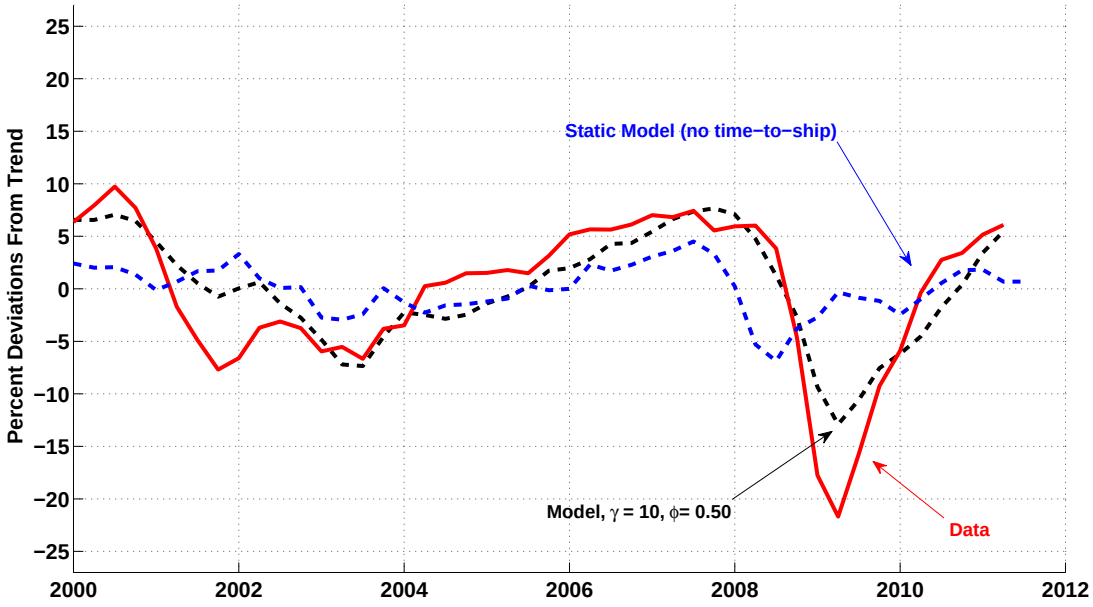


Figure 7: Trade Collapse: Model Predictions vs. Import Data, 2000-2011

the data and output from the model. The baseline model's root mean squared error is 0.037, 35 percent lower then the static model (0.057). As another point of comparison, the root mean squared error from the regression (2) is 0.037 as well. In other words, our calibrated model fits the data just as well as the best-fitting, linear regression of imports on absorption and prices. This shows that our model that can account well for cyclical fluctuations in U.S. imports data.

What are our model's implications for the 2008-2009 crisis? For the time period 2000-Q3 2011, Figure 7 plots import data, predictions from our model, and predictions from the static model. Our model performs very well over this time period. It picks up the decline in trade from 2000-2002/2003, the expansion in trade from 2003 to early 2008, the sharp drop from mid-2008 to 2009, and, finally, the expansion seen most recently. Figure 7 also plots the results from the static (no time-to-ship) model. It fails spectacularly on both the magnitude and the timing.

The key failure of the baseline model regards the magnitude during the collapse. The last row of Table 6 reports results regarding the percent deviation from trend during Q2 2009, which is the trough of the cycle. In the data, imports were 22 percent below trend. In the model, Q2 2009 is the trough of the cycle as well. However, at the trough, the model is only 13 percent below trend, nine percentage points away from the data.

There are meaningful ways of thinking about this discrepancy between the model and the data. Researchers have argued that there are mechanisms specific to the 2008-2009 crisis that are not in our model. Explicit mechanisms put forth are shocks in trade finance, as discussed in

Amiti and Weinstein (2011) or Chor and Manova (2010); inventory considerations discussed in Alessandria, Kaboski, and Midrigan (2010b); and input-output linkages and vertical specialization discussed in Bems, Johnson, and Yi (2010). Either of these mechanisms would complement our results and, perhaps, provide a complete account of the drop in trade.

The distinguishing feature of our explanation relative to this literature is that our model does not rely on specifics about the 2008-2009 crisis. As Figure 6(a) shows, our model that can account well for cyclical fluctuations in U.S. imports for 40 years of data. Trade finance, input-output linkages, compositional features, and inventories surely played some role in generating a large decline in trade during the 2008-2009 crisis. However, it remains to be shown that these mechanisms can account for cyclical fluctuations of international trade flows beyond the 2008-2009 crisis. Our model can meet this challenge.

5. Discussion

In this section, we discuss two issues regarding our analysis: general-equilibrium effects and implications of the IMRS in our model.

General-Equilibrium Effects. Our baseline exercise used a stochastic price process estimated from the data rather than solving for prices in general equilibrium. We explored versions of our model with prices determined in general equilibrium and found that our results are qualitatively preserved. This model is a two country version of the model outlined in Section 3 and is similar in sprit to the model of Cole and Obstfeld (1991) but with one-period time to ship.[19] However, the quantitative discrepancy arises because the model generates counterfactual terms-of-trade effects that work against our mechanism. Thus, the issue is not general equilibrium per se; rather it's that existing general-equilibrium models have empirically implausible effects that mitigate any intertemporal effects on international trade.

To see how general-equilibrium effects mitigate our quantitative results, consider a transitory positive shock to the domestic endowment. As in our model, the positive endowment shock increases the IMRS and pushes the agent to increase future consumption relative to present consumption, providing a force to increase in imports. However, the relative price of the home good decreases since its supply increases relative to that of the foreign good. This price change leads to an increase in the demand for the domestic good offsetting the impact from the change in the IMRS. Because relative prices moved in response to the endowment shock, the impact of

[19] We explored extensions with a two-country general-equilibrium production economy setup, as in Backus, Kehoe, and Kydland (1995). Plausible parameterizations can reverse our results; In the economy with capital, capital changes the choice to substitute presents consumption for future consumption, in response to a positive productivity shock. This lowers the IMRS, instead of rising it, and thus leads to a lower than unit income elasticity. This, however, relates to issues RBC models face in replicating features of asset prices (see, e.g., Jermann (1998)) and are solved with adjustment costs to capital accumulation.

the change in the IMRS is partially offset.

However, this strong correlation between international prices and output is inconsistent with the data. This is a generic problem of standard general-equilibrium models, as in Cole and Obstfeld (1991) or Backus, Kehoe, and Kydland (1995). Our approach was purposefully designed to sidestep these complications. Because we do not know the model that can achieve the correct correlations between prices and output, we instead estimated a price process from the data and fed it through our model to isolate the role of intertemporal substitution. This suggests, however, that a model that can deliver correlations between international prices and output consistent with the data will generate fluctuations in imports like those found in Table 4.

Implications of the IMRS. Driving our results are changes in the IMRS. This force is at the heart of understanding features of asset prices, such as the equity premium puzzle, the risk-free rate puzzle, etc. This raises natural questions about how variations in IMRS implied by our model relate to observable features of asset prices. Our model faces the same generic challenges that other models with CRRA preferences and standard stochastic processes for endowments face.

Two natural implications to examine in our model are the variation in the risk-free rate and the market price of risk. The risk-free rate is informative because the inverse of the expected IMRS should equal the gross return on a risk-free bond paying out in the aggregate consumption bundle, Q. The market price of risk, $\sigma(m_t)/E(m_t)$, is informative because it is bounded below by the Sharpe ratio, $E(R^e)/\sigma(R^e)$, where R^e is the excess return on equity relative to risk-free bonds as argued by Hansen and Jagannathan (1991).

Tallarini (2000) measures the standard deviation of the ex-post return on three-month U.S. Treasury bills (a presumed near risk-free asset) to be 0.80 percent and a Sharpe ratio of about 0.25. In our model, with $\gamma = 2$ and one-period time-to-ship, the standard deviation of the implied risk-free rate is 0.96 percent, but it has a market price of risk of 0.02. With $\gamma = 10$, the standard deviation of the risk-free rate is 5.4 percent, but the market price of risk is 0.12. Clearly, there is a tension between satisfying the low observed volatility in near-risk-free rates and the Hansen and Jagannathan (1991) bound. This should not be a surprise. It is just one of the challenges theories of asset prices face: Observed near-risk-free returns are not volatile, but the Hansen and Jagannathan (1991) bound suggests that the stochastic discount factor must be very volatile, and models with CRRA preferences and standard stochastic processes for endowments face challenges satisfying both.

If these challenges were met, would this alter our results? The answer is not clear. Possible routes for future research are extensions with recursive preferences and long-run risk, as in Bansal and Yaron (2004); habits, as in Campbell and Cochrane (1999); or diasters as in Rietz (1988) which are known to improve the empirical performance of asset-pricing models. However, we should caution that even if these extensions satisfied say the Hansen and Jagannathan

(1991) bound, it could miss features that are perhaps more critical for our results. In our model, its the covariance of the IMRS with macro-aggregates that is key to driving the cyclical properties of imports—not unconditional moments.

6. Conclusion

Our paper shows how incorporating dynamic, forward-looking features into international trade models improves their ability to explain the behavior of imports at business-cycle frequencies. The key premise is that international trade is a time-intensive activity, and, thus, variation in the rate at which agents are willing to substitute across time will affect how trade volumes respond to changes in income and prices. Quantitatively, we showed that our model can deliver the high income elasticity and low price elasticity of imports in U.S. time series data at business-cycle frequencies. Furthermore, we showed that our model can account well for both the collapse in U.S. imports during 2008-2009 and fluctuations over the past 40 years.

Several questions and avenues for future research remain open. While we discussed evidence on the volatility of trade by distance and mode, further analysis of this type may help to provide discipline regarding the mechanism put forth in this paper. Second, trade elasticities play critical roles in formulating predictions and recommendations for policy makers. Because our model has both theoretical consistency and statistical performance, exploring the model's ability to provide usable forecasts is an avenue for future research, as well.

References

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109(3), 659–84.

ALESSANDRIA, G., J. KABOSKI, AND V. MIDRIGAN (2010a): "Inventories, lumpy trade, and large devaluations," *American Economic Review*.

——— (2011): "U.S. Trade and Inventory Dynamics.," Discussion paper.

ALESSANDRIA, G., J. P. KABOSKI, AND V. MIDRIGAN (2010b): "The Great Trade Collapse of 2008-09: An Inventory Adjustment?," NBER Working Papers 16059, National Bureau of Economic Research, Inc.

AMITI, M., AND D. WEINSTEIN (2011): "Exports and Financial Shocks," *The Quarterly Journal of Economics*.

ANDERSON, J., AND E. VAN WINCOOP (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–192.

ANTRAS, P., AND C. F. FOLEY (2011): "Poultry in Motion: A Study of International Trade Finance Practices," Working Paper 17091, National Bureau of Economic Research.

BACKUS, D., P. KEHOE, AND F. KYDLAND (1995): "International business cycles: theory and evidence," in *Frontiers of Business Cycle Research*, ed. by T. Cooley. Princeton University Press.

BALDWIN, R. (2010): *The great trade collapse: Causes, Consequences and Prospects*. CEPR.

BANSAL, R., AND A. YARON (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 59(4), 1481–1509.

BEMS, R., R. JOHNSON, AND K. YI (2010): "Demand spillovers and the collapse of trade in the global recession," *IMF Economic Review*.

BOILEAU, M. (1999): "Trade in capital goods and the volatility of net exports and the terms of trade," *Journal of International Economics*, 48(2), 347–365.

CAMPBELL, J., AND J. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *The Journal of Political Economy*, 107(2), 205–251.

CAPELA, J. J. (2008): *Import/Export For Dummies*. For Dummies.

CASTRO, R. (2005): "Economic Development and Growth in the World Economy," *Review of Economic Dynamics*, 8(1), 195–230.

CHARI, V. V., P. KEHOE, AND E. R. MCGRATTAN (2007): "Business Cycle Accounting," *Econometrica*, 75(3), 781–836.

CHOR, D., AND K. MANOVA (2010): "Off the Cliff and Back? Credit Conditions and International Trade during the Global Financial Crisis," Working Paper 16174, National Bureau of Economic Research.

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.

DIEWERT, W. E. (1978): "Superlative Index Numbers and Consistency in Aggregation," *Econometrica*, 46(4), 883–900.

DJANKOV, S., C. FREUND, AND C. S. PHAM (2010): "Trading on time," *Review of Economics and Statistics*, 92(1), 166–173.

EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

EATON, J., S. KORTUM, B. NEIMAN, AND J. ROMALIS (2010): "Trade and the global recession," Working paper.

EHEMANN, C., A. KATZ, AND B. MOULTON (2002): "The chain-additivity issue and the US national economic accounts," *Journal of Economic and Social Measurement*, 28(1), 37–49.

ENGEL, C., AND J. WANG (2009): "International trade in durable goods: understanding volatility, cyclicality, and elasticities," Working paper.

FEENSTRA, R. (1994): "New product varieties and the measurement of international prices," *The American Economic Review*, 84(1), 157–177.

FISHER, I. (1922): *The making of index numbers*. Houghton Mifflin Company.

HANSEN, L., AND R. JAGANNATHAN (1991): "Restrictions on intertemporal marginal rates of substitution implied by asset returns," *Journal of Political Economy*, 99(2), 225.

HEATHCOTE, J., AND F. PERRI (2002): "Financial autarky and international business cycles," *Journal of Monetary Economics*, 49, 289–301.

HOOPER, P., K. JOHNSON, AND J. MARQUEZ (2000): "Trade elasticities for the G-7 countries," .

HOUTHAKKER, H., AND S. MAGEE (1969): "Income and price elasticities in world trade," *The Review of Economics and Statistics*, pp. 111–125.

HUGGETT, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 17(5-6), 953–969.

HUMMELS, D. (2007): "Transportation costs and international trade in the second era of globalization," *The Journal of Economic Perspectives*, 21(3), 131–154.

HUMMELS, D., AND G. SCHAUR (2012): "Time as a Trade Barrier," Working Paper 17758, National Bureau of Economic Research.

ISARD, P. (1975): *Dock strike adjustment factors for major categories of US imports and exports, 1958-1974*. Division of International Finance, Board of Governors of the Federal Reserve System.

JACKS, D., C. MEISSNER, AND D. NOVY (2010): "Trade costs in the first wave of globalization," *Explorations in Economic History*, 47(2), 127–141.

JACKS, D. S., C. M. MEISSNER, AND D. NOVY (2009): "The role of trade costs in the great trade collapse," in *The great trade collapse: causes, consequences, and prospects*, ed. by R. Baldwin. VoxEU.org.

JERMANN, U. (1998): "Asset pricing in production economies," *Journal of Monetary Economics*, 41(2), 257–275.

KEHOE, T. J., AND K. J. RUHL (2009): "How important is the new goods margin in international trade?," Staff Report 324, Federal Reserve Bank of Minneapolis.

KRUGMAN, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950–959.

——— (1989): "Differences in income elasticities and trends in real exchange rates," *European Economic Review*, 33(5), 1031–1046.

LEVCHENKO, A. A., L. T. LEWIS, AND L. L. TESAR (2010a): "The collapse of international trade during the 2008-2009 crisis: in search of the smoking gun," Working Paper.

——— (2010b): "The Role of Financial Factors in the Trade Collapse: A Skeptic's View," Working papers, Research Seminar in International Economics, University of Michigan.

MARQUEZ, J. (2002): *Estimating trade elasticities*. Springer.

MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

RIETZ, T. (1988): "The equity risk premium a solution," *Journal of monetary Economics*, 22(1), 117–131.

RUHL, K. (2008): "The International Elasticity Puzzle," *unpublished paper, NYU*.

SIMONOVSKA, I., AND M. WAUGH (2010): "Elasticity of Trade: Estimates and Evidence," Discussion paper, New York University.

TALLARINI, T. (2000): "Risk-sensitive real business cycles," *Journal of Monetary Economics*, 45(3), 507–532.

WHELAN, K. (2002): "A Guide to U.S. Chain Aggregated NIPA Data," *Review of Income and Wealth*, 48(2), 217–33.

Appendix

A. Data Sources

The data that we use for the estimation of the static CES import demand specification in Section 2 come from the Bureau of Economic Analysis' (BEA) National Income and Product Accounts (NIPA). As emphasized in the paper, our analysis focuses on imports and absorption of goods, excluding oil. The tables we use are:

- **Nominal Components:** Table 1.2.5 line 5 Nominal Goods GDP, Final Sales; Table 4.2.5 line 2 nominal exports of goods; Table 4.2.5 line 54 nominal imports of nonpetroleum goods.
- **Price Indexes:** Table 1.2.4 line 5 Fisher price index (100=2005) of goods gdp, final sales; Table 4.2.4 line 2 Fisher price index (100=2005) exports of goods; Table 4.2.4 line 54, Fisher price index (100=2005) of imports of nonpetroleum goods.
- **Quantity Indexes:** Table 1.2.3 line 5 Fisher quantity index (100=2005) of goods gdp, final sales; Table 4.2.3 line 2 Fisher quantity index (100=2005) exports of goods; Table 4.2.3 line 54, Fisher quantity index (100=2005) of imports of nonpetroleum goods.
- **Durable and Non-Durable Goods:** The same tables outlined above were used to construct the analogous data series for durable goods and non-durable goods. The only distinction are the line numbers—specifically lines 7 and 11 for nominal values, quantity and price indexes. Lines 48 and 49 for exports of durable and non-durable goods, lines 52 and 53 for imports.

As discussed in Section 2.1, the construction of real absorption and the associated price index is not as straightforward as this might seem. Real values in the U.S. NIPA accounts are chain-type indexes using an "ideal" chain index advocated by Fisher (1922). While these indexes have desirable properties, they are not additive across categories (see Ehemann, Katz, and Moulton (2002) and Whelan (2002) for detailed discussions). For our purposes, the implication is that one cannot compute real absorption simply by adding real goods GDP to real imports and subtracting real exports.

Our solution to this problem is to use a "Fisher of Fishers" approach suggested by Diewert (1978). The basic idea is to take the quantities indexes and the associated price indexes for the categories of interest and then compute Fisher indexes of these measures—hence the "Fisher of Fishers" name. For example, to construct real absorption and the associated price index, then the quantity indexes of goods GDP, (minus) goods exports, and goods imports and the price indexes are combined to create a Fisher quantity index and price index.

B. The Mapping From Model to Data

This section discusses how objects in our model relate to those in the data. The key issue regards adjusting observed time series given the timing friction to estimate the stochastic process for endowments and prices. Below, we discuss the one-period time-to-ship case and the intermediate time-to-ship case.

One-Period Time-to-Ship. To estimate the stochastic process for endowments and prices, we need to construct a time series to proxy for z_t, p_{xt}, and p_{yt}. Specifically, we need the variables in the data to reflect the timing at which they would be observed by the agent in the model, not the timing at which they are observed by U.S. statistical agencies. These agencies collect import data and the prices on arrival at the border. Yet, in the model, these prices are observed a quarter before. Thus, we adjust the data variables accordingly to make them consistent with the timing of our model. Below, we discuss the adjustment and measurement for p_y, z, and p_x in turn.

To measure p_{yt}, we used the observed price index of imports, but shifted one period back. Again, the reason is that is the price measured in Q1 2011 is really the price that the agent observed and on which he based his choice of imports in Q4 2010 according to our model. Thus, by shifting back the Q1 2011 price of imports, it will line up in the estimation with Q4 2010 real (adjusted) absorption and price index.

The timing assumption affects absorption, as well. Absorption from NIPA includes consumption of imports decided upon in the previous period. In contrast, we want domestic consumption today plus consumption of imports delivered tomorrow. To adjust for this, adjusted absorption is measured as

$$\begin{aligned} \text{Adjusted Absorption} &= \underbrace{\underbrace{p_{xt}x_t + p_{yt-1}y_{t-1}}_{C_t} + \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} - \underbrace{p_{yt-1}y_{t-1}}_{Imports_t}}_{GDP_t} - \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} + \underbrace{p_{yt}y_t}_{Imports_{t+1}} \quad (12) \\ &= p_{xt}x_t + p_{yt}y_t = p_{xt}z_t, \end{aligned}$$

with the last line showing that this process identifies the value of the endowment. Finally, to arrive at a real measure of the endowment, we can construct a quantity index for absorption, z_t, by using real GDP, minus real exports, plus real imports at $t+1$, and the associated price indexes using the "Fisher of Fishers" approach.

To proxy p_x, there are two approaches. The baseline approach uses the associated price index with our measure of absorption. A second approach computes a measure of domestic consumption and the associated price index. Specifically, we define domestic consumption as GDP

Table 7: Robustness, Elasticities: Data and Model, $\theta = 4$

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.48	2.00
Model, $\gamma = 2$	-2.64 $[-2.79, -2.48]$	2.00 $[1.90,\ 2.11]$
Model, $\gamma = 5$	-1.04 $[-1.24, -0.83]$	3.29 $[3.10,\ 3.50]$
Model, $\gamma = 10$	0.82 $[\ 0.45,\ \ 1.23]$	5.10 $[4.64,\ 5.64]$
Model, no time-to-ship	-3.82 $[-3.99, -3.63]$	1.02 $[0.90,\ 1.14]$
Model, $\gamma = \frac{1}{\theta}$	-4.02 $[-4.20, -3.83]$	1.04 $[0.98,\ 1.10]$

Note: Results are averages from 250 simulations, with each simulation being 178 periods long; values in brackets report 95-percent confidence intervals.

minus exports, which gives

$$\begin{aligned} \text{Domestic Consumption} &= \underbrace{\underbrace{p_{xt}x_t + p_{yt-1}y_{t-1}}_{C_t} + \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} - \underbrace{p_{yt-1}y_{t-1}}_{Imports_t}}_{GDP_t} - \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} \\ &= p_{xt}x_t, \end{aligned} \tag{13}$$

which is the value of consumption of the domestic good. With real values of GDP and exports and the price indexes, a price index can be constructed using a "Fisher of Fishers" approach. Quantitatively, we found little difference in the two approaches.

Intermediate Time-to-Ship. In this case, we followed the same general approach described above. The only difference is that measured imports are now a combination of imports decided upon today and yesterday, thus, all import series and the associated price indexes are adjusted to reflect this. Specifically,

$$\text{Adjusted Imports}_t = \varphi\text{Measured Imports}_t + (1-\varphi)\text{Measured Imports}_{t+1}. \tag{14}$$

Price indexes are adjusted similarly.

Table 8: Robustness, Orthogonal Endowment and Price Process

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.48	2.00
Model, $\gamma = 2$	-1.40 $[-1.46, -1.33]$	1.31 $[1.28,\ 1.35]$
Model, $\gamma = 5$	-1.26 $[-1.33, -1.18]$	1.88 $[1.81,\ 1.95]$
Model, $\gamma = 10$	-1.10 $[-1.21, -0.98]$	2.73 $[2.60,\ 2.86]$

Note: Results are averages from 250 simulations with each simulation being 178 periods long; values in brackets report 95 percent confidence intervals.

C. Robustness: Alternative Parameterizations

In this section, we report two alternative parameterizations. The first one reproduces Table 4 but with $\theta = 4$. The second alternative parameterization reports results when the stochastic process for endowments and prices are treated as independent AR(1) processes and estimated separately.

Table 7 presents the results with $\theta = 4$. The key observation is that both the income and price elasticities are significantly higher than in the baseline results. For example, with $\gamma = 2$, the income elasticity is two—the same as the $\gamma = 5, \theta = 1.5$ case presented in Table 4. Price elasticities are, higher, but they are significantly biased below the true price elasticity.

Table 8 presents the results with a different stochastic process for endowments. We estimated separate AR(1) processes for each series. Thus, there is no covariance at all between endowments and prices. For the income elasticity, the results are effectively the same. For the price elasticity, the results appear to be larger than those in Table 4. The reason is evident in the discussion at the end of Section 3.1 about how different covariance structures can either reinforce the results or offset them. Here, it appears that the covariance structure estimated from the data appears to reinforce the results specifically for the estimate of the price elasticity.