International Trade and Intertemporal Substitution

Fernando Leibovici

York University

Michael E. Waugh

New York University and NBER

August 2014

ABSTRACT

This paper studies the dynamics of international trade flows at business cycle frequencies. We show that introducing dynamic considerations into an otherwise standard model of trade can account for several puzzling features of trade flows at business cycle frequencies. Our insight 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 output and prices. We formalize this idea and calibrate our model to match key features of U.S. data. We find that, in contrast to standard static models of international trade, our model is quantitatively consistent with salient features of U.S. cyclical import fluctuations. We also find that our model accounts for two-thirds of the peak-to-trough decline in imports during the 2008-2009 recession.

Email: flei@yorku.ca, 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

The collapse of international trade during the recent financial crisis exposed the failure of standard trade models to explain the response of trade volumes to changes in economic activity during both normal and crisis episodes.[1] For instance, the empirical elasticity of imports to measures of economic output are well above one, yet standard models have a unitary income elasticity. Similarly, while the empirical elasticity of import volumes to measures of relative prices is well below one, typical calibrations of standard models use values that are well above one. Moreover, accounting exercises that use static trade models to measure deviations between predicted and observed fluctuations in imports find these deviations to be pro-cyclical.[2]

In this paper, we show that introducing dynamic considerations into a standard model of trade can account for these puzzling features. The dynamic consideration we focus on is a time-to-ship friction to import goods and its interaction with a finite intertemporal elasticity of substitution.[3] 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—the 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 breaks the unitary income elasticity, biases the estimated price elasticity relative to static trade models, and shows up as a time-varying trade friction. Quantitatively, we show that this insight is able to account well for these key features of U.S. cyclical import fluctuations.

We formalize these ideas by building a pure exchange, Armington model of trade. 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. International purchases are subject to an ad-valorem trade cost. The only other friction that agents face is that they must commit resources today for the delivery of imported goods in subsequent periods.

[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 that documents the high income elasticity for the U.S. and other countries at low 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.

[3]Hummels and Schaur (2013) document the time intensive nature of international trade and, thus, introducing a time-to-ship friction is a natural way to introduce dynamics to the decision to import.

We use the model to answer the following quantitative question: Given a stochastic process describing income and prices as 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 estimate the stochastic process for endowments and prices from U.S. data. We then feed this stochastic process into our model and walk through several exercises to study the quantitative importance of our proposed mechanism.

The first exercise studies the model implied income and price elasticities. To compute these measures, we simulate the model and estimate the model implied income and price elasticities using simulated data under different assumptions about the intertemporal elasticity of substitution, the elasticity of substitution across goods, and the shipping technology. We find that the model can quantitatively account for the high income elasticity and low price elasticity observed in U.S. time series data. For instance, with an intertemporal elasticity of substitution of 0.20 and an elasticity of substitution of 1.5, the income and price elasticities are 1.70 and -0.34; in the data, they are 1.99 and -0.26. 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.

The second exercise focuses on the ability of our model to account for the observed dynamics of U.S. imports. To discipline this exercise, we calibrate the preference parameters of our model to match features of the data over the first half of our sample. We then compute our model's predicted import series for the the un-targeted second half of our sample and compare it with data. We find that our model accounts well for the dynamics in U.S. imports by correctly capturing the overall magnitude and timing of cyclical fluctuations (see Figure 5(a).) And that the active intertemporal marginal rate of substitution is critical to accounting for these fluctuations.

Finally, we provide supporting evidence of the mechanism by examining some cross-sectional implications of our model. In particular, our model predicts that a country's bilateral imports should be more volatile when sourced from a partner with longer shipping times. This implication is a test of our model because static trade models (or our dynamic model with no active intertemporal marginal rate of substitution) predict that the volatility of imports is independent of to the time-to-ship/distance. Using data constructed by Hummels and Schaur (2013) and the World Bank on shipping times, we find that US imports from countries with higher than average shipping times are considerably more volatile than imports from countries with lower than average shipping times. We find these results to be supportive of the underlying mechanism at work.

Our results have much to say about the large drop in trade during the 2008-2009 crisis.[4] Our

[4]Much attention has focused on this episode. See, for example, the papers in Baldwin (2010); Alessandria,

explanation is simple: a very large 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. Quantitatively, we find that our model accounts two-thirds of the peak-to-trough decline in imports in U.S. imports (see Figure 6), when the parameters of the model are calibrated to pre-crisis events.

Our results complement alternative mechanisms proposed to explain the trade collapse. In particular, an active intertemporal marginal rate of substitution would surely amplify the role of financial fictions discussed in Amiti and Weinstein (2011) and Chor and Manova (2012), inventory considerations in Alessandria, Kaboski, and Midrigan (2010b), or the future value of manufactures as in Eaton, Kortum, Neiman, and Romalis (2013). The distinguishing feature of our story, however, is that it does not rely on specifics about the 2008-2009 crisis but applies broadly across time periods. This is a key implication of our second exercise which calibrated the model parameters to pre-crisis events. Moreover, the lack of reliance on specifics of 2008-2009 crisis is consistent with the evidence from Stock and Watson (2012) who find that the same factors that explained previous postwar recessions also explain the most recent recession.

Trade elasticities play critical roles in formulating predictions and recommendations for policy makers with regard to issues such as the effects of exchange rate fluctuations and policy (see, e.g., Houthakker and Magee (1969) and 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 important forecasting and policy questions.

2. Cyclical Features of International Trade Volumes

In this section, we document some key features of the cyclical fluctuations of imports, income, prices and their co-movement in U.S. time series data. The data features we describe are not 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 these three features of the data is important since we will later show that they can be simultaneously

Kaboski, and Midrigan (2010b); Amiti and Weinstein (2011); Bems, Johnson, and Yi (2010); Chor and Manova (2012); Eaton, Kortum, Neiman, and Romalis (2013); Jacks, Meissner, and Novy (2009); Levchenko, Lewis, and Tesar (2010a).

rationalized by a single underlying mechanism.

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 economic activity that we focus on is absorption, we use the terms income, output, and absorption synonymously throughout.[5]

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 or production functions.[6] In these models, the demand function for imports is given by:

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

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 is the price elasticity of imports, and ω_t is a "wedge," which we describe in more detail below.

We use the structure of equation (1) to summarize key features of the data. We do so in two ways. Our 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. These empirical elasticities inform us about the response of real imports to changes in income and import prices, and allow us to examine the extent to which standard models deviate from the relationship observed in the data.

The second exercise imposes the theoretical restrictions implied by equation (1), a unit income elasticity and an assumed value for the price elasticity, and uses the data on income and import prices to obtain a measure of predicted imports. Finally, we infer the wedge ω by comparing

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

[6] 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).

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 calibration of this parameter in the international business-cycle literature. Using larger θs, as in typical calibrations of international trade models, results 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 their associated price indices. Because these are important issues, we 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) report 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 (2013). To address these compositional issues, we focus on an absorption measure where most trade occurs (goods-only component of GDP). To address compositional issues within goods (i.e., durable vs. non-durable) emphasized by Boileau (1999) and Engel and Wang (2011), we perform the same analysis later in this section using only durable or non-durable goods.

Constructing real measures of these activities and their associated price indices is not as straightforward as it might seem. Real values in the U.S. NIPA accounts are chain-type indexes constructed 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"

Table 1: Empirical Price and Income Elasticities

Data	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$	R^2	# Obsv.
Goods GDP	-0.26 (0.13)	1.99 (0.14)	0.65	187

Note: Data are in logs and HP filtered over the time period from Q2 1967 to Q4 2013. Heteroskedastic robust standard errors are in parenthesis.

approach suggested by Diewert (1978). The basic idea is to take the real values and their 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 their associated price indexes starting in the second quarter of 1967 and ending in the fourth quarter of 2013. To deal with trends, we HP-filter the logarithm of these data with smoothing parameter 1600. Results using log-first-differences yield no significant differences.

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

Table 1 presents the estimated income and price elasticities of imports, using ordinary least squares to estimate equation (2). Figure 2(a) plots the result from the wedge accounting exercise. Below we outline three observations from these exercises.

O.1. Income elasticity > 1. The estimated income elasticity of imports is nearly two—i.e., a one-percent increase in absorption is associated with 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 contrasts with the implication of standard models of international trade. These models imply that a one-percent increase in absorption results in a one-percent increase in imports—i.e., these models feature a unit income elasticity. Thus, attempts at modeling trade volumes at business-cycle frequencies must confront this discrepancy between standard trade models and the data.

As mentioned above, the fact that the estimated income elasticity of demand for U.S. imports exceeds unity is not new, and dates back to Houthakker and Magee (1969).[7] Marquez (2002)

[7] Prominent explanations of the high income elasticity of import demand are largely based on expanding product variety (see, e.g., Krugman (1989) and Feenstra (1994)) and are best thought of as medium-/long-run explana-

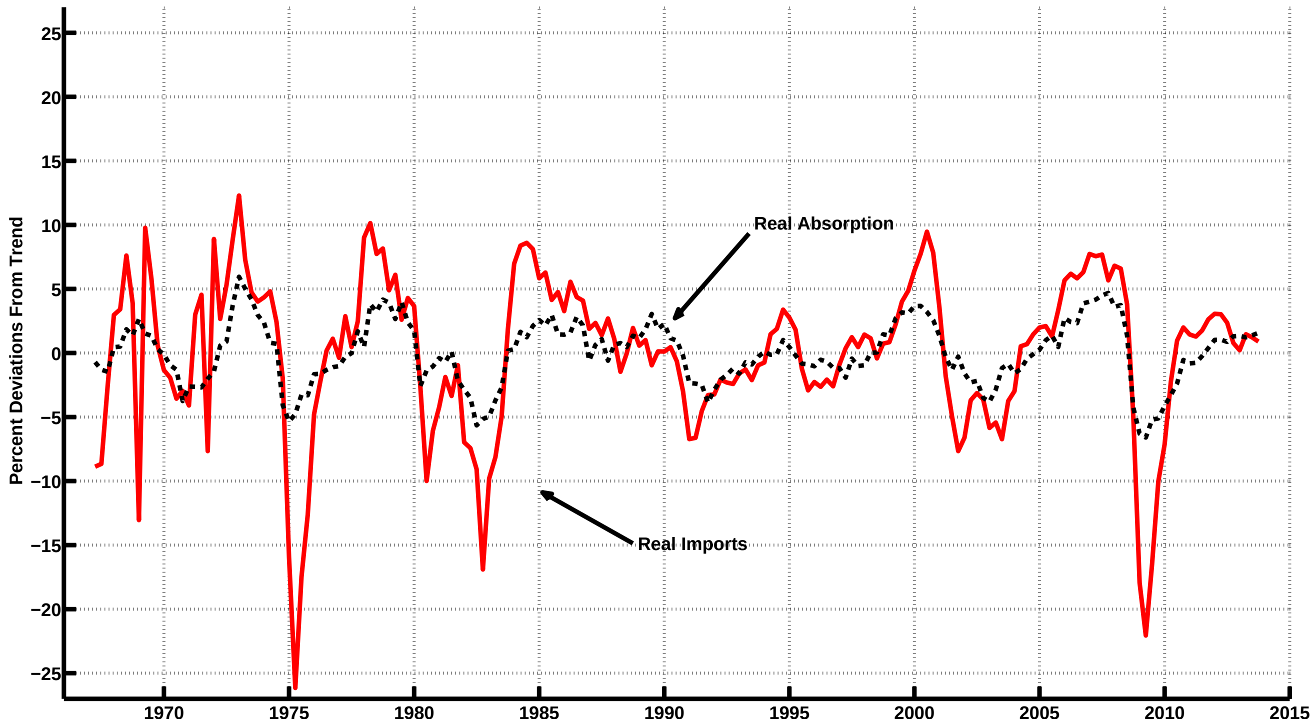

(a) **Real Absorption and Real Imports**

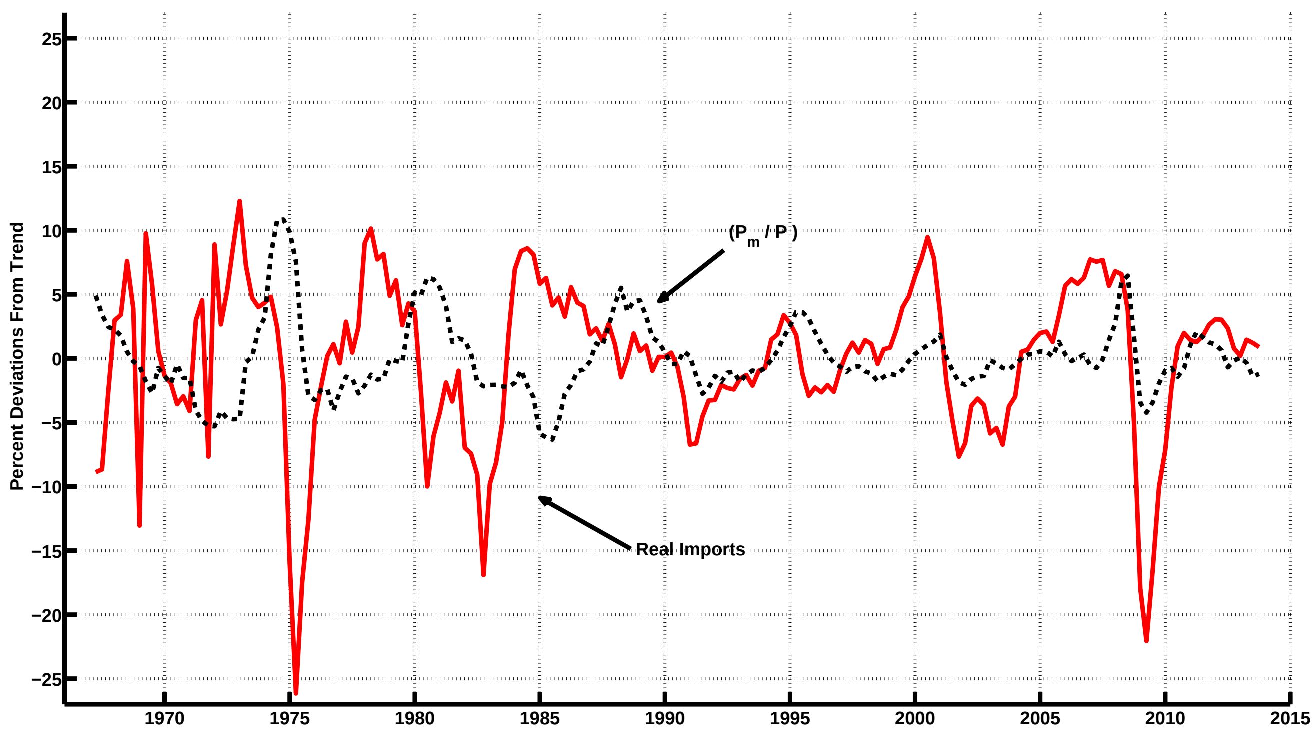

(b) **Relative Prices and Real Imports**

Figure 1: Absorption, Relative Prices, and Import Data

examines this feature of the data from a modern perspective and finds that it is robust to alternative econometric specifications, different frequencies, and commodity disaggregation. At a disaggregated level, Fitzgerald and Haller (2012) find a high income elasticity (and low price elasticity) at the plant level in census data from Ireland.

O.2. Low price elasticity. Our second observation is that the estimated import price elasticity is -0.26. 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.

While modelers have a choice over this parameter, typical calibrations/estimations of static trade models or international business-cycle models generally use values of it that are considerably larger. Moreover, estimates of this parameter based on static trade models and changes in trade flows during trade liberalizations typically suggest substantially higher values of it. Lower values, but still higher than we estimate, 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. Our last observation 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 and the imports 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 0.42. 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

---

tions. Quantitatively, Feenstra (1994) shows that the expanding product variety explanation can account for only a part of the high income elasticity. Ruhl (2008) shows that this margin is not quantitatively important at business-cycle frequencies. Kehoe and Ruhl (2013) measure changes in the extensive margin and find that it plays little role outside of significant structural transformations or trade liberalization.

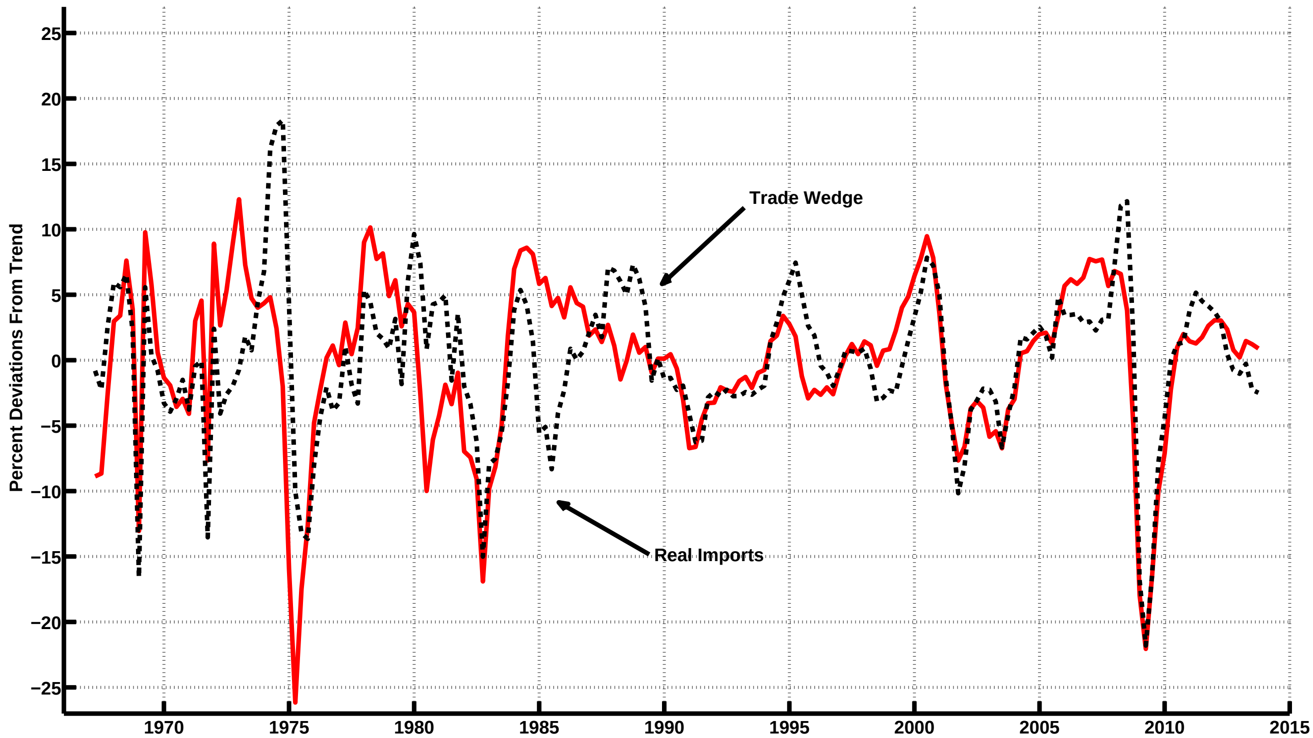

(a) **Wedges and Real Imports**

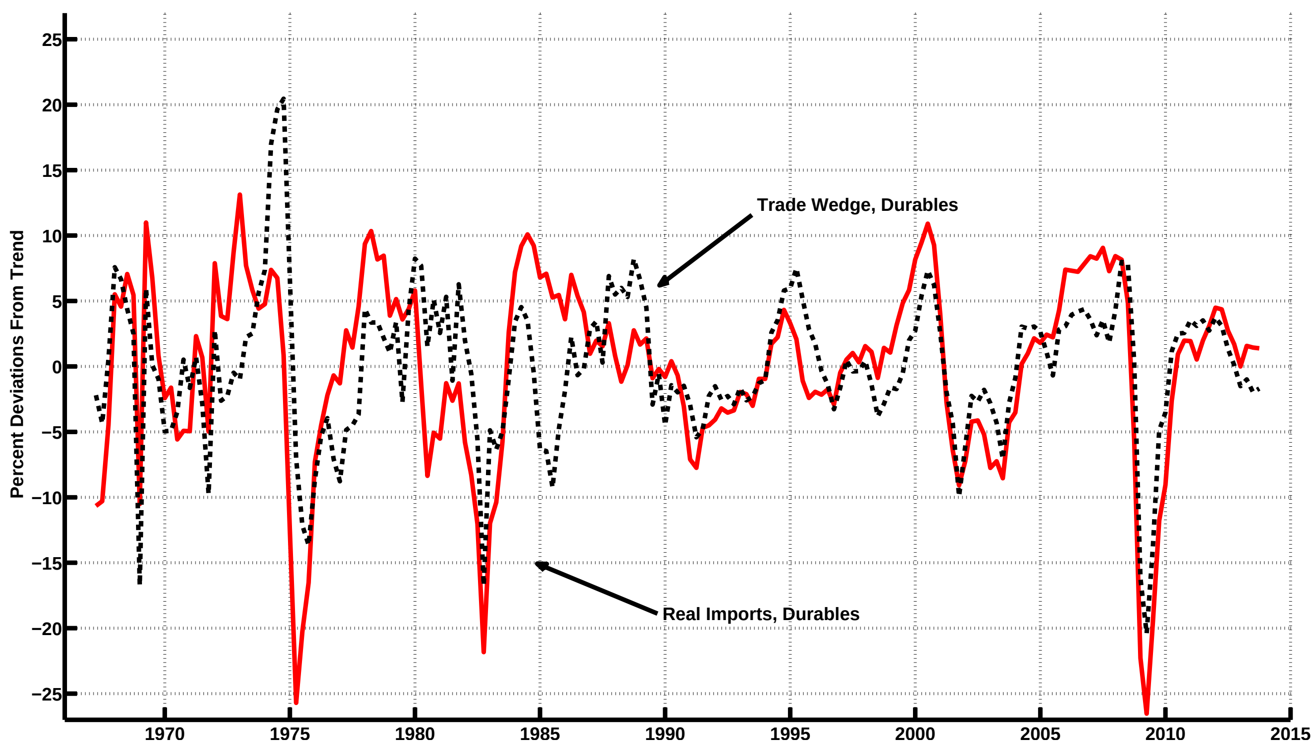

(b) **Wedges and Real Imports, Durables Only**

Figure 2: Wedges and Import Data

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 2008-2009 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 the role of inventories. Here, we show that observations **O.1-3** are robust to restricting attention to durables or non-durable goods, and accounting 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 (2011). The argument is based on the fact that a larger fraction of imports is classified as durable than in, say, absorption of total goods. This observation, combined with the fact that consumption of durables is more volatile than that of non-durables, 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 address this argument by re-estimating equation (2) restricting the import, absorption, and price data to include only durable goods or only non-durables. 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 of durables 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 0.80 and an R^2 of 0.42. 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

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

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

Data Series / Approach	Price Elasticity	Income Elasticity	Inventory Elasticity	R^2	# Obsv.
Durable Goods	-0.27 (0.13)	1.52 (0.11)	—	0.69	187
Non-Durable Goods	-0.11 (0.05)	2.72 (0.31)	—	0.45	187
Goods & Δ Inventories	-0.27 (0.12)	1.65 (0.11)	0.20 (0.02)	0.73	187

Note: Data are in logs and HP filtered over the time period from Q2 1967 to Q4 2013. Heteroskedastic robust standard errors are in parentheses.

opposite pattern. Thus, while the income elasticity is mitigated, the relative weighting of the income and substitution effects still conflict with what the data suggest.

Another implication of the durable goods explanation is the following: we should observe that the income elasticity of imports is higher when the share of durables in imports is large relative to the share of durables in total absorption—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 absorption.

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. We found a strong negative correlation (-0.63 and statistically different from zero) between the income elasticity and the relative share of durables in imports. This result goes against the positive correlation that the durables explanation implies.

Inventories. Another concern is that **O.1-3** arises because we abstract from changes in inventories.[9] 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).

[9] 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.

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 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 1.65, relative to 1.99 without controlling for inventories. Including inventories also improves the fit of the regression from an R^2 of 0.65 without inventories to 0.73 with inventories. These results suggest that inventory adjustments are a partial, but not a complete explanation, 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 breaks the unitary income elasticity, biases 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 infinitesimal 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 is the coefficient of risk aversion, and $\frac{1}{\gamma}$ is the intertemporal elasticity 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 given by $\theta = \frac{1}{1-\rho}$.

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, and can be either consumed domestically or used to acquire foreign goods through a centralized goods market. Goods cannot be stored.[10]

International trade consists of 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. We assume that goods cannot be re-exported.

Second, international purchases take time.[11] We model this as a time-to-ship friction, 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 $t+1$. Along with this assumption, we assume that goods must be paid for before they are delivered.[12] Underlying this international trading structure is an assumed enforcement technology that allows countries to coordinate the dynamic exchange of goods across international borders.

We abstract from trade in international financial assets.[13] This does not imply that the typical agent is in financial autarky. Because of the timing assumption on international trade, trade in goods 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 as 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 the 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:

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

[10] The shipping technology to trade goods internationally does actually provide a method to store goods, where the consumer could send part of his endowment away for delivery next period. However, given the size of τ, we conjecture that this would not be an equilibrium outcome.

[11] Empirically, Hummels (2007), Djankov, Freund, and Pham (2010),Alessandria, Kaboski, and Midrigan (2010a) and Hummels and Schaur (2013) document the time-intensive nature of shipping goods across borders.

[12] In reality, a variety of payment arrangements are used. The payment structure in the model is known as "cash in advance." See Capela (2011) 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.

[13] Evidence on the international asset structure that agents have access to is limited, and what is available points towards financial autarky; see, for example, Heathcote and Perri (2002).

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 imports ordered at date t (i.e., imports delivered at $t+1$) must equal the value of exports shipped at date t. Thus, our model features a "dynamic trade balance" condition that holds for every pair of contiguous time periods:

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

where we denote imports and exports with the index of the time period at which they arrive to, or depart from, the home country, as it is done by statistical agencies in the U.S. Then, we have that trade is not balanced period by period as in static trade models, and, thus, our model features 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 now focus on the problem of a consumer located in a particular home country. The consumer faces the following dynamic programming problem:

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

$$\text{subject to} \quad p_x x + \tau p_y y' \leq p_x z, \;\; \text{where} \quad S = \{z, p_y, p_x\}.$$

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 which we discuss in more detail below. The consumer takes the aggregate prices p_{yt} and p_{xt} as given and forms expectations of future prices based upon knowledge of the law of motion of S. The other state variable is last-period's orders of imports arriving for consumption this period. 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 for S, 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 $g(S,\nu)$ in (8).

The idea here is to 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, without having to set up a general equilibrium model that generates an equilibrium process of output and prices as observed in the data. Because we are not interested in world-economy outcomes or in performing counterfactuals, this approach allows us to sidestep this more involved alternative.

3.1. Qualitative Features

In this section, we study some qualitative features of the model to develop the intuition behind the mechanism that underlies our quantitative results in Section 4.

Below we derive a dynamic import demand equation and then use this equation to walk through examples as to how and why imports respond to orthogonal endowment or price shocks. To ease the exposition, we restrict attention to the special case in which the shocks are orthogonal and the non-diagonal terms of the VAR are set to zero. In the quantitative analysis we relax these assumptions and estimate the unconstrained VAR and the covariance structure of the shocks directly from US data. The key result is this section is that the our model is able to deliver an income elasticity of imports greater than one and a price elasticity of imports below the calibrated elasticity of substitution.

There is one caveat we must mention up front. The ability of or model to deliver these results does not ensure that it will empirically. The ability of the model to replicate the patterns seen in U.S. data very much depends on the covariance structure of the shocks to endowments and prices and hence is an empirical question.

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.[14] This term induces the demand for imports to depend 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.[15] The reason is that good x is not a perfect substitute for good y. Thus, the elasticity of substitution across goods shows up with the intertemporal elasticity. The special case when x and y are perfect substitutes (implying $\theta = \infty$) illustrates this point, as only the parameter γ shows 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 $\tilde{m}_{t+1} = \beta$. The intuition behind this is that when $\gamma = \frac{1}{\theta}$, the benefits from smoothing consumption are offset by the imperfect substitutability of the good through which this smoothing is done, such that the agent is indifferent between current and next period's consumption (up to factor beta).

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

[14] 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}.

[15] 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.

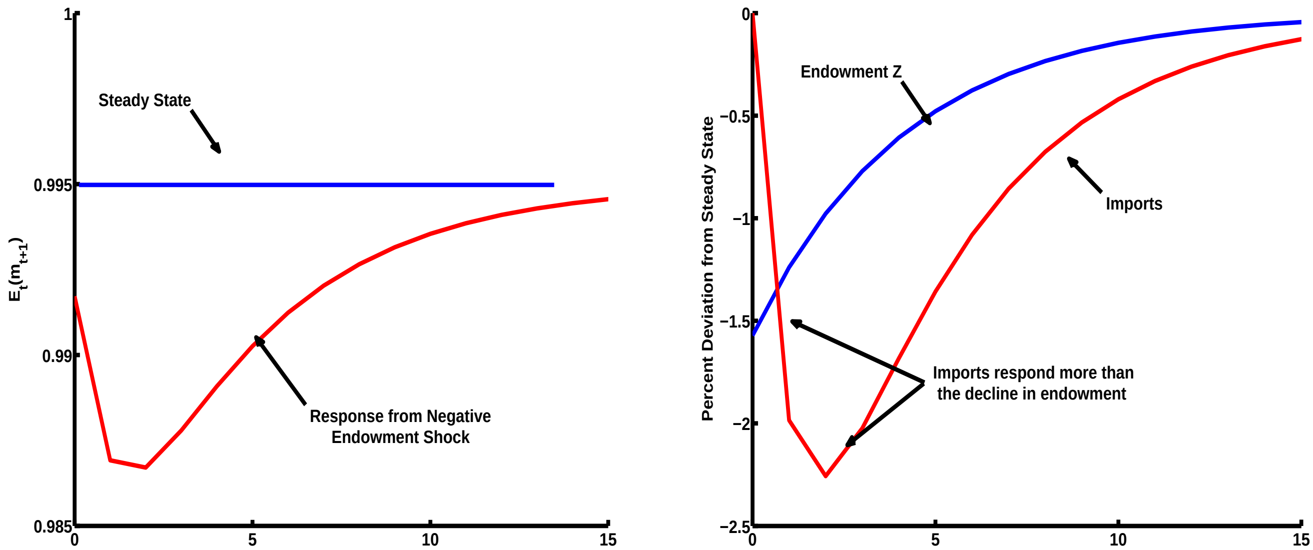

(a) Response in $E_t(\tilde{m}_{t+1})$ to a shock to z_t (b) Response in Imports to a shock to z_t

Figure 3: Imports Respond More Than Endowments.

Below, we describe how this mechanism can deliver an income elasticity greater than unity and an artificially low price elasticity. We focus on the empirically relevant part of the parameter space ($\gamma - \frac{1}{\theta} > 0$). We then restrict the stochastic process (for simplicity) to independent AR(1) processes describing the evolution of endowments and prices. As we discuss below, this simplifying assumption is not innocuous, however, it cleanly illustrates how the model works.

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 decrease. 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, marginal utility in that period is high. Moreover, because z_t is below its steady-state value, the agent rationally expects 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 $E_t(\tilde{m}_{t+1})$ lowers imports more than the drop in the endowment does.

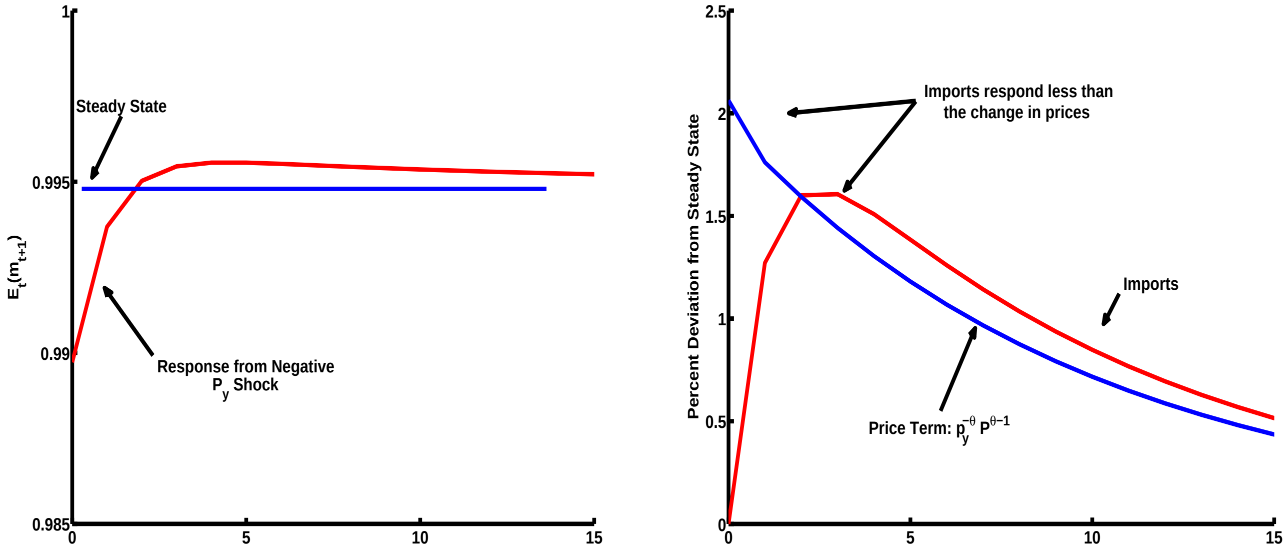

(a) Change in $E_t(\tilde{m}_{t+1})$ to a shock to p_{yt}

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

Figure 4: Imports Respond Less Than Prices.

Figures 3(a) and 3(b) trace these arguments out. Figure 3(a) plots the decline in $E_t(\tilde{m}_{t+1})$, when z_t is perturbed below its non-stochastic steady state and the economy is allowed to transit back. Figure 3(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 can also generate a price elasticity that is artificially below the elasticity of substitution θ. To gain 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.

As a result of the shock, imports increase by less than θ. The reason is that there are two forces working in opposite directions. First, a lower $p_{y,t}$ directly increases imports tomorrow by θ. This is the standard force from 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 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 4(a) and 4(b) illustrate these arguments. Figure 4(a) plots the decline in the stochastic discount factor when $p_{y,t}$ 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 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 4(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, potentially 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. 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 model's 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 that motivates our analysis is: How do the time-to-ship friction and finite intertemporal elasticity of substitution shape the decision to import, given a stochastic process describing output and prices as estimated from U.S. data?

To answer this question we estimate the endowment and price process from U.S. data and we then 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. In this section we focus specifically on how the model implied income and price elasticities depend on parameter values. The next section focuses on the ability of the model to account for the actual dynamics of U.S. imports.

4.1. Calibration

As a benchmark, we take a time period in the model to represent one quarter in the data, 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 it takes 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. We recognize that there is heterogeneity in the speed of the international delivery of goods, (see, e.g., Hummels (2007) and Hummels and Schaur (2013)) and, thus, we loosen this assumption in later in Section 4.3. Moreover, we provide additional evidence on time-to-ship in Section 6.

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 earlier, the case when $\gamma = \frac{1}{\theta}$ is interesting because $\tilde{m}_{t+1}$ is constant in this case 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) or Barro (2006)).

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 price elasticity estimated from simulated data 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 role 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 the discussion in Simonovska and Waugh (2014)).

The trade friction is calibrated so that imports are 15 percent of absorption. The top panel of Table 3 summarizes our calibration of preferences and technologies.

We estimate the stochastic process for endowments and prices (equation (9)) using HP filtered quarterly U.S. data on real absorption (z_t), absorption price index (p_{xt}), and the import price index (p_{yt}). 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 in the data (appropriately modified, as discussed below) and p_x to its associated price index.[16] Finally, import prices (appropriately modified, as discussed below) inform us about p_y. Appendix A provides the details on the construction of these series, and Appendix B

[16]An alternative 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.

Table 3: Summary of Calibration

Preferences and Technology	
Time Period	A quarter
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 — Time-to-ship model	
Transition Matrix	$A = \begin{pmatrix} 0.79^{***} & -0.14^{***} & 0.04 \\ 0.17^{***} & 0.94^{***} & -0.08 \\ -0.02 & 0.07^{***} & 0.72^{***} \end{pmatrix}$
Std. dev. of innovations	$\sigma_z = 0.014 \quad \sigma_{py} = 0.012 \quad \sigma_{px} = 0.004$
Corr. of innovations	$\text{corr}(z, p_y) = -0.09, \quad \text{corr}(z, p_x) = -0.24, \quad \text{corr}(p_y, p_x) = -0.15$
Estimated parameters for $\{z, p_y, p_x\}$ process — No time-to-ship model	
Transition Matrix	$A = \begin{pmatrix} 0.84^{***} & -0.11^{***} & -0.07 \\ 0.04 & 1.00^{***} & -0.69^{***} \\ 0.00 & 0.08^{***} & 0.77^{***} \end{pmatrix}$
Std. dev. of innovations	$\sigma_z = 0.011 \quad \sigma_{py} = 0.012 \quad \sigma_{px} = 0.003$
Corr. of innovations	$\text{corr}(z, p_y) = -0.09, \quad \text{corr}(z, p_x) = -0.37, \quad \text{corr}(p_y, p_x) = 0.07$

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

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 by consumers a quarter before. Thus, we adjust the data variables accordingly to make them consistent with the timing of our model. Appendix B provides the details.

The middle panel of Table 3 summarizes the estimated relationship between z_t, p_{yt}, p_{xt} corresponding to the timing in the time-to-ship model. Given that we contrast the outcomes of our model with those of a model without time-to-ship, we re-estimate the relationship between z_t, p_{yt}, p_{xt} corresponding to the timing in this static model, and use these estimates to compute its results. The bottom panel of Table 3 summarizes the estimated relationship for the no time-to-ship model.

Finally, in the simulations we measure output from our model to conform with National Income and Product Accounts (NIPA). That is, 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 from simulate time paths of imports, absorption, and prices from our model and then regressing imports on absorption and relative prices as described in equation (2). 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 1.16, 1.70, and 2.52. In the case of $\gamma = 5$, the income elasticity is near the 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 documented in the U.S. data outlined in observation **O.1**.

Table 4 also shows that the price elasticity estimated from simulated data lies below the true elasticity of substitution θ. Depending on the intertemporal elasticity of substitution, our model delivers a price elasticity of -1.07, -0.34, 0.64. 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 0.64. Again, 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.

We want to emphasize that there are several aspects of these results that are surprising and not predetermined. First, as discussed above, there exist covariance structures between endowments and prices that would not deliver these results. Thus, our model's ability to even qualitatively replicate observations **O.1** and **O.2** was not predetermined. It turns out, however, that the estimated empirical relationship between endowments and prices plus the economic environment are able to capture well the cyclical features of import data.

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

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.26	1.99
Model, $\gamma = 2$	-1.07 $[-1.16, -0.97]$	1.16 $[1.08, 1.24]$
Model, $\gamma = 5$	-0.34 $[-0.50, -0.17]$	1.70 $[1.55, 1.85]$
Model, $\gamma = 10$	0.64 $[0.37, 0.96]$	2.52 $[2.28, 2.75]$
Model, $\gamma = \frac{1}{\theta}$	-1.47 $[-1.54, -1.39]$	0.91 $[0.85, 0.95]$
Model, no time-to-ship	-1.50 $[-1.54, -1.47]$	1.00 $[0.98, 1.02]$

Note: Results are averages from 250 simulations, with each simulation being 187 periods long; values in brackets report 95-percent confidence intervals. Section 2 describes the data.

A second surprising result is that the same parameterization that generates an income elasticity that is close to the data is also the same parameterization generating a price elasticity close to the data. There is no reason to expect this outcome from our model. In other words, one mechanism—-systematic variation in the IMRS—*simultaneously* generates a high income elasticity and a low price elasticity.

The second to last row of Table 4 shows the results from the model when the intertemporal elasticity of substitution is set equal 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 effectively corresponds with the elasticity of substitution θ and we find an income elasticity close to unity, as in the static no time-to-ship model. This shows that the timing difference between the dynamic and static models is not driving the results.

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.

Table 4 reports results for an underlying elasticity of substitution, θ, set equal to 1.5. How do our results depend on this parameter? Appendix C presents the full details, but a summary is that increasing θ increases the income elasticity for all values of γ and results in larger biases in the estimated price elasticity. There are two reasons for this. First, an increase in θ changes the curvature on $\tilde{m}_{t+1}$ for a given γ parameter. Second, the import demand equation becomes

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

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$	Imports Volatility (%)
Data	-0.26	1.99	6.09
Model, $\varphi = 0$	0.64 $[\ 0.37,\ \ 0.96]$	2.52 $[2.28,\ 2.75]$	6.92
Model, $\varphi = 0.25$	0.14 $[-0.01, 0.32]$	2.21 $[2.07,\ 2.34]$	5.53
Model, $\varphi = 0.50$	-0.45 $[-0.52, -0.37]$	1.84 $[1.78,\ 1.90]$	4.18
Model, $\varphi = 0.75$	-1.04 $[-1.07, -1.00]$	1.44 $[1.41,\ 1.47]$	3.05

Note: Results are averages from 250 simulations, with each simulation being 187 periods long; values in brackets report 95-percent confidence intervals. Section 2 describes the data.

more sensitive to changes in $E_t(\tilde{m}_{t+1})$. Thus, these two forces result in movements in the IMRS playing a larger role in shaping import demand. Appendix C details the results.

4.3. Intermediate Time-to-Ship

The one-period shipping technology may be too stark of an 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.[17] 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 solve the problem in (8), with the modifications in (11). We then perform the same exercise described above. In the results, we focus on the case with $\gamma = 10$, varying the time-to-ship parameter φ.

Table 5 present the results. The second row reports the results with $\varphi = 0$, the one-period

[17] An alternative modeling strategy would be to allow for the "option" to pay for immediate delivery after endowments are realized as in Hummels and Schaur (2010).

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 a price elasticity close to the one in the data.

5. U.S. Import Fluctuations: 1967-2013

The results in Table 4 and 5 are encouraging in the sense that our model can easily account for cyclical properties of trade volumes in the U.S. with reasonable parameter values. This section focuses on the ability of our model to account for the actual dynamics of U.S. imports.

To discipline this exercise, we calibrate the model to match the empirical income and price elasticities observed in the data over the first half of our sample (i.e. the period from Q2 1967 to Q4 1989). We then ask: Can our model account for the dynamics of U.S. imports in the second half of our sample which includes the trade collapse (i.e. the period from Q1 1990 to Q4 2013)?

We answer this question by applying the Kalman smoother to the state-space representation of our model using U.S. data on absorption and prices over the entire Q2 1967-Q4 2013 time period. This allows us to generate the time series of imports implied by our model given the absorption and prices observed in the data, which we then contrast with the actual data on U.S. imports. Importantly, note that import data is only used to parameterize the model using data from first half of the sample—it plays no role in generating the implications of our model for the second half of the sample.

5.1. Calibration

As before, we estimate the the stochastic process for endowments and prices (equation (9)) using data on absorption and relative prices for the entire time period. Thus, the point estimates reported in Table 3 are the values used in this exercise. Estimating the stochastic process for only the second half of our sample yielded nearly identical results (in terms of calibrated parameter values, overall fit, trade collapse predictions, etc.) to those described below.

We choose the intermediate time-to-ship parameter to match the average number of days that it takes to ship goods into the U.S. from the time-to-ship data described in Section 6. Specifically, we compute the average time-to-ship across the different countries of origin, with observations weighted by the average value of imports across the sample. We find that the average time-to-ship is 33 days. From the lens of our model, we interpret this value as implying that almost two thirds of U.S. imports arrive contemporaneously, while a third of them arrive with a delay of

Table 6: Measures of Fit, Data and Model

Statistic	Data	Baseline Model	Static Model
Root Mean Square Error, Q1 1990 - Q4 2013	—	0.030	0.051
% Deviation From Trend Q1 2008 (Peak)	7.6	8.9	4.5
% Deviation From Trend Q2 2009 (Trough)	-22.1	-10.7	-0.46

Note: The first row presents the root mean square error between actual data and the model implied time series. The second row presents the percent deviation from trend for Q1 2008 and Q2 2009, the peak and the trough in data during the Great Recession.

one quarter, for an average of 33 days of time-to-ship delay. This leads us to set φ to 0.63.

We then calibrate γ and θ to match the empirical income and price elasticities over the first half of the sample: from Q2 1967 to Q4 1989. These elasticities are 1.93 and 0.40, respectively, over this period. We match these moments exactly with a parameterization that sets γ to 14.71 and θ to 1.53. This value of γ implies an intertemporal elasticity of substitution equal to 0.07, which is consistent with estimates from previous studies such as Hall (1988) and Vissing-Jorgensen (2002). Our estimated value of θ is very close to the value commonly used to calibrate international real business cycle models, as in Backus, Kehoe, and Kydland (1995).

5.2. Results

Figure 5(a) illustrates the results. It plots import data and predictions from our model for the entire time period Q2 1967-Q4 2013. Preference parameters were calibrated using data to the left of the vertical dashed line. The results to the right of the vertical dashed line correspond with the second half of the sample and, thus, can be viewed as an out of sample prediction of our model.

Our model performs very well. It tracks the data quite closely by capturing both the overall magnitude of fluctuations and the timing of peaks and troughs. In contrast to these outcomes, Figure 5(b) presents the results from the static model. As the figure illustrates, the static model does not fit the data well and it has severe problems regarding the timing and magnitudes in many instances.[18]

The first row in Table 6 provides a sense of fit. The first row reports the root mean squared

[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. suffered major shutdowns of U.S. ports due to dock worker strikes; see Isard (1975). Thus, these events appear to account for this discrepancy between the model and the data.

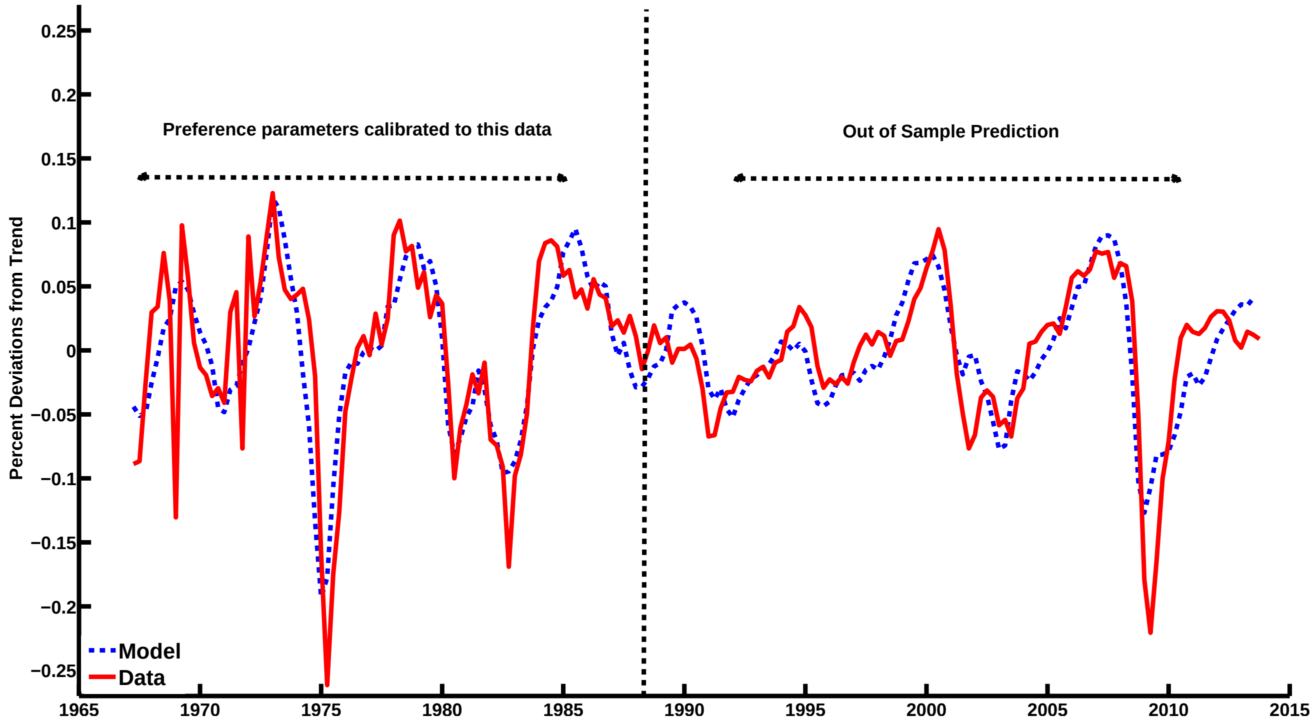

(a) **Model Prediction vs. Data**

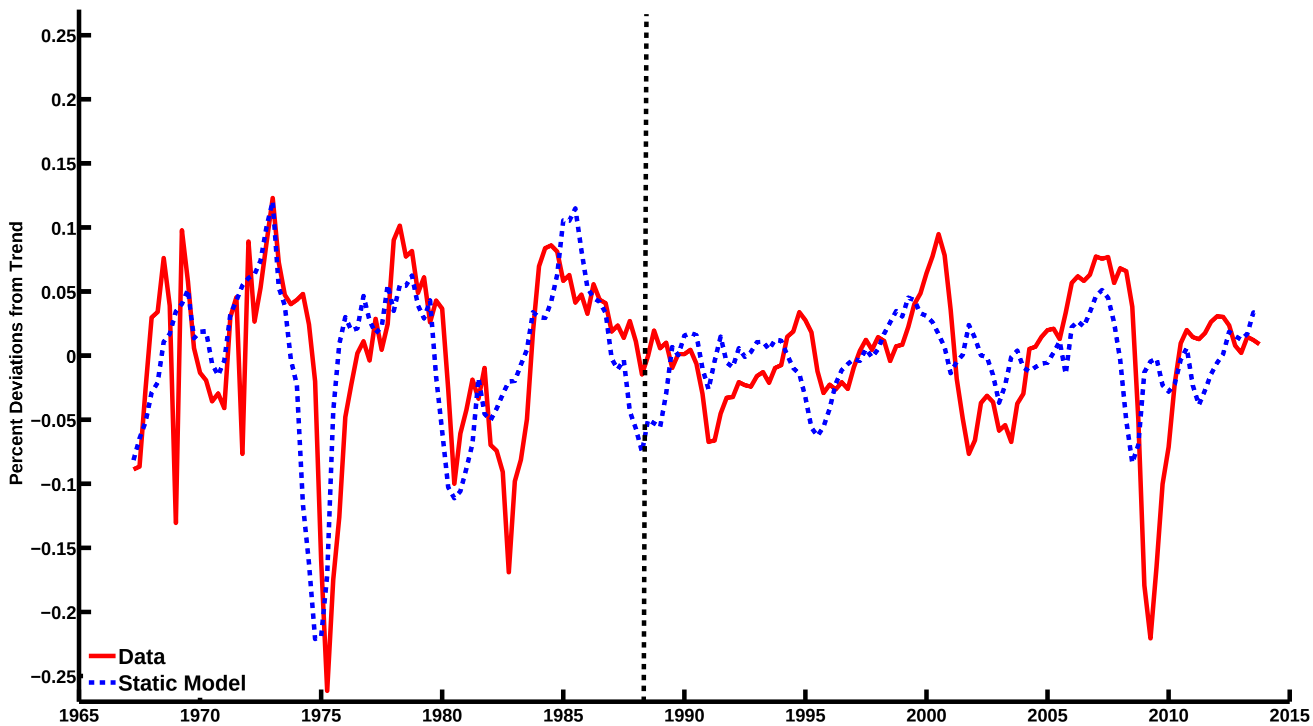

(b) **Static (No Time-to-Ship) Model Prediction vs. Data**

Figure 5: Model Predictions vs. Import Data, 1967-2011

error between the data for the second half of the sample and predictions from the model. The model's root mean squared error is 0.030, nearly 40 percent lower then the static model (0.051). As another point of comparison, the root mean squared error from the regression in (2) is 0.037. In other words, our calibrated model fits the data better than the best-fitting, linear regression of imports on absorption and prices. This shows that our model can account well for cyclical fluctuations in U.S. imports data.

Figure 6 zooms in on the last 14 years to illustrate the predictions from our model with respect to the 2008-2009 trade collapse. Our model 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.

The key failures of the baseline model regard the magnitude during the collapse and a one quarter miss in picking up the trough of the collapse. The last row of Table 6 reports the 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; the model, however, is only 11 percent below trend. Part of this miss is because the model predicts the trough to be Q1 2009 and 13 percent below trend. Overall, from peak to trough, our model predicts a 20 percent change in imports relative to the 30 percent change seen in the data.

There are meaningful ways of thinking about this discrepancy between the model and the data. Researchers have argued that there are important 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 (2012); inventory considerations discussed in Alessandria, Kaboski, and Midrigan (2010b); input-output linkages and vertical specialization discussed in Bems, Johnson, and Yi (2010); or shocks to the future value of manufactures as in Eaton, Kortum, Neiman, and Romalis (2013). Either (or all) 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 5(a) shows, our model can account well for cyclical fluctuations in U.S. imports for 40 years of data. Moreover, our model's ability to account for both the past 40 years and the most recent crisis is consistent with evidence regarding the sources of macroeconomic fluctuations. Stock and Watson (2012) use a dynamic factor model to analyze the contributions of various shocks to explaining the most recent and previous recessions. They find that the same factors that explained previous postwar recessions also explain the most recent recession.

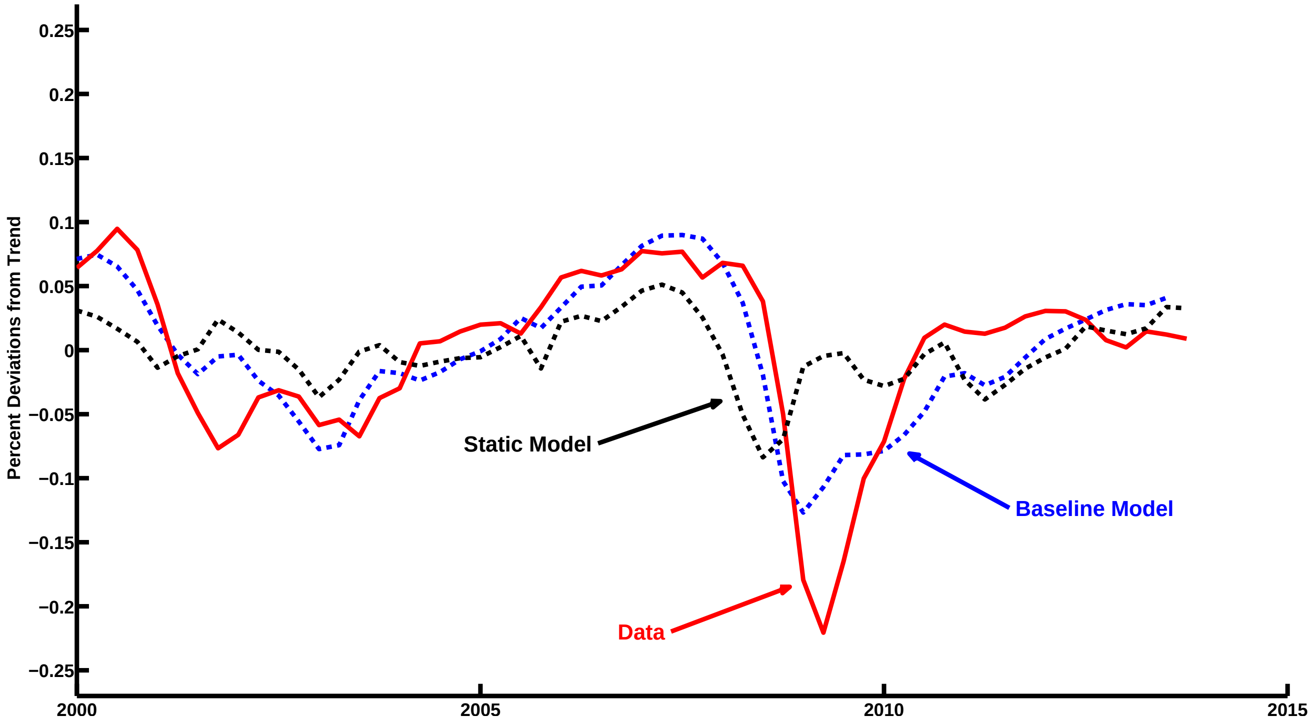

Figure 6: Trade Collapse: Model Predictions vs. Import Data, 2000-2014

6. Evidence: Time-to-Ship and Bilateral Import Volatility

This section examines some cross-sectional implications of our model. Our model predicts that a country's bilateral imports should be more volatile when sourced from a partner with longer shipping times. This implication is shown in the final column of Table 5; the volatility of nominal imports increases with the share $1 - \varphi$ (the effective length of time to ship) of imports that arrive in the following period. This implication is a "test" of our model because the static model (or the dynamic model with no active intertemporal marginal rate of substitution) predicts that the volatility of imports is independent of to the time-to-ship/distance.[19]

To explore this implication, we construct a measure of the time it takes to ship goods from a country of origin into the US by combining Hummels and Schaur's (2013) dataset on shipping times and the World Bank's Doing Business survey. Hummels and Schaur (2013) constructs a measure of the average time it takes to ship goods into the US from each country of origin, by mode of transportation, coast of arrival (east or west coast), and HS6 good categories. World Bank's Doing Business survey measures the "time necessary to comply with all procedures required to export goods" in each country of origin, as well as the "time necessary to comply with all procedures required to import goods" in the US. We construct our total measure of time-to-

[19]With ideal data (price indexes of US imports disaggregated by country of origin spanning a significant number of periods and countries), one could perform similar exercises to those in Sections 4 5 and compare and contrast the cross-sectional implications and/or try and "difference" out $E_t(\tilde{m}_{t+1})$ across destinations in equation 10. Unfortunately, these data are not available.

Table 7: Volatility of Imports by Time-to-Ship

Time-to-Ship	Imports Volatility (%)	# of countries
≤ 25 days	8.29	16
$25-50$ days	11.44	48
> 50 days	20.81	18

Note: Imports volatility measured as the standard deviation of the deviations of imports (log) around an HP-1600 trend. Observations weighted by the average imports volume across the time period.

ship by adding up these three variables.[20] We interpret this measure as the total additional time that it takes to purchase a good from international markets as opposed to domestic ones.

We compare these time-to-ship measures with quarterly data on US nominal imports disaggregated by country of origin. The data is obtained from the US Census, is not seasonally adjusted, and covers the period Q1 1992 - Q4 2013. For each country, the volatility of imports is computed as the standard deviation of the percentage deviation of imports around a Hodrick-Prescott trend with smoothing parameter set to 1600.

Table 7 reports the median volatility of imports with countries divided into three time-to-ship categories: (i) countries with average time-to-ship less than or equal to 25 days, (ii) between 25 and 50 days, and (iii) greater than 50 days. Observations are weighted by the average imports volume across the time period. Consistent with the implications of the model, countries with higher time-to-ship tend to have more volatile imports. Moreover, the magnitudes are economically significant: imports from countries with average time-to-ship greater than 50 days are more than twice as volatile as imports from countries with average time-to-ship under 25 days.

The first column of Table 8 formalizes this relationship by regressing import volatility on time-to-ship (in logs), weighting observations by the average imports volume across the time period. We find that the coefficient on our time-to-ship variable is statistically significant and economically large: a doubling of time-to-ship is estimated to increase the volatility of imports by 8 percentage points.

The second column of Table 8 studies the sensitivity of this relationship after controlling for variables that are used to explain bilateral trade flows: (i) a measure of the distance between

[20] Note that the average shipping times based on the Hummels and Schaur (2013) data are constructed by aggregating observations over the period 1991-2005, while the procedural times to export and imports from the World Bank's Doing Business dataset are constructed by aggregating observations over the period 2006-2013. Our results are robust to restricting attention to 2005 in the former dataset, and 2006 in the latter.

Table 8: Regression of Imports Volatility on Time-to-Ship

	Without controls	With controls
Time-to-ship (log)	0.080***	0.102***
Distance (log)	—	−0.023*
Common language	—	0.011
GDP per capita (log)	—	0.004
R-squared	0.18	0.20
Observations	82	82

Note: Imports volatility measured as the standard deviation of the deviations of imports (log) around an HP-1600 trend. Three asterisks denote statistical significance at 1% level, while one asterisk denotes significance at 10% level. Statistical significance based on heteroskedasticity-robust standard errors. Observations weighted by the average imports volume across the time period.

the different countries and the US, (ii) a dummy variable that is equal to one for countries with the same language as the US, and (iii) the countries' average level of GDP per capita over the sample period. Similar to the first column, the coefficient on our time-to-ship variable is statistically significant with and economically large.

Related findings have been previously documented in the literature. 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 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 findings.

Overall, we interpret these findings as evidence in support of the mechanism that we study in this paper. Standard models of international trade with static import decisions imply no systematic relationship between delivery times and import volatility. In the data, however, imports from countries with higher time-to-ship are systematically more volatile which is consistent with the cross-sectional predictions of our model.

7. Conclusion

Our paper shows how incorporating dynamic, forward-looking features into static 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 affects 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 presented 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*.

ALESSANDRIA, G., J. P. KABOSKI, AND V. MIDRIGAN (2010b): "The Great Trade Collapse of 2008–09: An Inventory Adjustment?," *IMF Economic Review*, 58(2), 254–294.

——— (2011): "US Trade and Inventory Dynamics," *American Economic Review*, 101(3), 303–07.

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.

BARRO, R. J. (2006): "Rare disasters and asset markets in the twentieth century," *The Quarterly Journal of Economics*, pp. 823–866.

BEMS, R., R. C. JOHNSON, AND K.-M. YI (2010): "Demand spillovers and the collapse of trade in the global recession," *IMF Economic Review*, 58(2), 295–326.

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. (2011): *Import/export for Dummies*. John Wiley & Sons.

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 (2012): "Off the cliff and back? Credit conditions and international trade during the global financial crisis," *Journal of International Economics*, 87(1), 117–133.

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 (2013): "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 (2011): "International trade in durable goods: Understanding volatility, cyclicality, and elasticities," *Journal of International Economics*, 83(1), 37–52.

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.

FITZGERALD, D., AND S. HALLER (2012): "Exporters and Shocks," Working paper.

HALL, R. E. (1988): "Intertemporal Substitution in Consumption," *The Journal of Political Economy*, 96(2), 339–357.

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

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 (2010): "Hedging price volatility using fast transport," *Journal of International Economics*, 82(1), 15–25.

HUMMELS, D., AND G. SCHAUR (2013): "Time as a Trade Barrier," *American Economic Review*, 103(7), 2935–59.

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.

KEHOE, T. J., AND K. J. RUHL (2013): "How important is the new goods margin in international trade?," *Journal of Political Economy*, 121(2), 358–392.

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 2008–09 crisis: in search of the smoking gun," *IMF Economic Review*, 58(2), 214–253.

LEVCHENKO, A. A., L. T. LEWIS, AND L. L. TESAR (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.

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

SIMONOVSKA, I., AND M. E. WAUGH (2014): "The elasticity of trade: Estimates and evidence," *Journal of International Economics*, 92(1), 34–50.

STOCK, J., AND M. WATSON (2012): "Disentangling the Channels of the 2007-2009 Recession," Discussion paper, National Bureau of Economic Research.

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

VISSING-JORGENSEN, A. (2002): "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *The Journal of Political Economy*, 110(4), 825–853.

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

Appendix: For Online Publication

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=2009) of goods gdp, final sales; Table 4.2.4 line 2 Fisher price index (100=2009) exports of goods; Table 4.2.4 line 54, Fisher price index (100=2009) of imports of nonpetroleum goods.
- **Quantity Indexes:** Table 1.2.3 line 5 Fisher quantity index (100=2009) of goods gdp, final sales; Table 4.2.3 line 2 Fisher quantity index (100=2009) exports of goods; Table 4.2.3 line 54, Fisher quantity index (100=2009) 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

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

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

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.26	1.99
Model, $\gamma = 2$	-2.49 [-2.70, -2.25]	1.77 [1.61, 1.92]
Model, $\gamma = 5$	-0.74 [-1.06, -0.35]	3.00 [2.68, 3.32]
Model, $\gamma = 10$	1.24 [0.66, 1.96]	4.73 [4.19, 5.30]
Model, $\gamma = \frac{1}{\theta}$	-4.06 [-4.30, -3.87]	0.92 [0.83, 0.99]
Model, no time-to-ship	-4.03 [-4.21, -3.88]	1.01 [0.93, 1.10]

Note: Results are averages from 250 simulations, with each simulation being 187 periods long; values in brackets report 95-percent confidence intervals.

minus exports, which gives

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

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,

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

Price indexes are adjusted similarly.

Table 10: Robustness, Orthogonal Endowment and Price Process

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.26	1.99
Model, $\gamma = 2$	-1.43 $[-1.53, -1.33]$	1.07 $[0.98,\ 1.17]$
Model, $\gamma = 5$	-1.32 $[-1.48, -1.17]$	1.49 $[1.33,\ 1.64]$
Model, $\gamma = 10$	-1.20 $[-1.45, -0.95]$	2.14 $[1.88,\ 2.39]$

Note: Results are averages from 250 simulations with each simulation being 187 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 9 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 10 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.