International Trade and Intertemporal Substitution

Fernando Leibovici

York University

Michael E. Waugh

New York University and NBER

June 2016

ABSTRACT

This paper studies the role of international trade delivery lags and variation in the intertemporal marginal rate of substitution in accounting for puzzling features of cyclical fluctuations of international trade volumes. 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 calibrate our model to match key features of U.S. data and discipline the variation in the intertemporal marginal rate of substitution using asset price data. We find that our model is quantitatively consistent with U.S. cyclical import fluctuations.

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

Standard trade models have difficulties explaining 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 is well above one, yet standard models imply a unitary income elasticity. Similarly, while the empirical elasticity of import volumes to measures of relative prices is well below one, typical calibrations 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]

We show that a time-to-ship friction to import goods and its interaction with a finite intertemporal elasticity of substitution can quantitatively rationalize these puzzling features of the data in an otherwise standard model of trade. Motivating the time-to-ship friction is the fact that international trade transactions involve nontrivial time lags between the order and delivery of goods. For example, Hummels and Schaur (2010) and Hummels and Schaur (2013) carefully document the time-intensive nature of trade and have shown how it shapes cross-sectional pattern of trade. Moreover, as we document below, the average time-to-ship of goods imported into the U.S. is 33 days, and is higher than a month and a half for a quarter of its trade partners.

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 the time-intensive nature of trade and variation in the intertemporal marginal rate of substitution breaks the unitary income elasticity, biases the estimated price elasticity, and shows up as a time-varying trade friction. Quantitatively, we find that this mechanism is able to account well for U.S. import fluctuations during both normal and crisis episodes.

We formalize these ideas by building a pure exchange Armington model of trade that is specifically designed to focus on the time-intensive nature of trade and variation in the intertemporal marginal rate of substitution. An agent within a country receives a stochastic endowment of its own nationally differentiated good. The aggregate consumption good is a composite of nationally differentiated goods, and the aggregator of the goods is of the constant elasticity of substitution class. Agents have time-separable preferences over an aggregate consumption

[1] Examples of models of this type are those of Krugman (1980), Eaton and Kortum (2002), Anderson and van Wincoop (2003), Melitz (2003), and 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. 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.

good. 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 the following period.

Our analysis proceeds in several steps. First, we *analytically* characterize the income and price elasticities in a special case of our model. Under plausible restrictions on preference parameters, we show that our model predicts an income elasticity of imports higher than unity and a price elasticity of imports lower than the elasticity of substitution between home and foreign goods.

Second, we quantify the model to assess the importance of this mechanism. The two key aspects of the calibration are how we discipline the behavior of the stochastic discount factor and the time-lag in imports. Our approach to disciplining the stochastic discount factor follows Cooper and Willis (2015) and infers the fluctuations in the stochastic discount factor that are attributable to the state variables of our model from data on asset returns. We then calibrate the preference parameters such that the model matches properties of the estimated stochastic discount factor. Our approach to disciplining the time-lag in imports is to allow a fraction of imports to arrive within the period, and the remaining fraction to arrive with delay. We then calibrate the fraction that arrive with delay to match the average shipping time in U.S. data.

Given the calibrated model, we first study the income and price elasticities implied by the model. To compute these measures, we simulate the model and estimate the model-implied income and price elasticities using simulated data, following the same approach that we use with real data. We find that the model quantitatively accounts for the high income elasticity and low price elasticity observed in U.S. time series data. Specifically, the implied income and price elasticities are 2.12 and -0.13; in the data, they are 1.99 and -0.26. Performing the same exercise, but removing the time-to-ship friction, we find the estimated income elasticity is effectively one and the price elasticity is the same as our calibrated elasticity of substitution.

We then focus on the model's ability to account for the actual time series of U.S. imports. To do so, we apply the Kalman smoother to compute our model's predicted import series given observed data on absorption and prices, 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 3(a).) The primary failure of the model concerns the trade collapse of 2008-2009. From peak to trough, our model predicts a 20 percent change in imports versus the 30 percent change seen in the data.

A fundamental feature of our analysis is the assumption that absorption and prices follow an exogenous stochastic process estimated to reproduce their dynamics in U.S. data. This approach allows us to sidestep well-known limitations of general equilibrium open economy models of business cycles to simultaneously account for the dynamics of prices and quantities. In our robustness section, we study the importance of our modeling approach within the context of the canonical international real business cycle model of Backus, Kehoe, and Kydland (1995) with time-to-ship. First, we show that the model of Backus, Kehoe, and Kydland (1995) with time-to-ship implies the same exact import demand equation that our analysis focuses on. However, when we calibrate this model, we find that the implied fluctuations in imports are very similar to those of a static trade model. The reason is that the Backus, Kehoe, and Kydland (1995) model has counterfactual predictions for the covariance of prices and absorption which counteracts movements in the stochastic discount factor. Because our model yields the same import demand equation but we estimate absorption and prices from data, we conduct our analysis in the context of a model that is consistent with the data along this dimension.

Finally, we provide evidence in support of the mechanism by examining some of the 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 predict that the volatility of imports is independent of time-to-ship and distance. Using data on shipping times constructed by Hummels and Schaur (2013) and the World Bank, we find that U.S. 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.

1.1. Related Literature

Our paper is motivated by a large literature, sparked by the collapse of trade during the 2008-2009 crisis, that has emphasized the limitations of standard models of international trade to account for cyclical trade fluctuations during both normal and crisis times (see, e.g. Engel and Wang (2011) and Levchenko, Lewis, and Tesar (2010a)).

We build on this literature by investigating the role of delivery lags and its interaction with a finite intertemporal elasticity of substitution. We interpret our findings as complementary to alternative mechanisms proposed to explain cyclical trade fluctuations as well as the collapse of trade during the 2008-2009 crisis. In particular, an active intertemporal marginal rate of substitution would surely amplify the role of financial frictions 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 (2015). The impact of vertical specialization, as investigated by Bems, Johnson, and Yi (2010), is also likely to be amplified by the interaction of time-to-ship with changes in the stochastic discount factor. Either (or all) of these mechanisms would complement our results and, perhaps, provide a complete account of trade fluctuations.

Closest to our work is Alessandria, Kaboski, and Midrigan (2010b), who study the role played by the adjustment of imported inventories on the collapse of U.S. imports during the crisis. In their model, imports are more inventory-intensive due to the combination of delivery lags and high fixed import costs. In contrast, we abstract from inventory considerations and nonconvex import costs and, instead, focus on the role that systematic variation in the stochastic discount factor plays in shaping fluctuations in trade.

Our findings also complement previous theoretical developments that emphasize the importance of dynamic considerations for understanding international trade, through the role of sunk costs as in Baldwin and Krugman (1989) or Alessandria and Choi (2007), as well as through the role of search and matching frictions specific to trade as in Drozd and Nosal (2012) and Eaton, Eslava, Krizan, Kugler, and Tybout (2009).

Finally, our paper is also related to a recent set of papers that emphasize the role that stochastic discount factors or "discount rates" play in shaping the business cycle properties of models of unemployment and investment. For example, in Hall (2014) and Kehoe, Midrigan, and Pastorino (2014) posted job vacancies are an investment by the firm and, thus, reductions in discount rates reduce the incentives to create a new job openings resulting in higher unemployment. The insight in this paper is similar; time-to-ship makes international trade look like an investment and, thus, variations in the discount rate applied to that investment affect the quantity demanded. The importance of the cyclically of discount rates has also been emphasized in the lumpy investment literature. In particular, we follow the Cooper and Willis (2015) in their approach to estimating the stochastic discount factor from asset return data. Consistent with Cooper and Willis (2015), Winberry (2015), and Beaudry and Guay (1996), we find a procyclical discount factor.

2. Cyclical Features of International Trade Volumes

In this section, we document 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 our quantitative exercise focuses on accounting for these features of the data.

Before proceeding, it is worthwhile to outline our language conventions. First, while we use the term "elasticity," the estimates we discuss are best thought of as simply summarizing the statistical properties of how imports, income, and prices behave in the time series. Second, although the measure of economic activity that we focus on is absorption, we use the terms income, output, and absorption synonymously throughout.[3]

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.[4] 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 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

[3]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.

[4]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).

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 (2015). 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), Engel and Wang (2011), and Bussiere, Callegari, Ghironi, Sestieri, and Yamano (2013), we perform the same analysis in Appendix C using only durable or non-durable goods.

Constructing real measures of these objects 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 the "Fisher of Fishers" 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.

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.

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 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.

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).[5] Marquez (2002) 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.

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.

[5] 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 explanations. 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.

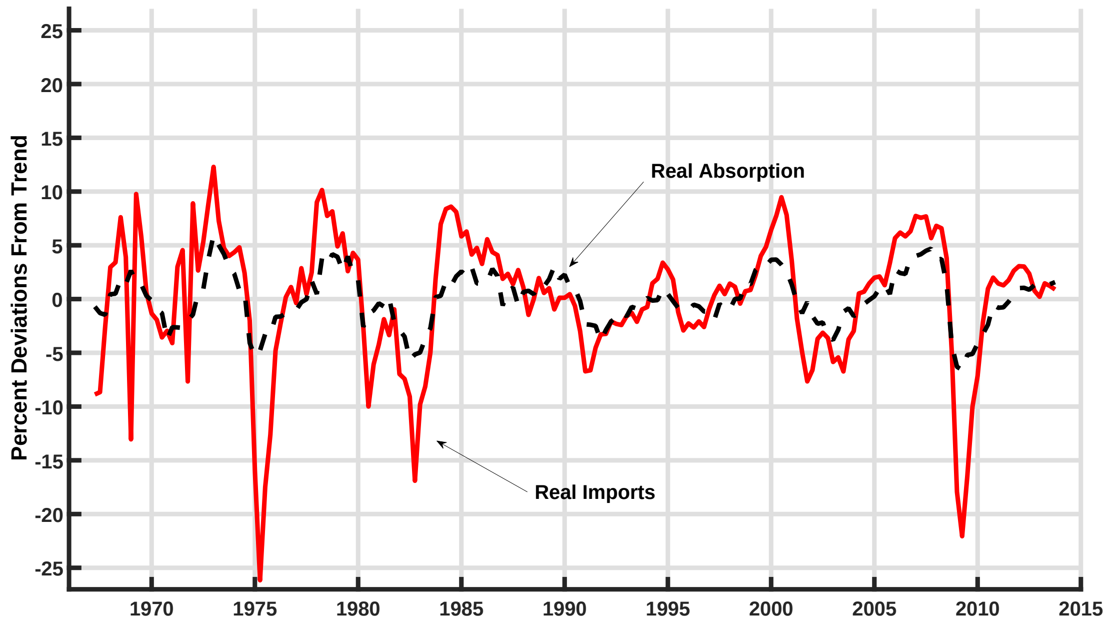

(a) **Real Absorption and Real Imports**

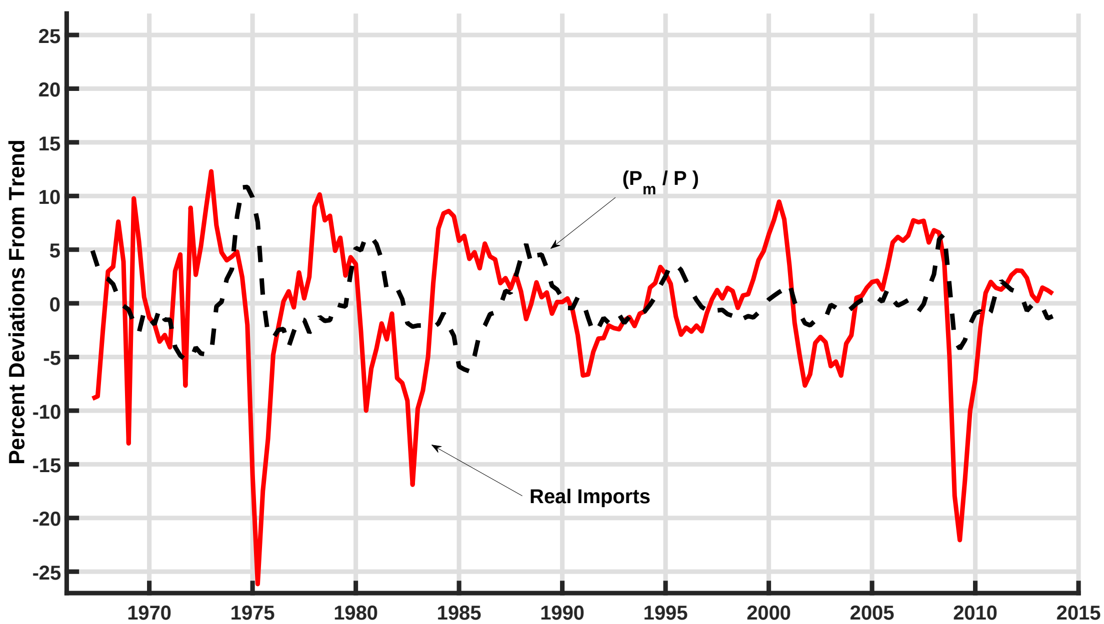

(b) **Relative Prices and Real Imports**

Figure 1: Absorption, Relative Prices, and Import Data

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.

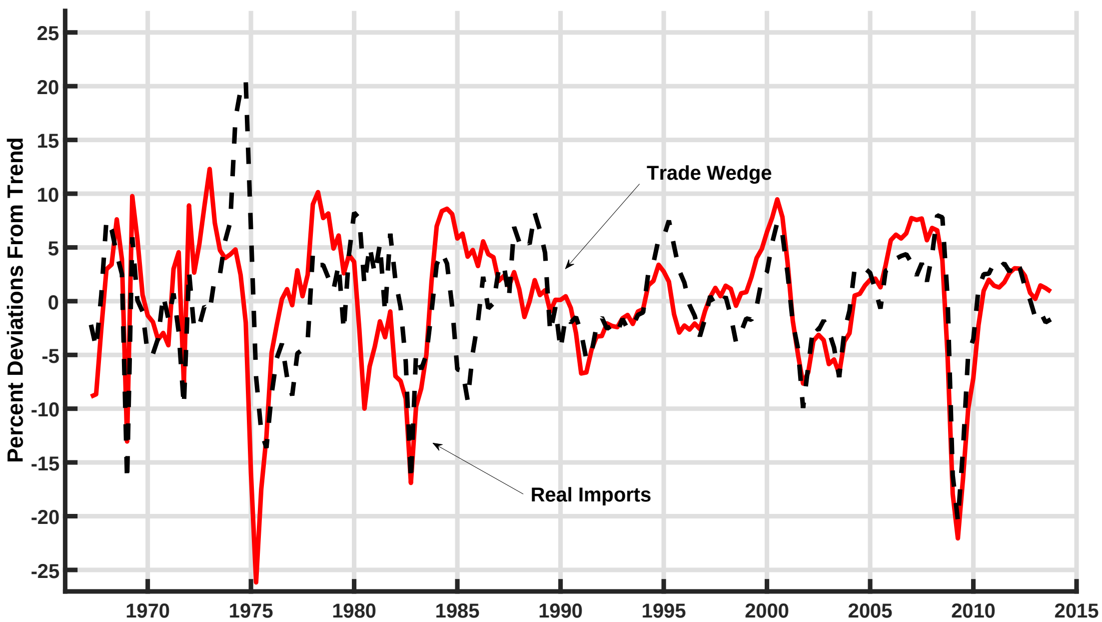

Figure 2: Wedges and Real Imports

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 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.

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

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

In the next sections, we argue that variation in intertemporal substitution can rationalize the features, **O.1-3**, of the data described above.

3. Model

This section describes a dynamic, pure exchange, Armington model of trade that is designed to answer the following quantitative question: given a stochastic process describing income and prices as in U.S. data, what are the implications for international trade volumes?

Countries and Preferences. We study a small open economy, which we refer to as home, populated by an infinitely-lived consumer who has utility function

$$\mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \, u(c_t) \tag{4}$$

where c_t is a consumption aggregate defined below, $\beta \in (0,1)$, and u is a strictly increasing, concave, and twice continuously differentiable period utility function. E_0 is the mathematical expectation operator conditional on information at date zero.

The composite consumption good c_t is a constant elasticity of substitution (CES) aggregate over two goods, x and y:

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

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

Endowments and Exchange. Every period, the consumer in the domestic economy receives a stochastic endowment of good x. This endowment can be either consumed domestically or used to acquire foreign goods y through a centralized international goods market. Goods cannot be stored.

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 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.[6] 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 foreign 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.[7] Underlying this international trading structure is an assumed enforcement technology that allows countries to coordinate the dynamic exchange of goods across international borders.

Then, the consumer in the home country faces the following budget constraint:

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

where p_{yt} and p_{xt} are the international prices that the consumer faces. The term z_t is the stochastic endowment of the x good that the domestic consumer 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.

We abstract from explicitly modeling trade in international financial assets. This does not imply that the agent is in financial autarky. As we discuss in Section (6.2) in the context of the general equilibrium model of Backus, Kehoe, and Kydland (1995), the proper interpretation of the endowment is that it includes any transfers of resources across countries. Thus, any trade in international financial assets are embedded in the stochastic endowment. Moreover, because we measure and estimate the stochastic endowment from data, our quantitative exercise will have taken into account trade in financial assets.

Defining the exogenous state variables as the vector $S = \{z, p_y, p_x\}$, we model the law of motion for S as a stationary VAR(1) process:

$$\log S_t = A \log S_{t-1} + \nu_t, \tag{7}$$

[6]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.

[7]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, evidence on the exact composition of payment structures of international trade is scarce. 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.

and the innovations ν_t are jointly normally distributed with mean zero and variance covariance matrix Σ.

The idea here is to model endowments and prices as following a stochastic process and estimate it from data. The motivation for this assumption is our desire to answer the following quantitative question in a simple and straightforward way: If an 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. In Section 6.2, we examine the implications of a general equilibrium extension of our model. Because we are not interested in performing counterfactuals, this approach allows us to sidestep this more involved alternative.

Therefore, the consumer solves the following dynamic programming problem:

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

$$\text{subject to } (6) \text{ and } \; S = \{z, p_y, p_x\}.$$

There are four state variables in the problem. The first three state variables are the consumer's endowment realization z and aggregate prices p_y and p_x which are summarized in the vector S. These evolve according to the law of motion described in (7). The other state variable is last-period's order 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.

3.1. 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 is expressed in equation 9 and summarized by Proposition 1.

Proposition 1 ***Dynamic Import Demand.*** *The demand for imports is given by:*

$$y_t = \frac{p_{xt} z_t}{P_t} \left[\frac{\tau}{E_t\left[m_{t+1}\right]} \frac{p_{yt}}{P_t} \right]^{-\theta} \tag{9}$$

where:

$$P_t = \left[p_{xt}^{1-\theta} + E_t\left[m_{t+1}\right] \left(\frac{\tau}{E_t\left[m_{t+1}\right]} p_{yt} \right)^{1-\theta} \right]^{\frac{1}{1-\theta}} \quad \textit{and} \quad m_{t+1} = \beta \frac{u'(c_{t+1})}{u'(c_t)}.$$

Equation (9) has several important features. 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) in a static trade model. This is a direct result of the time-to-ship assumption.

Second, $E_t(m_{t+1})$ enters equation (9). The term m_{t+1} is the intertemporal marginal rate of substitution (IMRS) or the stochastic discount factor (SDF). 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 stochastic discount factor is important for our purposes because it complicates the link between imports, endowments, and prices. That is, systematic variation in the SDF with endowments and prices implies that the model's income elasticity will differ from one and the price elasticity will differ from θ. Moreover, because $E_t(m_{t+1})$ shows up in the same places as the trade friction, variation in the SDF will appear as a time-varying trade wedge.

Finally, the relationship in (9) is not specific to our small open economy model, but belongs to a broader class of trade and international business cycle models. Specifically, we show in Section 6.2 that the dynamic demand function for imports in (9) is that in the international business cycle model of Backus, Kehoe, and Kydland (1995) with a time-to-ship friction.

3.2. Qualitative Results

Below we analytically derive the income and price elasticities for a special case of the model to develop intuition behind the mechanism driving the quantitative results in Section 5.

First, we focus on the case in which preferences are of the constant relative risk aversion type

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma},$$

where $1/\gamma$ is the intertemporal elasticity of substitution. Second, we restrict attention to the case in which the endowment and price shocks are orthogonal to each other and the non-diagonal terms of the VAR (equation (7)) are set to zero.

For the following discussion, define $\rho_z \in [0,1)$ as the auto-regression coefficient on the endowment process and $\rho_p \in [0,1)$ as the auto-regression coefficient on the price process.

3.2.A. Static Model

As a benchmark, we first characterize the income and price elasticities when the importing decision is not dynamic. That is, for a model in which import orders arrive for consumption within the same period. Proposition 2 summarizes the results.

Proposition 2 ***Income and Price Elasticities in the Static Model.*** *When the importing decision is not dynamic, the income and price elasticities are*

$$\frac{\texttt{Cov}\left(\log z_t, \log y_t\right)}{\texttt{Var}\left(\log z_t\right)} = 1, \tag{10}$$

$$\frac{\texttt{Cov}\left(\log p_{y,t}, \log y_t\right)}{\texttt{Var}\left(\log p_{y,t}\right)} = -\theta. \tag{11}$$

This proposition calculates theory-implied coefficients from a regression of log imports on endowments and prices. This says that if the static trade model is the data generating process, then the income elasticity that we estimate in the data should be one and the price elasticity should exactly reflect the elasticity of substitution across goods. As discussed in Section 2, these implications are inconsistent with the data.

3.2.B. Dynamic Model

Income Elasticity. For the dynamic model, we want to compute the same statistic in equation (10). To do so we make a couple of observations and then some approximating assumptions. First, notice that, in logs, imports measured at date t (that is, measured upon arrival in period t, ordered in period $t-1$) relate to objects at date $t-1$

$$\log z_{t-1} + \theta \log E_{t-1}(m_t) + \alpha_{1t} = \log y_t, \tag{12}$$

where α_{1t} is a collection of terms that we can abstract from since we are only focusing on the covariance of endowments and imports. The difficulty with this expression is that m_t depends on aggregate consumption and, in turn, endowments.

To provide some insight we assume that $c_t \propto z_t$. A rationalization for this approximation is that in the static model, aggregate consumption is proportional to the endowment realization. Thus, an interpretation of our approximation is that we are using analytical results from the static model to approximate aggregate consumption in the dynamic model. Numerical results show that this approximation is very good.

Given this assumption, we solve for $\log E_{t-1}(m_t)$ as a linear function of $\log z_{t-1}$ and constants which gives following import demand equation

$$\log z_{t-1} + \left[\theta\left(\rho_z - 1\right)\left(\frac{1}{\theta} - \gamma\right)\right]\log z_{t-1} + \alpha_{2t} \approx \log y_t, \tag{13}$$

where α_{2t} is a collection of terms we can abstract from and ρ_z is the auto-regression coefficient on the endowment process.

Endowment shocks have two effects on imports. The first term in (13) reflects the idea that when endowments increase, the agent one-for-one increases imports—this is the standard force in static trade models. Second, when endowments increase this affects the marginal utility of consumption at date t relative to $t-1$. A larger endowment today lowers the marginal utility of consumption today, increases the marginal utility of consumption tomorrow (because of the mean reversion in endowments). On net, this induces the agent to substitute the endowment into the future by importing more. This shows up in the second coefficient in (13). As long as the shocks are not permanent (i.e., $\rho_z < 1$) and that $\gamma > \frac{1}{\theta}$, then the coefficient on the second z_{t-1} term is positive. All together, this argument implies that a one percent increase in endowments results in a greater than one percent increase in imports.

Proposition 3 computes the regression coefficient of imports on endowments and summarizes the dynamic model's implications for the income elasticity.

Proposition 3 ***Dynamic Model: Income Elasticity Greater than Unity.*** *The income elasticity in the dynamic model is*

$$\frac{\texttt{Cov}\left(\log z_t, \log y_t\right)}{\texttt{Var}\left(\log z_t\right)} \quad \approx \quad \rho_z\left[1 + \theta\left(\rho_z - 1\right)\left(\frac{1}{\theta} - \gamma\right)\right] \quad > \quad 1. \tag{14}$$

with this elasticity being larger than in the static model as long as $\theta\gamma > \frac{1+\rho_z}{\rho_z}$.

The key result is that our economy can generate an income elasticity of imports that is greater than one. This is consistent with the data fact **O.1** in Section 2 and is distinct from the predictions of the standard static model outlined in Proposition 2. The key restriction is that the product of the elasticity of substitution across goods and the relative risk aversion parameter is larger than a function of the autocorrelation in endowments.

Price Elasticity. To compute the price elasticity in our dynamic model, we follow a similar approach as described above. In logs, imports depend on the price of imports in the following way

$$-\theta\log p_{y,t-1} + \theta\log E_{t-1}(m_t) + \alpha_{3t} = \log y_t, \tag{15}$$

where α_{3t} is a collection of terms that we can abstract from since we are only focusing on the covariance of relative prices and imports.

Again, the difficulty with this expression is that m_t depends on aggregate consumption and in turn, relative prices. We approximate m_t's dependence on $p_{y,t}$ by assuming that $c_t \propto p_{y,t}^{-\theta}$. Similar to the approximation for the income elasticity, a rationalization for this approximation is that in the static model, consumption is proportional to the relative price of imports adjusted by the elasticity of substitution. The weakness of this approximation is that it abstracts from the substitution pattern across goods that relative prices induce. However, we found numerically that it yields similar predictions to those from a full solution of the model.

Given this assumption, we solve for $\log E_{t-1}(m_t)$ giving the following import demand equation

$$-\theta \log p_{y,t-1} + (\theta\gamma - 1) \log p_{y,t-1} - \theta(\theta\gamma - 1) \log p_{y,t-2} + \alpha_{4t} \approx \log y_t, \tag{16}$$

where α_{4t} is a collection of terms we can abstract from. Equation (16) shows how there are two effects from a change in $p_{y,t-1}$ on imports. The first term in (16) is the standard effect in static trade models. If the price of imports decreases, then imports increase with elasticity θ. The second term in (16) is the dynamic effect.[8] A reduction in the price of imports makes agents wealthier in the future and this lowers their marginal utility of consumption tomorrow and, thus, induces the agent to intertemporally substitute a bit away from imports as long as $\theta\gamma > 1$.

Proposition 4 computes the regression coefficient of prices on endowments and summarizes the dynamic model's implications for the price elasticity.

Proposition 4 ***Dynamic Model: Low Price Elasticity.*** *The price elasticity in the dynamic model is*

$$\frac{\texttt{Cov}\left(\log p_{y,t}, \log y_t\right)}{\texttt{Var}\left(\log p_{y,t}\right)} \approx \rho_p \theta \left\{(\theta\gamma - 1) - 1\right\} - \rho_p^2 \theta(\theta\gamma - 1) \quad > \quad -\theta. \tag{17}$$

with this elasticity being greater than in the static model as long as $\theta\gamma > 1$ and $1 > \rho_p > 0$.

The important result here is that price elasticity of imports is artificially low (in absolute terms) relative to the elasticity of substitution θ. This is consistent with the data **O.2** in Section 2 which show that imports covary little with changes in relative prices. Similar to Proposition 3, the key restriction is that product of the elasticity of substitution across goods and the intertemporal elasticity of substitution is greater than one and that prices be positively autocorrelated.

[8] Why does p_{t-2} show up in this equation? Imports at date $t-1$ were decided upon at date $t-2$ and, thus, prices at date $t-2$ affect the marginal utility of consumption at date t.

4. Calibration

In this section, we describe how we calibrate our model. The quantitative question that guides our analysis is: How do the time-to-ship friction and finite intertemporal elasticity of substitution shape the dynamics of aggregate imports, given a stochastic process describing output and prices as estimated from U.S. data?

To answer this question we first parameterize the model by specifying the period utility function and a more realistic shipping technology. We then calibrate the parameters of the model using aggregate U.S. data to discipline the stochastic process for endowments and prices. And we use information on asset returns to discipline preference parameters and, thus, fluctuations in the stochastic discount factor.

4.1. Functional Form Assumptions

Preferences. For the quantitative analysis, we parameterize the period utility function to have external habits as in Campbell and Cochrane (1999). As shown above, CRRA preferences can deliver qualitative results consistent with the facts in Section 2. Quantitatively, however, CRRA preferences are unable to simultaneously have plausible calibrations for the intertemporal elasticity of substitution/risk aversion and match moments on the fluctuations in asset returns. Thus, we use preferences with external habits to satisfy both of these desires.

The preferences of the consumer are then represented by the following utility function

$$u(c_t, h_t) = \frac{(c_t - h_t)^{1-\gamma}}{1-\gamma}.$$

where h_t is a habit stock. The habit stock is parameterized to model the idea that current consumption is judged relative to past consumption. Specifically, define the surplus consumption ratio as

$$s_t = \frac{c_t - h_t}{c_t}, \tag{18}$$

and then define the law of motion for surplus consumption as

$$\log s_{t+1} = (1-\rho_s)\log \bar{s} + \rho_s \log s_t + \mu(s_t)\log\left(\frac{c_{t+1}}{c_t}\right) \tag{19}$$

$$\mu(s_t) = \frac{1}{\bar{s}}\sqrt{1 - 2(\log s_t - \log \bar{s})} - 1$$

where $\bar{s}$ is the steady-state level of surplus consumption and ρ_s denotes the persistence of surplus consumption. This is specification of habits follows Campbell and Cochrane (1999) and ensures that consumption is always larger than the habit stock. We further assume that the consumer does not take into account how their choice of consumption today affects their habit stock.

Shipping Technology. The one-period shipping technology used in Section 3 is too stark of an assumption for a model calibrated at 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.

To relax the one-period lag assumption we posit the following law of motion for how the consumption of imports relates to orders:

$$q(y_t, y_{t+1}) = \varphi y_{t+1} + (1-\varphi) y_t, \quad \text{and} \quad c_t = \left[x_t^\rho + q(y_t, y_{t+1})^\rho\right]^{1/\rho}, \tag{20}$$

where $q(y_t, y_{t+1})$ 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.[9] The law of motion in (20) 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.

4.2. Calibrating the Model

To calibrate the parameters of the model, we begin by partitioning the parameter space into three groups. The first group of parameters are set to standard values in the literature. The second group of parameters, specifically the share of imports that arrive with a lag and the stochastic process for endowments and prices, are estimated externally. The third set of parameters are calibrated to features of asset price data and the level of aggregate trade.

4.2.A. Predetermined Parameters

We begin by defining a time period in the model to represent one quarter in the data.

The set of predetermined parameters consists of risk aversion γ and the elasticity of substitution θ. The top panel of Table 3 summarizes our choice of predetermined parameters. We set $\gamma = 2$, which is a standard value commonly used to parameterize similar models.

The 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)).

[9] 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). Ravn and Mazzenga (2004), instead, assume that imports are shipped contemporaneously but chosen before the current shocks are realized.

Table 2: Stochastic Process for $\{z, p_y, p_x\}$

Transition Matrix	$A = \begin{pmatrix} 0.84^{***} & -0.12^{***} & 0.00 \\ 0.09^{**} & 0.98^{***} & -0.47^{***} \\ -0.01 & 0.08^{***} & 0.76^{***} \end{pmatrix}$
Std. dev. of innovations	$\sigma_z = 0.012 \quad \sigma_{py} = 0.012 \quad \sigma_{px} = 0.003$
Corr. of innovations	$\text{corr}(z, p_y) = -0.09, \quad \text{corr}(z, p_x) = -0.35, \quad \text{corr}(p_y, p_x) = -0.02$

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

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 will this value correspond with the price elasticity.

The top panel of Table 3 summarizes these choices.

4.2.B. Externally Estimated Parameters

We estimate the stochastic process for endowments and prices (equation (7)) via maximum likelihood 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.[10] 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 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

[10] 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.

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

Table 2 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 final parameter that is externally estimated is φ—the fraction of goods that arrive with delay. We estimate this parameter by calculating the average number of days that it takes to ship goods into the U.S. from the time-to-ship data described in Section 7. Average time-to-ship is computed 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 one quarter, for an average of 33 days of time-to-ship delay.[11] This leads us to set φ to 0.37.

4.2.C. Calibrated Parameters

The key set of parameters we calibrate are the discount factor β, the persistence of surplus consumption ρ_s, and the steady-state value of surplus consumption $\bar{s}$. These parameters are important because they control the properties of the stochastic discount factor and, thus, the behavior of the model. The calibration strategy we pursue is to infer the fluctuations in the stochastic discount factor from asset returns that are attributable to the state variables of our model. And then calibrate the preference parameters such that the model matches properties of the estimated stochastic discount factor.

To estimate the stochastic discount factor from asset returns, our strategy is to use the standard asset pricing condition that says that the product of the stochastic discount factor and the gross return on any asset must equal one in expectation. This condition then provides moment conditions and from which the stochastic discount factor can be estimated. Our implementation follows Cooper and Willis (2015) closely and the references therein.

The asset pricing condition we exploit says that given the gross return R^j_{t+1} on portfolio j and the stochastic discount fact the following condition must hold

$$\mathrm{E}_t\left(m_{t+1}R^j_{t+1}\right) = 1, \tag{21}$$

[11] This estimate is on the conservative end of 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.

where m_{t+1} is the same stochastic discount factor that shows up in equation (9). This simply follows from the first-order condition for the agents asset choice. Following Cooper and Willis (2015) and Zhang (2005), we assume that the stochastic discount factor relates to the state variables S_t and y_t in our model as:

$$\log m\left(S_t, y_t, S_{t+1}, y_{t+1}\right) = \alpha_0 + \boldsymbol{\alpha_1} \times \left[S_t - S_{t+1}, y_t - y_{t+1}\right]' \\ + \boldsymbol{\alpha_2} \times \left[(S_t - S_{t+1}) \times (S_t - \bar{S}),\ (y_t - y_{t+1}) \times (y_t - \bar{y})\right]', \tag{22}$$

where α_0 is a constant term, $\boldsymbol{\alpha_1}$ is a 4×1 vector of coefficients, $\boldsymbol{\alpha_2}$ is a 4×1 vector of coefficients, and variables with a bar are long-run averages. This specification follows Zhang (2005) and is motivated by our preference specification which display time-varying risk aversion.

Given the specification in (22), we estimate the αs using the moment condition implied by (21) via the Generalized Method of Moments. The interpretation of this procedure is that we are inferring the fluctuations in the stochastic discount factor from asset price data that are attributed to fluctuations in the state variables of our model.

The asset return data we use are the inflation adjusted returns on six different portfolios formed on size and book-to-market from Famma and French and the inflation adjusted rate on the 90-day Treasury bill. Data on the state variables are the same as described above. We implement the estimation by only estimating one coefficient for the terms of trade rather than a coefficient loading on each price separately. This reduces the number of parameters estimated to seven. Second, we use lagged values of the state variables as an instrument to provide additional moment conditions.

The first three moments in the lower right-hand part of Table 3 present the moments implied by the estimated (realized) stochastic discount factor. Consistent with the observed volatility in asset returns, our estimated stochastic discount factor is very volatile. Second, the realized stochastic discount factor is weakly correlated with absorption. We also find that the expected stochastic discount factor is positively correlated with absorption and imports at date $t + 1$.

Given the moments in the lower right-hand part of Table 3, we then pick the preference parameters such that the moments of the realized stochastic discount factor in our simulated model matches those in the data.[12] The final column in the right-hand part of Table 3 reports the model's fit. The one moment the model has difficulty matching is the correlation between absorption and the realized SDF. The lower left-hand part of Table 3 reports the calibrated parameters. The surplus consumption persistence and steady-state surplus consumption ratio are consistent with previous estimates in the literature by Campbell and Cochrane (1999), Verdelhan (2010), and Wachter (2006).

[12]Specifically, we choose these parameters to minimize the objective function MWM', where M is a row vector of the log-difference between each target moment and its model counterpart, and W is the identity matrix.

Table 3: Calibration of Time-to-Ship Model

Predetermined Parameters		**Externally Estimated Parameters**		
IES $1/\gamma$	0.50	Intermediate Time-to-Ship φ		0.37
Elasticity of Substitution θ	1.50	$\{z, p_y, p_x\}$ Stochastic Process		See Table 2
Calibrated Parameters		**Target Moments**	Data	Model
Discount Factor β	0.96	Avg. Realized SDF	0.988	0.984
Surplus Consumption Persistence ρ_s	0.99	Std. Dev. Realized SDF	0.212	0.217
Steady-State Surplus Consumption $\bar{s}$	0.11	corr(Realized SDF,Absorption)	0.104	0.255
Trade Cost τ	12.50	Imports/Absorption	0.219	0.219

Note: Realized SDF is estimated using (21), (22) and inflation adjusted Famma and French portfolios and the 90-day Treasury bill for the time period Q1, 1967 - Q3, 2013.

The final parameter that we calibrate is the trade cost τ. We pick this value to target the ratio of imports to absorption. The final line in Table 3 summarizes the result.

5. Quantitative Results

We perform two exercises with our quantitative model. First, we compute the model implied income and price elasticities by simulating the model and estimating the regression in (2) on simulated data. Second, we apply the the Kalman smoother to obtain 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.

5.1. Import Elasticities

To compute the model implied income and price elasticities, we simulate time paths of imports, absorption, and prices from our model, measuring absorption 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. Then, we use the simulated data to regress imports on absorption and relative prices as described in equation (2).

The results are reported in Table 4. The first row replicates the empirical income and price elasticities seen in Table 1. The middle row reports the results from our time-to-ship model, while the bottom row reports the results for the (static) model without time-to-ship.

Table 4 shows that our model implies an income elasticity greater than unity. Moreover, with an income elasticity of 2.12, our model can more than fully account for the difference between

Table 4: Import Elasticities

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.26 $[-0.44, -0.08]$	1.99 $[1.77, 2.20]$
Time-to-Ship Model	-0.13 $[-0.24, -0.01]$	2.12 $[2.00, 2.25]$
No Time-to-Ship Model	-1.50 $[-1.50, -1.50]$	1.00 $[1.00, 1.00]$

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.

the unit income elasticity implied by standard models and the 1.99 income elasticity estimated from U.S. data and 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. At -0.13, this elasticity is meaningfully below the calibrated elasticity of substitution of -1.5. Moreover, this result is quantitatively consistent with the observed price elasticity in the data. This result shows that variation in the stochastic discount factor is strong enough to rationalize the low price elasticity outlined in observation **O.2**—even though the true elasticity of substitution is -1.5.

The final row shows the results when dynamics are turned off. This is a standard static 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.

We want to emphasize that there are two aspects of these results that are surprising and not predetermined. First, the same parameterization that generates an income elasticity that is close to the data also delivers 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.

Second, as we discuss in Section 6, 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 both qualitatively and quantitatively.

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

To provide a comparison between observed data and our model, we apply 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 obtain 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.

Figure 3(a) illustrates the results. It plots import data and predictions from our model for the entire time period Q2 1967-Q4 2013. 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 3(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.[13]

The first row in Table 5 provides some metrics of fit. The first row reports the root mean squared error between the data and predictions from the model. The model's root mean squared error is 0.036, close to 35 percent lower than the static model (0.056). As another point of comparison, the root mean squared error from the regression in (2) is 0.037. That is, our calibrated model fits the data better than the best-fitting, linear regression of imports on absorption and prices.

The key failures of the model concern the magnitude during the trade collapse of 2008-2009. The last row of Table 5 reports the percent deviation from trend during Q2 2009, which is the trough of the cycle. In the data, imports were 22.1 percent below trend; the model, however, is only 11.5 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 (2015). Either (or all) of these mechanisms would complement our results and, perhaps, provide a complete account of the drop in trade.

[13] 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. We thank George Alessandria for pointing us in this direction.

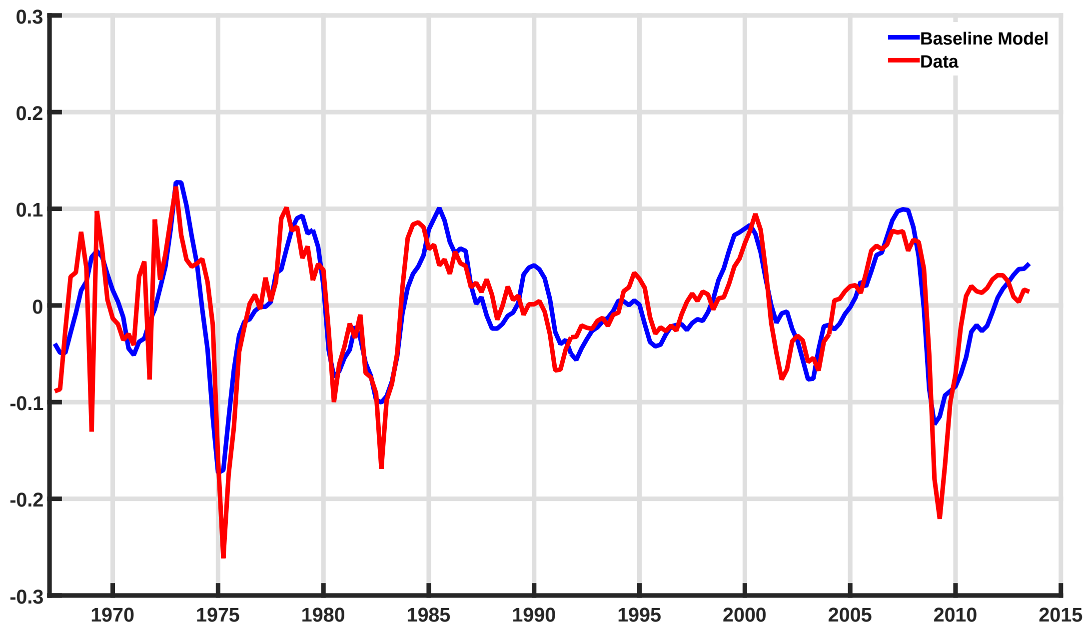

(a) **Baseline Model vs. Data**

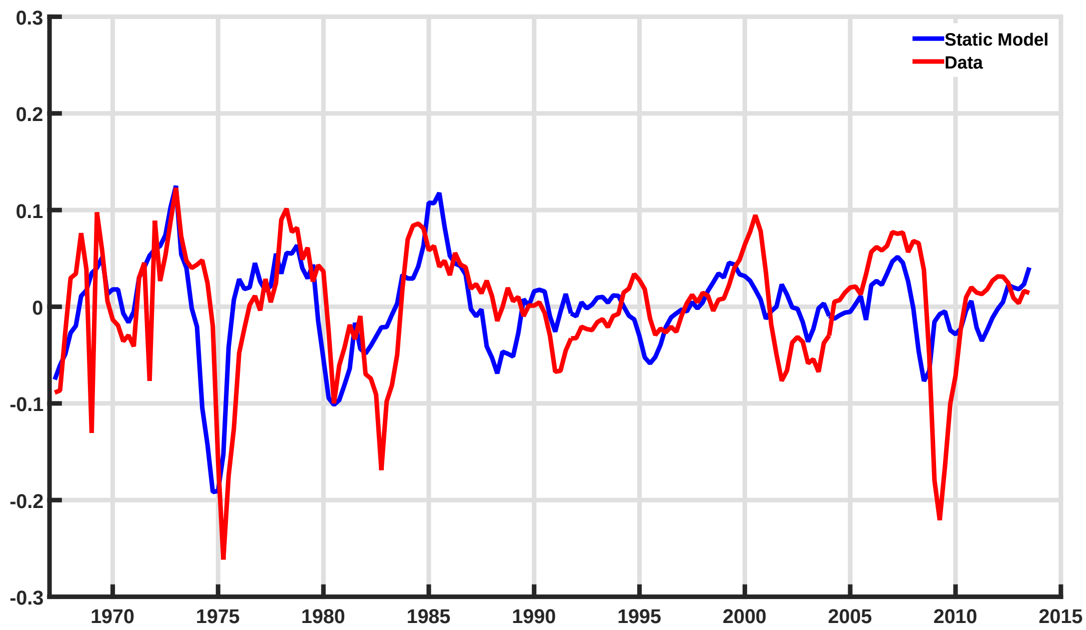

(b) **Static (No Time-to-Ship) Model vs. Data**

Figure 3: Model vs. Data, 1967-2014

Table 5: Measures of Fit, Data and Model

Statistic	Data	Baseline Model	Static Model
Root Mean Square Error, Q2 1967 - Q4 2013	—	0.036	0.056
% Deviation From Trend Q1 2008 (Peak)	6.8	8.1	-0.2
% Deviation From Trend Q2 2009 (Trough)	-22.1	-11.5	-0.8

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.

6. Robustness

This section examines the sensitivity of our findings to alternative modeling assumptions. Specifically, we investigate the role played by two key ingredients assumed throughout our analysis: habit preferences, and the exogenous process of absorption and prices estimated from U.S. data.

6.1. No Habits

Throughout our analysis in the previous sections, we introduced preferences with external habits to discipline the model's implications for the dynamics of the stochastic discount factor, a key ingredient of our mechanism. We now investigate the importance of this assumption for the quantitative implications of the model.

To do so, we recompute the model under the constraint that $h_t = 0$, while keeping all parameter values as in our baseline parametrization. We then simulate the model and use the simulated data to estimate the implied price and income elasticities following the approach of the previous section.

The results for the income and price elasticities are presented in the third row of Table 6 (in the first two rows we report the elasticities in the data and our baseline model, respectively). We find that the income and price elasticities are now significantly closer to the model without time-to-ship than to the empirical elasticities and our baseline model.

Indeed, the time-to-ship model without habits features a realized SDF that is one-tenth as volatile as our calibration target (0.023 vs. 0.212). In contrast, with habits, this moment is almost almost identical to the empirical target.[14] This shows that fluctuations in the stochastic discount factor, through external habits, play a key role in mapping the time-to-ship technology

[14]The other targeted moments do significantly change when habits are removed.

Table 6: Sensitivity Analysis, Import Elasticities

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.26 $[-0.44, -0.08]$	1.99 $[1.77,\ 2.20]$
Time-to-Ship Model	-0.13 $[-0.24, -0.01]$	2.12 $[2.00,\ 2.25]$
Time-to-Ship Model without Habits	-1.43 $[-1.45, -1.41]$	1.10 $[1.08,\ 1.12]$
BKK (1995) with Habits and Time-to-Ship	-1.51 $[-1.80, -1.26]$	1.08 $[0.96,\ 1.16]$

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.

into income and price elasticities that are close to their empirical counterparts.

6.2. Endogenous Absorption and Prices

Another fundamental feature of our modeling strategy is the assumption that absorption and prices follow an exogenous stochastic process estimated to reproduce their dynamics in U.S. data. This approach allows us to sidestep well-known problems or anomalies in general equilibrium open economy models of business cycles, such as Backus, Kehoe, and Kydland (1995). To evaluate the importance of this modeling approach, we contrast features of international trade fluctuations implied by our model with those of a standard general equilibrium international business cycle model extended to feature external habits and time-to-ship.

The general equilibrium model that we study is the two-country model of Backus, Kehoe, and Kydland (1995) and Heathcote and Perri (2002), with representative households having access to an internationally-traded one-period risk-free bond, and international trade subject to trade frictions modeled as an iceberg trade cost. In each of the countries, households consume domestic and foreign goods that are produced by representative firms that operate constant-returns-to-scale technologies that require labor and physical capital. Business cycle fluctuations are driven by Hicks-neutral country-specific shocks to the productivity of these representative firms. Time-to-ship and habits are modeled as above. For simplicity, we restrict attention to an economy with one-period time-to-ship (i.e. all orders of imports take one period to arrive).[15]

[15]With intermediate time-to-ship, the small deviations that we find between the model with endogenous ab-

See the Online Appendix for further details on the economic environment.

Before discussing the quantitative results, it is important to note that the general equilibrium IRBC model that we study has a dynamic import demand equation that is identical to the one implied by our model (equation 9). Proposition 5, summarizes this result.

Proposition 5 ***Isomorphism in Dynamic Import Demand.*** *In the model of Backus, Kehoe, and Kydland (1995) and Heathcote and Perri (2002) with time-to-ship, the demand for imports by final good producers in the home country is given by:*

$$y_{t+1} = \frac{p_{x,t} z_t}{P_t} \left[\frac{\tau}{E\left[m_{t+1}\right]} \frac{p_{y,t}}{P_t} \right]^{-\theta},$$

where:

$$P_t = \left[p_{x,t}^{1-\theta} + E\left[m_{t+1}\right] \left(\frac{\tau}{E\left[m_{t+1}\right]} p_{y,t} \right)^{1-\theta} \right]^{\frac{1}{1-\theta}}.$$

In Proposition 5, the variables and parameters have the same interpretation as in our model. However, in contrast to our model, z_t, $p_{x,t}$, and $p_{y,t}$ are endogenous objects that are determined in equilibrium by the structure of the economy. In particular, z_t consists of the sum of consumption and investment expenditures, while prices $p_{x,t}$ and $p_{y,t}$ are such that the domestic and foreign good markets clear. For a proof of the proposition and further details on these variables, see the Online Appendix.

Given the above proposition, this economy is just a general equilibrium extension of our model. Thus, we use it to evaluate the extent to which our findings depend on the estimation of absorption and prices from the data, as opposed to allowing absorption and prices to constitute equilibrium outcomes of the model.

To contrast the implications, we parameterize preferences and technologies following Backus, Kehoe, and Kydland (1995), except for β, τ, and the habits parameters which are calibrated following the exact same approach described in Section 4.

The bottom row of Table 6 presents the results. The main finding is that the income and price elasticities change substantially and are now almost the same as in the static model without time-to-ship—and far from the data. In a sense, this result mimics the finding in Backus, Kehoe, and Kydland (1994) who found little quantitative impact for business cycle statistics in their benchmark model with time-to-ship.

sorption and prices and its static counterpart would be further diminished.

Table 7: Business-Cycle Dynamics of Absorption and Prices

	corr(z,p_y)	corr(z,p_x)	corr(p_y,p_x)
Data	-0.17	-0.59	0.42
BKK (1995) with Habits and Time-to-Ship	0.15	0.23	-0.80

Why is there a discrepancy between our approach which estimates absorption and price from data and the general equilibrium BKK (1995) model? First, what it is not. Proposition 5 tells us it is not about a difference in the demand function for imports. Second, it is not because the stochastic discount factor has different properties. The results presented in Table 6 arise when the general equilibrium model is calibrated so that the properties of the stochastic discount factor are the same in both models.

The difference between our approach and the general equilibrium model is that the Backus, Kehoe, and Kydland (1995) model has counterfactual implications for the joint dynamics between absorption and prices. To show this, we compute the pairwise correlations among $\{z_t, p_{x,t}, p_{y,t}\}$, and contrast them with those implied by our model (which display the same dynamics as in U.S. data). We report these findings in Table 7.

The correlations implied by the BKK (1995) model are qualitatively and quantitatively different to the correlations across these variables observed in the data. In particular, we find that the BKK (1995) model implies a positive correlation between absorption and the price of the domestic good, which is negative in the data. Similarly, the model implies a negative correlation between the price of the domestic and foreign goods, which is positive in the data. Moreover, we find that absorption and import prices are positively correlated in the BKK (1995) model, while they are negatively related in the data.

These counter-factual dynamics reduce the impact of the time-to-ship friction on the implied trade elasticities. For instance, as in our model, an increase of absorption leads to a more than one-to-one increase of imports due to its impact on the SDF. As agents are better off today, they increase imports as a way to shift consumption into the future. However, the positive correlation between absorption and the price of imports implies that these effects are muted. The increase in the price of imports reduces the attractiveness of imports as a way to smooth consumption across time, both directly and through its impact on the SDF. Altogether, the counterfactual dynamics of absorption and prices lead to implied income and price elasticities that are far from those that we estimate in the data and close to those featured by static trade models.

Table 8: 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.

7. 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. Specifically, the model implies that 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 the time-to-ship/distance.[16]

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) construct 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-ship by adding up these three variables.[17]

We compare these time-to-ship measures with quarterly data on US nominal imports disaggre-

[16]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 5 5.2 and compare and contrast the cross-sectional implications and/or try and "difference" out $E_t(\tilde{m}_{t+1})$ across destinations in equation 9. Unfortunately, these data are not available.

[17]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 9: 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.

gated 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 8 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 9 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 9 studies the sensitivity of this relationship after controlling for variables that are commonly used to explain bilateral trade flows: (i) a measure of the distance between 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 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.

8. 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

ALESSANDRIA, G., AND H. CHOI (2007): "Do sunk costs of exporting matter for net export dynamics?," *Quarterly Journal of Economics*.

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.

BACKUS, D. K., P. J. KEHOE, AND F. E. KYDLAND (1994): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?," *American Economic Review*, 84(1), 84–103.

BALDWIN, R., AND P. KRUGMAN (1989): "Persistent Trade Effects of Large Exchange Rate Shocks," *The Quarterly Journal of Economics*, 104(4), 635–654.

BEAUDRY, P., AND A. GUAY (1996): "What do interest rates reveal about the functioning of real business cycle models?," *Journal of Economic Dynamics and Control*, 20(9), 1661–1682.

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.

BUSSIERE, M., G. CALLEGARI, F. GHIRONI, G. SESTIERI, AND N. YAMANO (2013): "Estimating Trade Elasticities: Demand Composition and the Trade Collapse of 2008-2009," *American Economic Journal: Macroeconomics*, 5(3), 118–51.

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.

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.

COOPER, R., AND J. WILLIS (2015): "Discounting: Investment Sensitivity and Aggregate Implications," Working paper.

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.

DROZD, L. A., AND J. B. NOSAL (2012): "Understanding international prices: Customers as capital," *The American Economic Review*, pp. 364–395.

EATON, J., M. ESLAVA, C. J. KRIZAN, M. KUGLER, AND J. TYBOUT (2009): "A search and learning model of export dynamics," *Penn State University mimeo*.

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

EATON, J., S. KORTUM, B. NEIMAN, AND J. ROMALIS (2015): "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.

HALL, R. E. (2014): "High Discounts and High Unemployment," Discussion paper, National Bureau of Economic Research.

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.

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, P., V. MIDRIGAN, AND E. PASTORINO (2014): "Debt constraints and employment," .

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.

RAVN, M. O., AND E. MAZZENGA (2004): “International business cycles: the quantitative role of transportation costs,” *Journal of International Money and Finance*, 23(4), 645–671.

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

VERDELHAN, A. (2010): “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *The Journal of Finance*, 65(1), 123–146.

WACHTER, J. A. (2006): “A consumption-based model of the term structure of interest rates,” *Journal of Financial economics*, 79(2), 365–399.

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

WINBERRY, T. (2015): “Lumpy Investment, Business Cycles, and Stimulus Policy,” Discussion paper.

ZHANG, L. (2005): “The value premium,” *The Journal of Finance*, 60(1), 67–103.

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

$$\begin{aligned}
\text{Adjusted Absorption} &= \underbrace{\underbrace{p_{xt}x_t + p_{yt-1}y_{t-1}}_{C_t} + \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} - \underbrace{p_{yt-1}y_{t-1}}_{Imports_t}}_{GDP_t} - \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} + \underbrace{p_{yt}y_t}_{Imports_{t+1}} \quad (23) \\
&= p_{xt}x_t + p_{yt}y_t = p_{xt}z_t,
\end{aligned}$$

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 10: 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

$$\begin{aligned}
\text{Domestic Consumption} &= \underbrace{\underbrace{p_{xt}x_t + p_{yt-1}y_{t-1}}_{C_t} + \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} - \underbrace{p_{yt-1}y_{t-1}}_{Imports_t}}_{GDP_t} - \underbrace{p_{xt}(z_t - x_t)}_{Exports_t} \qquad (24)\\
&= p_{xt}x_t,
\end{aligned}$$

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,

$$\text{Adjusted Imports}_t = \varphi \text{Measured Imports}_t + (1 - \varphi)\text{Measured Imports}_{t+1}. \qquad (25)$$

Price indexes are adjusted similarly.

C. 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), Engel and Wang (2011), and Bussiere, Callegari, Ghironi, Sestieri, and Yamano (2013). 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 11 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.[18]

The durable trade wedge still accounts for a lot of the variation in imports. Figure 4 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 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

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

Table 11: 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.

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.[19] Alessandria, Kaboski, and Midrigan (2010b) make this argument while studying the decline in trade flows during the 2008-2009 crisis; Alessandria, Kaboski, and Midrigan (2010a) argue that inventory considerations are important for understanding the dynamics of devaluations. An implication of Alessandria, Kaboski, and Midrigan's (2010b) model is that the regression equation in (2) should be augmented with the change in imported inventories (Alessandria, Kaboski, and Midrigan (2011) provide this derivation).

We followed this argument by augmenting the regression in (2) by including data on the real change in private inventories as an additional explanatory variable. Separate information on changes in inventories of imported goods is unavailable. The third row in Table 11 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

[19] 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.

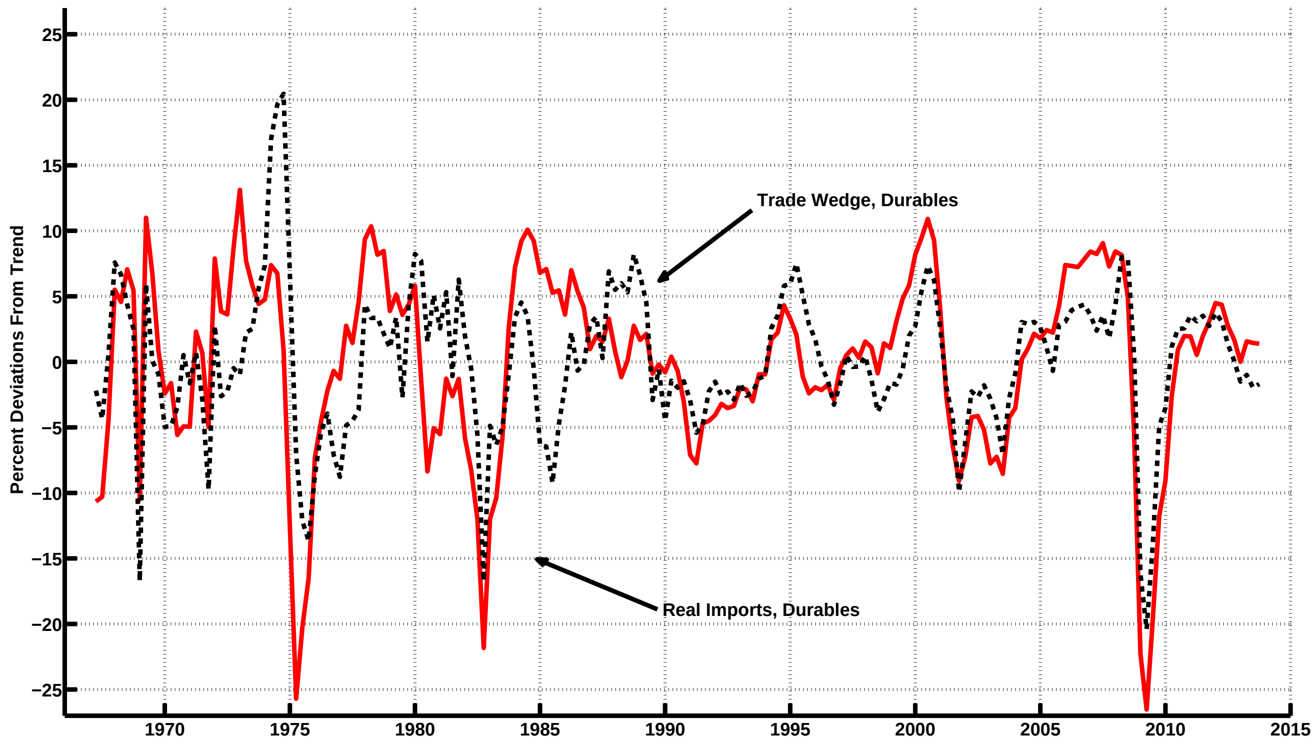

Figure 4: Wedges and Real Imports, Durables Only

elasticities observed at cyclical frequencies.