

International Trade and Intertemporal Substitution

Fernando Leibovici

Federal Reserve Bank of St. Louis

Michael E. Waugh

New York University and NBER

December 2018

ABSTRACT

This paper quantitatively investigates the extent to which variation in the intertemporal marginal rate of substitution can help account 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 use a standard small open economy model with time-intensive international trade, calibrated to match key features of U.S. data and disciplining the variation in the intertemporal marginal rate of substitution using asset price data. We find that variation in the intertemporal marginal rate of substitution helps rationalize puzzling features of import fluctuations and that this mechanism is quantitatively important during both normal and crisis times.

Email addresses: fleibovici@gmail.com, mwaugh@stern.nyu.edu. Fernando Leibovici gratefully acknowledges the financial assistance provided by the Social Sciences and Humanities Research Council of Canada (SSHRC) prior to his employment at the Federal Reserve Bank of St. Louis. The views expressed in this paper are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors. Code replicating the results of the paper is posted at https://github.com/mwaugh0328/trade_intertemporal_substitution.

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

This paper quantitatively explores the idea that variation in the intertemporal marginal rate of substitution can help account for these puzzling features of cyclical fluctuations of international trade volumes. The idea follows from two natural ingredients. First, time-to-ship and upfront payment frictions make the importing decision dynamic because resources today must be sacrificed for the delivery of goods tomorrow. Motivating the time-to-ship friction are previous studies which document that international trade transactions involve nontrivial time lags between the order and delivery of goods; previous studies also show that a nontrivial fraction of imports are paid before delivery, motivating the upfront payment frictions.³

Second, 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. Together, these ingredients lead to the insight that 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.

Our analysis proceeds in several steps. First, we outline the key facts behind fluctuations in trade volumes. While these results are not entirely new, we extend previous results from the literature to control for alternative mechanisms that we abstract from. Furthermore, these results provide the foundation for our quantitative analysis.

Second, we develop a relatively standard small open economy model. A country specializes

¹Examples of models of this type are those of Armington (1969), 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).

²Houthakker 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 (2010), 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.

³For example, Hummels and Schaur (2010) and Hummels and Schaur (2013) carefully document the time-intensive nature of trade and have shown how it shapes the cross-sectional pattern of trade. Similarly, Antras and Foley (2015), International Monetary Fund (2009), Asmundson, Dorsey, Khachatryan, Niculcea, and Saito (2011), International Monetary Fund (2011) document evidence on the payment arrangements under which import transactions are typically carried out.

in the production of a given good, and imports a good of a different type from the rest of the world. Domestic and imported goods are aggregated to produce a final good used for consumption and investment, and the aggregator of the goods is of the constant elasticity of substitution class. Agents take prices as given, and we model the evolution of relative prices and aggregate productivity as following a joint stochastic process. The key departure of our model from otherwise standard frameworks is that agents must commit resources today for the delivery of imported goods in the following period.

Third, we characterize the income and price elasticities in a special case of our model. We show that our model can predict an income elasticity of imports higher than unity and a price elasticity imports lower the elasticity of substitution between home and foreign goods. Moreover, these results motivate our calibration strategy by showing how the income and price elasticities of imports depend on the stochastic process of the intertemporal marginal rate of substitution.

Fourth, we quantify the model to assess the importance of this mechanism. The three key aspects of the calibration are how we discipline the behavior of the stochastic discount factor, the time-lag in imports, and the stochastic process that governs the dynamics of productivity and prices. 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. To discipline the dynamics of the stochastic discount factor we study an economy with external habits as in Campbell and Cochrane (1999); we 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. Finally, we estimate the joint process for productivity and prices to match salient features of the joint dynamics of prices and absorption observed 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 can quantitatively account, to a large extent, for the high income elasticity and low price elasticity observed in U.S. time series data. Specifically, the implied income elasticity is 1.25 v.s. 1.63 in the data; the implied price elasticity is -0.30 in the model and the data even though the structural elasticity of substitution is -1.30. 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 investigate 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).) In particular, the model accounts for a large share of the trade collapse of 2008-2009.

Fifth, we illustrate some key ingredients to our quantitative results. One important ingredient is having the model account for the observed dynamics of the stochastic discount factor. Habits and capital adjustment costs are key tools that allow our model to generate fluctuations in the stochastic discount factor and, in turn, deliver high income elasticities and low price elasticities. Another key ingredient to our findings is the existence of a mismatch between the timing of import payments and import deliveries such that a fraction of imports are paid before delivery. Third, we show that estimating the joint process for productivity and prices to match features of the dynamics of absorption and prices in the U.S. is also key to our results; our model implies income and price elasticities close to a static model under alternative parameterizations of the stochastic process. Finally, we also show that a two-country version of our model with endogenous prices implies counter-factual dynamics of absorption and prices and, thus, income and price elasticities close to a static model.

We conclude by providing 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. Our model is consistent with this evidence; 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 2010).

We build on this literature by investigating the role of delivery lags and payment frictions, and their 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 (2016). In contrast to studies in which interest rates affect trade flows via frictions in credit markets, our mechanism features a role for interest rates to affect trade even without financial constraints. 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. Similarly, Backus, Kehoe, and Kydland (1994a) and Ravn and Mazzenga (2004) have previously investigated the role of delivery lags on international business cycles; in contrast, we study their impact on the dynamics of international trade flows using a model that matches key features of asset prices and the joint dynamics of absorption and the price of traded goods.

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). Particularly related to this paper are Alessandria, Pratap, and Yue (2015), who emphasize the dynamic nature of export entry decisions in shaping the dynamics of aggregate exports in response to aggregate shocks.

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 (2017) and Kehoe, Pastorino, and Midrigan (2016) 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 cyclicity of discount rates has also been emphasized in the lumpy investment literature. In particular, we follow 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 salient 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 low frequencies; Ruhl (2008) on the low price elasticity; and Jacks, Meissner, and Novy (2009) and Levchenko, Lewis, and Tesar (2010) on the wedge analysis. However, summarizing these three features of the data is important since our quantitative exercise focuses on accounting for them.

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.⁴

To summarize the statistical properties of imports, income, and prices in U.S. time series data, we use a log-linear relationship relating imports to prices and income. The rationale for using this relationship comes from standard models of international trade based on CES preferences or production functions.⁵ In these models, the demand function for imports is given by:

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

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 σ 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

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

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 they allow us to examine the extent to which standard models deviate from the relationship observed in the data.

⁴Absorption is gross domestic product (GDP) plus imports minus exports. Because static trade models typically feature balanced trade, absorption corresponds with income. Hence, we use absorption and income synonymously.

⁵We 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).

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

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

This exercise is similar to that of Jacks, Meissner, and Novy (2009), and Levchenko, Lewis, and Tesar (2010). Following the arguments of Chari, Kehoe, and McGrattan (2007), this exercise is meaningful because *systematic* deviations between theory and data shed light on mechanisms through which underlying primitives operate. Specifically, if ω_t varies systematically with the business cycle, then this suggests that: (i) there are economic forces that are not reflected in equation (1); and (ii) any new mechanism posited to explain these deviations should operate through the wedge. We set $\sigma = 1.5$, which is a standard calibration of this parameter in the international business-cycle literature. Using larger σ s, 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 indexes. 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 of 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 (2016). 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) as emphasized by Boileau (1999), Engel and Wang (2011), and Bussiere, Callegari, Ghironi, Sestieri, and Yamano (2013), we extend the analysis below to control for differences in the share of durable goods between absorption and imports.

Constructing real measures of these objects and their associated price indexes is not as straightforward as it might seem. Real values in the U.S. NIPA are chain-type indexes constructed using an “ideal” chain index advocated by Fisher (1922). While these indexes have desirable proper-

Table 1: Empirical Price and Income Elasticities

Data	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$	R^2	# Obsv.
Goods GDP	-0.24 [-0.41, -0.07]	2.00 [1.80, 2.21]	0.65	205

Note: Data are in logs and HP filtered over the time period from Q2 1967 to Q2 2018. Values in brackets report 95-percent confidence intervals.

ties, 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 second quarter of 2018. To deal with trends, we HP-filter the logarithm of these data with smoothing parameter 1600. Results using log-first-differences yield no significant differences.


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

Figure 1(a) and 1(b) plot the data. 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 equal to 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.

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).⁶ Marquez (2002)

⁶Prominent 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 part of the high income elasticity. Ruhl (2008) shows that this margin is not quantitatively important at business-cycle fre-



(a) Real Absorption and Real Imports



(b) Relative Prices and Real Imports

Figure 1: Absorption, Relative Prices, and Import Data

examines this feature of the data from a modern perspective and finds that it is robust to alternative econometric specifications, different frequencies, and commodity disaggregation.

O.2. Low price elasticity. Our second observation is that the estimated import price elasticity is -0.24 . Figure 1(b) illustrates this finding. It plots the percent deviations from trend of relative prices $\frac{p_{m,t}}{P_t}$ and import data. Notice that prices and imports weakly correlate with each other negatively and, in some instances, even move in the same direction. Thus, the low price elasticity in Table 1 is not a surprise.

While modelers have a choice over this parameter, typical parameterizations of static trade models or international business-cycle models generally use values 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. 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

high R^2 . 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.

slope coefficient of 0.69 and an R^2 of 0.43. 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, the wedge must then decrease to rationalize the drop in imports. Furthermore, relative prices decreased and imports did not increase as predicted by the standard model—this reflects the low price elasticity. When the price elasticity is constrained to take on a standard value, this implies that the wedge must decrease even more. Thus, the fact that imports over-respond to income and under-respond to prices manifests itself as a pro-cyclical wedge when events like the 2008-2009 recession are analyzed in the context of a standard trade model.

2.3. Composition and Inventories

Recent papers on the collapse of trade during the 2008-2009 crisis have raised two issues: compositional mismatch between imports and absorption and the role of inventories. Here, we control for differences in the composition of absorption and imports as well as account for the dynamics of inventories.

Composition. One concern is that observations **O.1-3** reflect compositional effects of the sort described by many authors, e.g., Boileau (1999), Engel and Wang (2011), Eaton, Kortum, Neiman, and Romalis (2016), Levchenko, Lewis, and Tesar (2010); Bems, Johnson, and Yi (2013) provides a nice summary of compositional issues.

There are really two compositional issues. The first is the distinction between goods (of which imports are mostly comprised of) and services (which play a large role in GDP). As discussed above, we controlled for this issue by focusing on goods GDP rather than simply using aggregate GDP. The top row of Table 2 illustrates the importance of this choice. It shows that the income elasticity when we use aggregate GDP is three—50 percent larger than our baseline estimate.

The second composition issue is between durable and non-durable goods. 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 the absorption of durable goods is more volatile than that of total goods, suggests that an income elasticity larger than unity or pro-cyclical trade wedges may arise because of differences in durability between imports and goods GDP.

To explore this issue, we created a “synthetic” data series in which the composition of absorption matches the composition of imports. Specifically, we re-weight the individual components of absorption according to the observed trade weights.⁷ As an extreme example, if imports are only composed of durables, but absorption is some mix of durables and non-durables, then this re-weighting scheme matches up trade with only durable absorption. This scheme is closely related to the adjustments used in Levchenko, Lewis, and Tesar (2010). We perform a similar adjustment for the aggregate relative prices using the “Fisher of Fishers” approach described above. Alternative re-weighting schemes for prices yield similar results.

We then estimate equation (2) with our composition-adjusted measure of absorption and relative prices. The third row of Table 2 presents our results. It shows that the income elasticity of imports is well above one even after controlling for differences in the composition between imports and absorption.

Taken together, the top three rows of Table 2 suggest that compositional issues are important to understanding the high income elasticity observed at cyclical frequencies. With that said, even after controlling for composition, the income elasticity of imports is still well above one. In contrast, compositional issues are unable to explain away the low import price elasticity with all specifications resulting in estimates which are similar to that found in Table 1.

Inventories. Another concern is that **O.1-3** arises because we abstract from changes in inventories. 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 import dynamics in devaluations. An implication of Alessandria, Kaboski, and Midrigan’s (2010b) model is that regression equation (2) should be augmented with the change in imported inventories (Alessandria, Kaboski, and Midrigan 2011 provide this derivation).

We followed this argument by augmenting equation 2 by including data on the change of real private inventories as an additional explanatory variable.⁸ Separate information on changes in inventories of imported goods is unavailable. The third row in Table 2 reports the results. After controlling for changes in inventories, the income elasticity is 1.70, relative to 2.00 without controlling for inventories. These results suggest that inventory adjustments are a partial—but

⁷Using fixed weights leads to similar results. For example, setting the durable expenditure share to the average value in the trade data of 60 percent results in an income elasticity of 1.72.

⁸Normalizing the inventory series by the stock of inventories yielded nearly identical results.

Table 2: Price and Income Elasticities — Durables and Inventories

	Price Elasticity	Income Elasticity	Inventories	R^2
GDP	−0.19 [−0.30,0.01]	2.95 [2.69,3.21]	−	0.71
Baseline (Goods GDP)	−0.24 [−0.41,−0.07]	2.00 [1.80,2.21]	−	0.65
Controlling for Durables	−0.31 [−0.50,−0.13]	1.92 [1.73,2.12]	−	0.66
Controlling for Inventories	−0.25 [−0.40,−0.10]	1.70 [1.49,1.90]	0.03 [0.02,0.04]	0.72
Durables and Inventories	−0.30 [−0.46,−0.14]	1.63 [1.44,1.82]	0.03 [0.02,0.04]	0.73

Note: Data are in logs and HP filtered over the time period from Q2 1967 to Q2 2018. Values in brackets report 95-percent confidence intervals.

not a complete—explanation of the high income elasticities observed at cyclical frequencies.

The final row of Table 2 reports the estimation results controlling for the composition of imports and inventory adjustments. Consistent with the features of the data emphasized above, **O.1-3**, we estimate an empirical income elasticity equal to 1.63 and a price elasticity equal to -0.30 .

In the next sections, we propose an additional mechanism that can help rationalize the features, **O.1-O.3**. To focus on the mechanism we propose, our model abstracts from durable imports and inventories. Thus, in our quantitative analysis, we contrast the implications of our model with the empirical elasticities reported in the fifth row of Table 2, which control for these two un-modeled channels.

3. Model

This section describes a dynamic Armington model of international trade with endogenous production under international financial integration that is designed to answer the following quantitative question: Given a stochastic process estimated to match salient features of income and prices in the U.S., what are the implications for international trade volumes?

We study a small open economy, which we refer to as “home,” that is populated by four agents: a representative consumer, a representative final good producer, a representative producer of a

domestic good (x). In addition to the home country, there is the rest of the world. The rest of the world produces a foreign good (y) that is imported by the domestic economy, and which is used along with good x to produce final goods. Final goods are used for consumption and investment and are not traded internationally.

Household. The representative household is infinitely lived, and has preferences characterized by utility function

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, \ell_t), \quad (4)$$

where c_t denotes consumption of final goods, ℓ_t is the amount of leisure, $\beta \in (0, 1)$, and u is a period utility function that is strictly increasing, concave, and twice continuously differentiable in each of its arguments. \mathbb{E}_0 is the mathematical expectation operator conditional on information at date zero.

The household has two sources of income. On the one hand, the household owns the factors of production, capital and labor, which it supplies to producers of domestic good x , for which it gets compensated at wage rate w_t and rental rate of capital $r_{k,t}$. The household is endowed with a unit of time and chooses the fraction n_t to supply to the market as labor ($n_t = 1 - \ell_t$). On the other hand, the household owns the producers of final and domestic goods, from whom it earns profits (or finances losses) Π_t every period.

Earnings are used for two purposes. First, the household chooses the amount of final goods c_t to consume every period. Second, the household chooses how many final goods to invest i_t to increase its capital stock in period $t + 1$. In particular, capital accumulation is subject to a standard law of motion with linear depreciation rate $\delta \in (0, 1)$: $k_{t+1} = (1 - \delta)k_t + i_t$.

Finally, the household has access to international financial markets where it trades a one-period risk-free bond b_{t+1} . Following Heathcote and Perri (2002) these bonds are denominated in units of good x at interest rate r_t .

The household's problem can then be represented as:

$$\max_{\{c_t, n_t, i_t, k_{t+1}, b_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, 1 - n_t) \quad (5)$$

subject to

$$p_t c_t + p_t i_t + \frac{p_{x,t} b_{t+1}}{1 + r_t} = p_t w_t n_t + p_t r_{k,t} k_t + p_{x,t} b_t + \Pi_t$$

$$k_{t+1} = (1 - \delta)k_t + i_t,$$

where $p_{x,t}$ and p_t denote the price of good x and final goods, respectively.

International Trade. International trade is carried out by the producers of domestic and final goods. The producer of domestic goods exports good x to the rest of the world, while the final good producer imports good y from the rest of the world.

International trade is subject to three 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 the destination. We assume that goods cannot be re-exported.

Second, international purchases take time.⁹ We model this as a time-to-ship friction such that, if the final good producer purchases one unit of good y at date t , the foreign good arrives (and is only available as an input to produce final goods) at date $t + 1$.

Third, we assume that goods must be paid for in the same period in which they are ordered.¹⁰ Underlying this international trading structure is an assumed enforcement technology that allows countries to coordinate the dynamic exchange of goods across international borders.

Final Good Producer. The representative final good producer aggregates domestic good x and imported good y to produce final goods by operating a CES production technology:

$$G(x_t, y_t) = (x_t^\rho + y_t^\rho)^{\frac{1}{\rho}}, \quad (6)$$

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

The time-to-ship friction implies that the final good producer solves a dynamic problem, producing final goods in period t by combining domestic goods x_t ordered in t with imports y_t ordered in the previous period. The problem can then be represented as:

$$\max_{\{x_t, y_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \lambda_t [p_t G(x_t, y_t) - p_{x,t} x_t - \tau p_{y,t} y_{t+1}] \right\}, \quad (7)$$

where $p_{y,t}$ denotes the price of good y faced in period t (which arrives in period $t + 1$). Let Π_t^f denote the final good producer's profits or losses in period t , which are given by $p_t G(x_t, y_t) - p_{x,t} x_t - \tau p_{y,t} y_{t+1}$.

Given that the final good producer is owned by the household, it values profit streams across time and states of the world in the same way that the household does. Thus, profit streams corresponding to period t in a given state of the world are discounted according to discount factor

⁹Empirically, 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.

¹⁰In particular, we assume that import prices are known and determined before imports are shipped. See Capela (2011) for a description of how international transactions take place. In Section 5 we show that this assumption allows us to capture the payment arrangements observed in the data.

β^t and valued according to the Lagrange multiplier λ_t on the household's budget constraint of the corresponding period and state.

Domestic Good Producer. Domestic good x is produced by a representative domestic good producer with a Cobb-Douglas production technology:

$$x_t = z_t k_t^\theta n_t^{1-\theta}, \quad (8)$$

where θ is the capital share in production and z_t is an aggregate level of productivity that follows a stochastic process described below.

Then, the problem of the domestic good producer can be represented as:

$$\max_{k_{d,t}, n_{d,t}} p_{x,t} z_t k_{d,t}^\theta n_{d,t}^{1-\theta} - p_t w_t n_{d,t} - p_t r_{k,t} k_{d,t}, \quad (9)$$

where $k_{d,t}$ and $n_{d,t}$ denote the demand for capital and labor, respectively. Let Π_t^x denote the period- t profits or losses of the domestic good producer. Total profits Π_t are then given by $\Pi_t = \Pi_t^f + \Pi_t^x$.

Stochastic Process. We assume that aggregate productivity, the price of good x , and the price of good y follow a joint stochastic process. Defining the vector of exogenous state variables as $S_t = \{z_t, p_{y,t}, p_{x,t}\}$, we model the law of motion for S_t as a stationary VAR(1) process:

$$\log S_t = \Omega \log S_{t-1} + \nu_t, \quad (10)$$

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

This modeling choice is motivated by our desire to answer the following quantitative question in a simple and straightforward way: If an agent faces output and price dynamics as in the data, then what are the implications for imports? In the following sections, we estimate this flexible structure to capture salient features of the joint dynamics of output and prices observed in the data.

Finally, we assume that the interest rate r_t is constant over time and equal to $r = \frac{1}{\beta} - 1$.

Equilibrium. Given initial conditions $y_0 = \bar{y}$, $k_0 = \bar{k}$, and $S_{-1} = \bar{S}$, an equilibrium of this economy consists of state-contingent policy functions $\left\{ c_t, n_t, i_t, k_{t+1}, b_{t+1}, x_t, y_{t+1}, n_{d,t}, k_{d,t}, \Pi_t, \Pi_t^f, \Pi_t^x, \lambda_t \right\}_{t=0}^{\infty}$ and prices $\{p_t, w_t, r_{k,t}\}_{t=0}^{\infty}$ such that the following conditions hold:

1. Policy functions solve the household's problem;

2. Policy functions solve the final good producer's problem;
3. Policy functions solve the domestic good producer's problem;
4. Final good markets clear: $c_t + i_t = G(x_t, y_t)$ in every period and state;
5. Domestic good markets clear: $x_t + \tau x_{t+1}^* = z_t k_t^\theta n_t^{1-\theta}$ in every period and state;
6. Labor markets clear: $n_t = n_{d,t}$ in every period and state; and
7. Capital rental markets clear: $k_t = k_{d,t}$ in every period and state;

where x_{t+1}^* denotes exports of good x to the rest of the world, which are shipped in period t and arrive in $t + 1$. Given the price of good x follows an exogenous process, exports of good x to the rest of the world are such to ensure market clearing. In particular, for any given price of good x , final good producers pin down the domestic demand for this good, while domestic good producers pin down its supply; any excess supply is exported to the rest of the world.¹¹

3.1. Dynamic Import Demand

The key relationship in our model is the dynamic demand function for imports. After solving the representative final good producer's problem, we find that the domestic economy's demand for imports is expressed in equation (11) and summarized by Proposition 1.

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

$$y_{t+1} = A_t \left[\frac{\tau}{\mathbb{E}_t m_{t+1}^y} \frac{p_{y,t}}{P_t} \right]^{-\sigma}, \quad (11)$$

where:

$$A_t = \frac{p_{x,t} x_t + \tau p_{y,t} y_{t+1}}{P_t}, \quad P_t = [p_{x,t}^{1-\sigma} + (\mathbb{E}_t m_{t+1}^y)^\sigma (\tau p_{y,t})^{1-\sigma}]^{\frac{1}{1-\sigma}}$$

$$m_{t+1}^y = \beta \frac{\lambda_{t+1}}{\lambda_t} \frac{p_{t+1}}{p_t} \left[\frac{G(x_{t+1}, y_{t+1})}{G(x_t, y_t)} \right]^{\frac{1}{\sigma}},$$

where A_t is real aggregate absorption, P_t is the aggregate absorption price index,¹² and m_{t+1}^y is the stochastic discount factor used to price imports.

¹¹The underlying assumption here is that the rest of the world demands good x with a perfectly elastic demand function at price $p_{x,t}$. To the extent that there is excess demand for good x , the difference would be imported from the rest of the world; however, we restrict attention to parametrizations of the model in which this is never the case in equilibrium.

¹²Note that, given the above expression for P_t , measuring the ideal absorption price index implied by the model using data requires information on the price of domestic and foreign goods as well as on the expected value of the import-specific stochastic discount factor.

Equation (11) has several important features. First, the timing: Absorption and prices at date t affect imports delivered and used to produce final goods at date $t + 1$. This contrasts the contemporaneous effects of absorption 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, $\mathbb{E}_t m_{t+1}^y$ enters equation (11). The term m_{t+1}^y is a function of the intertemporal marginal rate of substitution or the stochastic discount factor. This term induces the demand for imports to depend on the agent's willingness to substitute final goods today (i.e. spending less on home goods today) for final goods tomorrow (i.e. imports arrive tomorrow).

This function of the stochastic discount factor is important for our purposes because it complicates the link between imports, absorption, and prices. That is, systematic variation in $\mathbb{E}_t(m_{t+1}^y)$ with absorption and prices implies that the model's income elasticity will differ from one and that the price elasticity will differ from σ . Moreover, because $\mathbb{E}_t m_{t+1}^y$ shows up along with the trade friction, variation in the SDF will appear as a time-varying trade wedge.

4. Qualitative Results

To develop intuition and motivate our calibration strategy in Section 5, we derive the income and price elasticities analytically for a special case of the model—an endowment economy with independent stochastic processes for endowments and world prices.

We do so by abstracting from (i) preferences for leisure and (ii) we restrict attention to an economy without capital ($\theta = 0$). Then, given the household supplies a unit of labor in every period and state of the world, output in this economy is solely determined by the productivity realization z_t . Finally, we restrict attention to an economy that operates under international financial autarky so that $b_{t+1} = 0$ in all periods and states of the world. Under these assumptions, we have that aggregate absorption A_t is equal to $\frac{p_{x,t} z_t}{P_t}$.

To derive analytical expressions of the income and price elasticities, we further assume that the stochastic process for absorption and world prices are independent of each other. Following standard OLS logic (see, e.g., Wooldridge 2015), this assumption implies that the income and price coefficients in a multiple regression framework are equivalent to the simple regression coefficient of say, imports on income. Thus, we can simplify matters and characterise the income and price elasticities by deriving formulas for the simple univariate regression coefficients.

4.1. 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 in the same period. For instance, in the static model we compute the income elasticity as the elasticity of y_t to

changes in A_t , assuming all other variables remain at their steady-state values. Computing this elasticity then boils down to estimating the regression coefficient of $\log A_t$ on $\log y_t$. Given this is a simple regression involving only two variables, it lends to a simple analytical solution. In the following sections we solve the model quantitatively to study the implied income and price elasticities when abstracting from the restrictive assumptions imposed in this section.

Proposition 2 summarizes the results for the static model.

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

$$\frac{\text{Cov}(\log A_t, \log y_t)}{\text{Var}(\log A_t)} = 1, \quad (12)$$

$$\frac{\text{Cov}(\log p_{y,t}, \log y_t)}{\text{Var}(\log p_{y,t})} = -\sigma. \quad (13)$$

This proposition shows that if the static trade model is the true 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.

4.2. Dynamic Model

Income Elasticity. For the dynamic model, we compute the same statistic in equation (12). However, in contrast to the static model, imports measured at date t (that is, measured upon arrival in period t , ordered in period $t - 1$) relate to objects at date $t - 1$:

$$\log A_{t-1} + \sigma \log E_{t-1}(m_t^y) + \alpha_1 = \log y_t, \quad (14)$$

where α_1 is a collection of terms that we can abstract from since we are only focusing on the covariance of absorption and imports.

Equation (14) illustrates that changes in endowments have two effects on imports. The first term reflects the idea that when endowments increase, the agent one-for-one increases imports—this is the standard force in static trade models. Second, an increase in endowments affects the marginal utility of consumption at date t relative to $t - 1$. This is the second term in (14). As an example, if a larger endowment today lowers the marginal utility of consumption today and increases the marginal utility of consumption tomorrow, then this induces the agent to substitute the endowment into the future by importing more.

This second term has several important implications. First, generically, our model will not

feature a unitary income elasticity because of changes in the stochastic discount factor. Second, the direction of the measured income elasticity (i.e. greater or less than one) depends on how the marginal utility changes with shocks. Third, this effect is amplified by the elasticity of substitution.

Proposition 3 summarizes these implications by mapping the income elasticity into moments of the stochastic process for endowments and the stochastic discount factor.

Proposition 3 *Dynamic Model: Income Elasticity $\neq 1$.* *The income elasticity is given by*

$$\frac{\text{Cov}(\log A_t, \log y_t)}{\text{Var}(\log A_t)} = \rho_z + \sigma \frac{\text{Std}(\log \mathbb{E}_{t-1} m_t^y)}{\text{Std}(\log A_t)} \text{Corr}(\log A_t, \log \mathbb{E}_{t-1} m_t^y), \quad (15)$$

where $\rho_z \in [0, 1)$ is the auto-regression coefficient on the endowment process. The key result is that our economy can generate an income elasticity of imports that is greater than one, depending on the dynamics of the stochastic discount factor relative to those of aggregate absorption. In particular, qualitatively, the key issue is that the stochastic discount factor is pro-cyclical. Quantitatively, the magnitude would depend upon the extent of the pro-cyclicality, the relative volatility of the stochastic discount factor, and the elasticity of substitution.

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

$$-\sigma \log p_{y,t-1} + \sigma \log \mathbb{E}_{t-1}(m_t^y) + \alpha_2 = \log y_t, \quad (16)$$

where α_2 is a collection of terms that we can abstract from since we are only focusing on the covariance of relative prices and 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 σ . The second term in is the dynamic effect—a change in prices today affects the marginal utility of consumption at date t relative to $t - 1$. The direction this force takes depends on details of the preferences. However, the intuition as to how this will work is that a reduction in the price of imports makes agents wealthier in the future, this lowers their marginal utility of consumption tomorrow and, thus, induces the agent to intertemporally substitute away from imports. Thus, the measured price elasticity will be lower (in absolute value) than σ .

Proposition 4 summarizes this logic by computing the regression coefficient of prices on imports.

Proposition 4 *Dynamic Model: Price Elasticity* $\neq -\sigma$. *The price elasticity is given by*

$$\frac{\text{Cov}(\log p_{y,t-1}, \log y_t)}{\text{Var}(\log p_{y,t})} = -\sigma \left[1 - \frac{\text{Std}(\log \mathbb{E}_{t-1} m_t^y)}{\text{Std}(\log p_{y,t-1})} \text{Corr}(\log p_{y,t-1}, \log \mathbb{E}_{t-1} m_t^y) \right]. \quad (17)$$

The key result here is that our economy can generate a price elasticity of imports that is artificially low (in absolute terms) relative to the elasticity of substitution σ . The important restriction in the expression above in getting a low price elasticity is that the stochastic discount factor is sufficiently volatile and co-moves positively with the price of imports.

Summary. These results show how variation in agents' stochastic discount factor may lead to imports that are more elastic to changes in income/absorption and less elastic to changes in relative prices. These implications stand in stark contrast to the predictions of standard, static trade models. Yet, these results are not guaranteed as (i) they depend on the correlation of the stochastic discount factor with other aggregate variables and (ii) these results abstract from other forces (labor supply, endogenous production, international borrowing or lending). Hence, the next section quantitatively evaluates the strength of this mechanism in a richer economy.

5. Calibration

In this section, we describe how we parameterize and calibrate our model. The quantitative question that guides our analysis is: How do the shipping and payment frictions interact with a finite intertemporal elasticity of substitution to shape the dynamics of aggregate imports in an economy consistent with the joint dynamics of output and prices observed in U.S. data?

Given the quantitative nature of our question, we begin by specifying additional features of the model required to match salient features of the data. In particular, we extend the model to feature a more realistic shipping technology, and we also introduce capital adjustment and bond-holding costs to better account for U.S. business-cycle dynamics.

We then calibrate the parameters of the model using aggregate U.S. data to discipline the stochastic process for productivity and prices. And we use information on asset returns to discipline preference parameters and, thus, fluctuations in the stochastic discount factor.

5.1. Quantitative Model

Preferences. We model the household's preferences between consumption and leisure to feature constant relative risk aversion over a Cobb-Douglas aggregate between them, following previous studies on international business cycles (Backus, Kehoe, and Kydland 1994a; Heathcote and Perri 2002; Corsetti, Dedola, and Leduc 2008).

Moreover, we parameterize the period utility function to have external habits in consumption as in Campbell and Cochrane (1999). While the preferences without habits can deliver results qualitatively consistent with the facts documented in Section 2, quantitatively they are unable to simultaneously feature plausible calibrations of the intertemporal elasticity of substitution/risk aversion and match moments on the fluctuations in asset returns. Thus, we use preferences with external habits in consumption to satisfy both of these desires.

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

$$u(c_t, h_t, n_t) = \frac{\left[(c_t - h_t)^\mu (1 - n_t)^{(1-\mu)} \right]^{1-\gamma}}{1 - \gamma},$$

where h_t is a consumption habit stock. The habit stock is parameterized to model the idea that the contribution of current consumption to utility depends on the previous period's consumption. Specifically, we define the surplus consumption ratio as

$$s_t = \frac{c_t - h_t}{c_t}, \quad (18)$$

and then define the law of motion for surplus consumption as

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

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

where \bar{s} is the steady-state level of surplus consumption and ρ_s denotes the persistence of surplus consumption. This 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.

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

Thus, we relax the one-period lag assumption by assuming that a fraction φ of the import orders made in period t arrive in t , while the remaining share $1 - \varphi$ of the import orders arrive in period

$t + 1$.¹³ The amount of imports available in period t to the final good producer is then given by:

$$\tilde{y}_t = (1 - \varphi)y_t + \varphi y_{t+1}, \quad \text{and} \quad G(x_t, \tilde{y}_t) = (x_t^\rho + \tilde{y}_t^\rho)^{\frac{1}{\rho}}, \quad (20)$$

where \tilde{y}_t is the consumption of imports in period t . The law of motion in (20) reflects the idea that some orders 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.

Capital Adjustment and Bond-Holding Costs. Finally, we introduce capital adjustment costs as well as bond-holding costs to better account for U.S. business-cycle dynamics.

Capital adjustment costs allow us to discipline the dynamics of investment over the business cycle. Previous studies have shown that such costs of adjustment are important to reconcile the dynamics of investment implied by international business-cycle models with the data (Mendoza 1991; Baxter and Crucini 1993; Backus, Kehoe, and Kydland 1994b), as well as to capture salient features of asset prices in production economies with habits in consumption (Beaudry and Guay 1996; Jermann 1998).

We model capital adjustment costs as quadratic in changes of the aggregate capital stock:

$$\Phi^k(k_{t+1}, k_t) = \frac{\eta_k}{2} (k_{t+1} - k_t)^2, \quad (21)$$

where $\eta_k \geq 0$ is a constant that controls the magnitude of this cost. We assume that the adjustment cost is denominated in final goods, so the household needs to purchase $\Phi^k(k_{t+1}, k_t)$ additional units of the final good to change the aggregate capital stock from k_t to k_{t+1} .

Bond-holding costs ensure the stationarity of the model (see, e.g., Schmitt-Grohé and Uribe 2003) and control the degree of international financial integration. On one extreme, if these costs are equal to zero, then the economy is perfectly integrated with the rest of the world conditional on the set of assets available. On the other end, if bond-holding costs are infinity, then the economy boils down to an economy under financial autarky.

We model bond-holding costs as in Schmitt-Grohé and Uribe (2003):

$$\Phi^b(b_{t+1}) = \frac{\eta_b}{2} b_{t+1}^2, \quad (22)$$

where $\eta_b \geq 0$ is a constant that controls the magnitude of this cost. We assume that the adjustment cost is denominated in units of the final good, so the household needs to purchase $\Phi^b(b_{t+1})$

¹³An alternative modeling strategy would be to allow for the “option” to pay for immediate delivery 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.

units of the final good to hold b_{t+1} units of the bond.¹⁴

5.2. 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 is set to standard values from the literature. The second set of parameters consists of the share of imports that arrive with a lag, which is estimated externally. The third set of parameters is estimated jointly to account for salient features of asset prices, business cycles, and the level of trade observed in U.S. data.

5.2.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 , the share of consumption μ in the utility function, the share of capital θ in the production function, and the depreciation rate δ . We set these parameters to standard values used in the literature to parameterize models of international business cycles (Backus, Kehoe, and Kydland 1994a; Heathcote and Perri 2002). The top panel of Table 3 summarizes our choice of predetermined parameters.

5.2.B. Externally Estimated Parameters

We estimate the share φ of imported goods that arrive in the same time period externally. To do so, we calculate the average number of days that it takes to ship goods into the U.S. We use data on shipping times by country of origin, mode of transportation (vessel or airplane), and coast of arrival (east or west coast of the U.S.); we describe this data in detail in Section 8. Average time-to-ship across these dimensions is computed weighting observations 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 (0.63) of U.S. imports arrive contemporaneously, while approximately one third (0.37) arrives with a delay of one quarter, for an average of 33 days of time-to-ship delay ($0.37 \times 90 = 33$).¹⁵ This leads us to set φ to 0.63.

Implied Payment Arrangement. International trade transactions are typically classified into one of two main groups depending on the timing of payment relative to the time that goods are

¹⁴The final-good market-clearing conditions are adjusted accordingly to allow for the use of final goods for capital adjustment and bond-holding costs.

¹⁵This 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.

Table 3: Calibration of Time-to-Ship Model

Predetermined Parameters		Externally Estimated Parameters		
IES $1/\gamma$	0.50	Intermediate Shipping Technology φ	0.63	
Consumption Share in Utility μ	0.34			
Capital Share θ	0.36			
Depreciation Rate δ	0.025			
Calibrated Parameters		Target Moments	Data	Model
Elasticity of Substitution σ	1.30	Empirical Price Elasticity	-0.301	-0.304
Capital Adjustment Cost η_k	2.11	Std. Dev. I / Std. Dev. GDP	3.117	3.061
Bond-Holding Cost η_b	0.86	Std. Dev. NX/GDP	1.111	1.099
Trade Cost τ	52.37	Imports/Absorption	0.231	0.231
Discount Factor β	0.97	Avg. Realized SDF	0.988	0.999
Surplus Consumption Persistence ρ_s	0.63	Std. Dev. Realized SDF	0.247	0.253
Steady-State Surplus Consumption \bar{s}	0.00035	corr(Realized SDF,Absorption)	0.088	0.090
Calibrated Stochastic Process $\{z, p_y, p_x\}$		Target Moments	Data	Model
Coefficient matrix $\Omega = \begin{pmatrix} 0.66 & -0.34 & -0.69 \\ 0.31 & 0.99 & 0.001 \\ -0.03 & 0.09 & 0.99 \end{pmatrix}$		corr(A_t, A_{t-1})	0.878	0.981
		corr($A_t, p_{y,t-1}$)	-0.270	-0.288
		corr($A_t, p_{A,t-1}$)	-0.597	-0.574
		corr($p_{y,t}, A_{t-1}$)	0.006	0.006
		corr($p_{y,t}, p_{y,t-1}$)	0.911	0.963
		corr($p_{y,t}, p_{A,t-1}$)	0.213	0.208
		corr($p_{A,t}, A_{t-1}$)	-0.506	-0.507
		corr($p_{A,t}, p_{y,t-1}$)	0.619	0.656
		corr($p_{A,t}, p_{A,t-1}$)	0.902	0.954
		Std. Dev. ν_z	0.0029	Std. Dev. A_t
Std. Dev. ν_{p_y}	0.0023	Std. Dev. $p_{y,t}$	0.031	0.026
Std. Dev. ν_{p_x}	0.0012	Std. Dev. $p_{A,t}$	0.009	0.008
corr(ν_z, ν_{p_y})	0.041	corr($A_t, p_{y,t}$)	-0.142	-0.148
corr(ν_z, ν_{p_x})	-0.995	corr($A_t, p_{A,t}$)	-0.598	-0.564
corr(ν_{p_y}, ν_{p_x})	0.055	corr($p_{y,t}, p_{A,t}$)	0.431	0.451

Note: Realized SDF is estimated using (23), (24) and inflation adjusted Fama and French portfolios and the 90-day Treasury bill for the time period Q2 1967 - Q1 2017.

delivered. Transactions in which importers pay for goods before delivery are typically referred to as “cash-in-advance”; transactions in which importers pay for goods after these are received are usually referred to as “open account.”

In our model, the share of cash-in-advance import transactions is given by the fraction of im-

ports that are paid one period before delivery, which is equal to $1 - \varphi$. Given a time-to-ship parameter equal to $\varphi = 0.63$, this implies that 37 percent of imports in our model are cash-in-advance transactions.

There is limited data available with information on the payment arrangements under which U.S. import transactions are carried out. With that said, the share of cash-in-advance import transactions in our model falls well within the range of existing estimates. Antras and Foley (2015) document that 42.2 percent of all exports of a big U.S. poultry exporter (41.3 percent of their total export value) were carried out under cash-in-advance between 1996 and 2009.

More generally, survey evidence compiled by the International Monetary Fund (IMF), the Bankers' Association for Finance and Trade (BAFT), and the International Financial Services Association (IFSA) among commercial banks in multiple countries (International Monetary Fund 2009; Asmundson, Dorsey, Khachatryan, Niculcea, and Saito 2011; International Monetary Fund 2011) estimate the share of cash-in-advance in international trade transactions around the 20 – 27 percent range over the period 2007-2011.

5.2.C. Calibrated Parameters

We estimate the rest of the parameters jointly following a simulated method of moments (SMM) approach that minimizes the distance between a set of moments of the data and their model counterparts.¹⁶

The set of parameters that we estimate includes the elasticity of substitution σ , the capital adjustment cost η_k , the bond-holding cost η_b , and the iceberg trade cost τ . Our estimation approach also features two other sets of parameters: the remaining preference parameters which largely determine the dynamics of the stochastic discount factor and the parameters governing the stochastic process for productivity and prices. The second and third panel of Table 3 report the estimated parameters as well as the target and model-implied moments.

We construct the model implied moments in the following way. First, we solve the model using a second-order approximation around its deterministic steady state. Then we simulate time paths of the relevant variables such that they conform with how we measured and treated them in the data. We then minimize the distance between model and data counterparts.

Below, we describe the target moments and how we measure them.

Standard, Aggregate Moments We target the following moments of U.S. data as targets: the empirical price elasticity estimated in the last row of Table 2, the standard deviation of real investment relative to the standard deviation of real GDP, the standard deviation of the net

¹⁶Specifically, 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.

exports to GDP ratio, and the average imports-to-absorption ratio.¹⁷

Stochastic Discount Factor. As discussed in Section 4, the stochastic properties of the stochastic discount factor is a key input into measuring the quantitative implications of the model. Our strategy is to use data on asset returns to infer the fluctuations of the stochastic discount factor that are attributable to the state variables of our model. We then use moments implied by the estimated stochastic discount factor in our calibration.

Our starting point to estimating the stochastic discount factor from asset returns is the assumption that there are no arbitrage opportunities across assets (as well as no transaction costs). Thus, we focus on the asset pricing condition which states 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 from which the stochastic discount factor can be estimated. Our implementation follows Cooper and Willis (2015) closely and the references therein.

We think this approach is valuable for two reasons. First, it puts strict discipline on the stochastic discount factor rather than simply letting it be an undiscipline outcome of the model. Second, we view it as a pragmatic approach in the presence of the well known challenges in the finance literature to finding *the* discount factor.

The asset pricing condition we exploit states that given the gross return R_{t+1}^j on portfolio j and the stochastic discount factor m_{t+1} the following condition must hold:

$$\mathbb{E}_t (m_{t+1} R_{t+1}^j) = 1. \quad (23)$$

Following Cooper and Willis (2015) and Zhang (2005), we assume that the stochastic discount factor relates to the value-of-import shipments y_t as well as to functions \tilde{S}_t of the other state variables in our model:

$$\begin{aligned} \log m \left(\tilde{S}_t, y_t, \tilde{S}_{t+1}, y_{t+1} \right) = & \alpha_0 + \alpha_1 \times \left[\tilde{S}_t - \tilde{S}_{t+1}, y_t - y_{t+1} \right]' \\ & + \alpha_2 \times \left[(\tilde{S}_t - \tilde{S}_{t+1}) \times (\tilde{S}_t - \bar{S}), (y_t - y_{t+1}) \times (y_t - \bar{y}) \right]', \end{aligned} \quad (24)$$

where α_0 is a constant term, α_1 is a $J \times 1$ vector of coefficients, α_2 is a $J \times 1$ 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 features time-varying risk aversion.

Given the specification in (24), we estimate the α s using the moment condition implied by (23) via the Generalized Method of Moments. The interpretation of this procedure is that we use

¹⁷In the data, we measure investment as total investment in equipment; all other variables are measured as described in Section 2. We measure the empirical price elasticity implied by the model as in the data; see Section 6 for details.

data on asset prices to infer the fluctuations of the stochastic discount factor that are attributable to fluctuations in the state variables of our model.

The asset return data that we use are the inflation adjusted returns on six different portfolios formed based on size and book-to-market value by Fama and French as well as the inflation adjusted rate on the 90-day Treasury bill over the period Q2 1967 to Q1 2017.¹⁸ We restrict attention to real absorption, the price of imports, and the price of absorption as proxies for the state variables (see Section 2 for a description of these variables).¹⁹ We implement the estimation by only estimating one coefficient for the terms of trade rather than a coefficient loading on each price separately; we thus have that $J = 3$, reducing the number of parameters estimated to seven. We also use lagged values of these variables as instruments to provide additional moment conditions.

The second panel of Table 3 presents 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. We also find that the realized stochastic discount factor is weakly correlated with absorption. Given these targets, we pick the preference parameters such that the moments of the realized stochastic discount factor in the model match those in the data.²⁰ The estimated parameters and the model’s fit are reported in the second panel of Table 3.²¹

Stochastic Process. We estimate the 15 parameters of the VAR process (equation 10) which determine the joint dynamics of productivity and prices in the model to match 15 moments on the joint dynamics of real absorption, the import price index, and the absorption price index from HP filtered quarterly U.S. data (see Section 2 for details of the data). In particular, we target the correlation of each of the three variables with their own lags as well as with the lags of the other variables (9 moments), the standard deviation of each of these variables (3 moments), and all contemporaneous correlations across them (3 moments).²² The estimated parameters as well as the target and model-implied moments are reported in the bottom panel of Table 3.

¹⁸The asset return data is publicly available from Kenneth French’s website: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. We adjust for inflation using the consumption price index for all urban consumers (all items less food and energy).

¹⁹An issue in the implementation of this approach is that we must modify the 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, some of these prices are observed by consumers a quarter before. Thus, we adjust the data variables accordingly to make them consistent with the timing of our model.

²⁰In our model, the stochastic discount factor m_{t+1} is given by $\frac{\beta \lambda_{t+1} p_{t+1}}{\lambda_t p_t}$.

²¹The 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).

²²We measure absorption and its associated price index as in the data; see Section 5 for details.

6. 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 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.

6.1. 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 NIPA. That is, in all our simulations, we collect data from our model and compute quarterly chain-type quantity and price indexes for absorption. Moreover, we measure imports upon arrival, as they are measured in the data.²³ 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 from Table 2 which control for composition and inventories. Given that we do not model these channels explicitly, we restrict attention to empirical income and price elasticities that control for them. The middle row reports the results from our time-to-ship model.

Two important results come from Table 4.

First, our model yields an income elasticity greater than unity. With that said, our mechanism only captures part of the high sensitivity of imports to income. With an income elasticity of 1.25, our model accounts for 40 percent of the difference between the unit income elasticity implied by standard models and the empirical income elasticity estimated from the data.

Second, the price elasticity lies well below the true elasticity of substitution σ . At -0.30 , this elasticity is meaningfully below the calibrated elasticity of substitution of -1.30 (second panel of Table 3). This result shows that variation in the stochastic discount factor is strong enough to rationalize the low price elasticity even though the true elasticity of substitution is estimated to be -1.30 .

We contrast these empirical elasticities with those implied by a static version of our model, a model without time-to-ship. To compute these elasticities, we re-estimate the model setting $\varphi = 1$ and keeping σ at its baseline value; the rest of the parameters are re-estimated as described in the previous section. Here, the income and price elasticities correspond with what the static model predicts—a price elasticity equal to the structural elasticity of substitution σ and a unit

²³In particular, we measure the quantity of imports in period t as \tilde{y}_t and its associated price as the (quantity-weighted) average price of imported goods $\frac{(1-\varphi)y_t}{\tilde{y}_t} p_{y,t-1} + \frac{\varphi y_{t+1}}{\tilde{y}_t} p_{y,t}$.

Table 4: Import Elasticities

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.30	1.63
Baseline Model	-0.30 [-0.36, -0.24]	1.25 [1.07, 1.49]
Static Model	-1.30 [-1.30, -1.30]	1.00 [1.00, 1.00]

Note: Data represents the inventory/durable adjusted elasticities. Results are averages from 250 simulations, with each simulation being 205 periods long; values in brackets report 95-percent confidence intervals. Section 2 describes the data.

income elasticity.


We want to emphasize that there are two aspects of these results that are surprising and not predetermined. First, the same parameterization that generates a price elasticity that is close to the data also delivers an income 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 7, there exist alternative parameterizations of the process governing the joint dynamics of productivity 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 parameterizing this process to capture the joint dynamics of absorption and prices observed in U.S. data allows our model to capture well the cyclical features of import data both qualitatively and quantitatively.


6.2. U.S. Import Fluctuations: 1967-2018

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-Q2 2018 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 the imports data and predictions from our model, for the entire time period Q3 1967-Q2 2018. Our model performs well, it tracks the data quite closely by capturing both the overall magnitude of fluctuations and the timing of peaks and



(a) Baseline Model vs. Data



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

Figure 3: Model vs. Data, 1967-2018

Table 5: Measures of Fit, Data and Model

Statistic	Data	Baseline Model	Static Model
Root Mean Square Error, Q3 1967 - Q2 2018	—	0.044	0.055
% Deviation From Trend Q1 2008 (Peak)	6.4	4.8	-0.7
% Deviation From Trend Q2 2009 (Trough)	-22.3	-15.5	3.0

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.

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 as well our model and it has severe problems regarding the timing and magnitudes in many instances.²⁴

Table 5 provides some metrics of fit. The first row reports the root mean square error between the data and predictions from the model. The baseline model's root mean squared error is 20% lower than the static model's (0.044 vs. 0.055).

Our model does particularly well in accounting for the dynamics of imports during the trade collapse of 2008-2009. The second and third rows of Table 5 report the percent deviation from trend during the peak in Q1 2008 and trough in Q2 2009. Our model correctly tracks the elevated levels of trade in Q1 2008. In contrast, the static model completely misses the peak. With regards to the trough, imports were 22.3 percent below trend in Q2 2009; our model is 15.5 percent below trend. As with the peak, the static model completely misses the trough.

Overall, from peak to trough, our model predicts a 20.3 percent change in imports relative to the 28.7 percent change seen in the data; in contrast, the static model implies a peak to trough change equal to 2.3 percent.

6.3. Discussion

The purpose of the quantitative exercise was to measure the importance of one particular mechanism: how variation in the IMRS affects imports and, in particular, can help rationalize puzzling features of import fluctuations. The results in Table 4 and Table 5 show that this mecha-

²⁴Our model can only partially account for the collapse of trade in the 1975 recession during the oil price crisis. Moreover, 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.

Table 6: Import Elasticities and Measures of the SDF

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Baseline Model	-0.30	1.25
Control for expected SDF	-1.10	0.99
Control for realized SDF	-0.32	1.19

Note: Results are averages from 250 simulations, with each simulation being 205 periods long; values in brackets report 95-percent confidence intervals.

nism is quantitatively important and that it is a powerful factor during crisis times. With that said, it is only a partial explanation—and this should not be a surprise. There are meaningful ways of thinking about the residual 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 in Amiti and Weinstein (2011) or Chor and Manova (2012); inventory considerations as discussed in Alessandria, Kaboski, and Midrigan (2010b); input-output linkages and vertical specialization as discussed in Bems, Johnson, and Yi (2010); or shocks to the future value of manufactures as in Eaton, Kortum, Neiman, and Romalis (2016).

These results also suggest approaches to better measure and forecast fluctuations in imports. Our model predicts that incorporating a measure of the **expected** SDF into the standard import demand regression of equation 2 would reduce biases in empirical income and price elasticities. Table 6 illustrates this point by estimating the income and price elasticities implied by the model when controlling for the model-implied SDF. As the second row shows, incorporating the expected SDF into the import demand specification leads to a near unitary income elasticity and much higher price elasticity. Furthermore, what is important for this specification is that it is the expected, forward looking estimate of marginal utility. When the realized value is used, it has essentially no impact.

How to measure the expected SDF? It's well known that this is a challenging task. However, proxies from financial markets such as stock prices or credit-spreads may yield valuable information. Empirically, we found success by incorporating the value of the real S&P 500 price index from Shiller (2018) into the import regressions discussed in Section 2. Following the lag structure implied by (11), we found that by using the stock index as a proxy for the expected SDF reduced the empirical income elasticity from two to 1.66 and increased the price elasticity from 0.3 to 0.5; adjusting for durables and inventories lowers the income elasticity to 1.29.

While not perfect, incorporating this simple proxy reduced the income elasticity and increased the price elasticity in the exact direction predicted by our theory.

7. Important Ingredients

There are essentially three key ingredients to the quantitative results. First, we chose functional forms on preferences and capital adjustment and bond-holding costs that allow the model to match the dynamics of the stochastic discount factor. Second, we consider an economy in which a fraction of import purchases are paid before goods are delivered (cash-in-advance). Third, we estimate a process for productivity and prices such that the model matches the joint dynamics of absorption and prices observed in U.S. data. This step contrasts with alternative specifications for the process behind productivity and prices or closing the model and endogenizing world prices. Below we show how deviations from these choices lead to counterfactual predictions and mitigate our quantitative results.

7.1. Habits

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. This was motivated by the observations around Propositions 3 and 4 that the dynamics of the stochastic discount factor play a critical role in shaping the response of imports to fluctuations in income and prices.

To illustrate the role played by habits, we compute the income and price elasticities implied by an economy without habits ($h_t = 0$ in all periods and states of the world), while keeping all other parameter values as in our baseline parametrization. The results for the income and price elasticities are presented in the third row of Table 7 (the first two rows report the elasticities in the data and the baseline model, respectively).

We find that the income and price elasticities are 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 97 percent *less* volatile as our calibration target (0.007 vs. 0.247). This shows that external habits allow us to match fluctuations in the stochastic discount factor and, thus, play a key role amplifying the income elasticities and dampening the price elasticities.

7.2. Capital Adjustment and Bond-Holding Costs

We also introduced capital and bond-holding costs. These costs are important to capture features of asset prices in production economies with habits in consumption (see, e.g., Beaudry

Table 7: Sensitivity Analysis, Import Elasticities

	Price Elasticity, $\hat{\alpha}$	Income Elasticity, $\hat{\beta}$
Data	-0.32	1.60
Baseline Model	-0.33	1.24
No Habits	-1.24 [-1.24, -1.24]	1.03 [1.02, 1.03]
No Capital Adjustment Costs No Bond-Holding Costs	-1.32 [-1.33, -1.30]	1.01 [1.00, 1.01]
No Cash-in-Advance	-1.32 [-1.32, -1.31]	1.01 [1.00, 1.01]
Alternative Stochastic Process	-1.30 [-1.30, -1.29]	1.00 [0.94, 1.06]
Two-Country Model	-1.26 [-1.30, -1.22]	0.94 [0.93, 0.95]

Note: Results are averages from 250 simulations, with each simulation being 205 periods long; values in brackets report 95-percent confidence intervals. Section 2 describes the data.

and Guay 1996; Jermann 1998). We now examine the importance of capital adjustment and bond-holding costs for our findings.

To do so, we compute the income and price elasticities implied by an economy without either of these costs ($\eta_k = \eta_b = 0$), keeping all other parameter values as in our baseline parametrization. The implied income and price elasticities are presented in the fourth row of Table 7.

As in the economy without habits, we find that the income and price elasticities are essentially the same as those of a static trade model. The issue is that without these costs the stochastic discount factor is substantially less volatile and pro-cyclical than the baseline model (its standard deviation is 0.141 vs. 0.247, and its correlation with absorption is 0.010 vs. 0.088). Consistent with the findings of Beaudry and Guay (1996) and Jermann (1998), these results show that capital adjustment and bond-holding costs play a fundamental role to generate plausible fluctuations of the stochastic discount factor.²⁵ And, in turn, activate the role that the stochastic

²⁵Both costs are important; the implied income and price elasticities are close to those implied by the model without time-to-ship if either $\eta_k = 0$ or $\eta_b = 0$.

discount factor plays in shaping import fluctuations.

7.3. Payment Technology

Another ingredient that is important is the payment technology. In our model, imports are paid for in the same period in which they are ordered, which leads to a mismatch between the timing of import payments and import deliveries. As described in Section 5.2.B, 37 percent of imports in our model are cash-in-advance transactions; a value which is well within the range of previous estimates in the literature.

To illustrate the role played by the payment technology, we extend the model to feature a richer set of payment technologies. In particular, we consider an extension of the model in which a fraction ψ of imports ordered in a given period is paid for in that same period; in the baseline model, $\psi = 1$.²⁶ Here, we set $\psi = \varphi = 0.63$ such that all import transactions are paid for in the period that goods are delivered. So the imports delivered today are paid for today, imports delivered tomorrow are paid for tomorrow.

We compute the income and price elasticities implied by this extension of the model, while keeping all other parameter values as in our baseline parametrization. The results are presented in the fifth row of Table 7. As in the economies without habits or without capital and bond-holding costs, we find that the income and price elasticities are essentially the same as those of a static trade model. This illustrates that the cash-in-advance payment technology is an important ingredient to our model. For our mechanism to be operative, resources need to be sacrificed today in exchange for the delivery of goods tomorrow.

7.4. Stochastic Process

Another fundamental feature of our modeling strategy is the assumption that productivity and prices follow an exogenous stochastic process estimated to reproduce the joint dynamics of absorption and prices observed in U.S. data. This approach allows us to sidestep well-known price anomalies featured by multi-country 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 our baseline results with two alternatives.

Alternative Stochastic Process. We first show the importance of estimating the stochastic process for productivity and prices to match the joint dynamics of absorption and prices observed in U.S. data by considering an economy with an alternative stochastic process.

The question we ask is: To what extent is it possible to parameterize the stochastic process to

²⁶See Section 4 of the Online Appendix for further details on this extension of the model.

imply price and income elasticities close to the static model in an economy that accounts for all target moments except the joint dynamics of absorption and prices observed in the U.S.? In other words, does there exist a stochastic process under which our model fails despite matching all other relevant moments?

The sixth row of Table 7 reports the price and income elasticities implied by an alternative parameterization of the stochastic process, re-estimating the rest of the parameters as described above, except for two differences. First, we do not target the joint dynamics of absorption and prices. Second, we target a unit income elasticity and a price elasticity equal to $-\sigma$.²⁷

We find that estimating the stochastic process to match the joint dynamics of absorption and prices in U.S. data plays a fundamental role in accounting for the empirical price and income elasticities. In particular, we find that there exist parameterizations of the stochastic process such that the implied price and income elasticities are arbitrarily close to their static-model counterpart, despite accounting for all target moments other than the joint dynamics of absorption and prices (see Online Appendix for details).

The implication of these exercises reemphasize the points made in Section 6—that our model need not deliver high income elasticities and low price elasticities. Features of U.S. data are key drivers of this conclusion.

Two-Country Model. We also examine the role of assuming an exogenous stochastic process for productivity and prices by investigating the implications of a two-country version of our model in which the prices of domestic and foreign goods are endogenous.

We parameterize the stochastic productivity process following Backus, Kehoe, and Kydland (1995), and estimate the rest of the parameters following the approach described in Section 5. The bottom row of Table 7 presents the results.²⁸

The main finding is that the income and price elasticities change substantially and are now close to those implied by the static model without time-to-ship—and far from the data. The key 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 (see Online Appendix for details). 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

²⁷See the Online Appendix for more details of this alternative parameterization of the model.

²⁸See the Online Appendix for a detailed description of the two-country model, as well as its parameterization.

reduces the attractiveness of imports as a way to smooth consumption across time, both directly and through its impact on the SDF. Altogether, the counter-factual 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.

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

This section examines some cross-sectional implications of our model. Our model predicts that, *ceteris paribus*, 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 the static model predicts that the volatility of imports is independent of time-to-ship/distance.²⁹

To explore this implication, we construct a measure of the average time it takes to ship goods into the U.S. by combining Hummels and Schaur's (2013) dataset on shipping times and the World Bank's Doing Business survey. Hummels and Schaur (2013) measure the average time it takes to ship goods into the U.S. from each country, by mode of transportation, coast of arrival (east or west coast), and HS6 good categories. For each country of origin, we use this data to compute the average across these dimensions weighting by imports volume. The World Bank's Doing Business survey measures the "time necessary to comply with all procedures required to export goods" in each country, as well as the "time necessary to comply with all procedures required to import goods" in the U.S. We construct our measure of average total time-to-ship into the U.S. by adding up these three variables.³⁰

We compare these time-to-ship measures with quarterly data on U.S. nominal imports disaggregated by country of origin. The data are obtained from the U.S. Census, are 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 eco-

²⁹With ideal data (price indexes of U.S. imports disaggregated by country of origin spanning a significant number of periods and countries), one could perform the exercises from Section 6 by country of origin and contrast the cross-sectional implications and/or try and "difference" out $\mathbb{E}_t(m_{t+1}^y)$ across destinations in equation 11. Unfortunately, these data are not available.

³⁰Note 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.

nomically 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.³¹

Table 8: Volatility of Imports by Time-to-Ship

Time-to-Ship	Imports Volatility (%)	# of Countries	Avg. Shipping Time
≤ 25 days	8.29	16	22.15 days
25 – 50 days	11.44	48	33.78 days
> 50 days	20.81	18	59.05 days

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.

Related findings have been previously documented in the literature. Levchenko, Lewis, and Tesar (2011) 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 sea. Both of these results are consistent with our findings.

Table 9 contrasts this evidence with the relationship between time-to-ship and imports volatility implied by the model. To do so, we consider three economies that differ in their value of φ and are otherwise identical to our baseline model; we think of these economies as the model-counterparts to the three country groups examined in Table 8. We choose the value of φ in each of these economies to match the average shipping time reported in Table 8 for each of the country groups. In particular, given an average shipping time equal to x days, we choose φ to equal $1 - \frac{x}{90}$.

The first two columns report the volatility of imports in the data and the model. While imports volatility is increasing in time-to-ship in the data and the model, imports are considerably more volatile in the data. This need not necessarily indicate an issue with our model; bilateral imports can be substantially more volatile than aggregate imports depending on the covariance structure across the different countries of origin.³²

Thus, in the last two columns we abstract from differences in the average level of imports

³¹In the Online Appendix we investigate the statistical significance of this relationship controlling for variables that are commonly used to account for bilateral trade flows.

³²Another potential source of higher import volatility in the data might be the lack of seasonal adjustment of the data series that we study in this section.

Table 9: Volatility of Imports by Time-to-Ship, Model vs. Data

Time-to-Ship	Imports Volatility (%)		Relative Volatility	
	Data	Model	Data	Model
22.15 days, $\varphi = 0.75$	8.29	2.17	1.00	1.00
33.78 days, $\varphi = 0.62$	11.44	2.81	1.38	1.29
59.05 days, $\varphi = 0.34$	20.81	4.35	2.51	2.00

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.

volatility between the data and the model and, instead, focus on the relationship between time-to-ship and imports volatility. To do so, we normalize the volatility of the group with lowest time-to-ship to equal unity. We find that the model and the data feature a surprisingly similar relationship between imports volatility and time-to-ship. In the model, the volatility of imports of the second country group is 1.29 times larger than the first group vs. 1.38 times larger in the data. Similarly, the volatility of imports of the third country group is 2 times larger than the first group in the model vs. 2.51 times larger in the data.

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; the cross-sectional predictions of our model are consistent with this evidence.

9. 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 account for a large fraction of 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. First, 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. As our preliminary work in Section 6.3 suggests, incorporating financial measures into standard import demand regression may improve their interpretation and forecasting performance. A second open issue is the interaction between our proposed mechanism and alternative mechanisms such as compositional changes and inventories. These alternative mechanisms involve forward looking behavior, thus we speculate that the explicit modeling of these forces together would lead to an improved understanding of fluctuations in trade volumes.

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.
- ALESSANDRIA, G., S. PRATAP, AND V. Z. YUE (2015): "Export dynamics in large devaluations," .
- 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 (2015): "Poultry in motion: a study of international trade finance practices," *Journal of Political Economy*, 123(4), 853–901.
- ARMINGTON, P. S. (1969): "A Theory of Demand for Products Distinguished by Place of Production (Une théorie de la demande de produits différenciés d'après leur origine)(Una teoría de la demanda de productos distinguiéndolos según el lugar de producción)," *Staff Papers-International Monetary Fund*, pp. 159–178.
- ASMUNDSON, I., T. W. DORSEY, A. KHACHATRYAN, I. NICULCEA, AND M. SAITO (2011): "Trade and trade finance in the 2008–09 financial crisis," .
- 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 (1994a): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?," *American Economic Review*, 84(1), 84–103.
- (1994b): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?," *The American Economic Review*, pp. 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.

- BAXTER, M., AND M. J. CRUCINI (1993): "Explaining saving–investment correlations," *The American Economic Review*, pp. 416–436.
- 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.
- (2013): "The great trade collapse," *Annu. Rev. Econ.*, 5(1), 375–400.
- 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.
- BUSSIÈRE, 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.
- CORSETTI, G., L. DEDOLA, AND S. LEDUC (2008): "International risk sharing and the transmission of productivity shocks," *Review of Economic Studies*, 75(2), 443–473.
- DI EWERT, 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 (2016): "Trade and the global recession," *The American Economic Review*, 106(11), 3401–3438.
- 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, cyclical, 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. (2017): "High discounts and high unemployment," *The American Economic Review*, 107(2), 305–330.
- 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.
- INTERNATIONAL MONETARY FUND (2009): "World Economic Outlook," Discussion paper.
- (2011): "The 6th IMF/BAFT-IFSA Survey, Key Findings and Observations," Discussion paper.
- ISARD, P. (1975): *Dock strike adjustment factors for major categories of US imports and exports, 1958-1974*. Division of International Finance, Board of Governors of the Federal Reserve System.
- JACKS, D., C. MEISSNER, AND D. NOVY (2010): "Trade costs in the first wave of globalization," *Explorations in Economic History*, 47(2), 127–141.
- JACKS, D. S., C. M. MEISSNER, AND D. NOVY (2009): "The role of trade costs in the great trade collapse," in *The great trade collapse: causes, consequences, and prospects*, ed. by R. Baldwin. VoxEU.org.
- JERMANN, U. (1998): "Asset pricing in production economies," *Journal of Monetary Economics*, 41(2), 257–275.
- KEHOE, P., E. PASTORINO, AND V. MIDRIGAN (2016): "Debt constraints and employment," Discussion paper, National Bureau of Economic Research.
- 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 (2010): "The collapse of international trade during the 2008–09 crisis: in search of the smoking gun," *IMF Economic Review*, 58(2), 214–253.
- (2011): "The Role of Trade Finance in the US Trade Collapse: A Skeptics View," *Trade Finance during the Great Trade Collapse*, p. 133.
- 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.
- MENDOZA, E. G. (1991): "Real business cycles in a small open economy," *The American Economic Review*, pp. 797–818.
- 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*.
- SCHMITT-GROHÉ, S., AND M. URIBE (2003): “Closing small open economy models,” *Journal of international Economics*, 61(1), 163–185.
- SHILLER, R. (2018): “U.S. Stock Markets 1871-Present and CAPE Ratio,” data retrieved from http://www.econ.yale.edu/~shiller/data/ie_data.xls.
- 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.
- WOOLDRIDGE, J. M. (2015): *Introductory econometrics: A modern approach*. Nelson Education.
- 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 throughout the paper come from the Bureau of Economic Analysis’ (BEA) National Income and Product Accounts (NIPA). 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 difference are the line numbers—specifically lines 7 and 11 for nominal values, quantity and price indexes of GDP. Lines 48 and 49 for exports of durable and non-durable goods, lines 52 and 53 for imports.
- **Inventories:** Table 5.7.6 line 1, change in real private inventories (chained 2009 dollars).

- **Investment:** Tables 5.3.4 and 5.3.5 lines 9 and 26 (nonresidential and residential equipment, respectively). We construct a real measure of total investment in equipment following the “Fisher of Fishers” approach described in the paper.