International Menu Costs and Price Dynamics*

Raphael Schoenle†
Brandeis University

September 2010

Abstract

In this paper, I analyze how the pricing behavior of firms systematically differs across domestic and export markets in terms of frequency, timing and size of price changes. First, I contrast domestic and export pricing decisions for the same products showing that (i) domestic producer prices change approximately twice as often as export producer prices, (ii) the probability of synchronized price adjustment across markets is 21% for upwards adjustments and 14% for downwards adjustments, (iii) the size of export price changes is substantially larger than the size of domestic price changes and (iv) there are strong seasonality effects in the data including a year-end synchronization effect. Second, I show that economic fundamentals such as inflation, productivity, demand, exchange rates and market structure can only partially explain adjustment decisions and cross-market synchronization. Third, I present a dynamic menu cost model of price-setting, and attribute the remaining unexplained part in adjustment decisions to differences in menu costs across countries. I calculate the implied export and domestic market menu costs from the data and estimate that export menu costs are 1.5% of period steady state revenues and three times as large as domestic market menu costs.

JEL Classification: E3; F41, F42; L16.

Keywords: International Price Dynamics, U.S. Producer Prices, Export Prices, Menu Costs, Synchronization.

---

*I would like to thank Alan Blinder and Marc Melitz for invaluable advice and encouragement. I would like to thank Pinelopi Koujianou Goldberg, Saroj Bhattarai, Thomas Chaney, Gauti Eggertsson, Oleg Itskhoki, Nobu Kiyotaki, Kalina Manova, Emi Nakamura, Chris Sims, Mu-Jeung Yang and and seminar and workshop participants at Brandeis University, Erasmus University Rotterdam, ESSEC, European School of Management and Technology, Federal Reserve Bank of Dallas, IMF Research Department, Royal Holloway University of London, and Princeton University for helpful suggestions and comments. This research was conducted with restricted access to the Bureau of Labor Statistics (BLS) data. The views expressed here are those of the author and do not necessarily reflect the views of the BLS. I thank my project coordinators, Ryan Ogden, and especially Kristen Reed, for substantial help and effort, as well as Greg Kelly and Rosi Ulicz.

†Email: schoenle@brandeis.edu; Tel: 617-680-0114; Fax: 781-736-2269; Mailing address: Department of Economics, MS 021, Brandeis University, 415 South Street, Waltham, MA 02453-2728.

1 Introduction

While it is well known that firms price to market, we know very little about the dynamics of price changes of the same product across countries. However, understanding the timing of export price-setting is key for understanding a number of issues in international macroeconomics such as the global spread of inflation, fluctuations in real exchange rates and optimal monetary policy in the open economy.

In this paper, I document several new facts about price-setting behavior in international markets, using producer price micro-data from the Bureau of Labor Statistics (BLS). The novelty of my analysis is to compare price dynamics directly from the perspective of the producer who sets prices both in the domestic and export markets. In a first step of analysis, I contrast price dynamics for the same products, defined at the six-digit level, and for a set of matched firms selling in the domestic and export market.[1] I show that the pricing behavior of firms systematically differs across markets in terms of frequency, timing and size of price changes. In a second step, I show that economic fundamentals such as inflation, productivity, demand, exchange rates and market structure can only partially explain adjustment decisions and cross-market synchronization. Finally, I present a dynamic menu cost model of price-setting, and attribute the unexplained part in adjustment decisions to differences in menu costs across countries. I estimate that export menu costs have to be approximately three times larger than domestic market menu costs to account for the empirical findings.

I start my empirical analysis by contrasting the frequency of price changes for the same products in the domestic and export markets. I find that domestic prices change 2.1 to 2.7 times as often as export prices of the same product. This difference in adjustment frequencies at the product-level corresponds to a median difference in the duration of price spells of 5 to 10 months. My estimates based on the data for matched firms imply a duration difference of 7.4 months. Using different methodologies to calculate the frequency of export price changes, I find that differences between domestic and export price-setting behavior are robust to the presence of missing values in the export data.

Second, while the fact that export prices change less frequently than domestic prices does not rule out that export prices adjust whenever domestic prices adjust, the data show a striking absence of synchronization in export and domestic price adjustment. The average monthly probability of upwards export price adjustment is 21% conditional on complete domestic upwards price adjustment of the same product. At the same time, the monthly probability of synchronized downwards price adjustment in the export market is 14%. I find almost identical results when I additionally break down the calculations by export destinations. There are two natural benchmarks to which to compare these synchronization

[1]Since the BLS domestic producer price data and the export price data have different sampling frames, I use a fuzzy matching algorithm to match firms in the domestic and export data. I am able to match 10% of the exporting firms.

results: First, since U.S. exports are denominated in USD in the data, the Law of One Price would imply perfect synchronization for the timing of price adjustment. A second benchmark is the fraction of other products, defined at the six-digit level, which adjust in the domestic market within the same four-digit category whenever a particular product fully adjusts in the domestic market: the probability of synchronized domestic upwards (downwards) adjustment of that particular product with other products in the same four-digit category is 48% (50%).

In addition to computing probabilities of export price adjustment conditional on complete adjustment of all goods in the same domestic product category, I regress domestic and export price adjustment decisions of each good in a given product category on the fraction of domestic price changes which have the same sign in that product category. I find that the estimated coefficient on the fraction of same-signed domestic adjustment is substantially larger in regressions when the dependent variable is the individual domestic adjustment decision than when it is the export adjustment decision. Therefore, regression results also show that export price adjustment decisions are not very synchronized with domestic price adjustment decisions. Because domestic and export prices change at different frequencies and not in a synchronized way, my findings suggest that we may have to carefully modify the standard modeling assumption of identical Calvo parameters or even simultaneous adjustment of prices across markets.

Third, I find that differences in price-setting behavior across markets also extend to the differences in the size of price changes: conditional on adjustment, the median absolute size of export price changes is 39% larger than the size of domestic price changes of the same product. The median percentage difference between export and domestic absolute sizes of price changes is 1.75%. When I break down price changes by destination, the median absolute size of export price changes is 30% larger than the size of domestic price changes. The median percentage difference between export and domestic absolute sizes of price changes is now 1.43%. Comparing the size of price changes for matched firms, I find even more pronounced differences: the median absolute size of export price changes, conditional on adjustment, is 2.5 times as large as the size of domestic price changes, conditional on adjustment. The median percentage difference between the size of export and domestic price changes conditional on adjustment is 3.10%.

Fourth, I document a new seasonality effect in the data in the form of year-end synchronization of export and domestic price changes. Conditional on domestic adjustment, the probability of export price adjustment, computed as the mean over all products in a given month and irrespective of the sign of adjustment, rises from a trough of 0.4 in March to values around 0.6 to 0.8 in October, November, December and January, with a slight intermittent peak in June. This immediately suggests that there may be seasonality in the international transmission of shocks or that for example measures of pass-through may have a highly seasonal component.

I provide several robustness checks that these findings are not driven by selection of firms or products into exporting: one concern about my findings could be that only a subset of domestic firms export as we know from theoretical and empirical work by Melitz (2003), Mayer et al. (2009), Goldberg et al. (2008) or Bernard et al. (2006). Moreover, exporting firms are more productive and larger, and could similarly differ with respect to their price adjustment behavior. This implies that a composition effect due to export selection could be driving results. I show robustness to this concern in several ways. First, I construct a synthetic domestic dataset which contains approximately the same goods in each product category as the export data. When I use this synthetic data in my comparison, I find even larger differences of domestic and export pricing behavior[2]. Second, while I have directly addressed selection of firms into exporting through my comparison of matched firms, I additionally run a robustness check on export selection of items within firms: firms tend to export their largest domestic sales items, their "hits." If these items were sampled in all export destinations, this could imply a composition effect in my analysis. Therefore, I compare only the largest domestic sales item to the items in each export destination of firms. I find that my results remain robust. Third, I check for composition effects using a simple regression of differences in domestic and export frequencies on product log export shares. The idea is that a small log export share is a proxy for a high chance that sampling among the universe of domestic goods can select a good different from the relatively smaller universe of exported goods.[3] I find no statistically significant relationship in the regression suggesting that composition effects are not driving my results. Finally, I present evidence from cover prices of The Economist that show how even prices of exactly identical items are subject to substantial differences in domestic and export price dynamics.[4]

Next, to understand what drives adjustment decisions and to explore the mechanisms behind synchronization across countries, I analyze a simple two-country dynamic menu cost model that relates adjustment decisions to economic fundamentals. Firms sell in both the domestic and export market, and set prices by trading off the cost of adjustment with profits foregone due to changes in demand, inflation, productivity and the real exchange rate. First, motivated by the state-dependent aspect of the model, I estimate a multinomial logit model of individual monthly adjustment decisions in the data. I find that economic fundamentals are significantly related to adjustment decisions, both in the domestic

[2]My results on synchronization previously discussed are based on an analysis of the synthetic data in the first place.

[3]The extreme case of a zero export share means that with certainty a different domestic good will be selected from the non-existing exported good. At the same time, a 100% export share means that the chance of sampling a different, non-exported domestic good has become smaller.

[4]In Schoenle (2009), I show that the substantial differences in price dynamics persist for 10 newspapers from different countries sold internationally, and 100 identical IKEA products sold in 20 countries. In all of these cases, firm and item export selection should not be responsible for differences in price dynamics.

and export data, and in some cases affect the decisions to adjust upwards and downwards with different signs, as predicted. Moreover, I find that strong effects of seasonality persist even after controlling for economic fundamentals.

Second, I use the model to estimate a structural equation that relates the decision to synchronize domestic and export adjustment to the size of productivity shocks and world inflation and demand shocks. I find strong evidence that these common shocks are driving synchronization of domestic and export price changes, as the model predicts. First, results show that the larger productivity shocks become, the larger is the strength of the estimated synchronization parameter. This holds in both parametric and semi-parametric estimations of the synchronization equation. Second, estimation shows that larger bilateral correlation coefficients of domestic and export market inflation, and country aggregate output growth rates are similarly associated with higher synchronization of price changes: when domestic and export market inflation (output growth) is more correlated, so are adjustment decisions. However, both the multinomial regressions that analyze pricing decisions in the first place and the structural test leave most of the variation in price adjustment and synchronization decisions unexplained.

Therefore, I attribute this unexplained residual to differences in domestic and export menu costs which I interpret more broadly as the cost of price adjustment, as work by Blinder et al. (1998) suggests. Domestic and export menu costs are free parameters in my model and I estimate them without imposing any restrictions a priori, using simulation-based maximum likelihood techniques while controlling for productivity, demand, nominal exchange rate, US and export-market inflation rates. I estimate menu costs based on data for the frequency of price changes, the size of price changes and on their joint distribution. I find that the likelihood functions are well-behaved, single-peaked and yield similar estimates in all three cases. The maximum likelihood estimates are approximately 0.4% of period steady state revenues for domestic menu costs and 1.5% for export menu costs.

My estimates are in the range of adjustment costs found by Zbaracki et al. (2004) who also provide anecdotal evidence for larger export menu costs involving higher managerial and customer relation costs. Relating my menu-cost estimates to country-characteristics, I find that adjustment costs are higher for more distant export markets and lower for markets which share English as the common language with the U.S.. These findings support a broader interpretation of larger export market menu costs as costs of price adjustment associated with higher managerial and customer relation costs.

My analysis complements several strands in the existing literature. On the empirical side, it is related to recent work by Gopinath and Rigobon (2008), Gopinath and Itskhoki (2010), Goldberg and Hellerstein (2009). Nakamura and Steinsson (2008a) and Neiman (2009) who also use BLS micro data on consumer, producer and export and import prices to study the behavior of prices, exchange-rate pass-through and the role of firm boundaries. I add to this literature a novel, disaggregated comparison of price dynamics from the perspective of the U.S. producer selling domestically and abroad. In addition, I explore how fundamentals determine pricing and importantly, synchronization decisions. Finally, I use the difference in export and domestic price-setting behavior to identify and estimate costs of international price adjustment based on a simple model.

My analysis also relates to work by Fitzgerald and Haller (2009) who study pass-through of export local-currency pricing by Irish producers, work by Gagnon (2009) who provides evidence for state-dependent pricing from Mexican consumer prices and studies of producer price-setting by the European Inflation Persistence Network (IPN), such as Vermeulen et al. (2007). In particular, I connect to the analysis of seasonality in Vermeulen et al. (2007) by documenting new seasonality effects, while complementing the early work on seasonality by Barsky and Miron (1989).

On the theoretical side, my findings have direct implications for modeling real exchange rate dynamics. For example, Benigno (2004a) shows how differences in adjustment frequencies across countries can generate realistic real exchange rate dynamics which most models fail to match. My findings directly support his modeling approach. More generally, the existence of differences between domestic and export price-setting behavior has important implications for key modeling assumptions in international macroeconomics. The standard assumption of identical Calvo parameters or even simultaneous adjustment of prices across markets may have to be carefully modified. In addition, if exports induce heterogeneity in the frequency of price changes, this new source of heterogeneity is likely to generate large real output effects similar to intra-country sectoral heterogeneity as in Carvalho (2006) or Nakamura and Steinsson (2008b). My findings also relate to the work of Benigno (2004b) and Aoki (2001) who show how optimal monetary policy should target inflation with weights according to the degrees of price stickiness. Finally, my results on seasonality connect to Olivei and Tenreyro (2007) who show how effectiveness of monetary policy is highly seasonal. My findings suggest that future research on pass-through should possibly focus on similar seasonality effects.

The paper proceeds as follows. Section 2 describes the BLS micro data underlying the U.S. producer and export price indices as well as other data sources used in my empirical analysis. Section 3 documents differences in price-setting behavior by comparing domestic and export pricing behavior. In Section 4, I present the dynamic menu cost model, implications based on numerical solution techniques and my empirical test strategy. Section 5 contains the results of the empirical test while Section 6 presents the results of the menu cost maximum likelihood exercise. Section 7 concludes.

2 Data

In this section, I briefly describe the data used in my analysis. In particular, I discuss aspects of the data related to my comparison of domestic and export prices. Appendix 9.4 contains a more detailed discussion and description of the data.

2.1 Domestic Price Data

As my data source of domestic producer prices, I use the micro data collected by the BLS to compute the U.S. Producer Price Index (PPI). These data are monthly transaction prices for individual "items" which are particular goods consistently defined over time, for example in terms of buyer or day of the month of the transaction. Prices are collected through a survey where items are selected according to a multi-stage design[5]. Importantly, weights of selection in this design are values of shipment for industries, firms and items within firms. From these data, I use data for the years from 1998 through 2005, excluding prices associated with intra-firm trade to avoid the problem of transfer prices.

2.2 Export Price Data

For the same time period, I obtain monthly, survey-based price data on individual exported items from the International Price Program (IPP) database of the BLS. Similarly to the PPI[6], the BLS samples these items in a three-stage design according to the value of shipments. The prices of these items tend to be allocative as discussed for example by Gopinath and Rigobon (2008), Neiman (2009) and Nakamura and Steinsson (2009). Moreover, I exclude from the data all prices associated with intra-firm trades. The preferred price basis in the data is "free alongside ship" (f.a.s.) or "free on board ship" (f.o.b.).[7] I concord the data to the six-digit NAICS level[8] to compare them with the PPI data. The appendix contains further details on the data.

One concern for my comparison with the PPI data could be that the IPP data contain a lot of imputed data. While this should not affect my analysis of price changes since I focus on price changes conditional on adjustment, imputation might affect my computation of frequencies of price changes. In

[5] For a detailed description of the sampling procedures, please see Appendix 9.4 and Chapter 14 of the BLS Handbook of Methods (US Department of Labor, 2008).

[6] For a detailed description of the IPP sampling procedures, see Chapter 15 of the BLS Handbook of Methods (US Department of Labor, 2008). The BLS additionally has a detailed description of the IPP data, the IPP Data Collection Manual, made available to the author.

[7] The first type of price basis accounts for 19.13% of the data, while fob (port) account for 1.11%, fob factory for 33.78% and fob without further specification for 6.13%, with the rest not further specified.

[8] I use the concordance available from the US Census Bureau Foreign Trade Data at http://www.census.gov/foreign-trade/reference/codes/index.html

particular, there are a lot of “pulled” prices in the data, about 19.82% of all prices. These pulled prices are prices carried over from last period, or “pulled” in BLS terminology. The existence of these pulled prices could imply that I may mistakenly detect a lot of artificial stickiness in the IPP data. Dropping imputed values which are flagged in the data as non-usable for calculation may have similar effects. However, as I argue in the data appendix, neither effort minimization, that is agents reporting no price change when there is one and thus inducing infrequent re-pricing and pulling prices forward, nor artificial stickiness due to non-tradedness nor the distribution of missing values relative to the incidence of price changes are likely to induce much artificial stickiness into results. In fact, after all manipulations of the data I find that the median time between any two price observations is exactly one month, equal to the sampling frequency, and the mean time between two observations is 1.3 months. I discuss the incidence of missing values in detail in the appendix and in addition, also take it into account in my calculations as described in the methodology section.

2.3 Newspaper Data

To complement the BLS data, I collect cover prices for The Economist. These prices are in local currency and cover 18 countries and have the advantage that there are no missing values in the time series. I collect these cover prices from each weekly print issue of The Economist for the time period from March 1990 to February 2001, thus for a total of 634 weeks. In addition, I obtain similar data of identical items sold in different export destinations: These data include the domestic and export sales price time series of 10 other newspapers from different countries of origin, price data from multiple years and countries in the annual IKEA catalogues and survey data on price-setting by European firms, collected through the efforts of the European Inflation Persistence Network (IPN).[9] I analyze these datasets in separate work (Schoenle, 2009) and refer to the results from that work where complementary.

2.4 Macro-Economic Data

I supplement the BLS micro price data with data on economic fundamentals. These include data on output, productivity, market concentration ratios, domestic and foreign inflation and exchange rates. I try to obtain these data at the monthly frequency whenever possible to match the monthly frequency of the BLS data. At the same time, I try to obtain these data at the most disaggregated levels available and then merge them by sector or destination country and date with the BLS data. In addition, I include a measure of geodesic distance from the U.S. in the data.

[9] I am grateful to Jonathan Haskel for sharing the IKEA data with me, and Ignacio Hernando, Fernando Martins and Harald Stahl from the Banks of Spain, Portugal and the Bundesbank for sharing their European Inflation Persistence Network data.

3 Differences in Domestic and Export Pricing Behavior

In this section, I document substantial differences in domestic and export market pricing behavior. First, I illustrate pricing differences using the data from The Economist. I find that on average domestic prices change more than twice as often as export prices but conditional on adjustment, export prices change by twice as much. In addition, I show that price changes are not synchronized across markets. Second, I document substantial differences in pricing behavior using the BLS micro-data on US producer prices. I find that US domestic prices change more than twice as frequently as export prices. In addition, synchronization between domestic and export price adjustment is low. Moreover, the size of export price changes conditional on adjustment is substantially larger than for domestic price changes. Finally, I document a new seasonality effects: export pricing synchronization increases towards the end of the year ("year-end synchronization"). These results hold both at the product-level and the firm-level and are robust to several checks for selection bias.

3.1 Methodology

Here, I briefly describe how I compute the frequency of price changes, my measures of synchronization, seasonal effects and the size of price changes for my data. In particular, I explain how I take into account missing values in the IPP data, how I calculate weights for the IPP-PPI comparison and how I compare the domestic and export data.

In the case of The Economist, first I compute the frequency of price changes for each sales destination as the mean of an indicator variable for price changes. Given that prices change relatively infrequently, I do not discard the first price spell and assume that the first observed price is associated with a price change[10]. I present results in terms of monthly frequencies and the average number of months a price spell lasts. Second, I compute the absolute size of price changes as the absolute log difference of prices for each destination, conditional on observing a price adjustment. Third, I calculate three measures of synchronization: the probability of adjustment in any export destination conditional on price adjustment in the UK, the probability of simultaneous adjustment in any two destinations, and the probability of adjustment in the UK conditional on price adjustment in any export destination. Since price changes might be slightly staggered, I compute the synchronization measures at a weekly, monthly and quarterly horizon.[11] I visually inspect the data to make sure that I am not missing any systematic staggering in price adjustment.

[10]Likely, this introduces a downwards bias in my estimate of price durations and makes my comparison more conservative since left-censored spells tend to be shorter.

[11]For example, differences in time zones might imply short adjustment lags.

In the case of the PPI and IPP data, I compute the frequency of price changes, the absolute size of price changes, measures of synchronization and seasonality in a very similar way but with special attention to the role of missing values in the IPP. In general, I compute my statistics first at the item-level, the smallest unit of observation in the data, then aggregate up in several steps to the six-digit product level and compare across PPI and IPP.[12]

First, I compute the frequency of price changes as the fraction of observed changes for each item, assuming that prices did not change when there are missing values. I drop left-censored spells but including them yields very similar results. Then, I aggregate up in the following way: First, I take the unweighted median of frequencies at the next higher classification level. Then, I take unweighted medians within each product category. In the case of exports, I compute this product-level frequency measure both for all export destinations pooled, as well as by separately examining each product-destination pair. I denotes these measures "Export I" and "Export II" in the tables[13]. To compare the resulting domestic and export frequencies, I compute the difference and ratios of export and domestic frequencies for each product or product-destination category. I proceed in an analogous to compute the absolute size of price changes conditional on adjustment.

Second, I calculate two measures of synchronization between domestic and export price adjustment: first, the monthly probability of export price adjustment conditional on adjustment of all goods in the same domestic product category and same direction, and second, the strength of synchronization obtained from a regression of upwards (downwards) price adjustment decisions of each exported good to the fraction of price upwards (downwards) changes in the domestic product category of this good. In both cases, I base my analysis on a synthetic subset of the domestic data which contains the same goods as the export data.[14] In the first case, I calculate monthly synchronization probabilities by computing the fractions of domestic and export upwards (downwards) price changes for each product-month-year, and additionally, for product-destination-month-years. Then, I compute the average fraction of upwards (downwards) export price changes conditional on observing full domestic upwards (downwards)

[12]This aggregation procedure is similar to that of Gopinath and Rigobon (2008) or Nakamura and Steinsson (2008a). Nakamura and Steinsson (2008a) begin computations at the cell-code which is a higher category than the item-level in the PPI. This is meant to take into account the possibility of no price changes in an item time series. At the same time, this induces a bias as in the case of ratio estimates in statistics. Empirically, however, this bias is small.

[13]In the case of the IPP data, when an item never changes price during its life-time, I follow Gopinath and Rigobon (2008) and impute the item's frequency of price change with the mean probability of item discontinuation within each six-digit category. This means taking the inverse of the total life of each item for items that at the end of the item life are either "out of scope, not replaced" or associated with firms "going out of business." Then, I average this number within each six-digit code and assign it to the items with no life-time price changes.

[14]I create the synthetic data by sorting goods j in each product category i by their sales value $S_{i,j}$. Then, I keep as many goods in each domestic category as in the corresponding export category: $n_i^D = n_i^X$. Since the BLS samples goods according to sales value, this should create a subset of the domestic data which contains approximately the same goods as the export data. I also use this synthetic dataset to conduct additional robustness tests.

adjustment of the same product. As a benchmark, I compute the fraction of other domestic products which change price in the same direction in the same four-digit category, conditional on full adjustment of the product under consideration.

In the second case, I obtain a measure of synchronization by estimating two sets of regression equations, one for upwards price adjustments, one for downwards price adjustments. In the upwards case, I estimate the following equations:

$$I(\Delta p^F_{i,j,t} \geq 0) = \beta_0 + \beta_1^{F+} f^{D+}_{i,t} + \epsilon_{i,j,t}$$
$$I(\Delta p^D_{i,j,t} \geq 0) = \beta_0 + \beta_1^{D+} f^{D+}_{i,t} + \epsilon_{i,j,t}$$

where $f^{D+}_{i,t}$ is the fraction of domestic goods j adjusting upwards in product category i at time t (excluding the item under consideration in the second equation). $I(\Delta p^k_{i,j,t} \geq 0)$ is an indicator variable for upwards adjustment of good j in product category i and market $k \in (D, F)$. Similar equations hold for downwards adjustment. I test the null hypothesis that synchronization is stronger for all goods within the domestic market than for the foreign export market: $H_0 : \beta_1^{D+} > \beta_1^{F+}$ and $H_0 : \beta_1^{D-} > \beta_1^{F-}$. I estimate export coefficients for all export destinations pooled, and additionally by treating destinations separately and then testing the null hypothesis.

Third, I complement the analysis of synchronization probabilities by considering the seasonal aspect of synchronization. I compute the fractions of export price changes conditional on complete domestic adjustment for each destination-product-month-year. Then, I calculate the median of these fractions across destinations and years and report the mean across all products in a given month.

Fourth, I compute two measures of frequencies of price changes which are robust to potential missing value problems in the IPP data. First, following the method suggested by Aucremanne and Dhyne (2004), I estimate the fraction of price changes in a given product-month and product-destination-month for which there are two consecutive observations so that a price change can correctly be inferred. Second, I take unweighted medians of this fraction across all months by product ("Exports I") or product-destination ("Exports II") and compare domestic and export product frequencies. In the tables, this first robust method shows up as "Method II."

As my second robust method, I estimate a constant hazard rate model for each item by maximizing the following log-likelihood function:

$$(1) \qquad L(j) = -\sum_i S^2_{i,j}\lambda_j + \sum_{UncensoredSpells} S_{i,j} ln(\lambda_j) + \sum_{RightCensoredSpells} S_{i,j} \ln(1 - e^{-\lambda_j(M_{i,j} - S_{i,j})})$$

where λ_j is the hazard rate for item j and $S_{i,j}$ denotes the ith price spell for that item. $M_{i,j}$ is the maximum possible length of a right-censored price spell computed as the number of months when a price different from the price at the beginning of the current spell is first observed. Following Neiman (2008), I assume that a price has not changed when it is the same before and after right-censoring. I also assume that every initial price marks the beginning of a price spell.[15] Again, I aggregate estimates up to the product categories and compare, referring to results under "Method III."

Finally, I calculate weights for my comparison by averaging PPI and IPP weights. Weights for the IPP data are derived from summing up to the six-digit product level the total value of shipments which are available at the classification level in the data but unfortunately not at the firm- or item-level. PPI total values of shipment at the six-digit level were provided by PPI staff. I calculate the weight for a product i by summing the respective domestic and export values of shipment, V_i^D and V_i^X and dividing by the total sum of shipments across all products:

$$w_i = \frac{V_i^D + V_i^X}{\sum_i (V_i^D + V_i^X)}. \tag{2}$$

where i is a six-digit product category.[16]

In the case of matched firms, I compute statistics in an exactly analogous manner to the product-level comparison: I compute item-level statistics and then aggregate up by taking the median across items within a firm. I compare unweighted domestic and export statistics since no export weights at the item- or firm-level are available. As an additional robustness check of composition and export selection effects, I compare only the largest domestic sales-item to export items of matched firms. Again, I take the median within firms of the respective statistics and then report medians across firms.

3.2 The Economist

The analysis of The Economist data presents succinct evidence for the differences in domestic and export pricing behavior of the same good: prices in export markets change substantially less frequently than the domestic price but when they change, export prices change by much larger percentages than domestic prices. At the same time, price changes of export and domestic markets are not very synchronized and staggered adjustment is not responsible for the lack of synchronization. There are no missing values underlying these results. In supplementary, individually collected datasets which I describe in Schoenle (2009) I similarly find that prices of identical items adjust systematically differently in the domestic and

[15] Estimating a constant hazard rate model to estimate the durations of price spells has the advantage of explicitly taking into account right-censoring. In particular, the hazard model puts higher weight on longer, uncensored price spells that contain "good" data. Similarly, it down-weights long, right-censored price spells.

[16] Since export shipment values are only available to me in total at the classification level but not by destinations, I calculate weights that are product- but not destination-specific. Because from a GDP point of view U.S. domestic sales dominate the distribution of production, as shown in Table 1 by the larger magnitude of PPI values of shipment compared to IPP shipments, destination-relative weights should be very similar if calculated.

the export markets.

Frequency I find that while the average monthly frequency of price changes in the U.K. home market is 11.67% which corresponds to an average duration of approximately 8.6 months, the export price changes occur approximately 2.5 times less frequently: I find that the export prices have a mean (median) frequency of 4.49% (4.40%) or durations of 22.22 (22.73) months. These are strongly statistically significantly different from the U.K. frequency and duration. Table 3 summarizes these results. Moreover, the table shows that export prices change less frequently for all foreign destinations within a wide range with average durations of 14 to 43 months.

Size of Price Changes I similarly find a large difference between domestic and export price changes: while the mean and median absolute domestic price change in the U.K. are 4.92% and 4.72%, the mean and median export price changes are approximately twice as large. As Table 3 shows, the mean of the country mean and median absolute price changes is at 10.96% and 10.06% and similarly, the median of the country mean and median price changes is at 11.40% and 9.15%. Again, these numbers are statistically significantly different from the domestic size of price adjustment.

Figure 1 illustrates the results on frequency and size of price changes: the dots represent the different countries and the associated mean absolute size and duration of price changes. The solid red lines go through the U.K. as point of origin while the dashed line plots the relationship implied by a linear regression of duration on absolute size of price changes. Two facts emerge: (i) the adjustments occur less frequently in the export destinations and all destinations except for three are subject to much larger average price changes (ii) the relationship between size of adjustment and duration is upwards-sloping. The less frequently prices change, the larger is the associated adjustment.

Synchronization Still, despite this pattern of less frequent but larger adjustment of exports compared to domestic sales, individual export price changes could be perfectly synchronized with domestic price changes. The results in Table 4 show that this is not the case: price changes are not synchronized across markets at all. In particular, the probability of export adjustment given domestic adjustment is 0.13 at a weekly horizon, 0.29 at a monthly and 0.54 at a quarterly horizon. Similarly, the probability of adjustment in the home market given adjustment in any export destination is 0.07 at the weekly, 0.15 at the monthly and 0.32 at the quarterly horizon. Finally, even when considering adjustment in any two markets, including the domestic market, the probability of synchronized adjustment is 0.38 at the weekly, 0.51 at the monthly and 0.54 at the quarterly horizon. Staggered adjustment cannot be driving much of these results since I have used three ways to compute synchronization probabilities in combination with intervals of different length and since I find resulting probabilities that are all much smaller than 1. When I additionally visually inspect adjustment patterns, I find no evidence of staggered adjustments.

3.3 BLS Data

The analysis of the BLS data shows that there are substantial differences in domestic and export market pricing behavior. These findings hold both at the product level and for matched firms. I first discuss the comparison at the product level, then the comparison for matched firms and finally, several robustness checks.

3.3.1 Frequency of Price Changes

I find large differences in the frequency and implied duration of domestically sold and exported products. Without yet directly comparing products, Table 5 summarizes the distribution of domestic and export frequencies and durations of the products used for comparison. For domestic prices, I find a monthly median frequency of price changes of 16% with an implied duration of 5.60 months.[17] Export prices change much less frequently: the median export frequency across all export destinations is 5.5%. The implied median duration is 17.84 months. When I calculate frequencies by product-destinations, I find very similar results, summarized under "Exports II" in the table. Now, the median frequency of export price changes is 6%. The implied median duration is 16.23 months. Weighted results have slightly higher frequencies and lower corresponding durations of price changes with a median domestic frequency of 18.5% and median export frequencies of 6.4% to 6.9%, as well as a median domestic duration of 5 months and median export durations of 14 to 15 months.

These values are quite similar to the estimates found in the recent literature that uses micro-price datasets. For US export prices, Gopinath and Rigobon (2008) report weighted median durations of 12.8 to 15.6 months, Neiman (2009) reports median and mean durations of 13 to 17 months and Nakamura and Steinsson (2009) report median frequencies of 7.2%.[18] For US domestic prices, Bils and Klenow (2004) find median domestic price durations of 4.3 to 5.5 months, Klenow and Kryvtsov (2008) 4 to 7 months, Nakamura and Steinsson (2008a) 7 to 9 months and Goldberg and Hellerstein (2009) 5 to 6 months.[19] Vermeulen et al. (2007) report a frequency of price change for European producer prices

[17] I compute implied durations given frequencies f as $d = -\frac{1}{\ln(1-f)}$.

[18] The slightly higher export durations associated with my main method are mostly due to my assumption of no price changes given intermittent missing values, but also due to my focus on a different time period, different methods of aggregation, weighting and a different subset of products. When I use my alternative methods to calculate export frequencies, I find weighted median durations of 12.4 to 14.5 for the method of Aucremanne and Dhyne (2004) and 10.1 months for the hazard rate estimates. The corresponding estimates to these two methods in Gopinath and Rigobon (2008) are 13.8 and 10 months.

[19] While my estimates for producer-price durations are lower than those in Nakamura and Steinsson (2008a), the reason for the difference is that I am not aggregating item-level statistics using finished good, intermediate good or crude material relative importance weights with the purpose of reporting estimates for these groupings. Relative importance weights can differ substantially, for example Farm Products receive a weight of 1.43% in the calculation of the BLS finished goods PPI, but 33.79% in the crude materials PPI. When I aggregate up to four-digit PPI codes and use relative importance weights

between 21% and 25%, or equivalently 4 to 5 months.

When I directly compare domestic and export price-setting product by product, I find large and statistically significant differences in the frequency and duration of price changes. Figure 2 gives a graphical summary of these results. Each point in the figure represents a product and its domestic and export frequencies of price adjustment. Almost all points lie below the 45-degree line and towards the export-price axis, illustrating how export prices tend to last substantially longer than domestic prices. In fact, as shown in Table 6, the median difference of export and domestic durations is 10.9 months for pooled destinations, and 9.3 months when I compare products destination by destination. When I use weights in the comparison, median differences range between 8.1 to 9.2 months.

In addition to comparing implied durations, I analyze the ratio of export to domestic frequencies. As Table 6 shows that the median ratio of export to domestic frequencies is 0.38 which is equivalent to domestic prices changing 2.7 times as often as export prices. The mean of the ratio across products is 0.54, equivalent to 1.85 more frequent adjustment in the domestic market. When I calculate the ratio separately for each product-destination pair, I find a median ratio of 0.40 and a mean ratio of 0.65, equivalent to 2.5 times and 1.6 times more frequent domestic price adjustments.

I find smaller but still large and statistically significant differences between frequencies of domestic and export price adjustment when I compute the frequency of price changes in the export data using the two alternative methods described in the methodology section. This suggests that missing values in the export data are not responsible for inducing as much artificial stickiness into the data as to be driving my results. First, as Table 7 shows, computations according to the method of Aucremanne and Dhyne (2004) ("Method II") imply a median difference of export and domestic durations of 7.4 months and a median frequency ratio of 0.47: domestic prices adjust approximately 2.13 times more frequently. When I consider product-destination pairs, I find a median difference of 6.8 months, as well as a median ratio of export to domestic frequencies of 0.42.

Results from duration-based hazard rate estimations also imply very large and significant differences in the timing of domestic and export price-setting, summarized under "Method III". I find a median ratio of export of domestic adjustment frequencies of 0.59, equivalent to 1.7 times more frequent domestic price adjustment. The implied median duration difference is 4.6 months. When I consider product-destination pairs, results imply a median difference of 5 months, as well as a median ratio of export to domestic frequencies of 0.56. Thus, while there is some variability in the magnitude of the difference of export and domestic durations across my three methods of frequency calculation, I find in all specifications that

as in Nakamura and Steinsson (2008a), I find a duration of 8.93 months for finished good producer prices replicating their findings. Moreover, the unweighted mean of the unweighted four-digit PPI frequencies reported in Table 23 of the online supplement to Nakamura and Steinsson (2008a) is 19.9%, implying an overall U.S. producer price duration of 4.5 months. Given my methodology of calculation, this number is most comparable to my PPI results.

median ratios of export to domestic frequencies lie robustly between 0.38 and 0.59. This is equivalent to domestic price changes that are between 1.7 and 2.7 times as frequent as export price changes.

3.3.2 Synchronization

My analysis shows that both upwards and downwards export pricing decisions are not very synchronized with domestic adjustment decisions. Moreover, I document a second new kind of seasonality effect in form of year-end synchronization.

First, as shown in Table 8, the mean probability of synchronized upwards price adjustment in the domestic and export markets is 20.81%. This is larger than the probability of synchronized downwards price adjustment, which is 14.14% across products. When I compute the synchronization probabilities first by product-destination pairs and then aggregate to the product-level, I find very similar results: the mean probability of upwards synchronization across products is 20.37%, the mean probability of downwards synchronization is 13.76%. The higher probability of upwards relative to downwards synchronization suggests that a model of adjustment decisions in the domestic and export market should include a common shock with an upwards trend as well as idiosyncratic shocks at the product-level.

There are two natural benchmarks to which to compare these synchronization results: First, since U.S. exports are denominated in USD in the data, the Law of One Price would imply perfect synchronization for the timing of price adjustment. Second, another benchmark is the probability of adjustment of other products in the same four-digit category conditional on adjustment of the product under consideration: this probability is 47.48% for upwards adjustment, and 50.31% for downwards adjustment. Clearly, the probabilities of synchronization are less than these benchmark probabilities. This suggests that domestic and export price changes are not only far from being perfectly synchronized but are also less synchronized than price changes of domestic products amongst each other within narrowly defined categories.

Second, when I estimate the extent of domestic and export market synchronization using regression analysis, I find very similar results as shown in Table 9. A monthly increase by one percentage point in the fraction of items which adjust upwards (downwards) in the domestic market is on average associated with a 0.10% (0.05%) higher probability of upwards (downwards) export price adjustment for an item in the export market and in the same product category. When I estimate the regression equations by product-destinations and then average first across destinations and then products, I find similar coefficients of 0.15% for upwards adjustments and of 0.04% for downwards adjustment. In addition, I use the regressions to obtain another benchmark: when I regress upwards (downwards) adjustment decisions of individual domestic items on the fraction of items which adjust upwards (downwards) in the domestic market[20], I find that a one percent increase in the fraction of items which adjust upwards (downwards) in the domestic market is associated with a 0.7% increase in the probability of domestic price adjustment of other items in the same product category. Thus, adjustment decisions tend to be significantly and substantially more synchronized within domestic products than across domestic and export markets for the same products.

Finally, in addition to a lack of synchronization I uncover a new kind of seasonality effect in the data. Figure 3 illustrates this effect. Conditional on full adjustment in the domestic market and irrespective of the sign of adjustment, the mean probability of export adjustment increases from a 40% trough in March towards values around 60% to 80% in October, November, December and January, with the exception of an intermittent peak of 62% in June. A very similar effect persists when considering the median synchronization probability.

3.3.3 Size of Price Changes

My results shows that export price changes, conditional on adjustment, are much larger than domestic price changes: in my data, the mean absolute domestic price change is 5.77% conditional on adjustment while the mean absolute export price change is 8.76%, and 11.83% when calculated separately for each destination and product combination. When I compare domestic and export prices directly by product, I find a median and mean difference of 1.75% and 3% between export and domestic price changes. When I use weights, the median and mean differences are 1.54% and 2.27%. Similarly, I find median and mean ratios of export to domestic price changes of 1.39 and 1.81, and of 1.19 and 1.57 using weights. Similar results hold when I consider product-destination pairs. Table 10 summarizes my results on differences between domestic and export price changes and Figure 4 shows how most domestic-export price points indeed tilt away from the 45-degree line towards the export price axis.

3.3.4 Matched Firms

An additional novelty of my analysis is a direct comparison of domestic and export price-setting behavior of identical firms. For these matched firms, I find that substantial and statistically significant differences in domestic and export pricing behavior persist. Results are robust and very similar quantitatively to the case of product comparisons. Moreover, this finding suggests that firm selection into exporting such as posited by Melitz (2003) is not driving differences in export and domestic price-setting behavior.

Using a fuzzy matching algorithm on PPI and IPP firm names[21], I match 381 firms selling 1950

[20]For this set of regressions, I calculate the fraction of domestic price adjustments in the product category excluding the adjustment decision of the item under consideration.

[21]Appendix 9.4 contains details of the algorithm which I use to identify matches.

items in the domestic data and 1594 items in the export data. Sales cover 91 countries and 181 product categories. The median number of exported items of the matched firms is 5.37 and the median number sold domestically is 6.42. Thus, an approximately equal number of items is sampled for each firm. The median number of employees in the unmatched and the matched data is similar with 145 and 129 employees, and both mean and median sales values are of similar magnitude. I compare characteristics of matched and unmatched firms in Table 2.

When I compare domestic and export price dynamics for these matched firms, I find almost identical results as when I compare domestic and export dynamics at the product level. Table 11 shows the results of the comparison. First, the median ratio of export to domestic frequencies is 0.484 which is significantly different from zero and one: domestic prices change 2.07 times as frequently as export prices. Second, the ratio of domestic to export implied durations is similarly 0.452. The associated implied durations are 5.77 months for domestic and 15.46 months for export prices. Third, the ratio of the absolute size of export to domestic price changes is 2.48, with an absolute size of 9.57% for export and 4.60% for domestic prices. Fourth, I compute the conditional probability of export price adjustment irrespective of the signs of price change and given full adjustment in the domestic market. I find a mean probability of synchronization of 0.45 across firms which is significantly different from zero and one.

Because these results are based on a comparison of identical firms, they present a robustness check of my comparison of products across the domestic and export market. A comparison of products which takes into account which firms sell in the domestic and the export market is important because the comparison could be subject to a compositional effect due to export selection of firms: analogously to the selection of firms into exporting according to their level of productivity as posited in Melitz (2003) and documented by Bernard et al. (2006) or Goldberg et al. (2008), exporting firms could, for example, differ in the dynamics of their productivity or demand processes. In the context of price-setting, this could induce differences in export price-dynamics. Since I find almost identical if not stronger results when I compare data of matched firms, this underlines the robustness of my findings to firm selection. I discuss several additional robustness checks in the following section.

3.3.5 Selection Bias

In addition to the robustness check on the selection of firms into exporting, I present a robustness check on export selection in my comparison of products, a robustness check on item selection within exporting firms and a robustness check based on export shares as a proxy of sampling selection. I also discuss supplementary evidence from related work.

First, I show that differences in domestic and export price adjustment at the product-level are not due to selection of different items and firms in the domestic and export data. Since there are more items in the domestic product categories than in the corresponding export categories and since exporting firms differ from firms which sell only domestically as discussed by Melitz (2003), my findings at the product-level could be due to a composition effect: there are different items and firms in each domestic and export product category and their price dynamics differ. I address this concern by using my synthetic domestic dataset to compare the frequency and size of price changes for each product.[22] I find that the ratio of export to domestic frequencies and the difference in implied durations are not statistically significantly different in the synthetic dataset compared to the full dataset. In fact, the difference between domestic and export market price-setting is even more pronounced. The reason is that the synthetic data contain items with relatively large sales values which tend to adjust more frequently. This is consistent with the finding of Goldberg and Hellerstein (2009) that weighting items by sales values increases the frequency of price changes in the PPI data. Differences in the size of domestic and export market price changes are similarly larger in the synthetic data. Table 12 summarizes the robustness check for frequencies.

Second, I show that export item selection is also not driving my results. Even though the IPP sampling procedure is meant to select distinct goods to be representative of the universe of exports, one could think that the same domestic item is sampled in different export markets. Therefore, when I compare items other than the largest domestic sales item to items in the export data, this could similarly introduce a composition effect into my comparison of the domestic and export data. I check the robustness of my findings against this concern in two ways.[23] First, I keep only the largest sales-item in each domestic product category and repeat my comparison of domestic and export frequencies of price adjustment. As Table 12 shows under the label of "Synth II", the difference in domestic and export price adjustment persists as in the full data and again is even stronger. Second, I keep only the largest domestic sales-item in each matched firm. Then, I compare statistics for this item to the statistics of exported items for each export destination within a firm. I find that my results remain robust as Table 13 shows: domestic price changes are 2.08 times as frequent but export price changes are 2.3 times as large, conditional on adjustment.

Third, I provide another robustness check that sampling selects different goods in the domestic and export data and therefore introduces a bias into my comparison. I check for composition effects using a simple regression of differences in domestic and export frequencies on product log export shares. The idea is that a small log export share is a proxy for a high chance that sampling among the universe of

[22]This robustness check relies on the fact that firms and items both in the PPI and IPP are sampled according to their total sales value. Therefore, when I sort the domestic data according to sales value and truncate them so that the number of items is the same for each product in the domestic and export data, the resulting domestic and export data should contain approximately the same goods and firms in each product category.

[23]This robustness check is based on the finding in the empirical trade literature that firms export their domestic sales "hit" to all countries. Most likely, this item would be the item sampled in all sales destinations.

domestic goods can select a good different from the relatively small universe of exported goods.[24] Since I find no statistically significant relationship between the difference of domestic and export frequencies and the log export share, this suggests that the selection in the sampling stage is not likely to drive my results. Figure 5 plots the difference of frequencies against the log export share and shows that there is no pattern or systematic relationship between the frequency difference and export shares.

Finally, I have presented evidence from cover prices of The Economist that show how even prices of exactly identical items are subject to differential dynamics between domestic and export markets, with more frequent and smaller adjustments in the domestic market, less frequent but larger adjustments in the export market and an absence of synchronization. Certainly, there is no firm and item selection bias driving these differences. In supplementary work summarized in Schoenle (2009), I show that the same differences in price dynamics persist for 10 newspapers from different countries sold internationally, and 100 identical IKEA products sold in 20 countries. In all of these cases, firm and item export selection should not be responsible for differences in price dynamics.

4 Model

In this section, I present a simple, dynamic menu cost model of price-setting: given a set of changing economic "fundamentals" in each market in each period, a firm has to trade off the decision to adjust prices with the cost of changing prices. The set of economic fundamentals include inflation, demand shocks and exchange rate shocks in the case of exports. Moreover, the firm faces a common productivity shock for both domestic and export production and demand is subject to a world shock component common to both countries.

I use the model for three purposes: First, I use the model to motivate a regression analysis of price-setting decisions and fundamentals. Second, I derive a structural equation between fundamentals and pricing synchronizations to understand the mechanisms underlying domestic and export market synchronization. Third, as the main exercise, I use the model to back out adjustment costs for domestic and export market price changes based on the empirical moments presented in the previous section.

4.1 Model Setup

The model is a partial-equilibrium model where a firm produces and sells a single product in two segmented markets. I denote the domestic market by H and the foreign export market by F. Throughout,

[24]The extreme case of a zero export share means that with certainty a different domestic good will be selected from the non-existing exported good. At the same time, a 100% export share means that the chance of sampling a different, non-exported domestic good has become smaller.

I assume that the firm takes as given the aggregate processes described below.

First, production technology is linear subject to the same technology shock for domestic and export market production:

$$Y_t = A_t L_t \tag{3}$$

where A_t denotes period productivity per unit of labor of the firm and L_t the total amount of labor used in production. I assume that log productivity follows an AR(1) process:

$$\ln A_t = \rho_A \ln A_{t-1} + \epsilon_{A,t} \tag{4}$$

where shocks $\epsilon_{A,t}$ are i.i.d. normal with $E[\epsilon_{A,t}] = 0$ and $var(A_t) = {\sigma_A}^2$. In line with the partial equilibrium approach, I normalize the real wage[25]:

$$\frac{W_t^H}{P_t^H} = \frac{\theta - 1}{\theta} \tag{5}$$

where W_t^H denotes the wage rate and P_t^H the domestic aggregate price level.

On the demand side, consumers in the domestic market demand amount c_t^H of the good:

$$c_t^H = C_t^H \left(\frac{p_t^H}{P_t^H}\right)^{-\theta} \tag{6}$$

where p_t^H denotes the price of the good produced and C_t^H domestic aggregate demand. Consumers in the export market demand amount c_t^F of the good:

$$c_t^F = C_t^F \left(\frac{p_t^F E_t}{P_t^F}\right)^{-\theta} \tag{7}$$

where p_t^F denotes the producer-currency export price of the good and P_t^F the foreign aggregate price level. C_t^F denotes aggregate foreign demand, E_t the nominal exchange rate with units of foreign currency over USD. With this formulation of demand, I go with the assumption of producer currency pricing since approximately 98% of all US export prices are set in USD and I have restricted my analysis to dollar-priced exports only. Gopinath and Itskhoki (2010) explicitly model endogenous invoice currency choice.

I assume that aggregate demand C_t^i in market $i \in (H, F)$ is exogenously given and follows an AR(1)

[25] In a full general equilibrium model, one could for example include a numeraire sector subject to frictionless international trade so that wages would be fixed.

process:

$$\ln C_t^i = \rho_C \ln C_{t-1}^i + \eta_{C,t}^i \tag{8}$$

where demand shocks $\eta_{C,t}^i$ can be thought of as a composite of a country-specific shock and a world shock. For computational simplicity, I identify the importance of this world shock with a single parameter, namely the correlation of domestic and export market shocks, $\rho_{H,F}$. I assume that the shocks have zero mean, $E[\eta_{C,t}^i] = 0$, variances $\sigma_{C,H}^2 = \sigma_{C,F}^2 > 0$ and correlation $\rho_{H,F}$. Note that $\rho_{H,F} = 0$ implies that demand shocks are uncorrelated and $\rho_{H,F} > 0$ means that demand shocks are more correlated, with common world shocks more important relative to market-specific shocks. While I do not model it explicitly, a similar world-shock could be modeled for inflation.

Finally, I also assume the following about the dynamics of price levels and the exchange rate. First, I assume that the domestic price level follows a random walk with a trend:

$$\ln P_t^H = \mu^H + \ln P_{t-1}^H + \eta_{P,t}^H \tag{9}$$

where shocks are i.i.d. with $E[\eta_{P,t}^i] = 0$ and $\sigma_{P,H}^2$. Second, to reduce the number of state variables in the export market, I re-write demand as $c_t^F = C_t^F \left(\frac{p_t^F E_t}{P_t^F}\right)^{-\theta} = C_t^F \left(\frac{p_t^F}{P_t^H}\frac{P_t^H E_t}{P_t^F}\right)^{-\theta} = C_t^F \left(\frac{p_t^F}{P_t^H} Q_t\right)^{-\theta}$ and assume that the real exchange rate is exogenously given by the following process:

$$\ln Q_t = \rho_Q \ln Q_{t-1} + \epsilon_{Q,t} \tag{10}$$

where shocks $\epsilon_{Q,t}$ are i.i.d. normal with $E[\epsilon_{Q,t}] = 0$ and $var(Q_t) = \sigma_Q^2$.[26]

Under these assumptions, period profits in the domestic market are given by:

$$\pi_t^H = \left(\frac{p_t^H}{P_t^H} - \frac{1}{A_t}\right) C_t^D \left(\frac{p_t^H}{P_t^H}\right)^{-\theta}$$

while period profits in producer currency in the export market are given by:

$$\pi_t^F = \left(\frac{p_t^F}{P_t^H} - \frac{1}{A_t}\right) C_t^F \left(\frac{p_t^F}{P_t^H} Q_t\right)^{-\theta}$$

The firm's dynamic problem is to choose when to adjust the price and if so, to what optimal level.

[26] I verify empirically that the innovations to the domestic price level, foreign aggregate output and the real exchange rate are approximately uncorrelated. I estimate an unconstrained vector auto-regression of these variables for 29 countries. Based on the estimated variance-covariance I find that the mean correlation of the domestic price level with foreign output is 0.023, with the real exchange rate is 0.021 and of the real exchange with output is 0.004. In addition, I also regress CPI-based real exchange rate time series on their components. I find that the nominal exchange rate accounts for 95.37% percent of the real exchange rate variation, foreign inflation for an additional 3.91% and U.S. inflation for 0.70%.

If the firm chooses to adjust prices, it has to pay K^D units of labor to adjust in the domestic market and K^X units of labor to adjust in the export market. I assume that menu costs represent a fraction of steady state revenues.[27] The problem for market i can be expressed recursively using the following two value functions for the adjustment and no-adjustment case:

$$V_i^{adj}(S_t^i) = \max_{p_t^i}\{\pi_t^i(p_t^i) - K^i + \beta \mathbf{E}[V_i(S_{t+1}^i)]\} \tag{11}$$

$$V_i^{noadj}(S_t^i) = \pi_t^i(p_{t-1}^i) + \beta \mathbf{E}[V_i(S_{t+1}^i)] \tag{12}$$

where

$$V_i(S_t^i) = \max\{V_i^{adj}(S_t^i), V_i^{noadj}(S_t^i)\} \tag{13}$$

and the states in Home and Foreign are given by

$$S_t^H = (A_t, \frac{p_{t-1}^H}{P_t^H}, C_t^H) \tag{14}$$

and

$$S_t^F = (A_t, \frac{p_{t-1}^F}{P_t^H}, C_t^F, Q_t) \tag{15}$$

where I have divided by the price level to obtain stationarity. In the case of adjustment, the past price is not a relevant state variable. In the case of non-adjustment, however, the period t profit on the right-hand side is a function of the price from period $t-1$ in terms of the time t price level. $\beta < 1$ is a discount factor which again in line with the partial equilibrium approach is assumed to be constant.

4.2 Model Dynamics

Here, I first solve the model numerically and discuss the relationship between economic fundamentals and price-setting. Second, I derive a structural equation that relates fundamentals to pricing synchronizations across markets. I discuss how to empirically test the structural equation.

To solve the model, I employ numerical methods: I use collocation methods, approximating each value function with a set of polynomials and solving a system of non-linear equations to find the basis coefficients of the polynomial. Collocation methods are for example described in Miranda and Fackler

[27]The proportionality assumption is guided by the fact that U.S. imports are similarly sticky to US exports, in particular US imports from small countries and U.S. exports to large countries.

(2002) and the appendix contains details of my implementation. Given the policy functions derived from these numerical solutions to the value functions, I use simulation to study the predictions of the model. In particular, I assume a monthly time frame which corresponds to the frequency of my micro-data and simulate, for a given set of parameter values, a set of shocks for many periods. I record price responses from the policy functions in each period and simulate many time series. I use these time series of prices and fundamentals to show first, how individual adjustment decisions relate to shocks and second, to derive a structural equation for synchronization of price-setting across markets.

In a first step, I estimate a multinomial logit model on the simulated data to relate individual adjustment decisions to shocks.[28] I regress the simulated adjustment decisions on the simulated model fundamentals X, that is changes in productivity, demand, the inflation rate and the exchange and impose a multinomial logit link function with 3 categories: $m = 0$ for no price change, $m = 1$ for a price increase and $m = -1$ for a price decrease. The multinomial logit model is described in detail for example by Agresti (1996). Denoting by Π_m the probability that decision m is taken, the probability under the multinomial logit link is given by

$$\Pi_m = \frac{e^{X\alpha_m}}{\sum_m e^{X\alpha_m}} \tag{16}$$

Since $\sum_m \Pi_m = 1$, the three sets of parameters are not unique. Therefore, I follow standard practice and choose category $m = 0$ as a baseline category:

$$\Pi_0 = \frac{1}{1 + \sum_{m\in(-1,1)} e^{X\alpha_m}} \tag{17}$$

I estimate the two remaining logit equations simultaneously. The logit model has the convenient property that the estimated coefficients take on the natural interpretation of the effect of the explanators on the probability of choosing the action of adjusting prices up or down over taking no action.

The model has the following predictions, summarized in Table 14: (i) an increase in the inflation rate is associated with a higher probability of upwards adjustment, and a lower probability of downwards adjustment, (ii) an increase in productivity is associated with a higher probability of downwards price changes and a lower probability of upwards price changes, (iii) higher consumption levels are associated with higher probabilities of both upwards and downwards adjustment. While (i) and (ii) are very intuitive, (iii) can easily be understood in terms of relative menu costs: an increase in market size means larger profits while menu costs stay constant and decrease relative to profit so that adjustment

[28]Estimating a multinomial model and as well as taking it to the data has the advantage over conventional probit analyses of micro data that it allows me to separately examine the relationship of upwards and downwards adjustment decisions and economic fundamentals.

is more frequent both upwards and downwards. (iv) a real exchange rate appreciation is associated with lower probabilities of upwards and downwards adjustment. I obtain (iv) because the exchange rate enters isomorphically to market size and similarly changes relative menu costs. More fundamentally, this prediction is due to my assumption of exporter producer-currency pricing and the assumption of exogenous real-exchange rate shocks. I discuss in the appendix how I test the relationship between individual adjustment decisions and economic fundamentals for an extended set X of fundamentals.

In a second step, I estimate a simple structural equation that relates the effect of common productivity and world shocks to synchronization in international pricing decisions, both using simulated and actual data. The purpose of this exercise is to examine how fundamentals relate to synchronization decisions and to test whether the mechanisms of the model are at work in the data. I estimate the following equation:

$$I(\Delta p^X_{i,t} \neq 0) = \beta_0 + \beta_1 I\{\Delta p^D_{i,t} \neq 0\} + \epsilon_{i,t} \tag{18}$$

relating the decisions of adjusting export prices to the decision of adjusting a domestic price. The key parameter is $\beta_1(M)$ where M denotes common shocks. I trace out $\beta_1(M)$ by estimation on simulated data and systematically increasing M. Figures 18 and 19 show the results: First, when inflation correlation or demand correlation increase, estimated β_1 monotonically increases. The probability of synchronized price-adjustment goes up. Since my simulation additionally includes common productivity shocks with exchange rate shocks being absent, I find that there is always a non-zero probability of synchronization even when demand or inflation are only modestly correlated across markets. In addition, since exchange rate variability is absent from the simulations, perfect correlation in shocks obviously implies perfect synchronization. Second, as Figure 19 shows, an increase in the importance of common productivity shocks similarly increases estimated β_1. Clearly, given the absence of perfect correlations of shocks across destinations, the model implies that firms set prices dynamically in each market and that the Law of One Price does not hold. Adding exchange rate variability only aggravates this prediction.

In the data, I test the structural equation (18) of the model relating to synchronization of price-setting in domestic and export markets in two ways. First, I estimate

$$f^F_{p,c,t} = \beta_0 + \beta_1 f^H_{p,t} + \beta_2 M_{p,c,t} + \beta_3 f^H_{p,t} * M_{p,c,t} + \epsilon_{p,c,t} \tag{19}$$

where $f^F_{p,c,t}$ and $f^H_{p,t}$ denote the monthly fraction of price changes in export destination c and the domestic market for each six-digit category p at time t, and $M_{p,c,t}$ denotes one of the following: (i) the standard deviation of productivity growth rates, σ_A, (ii) the correlation between US and export destination monthly aggregate output growth, $\rho^y_{H,F,c}$ or (iii) the correlation of monthly CPI inflation rates, $\rho^\pi_{H,F,c}$

The model predicts increased synchronization between domestic and export adjustment due to larger common shocks, that is $\beta_3 > 0$.

Second, I identify β_1 non-parametrically by separately estimating the following equation for the 1st and 99th percentiles of fundamentals (i), (ii) and (iii):

$$f^F_{p,c,t} = \beta_0 + \beta_1 f^H_{p,t} + \epsilon_{p,c,t} \tag{20}$$

That is, I consider only data $f^H_{p,t}$ and $f^F_{p,c,t}$ if the associated standard deviation of productivity shocks, the correlation between US and export destination monthly output growth, and the correlation of monthly inflation rates fall into the first or 99th percentile of their distributions. I expect a higher estimated coefficient β_1 and higher explanatory power for the 99th percentile when common shocks matter more, than for the 1st percentile.

5 Regression Analysis

Results from the regression analysis show that economic fundamentals are significantly related to pricing decisions both in the domestic and the export markets. Moreover, when I estimate the structural synchronization equation, I find that common shocks drive synchronization across markets. However, a large fraction of the variation in adjustment decisions is left unexplained.

5.1 Adjustment Decisions and Fundamentals

Here, I briefly describe the results of the estimation of the multinomial logit model with further details given in the appendix: when I estimate the multinomial model given by (16), I find that economic fundamentals such as inflation, demand, concentration ratios, productivity, firm size and exchange rate changes are significantly related to decisions to adjust prices. Similar results hold for both domestic and export market adjustments. In particular, I find that pricing decisions are not only related to fundamentals in a state-dependent but for some variables also in an asymmetric way as predicted by the model. For example, increases in inflation are related to higher probabilities of upwards and lower probabilities of downwards adjustment. However, a large fraction of the variation in adjustment decisions both for domestic and export market are left unexplained, with an R^2 of 23% for the domestic and 11% for the export pricing decisions.

5.2 Structural Estimation

When I estimate structural synchronization equation (18), I find that fundamentals are significantly related to synchronization of domestic and export market pricing decisions. I find particularly strong effects in my semi-parametric regression. However, while the mechanisms of the model seem to be at work, explanatory power is low and most of the variation in synchronization is left unexplained.

First, when I estimate (19) based on interactions of fundamentals and domestic adjustment, I find that larger productivity shocks, stronger domestic and export market output as well as inflation correlations are positively and significantly associated with synchronization of pricing decisions across markets. These results are summarized in Table 15. The coefficients on the interaction terms of fundamentals and US monthly adjustment fractions report the strength of the synchronization effect. I find that the largest effect is due to correlation in demand changes, with a coefficient of 0.5414: given a fraction of domestic adjustment, a 1% increase in output growth correlation is associated with an additional 0.54% higher probability of export price adjustment. The second largest coefficient is associated with inflation correlation, at 0.1044, and productivity takes the smallest coefficient with 0.0077. All of the coefficients are statistically significant but the productivity interaction term enters no longer significantly when I control for month fixed effects. As predicted by the model, the coefficients on the interaction terms are positive, significant and between 0 and 1.

Second, when I identify the synchronization coefficient semi-parametrically, I find particularly large differences in the synchronization effect of the common productivity shock and output and inflation correlation between the extremes of their distribution. I estimate (20) separately for the 1st and the 99th percentile of the productivity measure, output growth and inflation correlations. I plot the estimated synchronization coefficient β_1 in Figures 6 and 7: First, I find that β_1 increases from less than 0.05 for the first percentile to almost 0.3 for the 99th percentile, both for demand and inflation correlations. Second, β_1 increases from 0.2 to almost 0.8 when going from the first to the 99the percentile of the productivity measure. Strikingly, these estimates do not only take the positive sign predicted by the model but also lie in the range from 0 to 1 with values similar to the ones obtained in the simulations. Moreover, I find that explanatory power also increases the larger common productivity shocks and the stronger correlations become: for the productivity measure, explanatory power increases from 5.07% to 8.67% and for output and inflation correlations from 0.05% and 0.03% to 5.52% and 5.65%.

Economically, however, both kinds of regressions do not explain much of domestic and export market synchronization. I find that the R^2 is low, taking on values between 4% and 9%. When I try alternative specifications, I find similar results. Thus, while the mechanisms of the model clearly are at work, they cannot account for the main differences in domestic and export market pricing synchronization. Therefore, in the next section, I attribute residual differences in domestic and export pricing decisions to differences in domestic and export market menu costs and assess if this yields reasonable estimates of menu costs.

6 How Large Are Export Menu Costs?

In this section, I ask how large domestic and export menu costs have to be to account for differences in the dynamics of price adjustment across markets after I take into account differences in economic fundamentals. I use stochastic simulation techniques to find the maximum likelihood estimates for domestic and country-specific export market menu costs implied by my model. I estimate domestic market menu costs of 0.4% of period-steady state revenues, and median export market menu costs between 1.5% and 1.75% for all three likelihood exercises. To find the menu cost maximum likelihood estimates, I proceed in the following way: In a first step, I use stochastic simulation in combination with parameter grid search to obtain country-specific likelihood functions for menu costs.[29] I calibrate country-specific parameters by estimating them from the data and inserting them into the simulations. In a second step, I calculate maximum likelihood estimates based on the estimated likelihood functions.

6.1 Calibration and Maximum Likelihood Estimation

Thus, in a first step I obtain parameter values Ω_c for productivity, domestic and export market inflation and exchange rates for each country c from the data[30], attributing them to the model in the interpretation of a representative sector in each country. I assume away domestic and export market consumption variability. This assumption is likely to work in my favor since for my timeframe, the U.S. is subject to less variable demand shocks than other countries.[31] To calibrate the representative U.S. productivity process, I estimate an AR(1) of the four-factor productivity measure for each of the 459 sectors in the NBER productivity database, for the years 1984-1996. The associated implied monthly AR(1) coefficient is 0.96, with a standard deviation of 2.1%. To obtain the real exchange rate parameters, I construct monthly CPI-based real exchange rate time series from 1998-2005 for each country and estimate a set of AR(1) real-exchange rate processes. The results imply highly persistent real exchange rate dynamics

[29]While non-standard for maximum-likelihood estimation, employing stochastic simulation is necessary because no analytical expressions for the probability distribution of price adjustment decisions given fundamental shocks are available. For example, Gourieroux and Monfort (1997) discuss the use of stochastic simulation techniques to obtain maximum likelihood estimates. This combination of techniques is also very commonly used in bio-statistics.

[30]The countries in this exercise are: Belgium, Brazil, Canada, Switzerland, Chile, Germany, Denmark, Spain, France, Great Britain, Greece, India, Ireland, Israel, Italy, Japan, Korea, Mexico, the Netherlands, Russia, Sweden, Slovenia and Turkey.

[31]Using my data, I calculate a monthly standard deviation of U.S. aggregate demand shocks of 0.52% based on an AR(1) regression of monthly output. At the same time, export markets are subject to larger shocks: I compute a standard deviation of 0.64% for the E.U., 0.88% for Canada and 0.86% for Mexico, the largest U.S. trading partners.

as found in the literature, for the US-UK real exchange rate for example, the AR(1) coefficient is 0.9763 with a standard deviation of 1.83%[32]. I impose symmetry on the elasticity of demand in the domestic and export markets: $\theta_c = \theta = 4$[33]. For US inflation, I assume a monthly trend of $\mu^H = 0.21\%$ with a standard deviation of $\sigma_{P,H} = 0.37\%$ as used by Nakamura and Steinsson (2008b) to calibrate U.S. CPI inflation.

To obtain estimates of the likelihood functions, I choose a grid of menu costs between 0.05% and 6.5% of fractions of period steady state profits. Then, given parameter values, I simulate a time series for each menu cost K_i on this grid, recording the adjustment decisions and associated price changes. I compute the frequency of price change and the mean absolute size of price changes for the resulting series of adjustment decisions. I simulate many more time series for the same menu cost and parameter values, each time computing the frequency and mean absolute size of price changes. Using the resulting distributions, I calculate the likelihood $L(K_i|\Omega_c, f_c)$ of menu cost K_i given the frequency of price changes observed in the data, the likelihood $L(K_i|\Omega_c, \Delta p_c)$ given the size of price changes and the likelihood $L(K_i|\Omega_c, f_c, \Delta p_c)$ given both the frequency and size of price changes. Figure 9 shows the moments I am trying to match. I repeat the simulation exercise at all grid points to obtain an estimate of the likelihood function.

In a second step, I calculate maximum likelihood estimates. In the case of domestic sales, this is simply the maximum of the domestic likelihood over the menu cost grid, given the respective moments. In the case of exports, I additionally calculate a uniform export menu cost. Thus, I first compute the maximum likelihood estimate for each country based only on data on country-specific adjustment frequencies:

$$K^*_{c,1} = argmax_{K_i} L(K_i|\Omega_c, f_c)$$

Second, I calculate the estimates given only observed price changes for each country:

$$K^*_{c,2} = argmax_{K_i} L(K_i|\Omega_c, \Delta p_c)$$

Third, I compute the maximum likelihood estimates of the menu cost conditional on the both frequency and size of price changes for each country:

$$K^*_{c,3} = argmax_{K_i} L(K_i|\Omega_c, f_c, \Delta p_c)$$

I condition on frequency and size of price changes separately as well as jointly because it is not clear a

[32] Figure 8 summarizes the real exchange rate parameters used.

[33] The common parameter values in the literature range between 3 and 7. In particular, Broda and Weinstein (2006) estimate import demand elasticities for 73 countries and report median estimates of 3.7 to 3.2 at the three-digit level.

priori whether the model is able to match the larger set of moments or just the smaller sets. Finally, I compute a single maximum likelihood estimate based on pooling data on frequencies and size of price changes across export destinations, imposing a uniform export menu cost $K_c^* = K^*$. To this purpose, I sum the log of the likelihoods estimated at each grid point across countries. Then I compute the maximum across grid points:

$$K^* = argmax_{K_i} \sum_c \ln L(K_i|\Omega_c, f_c, \Delta p_c)$$

6.2 Results

I find that export menu costs are substantially larger than domestic menu costs: First, conditioning only on country-specific frequency data, estimates of export market menu costs range between 0.25% and 6% of period steady state profits. As shown in Table 16, I estimate a median of 1.5% and a mean of 1.97% with a standard error of 0.34%. All underlying likelihoods are well-behaved with a single peak. Figure 10 gives an example of this for US exports to France while Figure 13 plots the overall distribution of the estimates. Second, when I condition on price changes only, I find similar results. Now, estimates range between 0.25% and 3%, with both a mean and median estimate of 1.5%. Likelihoods are again well-behaved. Figure 14 shows the distribution of the estimates. The estimates take a tighter range now because implied mean price changes at the upper bound of menu costs are fairly large. Third, when I condition on both the frequency and size of price changes jointly, export menu costs are estimated to be quite similar. They range between 0.5% and 5.5% of steady state revenues. The mean and median estimates are 1.85% and 1.75%. These numbers lie between the estimates based on frequencies and price changes only, as expected. Figure 15 plots the joint distribution of the estimates. Finally, when I compute a single estimate for the export menu costs, I obtain both a median and mean maximum likelihood estimate of 2%. All these export estimates contrast with the lower domestic estimates shown in Figure 11: when I use data on the size of price changes, the frequency of price changes, or both, domestic menu cost maximum likelihood estimates are 0.4%, 0.45% and 0.4%, conditional on adjustment.

6.3 Interpretation

My estimates both for domestic and export menu costs lie in the range that Zbaracki et al. (2004) find in a detailed case study of costs of mainly domestic price adjustment, reporting price adjustment costs of 1.22% of total annual revenues. If domestic prices change approximately every four to six months, my estimation implies comparable total estimated annual menu costs of 0.8% to 1.2%. My estimates are also similar to those estimated in Nakamura and Steinsson (2008b) based on a general equilibrium domestic market menu cost model. Interestingly, my larger estimates of export menu costs are in line with the break-down of costs of price adjustment reported in Zbaracki et al. (2004): First, according to this case study information gathering, decision-making and communication costs are key parts of price adjustment costs. These costs are six times larger than the costs of changing actual price labels. Second, cost of communication and negotiation with customers are up to 20 times larger than physical menu costs. This wider interpretation of menu costs is also consistent with survey evidence by Blinder et al. (1998) that literal menu costs do not represent substantial price adjustment costs.

A natural interpretation of larger estimates of export menu costs might thus be that not the literal menu costs but the managerial and negotiation costs with foreign customers are substantially larger. The case study of Zbaracki et al. (2004) provides some anecdotal evidence that this may indeed be the case for actual price adjustment costs: sales in one export market of the company studied are subject to a 16% exchange rate depreciation and cause substantial profit losses before the company changes the relevant export price. The reason cited for the delayed response is that "it was too costly to open the doors to negotiation" with buyers in the foreign market. I test the idea of larger export managerial and negotiation costs more generally by relating my menu cost estimates to distance as well as common language with the U.S. in the export market. Figure 16 shows how more distant export markets are indeed associated with larger adjustment costs. In addition, I find that adjustment costs are larger when the language in the export market is not English and smaller if it is. Table 17 summarizes this comparison.

Thus, larger costs of adjusting export prices appear quite plausible given a wider interpretation of menu costs in terms of managerial and negotiation costs which may increase with distance or sales market barriers such as foreign languages.

7 Conclusion

In this paper, I present evidence how the pricing behavior of firms systematically and substantially differs across domestic and export markets in terms of frequency, timing and size of price changes. Using producer price micro data form the BLS, I contrast domestic and export pricing decisions for the same products and show that (i) domestic producer prices change approximately twice as often as export producer prices, (ii) the probability of synchronized price adjustment across markets is 21% for upwards adjustments and 14% for downwards adjustments, (iii) the size of export price changes is substantially larger than the size of domestic price changes and (iv) there are strong seasonality effects in the data including a new year-end synchronization effect. Second, I show that economic fundamentals such as inflation, productivity, demand, exchange rates and market structure can only partially explain adjustment decisions and cross-market synchronization. Third, I present a dynamic menu cost model of price-setting, and attribute the remaining unexplained part in adjustment decisions to differences in menu costs across countries. I calculate the implied export and domestic market menu costs from the data and estimate that export menu costs are 1.5% of period steady state revenues and three times as large as domestic market menu costs.

References

AOKI, K. (2001): "Optimal monetary policy responses to relative-price changes," *Journal of Monetary Economics*, 48, 55–80.

AUCREMANNE, L. AND E. DHYNE (2004): "How frequently do prices change? Evidence based on the micro data underlying the Belgian CPI," Research series 200404-1, National Bank of Belgium.

BARSKY, R. B. AND J. A. MIRON (1989): "The Seasonal Cycle and the Business Cycle," *Journal of Political Economy*, 97, 503–34.

BENIGNO, G. (2004a): "Real exchange rate persistence and monetary policy rules," *Journal of Monetary Economics*, 51, 473–502.

BENIGNO, P. (2004b): "Optimal monetary policy in a currency area," *Journal of International Economics*, 63, 293–320.

BERNARD, A. B., S. REDDING, AND P. K. SCHOTT (2006): "Multi-Product Firms and Trade Liberalization," CEP Discussion Papers dp0769, Centre for Economic Performance, LSE.

BILS, M. AND P. J. KLENOW (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985.

BLINDER, A. S., E. R. D. CANETTI, D. E. LEBOW, AND J. B. RUDD (1998): *Asking About Prices: A New Approach to Understanding Price Stickiness*, Russell Sage Foundation, New York.

BRODA, C. AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 121, 541–585.

CARVALHO, C. (2006): "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks," *The BE Journal of Macroeconomics (Frontiers)*, 2.

FITZGERALD, D. AND S. HALLER (2009): "Exchange Rates and Producer Prices: Evidence from Micro-Data," *Working paper, Stanford University.*

GAGNON, E. (2009): "Price Setting During Low and High Inflation: Evidence from Mexico," *Quarterly Journal of Economics*, 124, 1221–1263.

GOLDBERG, P. AND R. HELLERSTEIN (2009): "On Firms and Prices," Working paper, Princeton University and Federal Reserve Bank NY.

GOLDBERG, P., A. KHANDELWAL, N. PAVCNIK, AND P. TOPALOVA (2008): "Multi-product Firms and Product Turnover in the Developing World: Evidence from India," *Review of Economics and Statistics forthcoming.*

GOPINATH, G. AND O. ITSKHOKI (2010): "Frequency of Price Adjustment and Pass-through," *The Quarterly Journal of Economics*, 125.

GOPINATH, G. AND R. RIGOBON (2008): "Sticky Borders," *The Quarterly Journal of Economics*, 123, 531–575.

GOURIEROUX, C. AND A. MONFORT (1997): *Simulation-Based Econometric Methods*, Oxford University Press.

KLENOW, P. J. AND O. KRYVTSOV (2008): "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *The Quarterly Journal of Economics*, 123, 863–904.

MAYER, T., M. J. MELITZ, AND G. I. OTTAVIANO (2009): "Market Size, Competition, and the Product Mix of Exporters," Working paper, Paris School of Economics, Princeton University and Bocconi University.

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

NAKAMURA, E. AND J. STEINSSON (2008a): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

——— (2008b): "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model," *Working paper.*

——— (2009): "Lost in Transit: Product Replacement Bias and Pricing to Market," NBER Working Papers 15359, National Bureau of Economic Research, Inc.

NEIMAN, B. (2009): "Intrafirm Trades and International Macro Dynamics," Working paper, University of Chicago, Graduate School of Business.

OLIVEI, G. AND S. TENREYRO (2007): "The Timing of Monetary Policy Shocks," *American Economic Review*, 97, 636–663.

SCHOENLE, R. (2009): "Evidence on Export Price Dynamics," Work in progress, Princeton.

STIGLER, G. J. AND J. K. KINDAHL (1970): *The Behavior of Industrial Prices*, Columbia University Press, New York, N.Y.

VERMEULEN, P., E. GAUTIER, H. STAHL, M. DOSSCHE, R. SABBATINI, D. DIAS, AND I. HERNANDO (2007): "Price setting in the euro area: some stylised facts from individual producer price data," Discussion Paper Series 1: Economic Studies 2007, 03, Deutsche Bundesbank, Research Centre.

ZBARACKI, M. J., M. RITSON, D. LEVY, S. DUTTA, AND M. BERGEN (2004): "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *The Review of Economics and Statistics*, 86, 514–533.

8 Tables

8.1 Data Summary Statistics

Table 1: Characteristics of Products in Matched Data

	PPI	matched PPI	IPP	matched IPP
Products	675	354	371	354
Countries	1	1	161	156
Mean Firm Employment	2107	605	-	-
Median Firm Employment	145	139	-	-
Firms	28575	13285	3322	2694
Median # Items per Firm	78.06	82.05	7.71	7.85
Mean # Items per Firm	110.11	99.66	13.25	13.33
Items	121984	56603	15997	12043
Mean Value of Sales	$5 * 10^9$	$5 * 10^9$	$1.55 * 10^9$	$1.53 * 10^9$
Median Value of Sales	$2 * 10^{10}$	$1 * 10^{10}$	$6.24 * 10^7$	$5.04 * 10^8$

Characteristics of the BLS PPI and IPP data are for the sample period of 1998-2005. Products are defined at the six-digit NAICS level. Employment represents the number of employees in a firm at re-sampling. Firm sales values are based on reported values of shipment.

Table 2: Characteristics of Matched Firms

	PPI	PPI matched	IPP	IPP matched
Countries	1	1	161	91
Mean firm employment	2107	2062.211	2107	2062.211
Median firm size	145	129	145	129
Firms	28575	457	3322	381
Median # Items per Firm	78.06	4.67	7.71	3.34
Mean # Item per Firm	110.11	6.42	13.25	5.37
Item	121984	1950	15997	1594

Characteristics are based on firms matched using a name-based fuzzy matching algorithm described in the appendix. Data are for the period of 1998-2005. Employment represents the number of employees in a firm at re-sampling. Firm sales values are based on reported values of shipment.

8.2 Domestic and Export Pricing Dynamics

Table 3: Frequency and Size of Price Changes, The Economist

Sales Market	Frequency	Duration	Size of Price Changes			Spells
			Mean	Median	Std. Error	
Australia	4.40%	22.73	5.84%	5.62%	0.56%	5
Canada	3.11%	32.14	9.74%	10.00%	2.12%	4
Chile	5.28%	18.94	21.46%	13.33%	9.69%	6
France	7.00%	14.30	4.30%	4.26%	0.16%	9
Germany	3.89%	25.71	7.13%	7.42%	0.65%	5
Guyana	5.79%	17.27	12.95%	9.10%	6.15%	4
Hong Kong	4.40%	22.73	13.67%	14.84%	1.84%	5
India	5.28%	18.94	16.47%	14.29%	4.61%	6
Israel	2.84%	35.23	14.71%	14.71%		2
Italy	5.45%	18.37	6.46%	6.91%	0.84%	7
Netherlands	4.67%	21.43	3.97%	3.45%	0.57%	6
Panama	2.90%	34.53	4.55%	4.55%		2
Peru	5.79%	17.27	11.56%	9.15%	5.66%	4
Saudi Arabia	2.64%	37.88	21.52%	21.52%	7.89%	3
South Africa	7.04%	14.20	11.38%	10.06%	0.96%	8
Switzerland	3.52%	28.41	5.83%	7.14%	1.59%	4
US	2.33%	42.86	14.76%	14.76%	1.91%	3
UK	11.67%	8.57	4.92%	4.72%	0.75%	15
All exports	4.49%	22.22	10.96%	10.06%	1.38%	83
All exports†	4.40%	22.73	11.38%	9.15%	0.84%	83

The table is based on weekly cover prices of The Economist from 1990-2001. I compute frequency as the fraction of price changes. Duration reports the median duration of price spells. The size of price changes is based on log differences of price changes conditional on adjustment. Spells denotes the number of observed price spells for each destination. All exports denotes results when pooling the data of all export destinations and reporting means across countries. All exports† denotes results from pooling and reporting medians across countries.

Table 4: Probabilities of Synchronization

	Weekly	Monthly	Quarterly
Prob($\Delta p_c^X \neq 0 \mid \Delta p^{UK} \neq 0$)	0.13	0.29	0.54
	(0.091)	(0.125)	(0.144)
Prob($\Delta p_c^X \neq 0 \mid \Delta p_{-c}^X \neq 0$)	0.38	0.51	0.54
	(0.076)	(0.083)	(0.096)
Prob($\Delta p^{UK} \neq 0 \mid \Delta p_c^X \neq 0$)	0.07	0.15	0.32
	(0.048)	(0.070)	(0.102)

The table is based on weekly cover prices of The Economist from 1990-2001. I compute the three synchronization probabilities in the following way: (i) as the probability of at least one export price conditional on a price change in the U.K. (ii) as the probability of at least two export prices change conditional on any one export price changes and (iii) as the probability of a price change in the U.K. conditional on any change among export prices. I compute these probabilities at a weekly, monthly or quarterly horizon.

Table 5: Frequency and Implied Duration, Domestic and Export Sales

		Domestic	Export I	Export II
Unweighted				
Frequency	Mean	20.9%	11.5%	11.6%
	Median	16.3%	5.5%	6.0%
	Std. Error	0.9%	0.9%	0.3%
Implied Duration	Mean	6.64	20.53	20.49
	Median	5.60	17.84	16.23
	Std. Error	0.25	0.83	0.32
Weighted				
Frequency	Mean	28.9%	13.3%	12.7%
	Median	18.5%	6.4%	6.9%
	Std. Error	1.3%	1%	0.3%
Implied Duration	Mean	5.55	18.16	18.52
	Median	4.88	15.12	13.99
	Std. Error	0.24	0.79	0.34
	N	354	354	2731

The table is based on domestic PPI and export IPP micro data from the BLS. I compute the frequency of price change for each item in the data, then aggregate up by taking the unweighted median frequency within the next higher category and unweighted medians across categories within the six-digit products. Implied durations are given by $d = -\frac{1}{\ln(1-f)}$ where f is the relevant frequency of price change. I report means and medians for six-digit products matched across the PPI and IPP. Weights are combined product domestic and export sales values.

Table 6: Direct Comparison of Frequency and Duration

			Mean	Median	Std. Error	N
Unweighted						
	Exports I	r^f	0.5448	0.3796	0.0327	354
		Δd	14.01	10.89	0.81	354
	Exports II	r^f	0.6515	0.4066	0.0169	2731
		Δd	13.77	9.28	0.32	2731
Weighted						
	Exports I	r^f	0.5074	0.3560	0.0302	354
		Δd	12.67	8.10	0.74	354
	Exports II	r^f	0.5657	0.3315	0.0141	2731
		Δd	13.32	9.18	0.30	2731

The table is based on domestic PPI and export IPP micro data from the BLS. First, I compute the frequency of price change for each item in the data, then aggregate up by taking the unweighted median frequency within the next higher category and unweighted medians across categories within the six-digit products. I denote the ratio of export to domestic frequencies of price changes for each product by $r^f = \frac{f^F}{f^H}$, calculated by pooling all export destinations, Exports I, or for each product-destination pair separately, Exports II. Similarly, I calculate the difference in implied durations as $\Delta d = d^F - d^H$. Weights are combined product domestic and export sales values.

Table 7: Alternative Frequency Calculations

			Mean	Median	Std. Error	N
Exports I	Method I	r^f	0.5448	0.3796	0.0327	354
		Δd	14.01	10.89	0.81	354
	Method II	r^f	0.7277	0.4721	0.0430	354
		Δd	12.66	7.39	1.16	354
	Method III	r^f	0.7798	0.5915	0.0305	354
		Δd	5.89	4.56	0.71	354
Exports II	Method I	r^f	0.5673	0.4102	0.0284	2731
		Δd	11.20	9.33	0.58	2731
	Method II	r^f	0.6456	0.4155	0.0378	2715
		Δd	9.74	6.77	0.74	2715
	Method III	r^f	0.7359	0.5576	0.0292	2715
		Δd	5.06	4.95	0.36	2715

As in Table 6, I calculate ratios r^f and differences Δd of export and domestic frequencies of price adjustment and implied durations. Method I denotes item-level calculations of the frequency of price change as the fraction of an indicator variable for price adjustment. Method II and III present calculations robust to censoring issues in the IPP: Method II is based on the method of Aucremanne and Dhyne (2004) and Method III estimates hazard rates similar to Neiman (2009) and Gopinath and Rigobon (2008).

Table 8: Export-Domestic Synchronization Probabilities

	$\text{Prob}_{Up\|Up}$	$\text{Prob}_{Down\|Down}$
Product	20.81%	14.14%
	(2.16)	(2.17)
Product-Destination	20.37%	13.76%
	(2.19)	(2.24)
Benchmark	47.48%	50.31%
	(2.93)	(2.42))
N	768	67

The table shows the probabilities of export price adjustment conditional on adjustment of all goods in the corresponding domestic category. First, I compute the signed fractions of domestic and export price adjustment at each point in time and for each product category. Second, I condition on full adjustment in the domestic market. Third, I compute the signed mean conditional export probability for each product, then report the average across products. Similarly, I compute export fractions for each product-destination and aggregate up in the same way. The benchmark reports the fraction of other domestic products changing in the same direction within a four-digit category given full adjustment of a particular domestic product under consideration. Standard errors are given in brackets.

Table 9: Strength of Synchronization from Sychronization Regressions

	β_1^{D+}	β_1^{X+}	β_1^{D-}	β_1^{X-}	$\Delta\beta^+$	$\Delta\beta^-$
Product	0.71337	0.10137	0.70622	0.05407	-0.61200	-0.65215
	(0.01493)	(0.01609)	(0.01602)	(0.01417)	(0.02295)	(0.02333)
Product-Destination	0.71337	0.15174	0.70622	0.04437	-0.62312	-0.71052
	(0.01493)	(0.04869)	(0.01602)	(0.02192)	(0.02356)	(0.05003)

I report the strength of domestic-export synchronization β_1^{D+}, β_1^{D-}, β_1^{X+} and β_1^{X-} based on estimates of two sets of regression equations. In the case of synchronization of upwards price changes, I estimate $I(\Delta p_{i,j,t}^X \geq 0) = \beta_0 + \beta_1^{X+} f_{i,t}^{D+} + \epsilon_{i,j,t}$ and $I(\Delta p_{i,j,t}^D \geq 0) = \beta_0 + \beta_1^{D+} f_{i,t}^{D+} + \epsilon_{i,j,t}$ where $f_{i,t}^{D+}$ is the fraction of domestic goods j adjusting upwards in product category i at time t (excluding the item under consideration in the second equation). $I(\Delta p_{i,j,t}^k \geq 0)$ is an indicator variable for upwards adjustment of good j in product category i and market $k \in (D, X)$. Similar equations hold for downwards adjustment. *Product* denotes estimation irrespective of destination while *Product-Destination* denotes estimates computed first for each product-destination and then aggregated for each product. I multiply estimated coefficients by 100 and report standard errors in brackets.

Table 10: Direct Comparison of Size of Price Changes

			Mean	Median	Std. Error	N
Export I	Unweighted					
		r^f	1.81	1.39	0.09	342
		Δdp	3.00%	1.75%	0.31%	342
	Weighted					
		r^f	1.57	1.19	0.06	342
		Δdp	2.27%	1.54%	0.25%	342
Export II	Unweighted					
		r^f	2.33	1.30	0.10	2086
		Δdp	5.87%	1.43%	0.50%	2086
	Weighted					
		r^f	1.98	1.17	0.08	2086
		Δdp	4.56%	0.57%	0.45%	2086

The table compares the absolute size of price changes of domestic and export sales of products based on PPI and IPP data. First, I compute the mean absolute size of price changes at the item level, conditional on adjustment. Second, I aggregate up by taking unweighted medians in the next higher category and across categories within a six-digit product category. Exports I pools all exports by product, and Exports II computes results for product-destinations. r^f denotes the ratio of export to domestic price changes, Δdp the difference of export and domestic price changes. Weights are based on the value of domestic and export shipments.

Table 11: Firm-Level Comparison

	Frequency		Duration		$\lvert\Delta p\rvert$	
	Domestic	Exports	Domestic	Exports	Domestic	Exports
Median	15.91%	6.27%	5.77	15.46	4.60%	9.57%
	(0.64%)	(0.06%)	(0.26)	(0.16)	(0.45%)	(0.66%)
Ratio	0.4840		0.4520		2.48	
	(0.0356)		(0.0354)		(0.43)	
Difference	6.75%		7.38		3.10%	
	(0.85%)		(0.80)		(0.38)%	
N	381	381	381	381	381	381

The table presents a comparison of price-setting behavior based on matched firms in the BLS PPI and IPP data. Matches are obtained using a fuzzy matching algorithm. To compare, I first compute the item-level frequency and mean absolute size of price changes, conditional on adjustment. Second, I take unweighted medians within each firm for domestic and export sales. Third, I compute ratios and differences of export and domestic statistics exactly as in Table 6 and report unweighted medians.

Table 12: Robustness based on Synthetic Data

		Method I	Method II	Method III
r^f	Full Data	0.3796	0.4721	0.5915
		(0.0327)	(0.043)	(0.0254)
	Synth I	0.3339	0.4509	0.5480
		(0.0229)	(0.0372)	(0.0211)
	Synth II	0.3503	0.4696	0.5195
		(0.0332)	(0.0388)	(0.0334)
Δd	Full Data	10.89	7.39	4.56
		(0.81)	(1.16)	(0.41)
	Synth I	11.26	7.29	4.44
		(1.13)	(0.85)	(0.38)
	Synth II	10.49	5.93	4.17
		(1.63)	(1.00)	(0.42)

The table summarizes the ratio of export to domestic frequencies r^f and the difference in implied durations Δd. *Method I*, *Method II* and *Method III* denote the three methods of calculating frequencies of price change described in the text, based on the average number of price changes, the method of Aucremanne and Dhyne (2004) and a hazard model. *Synth I* and *Synth II* denote the synthetic datasets. To generate them, I sort domestic goods in each product category by sales value. Then, I truncate them such that there is the same number of goods in corresponding domestic and export categories, or that I retain only the largest domestic sales item.

Table 13: Robustness at the Firm-Level

	Frequency		Duration		$\|\Delta p\|$	
	Domestic	Export	Domestic	Export	Domestic	Export
Median	15.15%	6.27%	6.09	15.46	4.89%	9.57%
	(0.54%)	(0.06%)	(0.23)	(0.15)	(0.35%)	(0.66%)
Ratio	0.48		0.47		2.33	
	(0.05)		(0.48)		(0.17)	
Difference	5.94%		7.26		2.34%	
	(0.99%)		(1.12)		(0.68%)	
N	381	381	381	381	381	381

The table presents a comparison of price-setting behavior based on matched firms in the BLS PPI and IPP data. Matches are obtained from a fuzzy matching algorithm. In this robustness check, I first retain only the domestic item with the largest domestic sales value within each firm. Second, I compute ratios and differences of statistics for all export items and the domestic item. Third, I report the unweighted medians across items. Standard errors are given in brackets.

Table 14: Regression on Simulated Data

	Upwards Adjustment	Downwards Adjustment
Inflation	41.31	-73.78
ΔProductivity	-22.62	74.22
ΔDemand	223.51	308.63

Based on simulated time series data from the model, I estimate a multinomial logit model to relate adjustment decisions to inflation, changes in productivity and demand. Decisions are to adjust upwards or downwards with no adjustment as the base case. I report mean coefficients over many simulations.

Table 15: Test of Structural Synchronization Equation

	I	II	III	I	II	III
$f^H_{p,t}$	0.2046	0.1935	0.1725	0.4226	0.2808	0.2622
	(0.0115)	(0.0083)	(0.0077)	(0.0124)	(0.0091)	(0.5500)
$f^H_{p,t} * \sigma_A$	**0.0077**			**0.0019**		
	(0.0018)			(0.0018)		
σ_A	0.0039			0.0042		
	(0.0006)			(0.0006)		
$f^H_{p,t} * \rho^{\pi}_{H,F}$		**0.1044**			**0.1203**	
		(0.0228)			(0.0226)	
$\rho^{\pi}_{H,F}$		0.0373			0.0355	
		(0.0089)			(0.0089)	
$f^H_{p,t} * \rho^{g^Y}_{H,F}$			**0.5414**			**0.5500**
			(0.0554)			(0.0549)
$\rho^{g^Y}_{H,F}$			0.0088			0.0021
			(0.0206)			(0.0204)
F.E.	NO	NO	NO	YES	YES	YES
Intercept	0.0574	0.0855	0.0964	0.0360	0.0728	0.0817
	(0.0035)	(0.0033)	(0.0029)	(0.0050)	(0.0063)	(0.0062)
N	75446	38426	36182	75446	38426	36182
R^2	0.0425	0.0398	0.0425	0.0714	0.0571	0.0597

The table presents estimates for the strength of synchronization $\hat{\beta}_3$ obtained from estimating of $f^F_{p,c,t} = \beta_0 + \beta_1 f^H_{p,t} + \beta_2 M_{p,c,t} + \beta_3 f^H_{p,t} * M_{p,c,t} + \epsilon_{p,c,t}$ where $f^F_{p,c,t}$ and $f^H_{p,t}$ denote the monthly fraction of price changes in export destination c and the domestic market for each product p at time t, and $M_{p,c,t}$ denotes one of the following: (I) the standard deviation of productivity growth rates, σ_A, (II) the correlation between US and export destination monthly aggregate output growth, $\rho^y_{H,F,c}$ or (III) the correlation of monthly CPI inflation rates, $\rho^{\pi}_{H,F,c}$.

Table 16: Menu Cost Maximum Likelihood Estimates

	$K^*_{c,1}$	$K^*_{c,2}$	$K^*_{c,3}$	$K^*_c = K^*$
Mean	1.97%	1.50%	1.85%	2%
Median	1.5%	1.5%	1.75%	2%
Std. Error	0.34%	0.16%	0.34%	-
Maximum	6.00%	3.00%	5.50%	
Minimum	0.25%	0.25%	0.50%	
N	23	23	17	17

The columns contain the maximum likelihood estimates of menu costs based on different moments from the data. Estimates are based on the frequency of price changes ($K^*_{c,1}$), the absolute size of price changes ($K^*_{c,2}$) and the joint moments ($K^*_{c,3}$) for the following countries: Belgium, Brazil, Canada, Switzerland, Chile, Germany, Denmark, Spain, France, Great Britain, Greece, India, Ireland, Israel, Italy, Japan, Korea, Mexico, the Netherlands, Russia, Sweden, Slovenia and Turkey. I also impose a uniform export menu cost across countries ($K^*_c = K^*$).

Table 17: Menu Costs and Common Language

	$K^*_{c,3}$	$K^*_{c,2}$	$K^*_{c,1}$
Common Language			
No	2.02%	1.60%	2.15%
	(0.47%)	(0.20%)	(0.44%)
Yes	1.875%	1.25%	1.67%
	(0.125%)	(0.52%)	(0.44%)

Based on estimates from the calibration exercise, the table summarizes menu costs according to common language with the U.S.. The columns contain results for the three sets of moments: (i) both frequency and absolute size (ii) absolute size and (iii) frequency of price changes. The countries with common language are Ireland, Great Britain and Israel, as obtained from the CEPII classification.

9 Figures

9.1 Results from The Economist

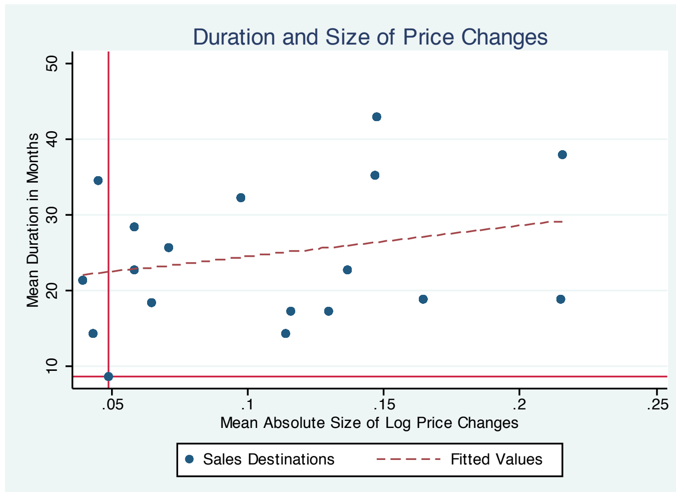

Figure 1: Duration and Size of Price Changes, The Economist

Each dot summarizes the duration and size of price changes for a sales destination of The Economist. The absolute size of price changes is calculated conditional on adjustment, the mean duration is the average of price spells. The graph is based on weekly cover prices of The Economist from 1990 through 2001 for the following sales destinations: Australia, Canada, Chile, France, Germany, Guyana, Hong Kong, India, Israel, Italy, Netherlands, Panama, Peru, Saudi Arabia, South Africa, Switzerland, U.S. and U.K.. The red coordinate lines go through the datapoint for the U.K..

9.2 Results from the BLS

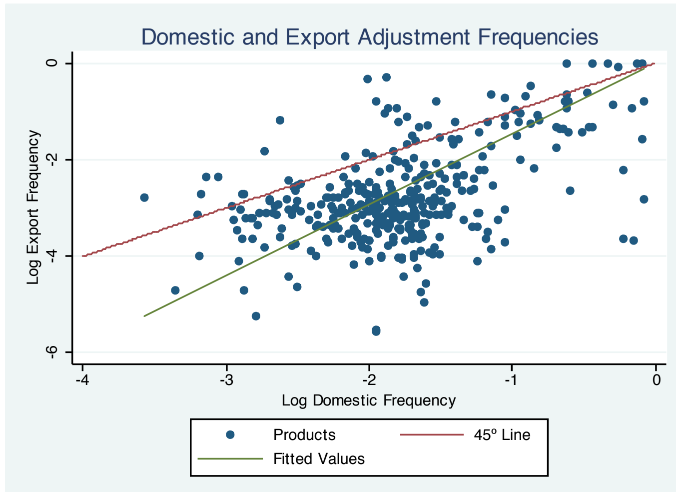

Figure 2: Domestic and Export Frequencies of Price Change

For each product, represented by a dot, I calculate the frequency of price change using the BLS PPI and IPP micro data. I first compute the frequency at the item-level and then take unweighted medians at the next higher classification level and within a six-digit product. The green line shows a simple linear fit of log export on log domestic frequencies.

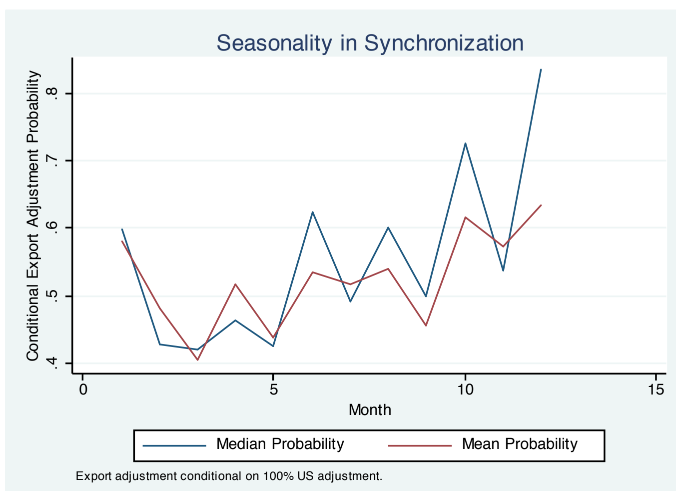

Figure 3: Seasonality in Synchronization

I compute synchronization probabilities by first computing the fraction of price changes in each domestic product-month-year and each export product-destination-month-year. Second, conditional on full adjustment in the corresponding domestic category, I take unweighted medians and medians over years and then destinations to arrive at a monthly conditional export adjustment probability.

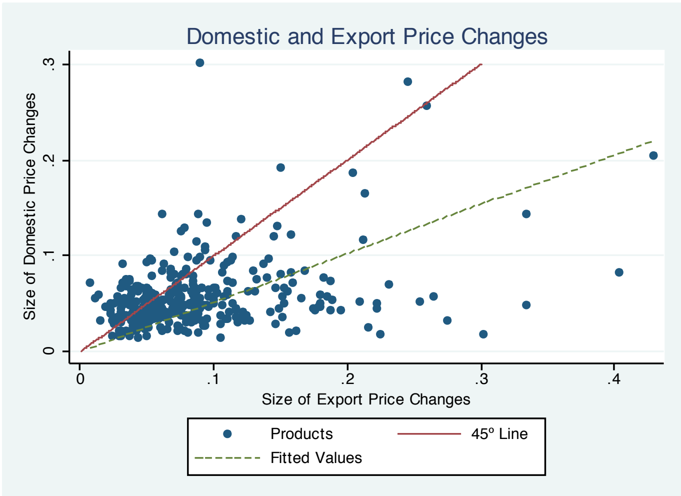

Figure 4: Size of Domestic and Export Price Changes, by Products

I first compute the mean absolute size of price changes at the item-level, then aggregate up to the product level as in the case of frequencies.

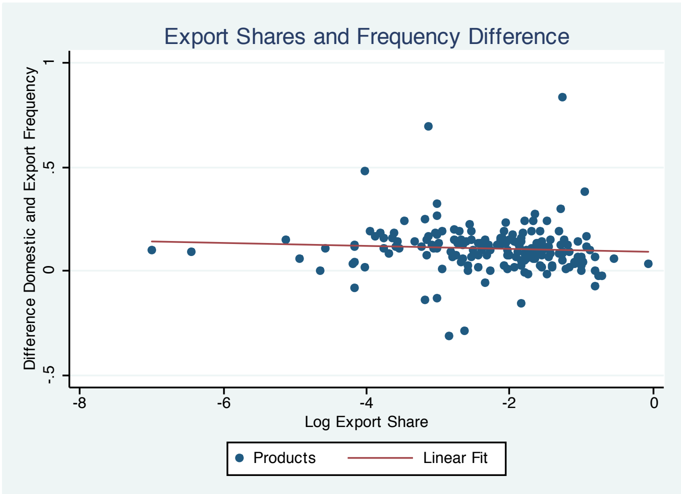

Figure 5: Frequency Differential and Export Share

I compute U.S. export shares using the Feenstra export data and output data from the NBER Productivity Database which I match at the four-digit SIC level. I use shares calculated for 1998. The difference in frequencies is based on unweighted median domestic and export frequencies.

9.3 Model Estimation

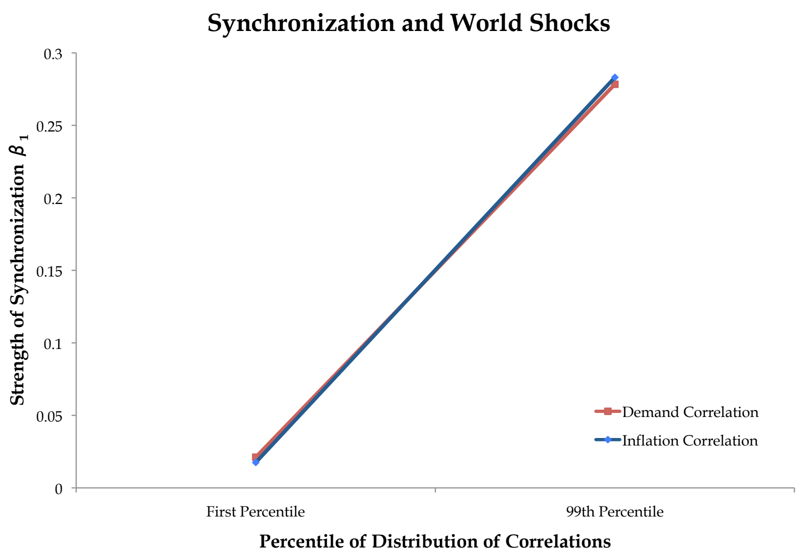

Figure 6: Estimated Strength of Synchronization and Correlation of Shocks

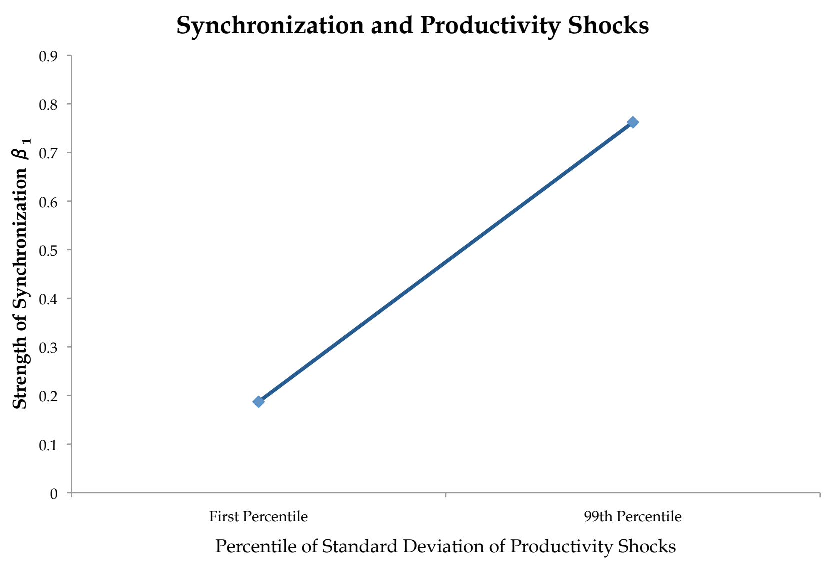

Figure 7: Estimated Strength of Synchronization and Productivity Shocks

9.4 Maximum Likelihood and Calibration Exercise

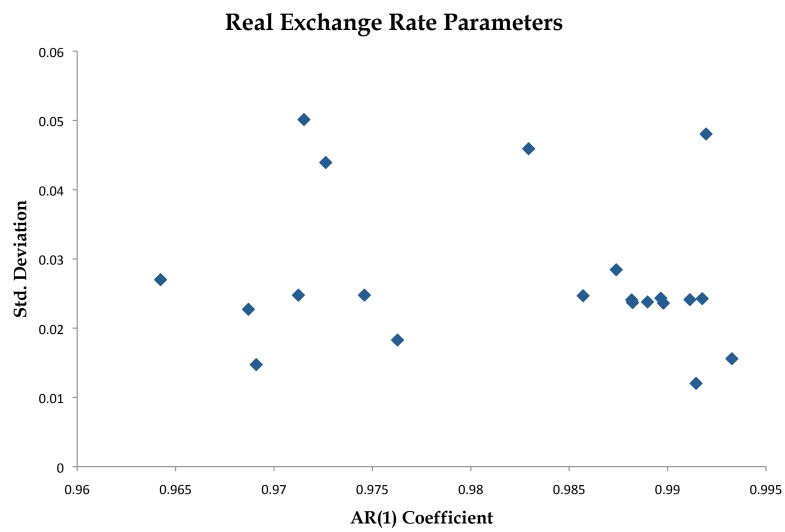

Figure 8: Real Exchange Rate Parameters

The graph shows the values of the real exchange rate parameters used in simulation, estimated from a monthly AR(1) of CPI-based log real exchange rates for each country: Belgium, Brazil, Canada, Switzerland, Chile, Germany, Denmark, Spain, France, Great Britain, Greece, India, Ireland, Israel, Italy, Japan, Korea, Mexico, the Netherlands, Russia, Sweden, Slovenia and Turkey.

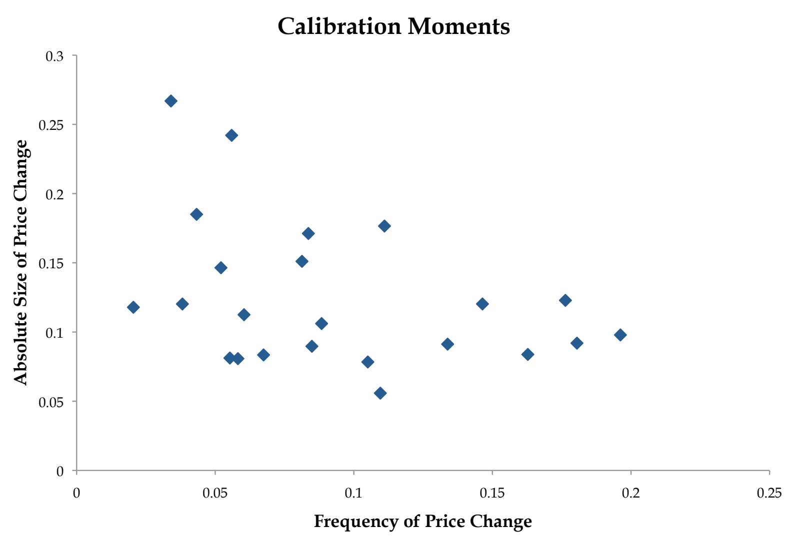

Figure 9: Calibration Moments

The moments of calibration are the frequency and absolute size of price change of the median product. Each dot represents one of the countries in the calibration exercise.

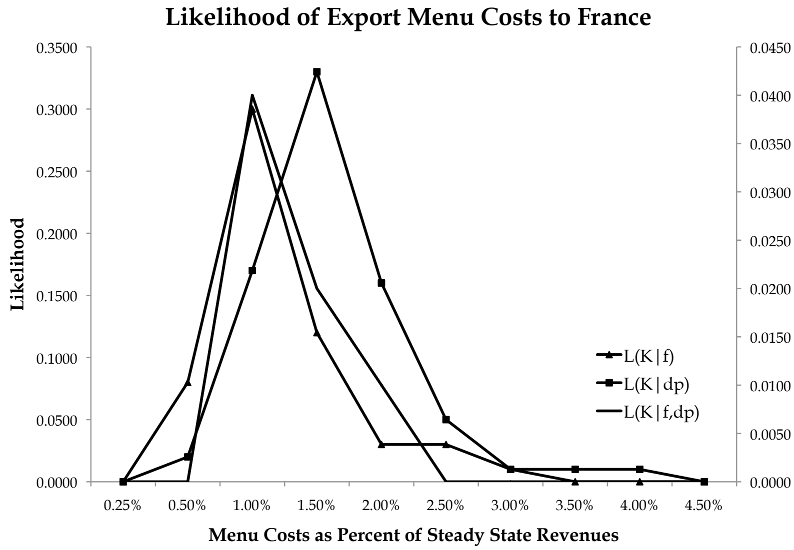

Figure 10: Likelihood of France Export Menu Costs

For the case of France, the graph shows the estimated likelihood of menu costs defined in terms of percent of steady-state revenues. The likelihood is estimated using stochastic simulation and conditional on frequencies, mean absolute size price change, or both.

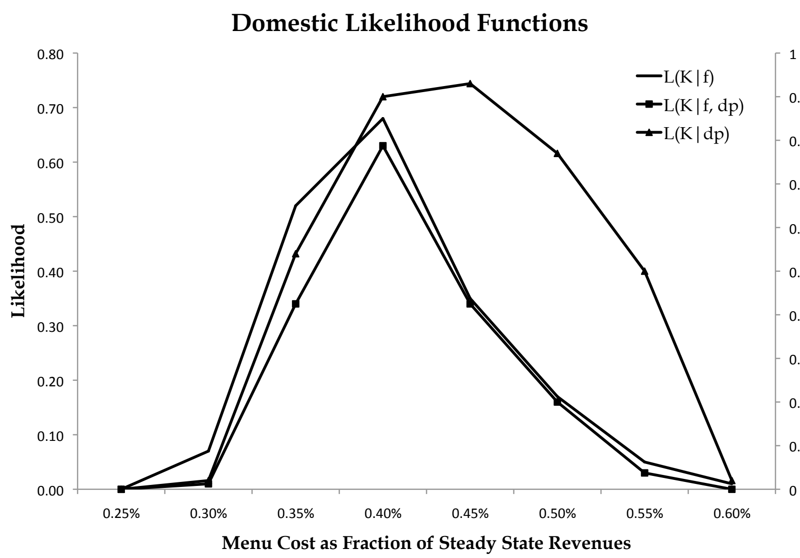

Figure 11: Likelihood of Domestic Menu Costs

The graph shows the estimated likelihood of menu costs defined in terms of percent of steady-state revenues. The likelihood is estimated using stochastic simulation and conditional on frequencies, mean absolute size price change, or both.

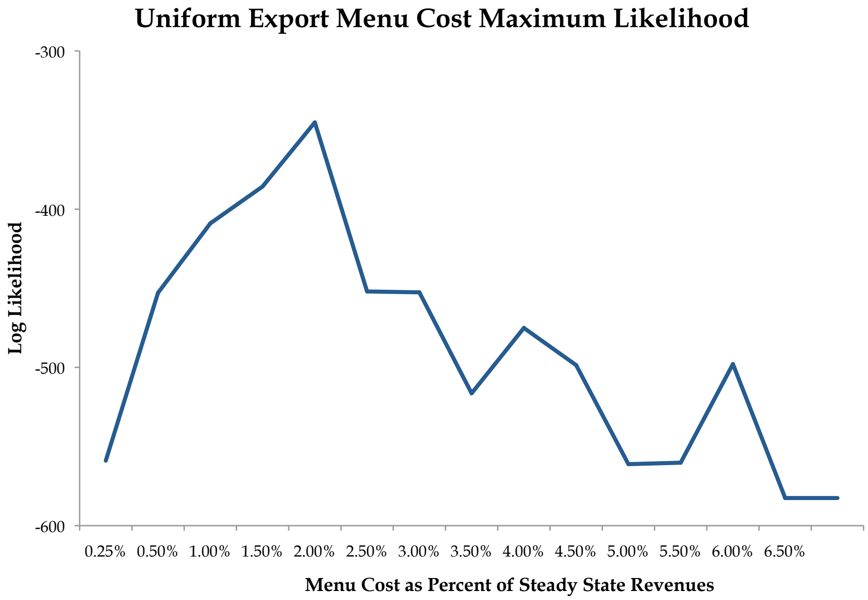

Figure 12: Uniform Export Menu Cost Maximum Likelihood

The graph shows the estimated likelihood of uniform menu export costs $K_c^* = K^*$. Estimates based on stochastic simulation.

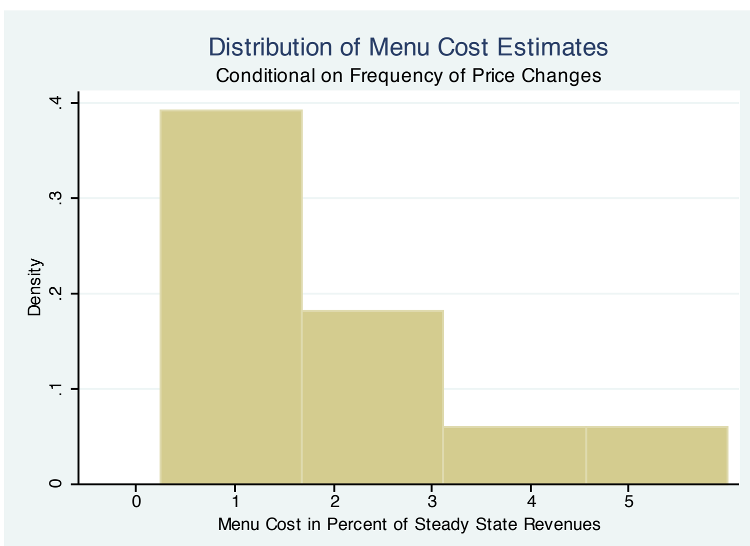

Figure 13: Distribution of Maximum Likelihood Estimates

The graph shows the distribution of maximum likelihood estimates of menu costs, estimated conditional on the frequency of price changes.

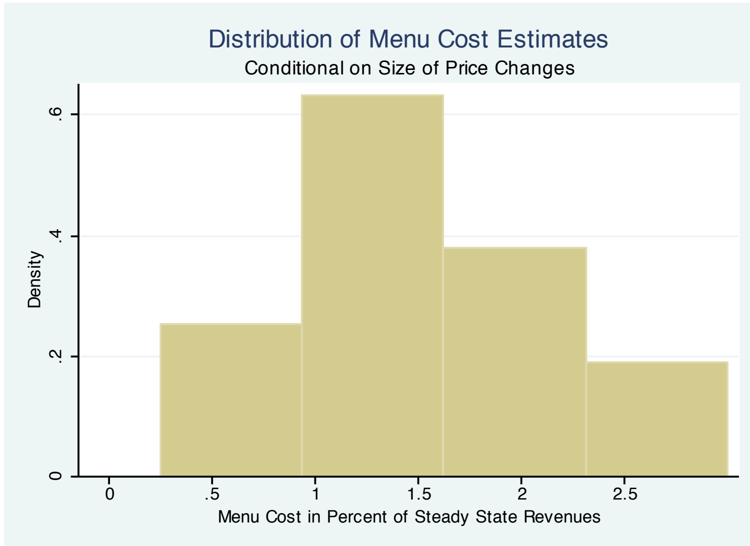

Figure 14: Distribution of Maximum Likelihood Estimates

The graph shows the distribution of maximum likelihood estimates of menu costs, estimated conditional on the mean absolute size of price changes.

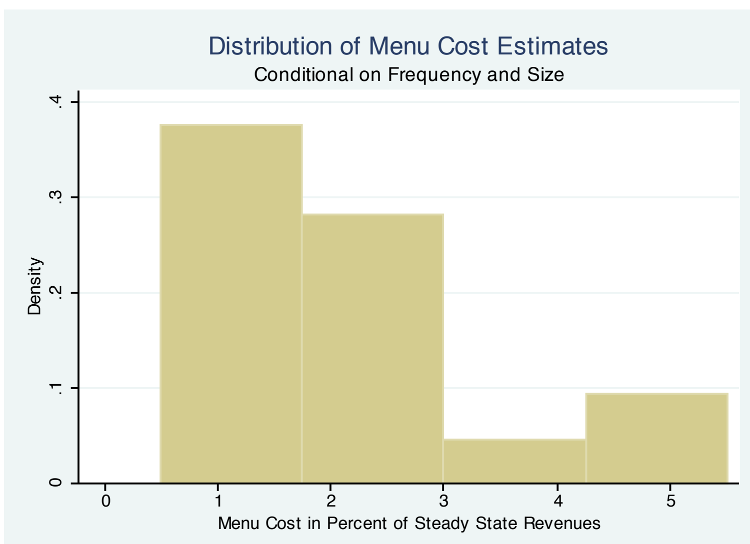

Figure 15: Distribution of Maximum Likelihood Estimates

The graph shows the distribution of maximum likelihood estimates of menu costs, estimated conditional on the frequency and mean absolute size of price changes.

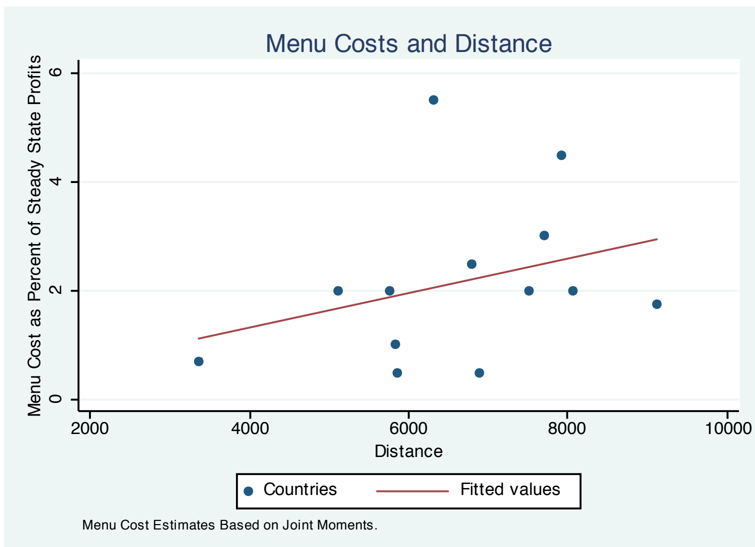

Figure 16: Menu Costs and Distance

The graph plots maximum likelihood estimates of menu costs against geodesic distance from the U.S.. Likelihood estimates are conditional on frequency and mean absolute size of price changes.

APPENDIX 1: Data Description

Domestic Producer Price Data

Here, I describe in detail the BLS PPI research database, how I manipulate the data and several robustness checks. First, the PPI contains a large number of monthly price quotes for individual "items", that is particular goods consistently defined over time. These items are selected to represent the entire set of goods produced in the US and are sampled according to a multi-stage design[34]. This sampling procedure takes three main steps: in a first step, the BLS compiles a sampling universe of all firms producing in the US using lists from the Unemployment Insurance System. Most firms are required to participate in this system and the BLS verifies and completes the sampling frame using additional publicly available lists, for example in the service sector. In a second step, "price-forming units" which are usually defined to be "production entities in a single location" are selected for the sample according to the total value of shipment of these units or according to their total employment. In a final series of steps called "disaggregation" a BLS agent conducts a field visit and selects the actual items to be selected into the sample. The goal of disaggregation is to select items that represent the most important product categories. Again, total values of shipment are used for selection.

In this last step, the BLS takes great care to obtain actual transaction prices. This emphasis on transaction prices goes back to a critique by Stigler and Kindahl (1970) when data was based on list and not transaction prices. In addition, the BLS also uniquely identifies an item according to its "price-determining" characteristics such as the type of buyer, the type of market transaction, the method of shipment, the size and units of shipment, the freight type and the day of the month of the transaction. Moreover, the BLS collects information on price discounts and special surcharges. Once an item has been sampled and uniquely identified according to its price-determining characteristics, the BLS collects monthly prices for that very same item and the same customer through a repricing form.

Despite this emphasis on transaction prices, there might be some concern about the quality of the price data: respondents have the option to report on the repricing form that a price has not changed. This might induce a bias in the price data towards higher price stickiness if respondents are lazy. Using the episode of the 2001 anthrax scare when the BLS exclusively collected prices by phone, Nakamura and Steinsson (2008a)[35] show that the frequency of price changes, controlling for inflation and seasonality, was the same in months when data were collected using the standard mail form as when the collection was done through personal phone calls.

Using these data, I choose my sample for analysis from the PPI data. First, I focus on prices for market transactions, eliminating all intra-firm trade prices. This removes 2.83% of the remaining data. Second, I choose the years from 1998 through 2005 as a time-frame. This has the advantage that consistent sampling methods, also for the IPP data, were applied during that time period. Third, I drop all time series where the buyer type is classified as "foreign buyer" during the entire time series. This removes 1.06% of the remaining prices. Fourth, I follow Neiman (2009) and Gopinath and Rigobon

[34]For a detailed description of the sampling procedures, see Chapter 14 of the BLS Handbook of Methods (US Department of Labor, 2008).

[35]See footnote 12 in Nakamura and Steinsson (2008a). This idea was first used in Gopinath and Rigobon (2008) where it is applied to export and import prices.

(2008) by dropping time series with fewer than six data-points. This affects only 0.37% of the data in the previous step. Fifth, I drop prices where the associated price changes are larger in magnitude than two log points. This applies to 0.006% of the data from the previous step. Finally, I concord four-digit SIC codes to six-digit NAICS codes.

Export Price Data

The IPP database similarly contains monthly, survey-based price data on individual exported items, but the IPP sampling frame is distinct from that of the PPI. The BLS compiles the IPP sampling frame using data on the total value and frequency of trade from the US Census Bureau or Customs Service and samples according to a three-stage sampling design[36]. In a first and second step, companies and then product groups are selected, both times proportionately to size. In the last sampling stage, individual items are chosen according to their shares of trade within the sampling unit.

Once an item is selected to be sampled for the dataset, item is "initiated" in the following way: when a field agent first contacts a firm to request a price, he has to obtain an actual transaction price for that item. Along with the price, the agent also records the item characteristics, similar to sampling for the PPI, such as item specification, the size and type of units shipped, the destination country and port of shipment, and details on the discount and tax structure, and whether these characteristics are price-determining or not. Then, repricing forms are used to gather prices of the same item each month as well as item characteristics if they are price-determining. The repricing forms ask for prices associated with transactions that occur on or as close to the 1st of each month as possible, and for transactions that possess the same or very similar characteristics.

If no transaction prices are available, the data contain only prices that meet certain quality criteria. While estimates given by reporting firms[37] are accepted as price records, list prices are not accepted after an item has been initiated unless actual discounts are applied and prices come from a verifiable transaction. Prices from "buy and sell agreements", "offset quantity arrangements", from the export of "used, rebuilt or remanufactured" items, "repaired or altered" items, leased equipment and military items are similarly not included in the sample in the first place. Under a strict set of six conditions, the BLS may accept average prices from repricing forms. The conditions for accepting average prices are that the item is a physically homogeneous good, that the buyer or seller is a homogeneous entity, that the destination country or region is the same for each of the transactions that are averaged, that the unit priced is the same for each of the transactions that are averaged, that the discount structure is constant and finally that averages or not computed over a period longer than a month.

Moreover, the BLS emphasis on consistently collectible price series means that the BLS asks for export prices that have a constant basis. Preferably, this price basis is "free alongside ship" (f.a.s.) or "free on board ship" (f.o.b.). The first type of price basis accounts for 19.13% of the data, while f.o.b. (port) account for 1.11%, f.o.b. factory for 33.78% and f.o.b. without further specification for 6.13%.

[36] For a detailed description of the IPP sampling procedures, see Chapter 15 of the BLS Handbook of Methods (US Department of Labor, 2008). The BLS additionally has a detailed description of the IPP data, the IPP Data Collection Manual, made available to the author.

[37] These are prices that firms would have charged if a transaction had occurred in the month when the re-pricing request was received.

28.77% of price bases are not further specified. In general, any price basis is deemed acceptable as long as it is consistently used for a series.

One potential concern for my comparison with the PPI data could be that the IPP data contain a lot of imputed data. While this should not affect my analysis of price changes since I focus on price changes conditional on adjustment, imputation may matter for my computation of frequencies of price changes. In particular, there are a lot of "pulled" prices in the data, about 19.82% of all prices. These pulled prices are prices that are carried over from last period, or "pulled" in BLS terminology. The existence of these pulled prices could imply that I may mistakenly detect a lot of artificial stickiness in the IPP data. Dropping imputed values which are flagged in the data as non-usable for calculation may have similar effects. I argue in the following that my final, resulting price data still represent data of good quality for my analysis.

First, evidence from the 2001 anthrax episode indicates that effort minimization is not a major driving factor of artificial stickiness in the data. That is, one might think that survey respondents try to minimize their reporting effort by indicating to the BLS agents that prices change less frequently than they actually do so that the BLS asks for re-pricing too infrequently and incorrectly "pulls" prices forward instead between re-pricings. However, when prices were collected through personal phone calls instead of using mail or fax re-pricing forms and thus raising the psychological barrier to responding with incorrect information, there was no substantial difference in response behavior and price durations as shown in Gopinath and Rigobon (2008). Consistent with such evidence, the reporting burden for individual respondents also seems fairly small: the median number of prices required from a price reporter is 3 and the mean number of items is 4.24. With particular regard to my comparison of PPI and IPP data, it also seems to be unclear why the reporting burden of given price should be much lower for domestic prices than for export prices.

Second, only a small fraction of imputed prices is due to the fact that items are not traded. In the data, prices are flagged as non-traded in the following way: "in all cases when a transaction for a particular item does not occur in a given repricing period, reporters mark the No Trade (NT) box and report an estimated price, if available, on the item's repricing form." I find that only 4.18% of all pulled prices and only 4.99% of all prices in the data are flagged as non-traded. Since the non-traded flag may under-report actual non-tradedness, I drop the entire price time series where at least ten percent of all prices are flagged as non-traded. This eliminates 13.31% of all prices. While I cannot determine if this procedure fully captures all non-traded data, the IPP sampling procedure that combines total value of shipments and frequency of trade to determine sampling probabilities may explain the relatively small percentage of non-traded data points to start with. As the IPP Data Collection Manual points out, "IPP generally collects data directly from companies which are active in international trade to ensure that the companies report actual transaction prices." Moreover, if non-traded prices are indicative of other hidden non-traded prices in the sense of a hidden Markov chain, then eliminating the entire time series when a certain threshold is reached may address the problem of non-traded prices.

Third, after I remove imputed and intra-firm prices and manipulate the data as described in this entire section, the existence of missing or unusable values is not likely to induce much artificial stickiness. Since missing values may be a problem if observed price changes are preceded by them[38], I compute

[38] In my methodology section, I explicitly describe how my methods deal with the missing values in the IPP data, beyond the data checking presented in the above.

the probability that price changes are preceded by missing values in my final dataset. I find that approximately 64% of all price changes are preceded by actual, good observations. Similarly, conditional on observing two identical subsequent prices, the probability that there is a missing value between them is only 7.8%. In fact, the median time between two price observations is exactly one month, thus equal to the sampling frequency, and the mean time between two observations is not much longer with 1.3 months. In addition, I find that price changes are much less frequently preceded by missing values for items that fall into the reference-priced and exchange-traded categories of Rauch (1997). While 56% of differentiated product price changes are not preceded by missing values, 73% of price changes of reference-priced items and 77% of price changes of exchange-traded items are not preceded by missing values. At the same time, the probability that there is a missing value between two unchanged price observations is 7.1% for differentiated goods but 8% and 12.7% for reference-priced and exchange-traded goods. This pattern is what one would expect if reference-priced or exchange-traded goods change prices more dynamically so they should endogenously and "correctly" be sampled more frequently when price changes are pervasive and important, but less frequently when prices do not change.

To create my final sample, I apply the following manipulations to the data. First, I eliminate all prices that are marked as non-usable by the BLS. Non-usable prices make up 29.83% of the remaining data. The largest fraction of these non-usable prices, 99.68%, are prices imputed by linear interpolation or imputed cell means, the rest consists of pulled prices. While the imputation does not bias IPP index movements, these prices are un-informative, non-transaction prices and I cannot use them for my analysis. Second, 0.25% of the remaining data are flagged as not usable for index calculation. I also exclude them from my sample. Third, I eliminate all prices that are estimated as is done in Gopinath and Rigobon (2008). This eliminates 4.56% of the remaining usable prices. Fourth, I concord the 10 digit Schedule B export codes to NAICS six-digit codes[39]. Fifth, I link old and new items into one series when items are replaced by closed substitutes according to the BLS[40]. Sixth, I eliminate prices associated with more than 2 absolute log price changes. This affects 0.04% of the remaining usable data. Lastly, I exclude price series with less than six usable prices, as in Neiman (2009) and Gopinath and Rigobon (2008).

Macro-Economic Data

I supplement the BLS micro price data with data on economic fundamentals such as output, productivity, market concentration ratios, inflation and exchange rates. I try to obtain these data at the monthly frequency whenever possible to match the monthly frequency of the BLS data. At the same time, I try to obtain these data at the most disaggregated level possible and then merge them by sector or destination country and date with the BLS data.

For the US, I obtain sector-level output series from the Federal Reserve Board "Industrial Production and Capacity Utilization" series. These monthly data are not seasonally adjusted and span the entire time period of the BLS data from January 1998 through December 2005. I compute monthly growth

[39] I use the concordance available from the US Census Bureau Foreign Trade Data at http://www.census.gov/foreign-trade/reference/codes/index.html

[40] The variable containing this information is replaced_item_id and I combine it with the "Out of scope (replaced)" discontinuation category, as in Gopinath and Rigobon (2008).

rates of output at the three-digit NAICS level and merge them with the BLS data at the three digit NAICS level. I add US inflation data in two ways: First, I compute monthly economy-wide average inflation using the PPI price data. Second, I merge in US CPI inflation from the OECD "Main Economic Indicators (MEI)", including both all-CPI-item inflation as well as CPI inflation excluding food and energy prices. To add a measure of market concentration, I draw on the eight-firm concentration ratios from the US 2002 Census of Manufacturing and merge them in at the six-digit level. For productivity data, I use the NBER Manufacturing Productivity Database. In particular, I use the annual four- and five-factor TFP annual growth rates available from 1958 through 1996. I compute the standard deviation of both growth rates for each of the 459 four-digit-SIC industries (1987 SIC) to obtain a long-run measure of the size of productivity shocks. Then, I concord the sector classifications to six-digit NAICS codes and merge the productivity measure into the BLS data.

To add foreign monthly output series, I use monthly aggregate production data from the OECD "Main Economic Indicators (MEI)." I compute the monthly growth rates and merge them by export destination and month with the BLS data. In order to have comparable monthly US aggregate data, I include the US in the previous step. In total, this yields an output measure for up to 37 destinations. I also add in total real GDP for 1998 as a measure of market size from the Penn World Tables 6.2 by multiplying real per capita GDP with total population for that year. In addition, I merge in CPI inflation from MEI, including both all-CPI-item inflation as well as CPI inflation excluding food and energy prices. This yields monthly inflation rates for 39 countries. I gather monthly exchange rate data from oanda.com. I compute the monthly average of the bid and the ask price and merge each currency series by country code into the BLS data. In the case of countries belonging to the Euro zone, I use their original currencies from 1998 until the introduction of the Euro and then continue the original series adjusting by the respective Euro conversion factors. Finally, since international concentration ratios are not easily available in comparable form to the US eight-firm and the six-digit level of disaggregation[41], I use US six-digit eight-firm concentration ratios, CR8, for all markets. I complete my data-set on economic fundamentals with data on geographic distance of each destination country from the US and common language with the US from the CEPII geography data-set.

APPENDIX 2: Model Solution and Simulation

Model Solution

In this appendix, I describe how I use numerical techniques to solve the maximization problem of the firm given by the following two equations:

$$V^a(S_t) = \max_{p_t}\{\pi(S_t) - K + \beta E[\max(V^a(S'_{t+1}), V^{na}(S'_{t+1}))]\} \tag{A-1}$$

$$V^{na}(S_t) = \pi(S_t) + \beta E[\max(V^a(S'_{t+1}), V^{na}(S'_{t+1}))]\} \tag{A-2}$$

[41]Eurostat, for example does not have available any firm concentration ratios or other measures of market concentration at all.

I take two steps to solve for the unknown value functions in the case of adjustment and non-adjustment, $V^a()$ and $V^{na}()$, with regards to the respective set of state variables S_t for the domestic and the export market. First, I approximate the value functions by projecting them into polynomial space as described for example in Miranda and Fackler (2000). Second, I solve for the coefficients of the polynomials that solve the non-linear system of equations given by the value functions.

In the first step, I approximate the value functions by projecting them into polynomial space. The domestic state variables are the current productivity shock A_t and last period's price in terms of the current price level, $\frac{p_{t-1}}{P_t}$ which matters only in the non-adjustment case. Therefore, $S_t^D = (A_t, \frac{p_{t-1}}{P_t})$. The state variables in the export case are the current productivity shock A_t, the current real exchange rate, Q_t and last period's producer-currency export price in terms of the current domestic price level, $\frac{p_{t-1}^X}{P_t}$ which again matters only in the export non-adjustment case. Therefore, $S_t^X = (A_t, \frac{p_{t-1}^X}{P_t}, Q_t)$. I solve for the two unknown value functions, $V^a()$ and $V^{na}()$ and the associated price policy functions, $p(S_t)$, by using projection methods to approximate the value functions and solving the system of non-linear equations given by (A-1) and (A-2). In particular, I approximate each value function by a set of higher order Chebychev polynomials and require (A-1) and (A-2) to hold exactly at a set of points given by the tensor product of a fixed set of collocation nodes of the state variables. This implies the following non-linear equation system, the so-called collocation equations, for each market:

$$\Phi^a c^a = v^a(c^a) \tag{A-3}$$

$$\Phi^{na} c^{na} = v^{na}(c^{na}) \tag{A-4}$$

where c^a and c^{na} are basis function coefficients in the adjustment and non-adjustment cases and Φ^a and Φ^{na} are the collocation matrices which are given by the value of the basis functions evaluated at the set of nodes. The right-hand side contains the collocation functions evaluated at the set of the collocation nodes. This is the same as the value of the right-hand side of the value functions evaluated at the collocation nodes but replacing the value functions by their approximations. In my solution, I use the same number of collocation nodes as for the order of polynomial approximation. Thus, I choose between 7-11 nodes for productivity state variable, 15-20 nodes for the real prices and 7-11 nodes for the real exchange rate. I pick the approximation range to be ± 2.5 times the standard deviation from the mean of the underlying processes. I solve for the right-hand side using Gaussian quadrature to calculate the expectations, with 11-15 points for the real price transitions due to inflation and calculate the expectations due to productivity and real exchange rate shocks exactly. In the adjustment case, I use a Nelder-Mead simplex method to find the maximum with an accuracy of the maximizer of 10^{-15}.

Second, I solve for the unknown basis function coefficients c^a and c^{na}. I express the collocation equations as two fixed-point problems

$$c^a = \Phi^{a-1} v^a(c^a) \tag{A-5}$$

$$c^{na} = \Phi^{na-1} v^{na}(c^{na}) \tag{A-6}$$

and update the coefficients iteratively until the collocation equations are satisfied exactly. I conduct two sensitivity analyses. First, given that zero menu costs imply flex-pricing, I verify that the approximate solution is "good" given the known analytical solution in this case. Figure 17 shows that optimal price policies and the price policies obtained by the approximation line up a 45-degree line. The norm of the error is of order 10^{-14} and errors are equi-oscillatory, as is a usual property of approximations based on Chebychev polynomials. Second, when I conduct the stochastic simulations in the maximum likelihood exercise, I find that the errors between the left- and right-hand sides of (A-1) and (A-2) at points other than the collocation nodes are on average of the order of 10^{-5}.

Stochastic Simulation

Using the policy the policy functions, I simulate time series to study the properties of the model. In the case of the synchronization equation, I simulate data and estimate the equation for two cases. In the first case, I hold the nominal exchange rate and consumption constant while the firm faces country-specific inflation shocks and a common productivity shock. I estimate equation (18) varying the correlation of domestic and export market inflation. I assume the following for parameter values: $\mu_P = 0$, $\theta = 4$, $\rho_A = 0.96$, $\sigma_A = 0.02$, $\sigma_P = 0.0037$, $\beta = 0.96^{(1/12)}$. In the second case, I vary the size of productivity shocks instead, given parameter values $\mu_P = 0.004$, $\theta = 4$, $\rho_A = 0.96$, $\sigma_P = 0.005$, $\beta = 0.96^{(1/12)}$. To estimate the multinomial model of price adjustment, I simulate time series assuming $\mu_P = 0.0021$, $\theta = 4$, $\rho_A = 0.96$, $\sigma_A = 0.02$, $\sigma_P = 0.0037$, $\rho_C = 0.43$, $\sigma_C = 0.02$, $\beta = 0.96^{(1/12)}$.

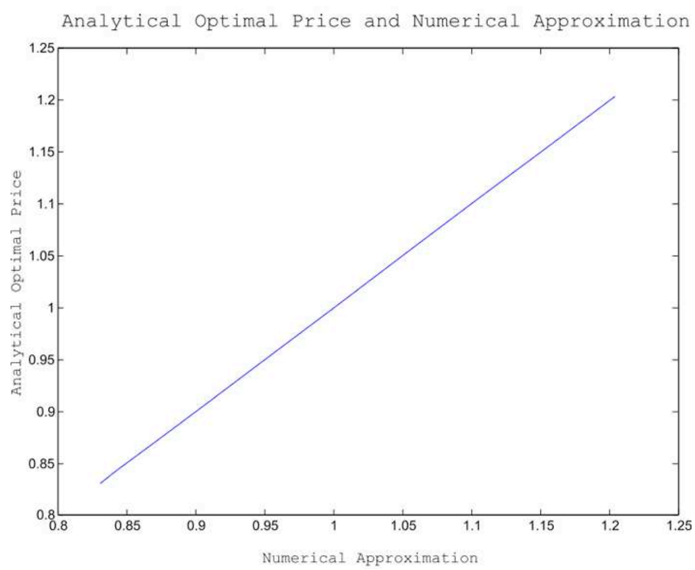

Figure 17: Analytical and Numerical Optimal Flex Price

I compute the numerical solution for optimal adjustment prices given productivity shocks and zero menu costs. I compare to the analytical solution known in the flex price case. Errors are of the order of 10^{-9}.

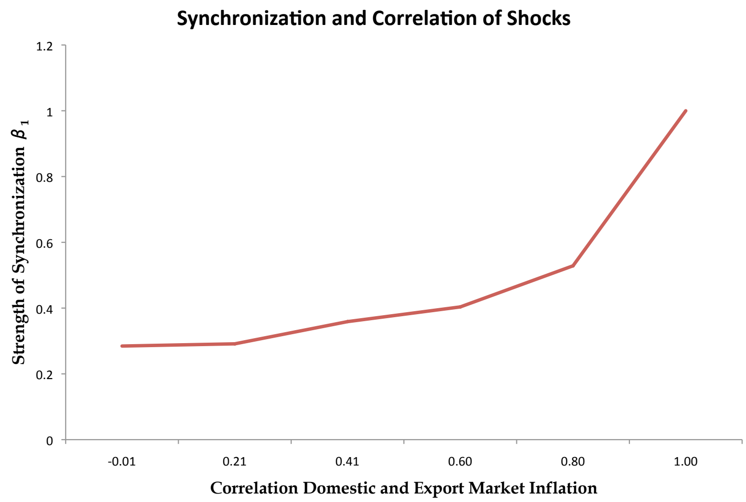

Figure 18: Predicted Strength of Synchronization and Correlation of Shocks

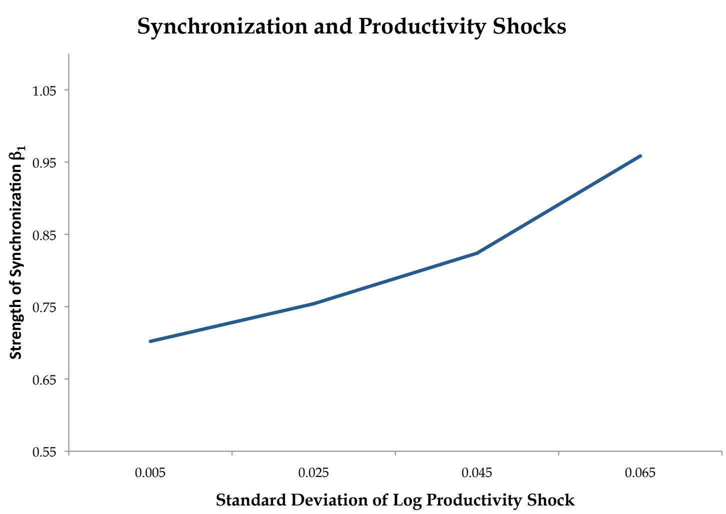

Figure 19: Predicted Strength of Synchronization and Productivity Shocks

APPENDIX 3: Fuzzy Matching Algorithm

Here I describe the fuzzy matching algorithm used to matched firms across the PPI and IPP. I match firms by their names, matching from the export data to the domestic data. Matching involves several steps.

First, I capitalize all firm names. Second, I remove the following punctuation marks from all text strings: ".", ",", ";", "/", "-", "(", ")", """', "[", "]", "{", "}", """. "/', "&" and "+". Third, I transform generic expressions in the data to a standardized, capitalized form. I summarize these standardizations in Table 18. Fourth, I drop all standardized plus additional generic terms, including: AG, AND, ASSN, ASSOC, BLDG, BROS, CO, CORP, DISTR, DIV, ENTERPRISE, ENTERPRISES, FDRY, FOOD, FOODS, GRP, INC, INC, IND, INTERAMERICAN, INTL, INTL, LP, LTD, MFG, MFG, N A, NA, OF, PRODUCTS, SERVICE, THE, USA. Fifth, I find all exact matches of firm names given these manipulations. I do not require firms to be in the same city or state. Sixth, I use a modified string similarity algorithm to get a measure of similarity between firm names which I did not match in the previous step.[42]

After I have obtained a list of matched firms and potential matches according to the string similarity algorithm, I manually verify and identify all matches between IPP and PPI which are "good" and contain no oddities unforeseen by the algorithm. This leaves me with approximately 381 matched firms which I can identify in the main PPI and IPP databases in my time period of interest. This is equivalent to a matching rate of 11.50% of exporting firms.

Table 18: Fuzzy Matching: String Manipulations

Original Term	Standardized Term
ASSOCIATES	ASSOC
ASSOCIATION	ASSN
BROTHERS	BROS
BUILDING	BLDG
COMPANY	CO
CORPORATION	CORP
DISTRIBUTION	DISTR
DIVISION	DIV
FOUNDRY	FDRY
GROUP	GRP
INCORPORATED	INC
INDUSTRIES	IND
INTER AMERICAN	INTERAMERICAN
INTERNATIONAL	INTL
INTL'L	INTL
LLC	INC
MANF	MFG
MANUFACTURING	MFG
NORTH AMERICA	NA
U S A	USA

[42] Bigram string comparisons compute similarity between strings based on the number of matching bigrams. For example, Dice's coefficient s is one such measure of similarity:

$$s = \frac{2C}{X_1 + X_2} \tag{A-7}$$

where C is the number of common bigrams, X_1 the number of bigrams in the first string and X_2 the number of bigrams in the second string. I use the Stata implementation given by reclink.

APPENDIX 4: Multinomial Logit Regressions

I estimate the multinomial model of price adjustment decisions by relating adjustment decisions to my measures of inflation, productivity, concentration ratios, demand changes, firm size, nominal exchange rate changes and market size. The base case is the decision not to change price. I report results from the estimated multinomial model in Tables 21 and 19 for the domestic case, and in Tables 22 and 20 for the export case. Tables 21 and 22 report relative risk ratios, where coefficients bigger (less) than 1 mean that an increase in the explanator makes adjustment more (less) likely. The other tables contain marginal effects associated with changes from the mean less one half standard deviation to the mean plus one half standard deviation. Results show that economic fundamentals are significantly related to the timing of pricing decisions both in the domestic and the export markets. However, a large fraction of adjustment decisions is left unexplained.

Adjustment in the Domestic Market

First, results show a statistically very strong, yet also asymmetric effect of inflation on adjustment: an increase of one standard deviation around the mean is associated a 3.75% higher probability of upwards adjustment, but only a 0.57% lower probability of downwards adjustment. The effect is much smaller in the full specification than the base case, dropping from an upwards 9.3% and a downwards -6.4% effect. This reduction is due to the inclusion of month fixed effects.

Second, I find an asymmetric effect of productivity: a one standard deviation increase in the productivity standard deviation around its mean is associated with a 0.02% increase in the upwards and an 0.44% increase in downwards adjustment probabilities. The effect is similar in the base case and the full specification. However, the coefficient for upwards adjustments is not significant.

Third, I find that a larger concentration ratio is statistically significantly and positively associated with both upwards and downwards price adjustment. A one standard deviation increase around the mean is associated with a 0.76% higher upwards and a 0.46% higher downwards probability of adjustment. This result is consistent across specifications.

Fourth, as shown in the full specification, I find that an increase in output is statistically significantly associated with higher upwards adjustment probabilities and lower downwards adjustment probabilities. The economic magnitude of these changes is small: a one standard deviation change is associated with a 0.1% higher upwards and a 0.15% lower downwards probability. The coefficient on upwards adjustment changes sign across specifications and both coefficients become smaller in magnitude when I add month fixed effects.

Finally, there is a statistically significant effect of firm size: a larger number of employees in a price-setting unit is associated with more frequent adjustment, both upwards and downwards. While the reported relative risk ratio is very small, the marginal effect of a change by one standard deviation around the mean is large with a 2.11% higher upwards and a 1.30% lower downwards probability. This difference in coefficients can be reconciled because the marginal effect calculation is driven by a few very large firms that pull up the mean and standard deviation.

The results show that adjustment decisions are highly significantly related to economic fundamentals and changes therein, both for upwards and downwards decisions. Finally, I note that month effects add approximately 5% of explanatory power when going to the full specification, increasing R^2 from 18% to 23%.

Adjustment in the Export Market

First, results show a statistically very strong, yet also asymmetric effect of producer price inflation on adjustment: an increase of one standard deviation around the mean is associated a 2.21% higher probability of upwards adjustment, and a 1.63% lower probability of downwards adjustment. The effect is much larger in the full specification than the base case, increasing from an upwards estimated effect of 0.43% and a downwards estimated effect of -0.21%. Similar to the domestic case, I find that this effect is due to the inclusion of month fixed effects. Overall, the effect of producer price inflation has the same sign as in the domestic case but larger coefficients.

Second, export market inflation is significantly associated with higher probabilities of downwards adjustment: a one standard deviation increase in foreign inflation around the mean is associated with a 0.20% increase in the downwards adjustment probability. In the case of upwards adjustment, the coefficient is positive but small with 0.02% and statistically insignificant.

Third, I find an asymmetric effect of productivity: a one standard deviation increase in the productivity standard deviation around its mean is associated with a -0.98% decrease in the upwards and a 0.17% increase in downwards adjustment probabilities. The effect is similar in the base case and the full specification. Compared to domestic pricing decisions, the effect of downwards adjustment is now negative and as the model predicts. At the same time, the coefficient for upwards adjustment is significant while for downwards adjustments it is no longer significant.

Fourth, I find that a larger concentration ratio is statistically significantly and positively associated with both upwards and downwards price adjustment, as in the case of domestic pricing decisions. A one standard deviation increase around the mean is associated with a 1.04% higher upwards and a 0.53% higher downwards probability of adjustment. This result is consistent across specifications and with the findings in the domestic case.

Fifth, results now show a positive sign on the output coefficient for both upwards and downwards adjustments as predicted by the model. A one standard deviation increase in output growth is associated with a 0.34% increase in the probability of upwards adjustment and a 0.01% increase in the probability of downwards adjustment. Compared to the domestic case, the sign in the case of downwards adjustment has changed from negative to positive. However, the coefficient for downwards adjustments has become insignificant. The difference between domestic and export results is due to the difference in the underlying data: while the domestic output series are at the six-digit level of disaggregation, I use monthly foreign data at the aggregate, economy-wide level. When I use U.S. monthly aggregate data instead in the domestic case, domestic price-setting is also related to output growth as predicted: an increase in demand is associated with both higher upwards and downwards probabilities of adjustment.

Sixth, I find that an appreciation of the US dollar versus the local currency is associated with a significantly higher probability of downwards adjustment, and a lower probability of upwards adjustment, but not significantly so. A one standard deviation in the nominal exchange rate is associated with a 0.20% higher probability of downwards and a 0.12% lower probability of upwards adjustment.

Finally, geographic distance and market size have significant, symmetric effects of similar magnitudes on both upwards and downwards adjustments. A one standard deviation increase in total real output is associated with a 0.67% higher probability of upwards adjustment and a 0.88% higher probability of downwards adjustment. This corresponds directly to the effects of a demand increase predicted by the model. At the same time, a one standard deviation increase in distance is associated with approximately 1% higher probability of upwards and downwards adjustment. The relative risk ratio for distance, however, is small and the economic magnitude is driven by Korea and Japan, the two countries most distant from the US market.

Last, as in the case of domestic adjustment decisions, the regressions only explain a fraction of the export adjustment decisions, approximately 11%.

Table 19: Marginal Effects, ± 1/2 Std. Dev., Domestic Case

	(I)		(III)	
	-	+	-	+
g_y	-0.569653	-0.101027	-0.146164	0.103048
CR8	1.16716	1.392618	0.464119	0.767348
σ_A	0.531632	0.055838	0.441243	0.016677
π	-6.368125	9.301845	-0.5674	3.754935
Employees			1.29497	2.10625

The table shows the marginal effects associated with one standard deviation changes in the key explanatory variables around the means, both for upwards (+) and downwards (-) domestic adjustment decisions in specifications (I) and (III) of Table 21.

Table 20: Marginal Effects by Bin, ± 1/2 Std. Dev., Export Case

	(I)		(III)	
	-	+	-	+
g_y	0.096779	0.802456	0.011092	0.336524
CR8	1.358724	2.159667	0.534014	1.034819
σ_A	0.40201	-0.877103	0.167836	-0.975988
π	0.211249	0.452945	-1.632464	2.206422
Δe	0.186492	-0.280011	0.196674	-0.123643
π^*	0.196885	0.089095	0.160605	0.014825
GDP	-	-	0.876845	0.665597

The table shows the marginal effects associated with one standard deviation changes in the key explanatory variables around the means, both for upwards (+) and downwards (-) export adjustment decisions in specifications (I) and (III) of Table 22.

Table 21: Multinomial Logit, Domestic Case

	(I)		(II)		(III)	
	-	+	-	+	-	+
g_y	0.081079	0.6413946	0.1703213	0.3480959	0.5212723	1.273738
	(0.0049905)	(0.0247052)	(0.0116487)	(0.0146979)	(0.0379921)	(0.0597399)
CR8	1.010895	1.008269	1.004004	1.003851	1.00472	1.004454
	(0.0001032)	(0.0000836)	(0.0001203)	(0.0000973)	(0.0001257)	(0.0001013)
σ_A	1.046834	1.000026	1.038715	0.9990898	1.041656	1.001631
	(0.0009763)	(0.0008307)	(0.0011182)	(0.000947)	(0.0011644)	(0.0009855)
π_{PPI}	0.1908589	4.89523	0.160182	5.248617	0.8995943	1.993551
	(0.0005901)	(0.0169714)	(0.0005317)	(0.0191415)	(0.0109895)	(0.0253917)
Employees			1.000057	1.000053	1.000063	1.000059
			(0.00000061)	(0.000000589)	(0.000000636)	(0.000000617)
Time Trend			1.000353	0.9998308	0.9999921	1.000012
			(0.00000293)	(0.0000023)	(0.00000405)	(0.00000313)
Ind. FE			YES		YES	
Month FE			NO		YES	
N	3479795		3479751		3479751	
R2	0.1391		0.1782		0.2293	

The table reports results from the estimation of a multinomial logit model for positive and negative price adjustment decisions. The base category is no price change. Coefficients show the estimated relative risk ratios, where values bigger (smaller) than 1 mean that the decision to adjust upwards or downwards is more (less) likely due to a change in the explanator and relative to the base category. Among the variables, g_y denotes the monthly log change in output, $CR8$ the eight-firm concentration ratio, σ_A the standard deviation of the sectoral four-factor productivity growth rate, π_{PPI} the monthly PPI inflation rate, and *Employees* the number of employees. Moreover, I include a linear time trend, industry fixed effects and month dummies with December as the base month.

Table 22: Multinomial Logit, Export Case

	(I)		(II)		(III)	
	-1	1	-1	1	-1	1
g_y	1.018808	1.089099	1.002388	1.045023	1.004474	1.047442
	(0.0106968)	(0.0101913)	(0.0118325)	(0.0113906)	(0.0121746)	(0.0116232)
CR8	1.013552	1.017163	1.006681	1.010344	1.006766	1.010439
	(0.0007773)	(0.0007045)	(0.0009981)	(0.0009172)	(0.0010007)	(0.0009182)
σ_A	1.021505	0.9573928	1.012742	0.9369535	1.009972	0.9349772
	(0.0048378)	(0.005051)	(0.0070172)	(0.0076143)	(0.0070002)	(0.0076078)
π	1.073997	1.123795	0.5501031	1.920271	0.5634607	1.95579
	(0.0270851)	(0.0258755)	(0.0457637)	(0.156598)	(0.0470416)	(0.1599026)
Δe	1.005814	0.9923809	1.008148	0.9947382	1.009252	0.9954236
	(0.0022192)	(0.0024643)	(0.0027206)	(0.0034465)	(0.0026903)	(0.003563)
π^*	1.004023	1.001729	1.003649	1.00037	1.00395	1.000497
	(0.0016621)	(0.0015033)	(0.001719)	(0.0016676)	(0.0017331)	(0.0016741)
Distance			0.9999763	0.9999753	0.9999284	0.9999438
			(0.00000465)	(0.00000434)	(0.00000877)	(0.00000781)
GDP					1	1
					(0.0000000000399)	(0.0000000000353)
Month FE			YES		YES	
Ind. FE			YES		YES	
R2	77329		67891		67891	
N	0.0144		0.1042		0.1051	

The table reports results from the estimation of a multinomial logit model for positive and negative domestic export price adjustment decisions. The base category is no price change. Coefficients show the estimated relative risk ratios, where values bigger (smaller) than 1 mean that the decision to adjust upwards or downwards is more (less) likely due to a change in the explanator and relative to the base category. Among the variables, g_y denotes the monthly log change in aggregate output, $CR8$ the US eight-firm concentration ratio, σ_A the standard deviation of US sectoral four-factor productivity growth rate, π the monthly US inflation rate, π^* the monthly export destination inflation rate, Δe the log change in the nominal exchange rate of the USD and the destination currency, *Distance* geodesic distance of the export destination and GDP destination country total real gdp. Moreover, I include a linear time trend, industry fixed effects and month dummies with December as the base month.