Upstream, Downstream: An Estimation of Armington Elasticities at Different Stages of Production

Paul Phillips

ECONOMICS WORKING PAPER SERIES
Working Paper 2024–04–A

U.S. INTERNATIONAL TRADE COMMISSION
500 E Street SW
Washington, DC 20436

April 2024

Upstream, Downstream: An Estimation of Armington Elasticities at Different Stages of Production
Paul Phillips
Economics Working Paper 2024–04–A
April 2024

Abstract

This paper uses panel data on variation in trade costs to estimate elasticities of substitution for NAICS industries and Harmonized System (HS) products. I build on previous studies that estimate these elasticities by computing elasticities both for upstream and downstream products. Results show that upstream elasticities are higher than downstream elasticities, indicating that upstream product groups are more substitutable for one another than downstream product groups. Highly aggregated downstream product groups have a lower elasticity of substitution than less aggregated downstream product groups, but this pattern is more ambiguous for upstream goods. Upstream goods and downstream goods both display some signs of aggregation bias, although this bias is more pronounced in HS products than in NAICS industries.

Paul Phillips
Research Division, Office of Economics
U.S. International Trade Commission
Paul.Phillips@usitc.gov

1 Introduction

The elasticity of substitution, or Armington elasticity, can have a large effect on structural estimates of the counterfactual changes brought about by alterations in trade policy. This parameter appears in constant elasticity of substitution (CES) utility functions, and describes the level of substitutability between alternatives, be they two different sectors or goods in the same sector that originate from different countries. A higher Armington elasticity indicates that consumers or firms may substitute the goods more easily for each other, and a lower Armington elasticity indicates that the goods are more complementary. Studies measuring Armington elasticities, as surveyed in Riker et al. (2021), vary widely in their methodologies and estimated results, and the literature lacks of consensus on how to estimate them precisely.

This paper contributes to the literature on estimating Armington elasticities by explicitly differentiating between 'upstream' and 'downstream' products.[1] An upstream product is used at an intermediate stage of production, or in the making of another product, while a downstream product is sold directly to consumers. For example, the World Integrated Trade Solution (WITS) categorizes HS 170250: Chemically pure fructose as an upstream product, while HS 170410: Chewing gum, whether or not sugar added would be considered a downstream product. Such distinctions can be important, as a litany of general equilibrium and partial equilibrium models contain multiple stages of production, and assuming an identical Armington elasticity at each stage of production may lead to misleading results. Due to data limitations, the majority of papers that incorporate multiple stages of production into a partial equilibrium or general equilibrium framework do not estimate Armington elasticities by sector at the upstream and downstream level.

In this paper, I use the trade cost framework from Riker (2020) to obtain regression estimates of elasticities both for NAICS and Harmonized System (HS) goods. Like Riker (2020), I perform my estimations with panel data on U.S. imports from the International Trade Commission's DataWeb. However, I also use a WITS classification of HS6 products as intermediate goods to divide each product group or industry into two sub–categories: upstream and downstream. I estimate Armington elasticities by product group at the HS6, HS4, and HS2 levels and by industry at the NAICS3 and NAICS4 levels. I then use the aggregation bias framework discussed in Schrammel and Schreiber (2023) to ascertain how this aggregation bias differs between upstream and downstream goods.

[1] Throughout the rest of this paper, I will use the term 'upstream' interchangeably with 'intermediate' and 'downstream' interchangeably with 'final good.'

I find that upstream products usually have a higher elasticity of substitution than downstream products, indicating lower complementarity. For HS data this paper echoes the findings of Schrammel and Schreiber that product groups at a higher level of aggregation have lower elasticities of substitution than product groups at a lower level of aggregation, but this conclusion does not hold with the NAICS data. The relationship between aggregation level and elasticity magnitude further varies depending on whether a product group is upstream or downstream, indicating that an aggregation across levels of upstreamness could introduce omitted variable bias into this relationship.

Section 2 briefly discusses the academic context into which this paper fits. In Section 3, I outline the structural equations and data used for estimation of Armington elasticities. Section 4 presents the estimated elasticities for upstream and downstream goods across all sectors, then discusses how the aggregation bias differs across these categories. Section 5 concludes.

2 Literature Review

This paper fits in with a branch of economics literature that estimates Armington elasticities for multiple sectors at a time. Ahmad and Riker (2019) estimate Armington elasticities for four–digit NAICS manufacturing industries using the markup formula for monopolistically competitive firms. Soderbery (2018, 2015) uses CES demand equations to estimate these elasticities at the eight– and ten–digit HTS level (2018) and the four–digit HS level (2015). Hertel et al. (2007) and Broda and Weinstein (2006) also estimate Armington elasticities at disaggregated HTS levels. As discussed in Ahmad et al. (2021), there is little consensus on the appropriate method to estimate Armington elasticities or on the appropriate values such elasticities should take. Moreover, the literature that estimates Armington elasticities has not distinguished between products that are used as intermediate goods and products that are largely consumed as final goods. The amalgamation of these two categories under one umbrella may introduce bias into the resulting elasticity estimates.

Partial equilibrium models that incorporate multiple stages of production include Guberman et al. (2024), Schreiber (2023), Powers and Schreiber (2023), and Riker (2022). These papers use models with upstream and downstream goods to gain additional information on how trade policy changes can affect prices or other outcomes. All of the above papers except Schreiber (2023) concentrate on one particular industry with its own specific set of upstream and downstream products, and so do not compare upstream and downstream elasticities for multiple different HS or NAICS codes at one time. Schreiber (2023) presents a model with different Armington elasticities at the upstream and downstream level, but does not estimate the model using data.

A much larger body of literature embeds multiple stages of production into a general equilibrium framework. Like with the partial equilibrium models, the addition of multi–stage production lends additional realism to the models and additional specificity to the discussion of how trade shocks might affect prices and industries in the United States. These general equilibrium models either assume that the elasticities of substitution for upstream goods and downstream goods are the same or else abstract from this discussion by modeling intermediate production using a Cobb–Douglas function. Alessandria et al. (2022) is in the former category, using a CES aggregator for intermediate manufactures that is exactly the same as the CES aggregator for consumers except for the presence of a home bias parameter. Alessandria and Choi (2014) also uses the same elasticity parameter for final goods and composite intermediate goods, while Caliendo and Parro (2015) and Caliendo, Dvorkin, and Parro (2019) model a CES function with elasticity η that is used both for consumption and in the production of more intermediate goods. Papers that model the first stage of production with a Cobb–Douglas function include Alessandria, Choi, and Ruhl (2021) and Sposi, Zhang, and Yi (2021). Kehoe, Ruhl and Steinberg (2018) model different elasticities between intermediate goods and between final goods, but the elasticities in their final goods CES aggregator depend only on country, not sector.

This study makes a novel contribution to the literature by putting together a database of Armington elasticities that explicitly distinguishes between upstream and downstream products. While this paper does not include a more detailed structural model like the papers discussed in the previous paragraph, the elasticities computed in this paper may provide an invaluable resource to future researchers seeking to calibrate such a model.

3 Methodology

I follow the trade cost estimation methodology detailed in Riker (2020) and used in Schreiber and Schrammel (2024) to estimate elasticities. I use United States import data from the USITC's Dataweb, subdivided by country of origin and customs district of entry. Assuming a CES demand framework with elasticity of substitution σ_j, the landed–duty paid value (LDPV) of imports of product j from country i going through port district d is given by

$$E_{idj} = \beta_{idj} E_j P_j^{\sigma_j - 1} (p_{ij} f_{idj})^{1-\sigma_j} f_{dj}^{1-\sigma_j} \tag{1}$$

where β_{idj} is a demand shift parameter, P_j is the aggregate price index for sector j, E_j is aggregate expenditure on goods of sector j, and p_{idj} is the producer price of imports of product j from country i. f_{dj} is the domestic trade cost factor, reflecting shipping rates and insurance charges of sending product j from port d to the consumer. I calculate the international trade cost factor, f_{idj}, as the ratio of the landed duty-paid value (LDPV) to the customs value (CV), or $\frac{\text{LDPV}}{\text{CV}}$. This measure would be one without trade costs, and a value higher than one reflects costs such as shipping rates, tariffs, and insurance charges. Taking logs of both sides of (1) gives us our regression equation:

$$\log(E_{idj}) = (1 - \sigma_j)\log(f_{idj}) + \alpha_{dj} + \alpha_{ij} \tag{2}$$

where α_{dj} and α_{ij} are, respectively, port– and country–fixed effects.

This paper departs from Riker (2020) and Schreiber and Schrammel (2020) by distinguishing between intermediate goods and final goods. The WITS provides a list of HTS6 products that it considers intermediate goods.[2] I cross–reference this list with a crosswalk between NAICS and HTS codes to obtain a list of intermediate and final products that comprise each NAICS4 industry. A median of 9.692 percent of imports in a given NAICS4 industry were upstream products, while a median of 55.12 percent of imports in a given HS4 category were upstream products. NAICS industries including baking, textile furnishings, motor vehicle manufacturing and furniture manufacturing had no corresponding intermediate products at all, while fabric mills, pesticide and fertilizer manufacturing, and aluminum manufacturing were composed almost entirely of intermediate products.

To determine elasticities of substitution for HS6 products, I simply regress the relevant landed duty–paid value on the relevant international trade cost factor. For NAICS4, NAICS3, HTS4, and HTS2, I regress the *aggregated* landed duty–paid values on international trade cost factors calculated using aggregated LDPVs

[2] The WITS divides HS products by stages of processing in order to distinguish between tariffs that may apply at different levels of the production chain, and I download all product categories categorized in 'UNCTAD–SoP2,' or intermediate goods.

and customs values. Aggregated values are simply the sum of trade in all HTS6 intermediate or final goods products that comprise a given HS4 or HS2 level, or alternatively all HTS6 intermediate or final goods products that map to a given NAICS4 or NAICS industry.

In the following section, I will summarize the elasticities I found and compare them across levels of upstreamness and levels of aggregation.

4 Results

4.1 Comparing Sector–Level Estimates

	Upstream			Downstream		
Classification	Min	Median	Max	Min	Median	Max
HS2	1.941 (.5168)	5.296 (1.219)	27.11 (7.487)	2.066 (.3404)	4.314 (.8350)	15.56 (3.843)
HS4	.0789 (.5365)	6.673 (2.084)	63.96 (11.43)	1.990 (.4038)	5.128 (1.327)	31.28 (11.19)
HS6	2.287 (.5462)	6.469 (2.304)	76.19 (32.21)	.0142 (.3932)	6.521 (1.835)	401.6 (94.39)
NAICS3	2.340 (.4000)	4.682 (.6305)	8.454 (1.629)	1.168 (.2783)	3.193 (.4123)	8.221 (.9193)
NAICS4	1.981 (.3537)	3.969 (.9153)	14.88 (2.385)	1.664 (.2783)	4.222 (.5878)	8.918 (1.580)

Table 1: Summary statistics of elasticity estimates

Table 1 shows summary statistics for upstream and downstream elasticities, broken up by measurement system (HS or NAICS), and then further by the level of aggregation. The numbers in parentheses represent minima, medians, and maxima of the standard error for regressions done for each classification. I have removed elasticity estimates that were negative or not significant at the ten percent level. In general, industries and groups composed of upstream products display higher elasticities of substitution than did groups composed of downstream products. For example, HS2 upstream product groups have a median elasticity of 5.399 compared to 4.314 for downstream sectors, while NAICS3 upstream industries have a median elasticity of 4.69 compared to 4.166 for downstream sectors. This result suggests that downstream goods are less easily substitutable for one another, or more complementary, than the intermediate products that go into production of these downstream goods.

Previous literature measuring these elasticities has found that more highly aggregated HS product groups display a lower elasticity of substitution than less highly aggregated HS product groups, indicating higher complementarity. This result also holds intuitively, since highly specified products would be more interchangeable. For example, one might suppose that a drill pipe of stainless steel (HS 730422) produced in the United States would be more substitutable for an imported stainless steel pipe than domestically produced articles of iron or steel (HS 73) would be substitutable for imported articles of iron or steel. However, this observed relationship among aggregation levels is not as clear in the upstream elasticities. The median upstream elasticity at the HS2 level is lower than either the HS4 upstream elasticity or the HS6 upstream elasticity, but the HS4 upstream elasticity is in fact slightly higher than the HS6 upstream elasticity. The NAICS4 upstream elasticity is lower than the NAICS3 upstream elasticity, while previous scholarship would lead us to expect the opposite pattern. This result suggests that the amalgamation of upstream and downstream products introduces some bias into the estimation of elasticities. U.S. firms might have stronger preferences between more aggregated sectors when it comes to purveying *final* goods, but not when the imports or domestic wares are to be used in the production of another good.

Table 2 summarizes the level of statistical significance for each classification system and level of upstreamness. Fewer than half of the HS elasticities are both positive and significant at the ten percent level, with NAICS elasticities displaying higher levels of significance because NAICS industries having a higher level of aggregation than HS product groups. Downstream elasticities are more likely to be statistically significant than those of upstream goods, as downstream products tend to have slightly more observations than upstream products. Finally, slightly over a third of NAICS upstream elasticity estimates are statistically significantly different from the equivalent downstream estimates when a t–test was performed. Taken together, these significance results imply that while the difference in substitutability between upstream goods and downstream goods may not be hugely relevant for most sectors, this difference is well worth paying attention to, especially in models with multiple stages of production.

Classification	Upstream	Downstream	Difference
HS2	.4459	.6146	
HS4	.2129	.3036	
HS6	.09714	.1489	
NAICS3	.6716	.7761	.3913
NAICS4	.8500	.8500	.3636

Table 2: Percent of a given category that is positive and statistically significant at the ten percent level

Elasticities for all HS sub–classifications follow a roughly normal distribution with an attenuated right

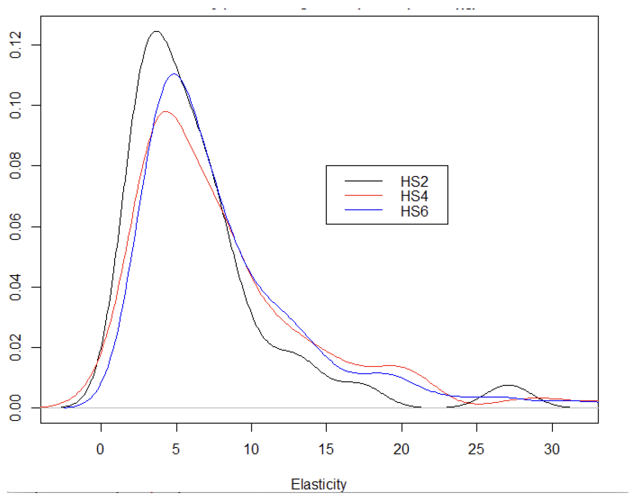

(a) Upstream Products

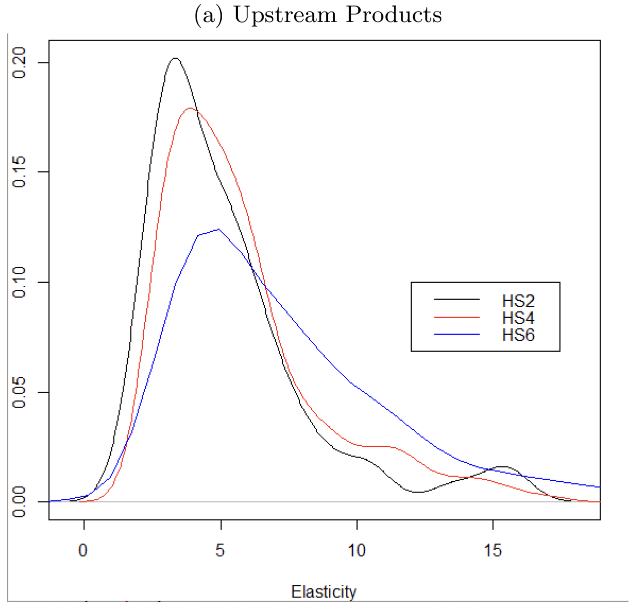

(b) Downstream Products

Figure 1: Elasticity Density Plots, Harmonized System

tail. The density plots shown in Figures 1 and 2 illustrate graphically how the median downstream elasticity is monotonically decreasing by level of aggregation and the median upstream elasticity is not. Furthermore, for downstream products but not for upstream products, HS6 elasticities display a higher level of dispersion than HS4 and HS4 elasticities display more dispersion than HS2, with a higher density of observations clustered around the median. This pattern does not hold for upstream products, where HS4 products have a higher level of dispersion from the median than HS6 products despite fewer observations and a higher level of aggregation.

The density of NAICS3 upstream elasticities shows an unusual binomial pattern, with elasticities clustered around three and also around five, the latter number of which is well above the median elasticity for all classifications. While the other density plots have a more clearly defined median, their distribution is nonetheless less clearly normal than the distributions shown in Figure 1. Downstream NAICS4 industries have an especially pronounced right tail, while both NAICS3 distributions have a cluster of observations around eight.

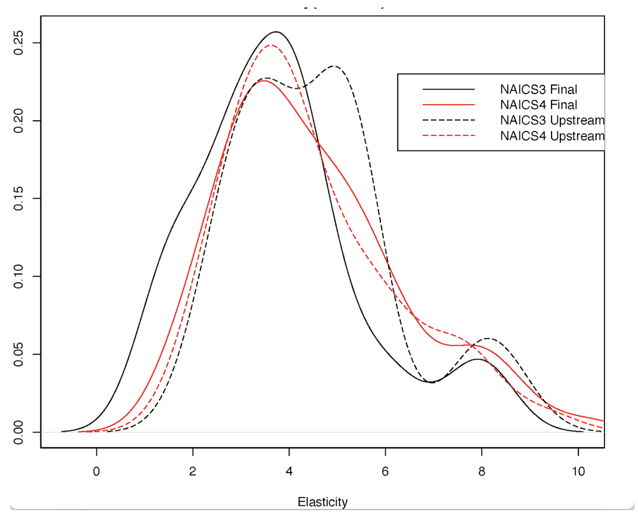

Figure 2: Elasticity Density Plot, NAICS

I will next discuss how upstream and downstream elasticities differ depending on the broad sector category being considered. Table 3 provides a side–by–side comparison of the observed elasticities for both HS two–digit product groups and NAICS three–digit industries that correspond to a particular sector. For example, the broader category of ‘prepared foodstuffs’ corresponds to NAICS industries 311–312 and HS product groups 16–24. I record the median observation among these industries’ elasticities as well as their standard deviation, which appears in parentheses below the median. For sectors with only one observation defined, I do not record any standard deviation.

The chemicals and transportation sectors have one of the highest Armington elasticities across both upstream/downstream goods and NAICS/HS categories, indicating that products in those sectors may be substituted without too much difficulty for products originating from another source. Base metals are fairly

			Upstream		Downstream	
Sector	NAICS3	HS2	NAICS	HS	NAICS	HS
Prepared foodstuffs	311–312	16–24	3.691 (1.454)	7.551 (.6630)	3.152 (1.901)	3.512 (1.450)
Textiles and apparel	313–316	50–67	2.900 (1.536)	2.769 (.2658)	3.846 (.8230)	3.557 (1.051)
Metals and base metals	331–332	72–83	4.612 (1.349)	5.949 (8.247)	2.950 (1.326)	3.791 (1.160)
Chemicals	325	28–38	5.176	4.749 (5.056)	6.253	5.087 (4.854)
Plastics, rubber Transportation	326 336	39–40 86–89	3.639 4.698	2.997 4.427	3.104 4.104	3.030 6.210 (5.213)
Wood, wood pulp, paper	321–322	44–49	3.454 (.6390)	2.193 (.3570)	2.810 (.1390)	3.025

Table 3: Elasticity estimates by broad sector category

substitutable when used as upstream goods, but much more complementary when consumed as final goods. Other groupings show comparable elasticities for their constituent upstream and downstream products.

NAICS estimates are generally similar to HS estimates for each category, with the exception of upstream prepared foodstuffs where the median NAICS elasticity is 3.691 and the median HS elasticity is 7.551, or nearly double. Note that HS codes record the physical characteristics of a product, while NAICS codes reflect the type of industry in which that product finds itself; this result therefore suggests that measuring foodstuffs by the sector to which they belong can result in underestimating the degree to which these foodstuffs may be substituted for foodstuffs from another source.

Table 4 shows the industry or product group with the minimum and maximum upstream and downstream

	Upstream		Downstream	
Classification	Minimum	Maximum	Minimum	Maximum
HS2	Cork, articles of cork	Tin, articles of tin	Perfumery, cosmetics	Misc. chemicals
HS4	Base metals clad with silver	Orthopedic appliances, splints, hearing aids	Festive articles, Carnival articles	Electric heaters, hand dryers
HS6	Plates, sheets of plastic	Pacemakers	Plastic bottles, flasks, carbuoys	Airplaces exceeding 15000 kg
NAICS3	Apparel	Computers, electronics	Beverages, tobacco	Primary metal manufacturing
NAICS4	Other nonmetallic minerals	Nonferrous metals (not aluminum)	Rubber products	Motor vehicles

Table 4: Sectors/Products with Minimum and Maximum Armington Elasticities

elasticity. In general, groupings with the minimum elasticities tend to be composed of raw materials, such as cork, rubber products, or other nonmetallic minerals, while groupings with the maximum elasticities are different types of electronic devices. This finding would imply that domestically produced electronic devices are easier to substitute for foreign imports than domestically—made manufactures of raw materials. However, as the level of aggregation changes, products or industries displaying the lowest or highest elasticity do not remain in the same broader groupings. For example, cork and articles of cork correspond to HS45, but base metals clad with silver (HS 7107) are not in the same two–digit grouping. This lack of continuity among minimum and maximum–elasticity sectors highlights the substantial heterogeneity that exists in elasticity estimates.

4.2 Aggregation Bias

Following the procedure in Schrammel and Schreiber (2023), I aggregate elasticities by trade weight and compare the results to the estimates from regressions performed with aggregated data in that product or industry group.[3] I then define 'aggregation bias' as the difference between the trade–weighted average elasticity and the estimated elasticity, or $\sigma_{\text{trade–weighted}} - \sigma_{\text{estimate}}$. I drop all elasticities from the aggregated set if they correspond to a product that is the sole product in its four–digit or six–digit aggregated grouping, as the inclusion of these values would bias results toward zero.[4]

Figures 3a) and 3b) show, respectively, the distribution of aggregation bias for HS6 and NAICS4 categories. Estimated elasticities of substitution are most often higher than aggregated elasticities for the HS data, while for the NAICS data the aggregation bias is closer to zero or even slightly positive. The increased precision among NAICS estimates may reflect these categories' being more aggregated to begin with. Another explanation is that aggregation bias is more pronounced when the objects being aggregated are products rather than industries.

For both classification systems, the bias among upstream elasticities is more dispersed than that of downstream elasticities, suggesting that trade–weighting of disaggregated estimates may be an especially important step when working with data on upstream products. The aggregation bias among HS4 elasticities is also less dispersed than that of HS2 elasticities, which is less surprising given that the four–digit elasticity

[3]For example, NAICS industry 313, textile mills, is composed of NAICS industries 3131–3133, representing respectively fibers, yarns and threads; fabrics; finished and coated textile furnishings. I can take a sum of the elasticities for these four–digit industries, weighted by their share in U.S. imports, and compare that sum to the estimated elasticity for NAICS 313.

[4]The inclusion of these values does not make a substantial difference to results.

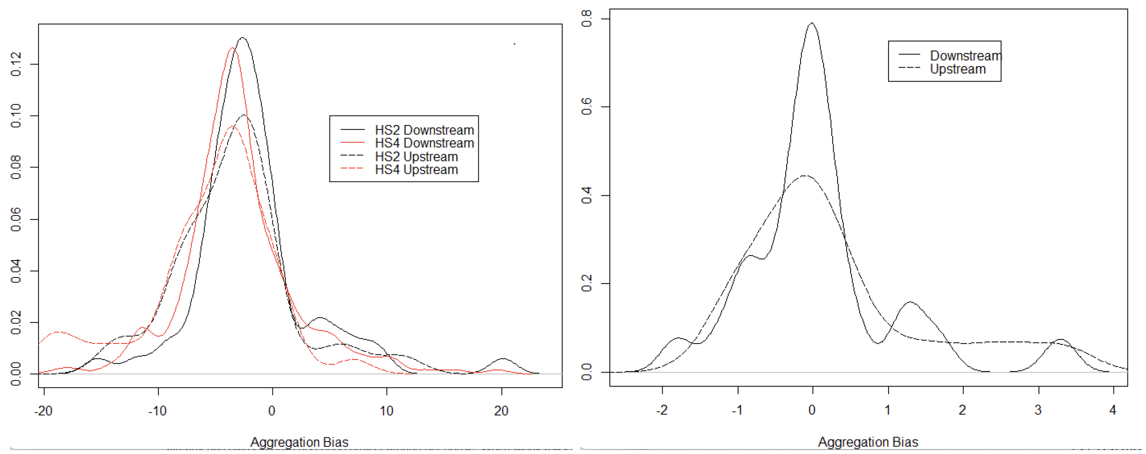

(a) Harmonized System

(b) NAICS

Figure 3: Aggregation Bias Density Plots

grouping would contain more observations and less precision.

Schrammel and Schreiber (2023) and Imbs and Mejean (2015) find that as the number of aggregation levels decreases, the trade–weighted estimates become closer to the corresponding Armington elasticity calculated with data from the larger product grouping. However, Figure 3a) shows the reverse pattern for both upstream and downstream elasticities; the median bias from aggregating from HS6 to HS2 is smaller in magnitude than the median bias from aggregating from HS6 to HS4.

		HS2	HS4	NAICS3
Upstream	Minimum	Live trees; bulbs, roots; cut flowers	Foliage; branches; grasses, mosses for bouquets	Wood products
	Maximum	Manmade staple fibers, yarns and woven fabrics	Articles of yarn, strip	Textile mill products
Downstream	Minimum	Inorganic chemicals	Sauces, condiments seasonings	Petroleum and coal products
	Maximum	Misc. chemical products	Rapeseed, colza or mustard oil and their fractions	Transport equipment

Table 5: Sectors/Products with Minimum and Maximum Aggregation Bias, in Magnitude

Table 5 displays the identity of the sectors or products with the highest and lowest *magnitude* of aggregation bias, or how far away the aggregation bias is from zero.[5] Among intermediate goods, groupings related to the textile industry contain the largest aggregation bias across all surveyed groupings, suggesting that

[5]I do not include summary statistics by broad sector here, as the standard deviations within each sector are so large that such measures are not very informative.

textile products are varied enough that any aggregated regression of them could yield misleading results. Among Harmonized System intermediate product groups, plant–related product groups showed the least aggregation bias. Beyond these observations, however, there are no clear patterns within the sectors displaying the most and least aggregation bias, a testament to the substantial heterogeneity that exists there.

5 Conclusion

In this study, I use country and port district–level panel data to estimate Armington elasticities for a variety of NAICS industries and HS product groups. Unlike other papers in the literature, I estimate these elasticities separately for intermediate, or upstream goods, and for final, or downstream, goods. I identify whether a given HS6 product is upstream or downstream using a data set from WITS, and I then aggregate these sorted sectors across NAICS and HS levels. I compare elasticity estimates across sectors and levels of aggregation, and compute aggregation bias for the more aggregated sectors by comparing trade–weighted summations with the elasticities computed by running regressions on the more aggregated data.

Median elasticities land in a range between three and seven, with upstream products often having higher Armington elasticities than downstream products. More aggregated final goods groupings show lower substitutability than less aggregated final goods groupings, but this relationship is not as straightforward with upstream goods. Among industries, the chemicals and transportation industries have the highest elasticities of substitution, while final goods in the metals and base metals industry have significantly more complementarity than upstream goods in this industry. Trade–weighted elasticity aggregates tend to be lower than elasticities estimated using the aggregated data, for HS products more than for NAICS industries. Interestingly, the aggregation of HS6 to HS2 results in lower bias for upstream and downstream products than the aggregation of HS6 to HS4.

The results from this paper imply that future structural models with multi–stage production should consider using different elasticity values depending on the stage of production, and that the upstreamness of a given aggregated product group or industry can introduce bias in the estimation of that grouping's elasticity. Several of the observations displayed in this paper contradict the conclusions drawn by previous literature studying Armington elasticities, and one potential explanation for these discrepancies is that the the omission of upstreamness measures from those papers introduces omitted variable bias. For example, the elasticity of substitution for final goods decreases with the level of aggregation, a conclusion also reached by Schrammel and Schreiber (2023), but the elasticity of substitution for upstream goods does not decrease as monotonically with the level of aggregation.

Future research should use more precise econometric techniques to examine exactly how the omission of upstream and downstream distinctions can result in biased elasticity estimates. Some of the stylized facts discussed in Section 4 defy expectations of how two variables would be related to one another, such as the magnitude of aggregation bias becoming smaller as the number of aggregation levels decreases. Furthermore, future work should consider more precise methods of estimating large numbers of elasticities at the upstream and downstream levels. Estimating these quantities with panel data on U.S. imports results in a large number of elasticities that are either statistically insignificant or impossible to estimate due to insufficient numbers of observations. As discussed in Sections 1 and 2, the Armington elasticity literature lacks a consensus on how to compute or estimate elasticities, and convergence towards such a method would give researchers and policymakers a powerful set of parameters to use in calibration.

6 References

Ahmad, S., Montgomery, C. and Schreiber, S. (2021). A Comparison of Sectoral Armington Elasticity Estimates in the Trade Literature, *Journal of International Commerce and Economics.*

Ahmad, S. and Riker, D. (2019). A Method for Estimating the Elasticity of Substitution and Import Sensitivity by Industry, USITC Economics Working Paper Series, 2019–05–B.

Alessandria, G., and Choi, H. (2014). Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization, *Journal of International Economics* 207–223.

Alessandria, G., Khan, S., Khederlarian, A., Mix, C., and Ruhl, K. (2022) The Aggregatie Effectsof Global Supply Chain Disruptions: 2020–2022 *NBER Working Paper 25965.*

Broda, C. and Weinstein, D. (2006). Globalization and the Gains from Variety, *Quarterly Journal of Economics* **121**(2): 541–585.

Caliendo, L., and Parro, F. (2015). Estimates of the Trade and Welfare Effects of NAFTA, *Review of Economic Studies* **82**: 1–44.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and Labor Market Dynamics: General Equilibrium

Analysis of the China Trade Shock, *Econometrica* **87**(3): 741–835.

Hertel, T., Hummels, D., Ivanic, M. and Keeney, R. (2007). How Confident can we be of CGE-based Assessments of Free Trade Agreements?, *Economic Modelling* **24**(4): 611–635.

Imbs, J. and Mejean, I. (2015). Elasticity Optimism, American Economic Journal: Macroe- conomics **7**(3): 43–83.

Kehoe, T., Ruhl, K., and Steinberg, J. (2018) Global Imbalances and Structural Change in the United States. *Journal of Political Economy* **126**: 761–796.

Powers, W. and Schreiber, S. (2023). Measuring the Impact of Trade Cost Volatility on Supply Chains: An Application to Plastics and Light Truck Manufacturing, *USITC Economics Working Paper Series, 2023–06–A*.

Riker, D. (2020). A Trade Cost Approach to Estimating the Elasticity of Substitution, *USITC Economics Working Paper Series, 2020–07–D.*

Schrammel, M. and Schreiber, S. (2023) How Does Data Aggregation Impact Elasticity of Substitution Estimation? Evidence from a New Elasticity Database, *USITC Economics Working Paper Series, 2023–10–B.*

Soderbery, A. (2015) Estimating Import Supply and Demand Elasticities: Analysis and Implications, *Journal of International Economics* **96**: 1–17.

Soderbery, A. (2018). Trade Elasticities, Heterogeneity, and Optimal Tariffs, *Journal of International Economics* **114**: 44–62.

Sposi, M., Yi, K., and Zhang, J. (2021) Trade Integration, Global Value Chains, and Capital Accumulation, *IMF Economic Review* **69**(3): 505–539.