Economic consequences of cabotage restrictions: The effect of the Jones Act on Puerto Rico

Russell Hillberry Manuel I. Jimenez*

December 23, 2022

Abstract

We estimate the economic burden placed on Puerto Rico (PR) by the "Jones Act," a US law that protects the US domestic maritime shipping market from foreign competition. We show that the supply of freight shipping to PR that satisfies Jones Act requirements lacks capacity for hauling general cargo and bulk commodities. In an empirical gravity framework, PR's imports of final goods reveal relatively greater substitution towards non-US sources among products that tend a) to be shipped by sea, b) to be physically heavy, c) not to be moved in containers. We use a structural gravity model and those product characteristics to estimate the tariff-equivalent trade costs the Jones Act imposes on US shipments of final goods. We use these estimates to calculate the compensating variation of Jones Act removal. Our preferred estimates suggest that the cost of final expenditure in PR would be \$1.4 billion (about 1.3%) lower per year without the Jones Act. The estimated annual burden on private consumption in PR is \$691 million (1.2%, or approximately \$203 per citizen). Ours are conservative estimates, for several reasons. Most importantly, we only consider effects of the Jones Act on purchases by final demand. The effects of the JA on upstream goods are not clearly visible in our estimates, which may reflect the evolution of the Puerto Rican industrial structure in ways that limit its reliance on inputs that are purchased from the U.S. and typically shipped by sea.

Keywords: Maritime Shipping, Cabotage, Jones Act, Gravity Model, Puerto Rico
JEL Codes: F13, F14, L91, R13, R48

---

[*] Department of Agricultural Economics, Purdue University, West Lafayette IN, 47906; e-mail: jimene19@purdue.edu. This research was financed by USDA NIFA grant #2020-67023-30963 and USDA Hatch Project IND-00134214. We thank the Puerto Rican Department of Treasury, the Department of Economic Development and Trade and the Planning Board and several local freight-forwarders for access to data and for their insights into Puerto Rico's external freight market. All errors are our own.

1 Introduction

The *US Merchant Marine Act of 1920* (the Jones Act) requires that maritime vessels moving goods from one US port to another must be US-built, US-owned and US-operated.[1] This protectionist policy raises the cost of intra-national maritime shipping in the US, and imposes a disproportionate burden on residents of US islands. These effects of the policy are generally understood as a qualitative matter, but there are few quantitative estimates of the economic burden the policy puts on US outlying areas. In this paper we estimate the economic burden the Jones Act (JA) imposes on residents of Puerto Rico (PR), an island territory of the United States.[2]

We begin our investigation with an exploration of data documenting ship movements in the Caribbean. Supplementary data from the same source reveal each ship's type, country of build, country of ownership and the flag under which it operates. We use these data to identify ships that call in Puerto Rican ports and satisfy JA regulations. We compare the characteristics of JA ships with others that call in PR, and with ships calling elsewhere in the Caribbean. This exercise reveals a striking pattern; the JA fleet serving PR contains no ships designed solely for the purpose of moving general cargo or bulk commodities. A small number of JA-compliant barges carry these kinds of freight, but larger-capacity general cargo / bulk shipping used elsewhere in the Caribbean is missing from this segment of the US cabotage market. Most of the waterborne trade between PR and the mainland US is carried by container ships dedicated to serving the US-PR market.

Our primary tool of analysis is a pooled product-level gravity model that we use to investigate cross-commodity variation in the size of PR's "home bias" towards imports coming from the US mainland.[3] Our hypothesis is that the JA causes Puerto Rican importers to substitute towards foreign sources when ordering products that depend more heavily on ocean shipping, and, in particular, on ships that carry non-containerized freight. This substitution should appear as reduced

---

[1] Specifically, the Jones Act requires that every vessel serving any US domestic route must be (1) at least 75% owned by US citizens; (2) built in the US; (3) crewed by US citizens or permanent residents; and (4) registered in the US (Merchant Marine Act, 46 U.S.C. §55102, 1920). At least 75% of the crew should be US citizens, and all of its officers and engineers (Beason et al. (2015), cited in Olney (2020)).

[2] It is likely that the JA has similar effects on other US islands and outlying areas, including Hawaii (HI) and Alaska (AK). The US does not collect the detailed intra-national trade data for AK and HI that we use here. There is similar data for other US possessions, but these island economies are quite small, more distant from the US, and not subject to the same breadth of JA restrictions that apply to PR.

[3] Home bias in this context is a measure of the excess intensity of shipments from the US mainland, relative to shipments imported from foreign locations, after controlling for geographic distance and other relevant variables.

home bias for those products that depend on ocean shipping, with even larger reductions among products not usually shipped in containers. In our pooled gravity model we interact the dummy variable indicating shipments from the US mainland with product characteristics that affect the transport mode used to ship the product. In order to avoid a potential bias generated by US-PR shippers' endogenous choice of transport mode, we use US import data (net of flows from Canada and Mexico) to calculate each product's a) vessel share of imports, b) weight-to-value ratio, and c) the containerized-share of imported shipments. In our sample of final goods, we find evidence consistent with the hypothesis that PR's home bias towards imports from the US mainland is smaller for products that are vessel-shipped, heavy, and not typically shipped in containers. These patterns are much less clear in shipments of upstream products, where PR's idiosyncratic production structure appears to obscure the forces we identify among final goods. In an effort to be conservative about attributing costs to the JA, we focus our attention on PR's imports of final goods.

We incorporate trade elasticities from Fontagné et al. (2022) in our empirical gravity model to estimate implied tariff-equivalent JA trade costs. The coefficients on the interaction terms between home bias, product characteristics and these elasticities provide product-level estimates of the relative trade cost penalty paid by goods that originate in the US. Relative trade costs become absolute when we normalize JA trade costs of a given air-shipped product to zero. To calculate *applied* JA tariff-equivalent trade costs, we multiply the structural gravity model's predicted tariff-equivalent JA costs by the share of PR's imports of the product from the US that actually moved by sea. Using this approach, we estimate that 87.0 percent of final products shipped from the US mainland have a positive JA-tariff equivalent; the simple average tariff-equivalent trade cost is 30.6 percent, while the trade-weighted average JA trade cost is 53.6 percent. We use an input-output table to map these estimates onto Puerto Rican consumption aggregates and calculate the compensating variation (CV) of removing JA restrictions on PR's imports from the US mainland. Our preferred estimates suggest that removing the JA would reduce the cost of final expenditure in PR by $1.4 billion (in 2016 dollars), which is approximately 1.3 percent of PR's final expenditure. We conduct CV calculations for different types of final expenditure. Focusing on household consumption, we estimate that, in return for removal of the JA, households in PR would accept $691 million in reduced expenditures. This amounts to about 1.2 percent of consumption expenditures, or $203 per citizen.

Our estimates include only the additional costs imposed by the JA on PR's *imports*, and consider only the distortions the JA places on shipments of final goods.[4] It is quite plausible that the JA is even more distortionary for imports of upstream products, some of which rely heavily on bulk shipping. We estimate the same specifications for upstream products as we do for final goods, but find little evidence that the effects of the JA on upstream goods reveal themselves in the same way we observe for final goods. We think this is due to at least two reasons. First, the JA is now more than 100 years old. Industries that process bulk-shipped goods have likely chosen to locate outside of PR rather than bear the ongoing burden of excessive freight costs on their inputs.[5] Second, PR has a history of significant tax advantages over competing locations in the US mainland, incentives that were most lucrative for certain manufacturing industries, especially pharmaceuticals and medical equipment. These incentives clearly influenced the industrial structure of PR, and have done so in ways that affect the relative demand for air- and ocean-shipped intermediates. Our CV estimates assume no impact at all on upstream sectors, and are thus quite conservative.

There is a relatively small academic literature on the JA. The paper that is closest to ours is Olney (2020), who shows evidence of substitution away from waterborne shipping among data on shipments arriving in U.S. coastal states. Our econometric exercise is similar, but with a few key differences. First, we focus on PR (an island) rather than coastal US states, and use data on Puerto Rican imports rather than freight movements destined for US mainland ports.[6] Our focus on Puerto Rican trade means that only two modes of transport are relevant in our data (air and sea), while Olney's data contain possibilities for easier substitution towards rail and road transport. We measure product characteristics (e.g. how likely is a shipment to move by sea?) in US import data (ex PR), while Olney uses the transport mode of the arriving shipment to make inferences

---

[4] The relative absence of bulk shipping may also impose a cost on PR's exports to the US mainland. Interviews with freight forwarders in PR indicated that the cost of moving containerized goods to the US is not overly burdensome because the ships are not full on their return voyage from PR. On the other hand, an ability to reach additional US ports by sea would reduce the overall transport costs of reaching inland customers on the US mainland.

[5] Hillberry and Hummels (2002) propose an explanation for excess home bias via a "production location effect" in which firms buying inputs choose to co-locate with their suppliers, thus obviating the need for costly trade. Hillberry and Hummels (2008) find evidence for this phenomenon in US Commodity Flow Survey data. In this context a production location effect would see industries that rely on heavy goods as inputs choose to locate nearer their input supplies in the US mainland, possibly avoiding location in PR altogether and making trade in those inputs unnecessary.

[6] Throughout the paper we will refer to shipments arriving to PR from the US as "imports." We will distinguish imports by their origin: US or rest of the world.

similar to ours.[7] Unlike Olney, we estimate a structural parameter, the tariff-equivalent cost of JA restrictions that explains the cross-product variation in estimated home bias that is attributable to the relevant product characteristics. These estimates reveal relatively higher tariff equivalents for products that are unusually heavy and/or unsuited for container shipping for other reasons (such as size or shape). We combine this information with evidence on Puerto Rican expenditure patterns to calculate the implied burden of the JA on PR.

We calculate the CV of the economic burden imposed by the JA on PR. Francois et al. (1996) use a computable general equilibrium (CGE) model of the entire US to measure the equivalent variation of removing the JA for the US economy as a whole. They calculate that the welfare cost of the JA to the US economy was approximately $3 billion in 1989.[8] These CGE estimates would include a significant burden of the JA operating through higher prices for upstream inputs. Our CV calculations consider only distortions to purchases by final demand, and rely solely on a cost/expenditure function, rather than a complete CGE model.[9]

There is also a consulting and/or policy literature on the JA.[10] The most relevant of these studies for our paper is John Dunham & Associates (2019), which notes that the JA is likely to put an especially large burden on the movement of heavy goods. We build on this insight, illustrating that the JA fleet serving PR relies heavily on containerized shipping. The John Dunham & Associates (2019) study estimates the cost differential for shipping products to PR on US and foreign routes using two different products over 10 different routes. These excess freight cost margins are imposed on every sector of the PR economy, and the effects of the JA on output, jobs and wages calculated in an input-output model. Instead, we exploit product-level variation in goods' physical weight,

---

[7] Since mode choice is endogenous to the presence of the JA, including the transport mode choice on the right hand side introduces a potential endogeneity problem.

[8] Estimates using this methodology also appear in USITC (1991, 1993, 1999, 2002). The estimated burden of the JA fell over time in these exercises, largely because the demand for intra-national waterborne freight movements in the US appeared to fall. PR was not considered part of the US economy in the CGE studies cited here.

[9] Protectionist measures like the JA also generate rents for domestic suppliers (US shipbuilders, owners of US ships, and US crews). Our approach assumes that these benefits accrue to residents of the US mainland, not to Puerto Ricans.

[10] See e.g., FED-NY (2012), GAO (1988, 2013), Grennes (2017), Kashian et al. (2017), Advantage Business Consulting (2019) and John Dunham & Associates (2019). There is also a Spanish language literature discussing the burden the JA puts on PR (e.g. Herrero-Rodríguez et al. (2003) and Valentin-Mari and Alameda-Lozada (2012)). Herrero-Rodríguez et al. (2003) reviews this literature and points to important early Spanish-language studies of the JA, including Pesquera (1965) and Quinonez-Dominguez (1990).

dependence on ocean-shipping and on container ships to measure the degree to which PR importers substitute away from US products that rely on non-containerized freight (such as general cargo and bulk ships). We estimate a structural gravity model and infer the implied tariff-equivalent costs of the JA. Rather than an input-output framework, we evaluate the costs of the JA with a framework more useful for economists, CV.

The remainder of the paper is organized as follows. Section 2 briefly reviews the unusual features of the Puerto Rican economy that are relevant to the exercise. Section 3 describes the data. Section 4 summarizes data on ship arrivals in the Caribbean, and compares characteristics of the JA fleet to other suppliers of freight services on Caribbean routes. Section 5 describes the theoretical framework and empirical estimation approach. In section 6 we report results from the gravity model, the JA tariff-equivalent trade costs and the CV of removing the JA. Section 7 concludes.

2 Structure of the Puerto Rican Economy

The effects of US sovereignty on the Puerto Rican economy are wide-ranging and consequential.[11] Separating the consequences of the JA from those of US sovereignty more broadly thus poses important challenges. In this section we offer a brief description of the Puerto Rican economy, focusing on one particular US policy that complicates our approach to measuring the effects of the JA, and probably leads our estimates to understate the burden the JA places on the PR economy.

Table 1 reports the gross output and employment shares of various sectors of the Puerto Rican economy. As is evident, manufacturing is a dominant sector. Manufacturing accounts for 45.2 percent of gross output and 11.8 percent of employment in PR.[12] Table 2 disaggregates the data on manufacturing activity. Pharmaceuticals and related manufacturing account for 64.6 percent of the manufacturing sector's gross output in PR, and 34.2 percent of its employment.[13] Medical equipment and supplies account for an additional 14.8 percent of manufacturing gross output and 23.7 percent of manufacturing employment, respectively. Manufacturing related to the production of

---

[11] For example, Marein (2020) finds that the average height of Puerto Rican males rose by 4.2 cm relative to their counterparts in the Caribbean during a 50-year period following US annexation.

[12] By comparison, in 2019 the manufacturing sector accounted for 15.8 percent of US gross output and 8.5 percent of US non-farm employment.

[13] While there appear to be no published data on the value added by the Puerto Rican Pharmaceuticals and Related Manufacturing sector itself, the sector does account for 98.7 percent of PR's gross output in the larger Chemical Products sector, which accounts for 72 percent of PR's total value added in manufacturing.

pharmaceuticals and medical equipment clearly represent a sizable share of total economic activity in PR.

The outsized importance of these two manufacturing sectors may be *partially* attributable to the JA. It is likely that both their inputs and their outputs have low weight-to-value ratios that make air shipping cost-effective. But a far more important reason these sectors are so dominant is a legacy of tax exemptions given to firms located in PR. The "Possession Tax Credit" - commonly known as Section 936 - offered qualifying firms a US federal tax credit equal to their total US tax liability.[14] The tax credit was especially advantageous for firms with intangible assets (such as patents).[15] While Section 936 was phased out by 2006, the continuing importance of the pharmaceuticals and medical equipment sectors in the island's economy is evidence that PR's existing industrial structure remains a legacy of Section 936. This matters for our study, for at least two reasons. First, a large share of PR's trade occurs in Pharmaceuticals (20.9 percent of imports and 74.0 percent of exports). Second, the composition of demand for imported intermediates is driven by the composition of output. The large skew of manufacturing towards pharmaceutical products (and other goods affected by Section 936) makes it difficult to observe effects of the JA on trade in intermediates. We exclude trade in pharmaceutical products from our regressions, but nonetheless find little implied effect of the JA on trade in upstream goods. This finding is striking because the estimated effects of the JA on PR's imports of *final* goods are both intuitive and sizable.

3 Data

We exploit three main sources of data. Comprehensive data on ship ports-of-call in the Caribbean help us to understand the relative supply of different types of JA and other shipping. We also estimate an empirical gravity model using data on PR's imports – from both foreign and US mainland sources. Once we have estimated tariff-equivalent trade costs attributable to the JA, we match them to data on final expenditures from a Puerto Rican input-output table. In this section we describe these three data sets, as well as some ancillary data that we use in our estimation.

---

[14] For a fuller discussion of Section 936, see GAO (1993).

[15] See Feliciano and Chen (2021), who study the effects of the ending of Section 936 on Puerto Rican manufacturing.

3.1 Port of call data

In order to better understand the relative and absolute supplies of different shipping services to PR, we purchased comprehensive data on freight vessels' ports of call in the Caribbean from Lloyd's List Intelligence (LLI). This is a commercial firm that provides these data to support "decisions in compliance, risk management, and operations."[16] The only other academic use of these data of which we are aware is Taylor (2021), who studies the impact of large oceangoing vessels on the reproduction rates of Southern Right Killer Wales. We report summary statistics from these data to guide and motivate our subsequent analysis of Puerto Rican import data.

The data offer comprehensive information on ports of call by freight hauling vessels in the Caribbean during the years 2004-2020. Each port of call record includes a unique vessel ID number. We also purchased from LLI data on important characteristics of the 16,188 freight-hauling vessels in our database. The following information is available for nearly every vessel: the vessel identification number, vessel type (bulk, containership, tanker, etc.), flag of registry, year and place of build, the vessel's owner, and its dead-weight tonnage (DWT).[17]

We use these data in two ways. First, the vessel characteristics data report information on the place-of-build, flag-of-registry and ownership, which allows us to identify ships that are JA compliant.[18] Second, we combine the vessel characteristics and the port-of-call data sets to offer a sketch of freight shipping in the Caribbean. We compare the observed supply of JA-compliant shipping services to the characteristics of the overall freight hauling fleet that calls in PR and the fleet that calls in other Caribbean ports.

---

[16] `https://www.lloydslistintelligence.com/`

[17] Some records do not contain DWT. Out of the total 16,188 vessels in the sample serving the Caribbean during 2004-2020, LLI data do not report the DWT of 202 vessels.

[18] Because tracking changes in ownership over time is somewhat onerous, we screen first for US-built and US-flagged ships. We then check by hand to verify that the firms that own the ships we initially identified are located in the US. The JA also requires domestic crewing of ships that are US-built, US-flagged and US-owned. Our data lack comprehensive data on the crews. We assume for these calculations that the ships arriving in PR that meet observable JA requirements are also US-crewed. This assumption only affects our initial summary statistics; it is irrelevant to our main results.

3.2 Puerto Rican import data

The empirical gravity model that we estimate relies on data documenting flows of imports into PR. These data are provided by the Instituto de Estadísticas de Puerto Rico (IEPR) for 2010-2017.[19] The IEPR constructs this data set by joining two data sets released by the US Census: (1) US merchandise trade imports; and (2) trade with US Possessions. The resulting data set reports information for monthly imports to PR disaggregated by HS10-digit product code, origin country (for foreign imports) and US customs district (for US-origin shipments).[20] The data report the value of imports (defined in FOB terms) and imported quantities (measured in kg.).[21]

The Puerto Rican import data are the source of the dependent variable in our pooled product-level gravity regressions. For independent variables, we calculate great circle distances from each country or US customs district to the port of San Juan.[22] For reasons that we describe later in the paper, we choose to parameterize export supply in the shipments' origin rather than relying on fixed effects to sweep out heterogeneity in supplying regions' product-level export supplies. To do this we calculate each origin's total export supply of a given HS6 product in a given year, and include it as a control in the gravity regression. The export supply measures are calculated using the BACI data for non-US origins and US export data for US origins.

Our analysis also relies on product characteristics. In particular, we exploit cross-product variation in characteristics that relate to demand for particular kinds of shipping and/or the freight rates that might be charged for transporting a given dollar value of that product. We use four product characteristics to predict reliance on particular kinds of shipping: a) the value share of a product's annual imports that moves by sea, b) the log weight-to-value ratio of imports in the commodity, c) the squared log weight-to-value ratio, and d) the value share of imports that are containerized. We calculate all these measures with US import data so that they are exogenous to the flows we observe involving PR.[23] All US export and import data files were retrieved from Peter Schott's

---

[19] We retrieved the data from `https://datos.estadisticas.pr/dataset/comercio-externo/resource/b4d10e3d-0924-498c-9c0d-81f00c958ca6`

[20] The variables we use to control for export supply and the trade elasticities we include to generate structural trade cost estimates are only available at the HS6-digit level. As a result, we aggregate our trade data to the HS6 level.

[21] The data do not report port of destination in PR, so for the purpose of calculating shipping distances we make a working assumption that all seaborne freight traffic goes through San Juan. Our port-of-call data from LLI shows this to be imperfect, but quite reasonable.

[22] We retrieve the GPS coordinates from `https://simplemaps.com/data/world-cities`

[23] We exclude imports from Canada and Mexico in these calculations so the US data we use reflect the air-vs-sea

web page.[24] We estimate the model separately over subdivisions of final and upstream goods. Our primary tool for separating final and upstream goods is the upstreamness measure from Antràs et al. (2012), though we also use the United Nations' BEC classification in a robustness check.[25]

Our gravity regressions also include trade policy measures that enter as control variables: a) the US statutory MFN tariff rate taken from USITC (2018), and b) dummy variables indicating countries that are members of a preferential trade agreement (PTA) with the US.[26] Finally, we include in the regressions (as data) product-level elasticities of substitution estimated in Fontagné et al. (2022). The inclusion of these elasticities in the estimation allows us to interpret the regression coefficients we estimate as directly informative of tariff-equivalent trade costs.

3.3 Puerto Rico's Production Structure

Calculation of the welfare costs of the JA requires information that is more comprehensive than what we have available to us in the trade flow data. A critical element for such calculations is data on PR's purchases of its own output, data that is not included in the trade data. We employ a Puerto Rican input-output (IO) table that separates expenditures on local output from expenditures on imports (including imports from US sources). The IO table was produced by the Junta de Planificación de Puerto Rico for the years 2006-2007.[27] This table reports final demand for every sector in PR's economy at the 4-digit NAICS code level. It also disaggregates final demand into several components: consumption, investment, exports, and governmental expenditure. The table also reports - for every purchasing NAICS code including final demands - expenditures on local production and on imported products, respectively. We match the trade data to the NAICS codes, and calculate US and rest of world (ROW) shares of imports for each NAICS code. This exercise allows us to calculate expenditure shares on purchases from PR, the US and ROW for each NAICS code.

---

choice that is available to shippers on PR routes. In this we follow Hummels and Schaur (2013) who use air and sea shipments to measure the value of time in the movement of US import shipments.

[24] `https://sompks4.github.io/sub_data.html`. All dollar values are deflated by the US consumer price index and expressed in 2019 dollars.

[25] We specifically use Revision 4 of the BEC classification and the conversion from HS2017 produced by United Nations (2022).

[26] We take the list of US PTAs from the US Trade Representative web site: `https://ustr.gov/trade-agreements/free-trade-agreements`.

[27] Unfortunately, the 2006-2007 table is the most recent table available for PR. Since this is the only data we have on purchases of PR's goods and factor value added, we use it, though we recognize this as a limitation.

4 Stylized Facts

In order to better understand the characteristics of the supply of JA-compliant shipping, we use the LLI port-of-call data to describe the supply of ocean shipping in the Caribbean. We first use data on vessel characteristics to identify the vessels that satisfy JA regulations. We compare the observed supply of JA-compliant shipping services to other ships that call in PR and to ships calling in other Caribbean ports. Finally, we use PR's import data to describe mode choices across source countries and across aggregations of up- and down-stream products.

4.1 Type of vessels serving in the Caribbean

In 2019, 3,155 freight-hauling vessels made a port of call somewhere in the Caribbean (See Panel A in Table 3).[28] Tankers accounted for 37.9 percent of these vessels, followed by bulk ships (21.5 percent), container ships (17.0 percent), and general cargo ships (14.7 percent).[29] Tankers and bulk ships provided most of the shipping capacity (70.6 percent of the total offered Deadweight Tonnage (DWT)). However, container and general cargo ships made more frequent ports of call, accounting for 51.2 percent of the total, compared to 31 percent for tankers and bulk carriers.

The type of vessels serving PR in 2019 were similar to those serving the broader Caribbean. In 2019, 363 vessels arrived in PR, accounting for approximately 11.5 percent of the vessels active in the Caribbean market (See Panel B in Table 3). Tankers accounted for 52.1 percent of the vessels calling in PR, followed by container ships (13.8 percent), bulk ships (13.2 percent), and general cargo ships (11.3 percent). Tankers and bulk ships also supplied most of the shipping capacity (73.8 percent of DWT). However, container ships arrived relatively more frequently in PR than in the broader Caribbean, making 41.7 percent of the port calls on the island.[30]

In order to distinguish JA-compliant ships from others serving the Caribbean market, we use the LLI data to identify vessels that were (1) built in the US; (2) US flagged and (3) US owned. Only 9 vessels that called in PR in 2019 satisfied these conditions (See Panel C in Table 3). The types of JA vessels serving PR were very different from others calling in PR, and from vessels operating

---

[28] The Caribbean in our sample includes all territories geographically located in the Caribbean sea in a (roughly rectangular) area contained by The Bahamas, Cayman Islands, Trinidad and Tobago and Montserrat.

[29] We use the year 2019 as the benchmark year instead of 2020 (the latest in the sample), because the global COVID-19 pandemic very likely affected market outcomes in 2020.

[30] The arrival rate of freight vessels in PR is approximately 5.5 per day; with approximately 4 of them arriving in San Juan.

in the broader Caribbean. Container ships accounted for 82.1 percent of the ports of call that JA vessels made in PR, and 80.8 percent of the reported DWT.[31]

The most notable fact about the JA fleet for our analysis is that the fleet of JA-compliant ships serving PR contain no bulk ships, tankers or general cargo vessels. JA-compliant barges carry bulk freight and general cargo, but there are no larger-capacity bulk or general cargo ships serving the US-PR market, even though such ships play an important role in Caribbean and PR shipping more broadly.[32] This fact motivates our subsequent exercises involving the gravity model of trade.[33]

4.2 Composition of Puerto Rico's Imports

Our primary tool of analysis is a gravity model of trade that pools across products. Before turning to this model, we report some aggregate statistics that inform the overall composition of PR's imports. We examine the composition of PR's aggregate imports and subdivide the data long multiple dimensions. Our interest is in characterizing differences in the composition of trade flows from the US and from ROW. We divide the data by mode of transport (air vs sea). We also separate products based on their position in vertical supply chains, using the 'Upstreamness' index proposed by Antràs et al. (2012).[34] We classify products with an upstreamness index of 1.3 or less as final goods, since these products are primarily purchased by final demand. Products with an upstreamness index greater than 1.3 are considered upstream goods that are purchased for further processing. We distinguish between final and upstream goods because product-level expenditures on upstream goods are highly sensitive to the industrial structure of the destination, and because PR's industrial structure is idiosyncratic among US regions and among Caribbean islands.

---

[31] One JA barge lacked data on DWT, so these figures overstate somewhat the share of containers in DWT calling in PR.

[32] Perhaps this should not be surprising. Brancaccio et al. (2020) shows that the operation of bulk ships in international trade typically follows a 'taxi' model. After delivering a cargo, a bulk ship looks to pick up another cargo nearby, delivering that cargo to whatever destination the cargo's owners prefer. This business model is effectively unavailable to JA ships if they are not competitive in international waters. If JA ships - once built - would be unable to compete effectively, then it is plausible that they will not be built at all. Containerships, on the other hand, are well-suited for the back-and-forth routes that are more common in U.S. domestic trade.

[33] The waiver of JA restrictions granted to a single non-JA compliant tanker that delivered 300,000 barrels of diesel fuel to PR after Hurricane Fiona in September of 2022 is probably indicative of the particularly high burden the JA imposes on bulk shipments. The most feasible large supplier of diesel fuel on short notice was the US, but there were no JA-compliant ships available to deliver it. See Page and Restuccia (2022) for a description of this particular episode.

[34] The Bureau of Economic Analysis (BEA) concordance between Input-Output Commodity codes and Foreign Trade Harmonized Codes of the year 2002 is used to merge this index to Puerto Rico's import database.

Table 4 shows the value and share of PR's imports in 2016 by type of goods (upstream and final), by mode of transportation (are and sea) and by the shipments' origin (US or ROW). Upstream products account for 90.5 percent of PR's total import value. 41.7 percent of PR's import value moved by oceangoing vessel. The US share in PR's total import value is 59.1 percent. 43.8 percent of PR's imports from the US are shipped by sea, as are 38.7 percent of imports from the ROW. 90 percent of final goods imports move by sea, while only 36.6 percent of upstream goods arrive by sea. Finally, the US accounts for 80.9 percent of PR's final goods imports, but only 56.8 percent of upstream imports.

Our primary takeaway from the data description exercises is that PR's imports of final goods are likely to be especially dependent on JA shipping. The US supplies most of PR's final goods imports, and final goods are heavily dependent on maritime shipping. We estimate empirical gravity models separately for both final and upstream products.

5 Theoretical Framework

Our approach to estimating the costs of the JA relies on a structural interpretation of the gravity model of trade. This framework has been used in similar efforts to quantify hard-to-measure trade costs or changes in trade costs.[35] Our effort to quantify the effects of the JA is complicated by two - partially offsetting - implications of US sovereignty. First, the JA should raise the relative costs of PR imports from the US mainland, causing importers to substitute away from US goods. These excess costs should be concentrated in products that are shipped by sea. But second, other effects of US sovereignty over PR should reduce trade costs on US imports (relative to imports from foreign sources), causing importers in PR to substitute in the other direction, towards goods sourced from the US mainland. This latter tendency - known in the literature as 'home bias' - we assume to be independent of the mode of transport used to deliver the product to the PR, while the effects of the JA are assumed to depend on the degree to which importing the product depends on waterborne transport. We exploit variation across commodities in the level of observable home bias, and ask whether products with characteristics that are likely to make them dependent on

---

[35] Hummels (1999) develops an approach to inferring the costs of distance and other geographic frictions in the gravity model. Anderson and Van Wincoop (2004) describe similar methods for inferring tariff equivalent trade costs of the US-Canada border. Head and Mayer (2014) offer a comprehensive guide to the use of the gravity model in international trade. Yotov et al. (2016) provide an accessible introduction to tools and methods used in this kind of research.

waterborne trade (and, in particular, on bulk/general cargo ships) have less observable home bias than other commodities.[36] Upon finding this pattern, we attribute to JA-related trade costs the relative difference in measured home bias that falls on goods most likely to be affected by the JA. When we move to using the structural model to estimate tariff-equivalent costs of the JA, we incorporate into the estimation (as data) external estimates of σ^k, the elasticity of substitution between varieties of a given product k.[37] The inclusion of external estimates of σ^k allow us to interpret the regression coefficients we estimate as product-level tariff-equivalents of the additional trade costs borne by waterborne shipments from the US. We attribute these inferred trade costs to the JA.

Our empirical estimates show a clear impact of the JA on downstream goods (goods purchased by final demand), but not on upstream goods. Our welfare analysis therefore focuses on the effects of JA distortions on final demand alone. We calculate the CV of the JA's effect on PR with an expenditure/cost function applied to spending on final demand.

5.1 Model Set-Up

There are several structural theories that can be used to motivate the empirical gravity model. The most suitable for our exercise is the version that assumes monopolistic competition, e.g. Helpman and Krugman (1985).[38] In the model, the representative agent in PR allocates their expenditures across commodities via an upper-level Cobb-Douglas Utility function, with expenditure shares for

---

[36] One might expect the effects of the JA to vary with distance, and so cause us to focus our investigation on cross-product variation in the distance elasticity of trade. These effects are plausible, but very difficult to cleanly identify. The US-PR shipping market appears to be completely segmented from the broader Caribbean shipping market. We focus instead on measuring the degree to which PR importers substitute away from US-origin products that move by sea, which is revealed by product-level variation in measured home bias. We count on the empirical gravity model to control for distance-related factors, and estimate flexibly with respect to distance in various specifications in order to check the robustness of our results.

[37] We are also able to use data on US tariffs to estimate, internally, a cross-commodity average σ. In a robustness exercise we use this parameter to infer a different set of JA costs, though our preferred estimates exploit the heterogeneous values of σ^k that were estimated externally.

[38] Several prominent theories - including this one - interpret bilateral variation in trade flows in the same way; the response of trade to geographic frictions is jointly determined by a bilateral trade cost and a trade elasticity. We choose the monopolistic competition model for two reasons related to our approach to identification: 1) our empirical approach requires us to account for variation across origins in the level of export supply, and 2) our approach to identifying trade costs relies on structural estimates that treat the trade elasticity as a demand side parameter. One could write a product-level Armington model with similar features, but we view Helpman and Krugman (1985) as better for illustrating a few subtleties, namely the possible endogeneity of the number of traded varieties. Our exercises will treat the number of traded varieties as exogenous to $\tau^k_{y,j}$, though it is also plausibly endogenous. In Melitz (2003), lower trade costs allow additional firms to select into serving a market. In Hillberry and Hummels (2002), locations of up- and down-stream firms trading intermediates are jointly endogenous, with trade in product k much less likely to exist at all when $\tau^k_{y,j}$ is large.

each commodity represented as α^k. Within commodity k, the consumer has constant elasticity of substitution (CES) preferences over varieties of k. Let $n^k_{j,y}$ be the number of producers in region j that export their variety of product k to PR in year y, and σ^k the elasticity of substitution between varieties of commodity k. The representative agent's utility in year y, U_y, is expressed as:

$$U_y = \prod_{k=1}^{K} \left[\sum_j (n^k_{j,y})(q^k_{j,y})^{\frac{\sigma^k - 1}{\sigma^k}} \right]^{\alpha^k \left(\frac{\sigma^k}{\sigma^k - 1} \right)} \tag{1}$$

where $q^k_{j,y}$ is the quantity per variety of product k purchased from location j in PR in year y.[39] The price of commodity k at origin j is $p^k_{j,y}$. The value of trade from region j to PR is $M^k_{j,y}$, which can be represented as the product of $q^k_{j,y}$, $n^k_{j,y}$ and $p^k_{j,y}$. The iceberg trade cost associated with moving goods from origin j to PR in year y is $\tau^k_{j,y}$. The price in PR of a variety of k purchased from origin j in year y is thus $p^k_{j,y} \times \tau^k_{j,y}$. Maximizing U_y in (1) given a level of annual expenditure E_y in PR and expenditure shares α_k returns a formula for the value of bilateral imports in commodity k, $M^k_{j,y}$ as follows:

$$M^k_{j,y} = n^k_{j,y} \times p^k_{j,y} \times q^k_{j,y} = n^k_{j,y}(\tau^k_{j,y})^{-\sigma^k} \left(\frac{p^k_{j,y}}{\widetilde{P}^k_y} \right)^{1-\sigma^k} \alpha^k E_y. \tag{2}$$

where $\widetilde{P}^k_y$ is the conventional Dixit-Stiglitz price index for good k in PR.[40] The focus of our interest shall be the form of the trade cost term, $\tau^k_{j,y}$. Following the literature on border effects, we consider $\tau^k_{j,y}$ to be a multiplicative form of trade costs that depends upon distance and international borders. Using an aggregate model (where $K = 1$), Anderson and Van Wincoop (2003) specify the trade cost function as follows:

$$\tau_{j,y} = dist^{\rho}_j (b_y)^{1-HOME_j} \tag{3}$$

where $dist_j$ is the distance from origin j to PR's main city and port, San Juan, ρ is the distance elasticity of trade costs, b_y is an estimable parameter equal to 1 plus the tariff-equivalent border cost associated with purchasing goods from outside the US in year y, and $HOME_j$ is an indicator that the product originated in the US. We amend (3) to include an effect of the JA on PR's imports

---

[39] We suppress a destination subscript here and elsewhere because the only destination in our sample is PR.

[40] Some authors, including Anderson and Van Wincoop (2003) write the exponent on $\tau_{j,y}$ as $1-\sigma$. This treatment assumes that trade value is measured in destination prices. Our data measures trade in FOB prices, so the $-\sigma$ exponent on the trade cost parameter is more appropriate. In this we follow Hummels (1999).

from the US. We assume that the JA imposes an additional tariff equivalent cost on PR's imports from the US, and that the tariff-equivalent JA cost varies across products according to product characteristics related to the manner in which the product is usually shipped. We specify another parameter like b,

$$JA_y^k = 1 + t_{JA,y}^k, \tag{4}$$

where JA_y^k is a vector of parameters to be estimated, and $t_{JA,y}^k$ a tariff-equivalent cost linked to the JA. The product-specific trade cost function now appears as:

$$\tau_{j,y}^k = dist_j^\rho (b_y)^{1-HOME_j} (JA_y^k)^{HOME_j}. \tag{5}$$

We exploit product level variation in the response of bilateral trade to the $HOME$ dummy to parameterize JA_y^k.[41]

5.1.1 Parameterizing $JA_{j,y}^k$

Conceptually, our approach to parameterizing JA trade costs is as follows. All ROW imports, regardless of transport mode, pay a common (average) tariff equivalent border cost, which takes the form $t_{HB,y}$ and is estimated via the parameter b_y. Goods that are shipped from the US by sea also pay a penalty (relative to goods shipped from the US by air). This trade cost takes the form $t_{JA,y}^k$ and is estimated by the parameter JA_y^k. A simple approach to parameterizing JA^k would be to assume a common JA trade cost that applies to all US-origin shipments that move by sea. In a pooled product-level gravity model, the effects of the JA would be identified through the coefficient on the interaction of the $HOME$ dummy and with an indicator that the goods moved by sea.[42] The $HOME_j$ coefficient (without an interaction term) would measure home bias toward U.S. products among air-shipped goods. The interaction terms would capture the reduction in home bias that is revealed among sea-shipped goods.

It is our view that the absence of bulk/general cargo carriers on US-PR routes is likely to cause important cross-product variation in the burden the JA imposes on sea-shipped goods from the US. Our approach to estimating these costs is a generalization of the method involving dummy variables

---

[41] Equation (5) is for purposes of illustration. When we move to estimation, we specify the function form of distance-related trade costs flexibly. Non-linearities in the effects of distance are allowed, as well as cross-product variation in the effects of distance.

[42] This is similar to the approach that Olney (2020) takes, though he does not attempt to estimate structural trade costs.

described above. To capture the burden the JA imposes on maritime shipments of product k, we specify a vector of product-specific characteristics, $\overrightarrow{Z}_y^k$. The elements of $\overrightarrow{Z}_y^k$ include an explicit measure of the degree to which transport of product k depends on maritime shipping, but also other characteristics that affect reliance on non-containerized shipping. We represent the log of the parameter JA_y^k as the inner product of the product characteristics $\overrightarrow{Z}_y^k$ and a vector of characteristic weights $\overrightarrow{\gamma}$, multiplied by the $HOME$ dummy:

$$ln(JA_y^k) = -\overrightarrow{\gamma}' \overrightarrow{Z}^k \times HOME_j. \tag{6}$$

In a reduced form gravity model, the elements of $\overrightarrow{\gamma}$ are not fully identified. The response of trade to any trade friction also relies on a trade elasticity. In our exercise this applies to the $HOME$ dummy itself as well to the responses that emerge from $ln(JA_y^k)$. Let β be the estimated coefficient on $HOME_j$ without an interaction, and $\overrightarrow{\beta}$ be the coefficients on the interactions of $\overrightarrow{Z}_y^k$ with $HOME_j$. In this case, response of trade to $HOME_j$ appears as:

$$\frac{\partial lnM_{j,y}^k}{\partial HOME_j} = \beta - \overrightarrow{\beta}' \overrightarrow{Z}^k. \tag{7}$$

As is familiar from the existing gravity literature, $\beta = \sigma b$; b can be identified only through the choice of σ.[43] Similarly, we interpret the parameters in $\overrightarrow{\beta}$ as the product of a trade elasticity and the predicted JA trade costs $\overrightarrow{\gamma}' \overrightarrow{Z}^k$.

5.2 Estimation

The key parameters of interest for what follows are the elements of $\overrightarrow{\beta}$, and their structural trade cost components $\overrightarrow{\gamma}$. We specify four $\overrightarrow{Z}_y^k$ variables that are intended to measure reliance on water-borne shipping, and bulk shipping in particular. The observable product characteristics we use are as follows: 1) the value share of US imports that travel by oceangoing vessel, Vsh_y^k, 2) the log of the median (across years) of the weight-to-value ratio of product k in US imports, $ln(WV^k)$, 3) the square of the logged median weight-to-value ratio $(ln(WV^k))^2$, and 4) the share of the product's US imports that were shipped in containers, $Ctnr_y^k$. We attribute to the JA the systematic variation in measured home bias that the pooled product-level gravity regression attributes to these characteristics. In subsequent regressions, we include estimates of σ^k in the regression, and interact them with geographic frictions. We interpret the coefficients on the interactions of σ^k with geographic

---

[43] Note that we will include in our regressions estimated values of σ^k. In this way we allow product specific responses to a common (or average) home bias parameter b.

frictions as structural trade cost parameters. Specifically, the inclusion of σ^k's in the regression allows the b term and the elements of $\overrightarrow{\gamma}$ to be identified. We use the estimates of $\overrightarrow{\gamma}$ to predict the tariff equivalent border cost of the JA attributed to product k.

Prior to turning to the structural model, we first estimate a reduced form gravity regression that links cross-product variation in estimated home biases to the characteristics of $\overrightarrow{Z}_y^k$. Before turning to the specification of this regression, it is useful to note some challenges that the data/estimation strategy imposes, and how we address them.

First, it is now conventional to estimate gravity regressions with vectors of fixed effects that sweep out important variation in the data. For example, many authors include some combination of origin-, destination-, product-, and time- fixed effects. In a cross-sectional regression with multiple origins and destinations, origin-product fixed effects control for systematic variation in the supply of a product, while destination-product fixed effects control for variation in expenditure levels and/or geographic remoteness of the destination.[44] In a time series context origin-product-year and destination-product-year fixed effects sweep out heterogeneous supplies and demands, as well as shocks to both supply and demand.

Our regressions use data from a single destination, Puerto Rico. This means that including destination-product fixed effects in the regression sweeps out cross-product variation in the data. Instead we attempt to parameterize import demand.[45]

Since our identification strategy relies heavily on the interaction between product characteristics and the $HOME$ dummy, the key threat to identification is if variation across products in the levels of demand (operating through the α_k's and the $\tilde{P}_y^k$'s) are correlated with the product characteristics of interest, the $\overrightarrow{Z}_y^k$'s. We address this problem in a manner that is conventional in applied econometrics; we include the product characteristics themselves (without the interaction) in the regression. In this way we control for cross-product variation in the level of demand that might

---

[44] See e.g. Anderson and Yotov (2016).

[45] In principle we could have included other destinations in the sample (especially other countries in the Caribbean), but since the US-PR flow would be the only domestic US flow this strategy would lead the $HOME$ dummy coefficient to compare US-PR flows to all flows to the Caribbean, rather than to ROW-PR flows. If the demand structure of PR was typical of the Caribbean, this might be preferable, but the effects of US sovereignty are likely to have made PR's import demand structure different than that of other Caribbean states.

bias the coefficients on the interaction terms.

Second, rather than sweep out variation in export supply with product-origin-year fixed effects, we include in the regression explicit measures of export supply (of a given commodity from a given origin in a given year). We do this because of our interest in the $HOME$ dummy variable, which would be co-linear with the usual full set of product-origin-year fixed effects. Instead, we fully parameterize export supply - using the total volume of exports of each product from each origin in each year - since these data are readily available in the trade data we have.[46] This approach allows us to estimate coefficients on the $HOME$ dummy, and on the interaction terms.

5.2.1 Model Specifications

Our reduced form regression model follows a Poisson Pseudo Maximum Likelihood (PPML) specification:

$$M^k_{j,y} = \exp\left[\delta\left(h^{-1}(X^k_{j,y})\right) + f\left(dist_j, \overrightarrow{Z}^k_y, \rho\right) + \beta HOME_j + \overrightarrow{\omega}\overrightarrow{Z^k_y} + \overrightarrow{\beta}\,\overrightarrow{Z}^k_y HOME_j\right] + \epsilon^k_{jy} \quad (8)$$

where $h^{-1}\left(X^k_{j,y}\right)$ is the inverse hyperbolic sine of the value of total exports of commodity k in time y from each region j, and δ the associated regression coefficient, $f\left(dist_j, \overrightarrow{Z}^k_y, \rho\right)$ is a flexible function of distance, product characteristics and parameters that controls for a region's distance to PR, and allows the effects of distance on trade to vary with product characteristics. $\overrightarrow{\omega}$ is a vector of estimated coefficients on the product characteristics themselves. β, $\overrightarrow{\beta}$ and $\overrightarrow{Z}^k_j$ are as described above. We include year fixed effects to account for annual shocks to the level of PR's import demand. In some specifications we also include the log of one plus the US MFN tariff, and a vector of dummy variables indicating that a country has a preferential trade agreements with the US. The coefficient on the MFN tariff provides an internal estimate of the elasticity of substitution, though the estimate assumes a common value of this parameter across products.

The reduced form specification in (8) is useful for illustrating the cross-product variation in home bias. But for measuring welfare we need to translate these estimates into trade costs. The basic problem is that the regression coefficients in (8) conflate the effects of trade costs and trade responses. We cannot infer trade cost parameters without an estimate of the trade elasticity. Our

---

[46] In some instances, the PR data report trade flows arriving from an origin, even though our corresponding data shows no exports of that product from that origin in that year. In these cases we add the PR trade flow total exports, and include a dummy variable indicating that we made this transformation.

solution to this problem is to incorporate into the estimation external estimates of the trade elasticity; we treat the estimates of σ^k in Fontagné et al. (2022) as data for the purpose of identification. We allow σ^k to enter into the regression itself. We also interact σ^k with the distance and the $HOME$ dummy, and with the interactions of these variables with $\overrightarrow{Z}^k_y$. According to our structural model, the coefficients on the interaction terms can be used to infer the trade costs that each friction imposes on each product k. Our new specification is as follows:

$$M^k_{j,y} = \exp\left[\delta\left(h^{-1}(X^k_{j,y})\right) + f\left(dist_j\sigma^k, \overrightarrow{Z}^k_y, \rho\right) + \gamma\sigma^k HOME_j + \overrightarrow{\omega}\overrightarrow{Z}^k_y + \varepsilon\sigma^k + \overrightarrow{\gamma}\overrightarrow{Z}^k_y\sigma^k HOME_j\right] + \epsilon^k_{jy} \tag{9}$$

where σ^k is the product-specific estimate of the elasticity of substitution from Fontagné et al. (2022), and ε is a reduced form estimate of the conditional correlation between σ^k and $M^k_{j,y}$. The key difference between this specification and that in (8) is that we have interacted σ^k with all of the geographic frictions, so that we can give a structural interpretation to the coefficient estimates. The coefficients of interest, γ and $\overrightarrow{\gamma}$ are structural equivalents to β and $\overrightarrow{\beta}$ (with $\gamma = -\frac{\beta}{\sigma^k}$ and $\overrightarrow{\gamma} = -\frac{\overrightarrow{\beta}}{\sigma^k}$). The γ term becomes lnb in equation (3), and $\overrightarrow{\gamma} \times \overrightarrow{Z}^k_y \times HOME_j$ produces a predicted distribution of tariff equivalent trade cost estimates of the JA. These estimates are not quite complete, because they are relative, rather than absolute measures of trade costs. We describe our process for turning relative estimates of JA trade costs into absolute values once we have our $\overrightarrow{\gamma}$ estimates in hand.

5.3 Welfare Analysis

Our welfare analysis focuses on quantifying the effects of JA distortions on PR's final demand. Our objective is to quantify the degree to which the JA requires higher levels of spending to obtain the same level of utility (for consumers) or output (producers). Our tool for this analysis is Compensating Variation (CV). Under the assumption that the rents that accrue to US shipbuilders, shipping companies and crews are received by agents on the US mainland, CV is an appropriate measure of PR's welfare loss due to the JA.

Consider an expenditure function $E(P_y, U_y)$ that reflects the minimized cost of purchasing an optimal consumption basket in PR in year y.[47] The expenditure function that is dual to the

---

[47] For final demand categories other than private consumption, we replace expenditure with cost and utility with output and conduct the same calculations.

utility function above is the product of a specific numerical level of utility $\overline{U}_y$, and the true cost of living index $P_y = \prod_l (\widetilde{P}_y^l)^{\alpha_l}$, where l indicates a NAICS sector that is an aggregate of the set of products $k \in l$.[48] The sub-indices $\widetilde{P}_y^l$ are CES aggregates of the delivered prices in sector l, $p_{j,y}^l \tau_{j,y}^l$), and the elasticity of substitution at the sector l level, σ^l.[49] Abstracting away from the possibility that changes in trade costs might induce increases in the number of varieties that PR purchases from the US mainland, for counterfactual analysis we define the price sub-index in sector l as:

$$\widetilde{P}_y^l = \left[\sum_j \theta_{j,y}^l (1 + \tilde{t}_{j,y}^l)^{1-\sigma^l} \right]^{\frac{1}{1-\sigma^l}} \tag{10}$$

where $\theta_{j,y}^l$ acts an Armington distribution weight, and $\tilde{t}_{j,y}^l$ is a trade cost in sector l that is set for removal in counterfactual analysis.[50] The set of regions j we consider in this analysis are US, ROW, and PR. In the case of US shipments $\tilde{t}_{j,y}^l$ is the JA tariff-equivalent; in the case of ROW shipments it is the trade-weighted average US tariff.

It is straightforward to calibrate this expenditure function. Let $S_{j,y}^l$ be the region $j's$ observed share of Puerto Rican purchases of sector l in year y. The presence of JA trade costs (in the case of the US mainland), and tariffs (in the case of ROW) means that the data shares are larger than the true distribution weights. The distribution weights can be uncovered by dividing the trade shares by the trade costs associated with an origin and product $\theta_{j,y}^l = {\frac{S_{j,y}^l}{1+\tilde{t}_{j,y}^l}}^{1-\sigma^l}$. The data required for this transformation - the values of $\tilde{t}_{j,y}^l$ and σ^l - are also applied where necessary in $\widetilde{P}_y^l$, and thus P_y. The α^l parameters are the observed expenditure shares from the IO table, whereas the initial values of $E(P_y, U_y)^0$ are the total expenditures observed in the table, and inflated by GDP growth to 2016.

Our calculation of CV is accomplished as follows. Let $\overline{U}_{jy}$ be the numerical value of utility associated with the initial price index P_y^0 and observed expenditures $E(P_y^0, \overline{U}_y)^0$. In counterfactual analysis we remove $\tilde{t}_{j,y}^l$ in (10) on US and ROW imports, respectively, assuming no change in prices

---

[48] Our data on PR's purchases of its own output (and value added) do not allow us to calculate welfare at the same level of aggregation as we use in the motivation for the structural regressions, so we replace superscript k with l.

[49] σ^l is a trade-weighted average of σ^k's.

[50] The link between the Armington and the monopolistic competition frameworks is: $\theta_{j,y}^l = n_{j,y} \left(\frac{p_{j,y}^l \tau_{j,y}^l}{\alpha^l (1+\tilde{t}_{j,y}^l)} \right)^{1-\sigma^l}$. This shift allows us to be more transparent in our calculation of CV. It also makes clear that we are not attributing to the JA any increase in product variety. This is another attempt to be conservative; the welfare losses from the JA are larger if high freight costs limit product variety in PR. This is likely but difficult to quantify without observing the results of a policy change.

at the origin.[51] Given new values of the price index P_y^1, we calculate an updated value of the expenditure function $E(P_y^1, \overline{U}_y)^1$. The compensating variation of the price change is calculated as

$$CV = E(P_y^0, \overline{U}_y)^0 - E(P_y^1, \overline{U}_y)^1 \tag{11}$$

We do this calculation for JA tariff-equivalent trade costs on US imports, and US tariffs on ROW imports. We also conduct the exercise for sub-components of final demand (Consumption, Investment, Government spending, etc.)

6 Results

We report regression results from a divided sample. One subsample consists of final goods, while the other subsample contains the remaining goods, those further upstream in the supply chain. We take this approach because PR has a highly unusual industrial structure, and it is likely that this has been shaped by the JA. In our primary set of estimates, we use the upstreamness measure of Antràs et al. (2012) to divide products into 'final' and 'upstream' goods. Specifically, we define as final goods all products that belong to an HS6 with an upstreamness index of 1.3 or less, and define products with upstreamness index values greater than 1.3 as upstream goods.[52]

6.1 Reduced form estimates

We report the results of reduced form gravity regressions for final goods in Table 5. All specifications include supply variables defined as above. All specifications also include both logged distance and the square of logged distance.[53] Our focus is on the coefficient on the $HOME$ dummy and on its interactions with the product characteristic variables ($\overrightarrow{Z}$) that predict shippers' choice of transport mode.

---

[51] The assumption of perfectly elastic supply to the Puerto Rican market is highly reasonable for US and ROW imports. Our calculations also assume no change in the prices of Puerto Rican goods. Puerto Rican prices might also be expected to fall with JA removal, since domestic suppliers would face greater competition. Falling domestic prices would raise our estimate of CV. But these additional welfare gains would be partially offset by reduced Puerto Rican income (absent any additional gains arising through comparative advantage). We take the assumption of no net change in Puerto Rican prices as a reasonable approximation that facilitates transparent calculations.

[52] We chose this threshold by inspection before estimating our regressions. Several industries near the 1.3 threshold but below it are clearly household consumption items (wine, apparel, frozen food). There are some industries that produce consumption items above the 1.3 threshold (books, cutlery), but most are products that are less obviously final goods (analytical laboratory instruments, support for oil and gas operations, miscellaneous electrical equipment, etc.).

[53] The second order term is included to allow the effects of distance on trade to vary over distance. The US is much closer to PR than are other developed countries with a similar export mix, so we wish to allow the effects of distance to taper - if the data suggest it - to reduce the chance that the assumption of a constant elasticity of distance biases the $HOME$ dummy coefficient. In subsequent specifications we also allow the effects of distance to vary across products as well as over distance.

Column 1 contains results from a simple specification focusing on the estimation of average home bias. The coefficient on the $HOME$ dummy is 2.22; this is the (cross-product) average effect of the $HOME$ dummy on logged bilateral trade, after controlling for variation in regional supplies and for flexibly-defined effects of distance. A product with the mean response of trade to $HOME$, $\beta = 2.22$ has imports from the US that are approximately $e^{2.22}–1 = 8.21$ times larger than from ROW.

Column 2 includes, as controls, the $\overrightarrow{Z}$ variables that we interact with $HOME_j$ in subsequent regressions. The coefficients in this column tell us whether the product characteristics help predict cross-product variation in the level of import demand for commodity k imports. All coefficients are significant; they jointly indicate that PR's total imports of a final product are relatively larger if the product is typically a) shipped by sea, b) heavier, and c) not containerized. The $HOME_j$ coefficient is basically unchanged when we include the $\overrightarrow{Z}$ variables in the regression.

Column 3 is the first specification that contains our variables of interest, the interaction of the $HOME$ dummy with the product characteristics. All the interaction coefficients are of the hypothesized sign, and all but the interaction of $HOME$ with $(\ln(WV^k))^2$ are statistically significant. Products typically shipped by sea have lower estimated home bias, which is consistent with the JA causing substitution away from US sources among sea-shipped goods. As predicted, home bias is also smaller in heavier products, and in products that are not typically shipped in containers. These results are consistent with the hypothesis that the JA places an even larger burden on products shipped in bulk carriers or general cargo ships. The hypothesized results are maintained in column 4, where we allow the effects of distance to vary across products. Column 5 shows that the results are robust to the inclusion of trade policy variables, including the US Most Favored Nation tariff on product k and a vector of dummy variables indicating that the supplying country is a member of a PTA with the US. The coefficient on the US MFN tariff variable can be interpreted as an estimate of the elasticity of substitution that is common across products.[54]

Results of the same regressions for the sample of *upstream* goods are reported in Table A1 in Appendix A. The column 1 and 2 regressions show a smaller (though still positive) home bias

---

[54] The estimate implies that $\sigma = 2.875$.

among upstream goods. the distance elasticitiy among upstream products is very large. The coefficients on the product characteristics themselves in column 2 show that PR's total imports of upstream products are lower when a product is vessel-shipped, relatively heavy and commonly shipped in containers. As Table 5 shows PR's imports of upstream products move primarily by air, and these estimated coefficients are consistent with most upstream products being air-shipped. This outcome may or may not be an effect of long-run exposure to JA trade costs. The inclusion of the interaction terms in column 3 sees the $HOME_j$ coefficient itself reduced to zero. Against this benchmark the coefficients on the interaction terms indicate negative home bias for vessel-shipped upstream goods, and positive home biases for heavy and container-shipped goods. Column 4 and 5 results are similarly puzzling, while column 5 results also suggest that products with larger US tariffs have *higher* foreign imports.[55] Overall, the results do not suggest a convincing interpretation of how the JA (or trade costs more generally) may have caused PR to substitute away from U.S. sources.[56] We think it likely that PR's idiosyncratic production structure - which is partially determined by high trade costs (as in Hillberry and Hummels 2002) - dominates within-product substitution responses to JA trade costs in determining the trade pattern. In order to be conservative in our assignment of costs to the JA, we hereafter constrain our analysis to estimates from the effects on final goods.

6.2 Structural estimates

We now turn to our structural model. The theoretical model implies that trade responses to geographic frictions can be decomposed into the product of a trade cost parameter and σ^k. In order to quantify trade costs, we incorporate external estimates of σ^k everywhere that a trade friction appears in the econometric model. We also include σ^k alone in the estimation model, in order to control for covariation of σ^k with the level of PR's imports of that commodity. The estimates of σ^k that we use in our primary results come from Fontagné et al. (2022), although we also employ estimates from Soderbery (2015) in a robustness check. Our primary results are

---

[55] Note that the signs on the distance and distance squared terms switch when the trade policy variables are included. This is likely related to the $HOME_j$ coefficient becoming significantly negative in column 5. The interactions of $\overrightarrow{Z}^k$ with distance all take the expected signs in columns 4 and 5, as do the interactions of $HOME_j$ with $Vshr^k_{j,y}$ and $Cntr^k_{j,y}$. The main problems for interpretation are the negative $Home_j$ coefficient (representing the air shipped benchmark) and the interactions of $\overrightarrow{Z}^k$ with the WV terms. These outcomes may be the result of the large share of air shipments in the upstream goods sample.

[56] In unreported results we find a sign pattern in the structural estimates that is similar to what we observe in the reduced form estimates.

reported in Table 6.

The results in column 1 offer a simple example of our method. Pre-multiplying $HOME_j$ by σ^k prior to estimation allows the associated regression coefficient to be interpreted as a measure of structural trade costs. The absence of interaction terms in this regressions means that the estimate $\widehat{\gamma} = 0.237$ is assumed to be common across commodities. Accounting for functional form, this estimate implies a tariff-equivalent trade cost associated with home bias of $\widehat{t}_{HB} = 0.267$.[57] In other words, our estimate is that the various commonalities that PR shares with the US mainland (a common legal system, a common currency, free movement of people, etc.) amounts to an equivalent tariff of 26.7 percent on foreign imports.[58] This estimate of $\widehat{t}_{HB}$ is biased downward because it ignores the counteracting effects of the JA on sea-shipped goods; we include it here as a guide to interpretation.

When we add $\overrightarrow{Z}$ in the regression in Column 2, $\widehat{\gamma}$ is largely unchanged. The estimate grows slightly in column 3, where we include the interaction of the $HOME$ dummy with $\overrightarrow{Z}$. In this case $\widehat{\gamma}$ should be interpreted carefully: the inclusion of interaction terms in the regression means that it no longer captures a cross-product average trade cost, it now represents an estimated cost for a product with particular characteristics. In this case, the coefficient captures the effect of an implied average tariff equivalent of home bias for a good that is air-shipped, is not typically containerized, and has a weight-to-value ratio of 1 (which means that $\ln(WV^k) = 0$). This estimate of $\widehat{t}^k_{HB}$ is somewhat larger than earlier estimates that applied to all commodities, a result that is expected since column 3 results control for the counteracting effects of the JA among sea-shipped goods.

It is the coefficients on the interaction terms that are of greatest interest. They retain the same intuitive sign pattern as in the reduced form regression. The coefficient estimate (-0.206) on $Vsh^k_y \times HOME$ implies that shipping the same product exclusively by vessel rather than exclusively by air implies an increase in $\widehat{t}^k_{JA}$ of $e^{0.206}$. If this coefficient were completely informative about the JA trade cost, it would imply an estimate of $e^{0.206} - 1 = 0.229$, a 23 percent tariff-equivalent effect of the JA. Heavier products have larger implied $\widehat{t}^k_{JA}$'s, while containerized

---

[57] $e^{0.237} - 1 = 0.267$.

[58] As a point of reference, Anderson and Van Wincoop (2004) estimate that the US-Canada border imposes a tariff-equivalent border cost of 47 percent.

shipments face substantially lower tariff-equivalent JA costs. Leaving aside the effects of product weight, full containerization nearly offsets the estimated costs attributed to shipping a product by sea: $e^{0.206-0.162} - 1 = 0.045$.[59]

In columns 4 and 5, we allow for further flexibility in the response of trade to distance, and to US trade policy variables. These change the magnitudes of the coefficients of interest, but the sign patterns remain robust. Since column 5 has the fullest set of controls, we use these results as our primary structural estimates of the distortions caused by the JA. Note that the inclusion of both flexible distance effects and effects of explicit trade policies produces a higher estimate of $\widehat{t}_{HB}$. The $HOME$ dummy coefficient of 0.37 implies $\widehat{t}_{HB}$ of approximately 45 percent (for non-containerized air-shipped goods with weight-to-value ratios of 1). The coefficients on the interaction terms also grew in magnitude, relative to column 3, which implies larger tariff-equivalent estimates of the JA.

The inclusion of continuous variables ($\ln(WV^k)$ and $(\ln(WV^k)^2)$) introduces a wrinkle into the prediction of $\hat{t}^k_{JA}$. The fitted values imply distributions of $\widehat{t}^k_{JA}$ across both air- and sea-shipped goods, with heavier products having relatively larger predicted trade costs for both modes, whether or not goods move by sea. There are two issues: First, one cannot identify an absolute value of $\widehat{t}^k_{JA}$ without benchmarking against a reference air-shipped product (which must have a particular value of WV^k). Second, the implied positive values of $\widehat{t}^k_{JA}$ among air-shipped products arising from variation in WV^k is not something we wish to attribute to the JA. We make two methodological choices to address these issues.

First, we choose as a benchmark a product that has the median weight-to-value ratio among air-shipped products arriving in PR. We calculate this value as 0.0247653 kg/$.[60] The formula for predicted tariff equivalent trade costs thus becomes:

$$\widehat{t}^k_{JA} = e^{-\left[\gamma_{Vsh} Vsh^k_y + \gamma_{WV}\left(ln(WV^k) - ln(0.0247653)\right) + \gamma_{WV2}\left((ln(WV))^2 - (ln(0.0247653))^2\right) + \gamma_{Ctnr} Ctnr^k_y\right]} - 1 \qquad (12)$$

Using (12) we calculate the values of $\widehat{t}^k_{JA}$ using estimates from columns 3-5 of Table 6. The column 5 estimates are our initial estimates of JA trade costs, and we show the distribution of these fitted

---

[59] 0.162 is the coefficient estimate on $Cntr^k_y \times HOME$ in column 3.

[60] Products that fit these criteria and have weight-to-value ratios in the neighborhood of this value are (1) men's suits (made of synthetic fibers, wool or fine animal hair); (2) women's suits (made of artificial fibers) and (3) toasters.

values in Figure 1, which plots the value of $\widehat{t}^{k}_{JA,2016}$'s against each product's weight-to-value ratio.[61] In the figure one can see the role that product weight plays in generating our estimates. As weight to value rises, the implied tariff equivalent rises, but at a decreasing rate. The heaviest products, relative to value, in the consumption sample are types of water (HS220190-*Non-mineral or aerated waters*, and HS220110-*Mineral or aerated waters*); both these products have predicted JA tariff-equivalents of nearly 100 percent. The products with the highest tariff-equivalents have somewhat lower weight-to-value ratios, but have product characteristics that lead them to be less frequently shipped in containers, raising their predicted JA-tariff equivalents. These products are HS110429-*Cereal grains of barley* (with a tariff-equivalent of 114.3 percent), HS200911-*Frozen orange juice* (113.8 percent) and HS870530-*Fire fighting vehicles* (99.0 percent). Nineteen products are not imported at all from the U.S., possibly because the JA makes shipping of these products from the U.S. uneconomical. The implied JA tariff equivalents for products not imported from the US in 2016 are marked separately in Figure 1. The figure also reveals a large number of products with negative predicted tariff-equivalent trade costs. Most of these are light-weight air-shipped products; to illustrate this point we shade data-points in proportion to their dependence on vessel-shipping in US-PR shipments. Some sea-shipped items with very low weight-to-value ratios have negative $\widehat{t}^{k}_{JA}$'s. We treat all negative $\widehat{t}^{k}_{JA}$ estimates as having predicted tariff-equivalent trade costs with zero, rather than negative, values.

The second adjustment to the tariff equivalents relates to air-shipped goods with positive values for $\widehat{t}^{k}_{JA}$. These arise among the air-shipped goods that are heavier than the median air-shipped product. Because we do not wish to assign JA distortions to goods that are air-shipped we undertake a subsequent calculation, $\bar{t}^{k}_{JA,j} = \widehat{t}^{k}_{JA,j} \times Vshr^{k}_{JA,j}$, where $Vshr^{k}_{JA,y}$ is the value share of PR's product k imports arriving from the US in an oceangoing vessel.[62] The values of $\bar{t}^{k}_{JA}$ that follow from this adjustment are our estimated JA trade costs going forward.[63]

---

[61] Each product is assigned a single weight-to-value ratio that is taken from the US data, even though those data stretch over multiple years. WV^{k} is defined as the median value of WV across the sample years. We take the median to minimize the possibility that measurement error in the annual data generates biased estimates in the regressions.

[62] Recall that the $Vshr^{k}_{j}$ used in the regressions comes from U.S. imports, not US-PR flows. We include the subscript JA here to indicate that we are using data from actual US PR flows in this adjustment.

[63] The application of $Vshr^{k}_{JA,y}$ to the predicted trade costs is an effort to be conservative in our estimates. We focus our efforts on quantifying the implicit distortion that causes PR importers to substitute towards rest of the world (ROW) products and away from US products. It is likely that the JA also causes US products to be shipped to PR by air rather than by waterborne transport. Arguably this distortion is evident in the positive JA tariff equivalents that we zero out because shipments - in fact - travel by air rather than by sea. While this is plausible, assigning

In Table 7 we report summary statistics for the distribution of estimates of $\bar{t}^k_{JA}$ for the year 2016. The first three rows contain estimates predicted from the columns 3-5 in Table 6. Our preferred estimates - labelled "All Controls" in Table 7 - come from the estimates in Column 5, which include flexibly defined distance terms and US trade policy controls. In these estimates, the simple average tariff-equivalent estimate of the JA is 30.6 percent, while the trade weighted average is 53.6 percent. 87 percent of final products have a positive JA tariff-equivalent trade cost.

The second row of Table 7 shows estimates of $\bar{t}^k_{JA}$'s calculated with the same methods, but calculated from estimates from Column 4 of Table 6. That regression excludes the trade policy controls. In this set of estimates, the simple average value of $\bar{t}^k_{JA}$ is 35.5 percent. The third row is constructed with estimates from column 3, which did not allow for flexibly defined distance in the regression. This specification produced rather lower values of $\bar{t}^k_{JA}$, with a simple average of 6.4 percent. We prefer the estimates in the first row of Table 7 because they contain the largest set of control variables.

6.3 Robustness exercises

We undertake a number of exercises to check the robustness of our results. Rather than report all of the regression estimates, we focus our reporting on the distributions of $\bar{t}^k_{JA}$ linked to each regression. These values are reported in the bottom half of Table 7. All of the results in Table 7 apply to estimates from a set of structural regressions among a sample of final goods. The general lessons are that a) the qualitative predictions of our hypothesis are robust among final goods, especially with respect to interactions involving products' vessel share and container share of shipments, b) the econometric specification matters for inferences about the sizes of $\bar{t}^k_{JA}$, and c) the values of σ^k are even more important for predictions of $\bar{t}^k_{JA}$.

In our first robustness exercise we divide the data into sub-samples differently, using the United Nations' BEC classification rather than the upstreamness index. We estimate the same empirical model separately for samples of goods the UN categorizes as Consumption, Capital and Intermediate goods, respectively. We report structural estimates for the sub-sample of Consumption goods in Appendix Table A2. In this sample, we find the same sign patterns as in the estimates for final

---

positive JA tariff equivalents to shipments that travel by air risks overstating the economic burden of the JA. We choose to be more conservative and treat goods travelling to PR by air as entirely unaffected by the JA.

goods in Tables 5 and 6, although the coefficients on the interaction of $HOME_j$ with the logged weight to value terms become statistically insignificant in columns 4 and 5. Both reduced form and structural estimates have the predicted sign pattern on the $\overrightarrow{Z} \times HOME_j$ interactions for all specifications involving Consumption goods. The samples of Capital and Intermediate goods have puzzling sign patterns like those we observe for upstream goods in Table A1. We assume that these estimates are also polluted by production location effects, and omit them in the interest of brevity.

The lower coefficients on the $\overrightarrow{Z}_y^k \times HOME_j$ interactions in the BEC Consumption sample imply lower estimates of tariff equivalent trade costs. The simple average estimate of $\bar{t}_{JA}^k$ is 11.3 percent, the trade weighted average 11.2 percent, and the median 14.0 percent.[64] Looking again at the estimates in Table A2 one notes that the sample size is much larger than in the relevant counterpart, Table 6. The BEC sample contains products that are further upstream than the set of final goods in Table 6. In this larger sample, the effects of the $\overrightarrow{Z}_y^k$ variables on predicted home bias are much weaker, which generates the compressed distribution of $\bar{t}_{JA}^k$ in Table 7. We note that the BEC has been criticized for not keeping up with technological change; consumption goods are now sometimes classified as intermediates and intermediates as final goods.[65] We therefore focus our remaining attention on the sample defined by products' position in the upstreamness index.

Returning to the original sample, we estimate a range of different econometric specifications to check robustness. So far, we have controlled for time-varying shocks by assuming they simply affect import demand in the aggregate; the main specification includes year fixed effects. We also estimate the model with year-product fixed effects, which allow for time-varying effects on import demand at the product level. This specification produces coefficient estimates on the interaction terms we study, even though the fixed effects mean that the coefficients on the $\overrightarrow{Z}$ variables alone are not reported because they are collinear with the fixed effects. In both the reduced form and the structural regressions, the sign pattern for the interaction terms is the same as in earlier specifications, though the magnitudes are different. We generate the distribution of imputed $\bar{t}_{JA}^k$'s from the specification with flexible distances and trade policy variables. These are reported in row 5 of Table 7, which shows a simple average $\bar{t}_{JA}^k$ of 11.8 percent and a weighted average of 25.2

---

[64] These lower tariffs would be applied to a greater share of Puerto Rican imports, offsetting the effects of the lower estimated tariffs on our CV calculations.

[65] See Sturgeon and Memedovic (2011) for a discussion of this issue.

percent.[66]

The somewhat lower estimates in this particular robustness check raise the question of which estimates are to be preferred. Normally, one might prefer an estimate from a specification with product-destination fixed effects, which would control for cross-product variation in α^k and $\tilde{P}^k_y$ if the sample also included PR-PR flows. There are two features of these data that lead us to prefer a specification that allows the $\overrightarrow{\omega}\overrightarrow{Z}$ terms to parameterize import demand. First, we lack detailed data on trade flows within PR. The structural parameters α^k and $\tilde{P}^k_y$ are shifters of *total* demand for the product k in PR, rather than shifters of *import* demand. The potential bias arising from this distinction would likely not be especially important if the data for each product contained imports from both US and ROW sources. There are, however, many products for which imports arrive from either the US or the ROW, but not from both regions. Consider the case of imports arriving only from US sources. Suppose a product with relatively high unobserved trade costs from ROW sees imports arrive only from the US. If JA trade costs also cause substitution towards domestic Puerto Rican sources, one will see relatively low values of total imports in this product. A product-destination fixed effect will interpret this outcome as the result of relatively low import demand for that product, rather than a result of high JA trade costs in the presence of higher levels of import demand. The $\overrightarrow{\omega}\overrightarrow{Z}$ terms in the preferred specification are, effectively, a model of PR's product-level import demands. They may or may not predict the level of import demand especially well (though the ω estimates on $Vshr^k_y$ and $Cntr^k_y$ are always highly significant). The relevant point is that the inclusion of the $\overrightarrow{Z}$ variables independently in the regression should produce an estimate of fitted import demand that will not bias downward estimates of JA trade costs in cases where those costs are idiosyncratically high.

Next we check robustness to our choice of Fontagné et al. (2022) as the source of structural estimates of σ^k. Soderbery (2015) produces a set of σ^k estimates for the US using a version of the Feenstra (1994) estimator. We estimate a set of structural regressions akin to those in Table 6, except that the σ^k estimates we include in the regression are Soderbery's, not Fontagné *et al.*'s. The sign pattern in the structural estimates is once again robust, but the magnitudes of the coefficients of

---

[66] In unreported results we estimate using fixed effects to control for demand shocks at the HS2-year rather than HS6-year level. The sign patterns are again consistent, and the distribution of $\bar{t}^k_{JA}$ for column 5 shows a simple average of 33.2 percent and a traded weighted average of 57.2 percent.

interest imply much larger $\hat{t}^k_{JA}$'s. This appears to be a mechanical result that comes from the fact that the Soderbery (2015) estimates of σ^k are generally lower than the estimates in Fontagné et al. (2022).[67] The predicted values of $\bar{t}^k_{JA}$'s implied by the Soderbery (2015) estimates are reported in Table 7. They are an order of magnitude larger than those implied by the Fontagné et al. (2022) estimates, with a simple average tariff equivalent of 310%. These are arguably implausible as *ad valorem* estimates of bilateral trade costs, since they imply that the additional transport costs due to the JA account for 3/4 of the delivered price for the good at the mean of the distribution.[68]

We also consider the implications of using the implied estimate of σ that is the coefficient estimate on the US tariff variable in column 5 of Table 5. That interpretation of the estimate implies that all commodities share the same elasticity of substitution. Since the estimate of $\sigma = 2.785$ is rather low, the implied values of $\overrightarrow{\gamma}$ are rather high, especially for the products most affected by the JA.[69] The mean estimate in this case is a 49.2 percent tariff equivalent. The maximum values are much higher than in the benchmark estimates that use heterogeneous σ^k's. The very high maximum values in the common-σ case likely arise because the products most affected by the JA are also commodities with high elasticities of substitution (e.g. types of water). In this instance, applying an average value of σ to all products biases upward the $\bar{t}^k_{JA}$ estimates for highly substitutable products.

6.4 Aggregating trade costs to NAICS sectors

A proper welfare analysis of the JA requires a model that allows consumers to choose products from the US mainland, ROW and from PR itself. Since we lack detailed data on intra-PR trade, we are unable to do these calculations at the HS6 level for which we have trade statistics. Fortunately, PR produces an input-output table that allows formal welfare analysis. The sectors in the table are defined at the 4-digit level of the North American Industrial Classification System (NAICS). We describe here the process of aggregating the $\tilde{t}^k_{JA}$ estimates to the 4-digit NAICS sector level.

---

[67] We speculate that the reason for this result is that the Feenstra (1994) estimator used by Soderbery (2015) is more reliant on time series variation than are the estimates in Fontagné et al. (2022), which exploit cross-sectional variation in a manner similar to Hummels (1999). Since short-run estimates are likely to be smaller than long-run responses - see Erkel-Rousse and Mirza (2002) - this would explain the discrepancy between the two sets of estimates. The JA is more than a century old, so long-run responses to trade costs are preferable.

[68] Another problem with the Soderbery (2015) estimates for our purposes is that there are many commodities without an estimate of σ^k. In these cases we are still able to estimate implied values of $\widehat{t}^k_{JA}$, by calculating the implied values predicted by the estimated $\overrightarrow{\gamma}$ coefficients and the product characteristics associated with those commodities. We have relatively low levels of confidence in these estimates, however, given the absence of σ^k.

[69] These estimates of $\overrightarrow{\gamma}$ are calculated by dividing the reduced form coefficients $\overrightarrow{\beta}$ (from Table 5) by the estimated value of σ, the coefficient on the MFN tariff in Table 5.

The table contains a square matrix of intermediate expenditures (each NAICS category's expenditures on all the others), a row of each sector's payments to domestic factors, and six columns of final expenditures on each NAICS sector's output. The categories of final expenditure include consumption, exports, and governmental expenditure.[70] Investment encompasses three columns, which we aggregate together.[71] Each cell of the table reports purchases from PR sources and from imports, where reported imports include imports from both domestic and foreign sources. We use a concordance of HS6 products to NAICS sectors provided by the US Census.

For our purposes, we must make three adjustments to the table to support welfare analysis. We need to divide PR's reported "imports" in the table between US and ROW sources. Second, we need to map up- and down-stream products onto NAICS level US imports. Finally, we have to reconcile the fact that the most recent IO table is from 2006-7, while our trade data are available for a later period of time.

We address the first and second problems jointly. In mapping the HS6 data to the NAICS codes we keep track of both the US and ROW share of imports in each NAICS code and the share of up- and down-stream products for each NAICS sector and origin-region. For each NAICS sector, we calculate US and ROW shares of up- and down-stream import value, respectively, and apply the average US and ROW shares for 2010-2017 to the import flows that the table reports for each NAICS code.

These calculations give estimated trade shares that we use in the welfare analysis. We must also generate JA-tariff equivalent trade costs at the NAICS level. Let $s_{k,l}$ be the share of each downstream HS product k in PR's imports from the US in NAICS sector l. We calculate trade-weighted average JA tariff equivalents by multiplying these shares by the values of $\bar{t}^k_{JA}$. Summing over these values within each NAICS sector generates a value of $\bar{t}^l_{JA}$ for each NAICS sector. In

---

[70] The architecture of the table allows us to split government expenditure into three categories: local, state and federal. Each of these subcategories appears as a purchasing sector in the IO table, and each sells its "output" to government purchases in final demand. We restructure the table so that each of the components of government expenditure eneters separately as final demand sector.

[71] The table splits investment into three categories: machinery and equipment, investment in construction, and changes in inventories. Changes in inventories can take positive and negative values, which can lead aggregate investment to be less than zero. We zero out changes in inventories, treating machinery and equipment investment and investment in construction as the only source of final demand expenditures in our calculations.

order to offer a comparison to US tariffs, we conduct a similar exercise that weights US MFN tariffs by shares of downstream imports of product k sector l.[72] We also calculate σ^l as a trade weighted average of σ^k's.

A final set of adjustments that must be made to these figures is to inflate the expenditure values in the 2006-2007 input-output table to 2016 values. In the absence of a more recent input-output table for PR, we must assume that the structure of input-output relationships is unchanged over time. One advantage to limiting our analysis to final goods is that we do not require an assumption that intermediate demands remained unchanged over this period. We are unable to account for changes across sectors l in either a) the share of each NAICS sector in each final demand category's expenditure, b) the relative sizes of final expenditure categories, and c) PR's share of total sales within a NAICS category. We simply adjust the level of nominal expenditure to 2016 by applying observed changes in GDP since 2006.[73] Puerto Rican nominal GDP grew by a factor of 1.195 between 2006 and 2016.

Table 8 provides information that is relevant to our welfare calculations.[74] Column 1 reports the share of each final demand category in total final expenditures. Household consumption accounts for nearly 52.6 percent of total final expenditures. Exports account for an additional 20.4 percent. (We assume that exporting firms purchase only upstream goods so domestic purchases for export do not affect PR's welfare in our calculations.) Investment and the three government spending categories account for just over 27 percent of total final expenditures.

We also calculate the share of U.S. imports for each final expenditure category; investment is the most dependent on US imports. We then report US trade-weighted averages of $\tilde{t}^l_{JA}$ for each expenditure category. The average JA trade cost for all of final demand is 8.0 percent; the value for private consumption is 9.2 percent.[75]

---

[72] Products from regions that are PTA partners with the US are assumed to face zero tariffs in this aggregation.

[73] We do not adjust for inflation, but rather report values in terms of 2016 nominal dollars. Our trade data for 2016 are reported in nominal dollars.

[74] All these calculations exclude the NAICS codes in PR's IO table that contain pharmaceutical production: NAICS 3251, 3254 and 3391. US Census concordances maps HS2 code 30 to these NAICS codes.

[75] Investment has the highest average tariff equivalent costs, 17.1 percent. The high average tariff-equivalent is primarily attributable to the sector's purchases of cars and of other motor vehicles. The large distortions in the Investment category apply to relatively small US import shares, so they are not overly important for our welfare calculations.

6.5 How much does the JA cost PR?

6.5.1 Back-of-the Envelope CV Estimates

Before we move to our CV estimates, we conduct a back-of-the-envelope calculation to provide a transparent illustration of the approximate magnitudes implied by our estimates. We apply average information from our sample to PR's final expenditure levels, in order to make our calculations transparent. Subsequent calculations exploit the rich cross-product heterogeneity present in our estimates. The latter estimates are our estimates of the actual welfare costs of the policy.

The parameter inputs into our aggregate calculation are as follows. With just a single aggregate sector, $\alpha^l = 1$. We calculate a trade-weighted average σ of 3.478.[76] Imports from the US face a JA tariff equivalent of 8 percent *ad valorem*. PR's expenditure shares on its own goods and services and on US and ROW imports are 81.6, 13.7, and 4.7 percent respectively. In the true cost of living index P_j, the associated distribution parameters θ_j for the products of PR and imports from the US and ROW are, respectively, 0.816, 0.166, and 0.049. In this setting the removal of JA trade costs $\bar{t}^k_{JA}$ from US products reduces P_j by 1.2 percent. A 1.2 percent reduction in P_j means that purchasers of final demand in PR could reduce their expenditure by an equivalent percentage, and still maintain their current standard of living. Our expansion of final expenditure puts it at an estimated value of \$107.3 billion in 2016. This implies a welfare cost of the JA of approximately \$1.3 billion in the back-of-the-envelope calculation.

6.5.2 CV Estimates from removing JA tariff equivalent and US tariff.

We apply this same approach to our disaggregated data, with $\bar{t}^k_{JA}$ and σ varying across NAICS sectors l. Using our preferred estimates of $\bar{t}^k_{JA}$, we calculate that final expenditure in Puerto Rico would be \$1.4 billion (about 1.3 percent) lower in 2016 without the JA. When we decompose this value into burdens on particular types of final expenditure, Table 9 indicates that consumption spending would be \$691 million (about 1.2 percent) lower per year, or \$203 per citizen annually. The highest burden is on investment, which could be maintained at existing levels with 3 percent lower expenditures if the JA were removed. In a dynamic framework, this implicit tax on investment would make the JA much more costly to PR.

Using the same approach, but considering the removal of US MFN tariffs for all goods arriving

---

[76] In this calculation, NAICS categories without trade have an assumed σ^l of 1.

from non PTA partners, we calculate that final expenditure could be US$133 million (about 0.1 percent) lower and produce the same level of utility without US tariffs. Table 10 indicates that the distortions caused by US MFN tariffs on all goods from non-PTA member countries cost Puerto Rican households approximately $92 million, or $27 per person per year. the cost of the JA for Puerto Rican households is 7.5 times as large as the cost of remaining US tariffs.

7 Conclusions

Standard international trade theory acknowledges a potential role for trade policy in meeting non-economic objectives such as national defense. The initial motivation for the JA was to ensure a market for domestically owned and produced ships that could serve the country in times of war. Although the merits of the JA as a national defense policy look shaky, it is not our purpose here to assess the national security benefits associated with the policy.

Our objective is to measure the economic burden associated with the policy. Relative to earlier estimates - notably the economy-wide welfare calculations done in various USITC reports - our estimates also highlight that the policy has important distributional consequences. Residents of US controlled islands bear a heavy share of the economic burden of the JA. We study the particular case of PR, which is large enough to collect suitable data for our exercise and close enough to the US to be highly dependent on the mainland for consumption goods.

We use data on ship movements in the Caribbean to characterize key differences between the JA-compliant fleet and other freight shipping in the region. Bulk shipping capacity is notably absent from the JA-compliant fleet, a finding that proves useful in guiding our subsequent approach to estimation. We hypothesize that products that are more likely to be moved by bulk shipping face heavier trade costs associated with the JA.

Noting the unusual industrial structure of PR, we estimate over separate subsamples of the data, which is divided into up- and down-stream products, respectively. In downstream products (those purchased primarily by final sources of demand), we find considerable evidence that trade responds as we hypothesized. Measured home bias is noticeably smaller in products that a) move by sea, b) are not containerized, and c) have high weight to value ratios. Among upstream products, our empirical gravity model finds much weaker implied effects of JA trade costs, an outcome we believe to be related to PR's unusual production structure. In order to be conservative in our estimates, we assume no JA trade costs on upstream goods, and focus our attention on the estimates from final goods.

In order to uncover structural trade cost estimates we incorporate into the estimation external estimates of the elasticity of substitution. These allow us to make inferences about the relative size of trade costs across products (that might also differ in the size of their elasticities of substitution). Using these estimates - and comparing to a reference product that is assumed to have no costs linked to the JA because it is air shipped - we are able to infer absolute trade costs. We calculate these at the product level, and characterize the cross-commodity distribution of tariff-equivalent trade costs associated with the JA.

Finally, we apply data from an input-output table for PR, which allows us to investigate the consequences of spending on local PR goods, to separate final and intermediate expenditures, and to investigate difference in the JA burden across subcategories of final demand. Our estimates suggest that the JA raised the cost of PR's final demand by \$1.4 billion. The estimated burden on final consumption is \$691 million, or approximately \$203 per PR citizen.

Our estimates are conservative in that they assume that the JA imposes no distortion on purchases of upstream goods. Since many upstream products are bulk-shipped this may be important. In our analysis of upstream goods we find no consistent evidence of JA effects on the trade pattern, which is why we assume no distortions. It seems likely that industries that rely on products that would see their input prices substantially affected by JA restrictions may not locate there at all, a behavior that would substantially increase - beyond our estimates - the economic burden of the JA on PR.

References

Advantage Business Consulting (2019). Analysis of Issues Related to Maritime Transportation to Puerto Rico. Technical report, Advantage Business Consulting, San Juan, PR.

Anderson, J. E. and E. Van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle. *American economic review 93*(1), 170–192.

Anderson, J. E. and E. Van Wincoop (2004). Trade costs. *Journal of Economic literature 42*(3), 691–751.

Anderson, J. E. and Y. V. Yotov (2016). Terms of trade and global efficiency effects of free trade agreements, 1990–2002. *Journal of International Economics 99*, 279–298.

Antràs, P., D. Chor, T. Fally, and R. Hillberry (2012). Measuring the upstreamness of production and trade flows. *American Economic Review 102*(3), 412–16.

Beason, S., D. Conner, N. Milonas, and M. Ruge (2015). Myth and conjecture: The cost of the jones act. *J. Mar. L. & Com. 46*, 23.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020). Geography, transportation, and endogenous trade costs. *Econometrica 88*(2), 657–691.

Erkel-Rousse, H. and D. Mirza (2002). Import price elasticities: reconsidering the evidence. *Canadian Journal of Economics/Revue canadienne d'économique 35*(2), 282–306.

FED-NY (2012). Report on the Competitiveness of Puerto Rico´s Economy. Technical report, Federal Reserve Bank of New York, New York.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.

Feliciano, Z. M. and M.-T. Chen (2021). Intangible assets, corporate taxes and the relocation of pharmaceutical establishments: The case of puerto rico. Technical report, National Bureau of Economic Research.

Fontagné, L., H. Guimbard, and G. Orefice (2022). Tariff-based product-level trade elasticities. *Journal of International Economics 137*, 103593.

Francois, J., H. Arce, K. Reinert, and J. E. Flynn (1996). Commercial Policy and the Domestic Carrying Trade: A General Equilibrium Assessment of the Jones Act. *Canadian Journal of Economics 29*(1), 181–198.

GAO (1988). The Jones Act: Impact on Alaska Transportation and U.S. Military Sealift Capability. Technical Report september, U.S. General Accountability Office - GAO, Washington D.C.

GAO (1993). Puerto Rico and the Section 936 Tax Credit. Technical Report GAO/GGD-93-109, U.S. General Accountability Office - GAO, Washington D.C.

GAO (2013). Puerto Rico: Characteristics of the Island's Maritime Trade and Potential Effects of Modifying the Jones Act. Technical Report March, U.S. General Accountability Office - GAO.

Grennes, T. (2017). An Economic Analysis of the Jones Act. Technical report, MERCAUTUS CENTER: George Mason University, Arlington, VA.

Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, Volume 4, pp. 131–195. Elsevier.

Helpman, E. and P. Krugman (1985). Market structure and foreign trade: Increasing return, imperfect competition, and the international economy. *Cambridge. USA.*

Herrero-Rodríguez, J. A., A. L. Soriano Miranda, and J. Valentín-Mari (2003). El Efecto del Régimen Actual del Comercio Exterior, en Relación al Transporte Marítimo, sobre la Economía de Puerto Rico. Technical report, Oficina de Asesores del Presidente de la Camara de Representantes de Puerto Rico.

Hillberry, R. and D. Hummels (2002). Explaining home bias in consumption: The role of intermediate input trade.

Hillberry, R. and D. Hummels (2008). Trade responses to geographic frictions: A decomposition using micro-data. *European Economic Review 52*(3), 527–550.

Hummels, D. L. (1999). Toward a geography of trade costs. *Available at SSRN 160533*.

Hummels, D. L. and G. Schaur (2013). Time as a trade barrier. *American Economic Review 103*(7), 2935–59.

John Dunham & Associates (2019). The Jones Act a Legacy of Economic Ruin for Puerto Rico. Technical report, John Dunham & Associates, San Juan, PR.

Kashian, R., J. Pagel, and I. Brannon (2017). *The Jones Act in perspective: A survey of the costs and effects of the 1920 Merchant Marine Act* (First ed.). Number 808. Honolulu, HI: Grassroot Institute of Hawaii.

Marein, B. (2020). Economic development in puerto rico after us annexation: Anthropometric evidence. *Economics & Human Biology 38*, 100892.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica 71*(6), 1695–1725.

Merchant Marine Act, 46 U.S.C. §55102 (1920).

Olney, W. W. (2020). Cabotage sabotage? the curious case of the jones act. *Journal of International Economics 127*, 103378.

Page, P. and A. Restuccia (2022). Biden administration grants waiver allowing diesel shipment into puerto rico. *The Wall Street Journal - WSJ*.

Pesquera, P. (1965). *Algunos Problemas que Confronta la Transportación Marítima entre Estados Unidos y Puerto Rico y sus Implicaciones para la Economía de Puerto Rico*. Master thesis, University of Puerto Rico.

Quinonez-Dominguez, M. (1990). *Las Leyes de Cabotaje: Un Análisis de Como Afecta los Precios de Algunos Productos Alimenticios*. Master, University of Puerto Rico.

Soderbery, A. (2015). Estimating import supply and demand elasticities: Analysis and implications. *Journal of International Economics 96*(1), 1–17.

Sturgeon, T. J. and O. Memedovic (2011). Mapping Global Value Chains: Intermediate Goods Trade and Structural Change in the World Economy. Technical report, Vienna.

Taylor, M. S. (2021). Trade, competitive exclusion, and the slow-motion extinction of the southern resident killer whales. Technical report, National Bureau of Economic Research.

United Nations (2022). Clasifications of economic statistics. https://unstats.un.org/unsd/classifications/Econ.

USITC (1991). The Economic Effects of Significant U.S. Import Restraints, Phase III: Services. Technical report, United States International Trade Commission - USITC, Washington D.C.

USITC (1993). The Economic Effects of Significant U.S. Import Restraints. Technical report, United States International Trade Commission - USITC, Washington D.C.

USITC (1999). The Economic Effects of Significant U.S. Import Restraints - Second Update. Technical report, United States International Trade Commission - USITC, Washington D.C.

USITC (2002). The Economic Effects of Significant U.S. Import Restraints - Third Update. Technical report, United States International Trade Commission - USITC, Washington D.C.

USITC (2018). Tariff - Dataweb.usitc.gov. Accessed: October, 2018.

Valentin-Mari, J. and J. Alameda-Lozada (2012). Economic Impact of Jones Act on Puerto Rico's Economy.

Yotov, Y. V., R. Piermartini, M. Larch, et al. (2016). *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*. WTO iLibrary.

8 Tables

Table 1: Economic Indicators of Puerto Rico's Economy per Sector

NAICS code	Sector	Sales, value of shipments, or revenue ($1,000)	%	Number of employees	%	Annual payroll ($1,000)	%
31-33	Manufacturing	85,263,498	45.2%	77,005	11.8%	2,854,815	17.3%
44-45	Retail trade	25,456,248	13.5%	132,033	20.3%	2,368,082	14.3%
42	Wholesale trade	19,110,386	10.1%	31,255	4.8%	1,114,254	6.7%
51	Information	14,705,619	7.8%	19,872	3.0%	841,980	5.1%
52	Finance and insurance	14,357,432	7.6%	31,928	4.9%	1,444,679	8.8%
62	Health care and social assistance	7,885,285	4.2%	84,933	13.0%	2,085,097	12.6%
54	Professional, scientific, and technical services	4,481,246	2.4%	34,559	5.3%	1,231,711	7.5%
72	Accommodation and food services	4,313,196	2.3%	82,815	12.7%	1,084,616	6.6%
56	Waste management and remediation services	3,197,260	1.7%	74,461	11.4%	1,445,280	8.8%
23	Construction	2,363,926	1.3%	20,215	3.1%	450,349	2.7%
48-49	Transportation and warehousing	2,345,789	1.2%	15,626	2.4%	424,650	2.6%
53	Real estate and rental and leasing	1,617,282	0.9%	13,101	2.0%	322,424	2.0%
81	Other services (except public administration)	1,106,981	0.6%	13,125	2.0%	250,707	1.5%
55	Management of companies and enterprises	1,040,588	0.6%	6,709	1.0%	342,976	2.1%
22	Utilities	626,611	0.3%	350	0.1%	21,840	0.1%
61	Educational services	418,303	0.2%	9,443	1.4%	149,688	0.9%
71	Arts, entertainment, and recreation	299,586	0.2%	3,811	0.6%	67,179	0.4%
21	Mining, quarrying, and oil and gas extraction	42,785	0.0%	482	0.1%	8,777	0.1%
	Total	188,632,020	100.0%	651,719	100.0%	16,509,101	100.0%

Source: U.S. Economic Census of Islands Area (2017).
Note: All figures are in U.S. current dollars, except by the number of employees. This figure corresponds to the actual number workers employed in every sector.

Table 2: Economic Indicators of Puerto Rico's Manufacturing Sector

NAICS code	Sector	Sales, value of shipments, or revenue ($1,000)	%	Number of employees	%	Annual payroll ($1,000)	%
3254	Pharmaceutical and medicine manufacturing	55,065,220	64.6%	977,279	34.2%	13,661	17.7%
3391	Medical equipment and supplies manufacturing	12,605,175	14.8%	676,710	23.7%	16,725	21.7%
3121	Beverage manufacturing	3,647,960	4.3%	87,190	3.1%	2,294	3.0%
3119	Other food manufacturing	1,143,133	1.3%	63,490	2.2%	1,997	2.6%
3353	Electrical equipment manufacturing	961,788	1.1%	97,711	3.4%	2,816	3.7%
3152	Cut and sew apparel manufacturing	478,824	0.6%	84,220	3.0%	5,862	7.6%
3118	Bakeries and tortilla manufacturing	475,038	0.6%	106,540	3.7%	6,316	8.2%
3256	Soap, cleaning compound, and toilet preparation manufacturing	452,632	0.5%	44,713	1.6%	1,016	1.3%
3345	Navigational, measuring, electromedical, and control instruments manufacturing	354,382	0.4%	77,901	2.7%	2,269	2.9%
3261	Plastics product manufacturing	291,499	0.3%	46,125	1.6%	1,689	2.2%
	Others	9,787,847	11.5%	592,936	20.8%	22,360	29.0%
	Total	85,263,498	100.0%	2,854,815	100.0%	77,005	100.0%

Source: U.S. Economic Census of Islands Area (2017).
Note: All figures are in U.S. current dollars, except by the number of employees. This figure corresponds to the actual number workers employed in every sector.

Table 3: Type of Vessels Serving the Caribbean (2019)

Panel A. All vessels in the Caribbean

Type of Vessel	Number of vessels	%	Number of Calls	%	DWT	%
Tanker	1,197	37.9%	8,600	24.0%	83,057,831	50.6%
Bulk	678	21.5%	2,434	6.8%	32,784,710	20.0%
Container ships	535	17.0%	11,650	32.6%	30,111,934	18.4%
General cargo	465	14.7%	6,651	18.6%	5,418,278	3.3%
Others	260	8.2%	5,288	14.8%	12,677,009	7.7%
No DWT reported	20	0.6%	1,165	3.3%		0.0%
Total	3,155	100%	35,788	100%	164,049,762	100%

Panel B. All vessels serving Puerto Rico

Type of Vessel	Number of vessels	%	Number of Calls	%	DWT	%
Tanker	189	52.1%	321	16.0%	9,330,379	66.4%
Container ships	50	13.8%	840	41.7%	1,042,222	7.4%
Bulk	48	13.2%	103	5.1%	1,999,501	14.2%
General cargo	41	11.3%	287	14.3%	467,487	3.3%
Others	30	8.3%	251	12.5%	1,205,813	8.6%
No DWT reported	5	1.4%	210	10.4%		0.0%
Total	363	100.0%	2,012	100.0%	14,045,402	100.0%

Panel C. Jones Act vessels

Type of Vessel	Number of vessels	%	Number of Calls	%	DWT	%
Container ships	4	44.4%	197	82.1%	118,949	80.8%
Barges, ferries, etc	3	33.3%	3	1.3%	28,073	19.1%
Others	1	11.1%	39	16.3%	163	0.1%
No DWT reported	1	11.1%	1	0.4%		0.0%
Total	9	100.0%	240	100.0%	147,185	100.0%

Note: DWT refers to the vessels' Deadweight Tonnage. In each case, the reported DWT corresponds to the sum of vessels' DWT per type of vessel. Number of Calls indicates the number of times a vessel stopped in a port. The difference between All vessels serving Puerto Rico and Jones Act vessels is that the latter are those that satisfy Jones Act conditions. In order to generate this table, the year of a shipping movement corresponds to the calendar year in which a vessel arrived at a port. No DWT reported correspond to vessels for which the DWT information is not available. In Panel C. the vessel without reported DWT is a barge.

Table 4: Composition of Puerto Rico's Imports - 2016

Type of Goods	Mode	Region	Import value	%
Final goods	Sea	USA	2,825,407,995	6.9%
		ROW	706,846,423	1.7%
	Air	USA	350,950,163	0.9%
		ROW	42,589,807	0.1%
Upstream goods	Sea	USA	7,814,495,558	19.0%
		ROW	5,799,527,269	14.1%
	Air	USA	13,321,970,188	32.4%
		ROW	10,279,286,376	25.0%
Total			41,141,073,779	100.0%

Note: Both panels show Puerto Rico's imported value and share for every combination, using the upstreamness index (UI) of Antràs et al. (2012) to classify the products between final (UI<=1.3) and upstream goods (UI>1.3).

Table 5: Reduced Form Gravity Estimates for Downstream (Final) Goods

	(1)	(2)	(3)	(4)	(5)
VARIABLES			$M^k_{j,y}$		
$\ln(dist_j)$	-1.593*	-2.206**	-2.873***	-5.415***	-4.599***
	(0.818)	(0.873)	(0.863)	(1.060)	(1.130)
$HOME$	2.220***	2.203***	1.688***	2.425***	1.739***
	(0.0639)	(0.0583)	(0.484)	(0.602)	(0.602)
Vsh^k_y		0.658***	2.304***	-10.50***	-11.21***
		(0.207)	(0.310)	(1.950)	(1.855)
$\ln(WV^k)$		0.392***	0.823***	1.112***	0.975***
		(0.0626)	(0.101)	(0.343)	(0.365)
$(\ln(WV^k))^2$		0.0134***	0.0121	-0.138***	-0.156***
		(0.00470)	(0.0280)	(0.0502)	(0.0520)
$Ctnr^k_y$		-1.060***	-1.935***	-0.171	1.165
		(0.138)	(0.237)	(1.960)	(2.012)
$Vsh^k_y \times HOME$			-1.679***	-2.193***	-1.538***
			(0.455)	(0.507)	(0.498)
$\ln(WV^k) \times HOME$			-0.674***	-0.718***	-0.617***
			(0.133)	(0.113)	(0.114)
$(\ln(WV^k))^2 \times HOME$			-0.0144	-0.0380***	-0.0260**
			(0.0283)	(0.0109)	(0.0115)
$Ctnr^k_y \times HOME$			1.190***	1.019***	0.722***
			(0.345)	(0.221)	(0.216)
$Vsh^k_y \times \ln(dist_j)$				1.686***	1.696***
				(0.275)	(0.259)
$\ln(WV^k) \times \ln(dist_j)$				-0.0324	-0.0273
				(0.0464)	(0.0485)
$(\ln(WV^k))^2 \times \ln(dist_j)$				0.0217***	0.0226***
				(0.00571)	(0.00584)
$Ctnr^k_y \times \ln(dist_j)$				-0.199	-0.328
				(0.252)	(0.258)
$\ln(1 + tar^k_y)$					-2.785***
					(0.788)
$(\ln(dist_j))^2$	0.0136	0.0535	0.0956	0.151**	0.0964
	(0.0576)	(0.0609)	(0.0602)	(0.0671)	(0.0741)
$IHST(\widetilde{X}^k_{j,y})$	0.679***	0.641***	0.643***	0.647***	0.651***
	(0.0308)	(0.0248)	(0.0249)	(0.0259)	(0.0263)
Constant	11.22***	15.25***	17.73***	33.48***	31.08***
	(2.781)	(2.990)	(2.806)	(3.969)	(4.098)
Observations	1,075,452	1,075,452	1,075,452	1,075,452	1,070,496
Year FE	YES	YES	YES	YES	YES
U.S. PTA's Dummy variables	NO	NO	NO	NO	YES
Pseudo R2	0.494	0.512	0.518	0.521	0.524
Average Weight to Value	0.121	0.121	0.121	0.121	0.121
Average Distance to USA (km)	3,701	3,701	3,701	3,701	3,701

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note: Estimates over the sample of HS6 products with values of the upstreamness index <= 1.3. The LHS variable in all models is the $M^k_{j,y}$, the total value imported in Puerto Rico i from place of origin j of product k in year y. All models are estimated using the PPML estimator on Puerto Rico's import data pooled across years, HS6 digit products and places of origin, with year fixed effects included in the estimation. Estimated coefficients correspond to the overall gravity effect, denoted above as $\bar{\beta}$. Pharmaceutical products are excluded from the estimation. IHST denotes the Inverse Hyperbolic Sine Transformation.

Table 6: Structural Gravity Estimates for Downstream (Final) Goods

	(1)	(2)	(3)	(4)	(5)
VARIABLES			$M^k_{j,y}$		
σ^k	-0.0722	0.115	0.215	0.587	0.225
	(0.352)	(0.359)	(0.388)	(0.407)	(0.473)
$\sigma^k \times \ln(dist_j)$	0.133	0.0670	0.0221	-0.0829	0.0332
	(0.106)	(0.109)	(0.118)	(0.132)	(0.152)
$\sigma^k \times HOME$	0.237***	0.230***	0.261***	0.452***	0.370***
	(0.0112)	(0.00863)	(0.0363)	(0.129)	(0.127)
Vsh^k_y		1.119***	3.282***	2.354***	2.319***
		(0.237)	(0.329)	(0.607)	(0.607)
$\ln(WV^k)$		0.482***	0.676***	0.101	0.0658
		(0.0563)	(0.0598)	(0.207)	(0.209)
$(\ln(WV^k))^2$		0.0237***	0.0538***	0.0275	0.0245
		(0.00441)	(0.00640)	(0.0219)	(0.0220)
$Ctnr^k_y$		-1.724***	-2.813***	-1.168***	-1.041***
		(0.158)	(0.190)	(0.350)	(0.358)
$\sigma^k \times Vsh^k_y \times HOME$			-0.206***	-0.590***	-0.528***
			(0.0304)	(0.125)	(0.122)
$\sigma^k \times \ln(WV^k) \times HOME$			-0.0537***	-0.111***	-0.106***
			(0.00910)	(0.0248)	(0.0245)
$\sigma^k \times (\ln(WV^k))^2 \times HOME$			-0.00522***	-0.00403	-0.00407
			(0.000695)	(0.00518)	(0.00467)
$\sigma^k \times Ctnr^k_y \times HOME$			0.162***	0.301***	0.281***
			(0.0300)	(0.0460)	(0.0496)
$\sigma^k \times Vsh^k_y \times \ln(dist_j)$				0.0581***	0.0505***
				(0.0188)	(0.0185)
$\sigma^k \times \ln(WV^k) \times \ln(dist_j)$				0.0149***	0.0146***
				(0.00492)	(0.00493)
$\sigma^k \times (\ln(WV^k))^2 \times \ln(dist_j)$				0.000246	0.000291
				(0.000785)	(0.000733)
$\sigma^k \times Ctnr^k_y \times \ln(dist_j)$				-0.0366***	-0.0356***
				(0.00767)	(0.00818)
$\sigma^k \times \ln(1 + tar^k_y)$					-0.548***
					(0.145)
$\sigma^k \times (\ln(dist_j))^2$	-0.0188**	-0.0140*	-0.0107	-0.00416	-0.0116
	(0.00804)	(0.00813)	(0.00872)	(0.00902)	(0.0105)
$IHST(\widetilde{X}^k_{j,y})$	0.677***	0.612***	0.606***	0.614***	0.616***
	(0.0333)	(0.0239)	(0.0232)	(0.0238)	(0.0242)
Constant	1.181**	3.765***	3.084***	1.633**	1.519*
	(0.566)	(0.430)	(0.450)	(0.832)	(0.827)
Observations	1,075,452	1,075,452	1,075,452	1,075,452	1,070,496
Year FE	YES	YES	YES	YES	YES
U.S. PTAs Dummy Variables	NO	NO	NO	NO	YES
Pseudo R2	0.451	0.479	0.489	0.497	0.499
Average Weight to Value	0.121	0.121	0.121	0.121	0.121
Average Distance to USA (km)	3,701	3,701	3,701	3,701	3,701

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Estimates over the sample of HS6 products with values of the upstreamness index <= 1.3. The LHS variable on all models is $M^k_{j,y}$, the total value imported in Puerto Rico i from place of origin j of product k in year y. All models are estimated using the PPML estimator on Puerto Rico's import data pooled across observations at the year, HS6 digit product, and place-of-origin level, with year fixed effects included in the estimation model. σ^k estimates from Fontagné et al. (2022) are interacted with geographic frictions, and enter separately in the regression itself. Pharmaceutical products are excluded from the estimation. IHST denotes the Inverse Hyperbolic Sine Transformation.

Table 7: Summary Statistics of Estimated Jones Act Tax Equivalent - 2016

Control of PR's Imports DA	# Obs.	Simple Average	Trade Weighted Average	Std. Dev	Minimum	Perc. 25%	Median	Perc. 75%	Maximum
Table 6 estimates									
FE: YEAR - All controls	609	30.6%	53.6%	23.8%	0.0%	8.4%	29.5%	46.5%	114.3%
FE: YEAR - No τ + No FTA's	609	35.5%	61.9%	27.4%	0.0%	9.8%	34.8%	54.0%	130.8%
FE: YEAR - No τ + No FTA's + No Dist $\times \overrightarrow{Z}$	609	6.4%	14.2%	5.9%	0.0%	1.8%	5.4%	9.1%	32.2%
Robustness									
FE: YEAR - All controls - BEC	1,099	11.3%	11.2%	5.2%	0.0%	8.9%	14.0%	14.7%	24.0%
FE: YEAR × product (HS6) - All controls	609	11.8%	25.2%	13.4%	0.0%	0.8%	7.4%	17.4%	80.7%
FE: YEAR × sector (HS2) - All controls	609	33.2%	57.2%	25.9%	0.0%	9.0%	32.1%	50.4%	128.5%
FE: YEAR - All controls + Soderbery σ	609	310.0%	555.3%	223.0%	0.0%	91.9%	335.8%	483.3%	917.9%
FE: YEAR - All controls + Common σ	609	49.2%	76.0%	42.6%	0.0%	10.8%	43.6%	76.0%	226.7%

Note: These statistics are calculated as the product of the predicted JA tariff equivalent for 2016 (calculated by equation (12)) and the vessel share of 2016 US-PR shipments in the corresponding product. Estimates are reported for numerous specifications of the structural regression (in equation (9)). All estimates rely on the σ estimates of Fontagné et al. (2022), except the row labeled "Soderbery σ" (for which we use Soderbery (2015)) and the common's σ case for which we use the coefficient on the US MFN tariff from Table 5. The number of observations when we use the UN BEC classification is greater, because this classification defines more products as consumption goods.

Table 8: Jones Act Tax Equivalent per Puerto Rico's Demand Component - 2016

	Share in Puerto Rico's Final Expenditure	Average Share of U.S. in PR Final Expenditure (per NAICS code)	JA Tariff Equivalent		U.S. tariff
			Heterogeneous $\widehat{\sigma}^k$	Common $\widehat{\sigma}^k$	
Final Demand	100.0%	16.8%	8.0%	12.1%	1.7%
Consumption	52.6%	14.4%	9.2%	14.4%	2.3%
Exports	20.4%	24.7%	4.7%	8.1%	1.6%
Investment	12.7%	29.8%	17.1%	22.6%	1.0%
Local Government	10.5%	3.4%	0.0%	0.0%	0.0%
State Government	2.6%	0.5%	0.1%	0.1%	0.0%
Federal Government	1.3%	3.0%	0.0%	0.0%	0.2%

Note: All estimates are calculated as trade-weighted averages of the estimated JA tariff equivalents summarized in Table 7. Heterogeneous $\widehat{\sigma}^k$ estimates represent JA tariff equivalent retrieved using the σ estimates of Fontagné et al. (2022) for all products. Common $\widehat{\sigma}^k$ estimates represent JA tariff equivalent that assume a common σ for all products. The U.S. tariff is also calculated as a weighted average of the tariffs for every NAICS code using as weights the share of the non-U.S. countries in the expenditure of every NAICS code sector. All calculations exclude NAICS codes 3251, 3254 and 3391 because these represent pharmaceutical products.

Table 9: Compensating Variation - Jones Act Distortion assuming Heterogeneous $\widehat{\sigma}^k$ - 2016

	Share in Final Demand	Total Value (U.S. million of 2016)	Value per capita (U.S. of 2016)	% Change vs No JA
Final Demand	100.0%	1,390	408	1.3%
Consumption	52.6%	691	203	1.2%
Exports	20.4%	289	85	1.3%
Investment	12.7%	403	118	3.0%
Local Government	10.5%	0	0	0.0%
State Government	2.6%	0	0	0.0%
Federal Government	1.3%	0	0	0.0%

Note: Compensating variation estimates are calculated using the estimated weighted average of the estimated JA tariff equivalent for every NAICS code. Specifically, it is calculated as the difference between the observed expenditure and the expenditure required to produce the same utility with the JA tariff equivalent tradecosts removed. All figures are in 2016 dollars; per capita estimates use PR's population in 2016. All calculations exclude NAICS codes 3251, 3254 and 3391 because these represent pharmaceutical products.

Table 10: Compensating Variation - U.S. Tariffs - 2016

	Share in Final Demand	Compensating Variation (U.S. million of 2016)	CV per capita (U.S. of 2016)	% Change vs No tariffs
Final Demand	100.0%	133	39	0.1%
Consumption	52.6%	92	27	0.2%
Exports	20.4%	34	10	0.2%
Investment	12.7%	7	2	0.1%
Local Government	10.5%	0	0	0.0%
State Government	2.6%	0	0	0.0%
Federal Government	1.3%	0	0	0.0%

Note: Compensating variation estimates are calculated as a weighted average using the average U.S. tariff for every NAICS code and the U.S. share in PR imports. Specifically, it is calculated as the difference between the observed expenditure and the predicted in a world without U.S. tariffs. All figures are in 2016 dollars and per capita using the population level from 2016. All calculations exclude NAICS codes 3251, 3254 and 3391, because they compile pharmaceutical products.

9 Figures

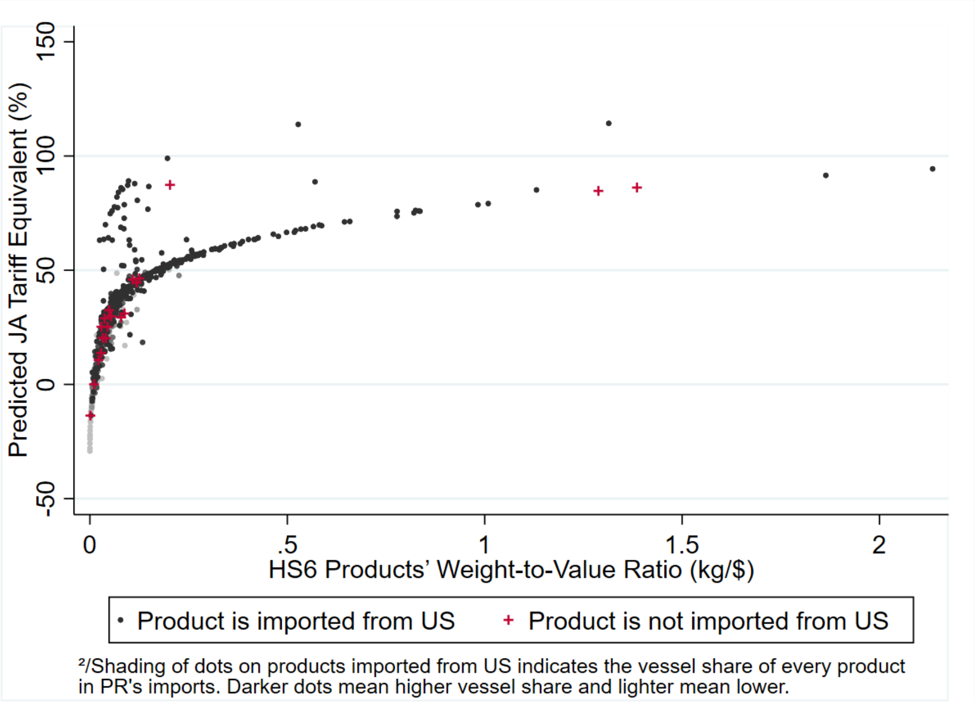

Figure 1: 2016 Jones Act Tax Equivalent estimates against Weight-to-Value Ratio

Note: Estimates of predicted HS6 product-level JA tariff-equivalents for 2016. All estimates are predicted by applying equation 12 and multiplying by 100. Predicted JA tariff-equivalents rely on parameter estimates from Column 5 of Table 6 and σ^k estimates from Fontagné et al. (2022). Product level weight-to-value ratios are calculated as the median of the weight-to-value ratios calculated among U.S. imports (net of imports from Canada and Mexico) for the years 2010-2017.

Appendixes

A Other regression results

Table A1: Reduced Form Gravity Estimates for Upstream Goods

	(1)	(2)	(3)	(4)	(5)
VARIABLES			$M^k_{j,y}$		
$\ln(dist_j)$	-4.047***	-2.995**	-2.649**	-2.668**	4.583***
	(1.087)	(1.201)	(1.257)	(1.224)	(1.582)
$HOME$	0.921***	0.774***	-0.0209	-0.0641	-1.237***
	(0.127)	(0.144)	(0.397)	(0.397)	(0.416)
Vsh^k_y		-1.282***	-0.989***	-4.546**	-13.28***
		(0.312)	(0.368)	(1.899)	(2.348)
$\ln(WV^k)$		-0.216***	-0.399***	2.334***	2.897***
		(0.0399)	(0.0692)	(0.368)	(0.484)
$(\ln(WV^k))^2$		-0.0476***	-0.0719***	0.270***	0.288***
		(0.00645)	(0.0101)	(0.0350)	(0.0440)
$Ctnr^k_y$		-0.465***	-1.737***	8.865***	12.21***
		(0.105)	(0.128)	(1.344)	(1.286)
$Vsh^k_y \times HOME$			-1.211***	-0.987**	-0.318
			(0.377)	(0.389)	(0.408)
$\ln(WV^k) \times HOME$			0.338***	0.183***	0.146**
			(0.0694)	(0.0582)	(0.0669)
$(\ln(WV^k))^2) \times HOME$			0.0531***	0.0361***	0.0369***
			(0.00917)	(0.00836)	(0.00913)
$Ctnr^k_y \times HOME$			3.598***	2.798***	2.854***
			(0.225)	(0.206)	(0.221)
$Vsh^k_y \times ln(dist_j)$				0.422**	1.426***
				(0.208)	(0.259)
$\ln(WV^k) \times ln(dist_j)$				-0.326***	-0.394***
				(0.0463)	(0.0604)
$(\ln(WV^k))^2) \times ln(dist_j)$				-0.0412***	-0.0440***
				(0.00511)	(0.00618)
$Ctnr^k_y \times ln(dist_j)$				-1.235***	-1.658***
				(0.164)	(0.156)
$\ln(1 + tar^k_{j,y})$					10.46***
					(0.569)
$(\ln(dist_j))^2$	0.198***	0.132*	0.112	0.101	-0.420***
	(0.0704)	(0.0775)	(0.0809)	(0.0794)	(0.106)
$IHST(\widetilde{X}^k_{j,y})$	0.864***	0.872***	0.868***	0.877***	0.902***
	(0.0237)	(0.0307)	(0.0311)	(0.0302)	(0.0353)
Constant	16.19***	13.02***	11.87**	12.55***	-11.39*
	(4.165)	(4.594)	(4.816)	(4.717)	(6.109)
Observations	6,174,735	6,174,735	6,174,735	6,174,735	6,174,735
Year FE	YES	YES	YES	YES	YES
U.S. PTAs Dummy Variables	NO	NO	NO	NO	YES
Pseudo R2	0.486	0.503	0.514	0.520	0.563
Average Weight to Value	0.517	0.517	0.517	0.517	0.517
Average Distance to USA (km)	3,701	3,701	3,701	3,701	3,701

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note: Estimates over the sample of HS6 products with values of the upstreamness index > 1.3. The LHS variable in all models is the $M^k_{j,y}$, the total value imported in Puerto Rico i from place of origin j of product k in year y. All models are estimated using the PPML estimator on Puerto Rico's import data pooled across years, HS6 digit products and places of origin, with year fixed effects included in the estimation. Pharmaceutical products are excluded from the estimation. IHST denotes the Inverse Hyperbolic Sine Transformation.

Table A2: Structural Gravity Estimates for BEC Consumption Goods

	(1)	(2)	(3)	(4)	(5)
VARIABLES			$M^k_{j,y}$		
σ^k	-2.020***	-1.847***	-1.713***	-1.618***	-1.386***
	(0.258)	(0.276)	(0.258)	(0.286)	(0.309)
$\sigma^k \times \ln(dist_j)$	0.683***	0.642***	0.599***	0.563***	0.508***
	(0.0702)	(0.0748)	(0.0705)	(0.0775)	(0.0842)
$\sigma^k \times HOME$	0.232***	0.236***	0.263***	0.328***	0.351***
	(0.00916)	(0.00912)	(0.0209)	(0.0641)	(0.0558)
Vsh^k_y		-1.929***	1.241***	5.205***	5.405***
		(0.391)	(0.378)	(1.217)	(1.162)
$\ln(WV^k)$		0.231***	0.413***	0.446***	0.308**
		(0.0328)	(0.0501)	(0.120)	(0.126)
$(\ln(WV^k))^2$		0.0182***	0.0283*	-0.00973	-0.0180
		(0.00433)	(0.0149)	(0.0293)	(0.0295)
$Ctnr^k_y$		1.916***	-0.125	-4.570***	-4.401***
		(0.330)	(0.250)	(0.970)	(0.949)
$\sigma^k \times Vsh^k_y \times HOME$			-0.434***	-0.200*	-0.209*
			(0.0706)	(0.115)	(0.113)
$\sigma^k \times \ln(WV^k) \times HOME$			-0.0407***	-0.0315**	-0.00749
			(0.00580)	(0.0123)	(0.0136)
$\sigma^k \times (\ln(WV^k))^2 \times HOME$			-0.00346***	-0.00692***	-0.00452***
			(0.000982)	(0.00105)	(0.00109)
$\sigma^k \times Ctnr^k_y \times HOME$			0.357***	0.0770	0.0483
			(0.0642)	(0.0952)	(0.100)
$\sigma^k \times Vsh^k_y \times \ln(dist_j)$				-0.0955***	-0.0946***
				(0.0309)	(0.0287)
$\sigma^k \times \ln(WV^k) \times \ln(dist_j)$				-0.00255	-0.00389
				(0.00246)	(0.00261)
$\sigma^k \times (\ln(WV^k))^2 \times \ln(dist_j)$				0.000630**	0.000435
				(0.000313)	(0.000316)
$\sigma^k \times Ctnr^k_y \times \ln(dist_j)$				0.106***	0.106***
				(0.0254)	(0.0245)
$\sigma^k \times \ln(1 + tar^k_{j,y})$					-1.390***
					(0.165)
$\sigma^k \times \ln(dist_j))^2$	-0.0568***	-0.0544***	-0.0515***	-0.0504***	-0.0467***
	(0.00474)	(0.00502)	(0.00477)	(0.00519)	(0.00565)
$IHST(\widetilde{X}^k_{j,y})$	0.489***	0.493***	0.485***	0.485***	0.476***
	(0.0115)	(0.0114)	(0.0107)	(0.0105)	(0.0103)
Constant	4.613***	4.923***	4.405***	4.981***	4.544***
	(0.219)	(0.205)	(0.331)	(0.533)	(0.459)
Observations	1,727,103	1,727,103	1,727,103	1,727,103	1,719,144
Year FE	YES	YES	YES	YES	YES
U.S. PTAs Dummy Variables	NO	NO	NO	NO	YES
Pseudo R2	0.395	0.400	0.408	0.411	0.418
Average Weight to Value	0.244	0.244	0.244	0.244	0.244
Average Distance to USA (km)	3,701	3,701	3,701	3,701	3,701

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note: Estimates over a sample of goods defined by the UN BEC "Consumption" classification. The LHS variable on all models is $M^k_{j,y}$, the total value imported in Puerto Rico i from place of origin j of product k in year y. All models are estimated using the PPML estimator on Puerto Rico's import data pooled across observations at the year, HS6 digit product, and place-of-origin level, with year fixed effects included in the estimation model. σ^k estimates from Fontagné et al. (2022) are interacted with geographic frictions, and enter separately in the regression itself. Pharmaceutical products are excluded from the estimation. IHST denotes the Inverse Hyperbolic Sine Transformation.