DISCUSSION PAPER SERIES

DP17118
(v. 2)

Triangle Inequalities in International Trade: The Neglected Dimension

Reto Foellmi, Christian Hepenstrick and David Torun

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

CEPR

ISSN 0265-8003

Triangle Inequalities in International Trade: The Neglected Dimension

Reto Foellmi, Christian Hepenstrick and David Torun

Discussion Paper DP17118
First Published 18 March 2022
This Revision 23 March 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Triangle Inequalities in International Trade: The Neglected Dimension

Abstract

The so-called triangle inequality (TI) in international trade should, theoretically, hold for any three countries to avoid cross-border arbitrage. When trade costs change, re-routing opportunities—as captured by the TI—might arise because a shipment through an intermediary becomes cheaper under adjusted trade costs. We show that the widely-used "exact hat algebra" approach, which does not require a calibration of trade costs, is unable to measure potential gains from re-routing. In addition, we show that standard empirical estimates of iceberg trade costs often violate the TI, and are therefore inconsistent with the theory. We propose an estimation routine that respects the TI and yields estimates that are consistent with the workhorse models. This measure of trade costs allows us to compute the impact of changes in trade barriers while complying with the TI. First, we compute the welfare gains using only "direct" changes in trade costs (the standard approach). Second, we update the trade cost matrix to allow for re-routing whenever the TI is violated. We show that welfare gains are often substantially different (and usually higher) when taking the TI into account.

JEL Classification: F10, F14, F17

Keywords: Trade Costs, re-routing, triangle inequality, Welfare

Reto Foellmi - reto.foellmi@unisg.ch
University of St. Gallen and CEPR

Christian Hepenstrick - christian.hepenstrick@snb.ch
Swiss National Bank

David Torun - david.torun@unisg.ch
University of St. Gallen

Triangle Inequalities in International Trade: The Neglected Dimension*

Reto Foellmi
University of St.Gallen
reto.foellmi@unisg.ch

Christian Hepenstrick
Swiss National Bank
christian.hepenstrick@snb.ch

David Torun
University of St.Gallen
david.torun@unisg.ch

March 21, 2022

Abstract

The so-called triangle inequality (TI) in international trade should, theoretically, hold for any three countries to avoid cross-border arbitrage. When trade costs change, re-routing opportunities—as captured by the TI—might arise because a shipment through an intermediary becomes cheaper under adjusted trade costs. We show that the widely-used "exact hat algebra" approach, which does not require a calibration of trade costs, is unable to measure potential gains from re-routing. In addition, we show that standard empirical estimates of iceberg trade costs often violate the TI, and are therefore inconsistent with the theory. We propose an estimation routine that respects the TI and yields estimates that are consistent with the workhorse models. This measure of trade costs allows us to compute the impact of changes in trade barriers while complying with the TI. First, we compute the welfare gains using only "direct" changes in trade costs (the standard approach). Second, we update the trade cost matrix to allow for re-routing whenever the TI is violated. We show that welfare gains are often substantially different (and usually higher) when taking the TI into account.

Keywords: trade costs · re-routing · triangle inequality · welfare

JEL Classification: F10 · F14 · F17

*We thank Fernando Alvarez, Alessandro Ferrari, Sam Kortum, Stefan Legge, Piotr Lukaszuk, Marc Muendler, Dennis Novy, Ralph Ossa, Michael Peters, Frank Pisch, Ariell Reshef, Ulrich Schetter, Mathias Thoenig, Josef Zweimüller, and the participants to several seminars for helpful comments and discussions.

The "Online Appendix" can be found at the bottom of this document.

1 Introduction

An important driver of globalization is the reduction in trade barriers of any kind. It is a long-standing question how large these are and how much they impede the flow of goods, services and ideas between countries. To determine the size of trade barriers in a world with many countries, we typically assume that trade costs are subject to a variant of the triangle inequality in Euclidean geometry. Trade costs between countries i and j must be lower than or equal to the trade costs between country i, an intermediate country k and country j. The reason for that is simple: a cost-minimizing producer will send a product through the cheapest transportation channel. A positive trade flow between countries i and j directly is only possible if the triangle inequality of trade costs holds.

However basic this assumption is, the literature typically does not account for potential violations of the triangle inequality (TI) when studying shocks to the trade cost matrix. After changes in trade costs, re-routing opportunities might arise if a shipment through an intermediary becomes cheaper under counterfactual trade costs, which would be captured by the TI. Neglecting violations of the TI has serious consequences for the validity of structural trade models, and therefore for their quantitative predictions of the (welfare) impact of trade liberalizations.

To address the TI problem, this paper proceeds in two steps. First, we document that standard estimates of trade costs violate the TI for many bilateral trade pairs and a large part of international trade volume. It would be inconsistent with the theory to use these estimates for counterfactual exercises. Second, as a solution to the aforementioned finding, we provide a model-consistent method to estimate trade costs while satisfying the TI in a standard quantitative trade model. If the inequality were not to hold the implied trade flow between two countries would be zero, since direct sourcing from a country where the TI is violated is not cost-minimizing. Hence, a basic premise to obtain trade costs that are logically consistent is an estimation of a trade model that is constrained by the respective TIs.

We must verify the TI whenever there is a change in trade costs. This holds also in cases where we focus on relative changes without measuring the level of trade costs—an approach labeled "exact hat algebra," which was first introduced by Dekle, Eaton, and Kortum (2007) and is now wide-spread in the literature.[1] To see this, consider a

[1]See, e.g., Costinot and Rodríguez-Clare (2014); Caliendo and Parro (2015); Eaton, Kortum, Neiman, and Romalis (2016); Caliendo, Dvorkin, and Parro (2019).

globalization step that results in an overall decrease of 10 percent of all bilateral trade costs. When using hat algebra, we do not need to know the initial level of trade costs. This procedure is only valid if the point of departure consists of trade costs that fulfill the TI. But even when the TI holds initially, an overall reduction of 10 percent in trade costs could create violations of the TI that are not taken into account if a standard hat algebra approach is applied. For example, let (initial) ad-valorem trade costs to export from country i to country j be equal to τ_{ij}. The triangle inequality for this trade relation reads $\tau_{ij} \leq \tau_{ik}\tau_{kj}$ for any potential intermediary country k. If we now assume that trade costs change by a factor $\gamma < 1$, such that new trade costs equal $\gamma\tau < \tau$, the TI holds only if $\tau_{ij} \leq \gamma\tau_{ik}\tau_{kj}$. If γ is sufficiently smaller than one, this inequality is violated, implying a profitable re-routing opportunity for country pair ij.[2]

To solve this problem, we propose to estimate trade costs using constrained regressions with constraints satisfying the TIs. We demonstrate our method building on the seminal contribution by Waugh (2010). Similar to his approach we use a gravity equation to estimate trade costs via standard OLS, but in addition we impose constraints to respect the TIs. The resulting bilateral trade costs satisfy the TI and cannot be smaller than one, which is different from standard estimates, as we will show. When we consider changes in trade costs, these estimates allow us to endogenously determine re-routing opportunities.

The specific examination of the TI has important quantitative implications for welfare. In one of our counterfactual exercises, we look at a global proportional drop in trade costs of 25%, as done by Eaton and Kortum (2012). Importantly, counterfactual iceberg trade costs are never allowed to fall below one. Our estimates suggest that, after this shock, about 13% of trade relations have at least one arbitrage opportunity, and would thus be better off by involving an intermediary. The trade-weighted share is almost twice as large. In a framework à la Eaton and Kortum (2002) focusing on aggregate trade flows, this has substantial implications for the welfare gains that arise from re-routing.[3] We find that most countries benefit from the option to re-route. The third of the sample with the largest welfare differences gains on average about 5% more real income. To put this into perspective, we scale gains from re-routing by

[2]Furthermore, we do not know whether trade costs hit their theoretical lower bound (i.e., $\tau_{ij} = 1$) after the trade cost shock if we do not estimate their initial levels.

[3]Qualitatively, the arguments we make also apply to models that feature variable and fixed trade costs, such as the Melitz (2003) model. However, quantitatively results may well differ, as our trade cost estimates are a compound measure of iceberg and fixed costs through the lens of a Melitz (2003)-Chaney (2008) framework. The welfare gains from re-routing that we measure after a shock to trade costs (cf. Section 4) would therefore not necessarily be the same in those models.

the well-known welfare losses that would arise from a complete trade shutdown. The average return in our sample from re-routing is equal to 17.5% of the net welfare loss from going to autarky. Our results suggest that accounting for arbitrage opportunities after shocks to trade costs becomes more important the larger the shock size. This indicates a sort of "local stability" of the trade cost matrix.

We show that our method can be easily applied to more complex models. We repeat our counterfactual exercises using the powerful multi-sector extension of the Eaton and Kortum (2002) model with input-output linkages by Caliendo and Parro (2015). The lesson from these additional analyses is twofold. First, in smaller samples re-routing is less likely to be profitable. This is intuitive, since arbitrage opportunities are more likely the larger our set of potential intermediaries. Second, although we observe less re-routing with fewer countries in our sample, involving intermediaries tends to be highly profitable when sectoral linkages are included in the analysis. After a 25% drop in bilateral sector-level trade costs, the average welfare gain increases by about 20% once we take the TI into account.

In economic geography, it is argued that an improvement in one segment of a transportation infrastructure may alter the shipping costs along the entire network (e.g., Allen and Arkolakis (2020) analyze changes in the U.S. highway network). The reason lies in the fact that, e.g., lower travel time between two cities may lead to changes of travel paths—re-routing—in the whole network. We argue in the paper that similar mechanisms are (implicitly) present in canonical international trade models when we consider the welfare effects of trade cost changes. A related study in this regard is that of Ganapati, Wong, and Ziv (2020), who show that entrepôts create an incentive to ship via indirect routes.

While the TI received little attention in quantitative trade models, it plays a central role in trade policy. An exporter from a third country i could have an incentive to ship a product to country j via country k, taking advantage of a favorable bilateral free trade agreement between countries k and j when there does not exist a similar treaty between i and j. Sophisticated rules-of-origin clauses in free trade agreements, typically based on value-added shares in each country, should preclude such trade deflection. However, Felbermayr, Teti, and Yalcin (2019) show that trade deflection to take advantage of tariff differentials is not profitable, as tariff differences do not outweigh increased transportation costs. Their argument underlies that the TI of trade costs should hold when we estimate trade costs, and that tariff differences (even if they could be exploited at all) are not sufficiently large to create violations of the TI.

The paper is organized as follows. In Section 2, we introduce the theory. Section 3 elaborates the estimation strategy, and Section 4 discusses counterfactual exercises. Section 5 concludes.

2 Theory

Consider the following simple theoretical structure. The world consists of N countries and many goods. Each of the countries purchases its products from the respective cheapest source in the world. This basic feature holds in a broad set of trade models; for instance, in a Ricardian trade model à la Eaton and Kortum (2002) (henceforth, EK).

Transportation sector. Trade is costly and takes place within a transportation sector that ships goods from country i to country j. The transportation sector is characterized by perfect competition. This sector uses the particular good as the only input, no labor is employed in transportation. Following the literature, we assume that transportation costs are proportional to the marginal costs of the good: if marginal costs of production in country i equal c_i, the marginal costs of producing a good in country i and selling the item to country j equal $\tau_{ij}c_i \geq c_i$, with $\tau_{ij} \geq 1$. Thus, we make the common assumption that international trade entails constant ad valorem trade costs—known as "iceberg" trade cost (Samuelson, 1954). Trade costs are a c.i.f. measure, hence they account for trade costs from border i to border j. Therefore it is consistent to set within-country trade costs $\tau_{jj} = 1$. In this setting, the triangle inequality (TI) of trade costs takes the familiar form[4]

$$\tau_{ij} \leq \tau_{ik}\tau_{kj}.$$

The TI requires that the trade cost of shipping the product via some intermediary k, τ_{ik} times τ_{kj}, may not be lower than the costs of shipping the product directly from country i to j, τ_{ij}. If the TI were violated, a cost-minimizing producer would choose

[4]If within-country trade costs are non-zero (e.g., for retail distribution), total trade costs between i and j would amount to $\tau_{ij}\tau_{jj}$. This leaves the TI unchanged as we would have to compare $\tau_{ij}\tau_{jj} \leq \tau_{ik}\tau_{kj}\tau_{jj}$ in this case, assuming that within-country distribution costs have not to be paid in intermediate country k. For an analysis of internal trade frictions see, e.g., Sotelo (2020).

to source the product through the trade cost minimizing path.[5] Assume there exists at least one intermediary k where $\tau_{ij} > \tau_{ik}\tau_{kj}$. We define

$$\tau_{ij}^{1} = \min_{k} \left\{ \tau_{ik}\tau_{kj} \right\}, \tag{1}$$

where $\tau_{ij}^{1} = \tau_{ij}$ when $\arg\min_k \left\{\tau_{ik}\tau_{kj}\right\} \in \{i,j\}$. Note that it suffices to look at three-countries triangle violations, and to replace trade costs that violate the TI with their updated value τ^1 iteratively. To see this, assume trade costs are lower if the good is shipped via two intermediaries k and k', but not if it is shipped via only one of them, i.e., $\tau_{ij} > \tau_{ik}\tau_{kk'}\tau_{k'j}$ but $\tau_{ij} < \tau_{ik}\tau_{kj}$ and $\tau_{ij} < \tau_{ik'}\tau_{k'j}$. Taking the former two inequalities together—obviously we get the identical result when using the latter inequality—, we see that the violation of the TI originates from a three-countries triangle violation. We obtain $\tau_{kj} > \tau_{kk'}\tau_{k'j}$, trade from k to j is cheaper via k'. In a violation where n intermediaries are involved, we apply the same reasoning to reduce the inequality to a violation where $n-1$ intermediaries are involved. As the number of countries is finite, the number of potential multi-intermediary paths for a trade flow between countries i and j is finite as well.[6]

Applying the procedure laid out in equation (1) to updated trade costs, we obtain $\tau_{ij}^{2} = \min_k \left\{\tau_{ik}^{1}\tau_{kj}^{1}\right\}$. We repeat this until all violations of the three-countries triangle inequality are eliminated, and therefore violations with $n > 1$ intermediaries are eliminated too. This leads to the following definition.

Definition 1 *We define* $\tilde{\tau}_{ij} = \min_k \left\{\tilde{\tau}_{ik}\tilde{\tau}_{kj}\right\}$*, where* $\forall\{i,j\}$ *it holds that* $\arg\min_k \left\{\tilde{\tau}_{ik}\tilde{\tau}_{kj}\right\} \in \{i,j\}$.

We will achieve the optimal $\tilde{\tau}$ after a finite number of iteration steps as in equation (1). The resulting updated trade cost between two countries follows an optimal trade trajectory along a finite set of countries. Formally, the optimal trade cost is a product of initial trade costs, i.e., $\tilde{\tau}_{ij} = \Pi_{u=1}^{u=n-1}\tau_{k_u k_{u+1}}$ with $k_1 = i$, $k_n = j$, and $n \geq 2$. We summarize this insight in the following proposition.

[5] In the words of Eaton and Kortum (2005, p. 172), "arbitrage would eliminate violations of the TI since k would emerge as an entrepôt, so that [using our notation] $\tau_{ij} = \tau_{ik}\tau_{kj}$." Note further that we are searching for a baseline trade cost matrix that is theoretically consistent rather than exploring reasons why violations might actually be present in the data. The presence of preferential trade agreements could be such a reason, as argued, e.g., by Mossay and Tabuchi (2015).

[6] Note further that multiple re-shipping via the same intermediary can never be optimal, as trade costs are theoretically bounded to satisfy $\tau_{ij} \geq 1$, and thus $\tau_{ij}^{x} \geq \tau_{ij}$ for any constant $x \geq 1$.

Proposition 1 *Consider a trade cost matrix T with elements $\{\tau_{ij}\}$, where $i, j \in \{1, 2, ..., N\}$. There exists a unique trade cost matrix $\widetilde{T}$ where all trade costs satisfy the triangle inequality $\tilde{\tau}_{ij} \leq \tilde{\tau}_{ik}\tilde{\tau}_{kj}$.*

Proof Existence is shown in the text. To see that $\widetilde{T}$ is unique, assume there exists a second matrix $\widetilde{T}'$ with at least one $\tilde{\tau}'_{ij} \neq \tilde{\tau}_{ij}$. Note that this would violate the definition of $\tilde{\tau}_{ij} = \min_k \{\tilde{\tau}_{ik}\tilde{\tau}_{kj}\}$. □

Counterfactuals. It is important to note what this implies for the structural estimation of trade models. Since the particular good is the input in the transportation sector, we can simply replace τ_{ij} by $\tilde{\tau}_{ij}$ in the welfare calculations of trade policies. In particular, it does not matter that the physical trade flow between i and j goes through intermediary k.[7] If $\tilde{\tau}_{ij} = \tilde{\tau}_{ik}\tilde{\tau}_{kj}$, updated trade costs are the same whether trade takes the direct route between i and j, or goes via the intermediary. If the total amount of trade between i and j equals X, all values between 0 and X for direct trade between i and j are consistent with the theory (with the remainder going through k). Hence, if we use the updated trade cost matrix, $\widetilde{T}$, we can continue to work with standard gravity equations in simulations, even when the initial matrix exhibits violations of the TI.

Corollary 1 *The gravity equations of trade flows remain functionally unchanged with the trade cost matrix $\widetilde{T}$ replacing T.*

What does this result imply for empirical studies of trade cost changes? Many counterfactual experiments make use of the fact that workhorse trade models allow us to estimate relative changes in economic outcomes as a function of relative changes in trade costs and initial trade shares only. There is thus no need to calibrate the level of trade costs (or size-related parameters such as the state of technology and population size). This simple and powerful tool is known as "exact hat algebra," where the word "exact" refers to the fact that no approximations are involved in the estimation process. Consistent with the theoretical setup, one assumes that the TIs are satisfied in the initial equilibrium. However, as the level of trade costs is only implicitly present in standard structural formulae, we do not know whether, after a

[7]Recall that the transportation sector is perfectly competitive, and thus there are no rents from being involved as an intermediary.

shock to the trade cost matrix, profitable re-routing opportunities arise.[8] Consider, e.g., a trade pair where the TI holds initially, i.e., $\tau_{ij} < \tau_{ik}\tau_{kj}$. Let us retake our thought experiment from the introduction and assume there is an overall reduction of bilateral trade costs to a factor $\gamma < 1$. Obviously, there exists a value of γ such that $\gamma\tau_{ij} > \gamma^2\tau_{ik}\tau_{kj}$. This implies that a simple hat-algebra approach, relying solely on changes in trade costs, could lead to logically inconsistent results, since violations of the TI must be taken into account in any simulation. In other words, potential second-order effects of trade cost changes on the transportation sector are neglected. Moreover, we cannot be sure whether γ is small enough for τ_{ij} to hit its lower bound (i.e., $\tau_{ij} = 1$). To ensure a model-consistent counterfactual that accounts for potential spillover effects, we therefore need to estimate the initial level of τ_{ij}, and after a shock update the trade cost matrix as outlined in Proposition 1, in case any TI is violated.[9] Before turning to the quantitative importance of TIs in Section 4, the next section describes our approach to estimating τ_{ij}.

3 Methodology

According to Section 2, we need to calibrate the initial trade cost matrix in order to verify whether violations of the TI arise after a shock to trade costs. In this section, we illustrate how we estimate a trade cost matrix that is consistent with the TIs. The underlying off-the-shelf EK frameworks are standard in the literature, and thus only briefly outlined in Appendix A.

TI violations. There are several contributions that estimate iceberg trade costs in variants of the EK framework (e.g., Eaton and Kortum, 2002; Waugh, 2010; Novy, 2013; Parro, 2013; Simonovska and Waugh, 2014; French, 2016). To our knowledge, however, we are the first to plug these estimates into the TIs to check that they are satisfied. We find that all of the listed methods generate a trade cost matrix that features violations of the TI. The share of trade relations which would be better off shipping via a third country varies between 8% and 50%. Moreover, the calibrated τ's

[8] Counterfactuals involving trade costs are unique in this respect. Other shocks such as changes in technology parameters will not encounter the issues discussed in this paper.

[9] Note that the widely-used reference point of welfare losses under complete autarky is an exception. With $\tau_{ij} \to \infty \; \forall i \neq j$, there cannot be any violation of the TI. See Arkolakis, Costinot, and Rodríguez-Clare (2012) for a detailed discussion on welfare losses from autarky implied by standard trade models. More generally, whenever all trade costs are raised by the same factor (i.e., $\gamma > 1$ in the example above), we will not encounter violations if the initial trade cost matrix was consistent with the TIs.

are not restricted to be greater than or equal to one, and indeed some estimates fall below this theoretical lower bound. Through the lens of the model, this would then imply that value is created—instead of lost—in transit. A trade cost matrix that does not fully satisfy the TIs is inconsistent with the theory, and thus cannot be employed to compute counterfactuals.[10]

Trade cost estimates. To overcome the encountered issues, we augment Waugh's (2010) influential method to ensure that the calibrated τ's will obey the TI, and satisfy $\tau_{ij} \geq 1 \; \forall i, j$. The reason we choose Waugh (2010) is twofold. First, these estimates help the model mimic empirically observed bilateral trade flows and prices. Second, the calibrated trade costs are entirely based on OLS coefficients, making the implementation very simple.

As detailed in Appendix A.1, the EK model yields a log-linear relationship between trade flows (scaled by home absorption) and trade costs, τ_{ij}, as well as importer- and exporter-specific terms

$$\log\left(\frac{\pi_{ij}}{\pi_{jj}}\right) = \log\left(T_i w_i^{-\theta}\right) - \log\left(T_j w_j^{-\theta}\right) - \theta \log\left(\tau_{ij}\right), \tag{2}$$

where π_{ij} is exporter i's share in importer j's total expenditure, $T_i w_i^{-\theta}$ is a measure of country i's productivity, and θ the so-called trade elasticity. Let $S_j^m := -\log\left(T_j w_j^{-\theta}\right)$, where the superscript m is chosen to indicate that this expression will be captured by the importer fixed effect (FE).

The functional form imposed to estimate trade costs is a key assumption. In line with Waugh (2010), trade costs are assumed to be a log-linear function of an exporter-specific shifter, ex_i, as well as bilateral dummy variables, that is,

$$-\theta \log\left(\tau_{ij}\right) = ex_i + \sum_{l=1}^{D} \beta^l d_{ij}^l, \tag{3}$$

where β^l is the coefficient on a dummy variable d_{ij}^l. Examples for dummy variables are common border, common language or distance bins. The shifter ex_i captures all exporter-specific aspects affecting trade costs that are not picked up by the dummy variables. Let $S_i^x := ex_i - S_i^m$, where the superscript x is chosen to indicate that this expression will be captured by the exporter FE.

[10] The TI violations found using Waugh's (2010) method are discussed in Appendix C, since our estimation will be based on this procedure. The estimated trade costs and violations that result from the other estimates are available upon request.

By plugging equation (3) into equation (2), and applying the definitions of S_i^x and S_j^m, we obtain our gravity regression

$$\log\left(\frac{\pi_{ij}}{\pi_{jj}}\right) = S_i^x + S_j^m + \sum_{l=1}^{D} \beta^l d_{ij}^l + \varepsilon_{ij}, \tag{4}$$

where ε_{ij} denotes the error term.[11] The dependent variable is directly observed in the data. The elements on the right-hand side of equation (4) are an exporter FE, an importer FE, and a set of bilateral dummy variables. Note that the exporter-specific component of trade barriers in equation (3) is simply the sum of the exporting country's exporter and importer FEs. In the absence of ex_i, the exporter and importer FEs would add up to zero according to the theory. Empirical deviations from this theoretical target will thus be attributed to trade costs. The higher the exporter-specific shifter, the lower trade costs are.

As in Waugh (2010), our baseline regression includes dummy variables indicating whether the country pair shares a common border, and whether the geographical distance between two countries lies within a specific distance interval (cf. Table E.1). In Appendix C, we run this regression on a sample of 127 countries in the year 2004. We show that this procedure can yield estimates that violate the TI, and sometimes fall below one. The data are described in Appendix B.

What leads to these violations of the theory? Let t_{ij} be the estimate of τ_{ij}, which is calculated as

$$-\hat{\theta}\log(t_{ij}) = ex_i + \sum_{l=1}^{D} \hat{\beta}^l d_{ij}^l,$$

where we set $\hat{\theta} = 4.14$, which is Simonovska and Waugh's (2014) preferred estimate, and $\hat{\beta}^l$ denotes the estimate of β^l. The TI takes a log-linear form, and can therefore be rewritten as

$$\log(t_{ij}) \leq \log(t_{ik}) + \log(t_{kj}).$$

Multiplying by $-\hat{\theta}$, and plugging in the expression above, we receive

$$0 \geq ex_k + \sum_{l=1}^{D} \hat{\beta}^l \left(d_{ik}^l + d_{kj}^l - d_{ij}^l\right) \quad \forall\ i, k, j. \tag{5}$$

[11] Alternatively, one could add the error term to the trade cost function in equation (3). In Appendix C, we show that this leads to even more TI violations. We chose the specification in equation (3) because this enables us to drastically reduce the number of constraints required in our constrained estimation, as explained below.

Inequality (5) nicely illustrates how TI violations may emerge.[12] To gain intuition, suppose that the triplet ikj does not share a common border, and all countries fall within the same geographical distance bin l. With these assumptions, inequality (5) reduces to $0 \geq ex_k + \hat{\beta}^l$. The additional distance cost from involving intermediary k is equal to $\hat{\beta}^l$ (distance has a negative impact on trade flows, i.e., $\hat{\beta}^l < 0$). If country k is a relatively strong exporter (i.e., ex_k is high), engaging this intermediary may outweigh the additional distance penalty, and thus yield estimates that suggest an arbitrage opportunity (in this case, $ex_k + \hat{\beta}^l > 0$).

Similarly, to have $t_{ij} \geq 1$, we need the following "lower-bound" condition to hold

$$0 \geq ex_i + \sum_{l=1}^{D} \hat{\beta}^l d_{ij}^l \quad \forall\, i, j. \tag{6}$$

In the modeling framework outlined above, inequalities (5) and (6) are necessary and sufficient conditions to make sure that the TI holds and the trade costs are limited from below. Next, we outline how to perform estimations consistent with the TIs.[13]

A model-consistent calibration. Standard gravity regressions that otherwise follow the theory closely neglect the additional implicit constraints (5) and (6). Our approach is a model-consistent calibration that imposes the latter. We minimize the sum of squared residuals subject to a set of linear inequality constraints. Accordingly, the solution to this convex optimization problem is a global minimum. However, the number of constraints is substantial. Let N be the number of countries in the sample. For every potential intermediary k, we have $(N-1) \times (N-2)$ TIs, as well as $(N-1)$ lower-bound conditions (inequality (6)). The theory therefore implicitly requires a total of $N \times (N-1)^2$ constraints—over two million in our sample with 127 countries. A "brute force" solution to the issues above would be to impose all of them in the estimation, which slows down the computation enormously.[14]

We can instead use the fact that ex_k does not vary across importers. Moreover, since the trade cost function contains exclusively dummy variables, t_{ij} may draw only

[12] Note that the long debated size of θ is irrelevant to the TI. It will matter, however, in the welfare calculations below.

[13] Recall from Section 2 that checking re-routing via multiple intermediaries is not necessary (as long as $t_{ij} \geq 1$), so that inequalities (5) and (6) suffice to guarantee that there are no arbitrage opportunities in the trade cost matrix. Requiring only one of the two inequalities is not enough.

[14] The rest of this section is only relevant to understanding our constrained estimation but in no way essential to understand the remainder of the paper.

a limited number of unique values.[15] For a given ex_k, the right-hand side of inequality (6) can take up at most 2^D unique values. Similarly, the "loss term" $d^l_{ik} + d^l_{kj} - d^l_{ij}$ has four potential combinations. Inequality (5) thus has at most 4^D unique values for a given ex_k. This yields a total of $N \times (4^D + 2^D)$ unique constraints, which is similar to $N \times (N-1)^2$ in our baseline regression with $N = 127$ and $D = 7$. However, we can substantially reduce the number of constraints, considering that most combinations are empirically irrelevant. For instance, the distance between countries i and j can obviously only fall within one specific interval. This implies that inequality (6) will take up at most 12 ($= 2 \times 6$) rather than 128 ($= 2^7$) unique values for a given ex_k. Out of these remaining 12 combinations several are again redundant, since countries share a border only with geographically close neighbors. In a similar vein, we can trim the 4^D unique values of inequality (5). For example, if countries i and j are in the top distance bin while i and k are in the bottom interval, we must have that the distance between k and j falls within one of the two top bins. Following this pruning procedure we heavily reduce the number of constraints from over two million to about ten thousand (or, by a factor of 200). The time required to compute the constrained regression drops from several hours to a few seconds. Importantly, most of these linear inequality constraints will not be binding in the estimation (only twelve in our main specification), and thus it is not an issue that the number of constraints is still close to the number of observations.[16]

Comparison. Let $\tilde{t}_{ij}$ be our *constrained* estimate of τ_{ij} that we obtain by running equation (4) while imposing the constraints described above. By construction, none of these estimates violates the TI, and we have $\tilde{t}_{ij} \geq 1\ \forall i, j$. To compare the constrained and unconstrained values, we set $t_{ij} = 1$ whenever we get $t_{ij} < 1$.

Table E.3 reports the coefficients from the unconstrained and constrained regressions, as well as their percentage impact on trade costs. The most apparent difference is that of the coefficients on the border dummy. The suggested benefit on trade costs more than halves to ensure that the TI remains satisfied. The impact of distance also flattens, with a notable drop in the gain from being in the lowest distance bin rather than in the second lowest one. By contrast, the values of ex_i stay very similar, although we see a decline for very strong exporters (cf. Figure D.2). The overall fit is

[15] Analogous to distance, we can easily accommodate continuous control variables by dividing them into bins, and including dummy variables that indicate into which bin an observation falls.

[16] The procedure (and constraints) is the same irrespective of whether we use OLS or PPML. We use OLS for all our main results.

comparable, in that the R-squared decreases by less than one percentage point.

These insights are reflected in Figure D.3, where we plot t_{ij} against $\tilde{t}_{ij}$. The constrained estimation pushes up low values of t_{ij}, whereas it leaves higher values largely unchanged, and, if anything, slightly decreases them. Intuitively, close countries cannot have trade barriers that are so low that involving a neighbor as an intermediary is beneficial. As a last step, we verify that the ad hoc solution of taking unconstrained estimates and computing post-estimation updated values following Proposition 1 does not yield the same results. In Figure O.1 of the Online Appendix, we show that there is a substantial difference between $\tilde{t}_{ij}$ and updated values of t_{ij}. Note that the updated value of t_{ij} is, by construction, always weakly smaller than t_{ij}, whereas this is clearly not the case with $\tilde{t}_{ij}$. In fact, constrained trade costs are generally closer to unconstrained ones than to the updated values.

4 Counterfactual Exercises

We want to understand the impact of the TI, working through potential re-routing of trade flows, on the welfare gains from trade liberalizations. For that aim, we conduct counterfactual experiments where we lower trade costs. We compute welfare gains (i.e., relative changes in real income) using hat algebra as explained in Appendix A, and we do so in two ways.[17] First, we shock the constrained trade cost matrix and do not verify whether TI violations arise (standard approach). Second, we check the TI and update trade costs whenever a new route allows for a cheaper delivery of the good. The relative change in trade costs is then the *effective* drop in trade barriers that materializes after the shock, which takes potential spillover (re-routing) effects into account. We do not allow trade costs to fall below one in either case.

To make our argument precise, we set the initial value of τ_{ij} equal to its constrained estimate, $\tilde{t}_{ij}$, so that the starting point features no arbitrage opportunity. For instance, let the counterfactual exercise be a proportional decrease in trade costs by 25%. The ratio of counterfactual to initial trade costs, denoted by $\hat{\tau}_{ij} = \tau'_{ij}/\tau_{ij}$, equals 75%, as long as the new trade cost τ'_{ij} does not fall below one. Formally, the first step sets $\hat{\tau}_{ij} = \max\{0.75, 1/\tau_{ij}\}\ \forall i \neq j$, and computes welfare gains as described in

[17] We do not explicitly account for trade imbalances in these experiments, since this dimension is not directly relevant to our discussion. A simple way to do so would be to include an additional parameter which captures net exports in the initial equilibrium (see Dekle et al., 2007).

Appendix A.[18] The second step verifies whether counterfactual trade costs, $\tau'_{ij} = \hat{\tau}_{ij}\tau_{ij}$, satisfy the TI. In case there are arbitrage opportunities, we follow Proposition 1 to obtain updated values, $\tilde{\tau}'_{ij}$. The welfare gains are then computed in the standard way but setting $\hat{\tau}_{ij} = \tilde{\tau}'_{ij}/\tau_{ij}$, i.e., we use the *effective* change in trade costs taking the option of re-routing explicitly into account.

Uniform drop in τ. We use the single-sector setup outlined in Appendix A.1, and decrease global trade costs proportionally by 25%—an exercise conducted by Eaton and Kortum (2012).

How many arbitrage opportunities arise? Table E.4 reports the share of trade pairs that faces at least one violation of the TI when trade costs drop by a quarter. Observe that some values would fall below one if we did not impose this lower bound. But even with this restriction, about 13% of trade pairs—or 24% of total trade—can supposedly decrease their delivery costs by involving a third country. The implied savings potential for these relations is meaningful, considering that trade costs can on average drop by a further 8%. Thus, the actual reduction in trade barriers is considerably larger for several trade relations when we account for potential second-order effects.

The welfare implications are depicted by Figure D.4. Intuitively, countries that are close to the main hubs in Table E.4, such as Spain or Malaysia, benefit most from re-routing. The results suggest that Spain could more than double its welfare gain by re-routing in response to this uniform globalization experiment. There are also countries that lose through this step, but the losses are usually negligible. As is well-known in the literature, the one-sector model generally yields relatively low gains from trade (see Costinot and Rodríguez-Clare (2014) and Ossa (2015) for a detailed discussion). It thus makes sense to put these numbers into perspective. The welfare loss from going to full autarky (i.e., $\tau_{ij} \to \infty \; \forall i \neq j$) is a widely-used benchmark, not least because calculating it is straightforward (see Appendix A.1).

In Table E.5, we list the 40 countries with the largest return from re-routing. Column (3) reports the relative real income loss that would result if a country were to completely shut down international trade, which serves as a benchmark. In Column (4), we scale the additional income gain from re-rerouting (i.e., Column (1) minus Column (2)) by the absolute value of Column (3). The extra gains are substantial relative to the

[18] This procedure slightly differs from the standard hat algebra approach, which would set $\hat{\tau}_{ij} = 0.75 \; \forall i \neq j$. We need to impose the theoretical lower bound, otherwise we would obtain extreme results when allowing for re-routing.

benchmark values. For instance, Spain can collect an additional 8% of its initial real income by exploiting arbitrage opportunities. This is double the reference number of a 4% loss in the case of a complete trade collapse. Although the magnitude of the latter strikes as unrealistic, it shows how, through the lens of the model, the implied welfare differences are sizable. These countries can on average extract almost half of their benchmark real income changes by involving intermediaries. This exercise nicely illustrates why it is important to consider potential spillovers even in counterfactuals that involve uniform trade cost changes.

Table E.6 shows that verifying the existence of arbitrage opportunities becomes more important the larger the shock to trade barriers is, suggesting that the trade cost matrix is "locally stable." Repeating the exercise from above for a 5% drop in global trade costs yields only few re-routing opportunities. By contrast, involving intermediaries creates a median net welfare change equal to 22.5% when we halve trade costs. On average, the net real income change from re-routing divided by the net loss from going to autarky is huge in the latter case, and non-negligible in the former scenario. Nevertheless, these statistics hide a lot of heterogeneity, as the more detailed results above show.

Asymmetric shock to τ. We also compute a counterfactual with non-uniform trade cost changes. Specifically, following Waugh (2010), we eliminate bilateral asymmetries in trade barriers by setting the counterfactual trade costs equal to the minimum of the (potentially different) bilateral trade costs; formally, $\tau'_{ij} = \min\{\tau_{ij}, \tau_{ji}\}$. A policy change that would come close to an equalization of bilateral trade costs would be the abolishment of artificial, policy-related bilateral distortions. Note that trade costs cannot fall below one due to this shock, but nonetheless the TI may not hold anymore.[19] In fact, Table E.7 shows that more than half of the trade relations now face at least one arbitrage opportunity, with an impressive extra reduction of 50% on average. The winners in this counterfactual scenario are relatively small countries, which is reflected

[19] How do re-routing opportunities arise? Suppose that for the triplet ikj we have $ex_k > ex_i \geq ex_j$. By eliminating bilateral asymmetries (which stem exclusively from the exporter-specific trade cost components), inequality (5) becomes

$$0 \geq \underbrace{(ex_k - ex_i)}_{>0} + ex_k + \sum_{l=1}^{D} \hat{\beta}^l \left(d^l_{ik} + d^l_{kj} - d^l_{ij} \right),$$

which will be violated if inequality (5) was (close to) binding to begin with. In words, country k is likely to become a hub if it is a relatively strong exporter.

in the low trade share that would benefit from re-routing. Therefore, weak exporters (i.e., with low ex_i) benefit not only from the direct trade cost reduction to strong partners, but also from an indirect decline through re-routing (see footnote 19).

Figure D.5 and Table E.8 are the counterparts of Figure D.4 and Table E.5, and they tell a similar story. Most countries benefit from the re-routing option, and those who do not lose relatively little. The return from involving intermediaries is substantial relative to the benchmark autarky losses. In general, welfare changes are large in this simulation due to the shock size, with the average drop in trade costs being equal to 23%.

Multi-sector model. We can easily repeat the two main simulations above within the multi-sector framework with sectoral input-output linkages outlined in Appendix A.2. The reasoning above on transportation and potential arbitrage opportunities applies entirely to multi-sector environments as well. The model yields a sector-level equivalent to the aggregate gravity equation. We can thus apply the method from Section 3 to estimate sector-specific iceberg trade costs, τ_{ij}^k, using sector-specific fixed effects and elasticities (i.e., having a sector-specific equation (4)). This in turn enables us to compute counterfactuals with and without the option of re-routing.

Our sector-level data cover only 34 countries (cf. Appendix B), and therefore there will be fewer potential intermediaries a country pair may select than in the single-sector counterfactuals above. This should dampen the impact of re-routing opportunities compared to the previous analyses. For completeness, Table O.5 reproduces the symmetric counterfactual in Table E.5 using the WIOD data instead of the ITPD-E data. Indeed the gains from re-routing after the symmetric shock to trade costs tend to be considerably smaller in the reduced sample.

What happens when the sectoral dimension is added? Trade cost changes affect the equilibrium in a more complex way when input-output linkages are taken into account, and therefore additional reductions in trade costs may have strong effects.[20] Table E.9 replicates Table E.5 using the multi-sector setup, i.e., we decrease all bilateral sector-level trade costs by 25%. The share of trade relations which face at least one arbitrage opportunity after the shock varies substantially across sectors, ranging from 0% to 66%, with a median value of 11%. The option to re-route has a sizable impact on welfare, with the average real income gain rising by about 20%. This is considerably

[20]The sectoral dimension tends to generate larger welfare changes in general, which is well-understood in the literature (see, e.g., Costinot and Rodríguez-Clare, 2014).

larger than the gains implied by the one-sector counterpart (see Table O.5).

The suggested benefit from re-routing after eliminating bilateral asymmetries in trade barriers is huge in both the single-sector (Table O.6) and the multi-sector version (Table E.10), although at a higher level in the latter case. This result is largely driven by the ROW becoming a hub—which ought to be expected, as 40% of bilateral trade comes from or flows to ROW, making it the strongest exporter, and thus an attractive intermediary (footnote 19). The ROW appears to lose substantially when we introduce the re-routing option in the single-sector model. Nevertheless, interpreting ROW as a hub is somewhat abstract, and we therefore tried excluding it from the sample. This reduces the gains from re-routing, but they nonetheless remain substantial (see Tables O.8 and O.10). In the symmetric counterfactual, ROW plays a way smaller role (see Tables O.7 and O.9).

In a nutshell, spillovers from trade cost changes appear to be an important factor to consider even in small samples. This seems to hold especially for analyses with multiple sectors.

5 Conclusion

We show how the triangle inequality (TI) of trade costs plays a crucial role for counterfactual analyses in general equilibrium trade models. We provide a method to estimate trade costs that satisfy the TI. Standard estimation methods do not take into account potential violations of the TI, and often yield estimates that imply arbitrage opportunities. The latter are inconsistent with the theory, and should thus not be employed in simulations.

We use our estimates to calculate the welfare gains that arise from re-routing after a shock to the trade cost matrix. Using both uniform and asymmetric reductions in trade costs, we show that counterfactual trade cost matrices feature many arbitrage opportunities. That is, several trade pairs could further reduce their delivery costs by involving an intermediary country after the shock. The effective drop in trade costs is therefore larger for these dyads than the "direct" reduction in trade barriers suggests. The TI—and thus re-routing—is quantitatively important, in that welfare gains are substantially higher for many countries when we check for arbitrage opportunities.

References

Allen, T. and C. Arkolakis (2020): "The Welfare Effects of Transportation Infrastructure Improvements," NBER Working Paper Series, No. 25487.

Alvarez, F. and R. E. J. Lucas (2007): "General Equilibrium Analysis of the Eaton–Kortum Model of International Trade," *Journal of Monetary Economics*, 54, 1726–1768.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.

Borchert, I., M. Larch, S. Shikher, and Y. V. Yotov (2021): "The International Trade and Production Database for Estimation (ITPD-E)," *International Economics*, 166, 140–166.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835.

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

Chaney, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98, 1707–21.

Costinot, A. and A. Rodríguez-Clare (2013): "Online Appendix to Trade Theory with Numbers: Quantifying the Consequences of Globalization," March. Available at: https://economics.mit.edu/files/9215.

——— (2014): "Trade Theory with Numbers: Quantifying the Consequences of Globalization," in *Handbook of International Economics*, Elsevier, vol. 4, 197–261.

Dekle, R., J. Eaton, and S. Kortum (2007): "Unbalanced Trade," *American Economic Review*, 97, 351–355.

Eaton, J. and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

——— (2005): "Technology in the Global Economy: A Framework for Quantitative Analysis," *Unpublished manuscript.*

——— (2012): "Putting Ricardo to Work," *Journal of Economic Perspectives*, 26, 65–90.

Eaton, J., S. Kortum, B. Neiman, and J. Romalis (2016): "Trade and the Global Recession," *American Economic Review*, 106, 3401–38.

Felbermayr, G., F. Teti, and E. Yalcin (2019): "Rules of Origin and the Profitability of Trade Deflection," *Journal of International Economics*, 121, 103248.

French, S. (2016): "The Composition of Trade Flows and the Aggregate Effects of Trade Barriers," *Journal of International Economics*, 98, 114–137.

Ganapati, S., W. F. Wong, and O. Ziv (2020): "Entrepôt: Hubs, Scale, and Trade Costs," CESifo Working Paper, No. 8199.

Head, K., T. Mayer, and J. Ries (2010): "The Erosion of Colonial Trade Linkages after Independence," *Journal of International Economics*, 81, 1–14.

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

Mossay, P. and T. Tabuchi (2015): "Preferential Trade Agreements Harm Third Countries," *The Economic Journal*, 125, 1964–1985.

Novy, D. (2013): "Gravity Redux: Measuring International Trade Costs with Panel Data," *Economic Inquiry*, 51, 101–121.

Ossa, R. (2015): "Why Trade Matters after All," *Journal of International Economics*, 97, 266–277.

Parro, F. (2013): "Capital-Skill Complementarity and the Skill Premium in a Quantitative Model of Trade," *American Economic Journal: Macroeconomics*, 5, 72–117.

Samuelson, P. A. (1954): "The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments," *The Economic Journal*, 64, 264–289.

Santos Silva, J. and S. Tenreyro (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.

Simonovska, I. and M. E. Waugh (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92, 34–50.

Sotelo, S. (2020): "Domestic Trade Frictions and Agriculture," *Journal of Political Economy*, 128, 2690–2738.

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. De Vries (2015): "An Illustrated User Guide to the World Input–Output Database: the Case of Global Automotive Production," *Review of International Economics*, 23, 575–605.

Waugh, M. E. (2010): "International Trade and Income Differences," *American Economic Review*, 100, 2093–2124.

Appendix

A Model Setup

The Ricardian frameworks we employ are picked off the shelf from well-known contributions in the literature. The only difference being our handling of trade cost (changes). First, we outline the one-sector Eaton and Kortum (2002) model, and then introduce the multi-sector version with input-output linkages.

A.1 Single-Sector EK

We briefly describe the Eaton and Kortum (2002) version with one sector and labor as the only input (see, e.g., Simonovska and Waugh, 2014). For a detailed discussion of the model consult these papers.

Setup. This static world consists of N countries, indexed by i and j. There is a single tradable final-goods sector, and a continuum of varieties indexed by $\omega \in [0, 1]$. Labor is supplied inelastically by a measure of L_i consumers. The representative agent has CES utility over a bundle of final goods with elasticity of substitution $\sigma > 1$.

Firms operate in perfect competition and employ a constant returns to scale technology, where the cost of producing one unit of good ω in country i is given by $w_i/z_i(\omega)$. w_i is the economy's wage rate, and $z_i(\omega)$ denotes country i's efficiency in producing good ω, which is independently drawn from a Fréchet distribution

$$F_i(z) = \exp\left(-T_i z^{-\theta}\right).$$

For one unit of the good to arrive in destination j, producers in country i need to ship $\tau_{ij} \geq 1$ units, as described in Section 2. We normalize $\tau_{jj} = 1\ \forall j$, and assume that the triangle inequality holds, i.e., $\tau_{ij} \leq \tau_{ik}\tau_{kj}\ \forall i, j, k$.

Due to perfect competition, country i charges for good ω its marginal cost of production and transportation to consumers in j, that is,

$$p_{ij}(\omega) = \frac{\tau_{ij} w_i}{z_i(\omega)},$$

and consumers pick the source which offers the lowest price. The productivity distribution coupled with this pricing behavior yields the CES price index[21]

$$P_j = \left[\sum_{k=1}^{N} T_k \left(w_k \tau_{kj}\right)^{-\theta}\right]^{-\frac{1}{\theta}},$$

and the gravity equation

$$\pi_{ij} = \frac{X_{ij}}{X_j} = \frac{T_i \left(w_i \tau_{ij}\right)^{-\theta}}{\sum_{k=1}^{N} T_k \left(w_k \tau_{kj}\right)^{-\theta}},$$

where π_{ij} is the so-called trade share, X_{ij} is the amount spent by country j's consumers on products from source i, and X_j is country j's total expenditure on final goods. In Section 3, we work with trade shares that are normalized by domestic absorption, that is,

$$\frac{\pi_{ij}}{\pi_{jj}} = \frac{T_i w_i^{-\theta} \tau_{ij}^{-\theta}}{T_j w_j^{-\theta}}.$$

To close the model, we assume that trade is balanced (i.e., every country has net exports equal to zero). As labor is the only source of income in this model, the set of wages satisfies the following condition

$$w_i L_i = \sum_{j=1}^{N} \pi_{ij} w_j L_j,$$

where we have $X_j = w_j L_j$.

Hat algebra. How does the equilibrium change after a shock to trade costs? Let the value of a variable x in the counterfactual equilibrium be equal to x', and let $\hat{x} := x'/x$. Using this "hat-notation," we can depict changes in outcomes as a function of relative trade costs, $\hat{\tau}_{ij}$, and initial expenditure levels, X_{ij}. This way we avoid calibrating the technology shifters, T_i. It is easy to show that the relevant set of equilibrium equations becomes

$$\begin{aligned} \pi'_{ij} &= \frac{\pi_{ij} \left(\hat{w}_i \hat{\tau}_{ij}\right)^{-\theta}}{\sum_{k=1}^{N} \pi_{kj} \left(\hat{w}_k \hat{\tau}_{kj}\right)^{-\theta}} \\ \hat{w}_i X_i &= \sum_{j=1}^{N} \pi'_{ij} \hat{w}_j X_j. \end{aligned}$$

After solving for the set of $\{\hat{w}_i\}_i$ that satisfies this equation system, we can calculate welfare changes (i.e., changes in real wages). More precisely, the gravity equation

[21] We abstract from constant terms which will be irrelevant to the entire discussion.

implies

$$\pi_{ii} = T_i \left(\frac{w_i}{P_i} \right)^{-\theta},$$

which means that welfare changes depend only on the trade elasticity, θ, and changes in the home share, π_{ii},[22]

$$\frac{\hat{w}_i}{\hat{P}_i} = (\hat{\pi}_{ii})^{-\frac{1}{\theta}}.$$

A useful benchmark is the real income loss from going to autarky (i.e., $\tau_{ij} \to \infty \; \forall i \neq j$). In that case we obtain

$$\frac{\hat{w}_i^a}{\hat{P}_i^a} = (\pi_{ii})^{\frac{1}{\theta}},$$

because we know that $\pi_{ii}^a = 1$, and thus $\hat{\pi}_{ii}^a = 1/\pi_{ii}$.

A.2 Multi-Sector EK

We use the multi-sector extension of the model above from Caliendo and Parro (2015). The only difference to Caliendo and Parro (2015) is that we do not explicitly model tariff revenues and trade imbalances to simplify exposition. Below we briefly sketch the setup.

Setup. We have K sectors, indexed by k and s, with sectoral input-output linkages. Labor is mobile across sectors. Households obtain utility from the consumption of final goods, C_i^k. Preferences are given by a Cobb-Douglas aggregator

$$u(C_i) = \prod_{k=1}^{K} \left(C_i^k \right)^{\alpha_i^k}, \quad \text{where} \; \sum_{k=1}^{K} \alpha_i^k = 1.$$

Firms produce varieties using two inputs, labor and a composite of intermediate goods. The value added share in sector k is denoted by $\gamma_i^k \geq 0$. The share of materials from sector k used in country i's sector s is denoted by $\gamma_i^{k,s} \geq 0$, with $\sum_{k=1}^{K} \gamma_i^{k,s} = 1 - \gamma_i^s$. The cost of an input bundle, c_i^k, is given by

$$c_i^k = w_i^{\gamma_i^k} \prod_{s=1}^{K} (P_i^s)^{\gamma_i^{s,k}},$$

[22] For more details on this type of counterfactuals in quantitative trade models, see Costinot and Rodríguez-Clare (2014). For a discussion on the equilibrium's uniqueness, see Alvarez and Lucas (2007).

where we again abstract from constant terms that will be irrelevant to the entire discussion.

Iceberg trade costs, τ_{ij}^k, and the parameters of the Fréchet distribution, T_i^k and θ^k, are sector-specific, which coupled with perfect competition yields the price of the composite intermediate good

$$P_j^k = \left[\sum_{i=1}^{N} T_i^k \left(c_i^k \tau_{ij}^k\right)^{-\theta^k}\right]^{-\frac{1}{\theta^k}},$$

and the gravity equation

$$\pi_{ij}^k = \frac{T_i^k \left(c_i^k \tau_{ij}^k\right)^{-\theta^k}}{\left(P_j^k\right)^{-\theta^k}}.$$

Finally, the aggregate price index is equal to

$$P_i = \prod_{k=1}^{K} \left(\frac{P_i^k}{\alpha_i^k}\right)^{\alpha_i^k},$$

and total household income, I_i, is given by

$$I_i = w_i L_i = \sum_{k=1}^{K} \gamma_i^k \sum_{j=1}^{N} \pi_{ij}^k X_j^k,$$

where X_j^k is total expenditure on goods from sector k, including households' final consumption and expenditure by firms on composite intermediate goods. Note that this model reduces to the one in Appendix A.1 if we set $K = 1$ and $\gamma_i = 1\ \forall i$.

Hat algebra. Analogous to the procedure above, we solve the model in relative changes after a shock to trade costs. The relevant system of equations is given by

$$\begin{aligned}
\hat{c}_i^k &= \hat{w}_i^{\gamma_i^k} \prod_{s=1}^{K} \left(\hat{P}_i^s\right)^{\gamma_i^{s,k}} \\
\hat{P}_j^k &= \left[\sum_{i=1}^{N} \pi_{ij}^k \left(\hat{c}_i^k \hat{\tau}_{ij}^k\right)^{-\theta^k}\right]^{-\frac{1}{\theta^k}} \\
\left(\pi_{ij}^k\right)' &= \pi_{ij}^k \frac{\left(\hat{c}_i^k \hat{\tau}_{ij}^k\right)^{-\theta^k}}{\left(\hat{P}_j^k\right)^{-\theta^k}} \\
\left(X_i^k\right)' &= \sum_{s=1}^{K} \gamma_i^{k,s} \sum_{j=1}^{N} \left(\pi_{ij}^s\right)' \left(X_j^s\right)' + \alpha_i^k \hat{w}_i I_i \\
\sum_{k=1}^{K} \sum_{j=1}^{N} \left(\pi_{ij}^k\right)' \left(X_j^k\right)' &- \left(\pi_{ji}^k\right)' \left(X_i^k\right)' = 0.
\end{aligned}$$

For a given guess of wage changes $\{\hat{w}_i\}$, we solve for (changes in) price indices and total expenditure, and then use the balanced trade condition to update the vector of relative wages, thereby adopting the tâtonnement algorithm from Alvarez and Lucas (2007). Relative changes in real wages, and thus welfare changes, are calculated as follows

$$\log\left(\frac{\hat{w}_i}{\hat{P}_i}\right) = -\sum_{k=1}^{K} \frac{\alpha_i^k}{\theta^k} \log\left(\hat{\pi}_{ii}^k\right) - \sum_{k=1}^{K} \frac{\alpha_i^k}{\theta^k} \frac{1-\gamma_i^k}{\gamma_i^k} \log\left(\hat{\pi}_{ii}^k\right) - \sum_{k=1}^{K} \frac{\alpha_i^k}{\gamma_i^k} \log\left(\prod_{s=1}^{K} \left[\frac{\hat{P}_i^s}{\hat{P}_i^k}\right]^{\gamma_i^{s,k}}\right).$$

For the multi-sector version, we require more data compared to the single-sector model above. Specifically, we need sector-specific trade shares, π_{ij}^k, input shares, $\gamma_i^{k,s}$, value added shares, γ_i^k, expenditure shares in final demand, α_i^k, and value added, $w_i L_i$. Following Costinot and Rodríguez-Clare (2014), we obtain all these values from the data described in Appendix B. The computation also requires sectoral trade elasticities, θ^k, which we take from Caliendo and Parro (2015) (see Appendix B).

B Data

Aggregate data. For the one-sector model, we use the International Trade and Production Database for Estimation (ITPD-E). The ITPD-E provides bilateral trade flows and production data retrieved from several administrative data sources (see Borchert, Larch, Shikher, and Yotov (2021) for details). Production data is used to calculate domestic trade as well as trade shares. The dataset covers the years 2000–2016, and up to 243 countries and 170 industries. For each country pair, we sum over all industries to obtain aggregate bilateral trade flows. We include only countries that appear as both importers and exporters in a given year, and we drop countries with a total population below one million. We merge China and Hong Kong, as well as Belgium and Luxembourg, which is often done in the literature due to large shares of re-exports. Our main results use the year 2004.

The gravity covariates such as distance and contiguity are taken from the widely-used CEPII Gravity database (Head, Mayer, and Ries, 2010).

Sector-level data. Despite its richness, the ITPD-E does not cover input-output linkages. Thus, our sector-level analysis uses the well-known 2013 release of the WIOD World Input-Output Tables (Timmer, Dietzenbacher, Los, Stehrer, and De Vries, 2015). This dataset covers 40 countries and a model for the "rest of the world" (ROW).

It includes 35 sectors and their input-output links over the period 1995–2011. We add Luxembourg to Belgium, and following Costinot and Rodríguez-Clare (2014), we merge several smaller countries to the ROW: Bulgaria, Cyprus, Estonia, Latvia, Lithuania, and Malta. According to Costinot and Rodríguez-Clare (2014), because of classification issues across countries, there are sectors for which countries have zero consumption and production. To avoid issues in this regard, we use the year 2008 and apply the same sector aggregation as Costinot and Rodríguez-Clare (2014), called "basic aggregation" in Table 3 of their online appendix (Costinot and Rodríguez-Clare, 2013). These steps leave us with 34 countries and 31 sectors. We use the sector-level trade elasticities that are listed in Table 3 of their online appendix (Costinot and Rodríguez-Clare, 2013), which were taken from Caliendo and Parro (2015, Table 1).

For the regressions, we use the median distance between a specific country in the sample and ROW-countries. The dummy variable indicating contiguity is set equal to zero between ROW and the other countries in the sample.

C Violations of the Triangle Inequality

In this section, we present violations of the TI found using Waugh's (2010) procedure to estimate τ, as discussed in Section 3. Let t_{ij} be the estimate of τ_{ij}. We calculate trade costs for the year 2004 using the data described in Appendix B, which gives us a sample of 127 countries. The regression results are reported in Table E.1.[23]

TI violations. For every pair, we plug t_{ij} into the TI, and verify whether this no-arbitrage condition is violated for at least one intermediary country. The results are presented in Table E.2. We find that 36% of trade relations could reduce their delivery costs by shipping via a third country, where most pairs have several potential intermediaries that are worth involving. Intuitively, the importers which would most benefit from an intermediary are geographically close to the top hubs (e.g., Switzerland and Vietnam). There are 34 relations where we obtain $t_{ij} < 1$, which would imply that trade costs could be pushed toward zero by re-shipping multiple times. But even when we set these t's equal to one, we see that 20% of relations would still have an arbitrage opportunity. This does not only concern small relations, as 28% of total trade is affected (when $t_{ij} \geq 1$). In Figure D.1, we plot the trade cost estimates against their

[23]The coefficients look very similar to those of Simonovska and Waugh (2014), who run the same regression on a sample with 123 countries in the year 2004.

"updated" values, where the latter are computed following Proposition 1. The potential savings for pairs with an arbitrage opportunity are substantial, with the average log-difference amounting to 23%.[24]

Are the violations driven by only a handful of hubs? In Panel B of Table E.2 we exclude the three intermediaries that would minimize trade costs for the most relations (Germany, Singapore, and Thailand), and rerun the regressions without them. The resulting estimates still feature many violations, despite a notable decline. Even when the ten top hubs are excluded—the U.S. being among them—about 6% of trade pairs would be better off involving an intermediary (cf. Panel C). It is important to note that these hubs are not only the intuitive suspects (such as the Netherlands or Singapore).

Robustness. We test different specifications to verify that the pattern we found is persistent. Table O.1 of the Online Appendix replicates the analysis in Table E.2, the only difference being that we add four dummy variables to the trade cost function which indicate whether a country pair (i) shares a common official language, (ii) a common currency, (iii) ever had a colonial link, or (iv) has a regional trade agreement. The results are very similar, with the share of TI violations slightly increasing.

In Table O.2 of the Online Appendix, we add the error term in equation (4) to the trade cost function. This table uses only trade relations with positive trade flows, as we cannot calculate trade costs for the remaining pairs in this specification.[25] We obtain even more TI violations when perfectly matching the observations with our trade cost estimates. In this specification, we may interpret the error term as a shock to trade costs rather than measurement error (Simonovska and Waugh, 2014).

In Table O.3 of the Online Appendix, we reproduce Table E.2 using PPML instead of OLS, as recommended by Santos Silva and Tenreyro (2006). The results are similar, although PPML appears to alleviate the issue. Especially the share of trade that is affected drops substantially. However, there are still implied arbitrage opportunities left even when we exclude the ten major hubs. The correlation between the OLS and PPML trade cost estimates is 0.89. Note that four of the ten top hubs in Table E.2 are

[24] Despite this relatively large average difference, one might still worry that potential savings are not statistically significantly different from zero. We can easily construct confidence bands around t_{ij} using the delta method. For pairs with an arbitrage opportunity, we find that the difference between indirect costs, $t_{ik}t_{kj}$, and direct costs, t_{ij}, is significantly smaller than zero at the five percent level in roughly 80% of the cases. Given that the literature focuses on point estimates and rarely works with confidence bands, we did not investigate this further.

[25] Alternatively, we could set $t_{ij} \to \infty$ for $\pi_{ij} = 0$. This would, however, mechanically increase the number of alleged arbitrage opportunities.

confirmed by this exercise. In line with OLS, results look similar (with slightly more violations) when we include additional covariates.

In Table O.4 of the Online Appendix, we see that a similar picture arises in other years. The share of violations appears to increase with the number of countries in the sample. This is intuitive: as more potential intermediaries become available, the TI is more likely to be violated.

D Figures

Figure D.1: Standard Trade Cost Estimates vs. Updated Trade Cost

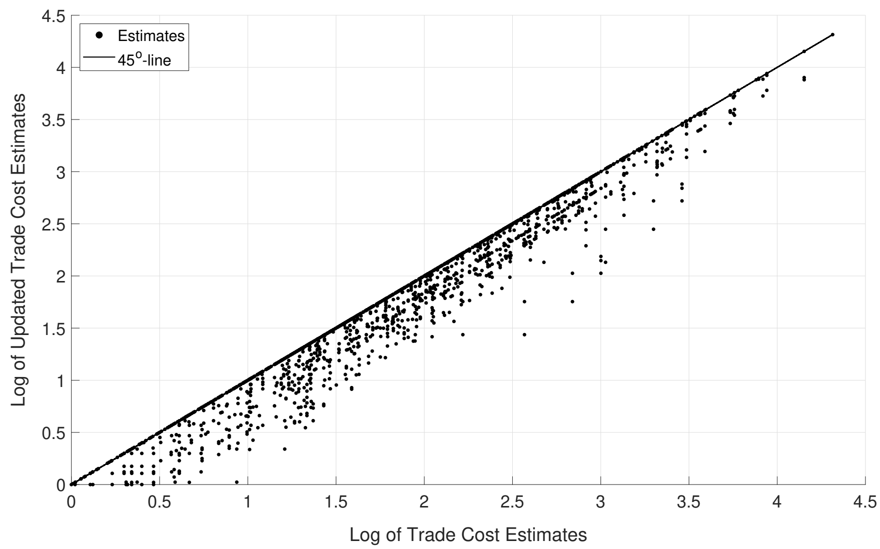

Notes. This figure depicts the log of estimated trade costs vs. their updated counterparts. As detailed in Appendix C, we estimate trade costs using Waugh's (2010) method. Whenever there is a violation of the triangle inequality, we obtain updated trade costs following Proposition 1. If an estimate falls below one, we set it equal to one (i.e., $t_{ij} \geq 1$).
Data. See Appendix B.
Graph. Authors' representation.

Figure D.2: Unconstrained vs. Constrained Values of ex_i

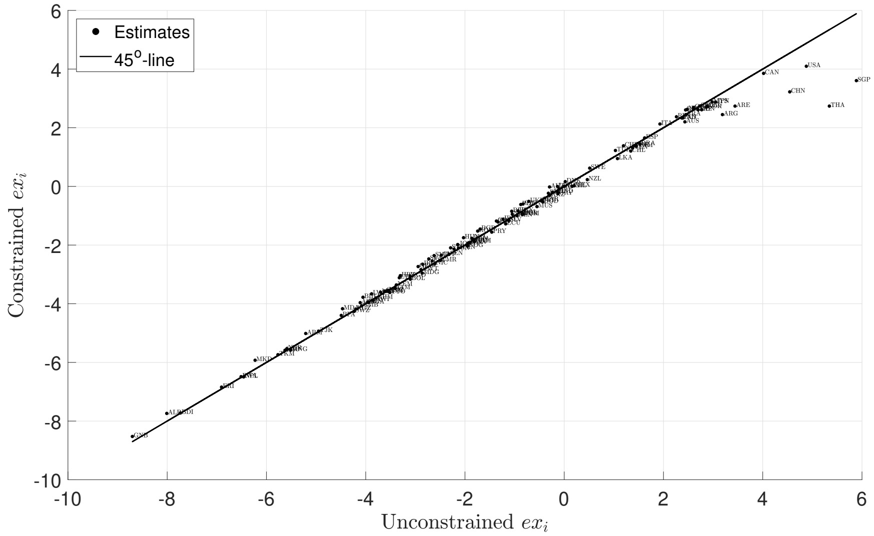

Notes. This figure depicts the unconstrained vs. constrained estimates of ex_i. As detailed in Section 3, we estimate trade costs via both unconstrained and constrained OLS. The regression is described by equation (4). The constrained version imposes inequalities (5) and (6) to ensure that estimates satisfy the TI. ex_i is the sum of country i's exporter and importer fixed effects.
Data. See Appendix B.
Graph. Authors' representation.

Figure D.3: Standard Trade Cost Estimates vs. Constrained Trade Cost

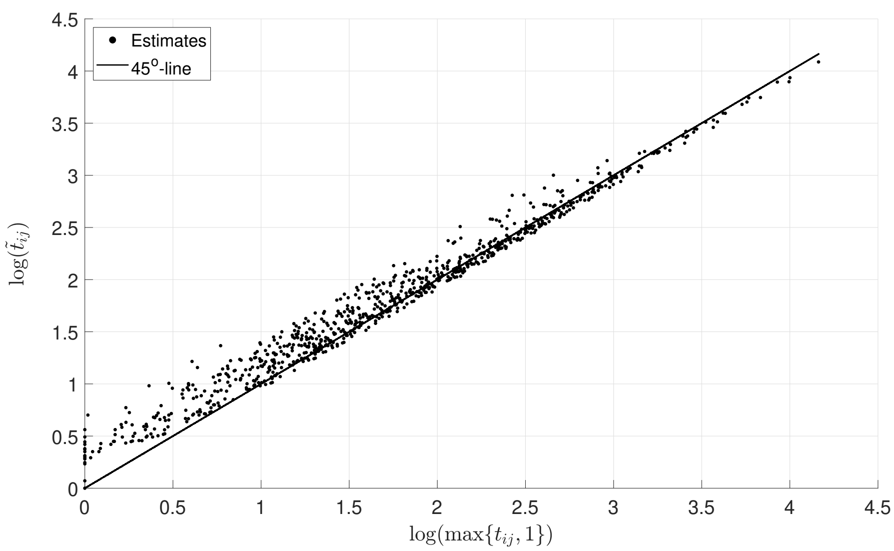

Notes. This figure depicts the log of estimated trade costs vs. their constrained counterparts. As detailed in Section 3, we estimate trade costs via both unconstrained and constrained OLS. The regression is described by equation (4). The constrained version imposes inequalities (5) and (6) to ensure that estimates satisfy the TI. t_{ij} ($\tilde{t}_{ij}$) are the unconstrained (constrained) estimates of τ_{ij}. If an unconstrained estimate falls below one, we set it equal to one (i.e., $t_{ij} \geq 1$).
Data. See Appendix B.
Graph. Authors' representation.

Figure D.4: Welfare Changes after 25% Drop in Global Trade Costs

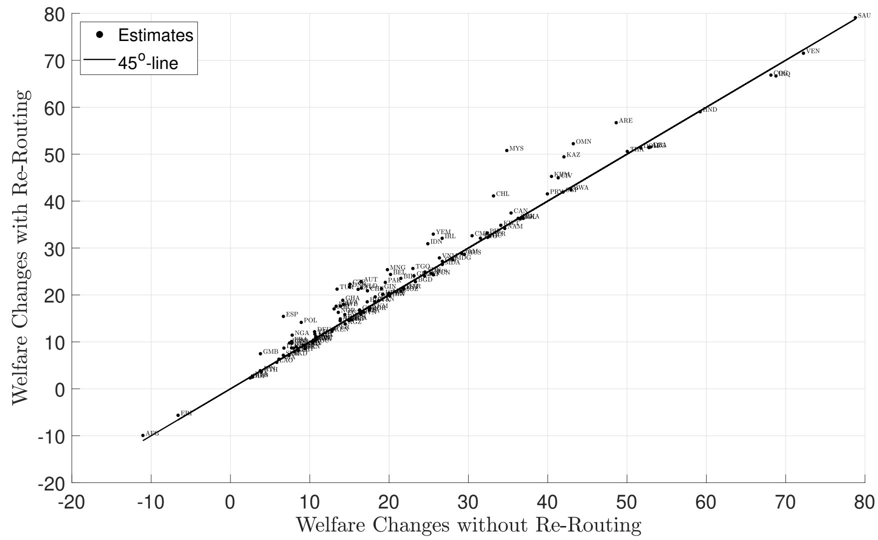

Notes. This figure depicts welfare changes (in %) after a drop in global trade costs by 25%. In one case, we do not allow for re-routing (standard approach), and in the other we allow for re-routing when the TI is violated. In both cases, we do not allow for trade costs to fall below one. The results are discussed in Section 4.
Data. See Appendix B.
Graph. Authors' representation.

Figure D.5: Welfare Changes after Eliminating Asymmetries in Trade Costs

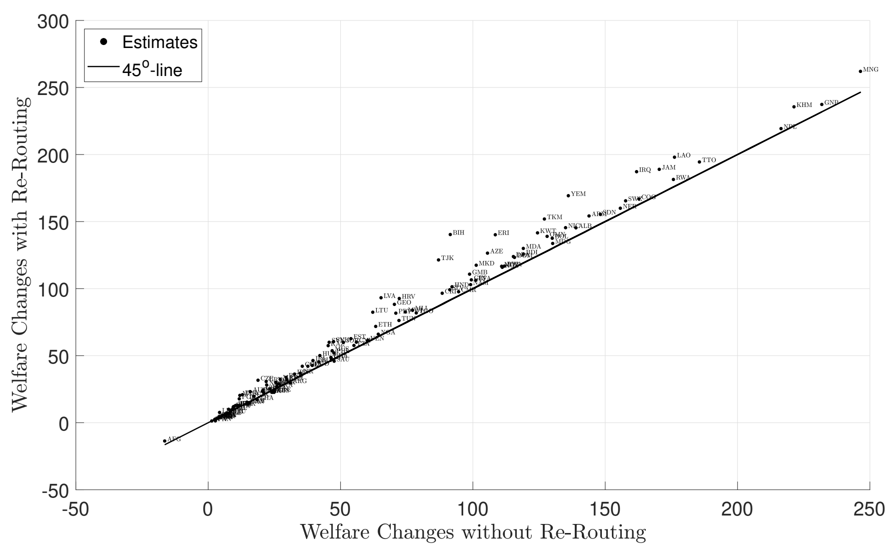

Notes. This figure depicts welfare changes (in %) after eliminating asymmetries in trade costs (i.e., setting $\tilde{t}_{ij} = \min\{\tilde{t}_{ij}, \tilde{t}_{ji}\}$). In one case, we do not allow for re-routing (standard approach), and in the other we allow for re-routing when the TI is violated. In both cases, we do not allow for trade costs to fall below one. The results are discussed in Section 4.
Data. See Appendix B.
Graph. Authors' representation.

E Tables

Table E.1: Standard Coefficients for Trade Cost Estimates

	OLS	PPML
Distance [0, 375)	-3.565***	-2.978***
	(0.352)	(0.621)
Distance [375, 750)	-4.323***	-3.099***
	(0.321)	(0.580)
Distance [750, 1500)	-5.415***	-3.964***
	(0.316)	(0.566)
Distance [1500, 3000)	-6.851***	-5.538***
	(0.311)	(0.582)
Distance [3000, 6000)	-7.861***	-6.246***
	(0.312)	(0.590)
Distance [6000, max.]	-8.532***	-6.455***
	(0.314)	(0.585)
Contiguity	1.260***	1.122***
	(0.118)	(0.215)
Nr. of observations	13,722	16,002

Notes. This table reports OLS and PPML results for the gravity equation (4) in Section 3. Robust standard errors are reported in parentheses. Regressions include importer and exporter fixed effects. Distance is weighted by population. Distance intervals are in miles. Trade flows are measured relative to domestic absorption. Dependent variable is absolute (log of) trade shares for PPML (OLS). The results are discussed in Appendix C.
*** significant at 1%-level.
Data. See Appendix B.
Results. Authors' computations.

Table E.2: TI Violations with Standard Trade Cost Estimates

	Baseline	Restrict $t \geq 1$
Panel A: Full sample		
% of trade pairs with arbitrage opportunity	36%	20%
... trade-weighted	51%	28%
Top 10 hubs (desc. order):		
DEU SGP THA ARG CHN ARE USA MYS GBR PAN		
Panel B: Excl. top 3 hubs		
% of trade pairs with arbitrage opportunity	23%	13%
Panel C: Excl. top 10 hubs		
% of trade pairs with arbitrage opportunity	10%	6%

Notes. This table reports violations of the triangle inequality implied by trade cost estimates, t, that we computed using Waugh's (2010) method. In the right-most column, we set $t = 1$ whenever we find $t < 1$. Panel A uses the full sample, the second row showing the share in total trade that has a potential arbitrage opportunity. Panels B and C exclude the three and ten main hubs, respectively, found in the full sample. A "hub" refers to an intermediary which allows for a cheaper delivery than the direct shipment between two countries. Details are given in Appendix C.
Data. See Appendix B.
Results. Authors' computations.

Table E.3: Constrained vs. Unconstrained Coefficients for Trade Cost Estimates

	Coefficients		%-effect on t_{ij} with $\theta = 4.14$	
	Unconstrained	Constrained	Uncons.	Cons.
Distance [0, 375)	-3.565	-4.479	136.6	195
Distance [375, 750)	-4.323	-4.768	184.1	216.4
Distance [750, 1500)	-5.415	-5.641	269.9	290.6
Distance [1500, 3000)	-6.851	-6.804	423.2	417.3
Distance [3000, 6000)	-7.861	-7.77	567.8	553.3
Distance [6000, max.]	-8.532	-8.396	685.3	659.9
Contiguity	1.26	0.579	-26.2	-13.1
R-squared	0.742	0.733		
Nr. of observations	13,722	13,722		

Notes. This table reports OLS results for the gravity equation (4) in Section 3 using the unconstrained (first and third columns) and constrained (second and fourth columns) versions. Regressions include importer and exporter fixed effects. The constrained version imposes inequalities (5) and (6) to ensure that estimates satisfy the TI. The left-most column replicates the first column in Table E.1. Distance is weighted by population. Distance intervals are in miles. Trade flows are measured relative to domestic absorption. Dependent variable is log of trade shares. t_{ij} is the estimate of τ_{ij}. The percentage effect on t_{ij} is equal to $100 \times (\exp(\text{coeff.}/(-\theta)) - 1)$. The method and results are discussed in Section 3.
Data. See Appendix B.
Results. Authors' computations.

Table E.4: TI Violations after 25% Drop in Global Trade Costs

	No restriction	Restrict $t \geq 1$
% of trade pairs with arbitrage opportunity	21%	13%
... trade-weighted	32%	24%

Distribution of potential savings:
Q1: 1%, Median: 5%, Mean: 8%, Q3: 15%, Max: 36%
Top 10 hubs (desc. order):
DEU SGP THA MYS ARE USA SAU FRA ITA CHN

Notes. This table reports violations of the triangle inequality that emerge after we decrease global bilateral trade costs by 25%. The right-most column sets the new trade cost values equal to one if they drop below one. The second row shows the share in total trade that has a potential arbitrage opportunity. The third row refers to reductions in trade costs from re-routing when there is an arbitrage opportunity (restricting $t \geq 1$). A "hub" refers to an intermediary which allows for a cheaper delivery than the direct shipment between two countries. The results are discussed in Section 4.
Data. See Appendix B.
Results. Authors' computations.

Table E.5: Welfare Changes after 25% Drop in Global Trade Costs (Selection)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
ARE	57%	49%	-45%	18%
AUT	23%	16%	-9%	68%
BEL	24%	20%	-13%	32%
BIH	24%	21%	-34%	6%
CAN	37%	35%	-22%	10%
CHE	21%	17%	-10%	35%
CHL	41%	33%	-18%	44%
CIV	45%	41%	-26%	14%
CMR	33%	30%	-15%	14%
CZE	22%	15%	-10%	72%
DNK	22%	15%	-8%	80%
ESP	15%	7%	-4%	208%
FRA	10%	8%	-5%	51%
GBR	10%	8%	-6%	34%
GHA	19%	14%	-15%	32%
GIN	21%	19%	-10%	22%
GMB	7%	4%	-29%	13%
IDN	31%	25%	-11%	56%
IRL	32%	27%	-12%	45%
ISR	18%	13%	-7%	59%
ITA	9%	7%	-4%	49%
JOR	17%	13%	-21%	19%
KAZ	49%	42%	-25%	29%
KHM	45%	40%	-21%	23%
LBN	10%	7%	-45%	5%
LTU	21%	16%	-12%	43%
LVA	18%	14%	-10%	36%
MNG	25%	20%	-15%	37%
MYS	51%	35%	-19%	84%
NGA	11%	8%	-3%	120%
NLD	21%	16%	-10%	51%
NOR	16%	14%	-6%	46%
OMN	52%	43%	-25%	36%
PAK	23%	20%	-18%	17%
POL	14%	9%	-6%	84%
PRY	42%	40%	-27%	6%
SWE	18%	14%	-7%	51%
TGO	26%	23%	-16%	17%
TUR	21%	13%	-9%	88%
YEM	33%	26%	-15%	49%
Average	26%	21%	-16%	45%

Notes. This table reports welfare changes after a decrease in global bilateral trade costs of 25%. The table is restricted to the 40 countries with the highest difference between Columns (1) and (2). Column (1) allows for re-routing after the shock. Column (3) reports the real income loss from going to full autarky (i.e., $\tau_{ij} \to \infty \ \forall i \neq j$). Column (4) shows the difference between Columns (1) and (2) scaled by the absolute value of Column (3). See Section 4 for further details.
Data. See Appendix B.
Results. Authors' computations.

Table E.6: Welfare Changes after Drop in Global Trade Costs (Averages)

Drop in trade costs	Distribution of net welfare changes from re-routing					
	Min.	Q1	Q2	Q3	Max.	Mean (scaled by net autarky losses)
5%	0%	0%	0%	0%	0.4%	0.2%
10%	0%	0%	0%	0.1%	1.9%	0.8%
15%	0%	0%	0.1%	0.3%	5.7%	2.6%
25%	0%	0.4%	0.7%	2.4%	15.9%	17.5%
40%	0%	1.8%	9.3%	18.3%	80.3%	108%
50%	0.4%	9.8%	22.5%	51%	249.9%	268.7%

Notes. This table reports the distribution of net welfare changes from re-routing after a decrease in global bilateral trade costs by 5%, 10%, 15%, 25%, 40%, or 50%. Q1, Q2, Q3 refer, respectively, to the first, second and third quartiles. The right-most column reports the average net real income change relative to the net loss from going to full autarky (i.e., $\tau_{ij} \to \infty \ \forall i \neq j$). See Section 4 for further details.
Data. See Appendix B.
Results. Authors' computations.

Table E.7: TI Violations after Eliminating Asymmetries in Trade Costs

	No restriction	Restrict $t \geq 1$
% of trade pairs with arbitrage opportunity	58%	58%
. . . trade-weighted	7%	7%

Distribution of potential savings:
Q1: 24%, Median: 45%, Mean: 50%, Q3: 70%, Max: 197%
Top 10 hubs (desc. order):
DEU USA SGP SAU ARE ZAF CAN CHN ARG ITA

Notes. This table reports violations of the triangle inequality that emerge after we eliminate asymmetries in trade costs (i.e., we set $\tilde{t}_{ij} = \min\{\tilde{t}_{ij}, \tilde{t}_{ji}\}$). The right-most column sets the new trade cost values equal to one if they drop below one. The second row shows the share in total trade that has a potential arbitrage opportunity. The third row refers to reductions in trade costs from re-routing when there is an arbitrage opportunity (restricting $t \geq 1$). A "hub" refers to an intermediary which allows for a cheaper delivery than the direct shipment between two countries. See Section 4 for further details.
Data. See Appendix B.
Results. Authors' computations.

Table E.8: Welfare Changes after Eliminating Asymmetries in Trade Costs (Selection)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|3|
ARM	154%	144%	-17%	61%
AZE	126%	106%	-14%	145%
BIH	140%	91%	-34%	142%
CRI	97%	88%	-12%	68%
CZE	32%	19%	-10%	128%
DNK	21%	13%	-8%	95%
DZA	124%	115%	-29%	30%
ERI	140%	108%	-15%	207%
EST	63%	54%	-12%	71%
ETH	72%	63%	-13%	67%
GEO	88%	70%	-25%	71%
GMB	111%	99%	-29%	42%
HND	101%	92%	-41%	23%
HRV	93%	72%	-12%	175%
HUN	50%	42%	-11%	69%
IRQ	187%	162%	-39%	65%
JAM	189%	170%	-24%	78%
JOR	60%	51%	-21%	42%
KAZ	57%	45%	-25%	48%
KHM	236%	221%	-21%	68%
KWT	142%	124%	-10%	178%
LAO	198%	176%	-17%	127%
LTU	82%	62%	-12%	171%
LVA	93%	65%	-10%	267%
MDA	130%	119%	-21%	53%
MKD	117%	101%	-7%	238%
MNG	262%	246%	-15%	104%
MWI	83%	74%	-7%	111%
MYS	20%	12%	-19%	44%
NIC	145%	135%	-36%	29%
OMN	139%	128%	-25%	44%
PRY	82%	71%	-27%	40%
SLV	99%	91%	-44%	18%
SVK	60%	47%	-12%	113%
SVN	60%	46%	-11%	136%
TJK	121%	87%	-10%	331%
TKM	152%	127%	-7%	368%
TTO	194%	186%	-14%	63%
URY	31%	22%	-15%	61%
YEM	169%	136%	-15%	218%
Average	113%	97%	-19%	110%

Notes. This table reports welfare changes after eliminating asymmetries in trade costs (i.e., we set $\tilde{t}_{ij} = \min\{\tilde{t}_{ij}, \tilde{t}_{ji}\}$). The table is restricted to the 40 countries with the highest difference between Columns (1) and (2). Column (1) allows for re-routing after the shock. Column (3) reports the real income loss from going to full autarky (i.e., $\tau_{ij} \to \infty \; \forall i \neq j$). Column (4) shows the difference between Columns (1) and (2) scaled by the absolute value of Column (3). See Section 4 for further details.
Data. See Appendix B.
Results. Authors' computations.

Table E.9: Welfare Changes after 25% Drop in Global Trade Costs (Multi-Sector)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
AUS	16%	15%	-16%	3%
AUT	30%	24%	-53%	11%
BEL	36%	33%	-56%	6%
BRA	12%	10%	-7%	20%
CAN	18%	15%	-32%	8%
CHN	13%	12%	-8%	22%
CZE	40%	33%	-36%	21%
DEU	18%	16%	-22%	6%
DNK	26%	20%	-41%	14%
ESP	16%	12%	-19%	21%
FIN	29%	25%	-21%	19%
FRA	14%	10%	-18%	17%
GBR	18%	14%	-23%	18%
GRC	16%	16%	-24%	1%
HUN	41%	34%	-56%	13%
IDN	23%	17%	-13%	48%
IND	13%	11%	-10%	26%
IRL	58%	42%	-38%	44%
ITA	15%	12%	-17%	17%
JPN	8%	6%	-5%	40%
KOR	23%	19%	-14%	26%
MEX	17%	14%	-19%	18%
NLD	28%	24%	-43%	8%
POL	26%	21%	-35%	15%
PRT	23%	21%	-39%	7%
ROU	23%	19%	-27%	14%
RUS	20%	18%	-38%	5%
ROW	15%	13%	-30%	4%
SVK	44%	38%	-53%	11%
SVN	36%	29%	-60%	11%
SWE	28%	23%	-26%	19%
TUR	14%	11%	-28%	12%
TWN	28%	28%	-18%	-4%
USA	7%	6%	-8%	12%
Average	23%	19%	-28%	16%

Notes. This table reports welfare changes after a decrease in global bilateral sector-level trade costs by 25%, calculated using the multi-sector model. Column (1) allows for re-routing after the shock. Column (3) reports the real income loss from going to full autarky (i.e., $\tau_{ij}^k \to \infty \; \forall i \neq j$, and k). Column (4) shows the difference between Columns (1) and (2) scaled by the absolute value of Column (3). See Section 4 for further details.
Data. See Appendix B.
Results. Authors' computations.

Table E.10: Welfare Changes after Eliminating Asymmetries in Trade Costs (Multi-Sector)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|3|
AUS	106%	21%	-16%	529%
AUT	126%	51%	-53%	143%
BEL	143%	51%	-56%	164%
BRA	143%	22%	-7%	1738%
CAN	137%	25%	-32%	354%
CHN	127%	18%	-8%	1387%
CZE	237%	47%	-36%	525%
DEU	109%	46%	-22%	282%
DNK	150%	40%	-41%	269%
ESP	88%	20%	-19%	347%
FIN	157%	34%	-21%	577%
FRA	74%	21%	-18%	290%
GBR	87%	38%	-23%	217%
GRC	178%	38%	-24%	583%
HUN	280%	59%	-56%	397%
IDN	195%	47%	-13%	1114%
IND	128%	29%	-10%	994%
IRL	199%	76%	-38%	320%
ITA	87%	20%	-17%	386%
JPN	56%	13%	-5%	920%
KOR	127%	26%	-14%	730%
MEX	176%	43%	-19%	701%
NLD	120%	49%	-43%	165%
POL	163%	42%	-35%	349%
PRT	185%	45%	-39%	363%
ROU	313%	70%	-27%	882%
RUS	297%	58%	-38%	624%
ROW	374%	353%	-30%	69%
SVK	327%	83%	-53%	456%
SVN	367%	133%	-60%	388%
SWE	134%	34%	-26%	379%
TUR	175%	56%	-28%	432%
TWN	120%	29%	-18%	500%
USA	56%	27%	-8%	352%
Average	169%	52%	-28%	527%

Notes. This table reports welfare changes after eliminating asymmetries in sector-level trade costs (i.e., we set $\tilde{t}^k_{ij} = \min\{\tilde{t}^k_{ij}, \tilde{t}^k_{ji}\}$), calculated using the multi-sector model. Column (1) allows for re-routing after the shock. Column (3) reports the real income loss from going to full autarky (i.e., $\tau^k_{ij} \to \infty$ $\forall i \neq j$, and k). Column (4) shows the difference between Columns (1) and (2) scaled by the absolute value of Column (3). See Section 4 for further details.

Data. See Appendix B.

Results. Authors' computations.

Online Appendix

O.A Figures

Figure O.1: Updated Trade Cost Estimates vs. Constrained Trade Cost

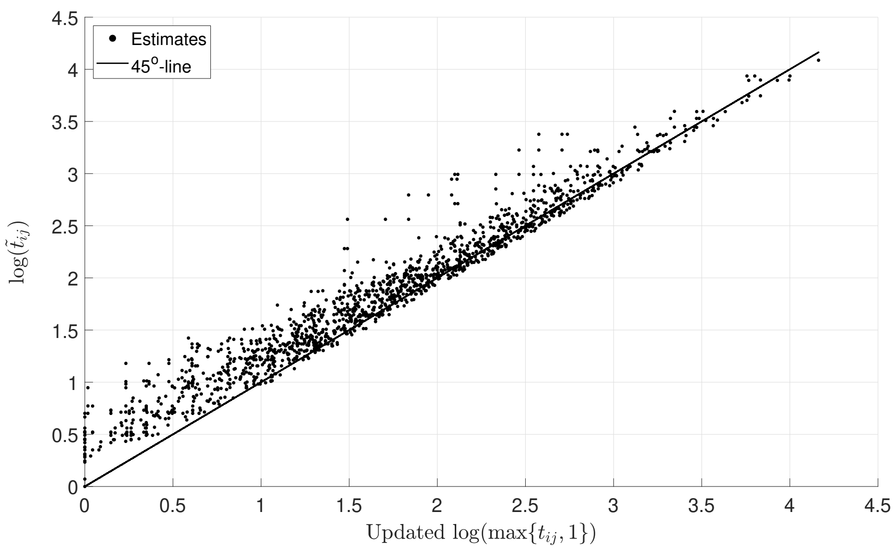

Notes. This figure depicts the log of "updated" unconstrained trade costs vs. their constrained counterparts. As detailed in Section 3, we estimate trade costs via both unconstrained and constrained OLS. The regression is depicted by equation (4). The constrained version imposes inequalities (5) and (6) to ensure that estimates satisfy the TI. t_{ij} ($\tilde{t}_{ij}$) are the unconstrained (constrained) estimates of τ_{ij}. If an unconstrained estimate falls below one, we set it equal to one (i.e., $t_{ij} \geq 1$). Whenever there is a violation of the triangle inequality, we obtain updated trade costs following Proposition 1.
Data. See Appendix B.
Graph. Authors' representation.

O.B Tables

Table O.1: Additional Covariates: TI Violations with Standard Trade Cost Estimates

	Baseline	Restrict $t \geq 1$
Panel A: Full sample		
% of trade pairs with arbitrage opportunity	41%	27%
... trade-weighted	62%	40%
Top 10 hubs (desc. order):		
DEU SGP THA ARG BEL ARE SAU USA GBR MYS		
Panel B: Excl. top 3 hubs		
% of trade pairs with arbitrage opportunity	28%	18%
Panel C: Excl. top 10 hubs		
% of trade pairs with arbitrage opportunity	15%	10%

Notes. This table reproduces Table E.2 adding four variables to the regressions that indicate whether a country pair (i) shares a common official language, (ii) a common currency, (iii) ever had a colonial link, or (iv) has a regional trade agreement. Details are given in the notes to Table E.2.

Table O.2: Error Term in τ: TI Violations with Standard Trade Cost Estimates

	Baseline	Restrict $t \geq 1$
Panel A: Full sample		
% of trade pairs with arbitrage opportunity	50%	48%
. . . trade-weighted	11%	9%
Top 10 hubs (desc. order):		
SGP ARE SAU THA RUS ZAF PAN ARG VEN USA		
Panel B: Excl. top 3 hubs		
% of trade pairs with arbitrage opportunity	47%	46%
Panel C: Excl. top 10 hubs		
% of trade pairs with arbitrage opportunity	39%	39%

Notes. This table reproduces Table E.2 but adding the error term in the regression equation (4) to the trade cost function (equation (3)). The table includes only trade pairs with positive trade flows, while the remaining trade relations are omitted. Details are given in the notes to Table E.2.

Table O.3: PPML: TI Violations with Standard Trade Cost Estimates

	Baseline	Restrict $t \geq 1$
Panel A: Full sample		
% of trade pairs with arbitrage opportunity	26%	16%
. . . trade-weighted	7%	5%
Top 10 hubs (desc. order):		
SGP ZAF SAU PAN THA ARE RUS KAZ KEN VEN		
Panel B: Excl. top 3 hubs		
% of trade pairs with arbitrage opportunity	10%	6%
Panel C: Excl. top 10 hubs		
% of trade pairs with arbitrage opportunity	2%	1%

Notes. This table reproduces Table E.2 using PPML instead of OLS. Details are given in the notes to Table E.2.

Table O.4: By Year: TI Violations with Standard Trade Cost Estimates

Year	% of trade pairs with arbitrage opportunity		Nr. of countries
	Baseline	Restrict $t \geq 1$	
2000	22%	13%	122
2001	28%	17%	124
2002	26%	14%	125
2003	27%	16%	126
2004	36%	20%	127
2005	37%	21%	127
2006	39%	22%	129
2007	37%	22%	131
2008	43%	27%	132
2009	43%	26%	133
2010	38%	23%	135
2011	47%	28%	136
2012	57%	36%	137
2013	61%	39%	137
2014	59%	37%	133

Notes. This table reproduces the top row in Table E.2 for several years. Details are given in the notes to Table E.2.

Table O.5: Welfare Changes after 25% Drop in Global Trade Costs (Single-Sector – WIOD Sample)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
AUS	6%	6%	-3%	0%
AUT	14%	13%	-7%	19%
BEL	19%	18%	-9%	12%
BRA	4%	4%	-2%	14%
CAN	10%	10%	-5%	8%
CHN	5%	5%	-2%	4%
CZE	15%	14%	-7%	23%
DEU	11%	11%	-5%	-1%
DNK	14%	13%	-7%	18%
ESP	6%	6%	-4%	0%
FIN	11%	11%	-5%	1%
FRA	6%	6%	-4%	0%
GBR	7%	7%	-4%	-1%
GRC	6%	6%	-5%	0%
HUN	17%	17%	-10%	-1%
IDN	7%	7%	-4%	0%
IND	5%	5%	-3%	2%
IRL	20%	20%	-10%	0%
ITA	6%	6%	-3%	0%
JPN	4%	4%	-2%	0%
KOR	10%	10%	-5%	0%
MEX	8%	8%	-4%	0%
NLD	16%	15%	-7%	8%
POL	9%	9%	-5%	1%
PRT	8%	8%	-5%	1%
ROU	7%	7%	-5%	0%
RUS	6%	6%	-3%	0%
ROW	10%	10%	-6%	0%
SVK	16%	16%	-9%	-1%
SVN	14%	14%	-8%	-1%
SWE	12%	12%	-6%	1%
TUR	5%	5%	-4%	1%
TWN	16%	16%	-8%	0%
USA	3%	3%	-2%	0%
Average	10%	10%	-5%	3%

Notes. This table reproduces Table E.5 using the WIOD data instead of the ITPD-E data, and thus reducing the sample to 34 countries. Details are given in the notes to Table E.5.

Table O.6: Welfare Changes after Eliminating Asymmetries in Trade Costs (Single-Sector – WIOD Sample)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
AUS	29%	4%	-3%	915%
AUT	45%	13%	-7%	467%
BEL	34%	7%	-9%	288%
BRA	29%	3%	-2%	1437%
CAN	31%	4%	-5%	609%
CHN	13%	2%	-2%	498%
CZE	53%	14%	-7%	539%
DEU	16%	8%	-5%	154%
DNK	58%	16%	-7%	617%
ESP	22%	3%	-4%	501%
FIN	57%	9%	-5%	897%
FRA	18%	3%	-4%	408%
GBR	19%	3%	-4%	409%
GRC	68%	20%	-5%	953%
HUN	63%	14%	-10%	500%
IDN	36%	6%	-4%	844%
IND	35%	5%	-3%	1036%
IRL	63%	15%	-10%	508%
ITA	20%	4%	-3%	448%
JPN	12%	3%	-2%	457%
KOR	20%	5%	-5%	287%
MEX	28%	6%	-4%	545%
NLD	35%	7%	-7%	372%
POL	37%	8%	-5%	544%
PRT	67%	14%	-5%	998%
ROU	61%	13%	-5%	892%
RUS	30%	6%	-3%	853%
ROW	44%	60%	-6%	-290%
SVK	83%	23%	-9%	655%
SVN	108%	45%	-8%	757%
SWE	44%	9%	-6%	581%
TUR	42%	9%	-4%	951%
TWN	35%	8%	-8%	347%
USA	6%	2%	-2%	194%
Average	40%	11%	-5%	593%

Notes. This table reproduces Table E.8 using the WIOD data instead of the ITPD-E data, and thus reducing the sample to 34 countries. Details are given in the notes to Table E.8.

Table O.7: Welfare Changes after 25% Drop in Global Trade Costs (Single-Sector – WIOD Sample excl. ROW)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
AUS	5%	5%	-2%	0%
AUT	13%	11%	-6%	23%
BEL	17%	16%	-9%	13%
BRA	2%	2%	-1%	0%
CAN	9%	9%	-4%	0%
CHN	4%	4%	-1%	0%
CZE	14%	13%	-7%	26%
DEU	10%	10%	-5%	-1%
DNK	11%	10%	-6%	21%
ESP	5%	5%	-3%	0%
FIN	9%	9%	-5%	0%
FRA	5%	5%	-3%	0%
GBR	6%	6%	-3%	-1%
GRC	3%	3%	-4%	0%
HUN	15%	15%	-9%	-1%
IDN	5%	5%	-2%	0%
IND	3%	3%	-1%	0%
IRL	18%	18%	-9%	0%
ITA	5%	5%	-3%	-1%
JPN	3%	3%	-1%	0%
KOR	8%	8%	-4%	0%
MEX	7%	7%	-4%	0%
NLD	15%	14%	-6%	10%
POL	8%	8%	-5%	0%
PRT	7%	7%	-5%	0%
ROU	6%	6%	-5%	-1%
RUS	5%	5%	-2%	0%
SVK	15%	15%	-8%	-2%
SVN	12%	12%	-7%	-1%
SWE	9%	10%	-5%	-1%
TUR	4%	4%	-2%	0%
TWN	14%	14%	-5%	0%
USA	2%	2%	-2%	0%
Average	8%	8%	-4%	2%

Notes. This table reproduces Table O.5 excluding “rest-of-the-world” (ROW) from the sample. Details are given in the notes to Table O.5.

Table O.8: Welfare Changes after Eliminating Asymmetries in Trade Costs (Single-Sector – WIOD Sample excl. ROW)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
AUS	1%	1%	-2%	0%
AUT	9%	7%	-6%	38%
BEL	4%	3%	-9%	11%
BRA	1%	1%	-1%	0%
CAN	3%	3%	-4%	0%
CHN	2%	2%	-1%	0%
CZE	11%	8%	-7%	46%
DEU	8%	8%	-5%	-3%
DNK	8%	6%	-6%	47%
ESP	1%	1%	-3%	0%
FIN	4%	4%	-5%	2%
FRA	1%	1%	-3%	2%
GBR	2%	2%	-3%	-1%
GRC	5%	4%	-4%	5%
HUN	7%	7%	-9%	0%
IDN	2%	2%	-2%	0%
IND	1%	1%	-1%	0%
IRL	8%	8%	-9%	3%
ITA	1%	1%	-3%	0%
JPN	1%	1%	-1%	0%
KOR	2%	2%	-4%	0%
MEX	4%	4%	-4%	0%
NLD	4%	4%	-6%	10%
POL	4%	4%	-5%	8%
PRT	7%	7%	-5%	6%
ROU	5%	5%	-5%	6%
RUS	2%	2%	-2%	0%
SVK	14%	14%	-8%	1%
SVN	22%	21%	-7%	4%
SWE	3%	3%	-5%	1%
TUR	2%	2%	-2%	8%
TWN	4%	4%	-5%	0%
USA	1%	1%	-2%	0%
Average	5%	4%	-4%	6%

Notes. This table reproduces Table O.6 excluding “rest-of-the-world” (ROW) from the sample. Details are given in the notes to Table O.6.

Table O.9: Welfare Changes after 25% Drop in Global Trade Costs (Multi-Sector – excl. ROW)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|(3)|
AUS	15%	15%	-13%	1%
AUT	27%	22%	-52%	10%
BEL	32%	29%	-55%	6%
BRA	11%	10%	-5%	15%
CAN	17%	14%	-31%	9%
CHN	12%	11%	-7%	27%
CZE	36%	30%	-35%	18%
DEU	16%	15%	-21%	7%
DNK	24%	18%	-40%	15%
ESP	16%	12%	-18%	21%
FIN	25%	21%	-20%	19%
FRA	13%	10%	-17%	16%
GBR	17%	13%	-22%	17%
GRC	13%	12%	-23%	1%
HUN	39%	32%	-55%	13%
IDN	22%	16%	-9%	73%
IND	13%	10%	-8%	39%
IRL	56%	40%	-37%	44%
ITA	14%	11%	-16%	18%
JPN	7%	5%	-3%	59%
KOR	21%	17%	-12%	34%
MEX	17%	14%	-18%	14%
NLD	26%	22%	-42%	8%
POL	24%	19%	-34%	14%
PRT	23%	21%	-38%	4%
ROU	21%	17%	-26%	13%
RUS	18%	16%	-36%	5%
SVK	40%	35%	-51%	11%
SVN	31%	26%	-58%	9%
SWE	24%	19%	-24%	18%
TUR	13%	10%	-26%	11%
TWN	26%	28%	-16%	-8%
USA	6%	5%	-8%	13%
Average	22%	18%	-26%	17%

Notes. This table reproduces Table E.9 excluding "rest-of-the-world" (ROW) from the sample. Details are given in the notes to Table E.9.

Table O.10: Welfare Changes after Eliminating Asymmetries in Trade Costs (Multi-Sector – excl. ROW)

Country	(1) With re-routing	(2) No re-routing	(3) Autarky loss	(4) ((1)-(2))/|3|
AUS	12%	8%	-13%	28%
AUT	50%	37%	-52%	24%
BEL	52%	47%	-55%	9%
BRA	14%	7%	-5%	131%
CAN	34%	20%	-31%	45%
CHN	11%	8%	-7%	44%
CZE	72%	29%	-35%	123%
DEU	51%	44%	-21%	30%
DNK	44%	26%	-40%	44%
ESP	19%	11%	-18%	42%
FIN	32%	15%	-20%	82%
FRA	21%	15%	-17%	38%
GBR	35%	30%	-22%	23%
GRC	37%	16%	-23%	94%
HUN	70%	32%	-55%	69%
IDN	21%	18%	-9%	39%
IND	16%	8%	-8%	107%
IRL	95%	70%	-37%	67%
ITA	21%	13%	-16%	50%
JPN	7%	5%	-3%	54%
KOR	16%	13%	-12%	28%
MEX	49%	33%	-18%	89%
NLD	67%	57%	-42%	25%
POL	40%	23%	-34%	50%
PRT	48%	25%	-38%	61%
ROU	65%	30%	-26%	134%
RUS	39%	17%	-36%	61%
SVK	98%	39%	-51%	116%
SVN	89%	43%	-58%	80%
SWE	31%	17%	-24%	61%
TUR	46%	33%	-26%	48%
TWN	19%	16%	-16%	16%
USA	25%	24%	-8%	19%
Average	41%	25%	-26%	59%

Notes. This table reproduces Table E.10 excluding "rest-of-the-world" (ROW) from the sample. Details are given in the notes to Table E.10.