

Only words? How Power in Trade Agreement Texts Affects International Trade Flows

Abstract

The texts of international trade agreements are often copy-and-pasted from previous treaties. Large and rich countries are more likely to act as “rule-makers” who have an internally consistent treaty network, while smaller and poorer countries are more likely to be “rule-takers” whose treaties resemble those of their respective partners, but not their own. But does the copy-and-pasting of trade agreements affect international trade flows? I use a new text corpus of machine-readable trade agreement texts, Texts of Trade Agreements (ToTA), to study whether rule-makers benefit more from trade agreements than rule-takers. I build indicators of textual similarity between agreements that capture how much an agreement resembles each party’s previous agreements and introduce them into a gravity model of international trade. I find that the exports of countries that had a greater influence on the treaty text increase more than their partners’, which confirms the hypothesis that rule-makers benefit more than proportionally from trade agreements. In an “average” power relation, the rule-maker increases its exports by about 15.8%, while the rule-taker only increases its exports by about 9.3%. The effect is driven by the power of the importer, suggesting that powerful countries effectively restrict access to their markets. This result is robust to controlling for tariffs, difference in GDP and GDP per capita, and to including trade agreement fixed effects. A regression with lags and leads of the trade policy variables suggests that, unlike the agreement per se, the balance of power within the agreement is not endogenous to anticipated trade flows. The results highlight that the copy-and-pasting of trade agreements has real economic effects, suggesting that countries should carefully assess text proposals when negotiating a trade agreement.

“The usual view is that international law is a check on state interests, causing a state to behave in a way contrary to its interests. In our view, the causal relationship between international law and state interests runs in the opposite direction. International law emerges from states’ pursuit of self-interested policies on the international stage. International law is, in this sense, endogenous to state interests. It is not a check on state self-interest; it is a product of state self-interest.”

Jack L. Goldsmith and Eric A. Posner
The Limits of International Law
Oxford University Press, 2005

1 Introduction

The texts of preferential trade agreements (PTAs) are often copy-and-pasted from previous treaties. To cite only the most prominent example, significant elements of the 1994 North-American Free Trade Agreement (NAFTA) have been included in more than 50 subsequent trade agreements signed by American and East Asian countries (Alschner et al. (2017b)). Some large countries, such as the United States or the European Union (EU), have highly consistent trade agreement networks, as their agreements with different partners closely resemble each other. They tend to act as “rule-makers” in international trade, while smaller countries are more likely to be “rule-takers”. But does this matter for the impact of trade agreements on international trade flows? This paper is the first to address the economic consequences of the copy-and-pasting of trade agreements, and it provides evidence that countries with greater influence on agreement texts experience larger increases in exports than their trading partners once the agreement enters into force.

I use a new corpus of machine-readable trade agreement texts, Texts of Trade Agreements (ToTA), to build text-as-data indicators of power in trade agreements. To capture the extent to which a country dominates a treaty, I compute the share of the text that is copy-and-pasted from earlier treaties concluded by the same country. My main variable of interest is relative power, which I define as the difference between the exporter’s and the importer’s influence on the text of the treaty in force between them.¹ I introduce this measure into a gravity model of international trade to assess the impact of power on exports, and find that this impact is positive. When a trade agreement resembles one partner’s previous agreements more than the other partner’s, the dominant country benefits from a more than proportional increase in exports. A number of robustness checks allow to corroborate this result and to gain further insights into the underlying mechanism. First, the result appears to be driven by the power of the importer, suggesting that rule-makers use their power in trade negotiations to restrict access to their markets. Second, results from including lags and leads of policy variables suggest that, while the conclusion of a trade agreement is likely to be endogenous to expected increases in trade flows, this is not the case for the power balance within the agreement.

It should be noted that, from an economic perspective, net exports should not necessarily be the main outcome of interest of a welfare-maximizing government. Imports can benefit both domestic producers and consumers, and modern trade agreements affect many other outcomes, such as the terms of trade, investment, trade in services and global value chain integration. Depending on the context, these may be more important for national welfare than gross trade flows. However, the public discourse about trade deficits suggests that increasing net exports is viewed as desirable by many politicians. As Rodrik (2018) formulates it in his discussion of trade agreements, “nations like it when their exports grow, but not so much when their imports expand. Effectively, nations trade market access: more of your imports in return for more of my exports”. If we accept the assumption that a central purpose of trade agreements is to increase trade between its members, and that policy-makers principally care about increasing their country’s exports, my findings imply that, in the negotiation of trade agreement texts, some countries are more able than others to exert power in a way that serves their interests.

¹Using asymmetric similarity variables that vary at the exporter-importer-year level represents a novel contribution to the literature on the text-as-data analysis of international economic agreements, which has so far only used unidirectional Jaccard similarities at the treaty-pair level.

In sum, the copy-and-pasting of trade agreements is not simply an efficient way of concluding trade negotiations more quickly, or an expression of a process of learning that benefits all partners equally. It has real economic effects, and it increases the exports of the dominant country more than proportionally. This finding is of particular relevance for small developing countries that negotiate trade agreements with larger partners. As they have less resources and bargaining power, they may be more likely to accept a trade agreement text proposed by a larger negotiating partner. If the text chosen by larger partners reflects legal provisions that are particularly advantageous to them, this may happen at the expense of smaller partners.

2 Literature Review

By assessing the impact of copy-and-pasting on international trade, this paper contributes to the economic literature on the impact of trade agreements as well as to the legal and political science literature on the design of trade agreement texts. With respect to the economic literature, I focus on how empirical applications have accounted for the heterogeneity of PTAs, as a key contribution of this article is to introduce new text-based variables that describe trade agreement characteristics.

Early applications studying the impact of trade agreements used dummy variables that captured the existence of a treaty between two countries in a given year without further classifying it. As trade agreements increased in number and complexity, the empirical literature started paying attention to differences in agreement design. Scholars have used variables capturing the type of agreement (e.g., Baier et al. (2015), who distinguish between free trade agreement, customs union, common market and economic union), binary variables distinguishing between deep and shallow agreements (Aichele et al. (2014)), the number of provisions contained in an agreement (Egger and Nigai (2015)), the count of legally enforceable or economically relevant provisions (Mattoo et al. (2017)), or of WTO-plus and/or WTO-extra provisions included in agreements Kohl et al. (2016). Studies such as Kohl (2014) and Baier et al. (2016) account for the heterogeneous impact of trade agreements by including dummies for each agreement. The latter decompose the variance in FTA effects and show that it is driven to approximately 35% by between-agreement, 36% by within-agreement-between-pairs, and 29% by within-pair variation. This observation is particularly relevant with respect to this paper, as one of the reasons why an agreement affects its members differently may be that some countries exert more influence on the treaty text than others. Alschner et al. (2017a) are the first to introduce text-based variables into an empirical study of international trade flows. We construct a variable capturing the textual similarity of different agreements to the TPP to illustrate the potential usefulness of text-as-data methods in assessing trade agreements, based on the assumption that similarly worded agreements might have a similar impact on trade flows.

While economists are primarily interested in the impact of trade agreements on trade flows and its determinants, lawyers and political economists have studied the design of trade agreements and its diffusion. Text-as-data methods² have proven particularly useful for this endeavor, as they allow to compute the textual overlap between agreements. Allee and Elsig (2015) compare the texts of 211 bilateral PTAs and conclude that “most PTAs take the overwhelming majority of their content verbatim from existing agreements” (mostly previous agreements of one of the partners). Alschner

²Two excellent overview articles of text-as-data in the social sciences are Grimmer and Stewart (2013), who discuss applications to political texts, and Gentzkow et al. (2017) who present examples from economics.

et al. (2017b) identify three different large clusters of trade agreements based on textual similarity, and document that large trading partners can follow one (United States) or several (European Union) treaty models. Allee and Lugg (2016b) show that large parts of the text of the Trans-Pacific Partnership (TPP) agreement are taken from previous trade agreements of the United States, and Alschner and Skougarevskiy (2016b) show that 81% of the TPP’s investment chapter correspond to previous investment agreements signed by the United States. However, not all trade agreements are copy-and-pasted; for instance, Allee et al. (2017) find that the Comprehensive Economic and Trade Agreement (CETA) between the European Union and Canada is, to a large extent, original.

One of the first text-as-data approaches explicitly addressing power in treaties is Spirling (2011), who studies U.S. Treaty Making with American Indians. He uses principal component analysis to uncover the underlying dimensions of the treaties and agreements between the U.S. government and American Indians. He then uses a random forests algorithm to predict the main dimension and studies the most important predictors (word stems), concluding that they capture the harshness of the treaties. He argues that harshness has increased over time, hand in hand with the rising economic and military power of the United States vis-à-vis the Native Americans. In Spirling’s application, the treaty partners are of different nature (the U.S. government versus Indian tribes), and the text of the treaty applies asymmetrically. My application differs from that setting in that the treaty partners in international trade agreements are of the same nature (independent states) from a legal point of view, and that the treaty text thus applies symmetrically to both. I therefore do not expect to be able to infer the relative power of treaty members from underlying treaty dimensions. Several authors have suggested that, in the context of international trade and investment agreements, textual similarity measures can reveal the extent to which a treaty text resembles a country’s preferred outcome. This argument is particularly compelling for bilateral investment treaties (BITs), where countries have so called “model BITs” that can be viewed as their preferred negotiation outcome.³ While no “model PTAs” exist, the idea that textual similarity to a country’s other trade agreements captures how much a treaty reflects the interests of this country is generally accepted in the respective literature. For example, the fact that CETA does not reflect the EU’s past agreements more than Canada’s is interpreted by Allee et al. (2017) as “contrast[ing] power-centric accounts of international negotiations”. Therefore, I use measures of textual similarity to reflect power in trade agreements, as will be explained in more technical detail in section 3.2.⁴

3 Data

3.1 Data sources

This paper uses data on bilateral trade flows from UN COMTRADE (HS 1988/92), data on tariffs from the UNCTAD Trade Analysis Information System (TRAINS) database, and data on GDP and GDP per capita from the World Bank Development Indicators database. Data on trade agreement texts is from the Texts of Trade Agreements (ToTA) corpus⁵, which contains texts and metadata of

³See, for example, Allee and Lugg (2016a), Alschner and Skougarevskiy (2016a), Alschner and Skougarevskiy (2016c) and Berge and Stiansen (2016)

⁴Note that I use of the word “power” differently than some of the above-mentioned articles, which discuss the economic or geopolitical power of a country as a determinant of similarity between its agreements, i.e. an “input” to negotiations. In contrast, I use “power” to refer to the “output” of the negotiations, assuming that the text of a trade agreement represents revealed evidence on the succesful exercise of power in the negotiation.

⁵See Alschner et al. (forthcoming); the corpus is available on <https://github.com/mappingtreaties/tota>.

448 trade agreements that have been notified to the WTO by its member states. For the analysis, I use data for all countries for which the necessary variables are available, and for the years between 2004 and 2016. I chose this period to strike a balance between using a large sample and having a relatively large share of observations with a PTA and one or more previous PTAs (which is necessary to compute the power indicators discussed below). The final dataset comprises 261 different trade agreements, and 2,972 changes in trade agreement within pairs (which comprises both concluding a first agreement and replacing one agreement by another).

3.2 Textual similarity

This paper uses different measures of textual similarity to capture the copy-pasting of treaties. Intuitively, the procedure resembles that of a plagiarism detection software that compares the content of two texts. First, I split each trade agreement text into 5-character gram components. For example, the word “free trade” would become “free ”, “ree t”, “ee tr”, “e tra”, “trad” and “trade”. Compared to the bag-of-words approach, which splits texts into the words that compose them, this method has the advantage of preserving word order, which has been argued to be more relevant for the study of legal texts. It has first been proposed for text classification by Lodhi et al. (2002), and applied to the study of legal texts by Spirling (2011) in his study of U.S. Treaty Making with American Indians. I follow Spirling (2011) in choosing a string length of 5, because string lengths between 4 and 7 performed best in the classification exercise using English language texts by Lodhi et al. (2002). As a robustness check, I run the main regressions with similarity measures based on 4-, 6- and 7-gram components. The results are robust to this variation.

Second, to compare two agreements, or more broadly, two sets of text, I compute measures of similarity that are based on assessing the share of elements that appear in both sets. Existing studies of textual similarity in trade and investment agreements have analyzed treaty-pair-level, undirected Jaccard similarities. Treaty-pair level Jaccard similarities are a useful tool to determine the extent of similarity between two treaties. However, as this paper ultimately estimates a gravity model at the exporter-importer-year level, I develop more flexible, asymmetric measures that vary at the country-treaty and exporter-importer-treaty level. In addition, I compare treaties to the entire set of past treaties of the same country, which allows me to capture the extent to which a country is dominating the text of a particular agreement.

All measures are based on the ratio model proposed by Tversky (1977), which generalizes different models of similarity proposed in the literature:

$$s_{trm}(A, B) = \frac{|A \cap B|}{(|A \cap B| + \alpha(A \setminus B) + \beta(B \setminus A))} \quad (1)$$

where $\alpha, \beta \geq 0$. I adapt α and β according to the situation studied in this paper, to capture power.

3.3 Copy-pasting and power

3.3.1 Absolute power

The similarity of trade agreement texts can indicate whether different member countries were rule-makers or rule-takers in trade negotiations. I assume that a country is a rule-taker, i.e. that it successfully exerted power, if a treaty closely resembles this country's previous agreements. To measure the influence of a country on a particular treaty, I compute the share of the text of the new treaty that is identical to text of any previous treaty signed by the same country. I call this measure "absolute power", as it captures the influence of a particular country on a treaty, but not the influence of its trading partners. Formally,

$$power(i, PTA_{i,T}^m) = \frac{|PTA_{i,T}^m \cap PTA_{i,t \leq T}^{n \neq m}|}{|PTA_{i,T}^m|} \quad (2)$$

The absolute power of country i in treaty PTA^m signed in year T is defined as the intersection between the set of unique 5-character gram components of the text of $PTA_{i,T}^m$ and the set of unique 5-character gram components of the texts of all other treaties signed by i in year T and earlier ($PTA_{i,t \leq T}^{n \neq m}$), divided by the number of unique 5-character gram components of $PTA_{i,T}^m$.⁶ In terms of the ratio model, this corresponds to a situation in which $\alpha = 1$ and $\beta = 0$. In other words, the higher the share of an agreement text that is copy-and-pasted from any of the country's previous agreements, the higher is that country's power in this particular agreement.

Figure 1a in the appendix shows the distribution of the mean of this measure across treaties. Figure 1b displays the distribution of absolute power across exporter-treaty combinations (which, by definition, is equivalent to the distribution over importer-treaty combinations). It should be noted that absolute power of exactly 0 implies that the exporter did not have any previous treaty, which is why the new treaty text cannot be similar to previous ones. Figure 1c displays the distribution of absolute power across all exporter-importer-treaty observations. Comparing the figures shows that, while a relatively large number of treaties are very similar to the exporter's previous treaties, those copy-pasted treaties tend to have less members, making the importer-exporter-year distribution less skewed.⁷

Figure 2a illustrates the correlation between the absolute power of the importer and the absolute power of the exporter. It is slightly positive, with a correlation coefficient of 0.05. This is partially driven by pairs in which the joint treaty is the first treaty for both partners, in which case the absolute power is 0 for both partners. Excluding pairs for which the treaty is the first for either partner yields figure 2b. The correlation is now slightly negative, with a correlation coefficient of -0.1. This relatively weak correlation implies that different scenarios coexist: pairs in which both partners have a high influence on the text, pairs in which neither partner has a high influence on the text, and pairs in which one partner has high and one partner has low influence.

⁶I include texts of treaties signed in the same year because the negotiations can have influenced each other.

⁷Note that the distribution of members per treaty is also relatively skewed. While 225 treaties in the final dataset are bilateral, the multilateral treaty with most member countries counts 40 members.

3.3.2 Relative power

The latter finding implies that absolute power at the exporter-treaty level may not be an ideal measure to assess the relationship between exporter and importer in the context of a treaty. A situation in which a new treaty closely resembles the exporter's but not the importer's previous treaties is not the same as a situation in which the new treaty resembles both partners' previous treaties. In other words, power in a treaty is relative. In addition, in a treaty with more than two members, country i may have more influence on the treaty than country j , but less than country k . This bilateral dimension of power within the same treaty should also be captured in the main measures of power used in the empirical analysis.

I therefore construct a measure of the relative power of exporter i over importer j in time T . It is defined as the difference between the share of the text of the (most important) treaty in force between i and j in T that is identical to the exporter's previous treaties and the share identical to the importer's previous treaties.⁸

$$power_{i,j,t} = power(i, PTA_{i,j,T}^m) - power(j, PTA_{i,j,T}^m) = \frac{|PTA_{i,j,T}^m \cap PTA_{i,t \leq T}^{n \neq m}|}{|PTA_{i,j,T}^m|} - \frac{|PTA_{i,j,T}^m \cap PTA_{j,t \leq T}^{n \neq m}|}{|PTA_{i,j,T}^m|} \quad (3)$$

Figure 3a shows the distribution of relative power at the treaty level. More precisely, it depicts the mean absolute power differential per treaty (the mean power differential at the exporter-importer-treaty level, by definition, would be 0). Figure 3b shows the distribution of relative power across exporter-importer-treaty observations. The graphs show that within many treaties and for many country pairs, the power differential is relatively balanced. However, some exporter-importer-country pairs show a considerable power differential.

Figure 4a illustrates absolute power by country, i.e. the mean of how much a country's treaties resemble its own previous treaties.⁹ The countries with the highest absolute power are the United States, Canada, Russia, Japan, and a number of Central American, European and African countries. Australia, New Zealand, and several countries in all continents have intermediate values. The countries with the lowest average values are mostly in South America, the Arab region and Asia.

Figure 4b illustrates relative power by country, i.e. the extent to which each country has more power over its trade agreements as compared to its trading partners.¹⁰ The countries with the

⁸In cases where country pairs have more than one agreement in a given year, I retain only the most pertinent agreement. For example, France and Finland, in 2016, were both members in the EC Treaty and in the EU - Chile treaty (and more than 30 other treaties between the EU and outside partners). I consider the EC Treaty to be the most pertinent one for this country pair.

⁹The mean is aggregated from the country-pair level. This implies that each treaty is weighted by the number of partners it applies to. However, only the most important treaty per country pair is accounted for. Otherwise, the absolute power of certain countries would appear much higher. Consider for example the case of the member states of the European Union: if the numerous EU treaties with outside partners were also considered when computing relative power between two EU member states, this would artificially inflate their absolute power.

¹⁰Again, only the most important treaty for each country pair is accounted for. In this case, for example, the power of countries in the European Union would be underestimated otherwise, as the value would be attenuated by the highly balanced power of member states among themselves in agreements with outside partners.

highest relative power are the United States, Canada, the members of the European Union, Russia, Australia, New Zealand, several South American countries and Sudan. Several Asian countries, including Japan, as well as Mexico, Chile, Mauritania and Turkey have intermediate power. With the exception of Norway and Iceland, the countries with the lowest average values are all in Latin America, the Arab region and Asia. A particularly noteworthy case is Central America. Trade agreements of countries in this region are highly consistent with their own past agreements, as shown in figure 4a. However, their trading partners are scoring even higher on this metric, which causes an overall negative relative power for Central America.

3.4 Descriptive Statistics

Table 8 (appendix) presents the descriptive statistics of the main variables to be used in the empirical analysis, at the importer-exporter-year level. Trade between partners ranges between 0 and 500 billion USD, with a mean of 500 million. 15% of country-pairs have a trade agreement notified to the WTO. The simple average of tariffs across products is 7.1%, while the trade-weighted average is 6.2%. The minimum tariffs are zero. The original data contained extreme values, some of more than 2000% tariffs, which are likely to be caused by reporting errors. Therefore, I excluded the highest 1% of tariff values from the analysis.¹¹ The maximum tariffs in my dataset are 34.52% (simple average) and 43.28% (weighted average). Countries' GDP ranges between 100 million USD (Kiribati in 2004) and 18,730 billion USD (USA in 2016), with a mean of 419 billion USD. GDP per capita ranges between 86 USD (Somalia in 2011) and 119,226 USD (Luxemburg in 2014).

Table 9 (appendix) shows summary statistics for country-pairs with a preferential trade agreement. In 16% of observations, the trade agreement is the first one for one of the partners. Agreement length in words ranges from 758 to 83,024, with a mean of 18,172.¹² The absolute power of the exporter (and, by definition, that of the importer) ranges from 0 to 1, with a mean of 0.71. At the mean, 71% of the agreement text is identical to text from the exporter's previous agreements. The mean of the absolute value of relative power is 0.28, with a minimum of 0 and a maximum of 0.98. For the mean observation, the difference between the share of text identical to each partner's previous agreements is 28%.

4 Empirical Strategy

4.1 Gravity model

To assess the impact of power in trade agreements on trade flows, I estimate a standard gravity model where X_{ijt} designates total goods trade flows from exporter i to importer j in time t . γ_{it} , δ_{jt} and ζ_{ij} capture exporter-year, importer-year and importer-exporter time fixed effects, and ϵ_{ijt} the error term. $policy_{ijt}$ represents trade policy measures in place between i and j at time t . It

¹¹The results on the coefficients of interest are robust to including the outliers.

¹²18,172 is the mean length in words when aggregating observations at from the importer-exporter-year level. At the agreement-level, the mean agreement length is 11,179 words. This is because longer agreements have more members, i.e. more exporter-importer-year observations, than shorter agreements.

comprises different sets of variables, all related to the trade agreements between i and j , and will be explained in detail in the next sub-section.

$$X_{ijt} = \beta * policy_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (4)$$

I estimate this model using the Poisson pseudo maximum likelihood (PPML) estimator proposed by Santos Silva and Tenreyro (2006), which allows to account for zero trade flows. I estimate it using the Stata command `ppml_panel_sg`, as described in Larch et al. (2017), which allows to include three-way fixed effects and cluster standard errors at the importer-exporter level. The baseline regressions use data on total trade in goods between 189 partners in the period from 2004 to 2016.¹³

4.2 Policy variables

First, as a baseline, I only include a dummy capturing whether exporter i and importer j have a trade agreement in English notified to the WTO in time t (PTA_{ijt}). The coefficient β_1 is expected to be positive and significant, and in line with values for PTA dummies found in the existing literature.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (5)$$

Second, I include the measure of relative power described in section 3.3.2. My main variable of interest is relative power, as I focus on situations in which there is a power differential between exporter and importer. Measures of absolute power are studied in a robustness check to support a more detailed interpretation of the main results (see section 5.1).

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{i,j,t} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (6)$$

$power_{i,j,t}$ captures the relative power of exporter i over importer j in the trade agreement in force between them at time t , as defined in section 3.2. If β_2 is positive and significant, this implies that a higher relative influence of one trading partner in a trade agreement increases this partner's exports more than proportionally.

4.3 Control variables

4.3.1 Tariffs

One concern may be that those countries that dominate agreement texts also offer preferential tariffs to developing countries in general. Thus, when they enter into a trade agreement with a developing country partner, it may be the case that the resulting reduction in tariffs is considerably larger for them than for the developing country, because the latter was already benefitting from preferential access before the entry into force of the agreement. Such a mechanism could drive the result that the exports of dominant countries increase more than proportionally after concluding a trade agreement. To address this concern, I include control variables capturing bilateral tariffs in the regression.

¹³A list of countries is included in the appendix.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{i,j,t} + \beta_3 * tariffs_{i,j,t} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (7)$$

As tariffs apply at the exporter-importer-year-product level, while the analysis of this paper is at the exporter-importer-year level, there is no ideal way to control for them. A simple average may overstate un-important goods, while a trade-weighted average is, by definition, endogenous to trade value. Hence, I use both the simple and weighted average (computed from the 6-digit level of the Harmonized System classification) and compare the results.

4.3.2 Differences in GDP

Larger and more developed countries tend to dominate the texts of trade agreements. It may be that the effects captured by the estimation reflect relative economic power, not the actual agreement texts. To verify this, I control for the difference in GDP and GDP per capita as a share of total GDP of exporter and importer.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{i,j,t} + \beta_3 * \frac{GDP_{i,t} - GDP_{j,t}}{GDP_{i,t} + GDP_{j,t}} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (8)$$

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{i,j,t} + \beta_3 * \frac{GDP_{i,t}^{pc} - GDP_{j,t}^{pc}}{GDP_{i,t}^{pc} + GDP_{j,t}^{pc}} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (9)$$

4.3.3 PTA fixed effects

As highlighted, among others, by Baier et al. (2016), the effects of PTAs are heterogeneous. The authors find that most heterogeneity in effects occurs within PTAs and that effects are often asymmetric within exporter-importer pairs, which is in line with the main argument of this paper. However, part of the heterogeneity is agreement-specific. If power imbalances are related to other agreement-specific characteristics, this might affect the results. To control for this possibility, I include agreement-level dummies for all PTAs in my sample.

$$X_{ijt} = \beta_1 * PTA_{ijt} + power_{i,j,t} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + PTA_{ijt}^{FE} + \epsilon_{ijt} \quad (10)$$

4.4 Results

Results from the regressions are reported in table 1. The baseline result (column 1) shows that trade agreements increase trade between member countries. The coefficient on the trade agreement dummy is significant at the 1%-level, and its size is comparable to results found in the literature.¹⁴

The coefficient on relative power (column 2) is positive and significant. This implies that a stronger relative influence of the exporter on the trade agreement text has positive impact on exports. In other words, if a country dominates the text of a trade agreement, its exports to the respective partner country increase more than proportionally. The positive impact of trade

¹⁴For example, Dai et al. (2014) find a coefficient of 0.118 when using the same methodology for a different sample.

agreements on exports is larger for rule-makers than for rule-takers. In economic terms, the coefficient is considerable. In the hypothetical case in which the trade agreement were to 100% copy-pasted from rule-maker i 's and to 0% from rule-taker j 's previous agreements ($power_{i,j,t} = 1$ and $power_{j,i,t} = -1$), the rule-maker's exports would increase by 24.6%, which is about twice as much as in a "perfectly balanced" trade agreement, while the rule-taker's exports would increase only by 1.6%. In the case of an "average" power relation, in which the power differential is 0.28 (see descriptive statistics), the "rule-maker" increases its exports by about 15.8%, while the "rule-taker" only increases its exports by about 9.3%.¹⁵

This result is robust to including all the control variables discussed in section 4.3. The coefficients on both the trade agreement dummy and the power measure remain highly significant and stable when including controls for tariffs (simple average in column 3, trade-weighted average in column 4). The sample is reduced in this case, as tariff data is not available for all countries. I run the baseline regression on the reduced sample to assess the effect of the change in sample (column 5). The size and significance of the coefficient is close to those from columns 3 and 4, which leads me to conclude that controlling for tariffs does not affect the main result. The results are equally robust to controlling for the share of difference in GDP (GDP in column 6, and GDP per capita in column 7), and to including all possible combinations of variables capturing tariffs and GDP (not reported here). When including treaty-fixed effects, the coefficient on power becomes smaller and less statistically significant, but it remains significant from an economic point of view. This result confirms the findings by Baier et al. (2016) that part of the heterogeneity in the impact of trade agreements occurs within agreements.

¹⁵Based on the coefficients from table 1, column 2: $e^{0.118} * e^{0.102*(0.28)} = 1.16$ and $e^{0.118} * e^{0.102*(-0.28)} = 1.09$

Table 1: Relative Power in PTA Texts and Trade Flows

X_{ijt}	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PTA_{ijt}	.117*** (.027)	.118*** (.026)	.114*** (.028)	.108*** (.029)	.119*** (.027)	.117*** (.026)	.117*** (.026)	-
$power_{ijt}$	-	.102** (.046)	.153*** (.055)	.147*** (.055)	.155*** (.055)	.093** (.046)	.099** (.047)	.061* (.034)
Tariffs								
<i>SimpleAverage</i>	-	-	-.003* (.002)	-	-	-	-	-
<i>WeightedAverage</i>	-	-	-	-.009*** (.002)	-	-	-	-
Difference in GDP (share)								
<i>GDP</i>	-	-	-	-	-	.198* (.120)	-	-
<i>GDPpercapita</i>	-	-	-	-	-	-	.224 (.142)	-
Fixed effects								
Ex-yr, im-yr, ex-im	✓	✓	✓	✓	✓	✓	✓	✓
Treaty	×	×	×	×	×	×	×	✓
Observations	358,745	358,745	247,388	247,388	247,388	358,745	358,745	241,200
R^2	0.995	0.995	0.995	0.995	0.995	0.995	0.995	0.995

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

5 Robustness checks

5.1 Absolute power

The variable capturing relative power in a trade agreement is constructed as the difference between the absolute power of the exporter and the importer. To study whether either of the two drives the impact of relative power, I include both absolute power variables separately into the regression:

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{i,j,t}^i + \beta_3 * power_{i,j,t}^j + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (11)$$

The results from this specification are reported in table 2. The coefficient on the trade agreement dummy is larger than in the main specification. The coefficient on the exporter's influence on the agreement text ($power_{i,j,t}^i$) is positive, but not significant. The coefficient on the importer's influence ($power_{i,j,t}^j$) is negative and significant. A stronger influence of the importer on the text of a trade agreement is associated with a smaller increase in imports than an average trade agreement. This finding may be interpreted as that rule-makers use their power to negotiate treaty texts that do

not increase access to their markets for their partners as much as they increase access to their partners' markets for themselves. When controlling for tariffs, the power of the exporter, which has a positive coefficient, becomes significant at the 10% level (column 3 and 4). As shown in column 5, this is driven by the sample, which may stem from a positive correlation between availability of tariff data and negotiation capabilities. The main results are consistent across columns 6 and 7 (controlling for differences in GDP and GDP per capita), but not column (7), in which part of the variation in the policy variables is absorbed by the treaty fixed effects.

Table 2: Absolute Power in PTA Texts and Trade Flows

X_{ijt}	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PTA_{ijt}	.117*** (.027)	.160*** (.044)	.199*** (.049)	.195*** (.051)	.202*** (.049)	.159*** (.045)	.160*** (.044)	-
$power^i_{ijt}$	-	.077 (.051)	.106* (.057)	.099* (.057)	.110* (.057)	.069 (.052)	.075 (.052)	.143 (.091)
$power^j_{ijt}$	-	-.128** (.048)	-.211*** (.059)	-.205*** (.060)	-.211*** (.059)	-.119** (.048)	-.127*** (.048)	.028 (.091)
Tariffs								
<i>SimpleAverage</i>	-	-	-.003* (.002)	-	-	-	-	-
<i>WeightedAverage</i>	-	-	-	-.009*** (.001)	-	-	-	-
Difference in GDP (share)								
<i>GDP</i>	-	-	-	-	-	.198* (.120)	-	-
<i>GDPpercapita</i>	-	-	-	-	-	-	.224 (.142)	-
Fixed effects								
Ex-yr, im-yr, ex-im	✓	✓	✓	✓	✓	✓	✓	✓
Treaty	×	×	×	×	×	×	×	✓
Observations	358,745	358,745	251,500	251,500	251,500	358,745	358,745	241,200
R^2	0.995	0.995	0.995	0.995	0.995	0.995	0.995	0.995

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

5.2 First agreements

The measure of relative power based on previous trade agreement texts underestimates the influence of countries who sign a trade agreement for the first time. By construction of the indicator, the absolute power of a country that signs its first agreement is equal to zero, as it does not have previous agreements from which it could have copied. However, it is unlikely that this country did not exert any power in the negotiation. This situation concerns 16% of importer-exporter-year pairs with a trade agreement in place. Underestimating the power of one trading partner is likely to bias the coefficients of the main variables of interest towards 0. To verify the expected bias I

distinguish between first and non-first agreements and include separate policy variables for relative power in either of those:

$$X_{ijt} = \beta_1 * PTA_{ijt}^{first} + \beta_2 * PTA_{ijt}^{nonfirst} + \beta_3 * power_{i,j,t}^{first} + \beta_4 * power_{i,j,t}^{nonfirst} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (12)$$

$power_{i,j,t}^{first}$ measures the relative power for agreements in situations where it was not the first agreement for either i or j . It takes the value 0 for agreements that were the first agreement for either partner. $power_{i,j,t}^{nonfirst}$ measures the relative power for agreements that were the first for either of the partners, and takes the value 0 for agreements that were not the first for either partner.

The results are displayed in table 3. The coefficient on relative power in first agreements is smaller than the overall relative power coefficient in the main specification, and it is not statistically significant. The coefficient on relative power in non-first agreements is larger than the overall relative power coefficient, and statistically significant at the 10% level. This decrease in significance may be due to the reduced number of observations for which this variable is non-zero, as compared to the main specification (table 1, column 2) in which all agreements are considered jointly. The coefficient sizes are in line with the assumption that the original measure underestimates the influence of countries who sign their first agreement. The coefficient on correctly measured (non-first) agreements is larger, as it was biased towards zero before. The coefficient on imperfectly measured (first) agreements is smaller, as the "true" power is smaller than what is measured. Overall, the true effect of power on exports may thus be slightly larger than the one obtained in the main specification (table 1).

Table 3: First versus non-first agreements

X_{ijt}	(1)	(2)
PTA_{ijt}^{first}	.134*** (.043)	.138*** (.040)
$PTA_{ijt}^{nonfirst}$.115*** (.026)	.118*** (.026)
$power_{ijt}^{first}$	-	.064 (.045)
$power_{ijt}^{nonfirst}$	-	.157* .092
Fixed effects		
Ex-yr, im-yr, ex-im	✓	✓
Treaty	×	×
Observations	358,745	358,745
R^2	0.995	0.995

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

5.3 Lags and leads

As with any policy, there is a concern that trade agreements and their texts may be endogenous to the outcome. First, countries may decide to negotiate a trade agreement with a specific partner if bilateral trade is expected to increase. Second, power in a trade agreement, as measured by the influence of each partner on its text, may also be endogenous to future exports. If a country expects a trade agreement with another partner to be particularly beneficial, it may invest more resources in negotiating the text and thereby dominate it. In this case, the coefficient of interest would be biased upwards. Conversely, a country that is expecting to substantially increase its exports (for example, a small exporter gaining access to a large market) may be willing to accept a text dominated by its partner, which would have the opposite effect and yield a downward-biased coefficient. In either case the result of the estimation could be affected by this type of expectations and motivations.

To test whether this is the case, I follow Mattoo et al. (2017) and include lags and leads of the trade policy variables in my regression. Results are reported in table 4. Columns 1-3 report the results for the baseline (trade agreement dummy) and columns 4-6 those for the relative power specification, with different combinations of up to four lags and leads of the respective policy variables.

Table 4: Lags and Leads

X_{ijt}	(1)	(2)	(3)	(4)	(5)	(6)
PTA_{ijt_lead4}	-	-	.064** (.026)	-	-	.065** (.025)
PTA_{ijt_lead3}	-	.054** (.025)	.006 (.013)	-	.055** (.025)	.006 (.013)
PTA_{ijt_lead2}	.074*** (.028)	.039* (.022)	.004 (.027)	.075*** (.028)	.039* (.022)	.004 (.026)
PTA_{ijt_lead1}	.016 (.023)	.014 (.034)	-.000 (.024)	.033 (.029)	.016 (.034)	.002 (.024)
PTA_{ijt}	-.001 (.015)	-.005 (.023)	.023 (.015)	-.000 (.015)	-.005 (.022)	.021 (.014)
PTA_{ijt_lag1}	.061** (.029)	.077*** (.029)	.046*** (.017)	.061*** (.028)	.075*** (.028)	.045*** (.017)
PTA_{ijt_lag2}	-.002 (.028)	-.003 (.023)	.014 (.014)	-.002 (.027)	-.001 (.022)	.016 (.013)
PTA_{ijt_lag3}	-	-.007 (.025)	.003 (.014)	-	-.006 (.024)	.002 (.014)
PTA_{ijt_lag4}	-	-	-.025 (.023)	-	-	-.023 (.021)
$power_{ijt_lead4}$	-	-	-	-	-	.102 (.090)
$power_{ijt_lead3}$	-	-	-	-	.071 (.080)	-.013 (.029)
$power_{ijt_lead2}$	-	-	-	.036 (.069)	-.012 (.036)	-.056 (.071)
$power_{ijt_lead1}$	-	-	-	.041 (.036)	.052 (.065)	.069 (.051)
$power_{ijt}$	-	-	-	.029 (.048)	.043 (.035)	.035 (.041)
$power_{ijt_lag1}$	-	-	-	.074** (.037)	.082** (.038)	.056** (.028)
$power_{ijt_lag2}$	-	-	-	-.027 (.040)	-.020 (.017)	-.012 (.014)
$power_{ijt_lag3}$	-	-	-	-	-.015 (.033)	.019 (.015)
$power_{ijt_lag4}$	-	-	-	-	-	-.034 (.026)
Fixed effects						
Ex-yr, im-yr, ex-im	✓	✓	✓	✓	✓	✓
Treaty	×	×	×	×	×	×
Observations	264,552	231,656	201,593	264,552	231,656	201,593
R^2	0.995	0.995	0.996	0.995	0.995	0.996

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

The coefficients on the trade agreement dummy and its leads and lags are very similar in both specifications. The earliest lead (here: four years before the agreement enters into force) is always positive and significant. This can be interpreted as a sign that countries tend to conclude a trade agreement when they expect trade to increase in the future. While the contemporaneous PTA dummy is insignificant, the first lag is always positive and significant, implying that the main impact of the agreement on trade takes effect after one year. The following lags have small and insignificant coefficients, meaning that the effect remains stable over time.

For the relative power variable, the coefficients on leads are all insignificant. This may be interpreted as cautious evidence that, while the conclusion of the agreement may be endogenous to trade flows, power in the agreement is not. For the lags, it is again the first lag that is significant across specifications. The fact that it is the same lag as for the dummy that is significant suggests that, together, these two variables capture the actual effect of the trade agreement, which occurs in the year after its conclusion and remains stable thereafter.

5.4 Different n-grams

As detailed in section 3.2, the textual similarity indicators for power are constructed by splitting each treaty into 5-character gram components. To test whether the length of components affects the results of my analysis, I construct equivalent measures based on 4-, 6- and 7-grams, and estimate the main regression specification using those. The results, reported in table 5, are close to those obtained using variables based on 5-gram components.

Table 5: Using different n-grams

X_{ijt}	(1) 4-grams	(2) 5-grams	(3) 6-grams	(4) 7-grams
PTA_{ijt}	.117*** (.026)	.118*** (.026)	.118*** (.026)	.119*** (.026)
$power_{ijt}$.088** (.039)	.102** (.046)	.105** (.051)	.103* (.055)
Fixed effects				
Ex-yr, im-yr, ex-im	✓	✓	✓	✓
Treaty	×	×	×	×
Observations	358,745	358,745	358,745	358,745
R^2	0.995	0.995	0.995	0.995

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

5.5 Excluding EU member countries

For a variety of reasons, one might suspect that the impact of trade agreements is different for the members of the European Union than for other countries. European market integration is much deeper than any trade agreement between other nations, which is why the hypotheses on trade agreements may not apply to internal EU trade. At the same time, the EU has a large number of

trade agreements with outside partners. While not all EU member states may have influenced those to the same extent, the power measures are by construction the same for all of them, as they adhere to the same commercial policy. This mechanism might affect the validity of results with respect to the EU members' external trade relations. To test whether the specificities of the EU affect the main results of my analysis, I estimate the main specification on a sample that excludes all EU member states. The results are reported in table 6. The coefficient on the trade agreement dummy is very close to the result obtained in the main specification (table 1, columns 1 and 2). The coefficient on power is slightly larger than in the main specification. This may be the case because the power measure overstates the power of less influential EU countries, causing a downward bias in the main specification.

Table 6: Excluding EU countries

X_{ijt}	(1)	(2)
PTA_{ijt}	.118***	.122***
	(.029)	(.029)
$power_{ijt}$	-	.152**
		(.059)
Fixed effects		
Ex-yr, im-yr, ex-im	✓	✓
Treaty	×	×
Observations	236,975	236,975
R^2	0.996	0.996

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

5.6 Agreement length

The text-based measures of power may be affected by the varying length of trade agreements, which has increased from a mean of about 5,000 words in the 1950s to a mean of about 50,000 words in the 2010s (Alschner et al. (2017b)). As a source, longer agreement texts provide more text to copy-paste from, and as a result, they provide more space for innovation. Therefore, the effect of power may differ according to agreement length. To test this, I distinguish between agreements that are shorter and those that are longer than the median agreement (11,179 words), and include separate policy variables for relative power for either group, similar to the distinction between first and non-first agreements in section 5.1.

$$X_{ijt} = \beta_1 * PTA_{ijt}^{short} + \beta_2 * PTA_{ijt}^{long} + \beta_3 * power_{i,j,t}^{short} + \beta_4 * power_{i,j,t}^{long} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \epsilon_{ijt} \quad (13)$$

$power_{i,j,t}^{short}$ measures the relative power for agreements that have up to the same number of words than the median agreement. It takes the value 0 for longer agreements. $power_{i,j,t}^{long}$ measures the relative power for agreements that consist of more words than the median agreement, and takes the value 0 for shorter agreements.

Results are reported in table 7. The coefficient on relative power in short agreements is smaller than the coefficient in the main specification, and it is not statistically significant. The coefficient on

relative power in long agreements is larger than the general coefficient, and statistically significant at the 10% level. This suggests that the core result is driven by longer agreements.

Table 7: First versus non-first agreements

X_{ijt}	(1)	(2)
PTA_{ijt}^{short}	.150*** (.033)	.151*** (.032)
PTA_{ijt}^{long}	.109*** (.028)	.113*** (.027)
$power_{ijt}^{short}$	-	.046 (.053)
$power_{ijt}^{long}$	-	.157* 081
Fixed effects		
Ex-yr, im-yr, ex-im	✓	✓
Treaty	×	×
Observations	358,745	358,745
R^2	0.995	0.995

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

6 Conclusion

This paper has addressed the phenomenon of copy-pasting of trade agreements texts through text-as-data analysis variables capturing the textual similarity between agreements. I constructed text-based variables of absolute and relative power in trade agreements and introduced them into a gravity model to study their impact on international trade flows. The descriptive statistics have shown that some countries, most of them large developed countries, have had a much larger influence on trade agreement texts than their respective trading partners. The results from the regression analysis provide evidence that it pays off to be a “rule-maker”. Countries that dominate trade agreement texts benefit from a larger increase in exports than their trading partners when the agreement enters into force. In an “average” power relation, the rule-maker increases its exports by about 15.8%, while the rule-taker only increases its exports by about 9.3%. In very unbalanced situations, which are rare but do occur, the rule-maker increases its exports by almost 25%, whereas rule-takers experience almost no increase.

The result is robust to including pertinent control variables, such as tariffs, differences in GDP and GDP per capita, and treaty fixed effects. Coefficients are larger for agreements which are not the first agreement ever concluded by either of the partners, in line with what one would expect due to the construction of the text-based power variable. They are also larger for longer agreements. Importantly, regression results based on separating the treaty-specific power of the exporter and the importer show that the asymmetric effect of PTAs is driven by the importer’s influence on the agreement text. This suggests that powerful countries are effectively restricting access to their markets for their partners. Specifications that include lags and leads of the main policy variables show that, unlike the entry into force of a trade agreement, the power balance within the agreement is not correlated with future trade flows, which alleviates concerns of potential endogeneity.

While this paper has shown that the copy-and-pasting of trade agreements yields relatively larger exports for powerful countries, this does not necessarily imply that it always harms their less powerful partners. Increasing one’s gross exports is not the only way to benefit from international trade, and further research is needed to determine whether the effects of copy-pasting are similar for other economic outcomes likely to be affected by trade agreements, such as the terms of trade, global value chain integration, trade in services, and investment. However, the present analysis draws attention to the fact that recycling previous trade agreement texts can have real economic consequences. This is of particular interest for small developing countries that have limited resources for international trade negotiations and need to decide how to allocate those optimally. Governments should devote resources to ensure that they negotiate not only optimal tariffs, but also optimal treaty texts; and international organizations and other partners should provide support and assistance to the most resource-constrained countries to ensure they can harness the opportunities presented by trade agreements with larger and more developed nations through effective negotiations.

References

- AICHELE, R., G. J. FELBERMAYR, AND I. HEILAND (2014): “Going Deep: The Trade and Welfare Effects of TTIP,” *CEPR Working Paper*, 5150.
- ALLEE, T. AND M. ELSIG (2015): “Are the Contents of International Treaties Copied-and-Pasted? Evidence from Preferential Trade Agreements,” Berlin, Germany, 43.
- ALLEE, T., M. ELSIG, AND A. LUGG (2017): “Is the European Union Trade Deal with Canada New or Recycled? A Text-as-data Approach,” *Global Policy*, 8, 246–252.
- ALLEE, T. AND A. LUGG (2016a): “Do BITs Reflect the Interests of Powerful States?” Atlanta, Georgia.
- (2016b): “Who wrote the rules for the Trans-Pacific Partnership?” *Research & Politics*, 3, 2053168016658919.
- ALSCHNER, W., J. SEIERMANN, AND D. SKOUGAREVSKIY (2017a): “The Impact of the TPP on Trade Between Member Countries: A Text-As-Data Approach,” *Asian Development Bank Institute (ADB) Working Paper*, 745.
- (2017b): “Text-as-Data Analysis of Preferential Trade Agreements: Mapping the PTA Landscape,” *UNCTAD Research Paper*, 5.
- (forthcoming): “Text of Trade Agreements (ToTA) – A structured corpus for the text-as-data analysis of preferential trade agreements,” *Journal of Empirical Legal Studies*.
- ALSCHNER, W. AND D. SKOUGAREVSKIY (2016a): “Mapping the Universe of International Investment Agreements,” *Journal of International Economic Law*, 19, 561–588.
- (2016b): “The New Gold Standard? Empirically Situating the TPP in the Investment Treaty Universe,” *The Journal of World Investment & Trade*, 17, 330–373.
- (2016c): “Rule-takers or Rule-makers? A New Look at African Bilateral Investment Treaty Practice,” *Transnational Dispute Management (TDM)*, 13.
- BAIER, S., J. BERGSTRAND, AND M. CLANCE (2015): “Preliminary examination of heterogeneous effects on international trade of economic integration agreements,” in *Trade Cooperation. The Purpose, Design and Effects of Preferential Trade Agreements*, ed. by A. Dür and M. Elsig, 355–373.
- BAIER, S., Y. YOTOV, AND T. ZYLKIN (2016): “On the widely differing effects of free trade agreements: Lessons from twenty years of trade integration,” *Drexel University School of Economics Working Paper Series*, 2016-15.
- BERGE, T. AND Y. STIANSEN (2016): “Negotiating BITs with Models. The Power of Expertise,” *PluriCourts Research Paper*, 16-13.
- DAI, M., Y. V. YOTOV, AND T. ZYLKIN (2014): “On the trade-diversion effects of free trade agreements,” *Economics Letters*, 122, 321–325.

- EGGER, P. AND S. NIGAI (2015): “Effects of deep versus shallow trade agreements in general equilibrium,” in *Trade Cooperation. The Purpose, Design and Effects of Preferential Trade Agreements*, ed. by A. Dür and M. Elsig, 374–391.
- GENTZKOW, M., B. T. KELLY, AND M. TADDY (2017): “Text as Data,” Working Paper 23276, National Bureau of Economic Research.
- GOLDSMITH, J. L. AND E. A. POSNER (2005): *The Limits of International Law*, New York: Oxford University Press.
- GRIMMER, J. AND B. M. STEWART (2013): “Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts,” *Political Analysis*, 21, 267–297.
- KOHL, T. (2014): “Do we really know that trade agreements increase trade?” *Review of World Economics*, 150, 443–469.
- KOHL, T., S. BRAKMAN, AND H. GARRETSEN (2016): “Do Trade Agreements Stimulate International Trade Differently? Evidence from 296 Trade Agreements,” *The World Economy*, 39, 97–131.
- LARCH, M., J. WANNER, Y. V. YOTOV, AND T. ZYLKIN (2017): “The Currency Union Effect: A PPML Re-assessment with High-Dimensional Fixed Effects,” *CESifo Working Paper*, 6464.
- LODHI, H., C. SAUNDERS, J. SHAWE-TAYLOR, N. CRISTIANINI, AND C. WATKINS (2002): “Text Classification Using String Kernels,” *Journal of Machine Learning Research*, 2, 419–444.
- MATTOO, A., A. MULABDIC, AND M. RUTA (2017): “Trade creation and trade diversion in deep agreements,” *World Bank Policy Research Working Paper*, 8206, 1–44.
- RODRIK, D. (2018): “What Do Trade Agreements Really Do?” *Journal of Economic Perspectives*, 32, 73–90.
- SANTOS SILVA, J. M. C. AND S. TENREYRO (2006): “The Log of Gravity,” *The Review of Economics and Statistics*, 88, 641–658.
- SPIRLING, A. (2011): “U.S. Treaty Making with American Indians: Institutional Change and Relative Power, 1784–1911,” *American Journal of Political Science*, 56, 84–97.
- TVERSKY, A. (1977): “Features of Similarity,” *Psychological Review*, 84, 327–352.

A Appendix

A.1 List of countries

Afghanistan, Albania, Algeria, Angola, Antigua and Barbuda, Argentina, Armenia, Australia, Austria, Azerbaijan, Bahamas, The, Bahrain, Bangladesh, Barbados, Belarus, Belgium, Belize, Benin, Bhutan, Bolivia, Bosnia and Herzegovina, Botswana, Brazil, Brunei, Bulgaria, Burkina Faso, Burundi, Cambodia, Cameroon, Canada, Cape Verde, Central African Republic, Chad, Chile, China, Colombia, Comoros, Congo, Rep., Costa Rica, Cote d'Ivoire, Croatia, Cuba, Cyprus, Czech Republic, Denmark, Djibouti, Dominica, Dominican Republic, Ecuador, Egypt, Arab Rep., El Salvador, Equatorial Guinea, Eritrea, Estonia, Ethiopia, Fiji, Finland, France, Gabon, Gambia, The, Georgia, Germany, Ghana, Greece, Grenada, Guatemala, Guinea, Guinea Bissau, Guyana, Haiti, Honduras, Hong Kong (China), Hungary, Iceland, India, Indonesia, Iran, Islamic Rep., Iraq, Ireland, Israel, Italy, Jamaica, Japan, Jordan, Kazakhstan, Kenya, Kiribati, Korea, Rep., Kuwait, Kyrgyz Republic, Lao PDR, Latvia, Lebanon, Lesotho, Liberia, Libya, Lithuania, Luxembourg, Macao, Macedonia, FYR, Madagascar, Malawi, Malaysia, Maldives, Mali, Malta, Marshall Islands, Mauritania, Mauritius, Mexico, Micronesia, Fed. Sts., Moldova, Mongolia, Morocco, Mozambique, Myanmar, Namibia, Nepal, Netherlands, New Zealand, Nicaragua, Niger, Nigeria, Norway, Oman, Pakistan, Panama, Papua New Guinea, Paraguay, Peru, Philippines, Poland, Portugal, Qatar, Romania, Russian Federation, Rwanda, Samoa, San Marino, Saudi Arabia, Senegal, Seychelles, Sierra Leone, Singapore, Slovak Republic, Slovenia, Solomon Islands, Somalia, South Africa, Spain, Sri Lanka, St. Kitts and Nevis, St. Lucia, St. Vincent and the Grenadines, Sudan, Suriname, Swaziland, Sweden, Switzerland, Syrian Arab Republic, Tajikistan, Tanzania, Thailand, Togo, Tonga, Trinidad and Tobago, Tunisia, Turkey, Turkmenistan, Uganda, Ukraine, United Arab Emirates, United Kingdom, United States, Uruguay, Uzbekistan, Vanuatu, Venezuela, Vietnam, Yemen, Rep., Zambia, Zimbabwe

A.2 Descriptive Statistics

Table 8: Summary statistics - by observation (exporter-importer-year)

	mean	sd	min	max
Trade in million USD	503	5,272	0	504,028
PTA	0.15	0.36	0	1
Tariff: Simple average	7.07	6.87	0	34.52
Tariff: Trade-weighted average	6.24	7.37	0	43.28
GDP (exporter), billion USD	419	1,517	0.10	18,730
GDP per capita (exporter), USD	14,982	20,313	86	119,226
Observations	368,740			

Table 9: Summary statistics - by observation, only with PTA

	mean	sd	min	max
First PTA for either partner	0.16	0.37	0	1
PTA length in words	18,172	18,146	758	83,024
Absolute power (exporter)	0.71	0.29	0	1
Absolute value of relative power	0.28	0.27	0	0.98
Observations	55,172			

A.3 Figures

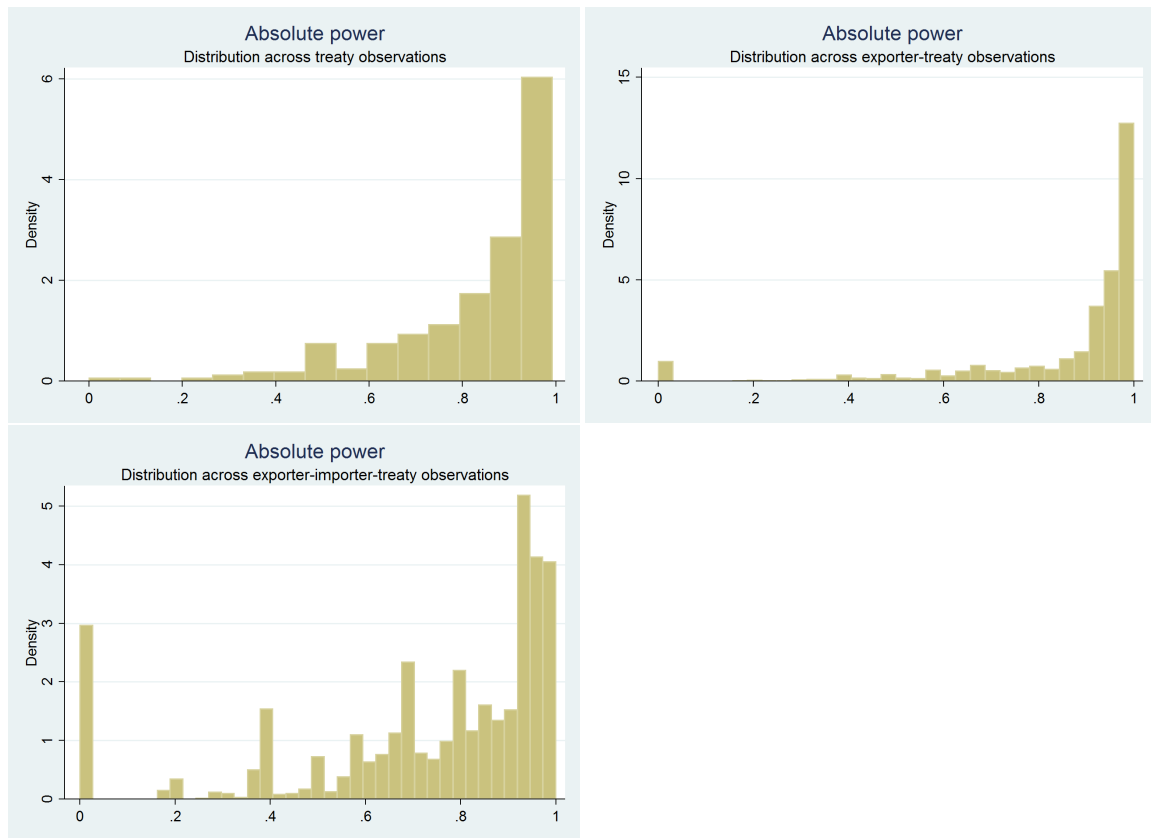


Figure 1: Absolute power

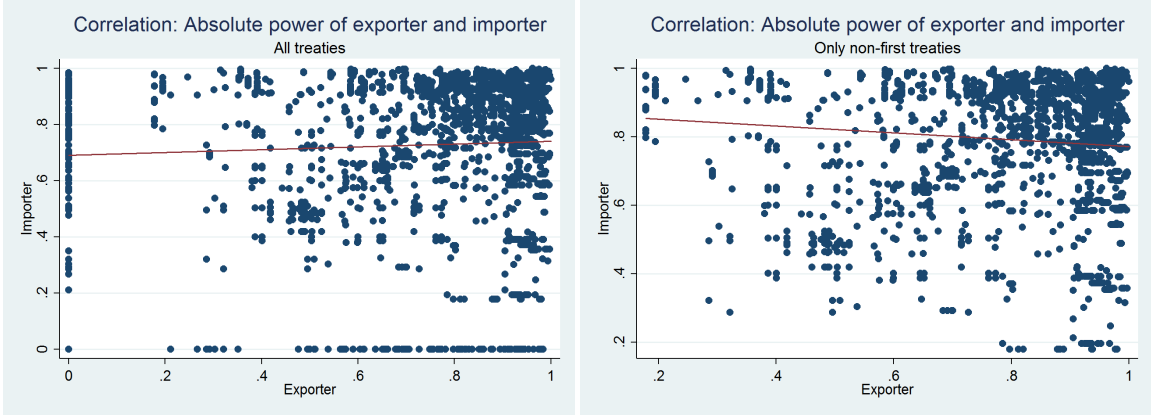


Figure 2: Correlation between absolute power of exporter and importer

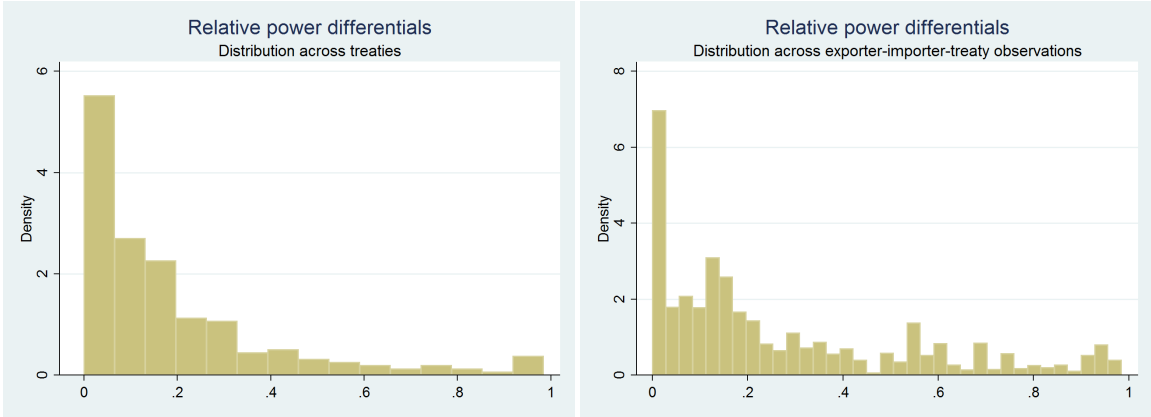


Figure 3: Relative power

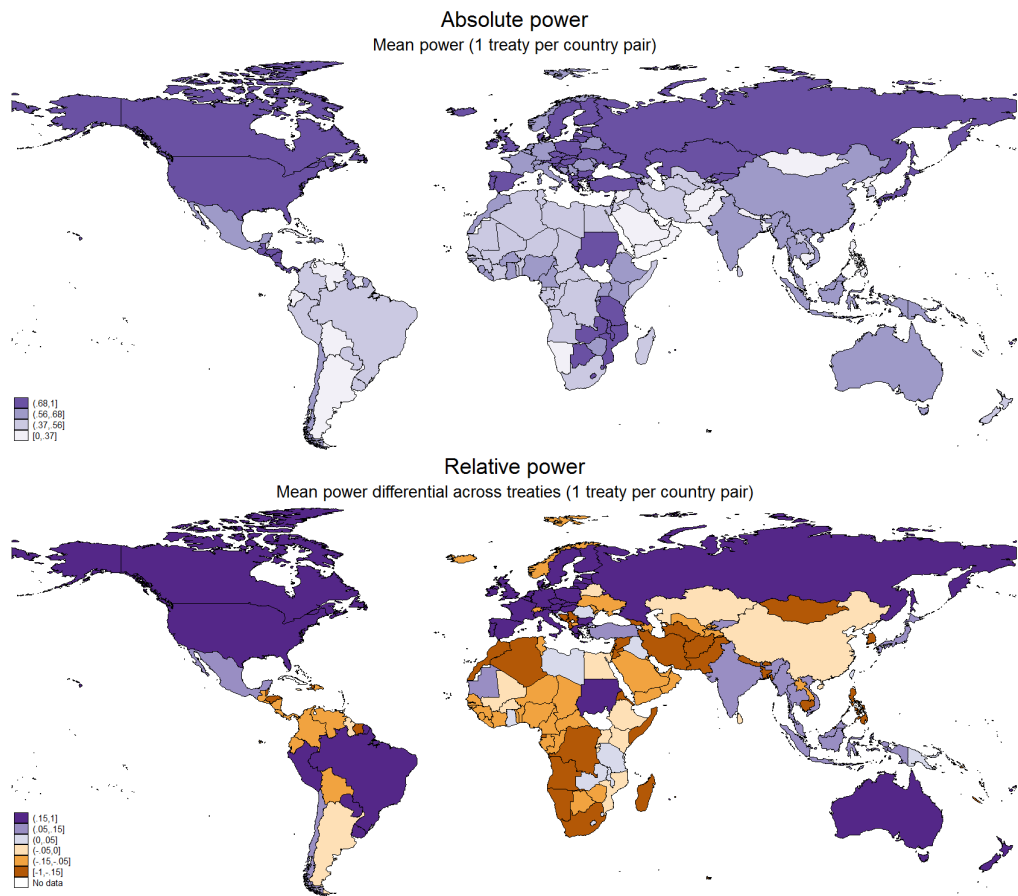


Figure 4: Power by country

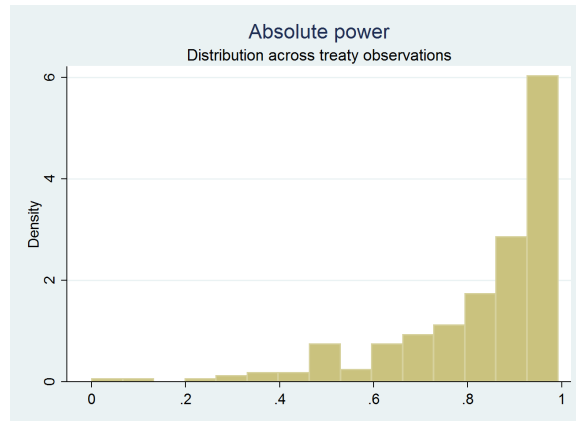


Figure 5: Absolute power 1




Figure 6: Absolute power 2

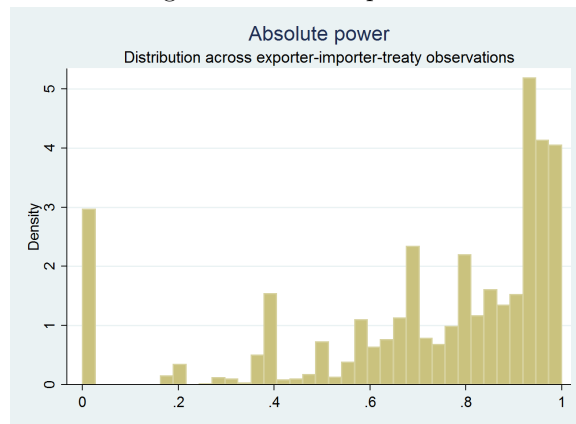


Figure 7: Absolute power 3

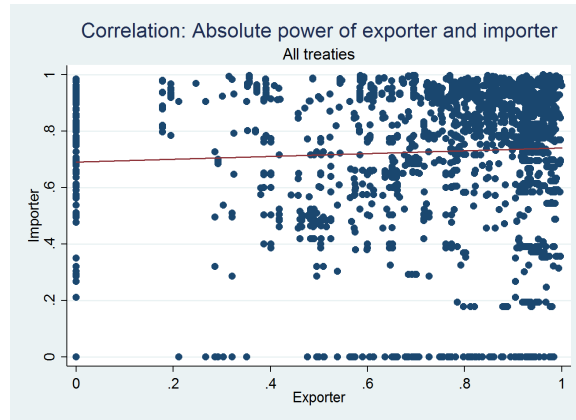


Figure 8: Correlation 1

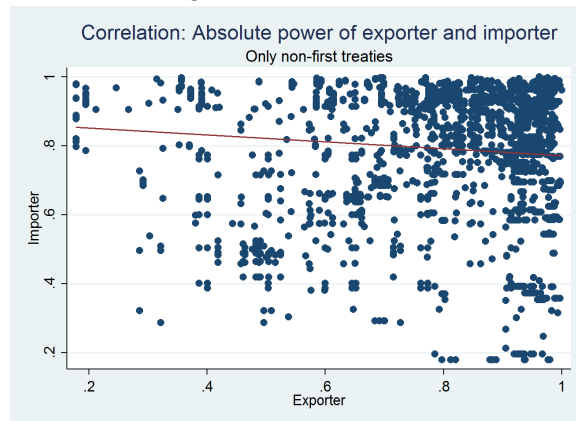


Figure 9: Correlation 2




Figure 10: Relative power 1




Figure 11: Relative power 2