

October 2018

UNCTAD Research Paper No. 26
UNCTAD/SER.RP/2018/8/Rev.1

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

Julia
Seiermann

Division on
International Trade
and Commodities,
UNCTAD

Julia.Seiermann@unctad.org

Abstract

The texts of international trade agreements are often recycled 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 partners more than their own. To study whether rule-makers benefit more from trade agreements than rule-takers, the paper uses a new text corpus of machine-readable trade agreement texts, Texts of Trade Agreements (ToTA). It builds indicators of textual similarity that capture how much an agreement resembles each party's previous agreements and introduce them into a gravity model of international trade. The results show that the exports of countries that had a greater influence on the treaty text increase more than their partners'. In an "average" situation, the rule-maker increases its exports by about 16.3%, while the rule-taker only increases its exports by about 8.8%. The effect is driven by the influence of the importer, suggesting that powerful countries effectively restrict access to their markets. The results demonstrate that the recycling of trade agreements has real economic effects, suggesting that countries should carefully assess text proposals when negotiating a trade agreement.

Key words: Trade, Text-as-data, International Trade Agreements



UNITED NATIONS

The findings, interpretations, and conclusions expressed herein are those of the author(s) and do not necessarily reflect the views of the United Nations or its official Member States. The designations employed and the presentation of material on any map in this work do not imply the expression of any opinion whatsoever on the part of the United Nations concerning the legal status of any country, territory, city, or area or of its authorities, or concerning the delimitation of its frontiers and boundaries.

This paper represents the personal views of the author(s) only, not the views of the UNCTAD secretariat or member States. The author(s) accept sole responsibility for any errors. Any citation should refer to the author(s) and not the publisher. This paper has not been formally edited.

Contents

Acknowledgements	2
1. Introduction	3
2. Literature Review	4
3. Data	5
4. Empirical Strategy	10
5. Robustness checks	13
6. Conclusion	19
References	20
Appendix	22

Acknowledgements

This note was prepared by Julia Seiermann (UNCTAD). The author thanks colleagues from UNCTAD's Trade Analysis Branch and seminar participants at The Graduate Institute for helpful comments and suggestions, and Wolfgang Alschner of University of Ottawa and Dmitriy Skougarevskiy of European University at St. Petersburg for collaborating to build the Texts of Trade Agreements (ToTA) corpus used in this paper.

1. Introduction

The texts of preferential trade agreements (PTAs) are often recycled 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, Seiermann and Skougarevskiy, 2017b). Some large countries, such as the United States of America or the European Union, 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 recycling 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.

This paper uses 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, it computes the share of the text that is identical to text from earlier treaties concluded by the same country. The main variable of interest is relative power, which is defined as the difference between the exporter’s and the importer’s influence on the text of the treaty in force between them.¹ This measure is introduced into a gravity model of international trade to assess the impact of power on exports. Its 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 a 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, the findings of this paper 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.

In sum, the recycling 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 fewer resources and less bargaining power, they are 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.

¹ Using 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.

2. Literature Review

By assessing the impact of recycling treaty texts 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, this literature review focuses 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, Bergstrand and Clance (2015), who distinguish between free trade agreement, customs union, common market and economic union), binary variables distinguishing between deep and shallow agreements (Aichele, Felbermayr and Heiland (2014)), the number of provisions contained in an agreement (Egger and Nigai (2015)), the count of legally enforceable or economically relevant provisions (Mattoo, Mulabdic and Ruta (2017)), or of WTO-plus and/or WTO-extra provisions included in agreements (Kohl et al. (2016)). Studies such as Kohl (2014) and Baier, Yotov and Zylkin (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, Seiermann and Skougarevskiy (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 endeavour, 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, Seiermann and Skougarevskiy (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 re-use text from previous agreements; for instance, Allee, Elsig and Lugg (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 United States Treaty Making with American Indians. It uses principal component analysis to uncover the underlying dimensions of the treaties and agreements between the United States government and American Indians. A random forests algorithm is applied to predict the main dimension and study the most important predictors (word stems), concluding that they capture the harshness of the treaties. The paper 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 United States government versus Native American tribes), and the text of the treaty applies asymmetrically.

² Two comprehensive 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.

This 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. It is therefore not expected that it is possible 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 European Union's past agreements more than Canada's is interpreted by Allee, Elsig and Lugg (2017) as "contrast[ing] power-centric accounts of international negotiations". Therefore, this paper uses 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

Data sources

This paper uses data on bilateral trade flows from United Nations 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) corpus,⁵ which contains texts and metadata of 448 trade agreements that have been notified to the WTO by its member states. The data used in the analysis includes all countries for which the necessary variables are available, and the years between 2004 and 2016. This period is chosen 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).

Textual similarity

This paper uses different measures of textual similarity to capture the recycling of treaties. Intuitively, the procedure resembles that of a plagiarism detection software that compares the content of two texts. First, each trade agreement text is split 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 a study of United States Treaty Making with Native Americans. This paper follows 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, the main regressions are run with similarity measures based on 4-, 6- and 7-gram components. The results are robust to this variation.

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

⁴ Note 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 successful exercise of power in the negotiation.

⁵ See Alschner, Seiermann and Skougarevskiy (2018); the corpus is available on <https://github.com/mappingtreaties/tota>.

Second, to compare two agreements, or more broadly, two sets of text, this paper computes 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 analysed 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, it develops more flexible, asymmetric measures that vary at the country-treaty and exporter-importer-treaty level. In addition, it compares treaties to the entire set of past treaties of the same country, which allows 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:

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

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

Text recycling and power

Absolute power

The similarity of trade agreement texts can indicate whether different member countries were rule-makers or rule-takers in trade negotiations. It is assumed 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, this paper computes the share of the text of the new treaty that is identical to text of any previous treaty signed by the same country. This measure is called "absolute power", as it captures the influence of a particular country on a treaty, but not the influence of its trading partners. Formally,

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

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 $PTA_i^{m,T}$ and the set of unique 5-character gram components of the texts of all other treaties signed by i in year T and earlier ($PTA_i^{n \neq m, t \leq T}$), divided by the number of unique 5-character gram components of $PTA_i^{m,T}$.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 identical to text from any of the country's previous agreements, the higher is that country's power in this particular agreement. For countries that join trade agreements after their initial entry into force, it is assumed that they did not influence the original text, i.e. $power(i, PTA_i^{m,T}) = 0$.

Figure 1a 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 can imply two situations: Either the exporter did not have any previous treaty, in which case the new treaty text cannot be similar to previous ones, or the exporter joined an existing treaty after its entry into force, in which case it could not influence the text. 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

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

similar to the exporter's previous treaties, those recycled treaties tend to have fewer members, making the importer-exporter-year distribution less skewed.⁷

Figure 1. Absolute power

Figure 1a

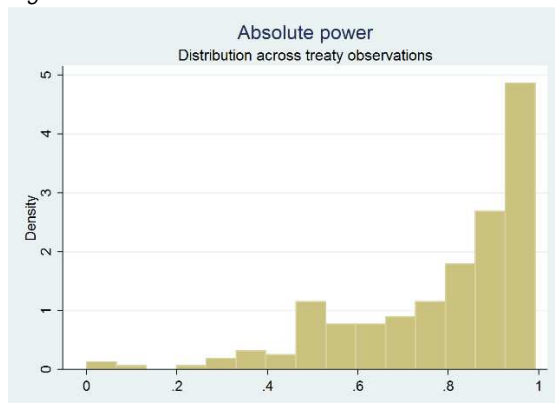


Figure 1b

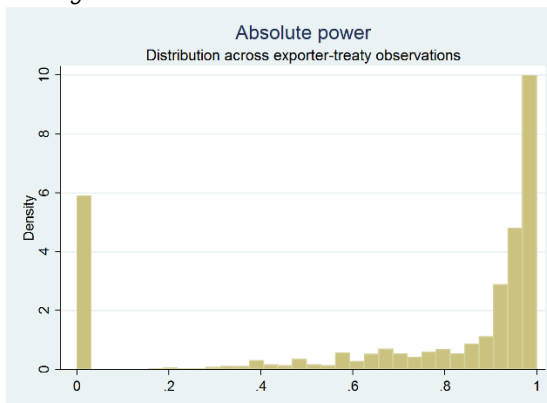
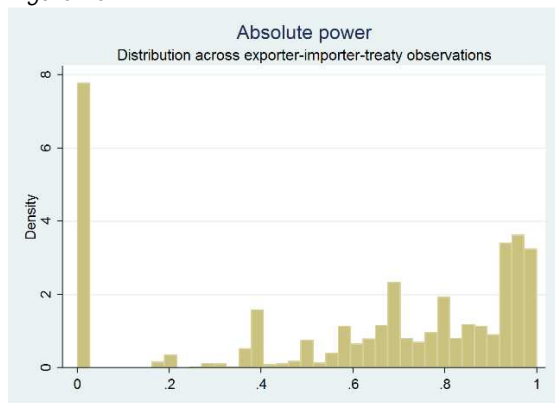



Figure 1c



Source: Author's computations, using data from Texts of Trade Agreements (ToTA).

Figure 2 illustrates the correlation between the absolute power of the importer and the absolute power of the exporter. It is positive, with a correlation coefficient of 0.49. Nonetheless, 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.

⁷ Note 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.

Figure 2. Correlation between absolute power of exporter and importer

Source: Author's computations, using data from Texts of Trade Agreements (ToTA).

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.

This paper therefore constructs 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

$$power_{i,j,t} = power(i, PTA_{i,j,t}^{m,T}) - power(j, PTA_{i,j,t}^{m,T})$$

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

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

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

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 3. Relative power

Figure 3a

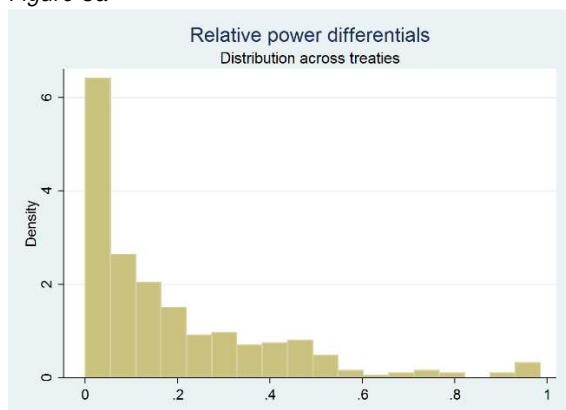
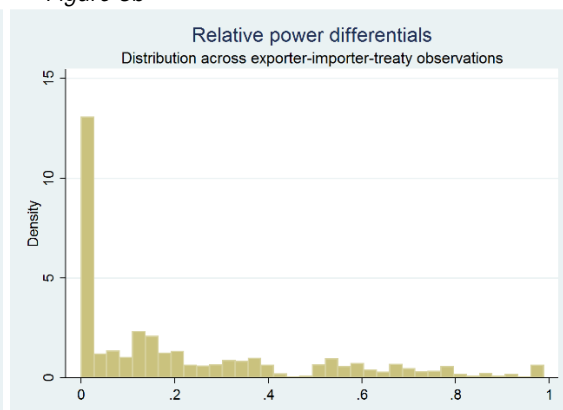


Figure 3b



Source: Author's computations, using data from Texts of Trade Agreements (ToTA).

Descriptive statistics

Table 9 (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 US\$, 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 7.2%, while the trade-weighted average is 6.1%. 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, the highest 1% of tariff values are excluded from the analysis.⁹ The maximum tariffs in the dataset used are 34.53% (simple average) and 43.3% (weighted average). Countries' GDP ranges between 100 million US\$ (Kiribati in 2004) and 18,730 billion US\$ (USA in 2016), with a mean of 419 billion US\$. GDP per capita ranges between 86 US\$ (Somalia in 2011) and 119,226 US\$ (Luxemburg in 2014).

Table 10 (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 18,177.¹⁰ The absolute power of the exporter (and, by definition, that of the importer) ranges from 0 to 1, with a mean of 0.55. At the mean, 55% of the agreement text is identical to text from the exporter's previous agreements. The mean of the absolute value of relative power is 0.25, 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 25%. North-North relationships are relatively balanced, with mean relative power of 0.16. For South-South relationships, the mean relative power is 0.23, close to the overall mean. North-South relationships are the most unbalanced on average, with a mean relative power of 0.33. Within North-South PTAs, developed countries have, on average, more influence on the text than developing countries, the mean relative power being 0.16 for developed and -0.16 for developing countries.

⁹ The results on the coefficients of interest are robust to including the outliers.

¹⁰ 18,177 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 11,762 words. This is because longer agreements have more members, i.e. more exporter-importer-year observations, than shorter agreements.

4. Empirical Strategy

Gravity model

To assess the impact of power in trade agreements on trade flows, this paper estimates 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 comprises different sets of variables, all related to the trade agreements between i and j , and will be explained in detail in the next subsection.

$$X_{ijt} = \beta * policy_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

The model is estimated using the Poisson pseudo maximum likelihood (PPML) estimator proposed by Santos Silva and Tenreyro (2006), which allows for zero trade flows. I estimate it using the Stata command `ppml_panel_sg`, as described in Larch et al. (2107), which allows for 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 2016.¹¹

Policy variables

First, as a baseline, a regression is run which only includes 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.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

Second, the measure of relative power described in section 3.3.2 is included. The main variable of interest is relative power, as this paper focuses 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).

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

$power_{ijt}$ 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. 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.

Control variables

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

¹¹ A list of countries is included in the appendix.

increase more than proportionally after concluding a trade agreement. To address this concern, control variables capturing bilateral tariffs are included in the regression.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{ijt} + \beta_3 * tariffs_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

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 unimportant goods, while a trade-weighted average is, by definition, endogenous to trade value. Hence, this paper uses both the simple and weighted average (computed from the 6-digit level of the Harmonized System classification) and compare the results.

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, this specification controls for the difference in GDP and GDP per capita as a share of total GDP of exporter and importer.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{ijt} + \beta_3 * \frac{GDP_{it} - GDP_{jt}}{GDP_{it} + GDP_{jt}} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{ijt} + \beta_3 * \frac{GDP_{it}^{pc} - GDP_{jt}^{pc}}{GDP_{it}^{pc} + GDP_{jt}^{pc}} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

PTA fixed effects

As highlighted, among others, by Baier, Yotov and Zylkin (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, agreement-level dummies for all PTAs are included in the sample.

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + PTA_{ijt}^{FE} + \varepsilon_{ijt}$$

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.¹²

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 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% identical to rule-maker i's and to 0% to rule-taker j's previous agreements ($power_{ijt} = 1$ and

¹² For example, Dai, Yotov and Zylkin (2014) find a coefficient of 0.118 when using the same methodology for a different sample.

$power_{ijt} = -1$), the rule-maker's exports would increase by 28.4%, which is about twice as much as in a "perfectly balanced" trade agreement, while the rule-taker's exports would decrease by 1.5%. In the case of an "average" power relation, in which the power differential is 0.25 (see descriptive statistics), the "rule-maker" increases its exports by about 16.3%, while the "rule-taker" only increases its exports by about 8.8%.¹³

This result is robust to including all the control variables discussed in this section. 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. The baseline regression is run on the reduced sample to assess the effect of the change in sample (column 5). The size and significance of the relative power 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, but it remains significant from both the statistical and economic point of view. This result confirms the findings by Baier, Yotov and Zylkin (2016) that part of the heterogeneity in the impact of trade agreements occurs within agreements.

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)	.115*** (.028)	.108*** (.028)	.119*** (.027)	.118*** (.026)	.118*** (.026)	-
$power_{ijt}$	-	.133*** (.047)	.161*** (.054)	.154*** (.055)	.163*** (.054)	.123** (.047)	.132*** (.047)	.083** (.035)
Tariffs								
$SimpleAverage_{ijt}$	-	-	-.003* (.002)	-	-	-	-	-
$WeightedAverage_{ijt}$	-	-	-	-.009*** (.002)	-	-	-	-
Difference in GDP								
GDP_{ijt}	-	-	-	-	-	.196 (.120)	-	-
GDP_{ijt}^{pc}	-	-	-	-	-	-	.225 (.142)	-
Fixed effects								
Ex-yr, im-yr, ex-im	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Treaty	No	No	No	No	No	No	No	Yes
Observations	358,745	358,745	247,391	247,391	247,391	358,745	358,745	358,745
R^2	.995	.995	.995	.995	.995	.995	.995	.995

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

¹³ Based on the coefficients from table 1, column 2: $e^{0.118} * e^{0.133*(0.25)} = 1.163$ and $e^{0.118} * e^{0.133*(-0.25)} = 1.088$.

5. Robustness checks

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, both absolute power variables are included separately in the regression:

$$X_{ijt} = \beta_1 * PTA_{ijt} + \beta_2 * power_{ijt}^i + \beta_3 * power_{ijt}^j + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

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 ($power_{ijt}^i$) is positive, but only significant at the 10%-level. The coefficient on the importer's influence ($power_{ijt}^j$) 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. The main results are consistent across most specifications (columns 3-7, controlling for tariffs and differences in GDP and GDP per capita), but not column 8, 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)	.174*** (.044)	.198*** (.049)	.193*** (.050)	.200*** (.048)	.172*** (.044)	.174*** (.044)	-
$power_{ijt}^i$	-	.101* (.053)	.114** (.056)	.107* (.056)	.117** (.06)	.092* (.053)	.100* (.053)	.168* (.093)
$power_{ijt}^j$	-	-.169*** (.050)	-.216*** (.058)	-.211*** (.058)	-.216*** (.057)	-.158*** (.050)	-.168*** (.051)	.007 (.088)
Tariffs								
$SimpleAverage_{ijt}$	-	-	-.003* (.002)	-	-	-	-	-
$WeightedAverage_{ijt}$	-	-	-	-.009*** (.002)	-	-	-	-
Difference in GDP								
GDP_{ijt}	-	-	-	-	-	.195 (.120)	-	-
GDP_{ijt}^{pc}	-	-	-	-	-	-	.225 (.142)	-
Fixed effects								
Ex-yr, im-yr, ex-im	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Treaty	No	No	No	No	No	No	No	Yes
Observations	358,745	358,745	247,391	247,391	247,391	358,745	358,745	241,200
R^2	.995	.995	.995	.995	.995	.995	.995	.995

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

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 re-used text. 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, this robustness check distinguishes between first and non-first agreements and include separate policy variables for relative power in either of those:

$$X_{ijt} = \beta_1 * PTA_{ijt}^{first} + \beta_2 * PTA_{ijt}^{nonfirst} + \beta_3 * power_{ijt}^{first} + \beta_4 * power_{ijt}^{nonfirst} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

$power_{ijt}^{first}$ measures the relative power for agreements that were the first for either for either i or j . It takes the value 1 for those, and the value 0 for agreements that were not the first for either partner. $power_{ijt}^{nonfirst}$ measures the relative power for agreements in situations where it was not the first agreement for either i or j . It takes the value 1 in those cases, and 0 for agreements that were the first agreement 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 5% level. 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 than in the main specification. This is consistent with the expectation that it was biased towards zero before. The coefficient on imperfectly measured (first) agreements is smaller, as the variable overestimates the power differential. 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 and non-first agreements

X_{ijt}	(1)	(2)
PTA_{ijt}^{first}	.134*** (.043)	.137*** (.040)
$PTA_{ijt}^{nonfirst}$.115*** (.026)	.118*** (.026)
$power_{ijt}^{first}$.104** (.045)
$power_{ijt}^{nonfirst}$.161** .092
Fixed effects		
Ex-yr, im-yr, ex-im	Yes	Yes
Treaty	No	No
Observations	358,745	358,745
R^2	.995	.995

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

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, this robustness check follows Mattoo, Mulabdic and Ruta (2017) and includes lags and leads of the trade policy variables in the regression. Results are reported in tables 4 and 5. Table 4 reports the results for the baseline (trade agreement dummy) and table 5 those for the relative power specification, with different combinations of up to four lags and leads of the respective policy variables.

The coefficients on the trade agreement dummy and its leads and lags are very similar in both specifications (tables 4 and 5). The earliest lead (two, three or four years before the agreement enters into force) is always positive and significant (column 1, 2 and 3, respectively). 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 (table 5). 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 impact of the trade agreement, which occurs in the year after its conclusion and remains stable thereafter.

Table 4. Lags and Leads: PTA dummy

X_{ijt}	(1)	(2)	(3)
<i>PTA_{ijt}_lead4</i>	-	-	.064** (.026)
<i>PTA_{ijt}_lead3</i>	-	.054** (.025)	.006 (.013)
<i>PTA_{ijt}_lead2</i>	.074*** (.028)	.039* (.022)	.004 (.027)
<i>PTA_{ijt}_lead1</i>	.016 (.023)	.014 (.034)	-.000 (.024)
<i>PTA_{ijt}</i>	-.001 (.015)	-.005 (.023)	.023 (.015)
<i>PTA_{ijt}_lag1</i>	.061** (.029)	.077*** (.029)	.046*** (.017)
<i>PTA_{ijt}_lag2</i>	-.002 (.028)	-.003 (.023)	.014 (.014)
<i>PTA_{ijt}_lag3</i>	-	-.007 (.025)	.003 (.014)
<i>PTA_{ijt}_lag4</i>	-	-	-.025 (.023)
Fixed effects			
Ex-yr, im-yr, ex-im	Yes	Yes	Yes
Treaty	No	No	No
Observations	264,553	231,656	201,593
R^2	.995	.995	.996

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

Table 5. Lags and Leads: Relative power

X_{ijt}	(1)		(2)		(3)	
PTA_{ijt_lead4}	-	-	-	-	.065**	(.025)
PTA_{ijt_lead3}	-	-	.056**	(.025)	.006	(.013)
PTA_{ijt_lead2}	.075***	(.028)	.039*	(.022)	.004	(.026)
PTA_{ijt_lead1}	.034	(.029)	.016	(.034)	.002	(.024)
PTA_{ijt}	-.000	(.015)	-.005	(.022)	.022	(.014)
PTA_{ijt_lag1}	.062**	(.028)	.073***	(.027)	.045***	(.017)
PTA_{ijt_lag2}	-.002	(.026)	-.001	(.021)	.016	(.014)
PTA_{ijt_lag3}	-	-	-.005	(.023)	.003	(.014)
PTA_{ijt_lag4}	-	-	-	-	-.021	(.020)
$power_{ijt_lead4}$	-	-	-	-	.100	(.095)
$power_{ijt_lead3}$	-	-	.073	(.084)	-.013	(.030)
$power_{ijt_lead2}$.039	(.072)	-.013	(.039)	-.071	(.078)
$power_{ijt_lead1}$.053	(.040)	.065	(.071)	.078	(.056)
$power_{ijt}$.037	(.051)	.043	(.034)	.024	(.042)
$power_{ijt_lag1}$.130**	(.056)	.134**	(.059)	.099**	(.044)
$power_{ijt_lag2}$	-.064	(.060)	-.017	(.027)	-.005	(.022)
$power_{ijt_lag3}$	-	-	-.061	(.054)	.002	(.025)
$power_{ijt_lag4}$	-	-	-	-	-.071	(.042)
Fixed effects						
Ex-yr, im-yr, ex-im	Yes		Yes		Yes	
Treaty	No		No		No	
Observations	264,552		231,656		201,593	
R^2	.995		.995		.996	

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

Different n-grams

As detailed in section 3, 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 the analysis, this robustness check constructs equivalent measures based on 4-, 6- and 7-grams, and estimates the main regression specification using those. The results, reported in table 6, are close to those obtained using variables based on 5-gram components.

Table 6. Using different n-grams

X_{ijt}	(1)	(2)	(3)	(4)
	4-grams	5-grams	6-grams	7-grams
PTA_{ijt}	.117***	.118***	.119***	.119***
	(.026)	(.026)	(.026)	(.026)
$power_{ijt}$.116***	.133***	.137***	.133**
	(.040)	(.047)	(.053)	(.057)
Fixed effects				
Ex-yr, im-yr, ex-im	Yes	Yes	Yes	Yes
Treaty	No	No	No	No
Observations	358,745	358,745	358,745	358,745
R^2	.995	.995	.995	.995

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

Excluding member countries of the European Union

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 European Union trade. At the same time, the European Union has a large number of trade agreements with outside partners. While not all European Union 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 European Union members' external trade relations. To test whether the specificities of the European Union affect the main results of the analysis, the main specification is estimated using a sample that excludes all European Union member states. The results are reported in table 7. 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 European Union countries, causing a downward bias in the main specification.

Table 7. Excluding European Union countries

X_{ijt}	(1)	(2)
PTA_{ijt}	.118***	.122***
	(.029)	(.029)
$power_{ijt}$	-	.151**
	-	(.060)
Fixed effects		
Ex-yr, im-yr, ex-im	Yes	Yes
Treaty	No	No
Observations	236,975	236,975
R^2	.996	.996

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

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, Seiermann and Skougarevskiy (2017)). As a source, longer agreement texts provide more text to recycle, and as a result, they provide more space for innovation. Therefore, the effect of power may differ according to agreement length. To test this, this robustness check distinguishes between agreements that are shorter and those that are longer than the median agreement (11,179 words). It includes separate policy variables for relative power for either group, similar to the distinction between first and non-first agreements reported in table 3.

$$X_{ijt} = \beta_1 * PTA_{ijt}^{short} + \beta_2 * PTA_{ijt}^{long} + \beta_3 * power_{ijt}^{short} + \beta_4 * power_{ijt}^{long} + \gamma_{it} + \delta_{jt} + \zeta_{ij} + \varepsilon_{ijt}$$

$power_{ijt}^{short}$ 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. $power_{ijt}^{long}$ 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 8. 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 5% level. This suggests that the core result is driven by longer agreements.

Table 8. Short and long agreements

X_{ijt}	(1)	(2)
PTA_{ijt}^{short}	.148*** (.033)	.147*** (.032)
PTA_{ijt}^{long}	.110*** (.027)	.113*** (.027)
$power_{ijt}^{short}$.076 (.058)
$power_{ijt}^{long}$.171** .072
Fixed effects		
Ex-yr, im-yr, ex-im	Yes	Yes
Treaty	No	No
Observations	358,745	358,745
R^2	.995	.995

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

6. Conclusion

This paper has addressed the phenomenon of recycling trade agreements texts through text-as-data analysis variables capturing the textual similarity between agreements. It 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 16.3%, while the rule-taker only increases its exports by about 8.8%. In very unbalanced situations, which are rare but do occur, the rule-maker increases its exports by about 28%, whereas rule-takers may even experience a slight decline in exports.

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 recycling 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 recycling 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", CESifo Working Paper, 5150.
- Allee, T. and M. Elsig (2015): "Are the Contents of International Treaties Copied-and-Pasted? Evidence from Preferential Trade Agreements", Working Paper prepared for presentation at the 8th Annual Conference on the Political Economy of International Organizations.
- 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?" Working paper prepared for presentation at the Annual Meeting of the International Studies Association.
- Allee, T. and A. Lugg (2016b): "Who wrote the rules for the Trans-Pacific Partnership?" *Research & Politics*, 3 (3).
- 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 (ADBI) Working Paper, 745.
- Alschner, W., J. Seiermann and D. Skougarevskiy (2017b): "Text-As-Data Analysis of Preferential Trade Agreements: Mapping the PTA Landscape", UNCTAD Research Paper, 5.
- Alschner, W., J. Seiermann and D. Skougarevskiy (2018): "Texts of Trade Agreements (ToTA): A structured corpus for the text-as-data analysis of preferential trade agreements", *Journal of Empirical Legal Studies*, 15(3).
- Alschner, W. and D. Skougarevskiy (2016a): "Mapping the Universe of International Investment Agreements", *Journal of International Economic Law*, 19, 561-588.
- Alschner, W. and D. Skougarevskiy (2016b): "The New Gold Standard? Empirically Situating the TPP in the Investment Treaty Universe", *The Journal of World Investment & Trade*, 17, 330-373.
- Alschner, W. and D. Skougarevskiy (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 Ø. Stiansen (2016): "Negotiating BITs with Models. The Power of Expertise", *PluriCourts Research Paper*, 16-13.
- Dai, M., Y. 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.
- Grimmer, J. and B. 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. Yotov and T. Zylkin (2017): "The Currency Union Effect: A PPML Re-assessment with High-Dimensional Fixed Effects", *CESifo Working Paper*, 646.
- 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.
- 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 Sciences*, 56, 84-97.
- Tversky, A. (1977): "Features of Similarity", *Psychological Review*, 84, 327-352.
-

Appendix

List of economies

Afghanistan, Albania, Algeria, Angola, Antigua and Barbuda, Argentina, Armenia, Australia, Austria, Azerbaijan, Bahamas, Bahrain, Bangladesh, Barbados, Belarus, Belgium, Belize, Benin, Bhutan, Bolivia (Plurinational State of), Bosnia and Herzegovina, Botswana, Brazil, Brunei Darussalam, Bulgaria, Burkina Faso, Burundi, Cambodia, Cameroon, Canada, Cabo Verde, Central African Republic, Chad, Chile, China, Colombia, Comoros, Democratic Republic of the Congo, Costa Rica, Côte d'Ivoire, Croatia, Cuba, Cyprus, Czechia, Denmark, Djibouti, Dominica, Dominican Republic, Ecuador, Egypt, El Salvador, Equatorial Guinea, Eritrea, Estonia, Ethiopia, Fiji, Finland, France, Gabon, Gambia, Georgia, Germany, Ghana, Greece, Grenada, Guatemala, Guinea, Guinea Bissau, Guyana, Haiti, Honduras, Hungary, Iceland, India, Indonesia, Iran (Islamic Republic of), Iraq, Ireland, Israel, Italy, Jamaica, Japan, Jordan, Kazakhstan, Kenya, Kiribati, Kuwait, Kyrgyzstan, Lao People's Democratic Republic, Latvia, Lebanon, Lesotho, Liberia, Libya, Lithuania, Luxembourg, Madagascar, Malawi, Malaysia, Maldives, Mali, Malta, Marshall Islands (the), Mauritania, Mauritius, Mexico, Micronesia (Federated States of), 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, Republic of Korea (the), Republic of Moldova (the), Romania, Russian Federation, Rwanda, Saint Kitts and Nevis, Saint Lucia, Saint Vincent and the Grenadines, Samoa, San Marino, Saudi Arabia, Senegal, Seychelles, Sierra Leone, Singapore, Slovakia, Slovenia, Solomon Islands, Somalia, South Africa, Spain, Sri Lanka, Sudan, Suriname, Swaziland, Sweden, Switzerland, Syrian Arab Republic, Tajikistan, Thailand, The Former Yugoslav Republic of Macedonia, Togo, Tonga, Trinidad and Tobago, Tunisia, Turkey, Turkmenistan, Uganda, Ukraine, United Arab Emirates (the), United Kingdom of Great Britain and Northern Ireland, United Republic of Tanzania, United States, Uruguay, Uzbekistan, Vanuatu, Venezuela (Bolivarian Republic of), Viet Nam, Yemen, Zambia, Zimbabwe, Hong Kong (China), Macao (China)

Descriptive statistics

Table 9. Summary statistics - by observation (exporter-importer-year)

	Mean	Standard Deviation	Minimum	Maximum	Observations
Trade in million US\$	503	5,272	0	504,028	367,477
PTA dummy	0.15	0.36	0	1	367,477
Tariff: Simple average	7.15	6.85	0	34.53	248,099
Tariff: Trade-weighted average	6.14	7.18	0	43.3	248,099
GDP (exporter), billion US\$	419	1,518	0.10	18,730	367,477
GDP per capita (exporter), US\$	14,952	20,298	86	119,226	367,477

Source: Author's computations, using data from UN Comtrade, Texts of Trade Agreements (ToTA), UNCTAD TRAINS, and the World Bank Development Indicators.

Table 10. Summary statistics - by observation, only with PTA

	Mean	Standard Deviation	Minimum	Maximum	Observations
First PTA for either partner	0.16	0.37	0	1	55,149
PTA length in words	18,177	18,146	758	83,024	55,149
Absolute power (exporter)	0.55	0.29	0	1	55,149
Absolute value of relative power					
<i>All pairs with PTAs</i>	0.25	0.27	0	0.98	55,149
<i>North-North pairs with PTAs</i>	0.16	0.29	0	0.98	12,306
<i>South-South pairs with PTAs</i>	0.23	0.27	0	0.98	19,179
<i>North-South pairs with PTAs</i>	0.33	0.23	0	0.98	23,664
Relative power in N-S PTAs					
<i>N-S PTAs, exporter=N</i>	0.16	0.37	-0.98	0.93	11,832
<i>N-S PTAs, exporter=S</i>	-0.16	0.37	-0.93	0.98	11,832
Observations	55,149				

Source: Author's computations, using data from Texts of Trade Agreements (ToTA).