

Alschner, Wolfgang; Seiermann, Julia; Skougarevskiy, Dmitriy

Working Paper

The impact of the TPP on trade between member countries: A text-as-data approach

ADB Working Paper, No. 745

Provided in Cooperation with:

Asian Development Bank Institute (ADBI), Tokyo

Suggested Citation: Alschner, Wolfgang; Seiermann, Julia; Skougarevskiy, Dmitriy (2017) : The impact of the TPP on trade between member countries: A text-as-data approach, ADB Working Paper, No. 745, Asian Development Bank Institute (ADBI), Tokyo

This Version is available at:

<https://hdl.handle.net/10419/179201>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc-nd/3.0/igo/>



ADB Working Paper Series

THE IMPACT OF THE TPP ON TRADE BETWEEN MEMBER COUNTRIES: A TEXT-AS-DATA APPROACH

Wolfgang Alschner,
Julia Seiermann, and
Dmitriy Skougarevskiy

No. 745
June 2017

Asian Development Bank Institute

Wolfgang Alschner is an assistant professor at the University of Ottawa. Julia Seiermann is a PhD student at The Graduate Institute in Geneva and an associate economic affairs officer at the United Nations Conference on Trade and Development. Dmitry Skougarevskiy is a PhD candidate at The Graduate Institute in Geneva and a researcher at the European University in St. Petersburg.

The views expressed in this paper are the views of the author and do not necessarily reflect the views or policies of ADBI, ADB, its Board of Directors, or the governments they represent. ADBI does not guarantee the accuracy of the data included in this paper and accepts no responsibility for any consequences of their use. Terminology used may not necessarily be consistent with ADB official terms.

Working papers are subject to formal revision and correction before they are finalized and considered published.

The Working Paper series is a continuation of the formerly named Discussion Paper series; the numbering of the papers continued without interruption or change. ADBI's working papers reflect initial ideas on a topic and are posted online for discussion. ADBI encourages readers to post their comments on the main page for each working paper (given in the citation below). Some working papers may develop into other forms of publication.

ADB recognizes "China" as the People's Republic of China; "Hong Kong" as Hong Kong, China; and "Korea" as the Republic of Korea.

We thank Veronika Zhirnova and Kseniia Tumasova for research assistance. The usual disclaimer applies.

Suggested citation:

Alschner, W., J. Seiermann, and D. Skougarevskiy. 2017. The Impact of the TPP on Trade Between Member Countries: A Text-As-Data Approach. ADBI Working Paper 745. Tokyo: Asian Development Bank Institute. Available: <https://www.adb.org/publications/impact-tpp-trade-between-member-countries-text-data-approach>

Please contact the authors for information about this paper.

Email: wolfgang.alschner@uOttawa.ca, julia.seiermann@graduateinstitute.ch, dmitriy.skougarevskiy@graduateinstitute.ch

Asian Development Bank Institute
Kasumigaseki Building, 8th Floor
3-2-5 Kasumigaseki, Chiyoda-ku
Tokyo 100-6008, Japan

Tel: +81-3-3593-5500
Fax: +81-3-3593-5571
URL: www.adbi.org
E-mail: info@adbi.org

© 2017 Asian Development Bank Institute

Abstract

This paper proposes a new method to predict the impact of preferential trade agreements (PTAs) on trade and welfare taking the Trans-Pacific Partnership (TPP) agreement as a case study. Relying on a novel dataset of treaty texts covering all trade agreements notified to the WTO, we first construct an indicator comparing existing PTAs to the TPP in terms of textual similarity. In a second step, we include this indicator into a gravity model of international trade in order to estimate the impact of different dimensions of PTA design on trade flows between TPP member countries. We derive predictions for two scenarios, the TPP with and without the United States, and compare the results to approaches using PTA incidence or depth indicators. At the aggregate level, our approach yields a slightly higher effect: In the scenario with the US, total trade between TPP partners is predicted to increase by 9.4% (162 billion USD in absolute terms), as opposed to the 2.6–6.4% (45–110 billion USD) obtained with conventional methods. Without the United States, the absolute increase would be much lower (67 billion USD), but the percentage increase of trade among the other 11 members higher (16%). A closer look at the exports of individual Asian TPP members reveals that yet more fine-grained variables are necessary to obtain reliable predictions at a more disaggregate level. Text-as-data methods offer the possibility to generate such variables through, for example, chapter- and article-level similarity measures.

Keywords/Brief Phrases: Trans-Pacific Partnership (TPP) agreement, mega-regional trade agreements, text-as-data analysis, trade agreement design

JEL Classification: F14, F13

Contents

| | | |
|-----|--|----|
| 1. | INTRODUCTION | 1 |
| 2. | LITERATURE REVIEW | 2 |
| 3. | MEASURING PTA DESIGN THROUGH TEXT-AS-DATA APPROACHES..... | 3 |
| 3.1 | PTA Dataset | 3 |
| 3.2 | Text-as-data Analysis of Textual Similarity..... | 4 |
| 3.3 | The TPP as a Benchmark | 5 |
| 3.4 | The TPP and the PTA Universe | 5 |
| 4. | EMPIRICAL ANALYSIS | 8 |
| 4.1 | Gravity Data – Sources | 8 |
| 4.2 | Approaches to Capture Policy Changes..... | 8 |
| 4.3 | Variables and Summary Statistics..... | 10 |
| 5. | RESULTS AND DISCUSSION..... | 11 |
| 5.1 | Regression Results..... | 11 |
| 5.2 | Impact of the TPP – Predictions under Different Scenarios..... | 13 |
| 5.3 | Discussion | 16 |
| 6. | CONCLUSION..... | 17 |
| | REFERENCES | 18 |
| | APPENDIX..... | 20 |

1. INTRODUCTION

Effective participation in international trade is vital to ensure sustained economic growth and achieve poverty reduction. With WTO negotiations in deadlock, countries are increasingly turning to preferential trade agreements (PTAs) to integrate their economies into the global trading system. An emerging generation of mega-regional PTAs, such as the Trans-Pacific Partnership (TPP) Agreement, the Comprehensive Economic and Trade Agreement (CETA) and the Transatlantic Trade and Investment Partnership (TTIP), is becoming ever more encompassing. These agreements bring together countries and regions that account for more than a quarter of international trade and foreign direct investment (FDI), and cover behind-the-border issues such as investment and competition, non-tariff barriers to trade such as divergent regulations and even broader questions such as environment and labor. Negotiators aim for these agreements to become the “gold standard” of trade agreements, which implies that they will influence the way trade rules are written within and outside the World Trade Organization (WTO). Due to their scope and comprehensiveness, the mega-regionals are expected to have a large impact on trade policy-making as well as on trade flows and welfare, not only in member countries but also non-members.

Existing studies of the impact of ambitious mega-regional agreements tend to focus on a limited set of PTA features - among other reasons, to make hand-coding efforts of PTA texts manageable. While those measures yield important insights, they may not capture all potentially meaningful differences in PTA content. In this paper, we use computational text analysis tools and demonstrate how it can complement existing methods to arrive at a more nuanced understanding of PTA design and its impact on trade flows. As a case study, we focus on the TPP, which is regarded as a particularly ambitious agreement in terms of both scope and membership.

Since the United States withdrew from the TPP in January 2017, the future of the agreement is uncertain. Some countries have proposed a “TPP minus one” approach (The Straits Times: 2017), while others consider the agreement substantially less interesting if it does not allow them to gain better access to the US market (Bloomberg: 2017). A joint statement after the March 2017 summit of TPP signatories and select other countries (People’s Republic of China, Republic of Korea, and Colombia) in Viña del Mar, Chile, does not provide new details, but confirms that their trade ministers are scheduled to meet again at the margins of the APEC meeting in May 2017. Furthermore, the Pacific Alliance (Chile, Colombia, Mexico and Peru), that hosted the Viña del Mar meeting, agreed “to begin trade negotiations with Asia-Pacific partners with the aim of quickly concluding comprehensive, balanced agreements that meet the high standards set by the TPP” (Muñoz: 2017). Therefore, the text of the TPP agreement continues to be an important reference for trade negotiations in the Asia-Pacific region.

In the first part of this paper, we describe a text-as-data approach to PTA analysis, and construct and discuss an indicator of textual similarity between the TPP and existing trade agreements. In the second part, we introduce a measure of textual similarity into a gravity model to estimate the effect of similarity to the TPP on trade flows and use those estimates to predict the effect of the TPP on trade flows between its member countries. We compare the results from the similarity specifications to those from existing approaches, which aim to capture the effect of the TPP through a PTA dummy, a dummy for different types of treaties, or a depth index, and discuss the differences between the approaches. As it is important for TPP member countries to assess the impact a similar treaty text without US participation would have on their trade, our

paper discusses predictions arising from two different scenarios: an agreement with the text of the TPP including (i) the initial twelve TPP members and (ii) the initial members without the US.¹ We present results on the impact on trade between all member states, as well as on exports of individual Asian member countries.

Our paper entails both a methodological contribution, suggesting the use of text-as-data based indicators for estimating the impact of trade agreements, and a contribution to the policy debate, providing estimates of the different impact of a TPP with and without the United States.

Our results show that

- The similarity of trade agreements in our dataset to the TPP ranges from less than 1% to 66%, the mean being 20%.²
- PTAs which resemble the TPP have a significantly higher positive impact on trade flows between members than those that do not.
- Predictions on the impact of the TPP on trade flows are slightly higher when using a similarity-based trade policy variable than with alternative approaches (9% versus 3–6%).
- The impact of the TPP on absolute trade flows is, as expected, much larger in the scenario in which the United States join the agreement. However, trade flows still increase considerably without US participation, and the relative increase is even larger in that scenario.
- A country-level analysis for the Asian TPP members reveals that the impact differs strongly between countries.
- Current approaches are still too coarse to yield reliable predictions at a disaggregate level, but text-as-data analysis offers opportunities to overcome these challenges, such as chapter- or article-level analysis and indicators based on the full matrix of similarities between PTAs.

2. LITERATURE REVIEW

A central question in the literature on the impact of preferential trade agreements on trade flows is how to account for the differing design of trade agreements—often referred to as their varying “depth”. Existing studies on the design of preferential trade agreements conceptualize the depth of trade agreements in a variety of ways. Some scholars measure depth through dummy variables capturing the type of agreement (FTA, customs union, common market, economic union; Baier, Bergstrand and Clance: 2015), others use a binary variable distinguishing deep from shallow agreements (Aichele, Felbermayr and Heiland: 2015), while again others refer to the number of provisions contained in an agreement (Egger and Nihai: 2015) as proxy for its depth.

¹ Another interesting scenario would include the initial TPP members (except the US) and the other Asia-Pacific countries present in Viña del Mar (People's Republic of China, Republic of Korea, and Colombia). We do not investigate this scenario because our dataset is limited to the years until 2012 due to the availability of data on internal trade flows. The years between 2013 and 2015 have witnessed a number of trade agreements between the above-mentioned countries, which our data would not be able to capture accurately.

² (Using a measure based on 5-character gram components, as explained more in detail in section 3.2).

Where studies resort to a more detailed mapping of treaty content through extensive hand-coding efforts they typically conceptualize depth in one of three ways. First, some studies use the WTO text as benchmark, identifying features in PTAs as WTO-equivalent, WTO-plus, WTO-minus or WTO-extra (e.g. Marceau: 2009, Hofmann, Osnago and Ruta: 2017; the latter also propose depth measures accounting for legal enforceability). A second approach distinguishes between different types of economic integration agreements (EIAs). Baier, Bergstrand and Feng (2015) use dummies to capture one-way preferential trade agreements, two-way preferential trade agreements, free-trade agreements and a category regrouping different types of “deeper” EIAs (customs unions, common markets and economic unions). They find that the “deeper” EIAs have a larger positive effect on trade flows than all other EIA categories. Third, other scholars derive an absolute measure of depth based on a set of content features, such as the Design of Trade Agreements Database (DESTA) (Dür, Baccini and Elsig: 2014). While all the above-mentioned approaches have already provided important insights, a more fine-grained analysis, making use of more of the information contained in the agreements’ texts, could yield an even more comprehensive understanding of trade agreement design and its impact on trade flows.

A major obstacle to the comprehensive analysis of treaty content – the resource-intensive and laborious manual mapping of agreements – is now beginning to disappear as modern technology opens up new avenues for the automated content analysis. Advances in computer science today allow for the automated extraction of information directly from texts without the involvement of human coders (Grimmer and Stewart: 2013, Spirling: 2012). Alschner and Skougarevskiy (2016a), for instance, use text-as-data methods to study the TPP’s investment chapter and compare it to existing investment agreements. They find that 81% of its text corresponds to prior American treaties but that, nonetheless, it goes further than most existing agreements in terms of investment protection, host state sovereignty and its investment arbitration architecture. Similarly, Allee and Lugg (2016) ask: “who wrote the rules of the Trans-Pacific Partnership?” and, using textual similarity analysis, conclude that the TPP displays a strong American handwriting. Text-as-data analysis thus offers new opportunities to explore the design of trade agreements and to link it to varying impacts on trade and welfare.

3. MEASURING PTA DESIGN THROUGH TEXT-AS-DATA APPROACHES

This paper builds on recent advances in text-as-data analysis to investigate the design of the TPP and to compare it to other trade agreements. We proceed in two stages. In this section, we will use text-as-data methods to extract legal content from the texts of PTAs and create a text-based indicator of similarity between each agreement and the TPP. In the next section, we will use the measures of textual similarity to the TPP to assess the impact of differences in PTA design on trade using a gravity model, and use the results to predict the impact of the TPP on trade between member countries in different membership scenarios.

3.1 PTA Dataset

For our analysis, we built a novel corpus of PTA full texts. As a source for our metadata, we used the WTO Regional Trade Agreement (RTA) Database. The RTA dataset comprises 476 treaties signed between 1948 and 2015 and notified to the WTO. It includes free trade agreements (FTAs), Custom Unions (CUs) and Regional

Economic Integration Organizations (REIOs). Of these agreements, 60% are currently in force. The remaining agreements have been signed and await ratification or have been replaced or suspended. In total, the corpus comprises PTAs signed by 202 states. To that data we add the TPP.

As a second step, we collected the texts of all agreements in the WTO RTA dataset. We deleted all annexes and schedules from these treaties to focus on the main body of each agreement. Through optical character recognition, text structure information, and manual work, we then transformed treaties from diverging formats into a unified marked-up text format (XML). This type of textual data allows us to capture the structure of agreements distinguishing between chapters, articles and paragraph as well as headers and full text.

As a third step, we divided those texts for which it was possible into chapters (70% of the English-language PTAs in our sample contain chapter-level information) and classified chapter headers by subject matter into 57 categories, such as “Investment”, “Financial Services” or “Competition”. This assignment of categories allows us to compare chapters that deal with the same subject matter but use differently worded chapter titles.

3.2 Text-as-data Analysis of Textual Similarity

We use this new text corpus to look at the textual similarity between the agreements in our dataset and their constituent chapters. Earlier work on the universe of investment treaties has shown that fully automated text-as-data analytics can successfully reveal legally meaningful differences among treaty texts shedding light on systematic treaty design variations (Alschner and Skougarevskiy: 2016b). We therefore follow Alschner and Skougarevskiy (2016b) in our operationalization of textual similarity. First, we split each document in our corpus into its underlying 5-character gram components.³ The word “free trade” would thus be split into the 6 components “free_”, “ree_t”, “ee_tr”, “e_tra”, “_trad” and “trade”. This technique is superior to alternative bag-of-words approaches, which count the occurrence of particular words, because it retains word order, which is particularly crucial in the context of legal documents (Spirling: 2012). Second, we check how many of such components overlap between documents calculating what is formally known as a Jaccard similarity coefficient – in our case a measure of textual similarity. The Jaccard similarity coefficient s_{ij} captures the share of overlap between two sets A_i and A_j – here, the two sets of 5-character grams that constitute the two treaties.

$$s_{ij} = \frac{|A_i \cap A_j|}{|A_i \cup A_j|}$$

Jaccard indices are of little value by themselves. Yet, through comparisons, they can be turned into a powerful analytical tool. Alschner and Skougarevskiy (2016b) have shown how Jaccard distances⁴ can be used to investigate treaty bargaining dynamics, trace diffusion of treaty elements over time, assess the novelty of treaty features or explore consistency and innovation in national treaty networks. In our context, Jaccard

³ We follow Spirling (2012) in setting the length of components to 5 characters. As pointed out by Lodhi et al. (2002), in the English language, “shorter or moderate non-contiguous substrings are able to capture the semantics better than the longer non-contiguous substrings”. Correlations between similarity indicators computed using different q are presented in appendix B.

⁴ Jaccard distance = 1-Jaccard similarity

similarity is used as a tool for assessing the impact of trade agreements: trade agreements that resemble each other are expected to have similar effects on trade.

Importantly, Jaccard indices can only be meaningfully compared among agreements in the same language. While the majority of texts in our dataset is in English (404) we also have a sizeable number of Spanish agreements (32). We can circumvent this language barrier by calculating similarities for each language separately and then compare ensuing counts. Fortunately, our benchmark treaty, the TPP, also exists in Spanish. We thus compare the English treaties in our dataset to the English version of the TPP and the Spanish treaties to the Spanish version. To verify whether English-language and Spanish-language similarity are indeed equivalent, we conduct a robustness check focusing on the sub-sample of treaties available in both languages (184). We compute the distance of each treaty's English version to the English version of the TPP, and of each treaty's Spanish version to the Spanish version of the TPP. The correlation between distance to the TPP in English and distance to the TPP in Spanish is 0.99, which supports our assumption that we can complete our English-language dataset with Spanish-language data without compromising the analysis.

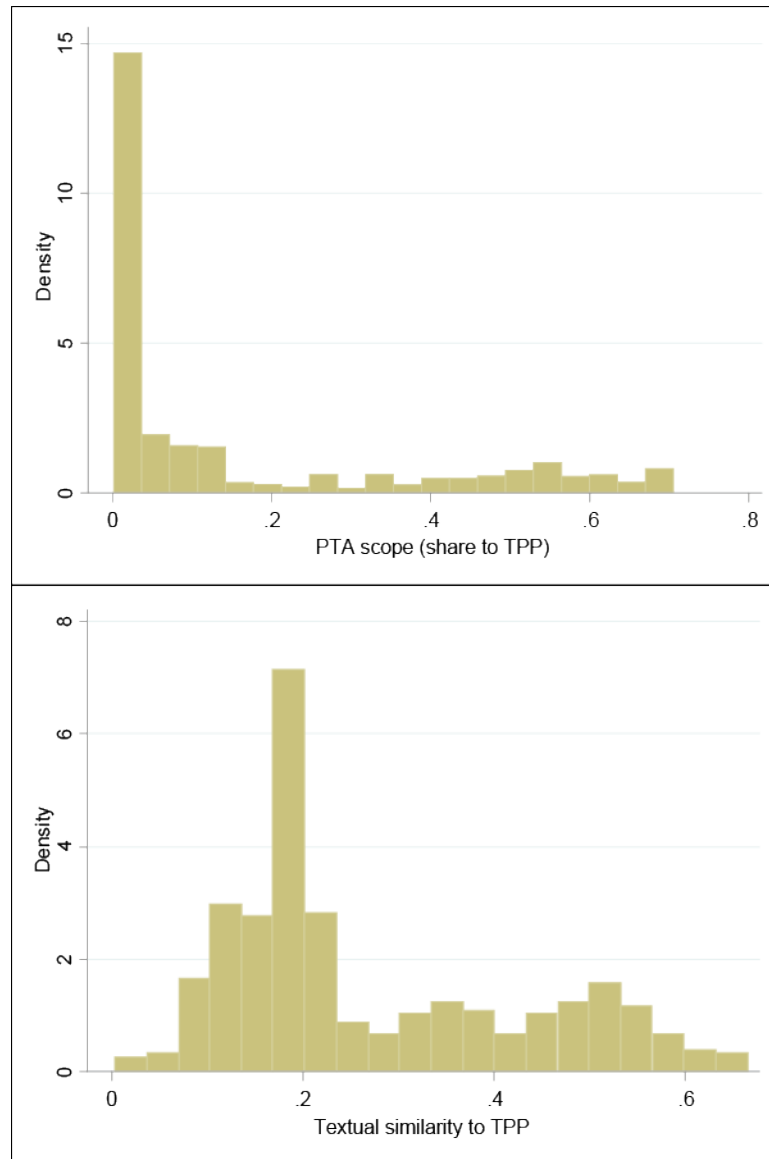
3.3 The TPP as a Benchmark

The former United States Trade Representative has described the TPP as a “high-standard, ambitious, comprehensive, and balanced agreement” (USTR: 2015). Our analysis suggests that this assessment is correct. The TPP contains chapters on 30 out of the 57 categories covered in PTAs more generally and is thus more comprehensive in scope than any other agreement in our dataset by that measure. This justifies using the TPP as a benchmark or “gold standard” (as it has sometimes been called by its proponents) to which we can compare other trade agreements.

3.4 The TPP and the PTA Universe

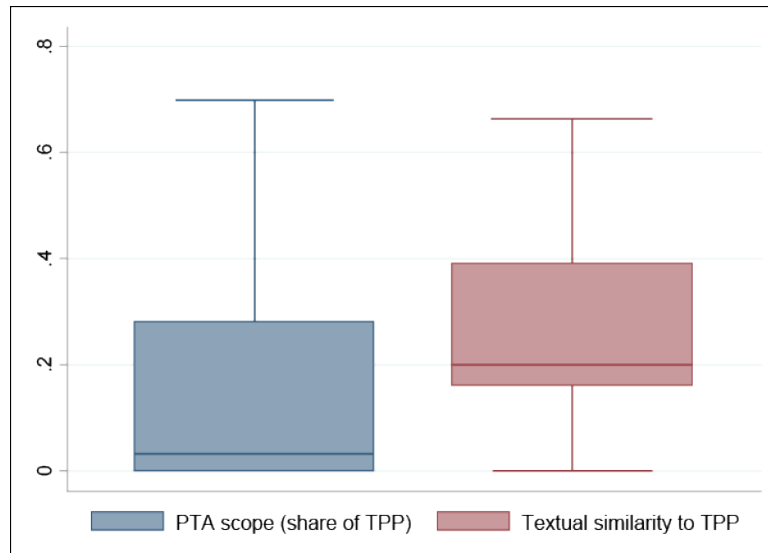
We use both an agreement's scope (the share of the 30 TPP categories it covers) and its textual similarity to compare other PTAs to the TPP benchmark. Empirically, the two measures are closely related: the correlation between them is 0.9⁵. Our preferred variables of comparison, however, is textual similarity to the TPP for three reasons. First, as shown in Figure 1, the distribution of scope is highly skewed. 169 of the 436 agreements in our database do not include any of the categories covered by the TPP. Textual similarity, on the other hand, is distributed more regularly – even the agreement least similar to the TPP has almost 5% textual overlap. Therefore, the textual similarity measure allows for a more nuanced analysis than the feature scope. Second, the aggregate count of categories included is not fine-grained enough to capture more nuanced variation within chapters. Third, the 57 categories we devised are generated based on information contained in PTA chapter titles. We thus risk omitting important variation not captured by these titles that are hidden within chapters or latently present in the agreements' structure. Textual similarity allows us to also capture this *a priori* unknown variation in our data.

⁵ The correlation is 0.87 when excluding agreements which do not cover any of the 30 TPP categories.

Figure 1: Distribution of PTAs Scope and Similarity to TPP

Source: Authors' calculations.

The universe of PTAs is very diverse when it comes to similarity to the TPP. As can be seen in Figure 2, the median PTA possesses a 20% textual similarity to the TPP and only 3% of its feature scope. Differently put, large parts of the PTA universe are made up of agreements which are not very similar to the TPP. On the other end of the spectrum, we have an upper quartile of agreements that are relatively close to the TPP textually (upper bound 66%) and in terms of features (upper bound 70%).

Figure 2: Boxplot Representation of PTAs' Scope and Similarity to TPP

Source: Authors' calculations.

Table 1: Top-25 PTAs in Their Textual Similarity to TPP

| RTA Name | Year of Signature | Textual Similarity |
|---------------------------------|-------------------|--------------------|
| US – Peru | 2006 | 66% |
| US – Colombia | 2006 | 66% |
| US – Republic of Korea | 2007 | 65% |
| US – Panama | 2007 | 65% |
| CAFTA-DR | 2004 | 65% |
| US – Oman | 2006 | 62% |
| Republic of Korea – Australia | 2014 | 62% |
| US – Australia | 2004 | 62% |
| US – Morocco | 2004 | 62% |
| Australia – Chile | 2008 | 60% |
| US – Singapore | 2003 | 60% |
| Canada – Colombia | 2008 | 59% |
| Canada – Republic of Korea | 2014 | 59% |
| Canada – Peru | 2008 | 59% |
| Panama – Peru | 2011 | 59% |
| Republic of Korea – New Zealand | 2015 | 59% |
| US – Bahrain | 2005 | 59% |
| NAFTA | 1992 | 59% |
| Peru – Republic of Korea | 2011 | 58% |
| Nicaragua – Chinese Taipei | 2006 | 58% |
| Canada – Panama | 2010 | 57% |
| Costa Rica – Peru | 2011 | 56% |
| EU – Central America | 2012 | 56% |
| Mexico – Nicaragua | 1997 | 56% |
| EU – Colombia and Peru | 2012 | 56% |

Source: Authors' calculations.

Once we take a closer look at the agreements most textually similar to the TPP (Table 1), we see the strong presence of American agreements already noted in Alschner and Skougarevskiy (2016b) and Allee and Lugg (2016). Yet, the table also suggests that there is merit in considering the entire universe of PTAs rather than only the agreements signed by TPP parties as Allee and Lugg do. The Nicaragua – Chinese Taipei PTA (2006) or the EU-Central America (2012) are both close to the TPP, with 58% and 56% textual overlap, respectively, yet they are not signatories to it. This suggests that the textual genesis of the TPP needs to be considered in a broader context of treaty design convergence in which a wider range of countries adopted similarly worded agreements.

With these descriptive measures of treaty design in mind, we now turn to the impact of varying agreement design on trade flows. In the remainder of this paper, we use the textual similarity measure obtained to estimate the impact of the TPP on trade using a gravity model. We compare the results from using trade policy indicators based on textual similarity to the TPP to those obtained through existing methods, and discuss the results as well as the caveats of the different approaches, further steps in our research, and policy implications from our findings.

4. EMPIRICAL ANALYSIS

4.1 Gravity Data – Sources

In addition to the data described in section 3, we use a number of other data sources to construct a gravity dataset covering 90 countries.⁶ It includes all TPP signatories except Brunei (for which data is not available). A list of all countries is included in the annex. Data on trade flows is obtained from UN COMTRADE. Internal trade flows, i.e. production, are from several sources. Manufacturing production is from UNIDO INDSTAT2, complemented with UNIDO INDSTAT4, the World Bank Trade, Production and Protection database and CEPII Trade, Production and Bilateral Protection (TradeProd). Agricultural production is from FAOSTAT. Trade policy indicators are from the Database on Economic Integration Agreements by Baier and Bergstrand and the Design of Trade Agreements (DESTA) database. Our dataset covers the period between 1993 and 2012; however, in the analysis, we use data at three-year intervals between 1994 and 2012, as the use of data with intervals is recommended for estimating gravity equations (Piermartini and Yotov (2016)).

4.2 Approaches to Capture Policy Changes

To estimate the impact of the TPP on trade flows, we use a gravity approach and, in particular, the Poisson pseudo maximum likelihood (PPML) estimator proposed by Santos Silva and Tenreyro (2006). We use importer-time, exporter-time, and exporter-importer fixed effects.

⁶ A list of all countries is included in appendix A.

We use five approaches that capture the policy change of moving from “no TPP” to “TPP” in different ways. The first three approaches are similar to those currently used in the literature, whereas the fourth, and fifth approach are completely novel approaches, based on our measure of textual similarity discussed in section 3. Summary statistics of the different trade policy variables are presented in section 4.3.

4.2.1 PTA versus No PTA

The first approach is the most basic: trade policy is measured by a dummy for the existence of a PTA between a country-pair:

$$X_{ijt} = \beta_1 PTA_{ij} + \gamma_{it} + \delta_{jt} + \zeta_{ij} \quad (1)$$

γ_{it} , δ_{jt} and ζ_{ij} are exporter-time, importer-time and exporter-importer fixed effects, respectively.

The trade policy variable PTA_{ij} combines the information from the Baier/Bergstrand database, the DESTA database, and our own database. It takes the value 1 if there exists a reciprocal preferential trade agreement, free trade area, customs union, common market or economic union between i and j at time t , and 0 otherwise.

4.2.2 Dummies for Agreement Type

The second approach corresponds to those used by Baier, Bergstrand and co-authors and includes dummies for different types of economic integration agreements as contained in the Baier/Bergstrand dataset.

$$X_{ijt} = \beta_1 OWPTA_{ijt} + \beta_2 TWPTA_{ijt} + \beta_3 FTA_{ijt} + \beta_4 CU_{ijt} + \beta_5 CM_{ijt} + \beta_6 EUN_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} \quad (2)$$

$OWPTA_{ijt}$ captures one-ways preferential trade agreements, $TWPTA_{ijt}$ two-way preferential trade agreements, FTA_{ijt} free-trade agreements, CU_{ijt} customs unions, CM_{ijt} common markets and EUN_{ijt} economic unions.

4.2.3 DESTA Depth Index

In the third approach, we use an additive depth index from the DESTA database. It is based on seven key provisions (a provision foreseeing the complete abolishment of tariffs, and provisions services, investment, standards, public procurement, competition and intellectual property rights) and counts, for each agreement, whether it contains substantive provisions in these areas.

$$X_{ijt} = \beta_1 D_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} \quad (3)$$

D_{ijt} captures the DESTA depth index, which lies between 0 and 7.

4.2.4 Textual Similarity with TPP

The fourth approach incorporates an index of similarity between existing agreements and the TPP, based on our measure of similarity to the TPP.

$$X_{ijt} = \beta_1 \text{similarity}_{ijt} + \beta_2 \text{PTA_nodist}_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} \quad (4)$$

similarity_{ijt} is a continuous variable which is defined as $\text{similarity}_{ijt} = 1 - \text{distance}_{ijt}^{\min}$, where $\text{distance}_{ijt}^{\min}$ is the minimum of the similarities between all the PTAs both i and j belong to at time t and the TPP. PTA_nodist_{ijt} is a dummy which takes the value 1 if an existing trade agreement is not included in our similarity dataset because it has not been notified to the WTO. It is included to avoid confounding the effect of no similarity to the TPP with the effect of missing data.

4.2.5 Textual Similarity with TPP – Quintile Dummies

The fifth approach addresses the problem that the index of similarity to the TPP may be adequate for *ex post* estimation, but not for prediction in our context. The reason is that the variable similarity_{ijt} is distributed between 0 and 0.66, but that the value it would take for country-pairs which are members of the TPP if they were to adopt it is 1. Using this value for prediction based on coefficients from a regression using only observations between 0 and 0.66 risks to yields substantially biased estimations. To obtain realistic predictions of the impact of the TPP while still using the similarity data, we construct dummies capturing quintiles of similarity to the TPP instead.

$$X_{ijt} = \beta_1 \text{sim}_{q1}_{ijt} + \beta_2 \text{sim}_{q2}_{ijt} + \beta_3 \text{sim}_{q3}_{ijt} + \beta_4 \text{sim}_{q4}_{ijt} + \beta_5 \text{sim}_{q5}_{ijt} + \beta_6 \text{PTA_nodist}_{ijt} + \gamma_{it} + \delta_{jt} + \zeta_{ij} \quad (5)$$

sim_{q1}_{ijt} takes the value 1 if an agreement is in the first quintile of similarity to the TPP, i.e. if it is among the 20% of treaties which are least similar to the TPP. sim_{q2}_{ijt} takes value 1 if an agreement is in the second quintile, etc. PTA_nodist_{ijt} again controls for the effect of agreements which are not contained in our dataset.

4.3 Variables and Summary Statistics

The summary statistics (Table 2) reveal that 35% of the country pairs in our dataset appear to be linked through a preferential trade agreement in at least one of the three datasets we are combining (Baier/Bergstrand, DESTA and our own dataset). 10% of country pairs are linked through a treaty which does not appear in our dataset, meaning that the agreement in question has not been notified to the WTO. In the Baier/Bergstrand classification, the most common type of agreement through which country-pairs are linked is one-way PTAs (13.8%), followed by FTAs (9.7%). 13% of agreements do not appear in the Baier/Bergstrand database. On average, country pairs are linked by an agreement of depth 0.869, as measured by the DESTA depth index, and of similarity to TPP 0.07. (Note that these variables are substantially larger for country-pairs for which PTA=1: mean depth is 2.5 and mean similarity 0.2). Regarding the similarity quintiles, the most common categories are quintile 3 (linking 10% of observations) and 4 (4.4% of observations).

Table 2: Summary Statistics (Country-pair Year Level)

| Variable | Mean | Std. Dev. | Min | Max |
|------------------------------------|--------|-----------|-----|---------|
| PTA | 0.351 | 0.477 | 0 | 1 |
| PTA not in our data | 0.100 | 0.300 | 0 | 1 |
| One-way PTA | 0.138 | 0.345 | 0 | 1 |
| Two-way PTA | 0.035 | 0.184 | 0 | 1 |
| FTA | 0.097 | 0.297 | 0 | 1 |
| Customs Union | 0.005 | 0.070 | 0 | 1 |
| Common Market | 0.038 | 0.191 | 0 | 1 |
| Economic Union | 0.012 | 0.108 | 0 | 1 |
| PTA not in Baier/Bergstrand data | 0.133 | 0.340 | 0 | 1 |
| Depth indicator | 0.869 | 1.668 | 0 | 7 |
| Similarity to TPP | 0.070 | 0.136 | 0 | .664404 |
| Similarity quintile 1 ⁷ | 0.019 | 0.135 | 0 | 1 |
| Similarity quintile 2 | 0.016 | 0.124 | 0 | 1 |
| Similarity quintile 3 | 0.099 | 0.298 | 0 | 1 |
| Similarity quintile 4 | 0.074 | 0.262 | 0 | 1 |
| Similarity quintile 5 | 0.044 | 0.205 | 0 | 1 |
| N | 56,700 | | | |

Source: Authors' calculations.

5. RESULTS AND DISCUSSION

5.1 Regression Results

This section presents and discusses the results from the partial equilibrium estimation and compares the quantitative outcomes of the predictions based on its results.

Table 3 summarizes the results from the regressions using approaches 1–5. The PTA dummy (column 1) is positive and significant at the 5% level. The introduction of a PTA increases trade between two countries by 13.5%. The coefficients on different types of economic integration agreements (column 2) are larger and more significant for “deeper” agreements. Both one-way and two-way preferential trade agreements do not have a significant impact on trade flows, FTAs have a relatively small impact that is significant at the 5% level, whereas customs unions, common markets and economic unions have a large impact that is significant at the 0.1% level, increasing trade by between 80% and almost 100%. The DESTA depth indicator (column 3) has a positive and significant impact. A one-step increase in depth (on a scale from 0 to 7) leads to 3.4% higher trade. Column 4 reports the results from the first specification based on the textual similarity index. The coefficient on the indicator for similarity to the TPP is 0.63 and significant at the 0.01% level. This number can be understood as follows: Moving from no regional trade agreement (similarity=0) to the TPP (similarity=1) would increase trade by 63%. However, as this represents an out-of-sample prediction, the coefficient should be interpreted with caution. In specification 4, PTAs that are not included in our dataset increase trade by, on average, 11.7%. Column 5 reports the

⁷ The number of observations for each similarity quintile is not the same, because the quintiles were computed at the treaty-level, while the data displayed in the summary statistics is at the importer-exporter-year level.

results of using dummies for the quintiles of similarity to the TPP. The coefficient on the first quintile is negative and significant – agreements which do not resemble the TPP seem to decrease trade by 30%. The coefficients on quintiles 2 to 4 are insignificant, whereas, the coefficient on quintile 5 is large and significant at the 1% level. Moving to a trade agreement that closely resembles the text of the TPP increases trade between two countries by 33%. PTAs not included in our dataset do not have a significant impact in specification (5).

Table 3: Partial Equilibrium Results

| Trade | (1) | (2) | (3) | (4) | (5) |
|----------------------------------|----------|----------|----------|----------|----------|
| RTA | 0.135* | | | | |
| | (0.0581) | | | | |
| One-way PTA | | 0.0672 | | | |
| | | (0.0648) | | | |
| Two-way PTA | | 0.151 | | | |
| | | (0.0786) | | | |
| FTA | | 0.181* | | | |
| | | (0.0756) | | | |
| Customs Union | | 0.820*** | | | |
| | | (0.107) | | | |
| Common Market | | 0.804*** | | | |
| | | (0.0905) | | | |
| Economic Union | | 0.998*** | | | |
| | | (0.115) | | | |
| RTA not in Baier/Bergstrand data | | 0.0448 | | | |
| | | (0.0551) | | | |
| Depth indicator | | | 0.0337* | | |
| | | | (0.0131) | | |
| Similarity to TPP | | | | 0.631*** | |
| | | | | (0.142) | |
| RTA not in our data | | | | 0.117** | 0.0775 |
| | | | | (0.0442) | (0.0496) |
| Similarity Quint1 | | | | | -0.302** |
| | | | | | (0.113) |
| Similarity Quint2 | | | | | 0.0680 |
| | | | | | (0.108) |
| Similarity Quint3 | | | | | 0.0513 |
| | | | | | (0.0765) |
| Similarity Quint4 | | | | | 0.0967 |
| | | | | | (0.0580) |
| Similarity Quint5 | | | | | 0.330*** |
| | | | | | (0.0676) |
| N | 55,930 | 55,930 | 55,930 | | 55,930 |

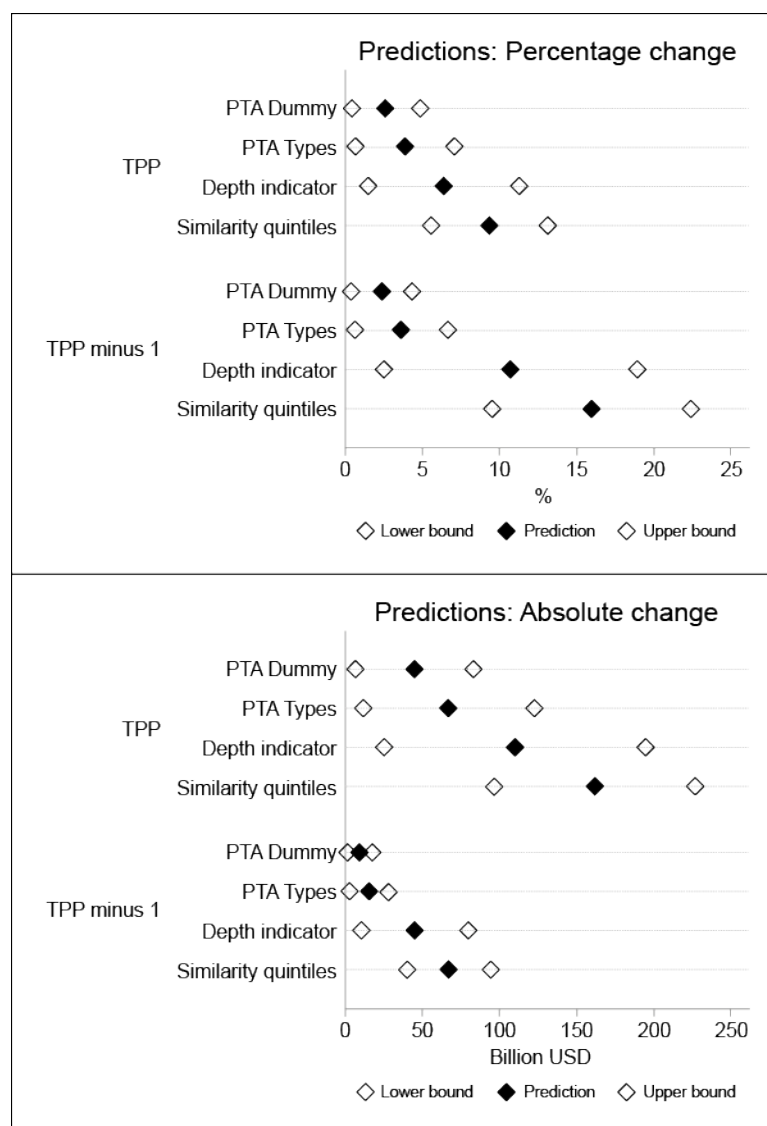
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

5.2 Impact of the TPP – Predictions under Different Scenarios

5.2.1 Total Trade between Member Countries

To predict the potential impact of the TPP, we use the results from the partial equilibrium analysis to predict the change in trade flows which would occur between TPP members if they decided to move from their current situation (some of them already have PTAs between each other, others do not) to the TPP. We use trade flows from 2012 as a basis for this prediction, as this is the most recent year in our dataset. From a methodological point of view, the simple similarity indicator cannot be used in this context, as setting it to 1 (representing a treaty with the text of the TPP) would imply an out-of-sample prediction. Therefore, we only report results from the quintile specification based on our similarity indicator.

Figure 3: Predicted Impact of the TPP on Trade between Members



Source: Authors' calculations.

As discussed in the introduction, the US withdrew from the TPP, and as a result, the remaining signatories are assessing whether and how to move forward. Therefore, we consider two potential scenarios: a treaty between (i) the initial twelve TPP members and (ii) the initial members without the US (TPP minus 1). The text of a potential future treaty is unknown, which is why we base our predictions in all three scenarios on the assumption of a “deep” treaty resembling the TPP. (In terms of the policy variables, this corresponds to a free trade area of depth 7, in quintile 5 of similarity to the TPP.)

Figure 3 summarizes the results from this simulation, presenting the aggregate impact of moving to the TPP on trade between members. It provides information on percentage change (left) and absolute change (right) to illustrate the impact of differing membership in the different scenarios. In practical terms, a large percentage increase in trade between two small countries can translate into a small absolute increase. The solid diamond indicates the point estimate, whereas the hollow diamonds represent the 95% confidence interval.

As Figure 3 illustrates, the predicted impact of the TPP or a comparable treaty with different membership varies strongly with the policy variables used. In both scenarios, the predicted change is lowest when using the PTA dummy, followed by the PTA types. The depth indicator predicts a considerably higher increase, and the similarity quintiles an even higher one. In the case of the TPP with its original members, the PTA dummy predicts an increase of trade flows between members of 2.6% (45 billion USD), the PTA type dummies 3.9% (67 billion USD), the depth index 6.4% (110 billion USD) and the similarity quintiles 9.4% (162 billion USD). While a TPP without the US would yield a similar or even higher percentage increase of trade (depending on the approach used), the absence of the largest TPP market translates into much lower absolute increases predicted for this scenario.

5.2.2 Exports of Asian Member Countries

To gain a closer insight into the implications of the different scenarios for Asian TPP members, Figures 4 and 5 display the predicted impact for Japan, Malaysia, Singapore and Viet Nam.

Figure 4 summarizes the prediction for exports from these countries to TPP (TPP minus 1) member states. As in the aggregate results discussed above, the predicted impact is always lowest when using the PTA dummy, followed by the PTA type dummies. However, comparing different countries’ results from the depth index and the similarity quintiles shows a different pattern. The similarity quintiles predict a higher increase in the exports of Japan and Viet Nam than the depth indicator, while the opposite is true for the exports of Malaysia and Singapore. This can be explained by the existing trade integration of these countries with TPP partners: Japan and Viet Nam are already linked to important importers of their goods by agreements that are “deep” as measured by the DESTA depth index, but are not in the highest quintile of similarity to the TPP. Japan has bilateral trade agreements with several Asian and Latin-American TPP members (among them, Viet Nam). All of those are relatively “deep” (6 or 7 out of 7); however, only those with Peru and Chile are highly similar to the TPP. The predictions for Viet Nam are driven by exports to Japan and the United States.⁸ Malaysia and Singapore, on the other hand, are linked to each other and various important trade partners through the ASEAN-Australia-New Zealand agreement, which

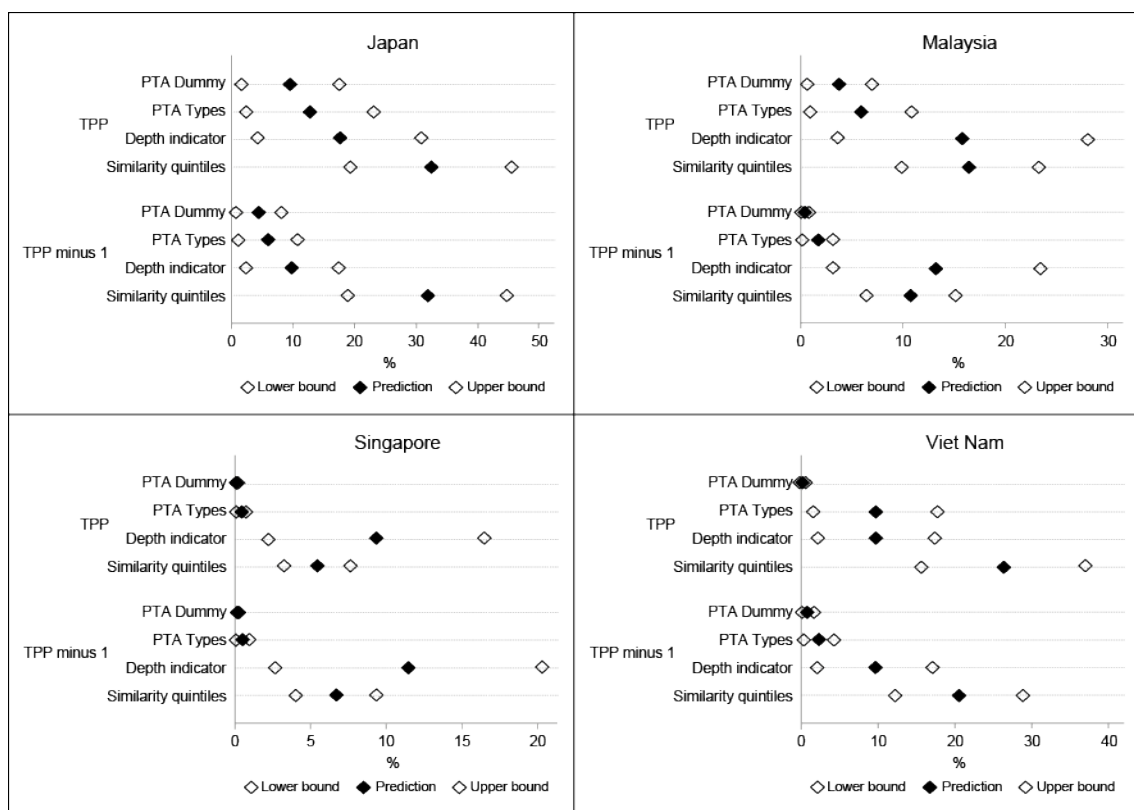
⁸ The difference in results between the depth index and the similarity quintiles approach for Viet Nam’s exports is to a large extent driven by exports to the US. While the DESTA dataset contains a trade agreement between the US and Viet Nam of depth 4, this agreement is not in our dataset, as it has not been notified to the WTO.

is not particularly “deep” as measured by the DESTA indicator, but has a similar text than the TPP.

These observations illustrate that the two variables capture different aspects of integration. Each turns out more important for some countries than for others, indicating that further, more fine-grained analysis is required to obtain adequate predictions at the disaggregate level.

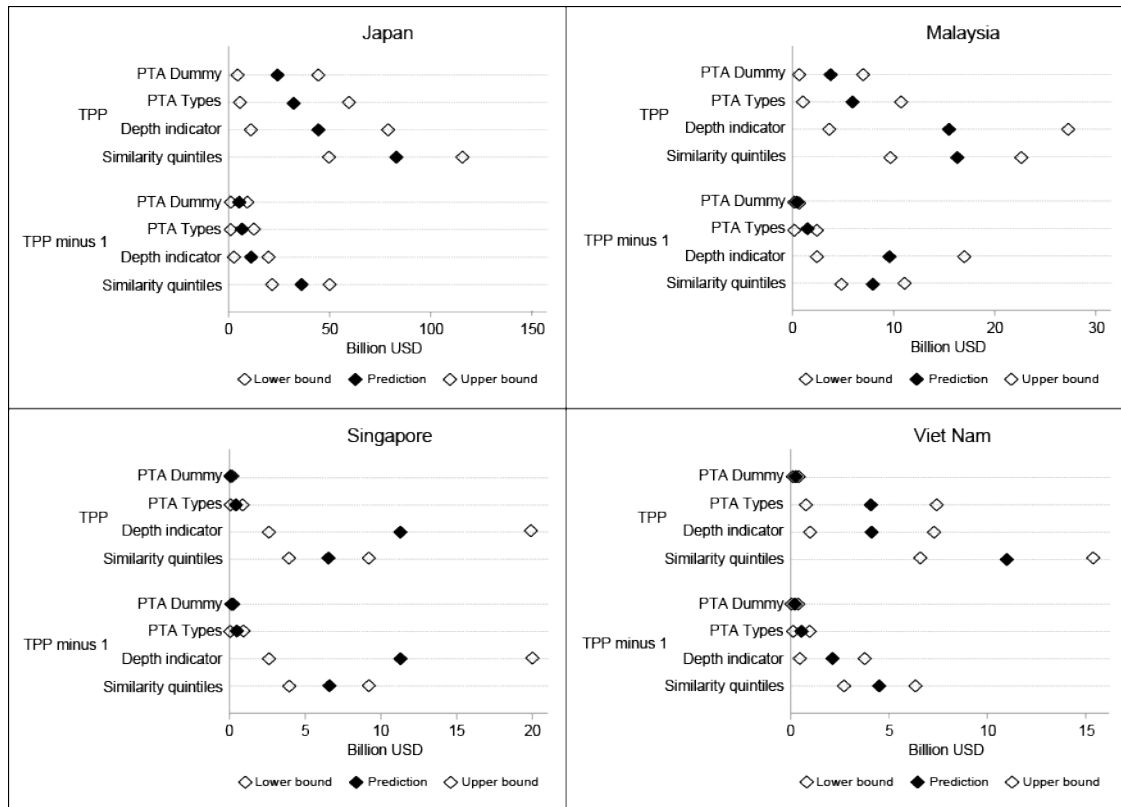
Figure 5, which illustrates the predictions in absolute terms, confirms this assessment. Most strikingly, the absolute increase in Singapore’s exports is predicted to be the same under the TPP with and without the United States. Singapore is already highly integrated with the United States through the US-Singapore agreement, which has the maximum level of depth as measured by the DESTA indicator and closely resembles the TPP. The variables used here are not yet fine-grained enough to capture the change the TPP would represent for trade integration between these two countries. Therefore, no increase in trade flows between the United States and Singapore is predicted with the TPP, which would very likely not have been the case in reality.

Figure 4: Predicted Impact of the TPP on Exports from Asian Members to Member Countries (Percentage Change)



Source: Authors' calculations.

Figure 5: Predicted Impact of the TPP on Exports from Asian Members to Member Countries (Absolute Change)



Source: Authors' calculations.

5.3 Discussion

Before discussing the results in detail, we would like to remind the reader that there are a number of trade agreements that have not been notified to the WTO and are thus not (yet) included in our dataset. An estimation with a complete textual database would yield more precise estimates. However, as our current approach controls for this issue by including a dummy for non-notified agreements, it is unlikely that the notable differences between the results from the different approaches are driven by this sample issue at the aggregate level. At the exporter, and even more the country-pair level, however, predictions may be biased due to this characteristic of the dataset. The differences at the aggregate level are explained by the fact that some approaches capture the change from existing trade policy to the TPP more precisely than others.

The PTA dummy is particularly ill-suited to capture the change from the current situation to the TPP: As 80 of the 110 TPP country-pairs already had a PTA between them in 2012, the simulation predicts no change for those pairs at all. A similar logic holds for the approach using different types of economic integration agreements: as 64 of 110 pairs already have an FTA, which is the category the TPP would be classified into, there would not be any change for more than half of the TPP country-pairs. The DESTA depth indicator captures policy change much better, as only 20 TPP country-pairs already reach the highest possible level (7) in 2012. 30 country-pairs are at depth 6, 44 at depth 0, and the remaining 16 are distributed between 2 and 4. Hence, a quite significant change is captured for approximately half the TPP country-pairs. Finally, the

approach using quintiles of similarity captures a similar extent of change, with 56 TPP country-pairs already in the fifth quintile and 54 in lower quintiles.

Overall, the aggregate predictions and the closer look at the exports of Asian TPP members show, as expected, that for some of them, the withdrawal of the United States from the TPP represents a considerable drop in the expected gains from the agreement in terms of exports. This is in particular the case for Japan and Malaysia, which did not have a PTA with the United States before. The case of Singapore is very different, as it is already strongly integrated with the United States. However, the relative and absolute gains are still considerable without the United States, which is why the 11 members may consider it worth pursuing an agreement with each other.

The country-level analysis has also highlighted that the different trade policy variables, including the treaty-level similarity index introduced in this paper, are still too coarse to capture all meaningful variation in trade integration between countries. A more detailed analysis is necessary, in particular if one wants to go beyond aggregate predictions and focus on trade relations between particular country pairs. While the similarity index represents a first step how textual analysis can contribute to a better understanding of the impact of trade agreement design on trade flows, further research is needed to explore its full potential. This includes, but is not limited to, the use of more disaggregate similarity indicators (at the chapter and article level) and indicators based on similarity between all RTAs in our dataset (as opposed to using the TPP as a reference).

6. CONCLUSION

This paper proposes to use text-as-data analysis to gain a better understanding of trade agreement design and its impact on international trade flows and illustrates a simple method to do so using the TPP as a case study. Relying on a novel textual database of preferential trade agreements, we construct an indicator of textual similarity to the TPP. We introduce a variable based on this indicator into a gravity model and generate predictions on trade flows for two scenarios, one in which the United States joins the TPP and one in which it does not. We compare our results to those using existing variables capturing the existence and nature of trade agreements between countries.

Across methods, our findings suggest that the aggregate impact of the TPP on trade between member countries is, as expected, larger in absolute terms when the United States participate in the agreement. However, trade increases remain significant without US participation, and the relative increases between the other members are even larger. This result holds across the different approaches, and implies that the remaining member countries may consider it worthwhile to conclude a similar agreement without the United States. An exporter-level analysis for the Asian TPP member countries reveals that the differences in predictions between approaches differ from country to country. This finding highlights that the current approaches, including the treaty-level similarity indicator introduced in this paper, are not yet fine-grained enough to produce reliable predictions at a disaggregate level. Text-as-data analysis offers further avenues to overcome this problem, including, but not limited to, chapter- and article-level analysis and indicators based on the full matrix of similarities between all the treaties in our dataset.

REFERENCES

- Aichele, R., G. Felbermayr and I. Heiland (2015): Going Deep: The Trade and Welfare Effects of TTIP. CESifo Working Paper No. 5150.
- Allee, T. and A. Lugg (2016): Who wrote the rules for the Trans-Pacific Partnership? Research and Politics, July–September 2016, pp. 1–9.
- Alschner, W. and D. Skougarevskiy (2016a): The New Gold Standard? Empirically Situating the Trans-Pacific Partnership in the Investment Treaty Universe. The Journal of World Investment & Trade. 17 (2016), pp. 330–373.
- Alschner, W. and D. Skougarevskiy (2016b): Mapping the Universe of International Investment Agreements. Journal of International Economic Law, 19 (2016), pp. 561–688.
- Baier, S.L., J.H. Bergstrand and M.W. Clance (2015): Preliminary examination of heterogeneous effects on international trade of economic integration agreements. In: Dür, A. and M. Elsig (2015). Trade Cooperation. The Purpose, Design and Effects of Preferential Trade Agreements, pp. 355–373.
- Baier, S.L., J.H. Bergstrand and M. Feng (2015): Economic integration agreements and the margins of international trade. Journal of International Economics 93(2014), pp. 339–350.
- Bloomberg (2017): Malaysia Reacts Coolly to Prospect of Trade Pact Minus U.S.. (13 March 2017). <https://www.bloomberg.com/politics/articles/2017-03-13/malaysia-reacts-coolly-to-prospect-of-tpp-trade-pact-minus-u-s>. Accessed 27 April 2017.
- Dür, A., L. Baccini, and M. Elsig (2014): The Design of International Trade Agreements: Introducing a New Database. Review of International Organizations 9(3), pp. 353–375.
- Egger, P. and S. Nihai (2015): Effects of deep versus shallow trade agreements in general equilibrium. In: Dür, A. and M. Elsig (2015): Trade Cooperation. The Purpose, Design and Effects of Preferential Trade Agreements, pp. 374–391.
- 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(3), pp. 267–297.
- Hofmann, C., A. Osnago and M. Ruta (2017): Horizontal Depth. A New Database on the Content of Preferential Trade Agreements. World Bank Policy Research Working Paper 7981.
- Lodhi, H., C. Saunders, J. Shawe-Taylor, N. Cristianini, and C. Watkins (2002): Text classification using string kernels. Journal of Machine Learning Research 2, pp. 419–444.
- Mantel, N. (1967): The detection of disease clustering and a generalized regression approach. Cancer Research 27 (2), pp. 209–220.
- Marceau, G. (2009): News from Geneva on RTAs and WTO-plus, WTO-more and WTO-minus. In: Proceedings of the 103rd Annual Meeting 103, pp. 124–128.
- Piermartini, R. and Y. Yotov (2016): Estimating trade policy effects with structural gravity. WTO Working Paper ERSD-2016-10.

- Santos Silva, J.M.C. and S. Tenreyro (2006): The Log of Gravity. *The Review of Economics and Statistics* 88(4), pp. 641–658.
- Spirling, A. (2012): U.S. Treaty Making with American Indians: Institutional Change and Relative Power, 1784–1911. *American Journal of Political Science* 56(1), pp. 84–97.
- The Straits Times (2017): Australia, Japan lobby for TPP-11. (21 April 2017). <http://www.straitstimes.com/asia/east-asia/australia-japan-lobby-for-tpp-11>. Accessed 27 April 2017.
- United States Trade Representative (2015): Summary of the Trans-Pacific Partnership Agreement. USTR Press Release (4 October 2015). <https://ustr.gov/about-us/policy-offices/press-office/press-releases/2015/october/summary-trans-pacific-partnership>. Accessed 27 April 2017.

APPENDIX

A. List of Countries in the Dataset

Albania, Australia, Austria, Azerbaijan, Bangladesh, Bulgaria, Brazil, Canada, Switzerland, Chile, People's Republic of China, Cameroon, Congo, Republic of the, Colombia, Costa Rica, Cyprus, Czech Republic, Germany, Denmark, Algeria, Ecuador, Egypt, Eritrea, Spain, Estonia, Ethiopia, Finland, Fiji, France, United Kingdom, Greece, Hong Kong, China, , Hungary, Indonesia, India, Ireland, Iran, Iceland, Israel, Italy, Jordan, Japan, Kazakhstan, Kenya, Kyrgyz Republic, Republic of Korea, Lebanon, Sri Lanka, Lithuania, Latvia, Morocco, Moldova, Madagascar, Mexico, Macedonia, Malta, Mongolia, Mauritius, Malawi, Malaysia, Netherlands, Norway, Nepal, New Zealand, Pakistan, Peru, Philippines, Poland, Portugal, Paraguay, Romania, Russian Federation, Saudi Arabia, Senegal, Singapore, Slovak Republic, Slovenia, Sweden, Thailand, Tajikistan, Trinidad and Tobago, Tunisia, Turkey, Tanzania, Ukraine, Uruguay, United States, Viet Nam, Yemen, South Africa

B. q-gram Jaccard Distances – Correlations between Different q

Table A.1: Correlations between English-language RTA Full Text Distance Matrices

| <i>q</i> | 3 | 4 | 5 | 6 | 7 |
|----------|-------|-------|-------|-------|---|
| 3 | 1 | | | | |
| 4 | 0.977 | 1 | | | |
| 5 | 0.932 | 0.985 | 1 | | |
| 6 | 0.884 | 0.955 | 0.992 | 1 | |
| 7 | 0.840 | 0.923 | 0.974 | 0.995 | 1 |

Note: this table reports Pearson correlations between elements of *q*-gram Jaccard distance matrices of 415 English-language RTA full texts computed under various *q*. Permutation-based Mantel (1967) tests were performed on each correlation coefficient; all tests can reject the null of zero correlation with *p*-value < 0.001.

Table A.2: Correlations between English-language RTA Chapter Text Distance Matrices

| <i>q</i> | 3 | 4 | 5 | 6 | 7 |
|----------|-------|-------|-------|-------|---|
| 3 | 1 | | | | |
| 4 | 0.981 | 1 | | | |
| 5 | 0.943 | 0.986 | 1 | | |
| 6 | 0.888 | 0.950 | 0.988 | 1 | |
| 7 | 0.826 | 0.903 | 0.959 | 0.991 | 1 |

Note: this table reports Pearson correlations between elements of *q*-gram Jaccard distance matrices of 2,837 English-language RTA chapter texts computed under various *q*. Permutation-based Mantel (1967) tests were performed on each correlation coefficient; all tests can reject the null of zero correlation with *p*-value < 0.001.