

JULY 2017

UNCTAD Research Paper No. 5
UNCTAD/SER.RP/2017/5/Rev.1**Wolfgang
Alschner**

University of Ottawa

**Julia
Seiermann**Division on
International Trade in
Goods and Services,
and Commodities
UNCTAD**Dmitriy
Skougarev
skiy**European University
in St. Petersburg &
Graduate Institute of
International and
Development Studies

Text-as-data analysis of preferential trade agreements: Mapping the PTA landscape

Abstract

Preferential trade agreements (PTAs) form an intricate web that connects countries across the globe. In this article, we introduce a PTA text corpus and research tools for its fine-grained, automated analysis. Recent computational advances allow for efficient and effective content analysis by treating text as data. We digitize PTA texts and use textual similarity tools to assess PTA design patterns on the global, national, and chapter level. Our descriptive analysis reveals, inter alia, that PTAs are more heterogeneous as a group than, for instance, bilateral investment agreements, but that they converge in regional or inter-regional clusters of similarly worded agreements. Following our descriptive account, we provide three concrete, interdisciplinary examples of how text-as-data analysis can advance the study of trade economics, politics, and law. In trade economics, similarity measures can provide more detailed representations of PTA design differences. These allow researchers to capture more meaningful variation when studying the economic impact of PTAs. In trade politics, scholars can use treaty similarity to trace design diffusion more accurately and test competing explanations for treaty design choices. Finally, in trade law, similarity measures offer new insights into the processes of normative convergence between legal regimes such as trade and investment law.

Key words: Text-as-data, preferential 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
Introduction.....	3
1. Investigating Treaty Design: From Hand-Coding to Text-as-Data.....	4
2. Mapping the PTA Landscape.....	6
3. Trade Economics - Going Deeper than "Depth".....	17
4. Trade Politics - Detecting Design Diffusion.....	20
5. Trade Law - Tracking Legal Convergence.....	23
Conclusion	26
References	27

Acknowledgements

The authors are grateful to seminar participants in Villars, Geneva and Lausanne for useful comments and suggestions on previous versions of the paper.

We thank Veronika Zhirnova and Kseniia Tumasova for research assistance and gratefully acknowledge the funding support of the SNSF project “Convergence versus Divergence? Text-as-data and Network Analysis of International Economic Law Treaties and Tribunals” (Grant Number: 162379).

Introduction

Preferential trade agreements (PTAs) have proliferated globally since the 1990s and now form a complex network of agreements connecting virtually every country in the world (Pauwelyn and Alschner 2015). While their count can diverge depending on the definition that is used, the World Trade Organization's (WTO) repository of regional trade agreements (RTAs) notified to the WTO lists 291 PTAs in force, which can be considered a conservative lower bound.¹ Many of these PTAs are hundreds of pages long and contain an intricate web of detailed obligations. Due to its growing size and complexity, the PTA universe has become increasingly hard to navigate for scholars, policy-makers, negotiators, and litigators. At the same time, the need for understanding PTAs to better assess their impact, importance, and evolution has never been greater with continued deadlock in WTO Doha Round negotiations and the emergence of new mega-regional deals such as the Comprehensive Economic and Trade Agreement (CETA) between the European Union and Canada and the on-going negotiations of a Tripartite Free Trade Area in Africa or of the Regional Comprehensive Economic Partnership (RCEP) in South East Asia.

In this article, we introduce a new structured corpus of digitized PTA full texts drawn from the WTO RTA database, and apply text-as-data tools to map the design of PTAs. We argue that textual similarity measures are particularly suitable to capture fine-grained treaty design differences and find that the term PTA regroups a set of very heterogeneous agreements, which vary systematically in scope, content and language. Moreover, in contrast to bilateral investment treaties (BITs), which largely follow the design of country-specific model agreements of developed countries, PTAs cluster in regional or inter-regional groups of similarly worded agreements. These agreements furthermore display a greater degree of customization amidst standardization than BITs. Finally, even PTAs that are similar in overall design typically display important textual variation in specific chapters.

In addition to providing a descriptive account of the design variation of PTAs, this article illustrates how similarity measures can be harnessed for specific disciplinary research applications. In trade economics, textual similarity can complement existing measures of depth to capture more of the meaningful variation of treaty design when estimating the impact of PTAs on trade flows, as we highlight for the Trans-Pacific Partnership (TPP) Agreement after the withdrawal of the United States of America. In trade politics, textual similarity metrics can help trace the diffusion of treaty design, which we exemplify by tracking the diffusion of NAFTA throughout the Pacific Rim. Finally, in trade law, textual similarity can shed new light on normative convergence processes. Investment chapters in PTAs, for instance, are usually perceived to be contributing to the convergence of trade and investment law, but, as our analysis shows, potentially also fragment the two fields.

1. Investigating Treaty Design: From Hand-Coding to Text-as-Data

The empirical study of the design of trade agreements has made major advances over the past decade. From quantitative studies that treat PTAs as a 0 or 1 depending on whether they are absent or present in an inter-state dyad, we have moved towards an ever more detailed assessment of the design of ever more treaties. Projects relying on traditional hand-coding have spearheaded these efforts. Yet, semi- or fully-automated methods that treat PTA texts as data have now become available. They have the potential of adding new scale to the empirical investigation of treaty design by analysing vast amounts of information more efficiently and effectively.

¹ WTO RTA database, available at: <http://rtais.wto.org/UI/PublicMaintainRTAHome.aspx>. Even though the WTO has adopted the term "RTA", the database also contains inter-regional preferential trade agreements. In the WTO context, the abbreviation PTAs is used for unilateral preferential tariff arrangements accorded by developed to developing countries. In this article, we use the term "preferential trade agreement" for reciprocal trade agreements between two or more partners.

1.1 Text-as-data as the new frontier in PTA content analysis

The past decade has seen a push towards a more fine-grained understanding of design variation across preferential trade agreements. In 2010, Horn, Mavroidis and Sapir investigated the content of 28 PTAs concluded by the European Union and the United States of America across 52 subject areas distinguishing between “WTO-plus” (where PTAs provide more ambitious commitments than the WTO agreements) or “WTO-extra” (policy areas covered in a PTA but not in the WTO agreements) and their legal enforceability (Horn, Mavroidis, and Sapir 2010). Subsequent work by the WTO in 2011 (WTO 2011) and, more recently, by Hofmann, Osnago, and Ruta has extended this mapping to all PTAs notified to the WTO until 2015 (Ruta, Hofmann, and Osnago 2017). Parallel efforts have been undertaken by Baccini, Dür, and Elsig to study PTAs in their own right (rather than in reference to the WTO baseline). As part of their Design of Trade Agreements (DESTA) project, researchers have coded over 600 PTAs across more than 100 content variables (Dür, Baccini, and Elsig 2014). Together, these studies have provided unprecedented insights into the design of trade agreements and spurred further research into the causes of varying PTA design as well as its effects (e.g. Dür, Baccini, and Haftel 2015; Pauwelyn and Alschner 2015; Felbermayr, Aichele, and Heiland 2016; Baccini, Pinto, and Weymouth 2017).

Hand-coding of PTAs has both advantages and disadvantages. On the one hand, if done by subject matter experts, hand-coding can capture nuanced differences in treaty design that are difficult to spot for laymen or computers. Moreover, hand coders can, through their background knowledge, differentiate between meaningful variation in language and mere stylistic variation more easily than computers. On the other hand, hand-coding is labourious, expensive and time consuming and thus less scalable. Resource constraints limit the number of variables that can be coded and make it cumbersome to extend the coding to new treaties or new variables. Hand-coding is also a deductive enterprise as researchers identify the content variation they want to investigate *ex ante* or after a sample coding. As a result, content variables outside the contemplation of the original coders or the test sample risk to be systematically omitted.

To overcome these shortcomings and to add a new scale to investigations, empirical scholars in the social sciences increasingly turn to semi- or fully automated means of content analysis as part of the emerging field variously known as “text-as-data” research (Grimmer and Stewart 2013; Gentzkow, Kelly, and Taddy 2017; Alschner, Pauwelyn, and Puig 2017), “computational social science” (Lazer et al. 2009) or the “digital humanities” (Berry 2012). What these approaches have in common is that they rely on recent advances in computer science to process social or textual data. The standardized nature of the legal language of international treaties provides particularly fertile ground for the deployment of text-as-data approaches from rule-based key word searches (Manger and Peinhardt 2017) to similarity measures akin to those employed in plagiarism detection software (Alschner and Skougarevskiy 2015) or machine learning applications (Alschner and Skougarevskiy 2016a).

Text-as-data approaches have three advantages over traditional hand-coding. First, provided that properly formatted and machine-readable text exists, text-as-data tools can analyse thousands of pages in a matter of minutes or even seconds. They are thus scalable and highly efficient when dealing with design variation in large text corpora. Second, text-as-data measures are easily adjustable and extendable making them versatile, effective, and attractive for use across projects that investigate similar data, but from different research or disciplinary angles. Third, text-as-data measures are not tied to any particular conceptual framework. They can thus be employed in concert with a multitude of theoretical lenses as well as inductively to reveal latent patterns in texts that were previously unknown. Given these advantages and potential, text-as-data can be seen as a promising new frontier in the empirical study of treaty design.

1.2 Textual similarity analysis of PTAs

Text-as-data analysis comes with an entire toolkit of approaches that can be applied to the study of PTA design (Grimmer and Stewart 2013). For instance, dictionary methods used to study tone and sentiment of texts by compiling lists of positively or negatively connoted words can be repurposed to detect characteristics in legal texts such as their legal bindingness. Supervised machine learning can be harnessed to automatically code treaty features after having been trained on a sample of hand-coded data. Unsupervised machine learning

tools, in turn, can be employed to automatically cluster treaties in groups of similar agreements. Which approach (or combination thereof) is chosen depends on the research question pursued and the level of sophistication required.

In the interest of parsimony, we restrict ourselves to one simple but powerful text-as-data method in this article: textual similarity. According to Alschner and Skougarevskiy textual similarity metrics can “reveal latent treaty design differences and trends that we did not know to exist nor actively looked for beforehand” (Alschner and Skougarevskiy 2016b: 564).² To illustrate why detecting latent design differences may be important, just ask yourself how many variables it takes to account for meaningful design variation in PTAs. DESTA uses over 100 variables to code the content of entire PTAs (Dür, Baccini, and Elsig 2014). The mapping project of the United Nations Conference on Trade and Development (UNCTAD), in contrast, already relies on over 120 variables just to code design variation of individual investment chapters within larger PTAs.³ So how many variables does it take to comprehensively map the design of an entire PTA – 500, 1000 or more? Not only is this question difficult to answer in the abstract, but responses are likely to vary between disciplines as researchers differ on the design variation they consider to be of theoretical or practical importance. Textual similarity allows us to side-step this question by using textual variation as a proxy for fine-grained design distinctions that can then be adapted to specific disciplinary research contexts.

The intuition underlying textual similarity assessments of treaty design is simple. Where language overlaps, PTAs follow similar design. Where language diverges, so does PTA design. Prior research suggests that textual similarity is a useful proxy for meaningful treaty design differences. With respect to bilateral investment treaties, Alschner and Skougarevskiy demonstrate that such metrics can expose power asymmetries in rule-making, track consistency and innovation in national treaty programs, trace diffusion of treaty elements and assess a treaty’s novelty (Alschner and Skougarevskiy 2016b). In the PTA context, similar techniques have been employed to show that the United States of America wrote the rules of the TPP (Alschner and Skougarevskiy 2016c; Allee and Lugg 2016), to highlight innovation in the Comprehensive Trade and Investment Partnership (CETA) (Allee, Elsig, and Lugg 2017a) or to account for the presence of WTO references in PTAs (Allee, Elsig, and Lugg 2017b).

There are a number of reasons why textual similarity metrics are a particularly effective proxy for design variation in the legal domain. First, legal language tends to be more uniform than natural language. Contracts, statutes and treaties are highly formulaic, structured and standardized documents full of terms of art that carry their own legal meaning. While this makes such documents harder to read for humans, they become more amenable to computational analysis as they tend not to vary for stylistic reasons alone.

Second, lawyers place a premium on predictability given that legal language has legal consequences. As a result, they employ tried and tested (“boilerplate”) language (Baird 2006). Even when they innovate, legal drafters tend to try to build on existing formulations rather than inventing new terms from scratch (see e.g. Vandeveld 2009, 108). Moreover, specifically in relation to PTAs, multilateral treaties, such as the WTO agreements, provide important reference language that is often replicated in PTA texts, adding to their homogeneity and path dependency (Allee, Elsig, and Lugg 2017b).

Third, the texts of PTAs emerge from a domestic political and institutional context that privileges textual uniformity and continuity across consecutive agreements. Legislative directives, such as the Trade Promotion Authority bill passed by the United States Congress in 2002, create clear expectations on what trade agreements should and should not contain, limiting the flexibility of the Executive.⁴ Furthermore, bureaucracies tend to rely on the language of existing treaties when negotiating new ones (Puig 2013). All these reasons make uniformity among PTA texts common – and deviations meaningful. Textual similarity is thus a particularly promising instrument from the toolkit of text-as-data methods to study treaty design variation across PTAs.

² See also www.mappinginvestmenttreaties.com.

³ UNCTAD’s IIA University Mapping Project, available at: <http://investmentpolicyhub.unctad.org/Upload/Documents/UNCTAD%20IIA%20MAPPING%20PROJECT%202013-2014.pdf>.

⁴ On the role of parliaments in creating templates for negotiations see Hornbeck and Cooper (2011) (for the US); Fecák (2011) (for Czechia).

At the same time, we acknowledge that textual similarity measures also have obvious limitations. In contrast to theoretically derived concepts describing institutional design such as a treaty's degree of commitment, flexibility or delegation (Koremenos, Lipson, and Snidal 2001; van Aaken 2009), there are no prior expectations as to how similar PTAs should be or what dissimilarity means in legal or policy terms. That is because similarity is a relational concept, which requires comparison and context to become meaningful. In consequence, textual similarity assessments are particularly useful and intuitive when studying relational processes with clear benchmarks such as negotiation outcomes based on initial positions, diffusion of texts or the novelty of language (see e.g. Berge and Stiansen 2016). Another shortcoming of textual similarity is that textual variation does not always imply a variation of meaning. Some linguistic choices have large repercussions ("shall" vs "should") while others do not. Textual similarity measures thus benefit from being combined with other text-as-data tools, hand-coding or contextual legal analysis to guide the interpretation of textual differences.

In short, textual similarity is an imperfect proxy for treaty design variation. At the same time, it is also a useful and easily scalable one. It can reveal meaningful latent design variations difficult to detect through traditional methods and can conveniently complement legal analysis and detailed feature coding. By mapping the PTA universe and showcasing a set of interdisciplinary research applications, we hope to illustrate the potential textual similarity assessment has to offer.

2. Mapping the PTA Landscape

Turning now to our empirical mapping of PTAs, we will first introduce our text corpus. We then characterize our data through more traditional proxies of treaty design such as PTA type, treaty length and feature content before investigating the textual similarity between the agreements of our corpus. We find that PTAs are heterogeneous as a group – much more so than bilateral investment treaties. This variation in terms of agreement type, scope, content and language provides important opportunities for empirical studies on PTA design in trade economics, politics, and law as we show in the next section.

2.1 Dataset

To map the PTA landscape, we have created a text corpus based on the WTO RTA database, the primary official repository of preferential trade agreements.⁵ That dataset comprises 447 PTAs that WTO members have notified to the organization. These agreements have been signed between 1948 and 2015 and involve 202 parties. 60% of them are currently in force; the remaining treaties have been signed and await ratification or have been replaced or suspended. To that list, we added the Transpacific Partnership Agreement (TPP) to account for important recent trends in PTA-making.⁶ Of these treaties, 414 are in English full texts, 33 in Spanish and one in French.

The WTO RTA database comprises four types of PTAs⁷ (with multiple assignments for the same PTA being possible in some combinations): (1) "Customs Unions" (CUs) pursuant to Article XXIV:8(a) of the GATT or Article 2(c) of the Enabling Clause (between developing countries); (2) "Goods Free Trade Agreements" (FTAs) pursuant to Article XXIV:8(b) of the GATT or Article 2(c) of the Enabling Clause (between developing countries); (3) "Services FTAs" pursuant to Article V of the GATS (listed in the RTA database as "Economic Integration Agreement"); and (4) "Partial Scope Arrangements" (PSAs) between developing countries liberalizing trade on a more limited number of products based on Article 2(c) of the Enabling Clause.

For each of these treaties, the WTO RTA database contains meta-data and links to their full texts, which we exploit to build our corpus. After collecting all full texts from the stipulated sources and manual search in cases

⁵ Regional Trade Agreements Information System, available at <http://rtais.wto.org/UI/PublicMaintainRTAHome.aspx>.

⁶ We intend to add further agreements outside the WTO RTA dataset in the future.

⁷ See the WTO RTA Database User Guide: http://rtais.wto.org/UserGuide/RTAIS_USER_GUIDE_EN.html.

of broken links, we deleted annexes and schedules 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 and machine-readable marked-up text format (XML). This type of digitized textual data allows us to capture the structure of agreements, distinguishing between chapters, articles and paragraphs as well as headers and full text for our subsequent analysis, and permits the deployment of text-as-data tools.

The WTO RTA database only lists those agreements that have been notified by WTO members. The Secretariat itself acknowledges that over 70 known PTAs that WTO members should ordinarily have notified are missing from the list.⁸ Yet, patterns existing in the WTO RTA dataset are likely to be indicative of structures present in the wider universe of non-notified PTAs. Pending future efforts to expand our corpus, the sample of the WTO RTA dataset, thus provides a useful starting point to study the PTA landscape.

2.2 The scope, content and evolution of PTAs

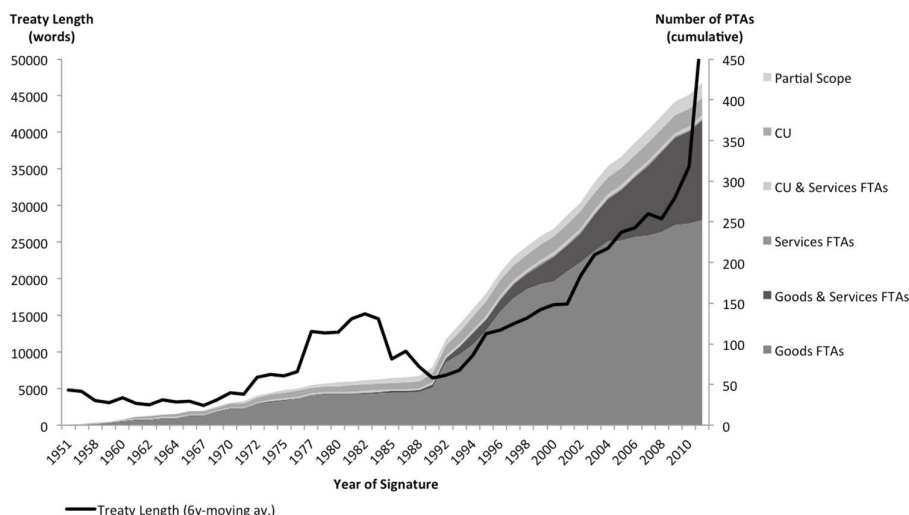
In this section, we offer a preliminary characterization of the scope, content and design evolution of the PTAs in our database. We begin by looking at treaty length. The number of words in an agreement is a first indicator of a treaty's coverage and complexity. Treaty length varies systematically across PTA types. As can be seen in Table 1, Partial Scope Agreements are shortest, with an average of 2545 words per treaty, followed by Customs Unions and Goods FTAs. Agreements that contain services components are wordier and Goods & Services FTAs are by far the longest agreements in our dataset.

Table 1. PTA type counts, length and share of texts with chapter structure

Type of PTA	# of PTAs (en, sp, fr)	mean PTA length (words)	% of PTAs with Chapters
Partial Scope Arrangements	18 (15,3,0)	2545	33%
Customs Unions	22 (19,2,1)	5672	83%
Customs Unions & Services FTAs	7 (7,0,0)	10095	71%
Goods FTAs	256 (255,1,0)	5063	59%
Services FTAs	2 (2,0,0)	15095	100%
Goods & Services FTAs	142 (116,16,0)	36360	96%

Average treaty length has increased tenfold over the past 60 years, from a mean of about 5000 words in the 1950s to over 50 000 words in the 2010s. As illustrated by Figure 1, this increase in length is driven in large parts by the conclusion of wordier Goods & Services FTAs starting in the early 2000s, which continued to proliferate while the number of existing Customs Unions stagnated and Goods FTA grew only modestly.

⁸ Secretariat Note to the Committee on Regional Trade Agreements, List of RTAs which have appeared in factual presentations and have not yet been notified to the WTO, 21 June 2016, WTO Doc Ref. WT/REG/W/104.

Figure 1. Evolution of PTAs per type and increase in length

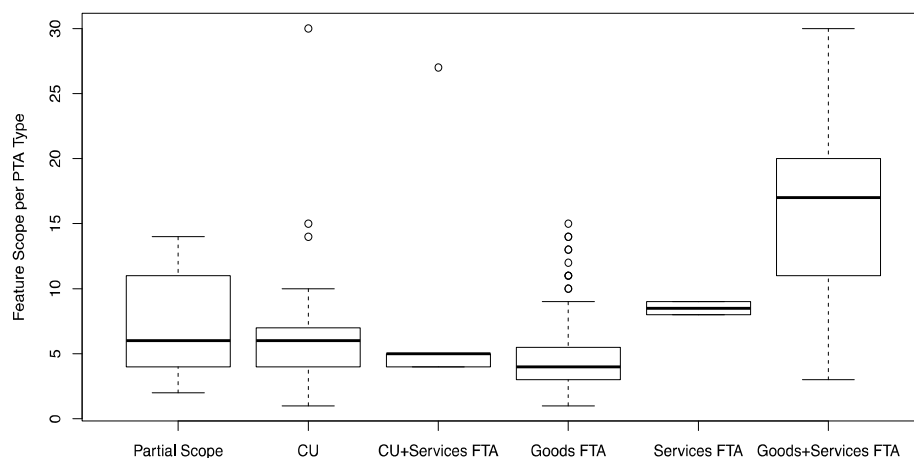
Treaty length, however, tells us little about treaty content. To investigate content, we make use of the fact that contracting states structure most PTAs along subject-specific chapters, e.g. on services, intellectual property rights or investment, to better manage their complexity. This structure is present in 319 of our 447 English-language PTA texts (71%).⁹ In total, these PTAs contain 2688 chapters, which we classify by subject matter into 73 substantive categories, such as “Investment”, “Financial Services”, or “Competition”.¹⁰ Each chapter is assigned only one such category. Aggregated by treaty this chapter-level coding provides a snapshot of the content of a PTA.

To compare the scope of content features across PTA types, we count the number of unique categories in each treaty. Plotted in Figure 2, the feature coding for PTAs with chapters corroborates the scope differences between PTA types, but adds further nuances. Goods & Services FTAs cover considerably more issues than all other types of PTAs with a median of 17 features per treaty. In comparison, the median of Goods FTAs is lowest with 4 features, while Partial Scope Arrangements, Customs Unions, and Services FTAs have slightly higher, but still single-digit median scores. While generally word counts and feature counts of PTAs are strongly and positively correlated,¹¹ this relationship holds less well with respect to Partial Scope Agreements. This is due, in part, to the fact that word counts understate the coverage of agreements, which, like the Mercosur–India PTA (2004) or the El Salvador–Cuba PTA (2011), reaffirm WTO commitments (e.g. on trade remedies, sanitary measures or technical barriers) or incorporate them by reference rather than spelling them out.

⁹ For the purpose of this paper we consider a “chapter” to be the highest level in the hierarchy of an agreement. Most agreements refer to this level as “Chapter” in the text, while others use the words “Title”, “Part” or “Section”. This definition is a simplification to streamline divergent PTA structures as some agreements include further segmentations.

¹⁰ We aggregate subject matters by area of economic integration. PTAs may thus contain multiple chapters on the same subject, if they relate to political or institutional rather than economic features. For instance, the EC Treaty has several chapters, which we regrouped under the category of “institutions” as the treaty creates several distinct bodies. Such multiple occurrences are nevertheless rare in our dataset and more frequent in customs unions than FTAs.

¹¹ Feature counts are highly, positively correlated with treaty length (0.82) for those treaties for which we have both chapter and treaty length data, suggesting that both measures capture related aspects of PTA scope.

Figure 2. Boxplots of chapter-level content features per PTA type

Shifting from scope to content, we can also use our chapter-level feature coding to characterize treaty differences in substance. Goods FTAs primarily cover chapters on trade in industrial products plus agriculture and provide for few if any treaty-based institutions. Customs Unions often accompany provisions on a common external tariff with the creation elaborate supra-national institutions and/or a common internal market.¹² Goods & Services FTA, in contrast, set out very modest, intergovernmental institutions and dedicate most chapters to spelling out trade and investment disciplines. This difference between Customs Unions and Goods & Services FTAs is, however, not a difference in the depth of economic liberalization. As the European integration project suggests, supra-national institutions can fill otherwise vague commitments with substance and incrementally expand them to cover new areas. Rather, they exemplify varying approaches to economic integration. As Baccini et al. put it, aside from “institutions-based integration, as exemplified by the European Union [or other regional integration organizations formed around customs unions], one can identify PTAs that engage in rules-based integration, exemplified by NAFTA [or goods & services FTAs more generally]” (Dür, Baccini, and Haftel 2015: 171).

In addition to corroborating known distinctions between types of PTAs, our chapter-level content mapping also highlights cross-cutting issues that expose the limitations of thinking about PTAs exclusively through the lens of the WTO’s typology. Developing countries, for instance, enjoy flexibility whether to notify their PTAs as a free trade agreement under GATT Article XXIV or under the Enabling Clause, in which case the WTO RTA database considers it a partial scope agreement. The already mentioned Mercosur–India PTA (2004) and the El Salvador–Cuba PTA (2011) illustrate that this self-classification can be misleading: with 14 and 11 content features respectively, both PTAs more closely resemble the scope and chapter coverage of the Goods & Services FTA classification. Similarly, chapters on coordination in foreign and security policy matters, which are associated with regional economic integration projects such as the European Union, COMESA, or the East African Community, are also found in several Goods and Services FTAs, which the European Union has concluded with Georgia, Ukraine, and the Republic of Moldova, showing that content overlaps between different PTA types.

In short, while the WTO typology or simple word counts capture some variation in the design of PTAs, they do not reflect the important heterogeneity amongst PTAs. In the same vein, our chapter-level coding only provides a glimpse of wider PTAs differences. We thus turn to textual similarity to assess PTA design variation in a more fine-grained manner.

¹² The TPP (a goods and services FTA concluded in 2015) and the COMESA treaty (a customs union concluded in 1993) are the two PTAs that cover most content areas (i.e. 30 out of 74 features), yet their content is strikingly different. While in the COMESA treaty, one third of chapter features deal with administrative, cooperative or institutional questions, setting up, amongst others, a supranational Court of Justice, in the TPP only one sixth of chapters is devoted to these issues and its institutions are of a more diplomatic nature and its dispute settlement ad hoc.

2.3 Text-as-data similarity assessment

To operationalize the notion of textual similarity and difference between PTAs we follow Alschner and Skougarevskiy (2016b: 564). We first split each treaty in our corpus into its underlying 5-character gram components.¹³ To illustrate, the word “free trade” would become “free_”, “ree_t”, “ee_tr”, “e_tra”, “_trad” and “trade”. The advantage of this technique as compared to word frequency counts is that word order is preserved – something particularly important for the study of legal documents (Spirling 2012). We then calculate the share of unique 5-character gram components that overlap between two treaties. The resulting number between 0 (no overlap) and 1 (100% overlap) is formally known as a Jaccard similarity – in our case a measure of textual similarity.¹⁴ Since working with multiple languages would artificially inflate Jaccard distances, we exclusively work with English-language PTAs (414 treaties) for the purpose of this section, which we standardize to British English.

Jaccard similarities are of little value by themselves, but they can be turned into a powerful analytical tool once they are used to compare similarities across treaty-pairs. In the remainder of this section we explore patterns of similarity between PTAs across three levels: (1) the global PTA landscape, (2) national PTA networks, and (3) PTA chapters on the same issue. Where appropriate, we draw comparisons to the neighbouring field of bilateral investment treaties (BITs) to highlight structural similarities and differences between these neighbouring international economic law regimes.

2.3.1 Global level: Regional clusters rather than country models

We begin our assessment with the entire landscape of English-language PTAs in our dataset. Figure 3 displays them in a heat map that color-codes textually similar treaty-pairs dark and textually different treaty-pairs light. The diagonal compares every treaty with itself. To inductively detect groups of similarly worded agreements, we order the heat map by similarity.¹⁵ The ordering algorithm groups textually similar agreements together and textually different agreements apart in order to reveal latent patterns in our data.

In Figure 3, three sets of mega-clusters and a fourth group of more heterogeneous agreements emerge. First, on the upper left, we can observe what we term the “Eastern European Cluster”. The 129 agreements in that group were mostly signed in the 1990s and primarily consist of Goods FTAs (see Table 2) signed amongst countries in Eastern Europe or between regional blocks such as the European Union and Eastern European partners. PTAs with chapter-level elements within that cluster contain on average between 3 and 4 chapters, which deal with industrial and agricultural goods supplemented by a chapter on general provisions or treaty administration. The mean pairwise similarity of agreements within that cluster is 59.8%.

The second large cluster, in the centre of Figure 3, is made up of 136 agreements, primarily Goods & Services FTAs concluded by countries from all continents in the 2000s. We hence call it the “Modern Global Cluster”. With an average similarity of 45.6%, it is more heterogeneous than the “Eastern European Cluster”, but still comprises a set of relatively similar PTAs. On average, these contain 14 chapters that mostly cover areas of economic cooperation from trade in goods and services, to competition, investment and intellectual property rights, set up modest institutions such as inter-state committees, and provide for ad hoc dispute settlement mechanisms. The “Modern Global Cluster” is where most treaty-making activity has taken place in recent years. It regroups all of Japan’s PTAs, most of the PTA signed by the United States of America including NAFTA

¹³ We follow Spirling (2012) in his choice of 5-character grams. For a correlation table on different string lengths on the same dataset, see Alschner, Seiermann and Skougarevskiy (2017).

¹⁴ Formally, for each treaty text i in our corpus let $Q(i, q)$ be a set of unique substrings of q consecutive characters, also called q -grams. Then the q -gram Jaccard distance function between treaty texts i and j is defined as $J(i, j, q) = 1 - |Q(i, q) \cap Q(j, q)| / |Q(i, q) \cup Q(j, q)|$, where $| \cdot |$ indicates set cardinality (Van der Loo 2014: 118). This function only obeys the properties of symmetry ($J(i, j, q) = J(j, i, q)$), non-negativity ($J(i, j, q) \geq 0 \forall i, j, q$), and triangle inequality ($J(i, z, q) \leq J(i, j, q) + J(j, z, q) \forall i, j, z, q$). However, identity property is not satisfied (Ukkonen 1992: 193). This means that q -gram Jaccard distance can be zero for documents which are not identical. In practical applications, this problem is less pronounced for big strings and large q .


¹⁵ We are grateful to Dr. Björn Sprungk for his help in developing the sorting algorithm.

or the TPP, recent PTAs concluded by the European Union and a majority of PTAs that connect states in South East Asia. Scholars and practitioners interested in recent innovations and more ambitious PTAs should thus devote special attention to that cluster.

Third, towards the lower right of the heat map, we see the “Central Asian Cluster” comprising 52 PTAs of countries that used to be part of the Soviet Union. They display an average treaty similarity of 33.2%. While comparable to the “Eastern European Cluster” in the sense that most agreements were concluded in the 1990s and are Goods FTAs, they are less than half as long as their Eastern European counterparts and most of them do not have chapters. PTAs in the Central Asian cluster are thus less ambitious in scope.

Finally, all other 97 PTAs either form small clusters to be investigated in future work or are very heterogeneous agreements, which together we regroup as “Other”. The bulk of PTAs in that group are customs unions and partial scope arrangements, which have been primarily concluded before the 1990s and, for those that contain chapters, have a medium scope of 7 chapter features, which aside from dealing with economic integration provide disciplines on treaty administration, treaty-based institutions and inter-state cooperation.

Figure 3. Global PTA heat map ordered by textual similarity



Note: This figure shows a heat map of Jaccard similarities between all English PTA full texts. Treaties are ordered by their similarity. Dark indicates high similarity between two treaties and light low similarity. For a more detailed investigation, we provide an interactive version of this figure on our website: <http://mappinginvestmenttreaties.com/rta/heatmap.html>

Table 2. PTA type by "cluster"

Type of PTA	"Eastern European Cluster" Majority 1990s Scope: Medium (≈4 chapters)	"Modern Global Cluster" Majority 2000s Scope: Broad (≈14 chapters)	"Central Asian Cluster" Majority 1990s Scope: Limited (no chapters)	Other Majority Pre-1990s Scope: Medium (≈7 chapters)
Customs Union			3	16
Customs Union & Service FTA				7
Service FTA		1		1
Goods FTA	125	32	49	49
Goods & Service FTA	4	102		10
Partial Scope		1		14

This structure of the PTA landscape differs markedly from what we observe in the universe of BITs. First, with 43.0% overall textual similarity BITs are more homogenous than PTAs, that display a 32.2% overall textual similarity. This is partly explained by the fact that the former are special-purpose vehicles, while the latter regroup a set of very different agreement types.

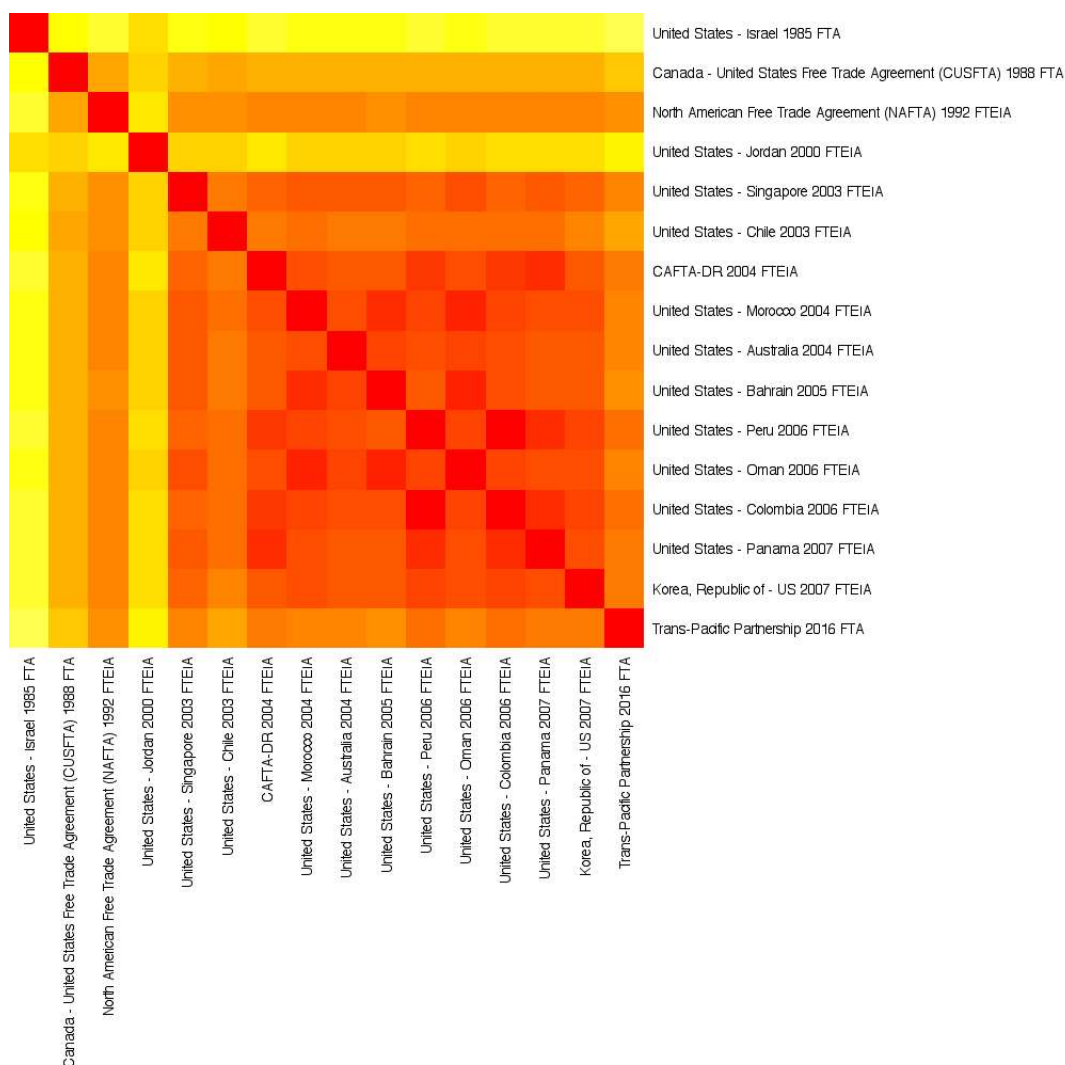
Second, and more surprisingly, whereas prior research found that power asymmetries are the major determinant for treaty design in the BIT universe, this does not seem to be the case with PTAs to the same extent. In relation to BITs, rich developed states act as rule-makers in bilateral treaty negotiations while developing states act as rule-takers opting into the model treaties put forth by their Northern counterparts (Alschner and Skougarevskiy 2016b; Allee and Peinhardt 2010, 2014). As a result, there is a strong positive correlation between a country's economic development and the consistency of its national treaty networks (Alschner and Skougarevskiy 2016b). We found no equivalent correlation between treaty consistency and economic power in relation to PTAs.

One interpretation of this difference between BITs and PTAs in light of Figure 3 is that states follow regional or even global templates or practices when it comes to PTAs, while they use national templates in the BIT context. As a result, poorer states that are treaty-design-takers when it comes to BITs can achieve consistent agreements by jumping on the bandwagon of regionally or globally converging practice in relation to PTAs. The "Eastern European Cluster" of PTAs is a case in point – highly consistent agreements being signed amongst countries with relatively low economic bargaining power in the shadow of a convergent regional practice.

2.3.2 National level: Customization and standardization

Another important reason economic power is not strongly correlated with the consistency of a country's PTA design is the fact that powerful states adapt their PTAs strategically to their negotiation partner rather than employing the one-size-fits-all approach used for signing investment agreements. In the BIT universe, countries draft model agreements in order to achieve consistent agreements over time (Berge and Stiansen 2016). The most impressive example is the United Kingdom, whose BITs have an average textual similarity of 70% across 35 years of treaty making and more than a hundred agreements (Alschner and Skougarevskiy 2016b, 576). In the PTA universe, in contrast, states more often adapt their agreements to the specific (foreign) policy concerns governing the underlying trading relationship. Treaties the United States of America signed with Israel or Jordan look different than those it signed with Canada (Figure 4); similarly, the agreements the European Union pursued with accession countries differ systematically from the agreements it signed as part of its neighbourhood policy with non-accession North African states at the same time (Figure 5).

Figure 4. Heat maps of 5-*q*-gram Jaccard distances between English-language PTA texts signed by the United States of America

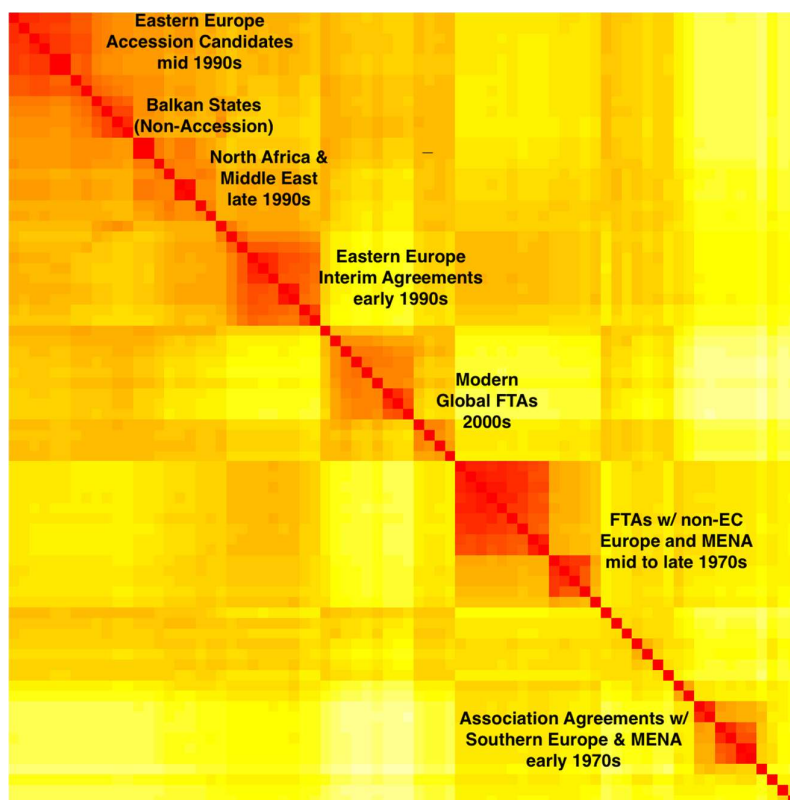


Note: This figure shows a heat map of pairwise Jaccard distances between all PTA full texts to which the United States is a party. Treaties are ordered by date of signature.

Notwithstanding this tendency of customization, we also observe a counter-trend of standardization. In the United States of America, the Trade Promotion Authority (TPA) legislation set the mandate for PTA and BIT negotiations in 2002 streamlining the design of subsequent agreements (Hornbeck and Cooper 2011). Consequently, treaties concluded in later years feature much higher similarity, as can be seen in Figure 4. In the European Union context, we see that treaty practice differed between groups of signatories depending on what type of economic cooperation was envisaged and what political relationship was involved. Within each of these groups, however, we also observe an important degree of standardization.

Hence, whereas BITs tend to follow a one-size-fits-all approach, PTAs seem to be more prone to differentiation to accommodate varying political and economic objectives while at the same time pursuing consistency through standardization in agreements where objectives overlap.

Figure 5. Heat maps of 5-gram Jaccard distances between English-language PTA texts signed by the European Union / European Communities



Note: This figure shows a heat map of pairwise Jaccard distances between all PTA full texts to which the European Union / European Communities is a party. Treaties are ordered by treaty similarity


2.3.3 Chapter level: Systematic variation between subject areas

A final important difference between BITs and PTAs in terms of textual similarity is that BITs relate to one subject matter – investment – while PTAs cover a host of different issue areas including investment. It is thus instructive to compare PTAs not only on the treaty-, but also on the chapter-level. As Figure 6 suggests there is significant variation across the specific subject matters contained in PTAs. Whereas some subject matters are relatively homogeneous across treaties, such as textile, telecommunications, financial services or transparency, which all display a mean similarity score of more than 50%, other issue areas are very diverse. The latter group includes transportation, membership or economic integration chapters, with less than 15% of mean textual overlap.

Studying chapter-level differences can help fine-tune the comparison between treaties. While two PTAs could appear very different overall, they may contain specific chapters that are rather similar. Vice versa, whereas PTAs may be very similar overall, some of their chapters may be very different. The United States of America-Morocco and United States of America-Australia FTA, for instance, were signed in the same year and display a textual similarity of 79%, yet, for example, their investment chapters are only 60% similar since the former contains investor-state arbitration whereas the latter does not. What chapters drive differences between treaties can be explored in detail on our accompanying website.¹⁶

¹⁶ <http://mappinginvestmenttreaties.com/rta/heatmap.html>

Figure 6. Mean textual similarity across chapter categories (excluding features present in 5 treaties or less)

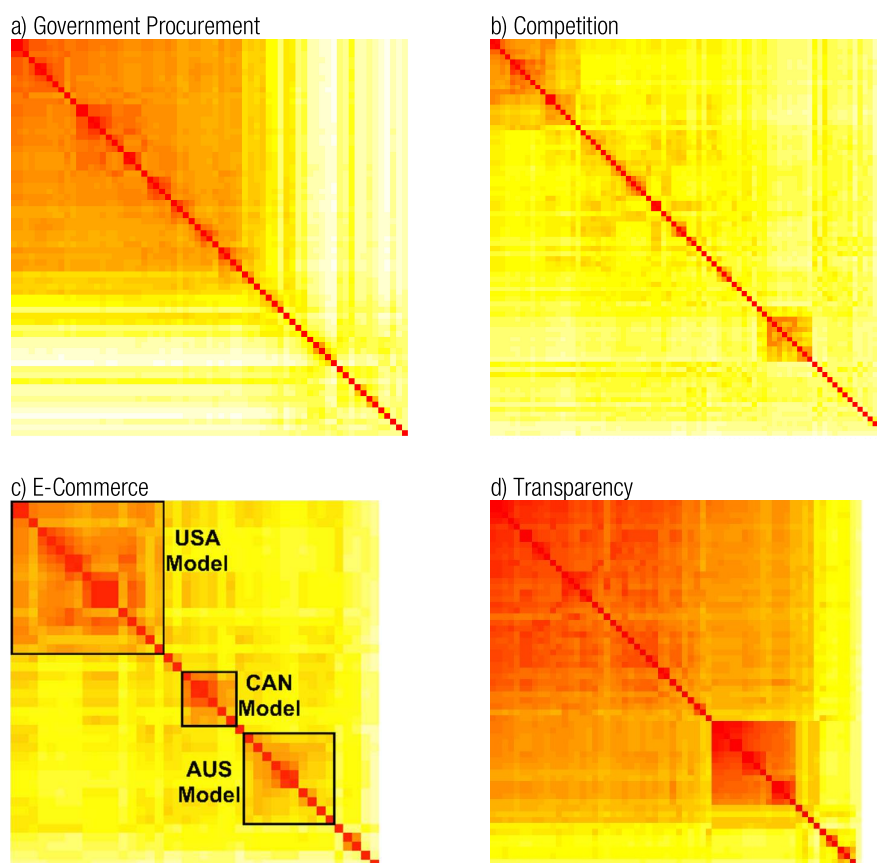


Note: The number in brackets after each chapter category indicates the number of treaties in which this category occurs.

Aside from adding nuance to differences between treaty-pairs, chapter-level comparisons can reveal systematic variation across subject matters. Competition and government procurement chapters, for example, are typically included in the same set of agreements.¹⁷ One may thus suspect that their design variation is correlated. Subfigures 7a) and b), however, illustrate that both types of chapters behave very differently. While the majority of government procurement chapters follows a common design (perhaps due to the unifying presence of parallel WTO language), the vast majority of competition chapters are textually dissimilar. In other instances, competing treaty design choices fragment one issue area into clusters, while converging practice unifies another. As illustrated in Subfigures 7c) and d), the United States of America, Canada, and Australia all have different approaches when it comes to the regulation of e-commerce resulting in three distinct design clusters. Transparency chapters, in contrast, follow a similar design across PTAs in our dataset. Hence, the comparison of chapters reveals that design can converge in one issue area, but diverge in another, including where the same set of treaties is involved.

¹⁷ Both features appear predominantly in Goods and Services FTAs of the "Modern Global Cluster". Of these FTAs, 70 have competition and 63 have government procurement chapters. Importantly, in roughly 3 out of 4 of these FTAs, both chapters occur in the same treaty.

Figure 7. Textual similarity between chapter level features, ordered by similarity



Note: This figure shows a heat map of pairwise Jaccard distances between PTA chapters on a given subject matter. Treaties are ordered by treaty similarity.

3. Trade Economics - Going deeper than "Depth"

In the realm of applied trade economics, scholars seek to understand the effect of trade agreements on economic indicators such as trade flows and, at the end of the day, welfare. They thereby strive to answer a host of highly policy-relevant questions. What treaty design is most trade enhancing? Do PTAs create or divert trade? In recent years, research has moved beyond binary measures that indicate the absence or presence of a trade agreement to account for the insight that differences in PTA design will have varying impacts on trade flows (see e.g. Baier, Bergstrand, and Clance 2015; P. Egger and Nihai 2015). In this endeavour, empirical trade economics can rely on textual similarity to complement feature coding of agreements in order to provide more nuanced proxies for design differences.

3.1 Why current "depth" measures may not be deep enough

The primary means by which trade economists and political economy scholars operationalize variation across PTAs is by measuring their "depth" (see e.g. Dür, Baccini, and Elsig 2014; P. Egger and Nihai 2015; Felbermayr, Aichele, and Heiland 2016). It is assumed that PTAs can be grouped alongside a spectrum. On

one end of it appears a PTA that liberalizes hardly any trade amongst its signatories. On the other end exists a PTA that leads to tight economic integration between participating territories. While seemingly intuitive, we argue that this monolithic concept of PTAs' depth risks oversimplifying the more complex relationship between treaties and trade flows. This, in turn, should motivate researchers to go deeper than depth and model PTA design differences more accurately to capture their impact more precisely.

In empirical economic research, the concept of depth has come to regroup a host of approaches measuring different aspects of PTA design. Baier, Bergstrand, and Clance, for instance, implicitly follow Balassa's theory of stages in economic integration which conceives depth as a progression from a free trade agreement to a customs union, to a common market, and finally to an economic union, i.e. in terms of agreement type (Baier, Bergstrand, and Clance 2015). In their adaptation, the European Union's Euro Zone first set out in the Treaty of Maastricht or the Economic and Monetary Community of Central Africa (CEMAC) are amongst the deepest PTAs in existence. Dür, Baccini, and Elsig, in contrast, as part of the Design of Trade Agreements (DESTA) project, conceive of depth as an agreement's substantive scope (Dür, Baccini, and Elsig 2014). Each PTA is investigated for the presence of seven features: a tariff reduction to zero and the existence of provisions on services, intellectual property, investment, standards, government procurement and competition. In their ranking the Maastricht Treaty has a depth of 4 and the CEMAC treaty only of 2, while, for instance, NAFTA is amongst the deepest agreements, scoring a 7.

The main problem of such depth classification, however, is not the inconsistency between studies operationalizing depth in contradictory ways, but the concept of depth itself. As it is currently used, depth reduces meaningful PTA design variation to a uni-dimensional scale from shallow to deep. Yet, PTA design actually varies meaningfully in multi-dimensional ways, as we showed above. Due to this multi-dimensionality, Baier et al. and Dür et al. can disagree on what the "deepest" PTA is, yet still both be right, since they are measuring different, yet equally meaningful aspects of PTA design variation. Similarly, within each ranking, multi-dimensionality of PTA design variation creates awkward results. In DESTA, for instance, the European Union's founding Treaty of Maastricht displays a lower level of depth than the agreements the block has signed subsequently with, for instance, Caribbean states (the European Union-CARIFORUM EPA). Yet, few economists would consider the level of economic integration to be higher between these states and the European Union than within the European Union itself. Similar situations occur in the work of Baier, Bergstrand, and co-authors who classify PTAs with limited economic integration in terms of substantive scope, such as the Eurasian Economic Community, as deeper than high economic integration treaties such as NAFTA. These issues arise when meaningful multi-dimensional variation is reduced to a uni-dimensional metric. To precisely estimate the impact of PTAs on trade flows and welfare, we thus argue that any measure of treaty design has to capture the multi-dimensionality of its meaningful variation. In a sense, we must go "deeper" than depth.

Textual similarity offers a way to get around the problem of creating coherent and comprehensive ex ante measures of "depth" by capturing multi-dimensionality in PTA variation inductively. As a fine-grained measure of PTA design, it picks up content variation that is associated with both the Baier et al. typology and DESTA's substantive coverage classification, while also including elements that fall through their coarser lenses. Even though similarity is thus no substitute to a theoretically-informed and multi-dimensional account of what treaty design matters for trade flows it offers a way to side-stepping that question in selected research contexts.

3.2 Predicting the impact of the TPP-1 on trade flows

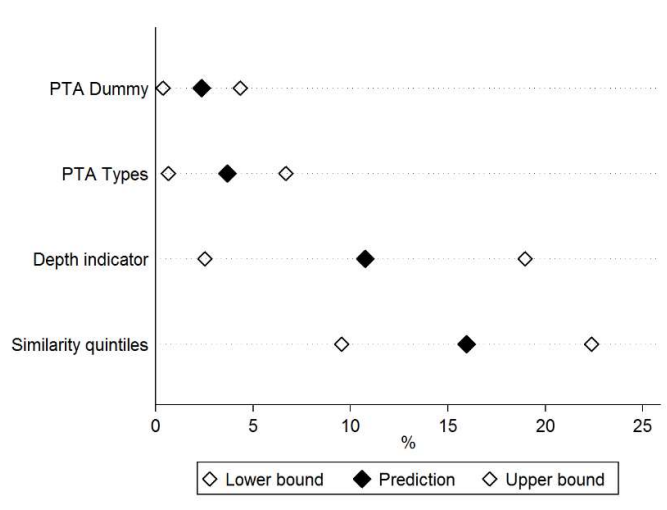
To illustrate how variables derived from text-as-data methods can be used in econometric studies, we offer a simple example using textual similarity to predict the impact of the TPP-11 (a PTA with the text of the TPP between the 11 TPP members remaining after the withdrawal of the United States of America) on trade flows.¹⁸ The intuition is that we expect PTAs that resemble each other to have a similar impact on trade. If we want to predict the impact of the TPP, we can approximate it by looking at the impact of agreements similar to the TPP. To compare our results, we also include three alternative proxies for treaty design in our standard gravity framework.

¹⁸ The example is taken from (Alschner, Seiermann and Skougarevskiy 2017).

Our dataset comprises international and internal trade flows of 90 countries at three-year intervals over the period of 1994-2012. We estimate a gravity equation including importer-year, exporter-year and exporter-importer fixed effects using a Poisson pseudo maximum likelihood (PPML) estimator.¹⁹ The trade policy change of adopting the TPP is captured by different variables in four different approaches. First, we use a dummy for the existence of a PTA between two countries. Second, we follow Baier, Bergstrand, and co-authors and include a set of dummies for different agreement types (one-way preferential agreement, two-way preferential agreement, free-trade agreement, customs union, common market and economic union). Third, we include the DESTA depth index of integration between, which lies between 0 and 7. Finally, we use a measure capturing the Jaccard similarity of the PTA in force between two countries to the TPP. More specifically, we include dummies for quintiles of similarity to the TPP in our gravity equation.²⁰ For example, the United States of America-Peru PTA, which entered into force in 2009, is most similar to the TPP, with a coefficient of 66%. Therefore, the dummy for the fifth similarity quintile takes the value 1 for the country pair United States of America-Peru in 2009 and later years.²¹

The results, including the 95% confidence interval, are presented in Figure 8. The predicted impact of the TPP varies substantially depending on the approach used. While the PTA dummy predicts an increase in trade of only 2.6%, the agreement type dummies yield 3.9%, the depth index 6.4% and the similarity quintiles 9.4%. This divergence can be explained by the fact that the variables are not equally fine-grained and that they capture different aspects of the policy change of moving from the current situation to the TPP.

Figure 8. Impact of TPP-11 on trade between member states (in %)



Future extensions using text-as-data methods will allow us to construct a variety of more disaggregate variables, such as chapter-level and article-level similarity, which will allow us to uncover the elements that drive these differences and yield more precise estimates.

¹⁹ We use the Stata command `ppml_panel_sg`, as described in (Larch et al. 2017).

²⁰ In this simple example, we do not account for the problem of joint determination of trade flows and trade agreements, as addressed, among others, in (Egger and Tarlea 2017).

²¹ We use similarity quintiles as using the similarity coefficient directly would imply making out-of-sample predictions. Similarity to the TPP ranges from 5% to 66% in the sample of observations, while the TPP has a similarity of 100% to itself.

4. Trade Politics - Detecting Design Diffusion

The diffusion of trade-related norms, policies and institutions across states has become a core area of research for international political economy scholars. Why did developing countries follow their Northern counterparts in launching trade and investment treaty programs? Under what conditions do states reproduce existing treaty design and when do they innovate? Which actors and factors drive the diffusion, evolution and sometimes revolution of treaty design choices? What all these questions have in common is that treaty design is not the explaining (independent) variable, the effect of which is measured, but rather an outcome (dependent) variable that is to be explained. Text-as-data approaches can help in that endeavour.

4.1 Understanding design diffusion

The principal challenge for scholars studying treaty design diffusion is to differentiate between competing explanations for the proliferation of a given treaty design (Dobbin, Simmons, and Garrett 2007; Gilardi 2012). Is coercion, competition, learning, emulation, or something else at work? When a PTA is expressed in binary terms as a 0 or 1, meta-information on treaty signatories or year of conclusion provides the only clues for a treaty's emergence. However, once the varying design of agreements is taken into account new evidence emerges that can assist in buttressing or refuting competing explanations of policy diffusion. Investigating variation in dispute settlement provisions of BITs, for instance, Allee and Peinhardt show that investors' home states' power politics explains investment treaty design better than the investors' host states' need to make commitments credible (Allee and Peinhardt 2014). Similarly, text-as-data analysis can help empirically disentangle competing explanations by adding new nuance to the study of treaty design variation.

In the universe of PTAs, similarity across treaties has already been fruitfully applied to study PTA design and its diffusion. Studies have, for instance, compared the Transpacific Partnership (TPP) to its overlapping agreements finding that it bears a striking resemblance to prior treaties concluded by the United States of America in the early 2000s (Alschner and Skougarevskiy 2016c; Allee and Lugg 2016). Power, path dependency, and emulation around an American model thus emerge as more compelling explanations of the TPP's design than learning or rational design adjustments. Other works have challenged the conventional wisdom of conceiving PTAs as a counterweight to the WTO by showing that WTO references and language in PTAs have actually increased over time and that those countries signing more PTAs are also the ones tying their agreements most closely to the WTO (Allee, Elsig, and Lugg 2017b). This emulation suggests that states, at least in part, perceive PTAs as a means to strengthen rather than to undercut WTO norms.

4.2 Diffusion of NAFTA treaty design

To further illustrate how textual similarity of PTAs can advance the study of treaty design and its diffusion, we investigate the proliferation of NAFTA language and offer some tentative explanations. When it was signed in 1992, NAFTA was amongst the most far-reaching free trade agreements ever concluded, covering issues from trade in goods and services to investment, intellectual property rights, and competition in addition to also containing side agreements on labour and the environment. Early commentators saw in NAFTA a milestone that was to set the standard for future PTAs (Price 1993: 736; Deng and Braun 1997: 333). They were not far off the mark. Scholars in both law and political science have since noted that NAFTA design elements have diffused across the globe; some speak of a "NAFTA model" (Estevadeordal et al. 2005; Dür, Baccini, and Haftel 2015), a "convergence" around NAFTA (Lavranos 2013), or a "NAFTA-ization" of treaty design (Berger 2013).

While the proliferation of NAFTA-inspired treaty design elements is widely acknowledged, the path and the drivers of this diffusion remain subject to debate. For Grugel, for instance, the spread of NAFTA design elements is linked to power and hegemony of the United States of America as the country attempts an "implantation" of its preferred set of neoliberal rules abroad especially in Latin America (Grugel 2004). For Baccini et al. rational design rather than power explains the growth in NAFTA design. They argue that NAFTA represents first and foremost an alternative regional integration model that, in distinction to the institutions-based approach

epitomized by the treaties of the European Union, lacks supranational bodies and is instead based on often more precise and enforceable rules and that States chose the NAFTA model amongst existing templates because it best fit their interests (Dür, Baccini, and Haftel 2015). Finally, Puig conjectures that the diffusion of NAFTA treaty design is driven partly by bureaucratic inertia, at least in the case of Mexico, as negotiators consider it convenient to “specialize” in a single treaty model that is subsequently replicated in negotiations with third states (Puig 2013).

The question whether any of these explanations convincingly describes NAFTA's diffusion can be approached empirically. To trace NAFTA's diffusion, we look for treaties in our dataset that overlap textually to at least 55% with NAFTA in terms of Jaccard similarity.²² We find 50 PTAs in our dataset that meet this criterion and order them by date of signature to create a diffusion chronology. What we find is striking: (1) the United States of America played no direct role in NAFTA's treaty design diffusion in the first ten years of its existence and (2) the most important promoters of NAFTA treaty design were non-NAFTA countries through South-South agreements.

The chronology suggests that NAFTA was diffused in four waves, which we display in Figure 9. NAFTA treaty design first proliferated abroad from the mid 1990s to 1998 through Mexico and Canada as original NAFTA parties. Mexico signed FTAs with Costa Rica and Colombia (1994), Nicaragua (1996) and Chile (1998), all of which closely followed NAFTA treaty design with a textual similarity between 63% and 72%. In the same period, Canada only signed one FTA modelled on NAFTA with Chile in 1996 (68% similarity), while the United States of America signed no PTAs at all.

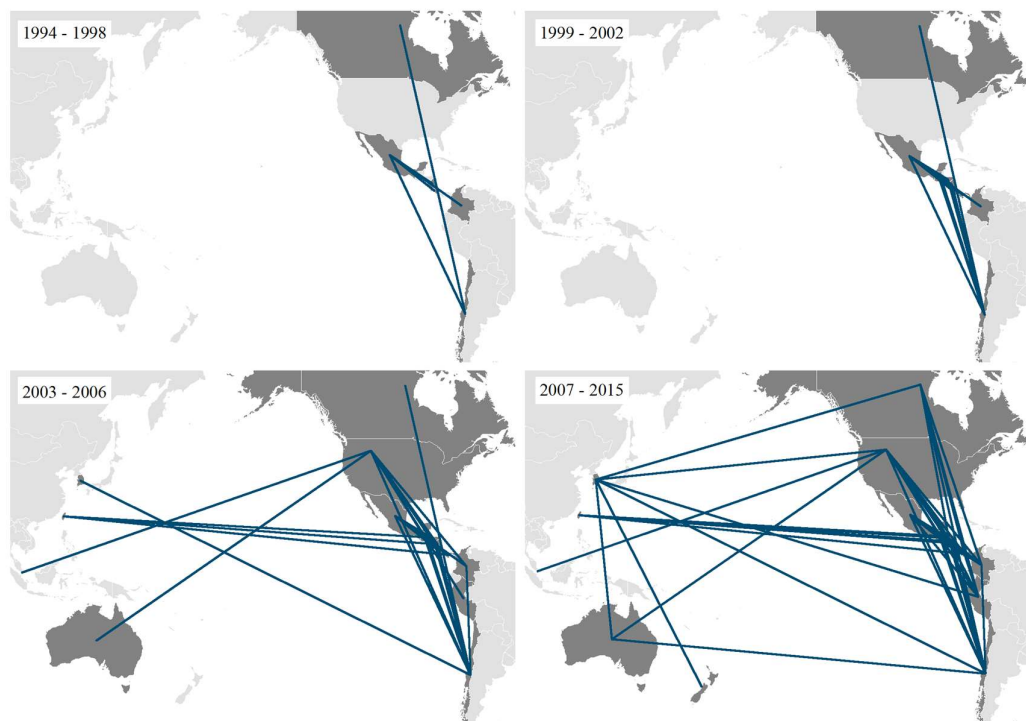
The second wave saw a continuation of that practice with Mexico concluding new PTAs with Honduras, Guatemala and El Salvador (2000). More importantly, however, the Latin American partner states of the first wave started to sign South-South FTAs modelled on NAFTA. These agreements not only exposed new Latin American states to NAFTA treaty design, such as Panama through its FTA with Central American countries in 2002, but also entrenched the NAFTA model further with Mexico's first wave partners. At the end of 2002, NAFTA-type PTAs connected most Latin American states bordering the Pacific.

The third wave starting in 2003 brought the NAFTA model to Asia with the Panama–Taiwan Province of China FTA and the Chile–Republic of Korea FTA, both signed in 2003 and displaying a 60% treaty similarity to NAFTA. That third wave also marked the return of the United States of America as promoter of NAFTA-inspired treaty design concluding a series of agreements in the mid-2000s starting with Singapore in 2003 (58% textual similarity with NAFTA).

The final wave from 2007 onwards saw East Asian states signing NAFTA inspired deals amongst themselves. Specifically Taiwan Province of China and the Republic of Korea were active in signing PTAs modelled on NAFTA with other states in South East Asia. At the same time, new treaties between North and South America as well as between the Americas and South East Asia fully established NAFTA-inspired treaty design in the Pacific Rim region paving the ground for the later Transpacific Partnership, which overlaps with NAFTA to 58%.

These findings force us to reconsider prevailing explanations relating to NAFTA design diffusion. First, attributing the proliferation of NAFTA elements to US' coercion is unconvincing, since the United States of America played little part in the initial diffusion of NAFTA to South America. Certainly, hegemony may have cast a shadow, as talks on a Free Trade Area of the Americas involving the United States of America were ongoing in the early 2000s. Moreover, Latin American partners such as Chile, that had earlier even been invited to join NAFTA, had a strong prospect of getting a future deal with the United States of America (Gantz 2003). Yet, the absence of active leadership of the United States of America (including for domestic reasons, as “Fast Track” authority for PTA negotiations was only granted in 2002) weighs against claims of a coercive “implantation” of treaty design.

²² We impose this threshold in order to ensure that textual similarity is more than incidental.

Figure 9. Diffusion of NAFTA treaty design

Note: This figure shows the diffusion of NAFTA treaty design. A link is established between two countries if they sign a PTA that overlaps textually with NAFTA to at least 55%. Once a country has signed a NAFTA-inspired treaty, we color-code the country grey.²³

Second, interest-based considerations are equally unlikely to explain the enthusiastic endorsement of NAFTA by South American states. Copy-and-pasting NAFTA was arguably not the optimal design choice to further economic integration amongst developing countries. For instance, while NAFTA's chapter on intellectual property rights may have well reflected interests of the United States of America of protecting right holders in jurisdiction with weak IP enforcement, it ran counter the policy of many South American states, partly pursued at the WTO, to privilege access to medicine over monopoly rights (Correa 2006). It is thus not clear why Latin American states would want to transpose a North-South Agreement such as NAFTA to govern their South-South relations. Yet a significant part of the diffusion of NAFTA design in the Americas and in East Asia occurred through non-NAFTA parties signing South-South agreements.

If power and state interests cannot convincingly account for the diffusion patterns we observe, alternative explanations are worth pursuing. Similar to Poulsen's account of states stumbling into BITs without appreciating their full costs (L. Poulsen 2013; L. N. S. Poulsen 2015), and Puig's argument of bureaucratic inertia, Latin American states may have endorsed NAFTA treaty design out of bounded rationality even if it was not optimally tailored to their needs. Path dependency, default rules and emulation may thus help explain why states, such as Panama, not having signed any PTA with one of the NAFTA countries directly, became a champion of its design nevertheless and helped diffuse it across the Pacific. Of course, further investigation is necessary to evaluate these points further. Our example is merely supposed to illustrate how a more fine-grained representation of treaty design can challenge prevailing accounts of NAFTA design diffusion and open the door for the testing of alternative theories.

²³ We gratefully acknowledge Friederike Alschner's work in preparing this visualization.

5. Trade Law - Tracking Legal Convergence

Finally, turning to the study of the trade regime from a legal perspective, one of the most important debates amongst legal scholars concerns international trade law's interaction with other areas of international law (Pauwelyn 2004; Simma and Pulkowski 2006). Do fields of international law converge or are we facing a set of increasingly fragmented self-contained regimes? Text-as-data analysis of PTAs has important empirical insights to offer on that front.

5.1 PTAs and the fragmentation debate

PTAs are in many ways in the midst of international law's fragmentation debate. By definition, they pose a risk of fragmenting the body of international trade law that is otherwise multilateralized through the WTO (Kwak and Marceau 2003). At the same time, PTAs are an important source of convergence, as they often encompass chapters relating to investment, competition, the environment or labour and thereby explicitly connect the trade regime with its neighbouring fields of law (Housman and Orbuch 1992). In fact, while the WTO's rulemaking arm stagnates, PTAs have become a laboratory for how trade law interacts with other areas.

The convergence or divergence of trade and investment rules is one aspect of that larger conversation. While much of that debate is normative in nature (Antoni and Ewing-Chow 2013; Alford 2014; Puig 2015), scholars have begun to approach it empirically. Trade and investment law interacts in litigation through cross-citations (Charlotin 2017), judicial borrowing (Alford 2014) or parallel claims (Puig 2016; for a different example, see Antoni and Ewing-Chow 2013), but also through rulemaking, as PTAs increasingly contain investment chapters. Using our similarity metric, we focus on the latter interaction and investigate whether PTA investment chapters normatively converge with or diverge from the rules found in bilateral investment treaties.

5.2 Tracing convergence between BITs and PTAs investment chapters

Commentators disagree on whether there are systematic, normative differences between BITs and PTA investment chapters. For Alford, BITs, not unlike the WTO, set a baseline of protection and liberalization commitments, whereas PTAs, by regulating value chains in a more integrated manner, go beyond that minimum standard and foster deeper economic integration (Alford 2014: 12). Baetens, on the other hand, notes that the content of modern BITs is often virtually identical to that of investment chapters while acknowledging that the inclusion of investment provisions in a broader PTA may affect the interpretation of these clauses (Baetens 2013). Finally, Kotschwar argues that there is greater variation in the content of investment chapters than in BITs, as countries tailor their investment provisions more carefully to their negotiating partners in PTAs (Kotschwar 2009, 374). As an empirical matter, it remains unclear whether BITs and PTAs differ systematically in content.

When comparing BITs and investment chapters in PTAs, two caveats have to be borne in mind. First, investment disciplines in PTAs are often not limited to the investment chapter. Core provisions commonly found in BITs, such as a mechanism for inter-state dispute settlement, special rules on taxation, or security and public policy exceptions are typically relegated to their own chapters in PTAs.²⁴ Such structural differences will slightly reduce the degree of similarity between PTA investment chapters and BITs.²⁵ Second, BITs have been signed from the late 1950s onwards, whereas PTAs with fully-fledged investment chapters have only been around for the past 30 years. Hence, as a group, PTA investment chapters necessarily embody a more modern approach to investment policy-making than BITs.²⁶

Our dataset contains 80 PTA investment chapters, which are displayed in Figure 10 as a heat map ordered by treaty similarity. The figure reveals three different groups of treaties. A first group comprising the majority of


²⁴ Investment-related considerations further afield such as investment in intellectual property rights or in services typically also benefit from a self-standing chapter.

²⁵ One way to compensate this effect is to compare all investment-related provisions in PTAs to BITs rather than only PTA investment chapters.

²⁶ Some states, such as Chile or New Zealand, even departed entirely from BITs and instead focus on protecting investors exclusively through dedicated chapters in PTAs.

agreements consists of a relatively homogenous set of treaties, signed by the United States of America, Canada, Japan, Australia and several other East Asian and Latin American countries. On closer inspection, these PTA chapters strongly resemble the structure of BITs: they all contain provisions protecting investors against discrimination, unfair treatment or expropriation and, with few exceptions, e.g. the Australia- United States of America investment chapter, give them the right to initiate investor-state arbitration claims. A second group of investment chapters belongs to PTAs signed by EFTA countries that contain a more limited set of investment provisions. These chapters provide national treatment, free transfer of funds and entry of investment-related personnel subject to reservations and exceptions, but no investor-state arbitration.²⁷ The third group of chapters have little in common with each other and comprise a range of framework agreements launching negotiations on future investment disciplines,²⁸ investment promotion goals,²⁹ references that incorporate or acknowledge overlapping BITs,³⁰ or rudimentary protection provisions in pre-1960 Friendship, Navigation and Commerce treaties.³¹

Figure 10. Heat maps of 5-*q*-gram Jaccard distances between English-language Investment Chapters in PTAs



Note: This figure shows a heat map of pairwise Jaccard distances between PTA investment chapters. Treaties are ordered by treaty similarity.

Shifting to a comparison between investment chapters of the first group and BITs, do states pursue the same treaty design in both instances? The United States of America, whose treaty network is displayed in Figure 11, started with a temporary misalignment between its BIT and PTA policy on investment after NAFTA.³² The 2004 model BIT then converged practice to the point that BITs and PTAs of the United States of America reach near identity in treaty design terms. Japan exhibits a similarly close alignment. The country changed treaty design in 2002 and concluded relatively similar BITs and PTA investment chapters ever since, albeit not as similar as

²⁷ It should be noted that this excludes the PTAs EFTA signed with Singapore and Ukraine which contain investor-state arbitration provisions and which our similarity algorithm clusters in the first group.

²⁸ See e.g. ASEAN-Japan PTA Chapter 7.

²⁹ See e.g. EFTA-Tunisia PTA Chapter IV.

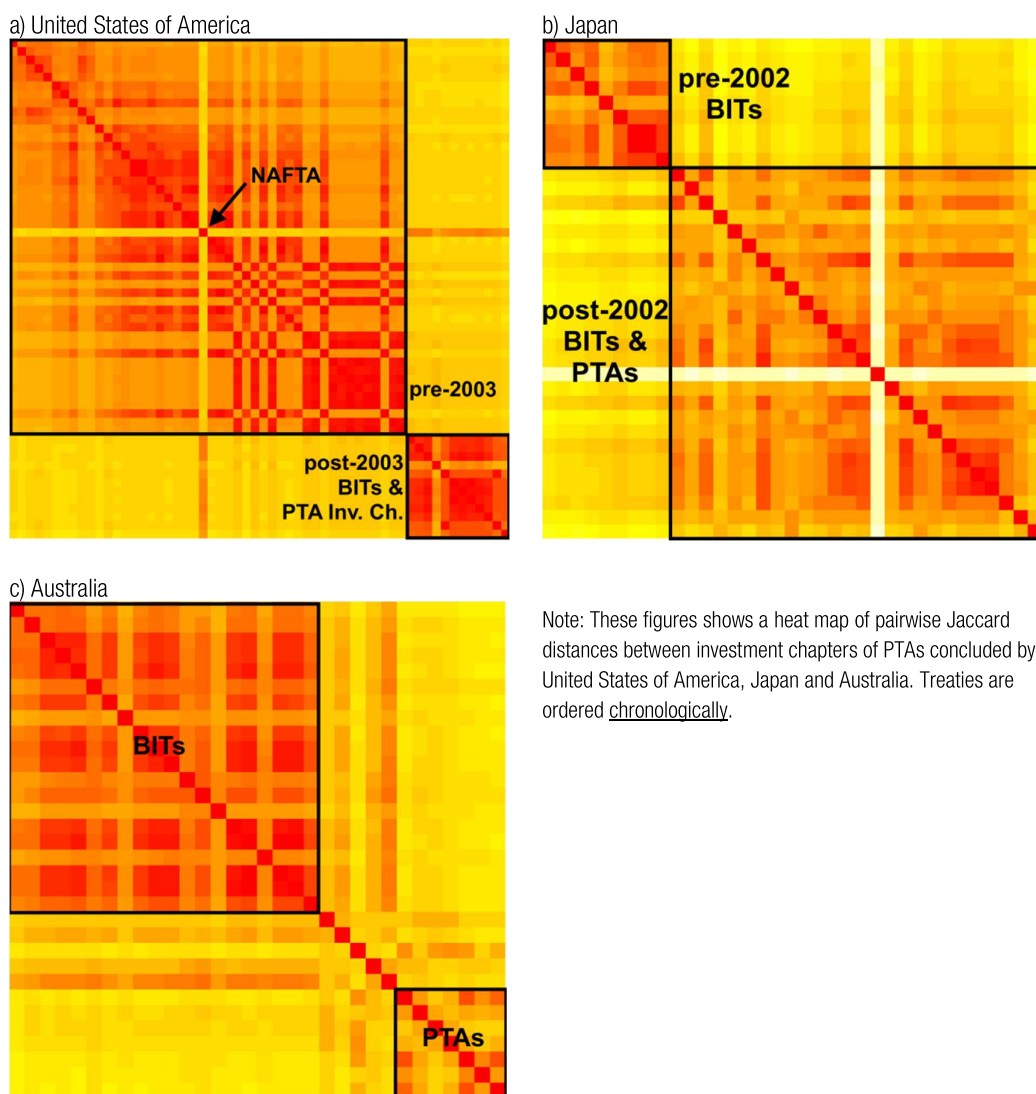
³⁰ See e.g. China-Iceland PTA Chapter 8.

³¹ See e.g. El Salvador-Nicaragua PTA Chapter VI.

³² (On this point, see Alschner 2017)

those of the US. In other countries, the emergence of PTA investment chapters has been a source of divergence. In Australia, the 2004 FTA with Singapore marked the end of what was hitherto a relatively consistent BIT network. Thereafter, Australia continued to conclude new BITs that followed earlier BITs in design, while simultaneously diverging from that design in new FTAs, which as a group were more heterogeneous than the country's prior BITs.

Figure 11. Heat maps of 5-gram Jaccard distances between English-language Investment Chapters in PTAs and BITs



This brief comparison yields several crucial insights on the impact of investment chapters on normative convergence. First, aside from minor structural differences, there is no reason to a priori differentiate between BITs and PTAs investment chapters in terms of scope or content. As the case of the United States of America shows, both types of treaties can converge almost perfectly around the same norms. Second, the emergence of PTA investment chapters has given rise to more divergence in the treaty networks of countries like Australia – between BITs and PTAs and even amongst PTAs.

The reason for this divergence is likely not to be the difference between agreement types but between agreement partners. BITs were primarily concluded between developed countries as rule-makers and developing countries as rule-takers, generally resulting in consistent treaty networks for the former and

inconsistent treaty networks for the latter (Alschner and Skougarevskiy 2016b). PTAs escape this North-South logic and are also concluded amongst developed countries. While the world's superpower can still achieve consistent treaties in that setting, less powerful states cannot. As relative bargaining power wanes, divergence within national treaty networks grows, as we saw in the case of Japan and, more strikingly, Australia. Hence, while showing evidence for convergence between investment chapters and BITs, our analysis also points to an increased fragmentation of investment rules as PTA investment chapters diverge from one bargain to another and thus add more variety and complexity to investment rules.

Conclusion

Using recent advances in text-as-data techniques on a newly created structured text corpus of treaties from the WTO RTA dataset, this article traced textual similarity across PTAs to render latent treaty design patterns visible. On the global level, it identified that PTAs are divided along regional or inter-regional clusters. On the national level, it highlighted that states do not follow a one-size fits all approach when designing PTAs but rather adjust their agreements to their negotiation partner while preserving a certain degree of standardization. On the chapter-level, it detected an important degree of variation across chapters meriting a study of PTAs on the sub-treaty level. Finally, the article went beyond describing PTA design by showcasing how textual similarity can be integrated into broader inferential or normative research in trade economics (e.g. measuring the impact of PTA design on trade flows), politics (e.g. tracing and explaining PTA design diffusion) and law (e.g. investigating normative convergence between the trade and investment regime). In terms of future work, we plan to add new PTAs to the dataset, explore similarity patterns further at the chapter and article level and apply additional text-as-data analytics to our treaty corpus.

References

- Aaken, A. van. 2009. "International Investment Law Between Commitment and Flexibility: A Contract Theory Analysis." *Journal of International Economic Law* 12 (2): 507–38. doi:10.1093/jiel/jgp022.
- Alford, Roger P. 2014. "The Convergence of International Trade and Investment Arbitration." *Santa Clara Journal of International Law* 12 (1): 35.
- Allee, Todd, Manfred Elsig, and Andrew Lugg. 2017a. "Is the European Union Trade Deal with Canada New or Recycled? A Text-as-Data Approach." *Global Policy*, March, n/a-n/a. doi:10.1111/1758-5899.12420.
- . 2017b. "The Ties between the World Trade Organization and Preferential Trade Agreements: A Textual Analysis." *Journal of International Economic Law* 20 (2): 333–63. doi:10.1093/jiel/jgx009.
- Allee, Todd, and Andrew Lugg. 2016. "Who Wrote the Rules for the Trans-Pacific Partnership?" *Research & Politics* 3 (3): 2053168016658919. doi:10.1177/2053168016658919.
- Allee, Todd, and Clint Peinhardt. 2010. "Delegating Differences: Bilateral Investment Treaties and Bargaining Over Dispute Resolution Provisions." *International Studies Quarterly* 54 (1): 1–26. doi:10.1111/j.1468-2478.2009.00575.x.
- . 2014. "Evaluating Three Explanations for the Design of Bilateral Investment Treaties." *World Politics* 66 (01): 47–87. doi:10.1017/S0043887113000324.
- Alschner, Wolfgang. 2017. "The Impact of Investment Arbitration on Investment Treaty Design: Myth Versus Reality." *Yale Journal of International Law* 42 (1).
- Alschner, Wolfgang, Joost Pauwelyn, and Sergio Puig. 2017. "The Data-Driven Future of International Economic Law." *Journal of International Economic Law* 20 (2): 217–31. doi:10.1093/jiel/jgx020.
- Alschner, Wolfgang, Julia Seiermann, and Dmitriy Skougarevskiy. 2017. "The Impact of the TPP on Trade Between Member Countries—A Text-as-Data Approach." <https://www.adb.org/sites/default/files/publication/321776/adb-wp745.pdf>.
- Alschner, Wolfgang, and Dmitriy Skougarevskiy. 2015. "Treaty Texts as Data—Developing New Tools for Negotiators and Litigators to Compare Bilateral Investment Treaties." In *Legal Knowledge and Information Systems: JURIX 2015: The Twenty-Eighth Annual Conference*, edited by Antonino Rotolo, 141–44. IOS Press.
- . 2016a. "Can Robots Write Treaties? Using Recurrent Neural Networks to Draft International Investment Agreements." In *Legal Knowledge and Information Systems: JURIX 2016*, edited by Floris Bex, 119–24. IOS Press.
- . 2016b. "Mapping the Universe of International Investment Agreements." *Journal of International Economic Law*, CTEI Working Papers, 19 (8).
- . 2016c. "The New Gold Standard? Empirically Situating the Trans-Pacific Partnership in the Investment Treaty Universe." *Journal of World Investment and Trade* 17 (3).
- Antoni, A., and Micheal Ewing-Chow. 2013. "Trade and Investment Convergence and Divergence: Revisiting the North American Sugar War." *Latin American J. Int'l Trade L.* 1: 315.
- Baccini, Leonardo, Pablo M. Pinto, and Stephen Weymouth. 2017. "The Distributional Consequences of Preferential Trade Liberalization: Firm-Level Evidence." *International Organization*, April, 1–23. doi:10.1017/S002081831700011X.
- Baetens, Freya. 2013. "Preferential Trade and Investment Agreements and the Trade/Investment Divide: Is the Whole More than the Sum of Its Parts?" In *Preferential Trade and Investment Agreements: A New Ordering Paradigm for*

International Investment Relations?, edited by Rainer Hofmann, Christian Tams, and Stephan W. Schill. Frankfurt University Press.

Baier, Scott L., Jeffrey H. Bergstrand, and Matthew W. Clance. 2015. "Heterogeneous Economic Integration Agreement Effects." CESifo Working Paper Series 5488. CESifo Group Munich. https://ideas.repec.org/p/ces/ceswps/_5488.html.

Baird, Douglas G. 2006. "The Boilerplate Puzzle." *Michigan Law Review* 104 (5): 933–952.

Berge, Tarald Laudal, and Øyvind Stiansen. 2016. "Negotiating BITs with Models. The Power of Expertise." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2851454.

Berger, Axel. 2013. "Investment Rules in Chinese PTIAs - A Partial 'NAFTA-ization.'" SSRN Scholarly Paper ID 2171765. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=2171765>.

Berry, David M., ed. 2012. *Understanding Digital Humanities*. London: Palgrave Macmillan UK. doi:10.1057/9780230371934.

Charlotin, Damien. 2017. "The Place of Investment Awards and WTO Decisions in International Law: A Citation Analysis." *Journal of International Economic Law* 20 (2): 279–99. doi:10.1093/jiel/jgx010.

Correa, Carlos María. 2006. "Implications of Bilateral Free Trade Agreements on Access to Medicines." *Bulletin of the World Health Organization* 84 (5): 399–404.

Deng, Shengliang, and Katherine Braun. 1997. "China - NAFTA Trade: What Does the Future Hold?" In *China and the Asia Pacific Economy*, edited by C. H. Chai, Y. Y. Kueh, and Clement Allan Tisdell, 329. Nova Publishers.

Dobbin, Frank, Beth Simmons, and Geoffrey Garrett. 2007. "The Global Diffusion of Public Policies: Social Construction, Coercion, Competition, or Learning?" *Annual Review of Sociology* 33 (1): 449–72. doi:10.1146/annurev.soc.33.090106.142507.

Dür, Andreas, Leonardo Baccini, and Manfred Elsig. 2014. "The Design of International Trade Agreements: Introducing a New Dataset." *The Review of International Organizations* 9 (3): 353–375.

Dür, Andreas, Leonardo Baccini, and Yoram Z. Haftel. 2015. "Imitation and Innovation in International Governance: The Diffusion of Trade Agreement Design." In *Trade Cooperation*, edited by Andreas Dür and Manfred Elsig. Cambridge: Cambridge University Press. doi:10.1017/CBO9781316018453.008.

Egger, Peter H., and Filip Tarlea. 2017. "Comparing Apples to Apples: Estimating Consistent Partial Effects of Preferential Economic Integration Agreements." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2934206.

Egger, Peter, and S. Nihai. 2015. "Effects of Deep versus Shallow Trade Agreements in General Equilibrium." In *Trade Cooperation*, edited by Andreas Dür and Manfred Elsig, 374–91. Cambridge: Cambridge University Press. doi:10.1017/CBO9781316018453.024.

Estevadeordal, Antoni, Kati Suominen, Pablo Sanguinetti, and Alberto Trejos. 2005. "Rules of Origin in Preferential Trading Arrangements: Is All Well with the Spaghetti Bowl in the Americas?" *Economía* 5 (2): 63–103.

Fecák, Tomáš. 2011. "Czech Experience with Bilateral Investment Treaties: Somewhat Bitter Taste of Investment Protection." *Czech Yearbook of International Law* 2: 233–67.

Felbermayr, Gabriel, Rahel Aichele, and Inga Heiland. 2016. "Going Deep: The Trade and Welfare Effects of TTIP Revised." Ifo Working Paper Series Ifo Working Paper No. 219. Ifo Institute - Leibniz Institute for Economic Research at the University of Munich. http://econpapers.repec.org/paper/cesifowps/_5f219.htm.

Gantz, David A. 2003. "The Evolution of FTA Investment Provisions: From NAFTA to the United States-Chile Free Trade Agreement." *Am. U. Int'l L. Rev.* 19: 679.

- Gentzkow, Matthew, Bryan T. Kelly, and Matt Taddy. 2017. "Text as Data." National Bureau of Economic Research. <http://www.nber.org/papers/w23276.pdf>.
- Gilardi, Fabrizio. 2012. "Transnational Diffusion: Norms, Ideas, and Policies." *Handbook of International Relations* 2: 453–477.
- Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis*. doi:10.1093/pan/mps028.
- Grugel, Jean B. 2004. "New Regionalism and Modes of Governance — Comparing US and EU Strategies in Latin America." *European Journal of International Relations* 10 (4): 603–26. doi:10.1177/1354066104047850.
- Horn, Henrik, Petros C. Mavroidis, and André Sapir. 2010. "Beyond the WTO? An Anatomy of EU and US Preferential Trade Agreements." *The World Economy* 33 (11): 1565–1588. doi:10.1111/j.1467-9701.2010.01273.x.
- Hornbeck, J. F., and William H. Cooper. 2011. *Trade Promotion Authority (TPA) and the Role of Congress in Trade Policy*. Congressional Research Service.
- Housman, Robert F., and Paul M. Orbuch. 1992. "Integrating Labor and Environmental Concerns into the North American Free Trade Agreement: A Look Back and a Look Ahead." *American University Journal of International Law and Policy* 8: 719.
- Koremenos, Barbara, Charles Lipson, and Duncan Snidal. 2001. "The Rational Design of International Institutions." *International Organization* 55 (04): 761–99. doi:10.1162/002081801317193592.
- Kotschwar, Kotschwar. 2009. "Mapping Investment Provisions in Regional Trade Agreements: Towards an International Investment Regime?" In *Regional Rules in the Global Trading System*, edited by Antoni Esteveadeordal, Kati Suominen, and Robert Teh. Cambridge University Press.
- Kwak, Kyung, and Gabrielle Marceau. 2003. "Overlaps and Conflicts of Jurisdiction between the World Trade Organization and Regional Trade Agreements." *Canadian Yearbook of International Law* 41: 83.
- Larch, Mario, Joschka Wanner, Yoto Yotov, and Thomas Zylkin. 2017. "The Currency Union Effect: A PPML Re-Assessment with High-Dimensional Fixed Effects." 2017–7. School of Economics Working Paper Series. LeBow College of Business, Drexel University. https://ideas.repec.org/p/ris/drxlwp/2017_007.html.
- Lavranos, Nikos. 2013. "The New EU Investment Treaties: Convergence towards the NAFTA Model as the New Plurilateral Model BIT Text?" SSRN Scholarly Paper ID 2241455. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=2241455>.
- Lazer, David, Alex (Sandy) Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, et al. 2009. "Life in the Network: The Coming Age of Computational Social Science." *Science (New York, N.Y.)* 323 (5915): 721–23. doi:10.1126/science.1167742.
- Manger, Mark S., and Clint Peinhardt. 2017. "Learning and the Precision of International Investment Agreements." *International Interactions* 0 (0): 1–21. doi:10.1080/03050629.2017.1311258.
- Pauwelyn, Joost. 2004. *Conflict of Norms in Public International Law: How WTO Law Relates to Other Rules of International Law*. Repr. 2004. Cambridge Studies in International and Comparative Law. Cambridge [etc.]: Cambridge University Press.
- Pauwelyn, Joost, and Wolfgang Alschner. 2015. "Forget About the WTO: The Network of Relations between Preferential Trade Agreements (PTAs) and 'Double PTAs'." In *Trade Cooperation*, edited by Andreas Dür and Manfred Elsig. Cambridge: Cambridge University Press. doi:10.1017/CBO9781316018453.024.
- Poulsen, Lauge. 2013. "Bounded Rationality and the Diffusion of Modern Investment Treaties." *International Studies Quarterly*. doi:10.1111/isqu.12051.

Poulsen, Lauge N. Skovgaard. 2015. *Bounded Rationality and Economic Diplomacy: The Politics of Investment Treaties in Developing Countries*. Cambridge University Press.

Price, Daniel M. 1993. "Overview of the NAFTA Investment Chapter: Substantive Rules and Investor-State Dispute Settlement, An." *The International Lawyer* 27: 727.

Puig, Sergio. 2013. "Does Bureaucratic Inertia Matter in Treaty Bargaining-Or, toward a Greater Use of Qualitative Data in Empirical Legal Inquiries." *Santa Clara J. Int'l L.* 12: 317.

———. 2015. "The Merging of International Trade and Investment Law." *Berkeley Journal of International Law* 33: 1.

———. 2016. "Tobacco Litigation in International Courts." *Harvard International Law Journal* 57 (2): 383–432.

Ruta, Michele, Claudia Hofmann, and Alberto Osnago. 2017. "Horizontal Depth : A New Database on the Content of Preferential Trade Agreements." WPS7981. The World Bank.
<http://documents.worldbank.org/curated/en/700771487791538589/Horizontal-depth-a-new-database-on-the-content-of-preferential-trade-agreements>.

Simma, Bruno, and Dirk Pulkowski. 2006. "Of Planets and the Universe: Self-Contained Regimes in International Law." *European Journal of International Law* 17 (3): 483–529. doi:10.1093/ejil/chl015.

Spirling, Arthur. 2012. "U.S. Treaty Making with American Indians: Institutional Change and Relative Power, 1784–1911." *American Journal of Political Science* 56 (1): 84–97. doi:10.1111/j.1540-5907.2011.00558.x.

Ukkonen, Esko. 1992. "Approximate String-Matching with q-Grams and Maximal Matches." *Theoretical Computer Science* 92 (1): 191–211. doi:10.1016/0304-3975(92)90143-4.

Van der Loo, Mark PJ. 2014. "The Stringdist Package for Approximate String Matching." *The R*, 2.

Vandeveld, Kenneth J. 2009. *U.S. International Investment Agreements*. Oxford: Oxford University Press.

WTO. 2011. *World Trade Report 2011: The WTO and Preferential Trade Agreements: From Co-Existence to Coherence*. Geneva: WTO. https://www.wto.org/english/res_e/booksp_e/anrep_e/world_trade_report11_e.pdf.
