

The Data-Driven Future of International Economic Law

Wolfgang Alschner*, Joost Pauwelyn** and Sergio Puig***

ABSTRACT

The availability of more data and new ways of analyzing it is changing the way we do empirical legal research. With the help of modern technology we can study adjudicators, awards and agreements in greater numbers, less time and more detail opening the doors for new research questions, theory building and legal technology applications for scholars and practitioners. This introduction to the *Journal of International Economic Law* Special Issue on new frontiers in empirical legal research provides a first take on this data-driven future. It distinguishes data-driven research from more traditional methods by pointing to (1) its “data first” attitude, (2) its ambition to look at all the available data rather than subsamples thereof and (3) its focus on computing rather than reading or counting. Data-driven research comes with new promises, but also challenges and limitations. While it allows researchers to uncover latent structures, debunk past myths and even forecast the future, it also requires new skills and competencies including an ability to tell patterns from noise in inductive data analysis. We argue that the time is ripe to overcome these challenges and to seize the opportunities of the new data-driven frontier in empirical legal scholarship.

I. INTRODUCTION

Since the 1980s, legal academia has slowly embraced empirical legal research.¹ Empirical legal studies are now common in international law scholarship too and there is a growing trend of empiricism in this *Journal's* sub-field of international economic law.² Its economic nature, the wealth of treaties and cases and the

* Assistant Professor, Common Law Section, University of Ottawa. Email: wolfgang.alschner@uOttawa.ca.

** Professor, Graduate Institute of International and Development Studies, Geneva and Murase Visiting Professor, Georgetown University Law Center. Email: joost.pauwelyn@graduateinstitute.ch.

*** Associate Professor and Director, International Economic Law and Policy Program, James E. Rogers College of Law, University of Arizona. Email: spuig@email.arizona.edu.

1 See generally Shari Seidman Diamond and Pam Mueller, ‘Empirical Legal Scholarship in Law Reviews’, 6 *Annual Review of Law and Social Science* 581 (2010), at 589. Lee Epstein and Gary King, ‘The Rules of Inference’, 69 *University of Chicago Law Review* 1 (2002) (providing a critique of empirical legal research).

2 See generally Gregory Shaffer and Tom Ginsburg, ‘The Empirical Turn in International Legal Scholarship’, 106 *American Journal International Law* 1 (2012) (discussing the growth of empirical research of international law); see also Emilie M. Hafner-Burton, David G. Victor and Yonatan Lupu, ‘Political Science Research on International Law: The State of the Field’, 106 *American Journal of International Law* 47 (2012) (providing a review of political science research relevant to international law); Beth Simmons, ‘Treaty Compliance and Violation’, 13 *Annual Review of Political Science* 273 (2010) (reviewing empirical

existence of compulsory dispute settlement make the international law of trade, investment, taxation, finance, and intellectual property fruitful areas for empirical research with both conceptual and practical applications.

Changes in information technology are now beginning to open a new frontier in that empirical international economic law scholarship. The Internet provides scholars with unprecedented access to data on treaties, cases, and adjudicators collected by states, researchers, and international organizations. New tools for their analysis are emerging too. The exponential growth of computing power has fuelled advances in data analytics, artificial intelligence, and machine learning allowing small teams of researchers to analyze massive amounts of information in little time and at little cost. As data grows big and the means for its analysis turn computational, we are entering the data-driven future of empirical international economic law scholarship.

This Special Issue is a first attempt at exploring this new frontier. Generously funded by the Swiss National Science Foundation (SNF) in the context of the Project ‘Convergence versus Divergence? Text-as-data and Network Analysis of International Economic Law Treaties and Tribunals’, we convened an exploratory workshop at the Graduate Institute in Geneva in May 2016 to brainstorm on three related challenges of data-driven research in international economic law: (i) data: its collection, availability, sharing, and related best practices, (ii) methods: their diversity and relative effectiveness to tackle specific questions and how to most efficiently conduct interdisciplinary research, and (iii) applications: that is, how to demonstrate and showcase the practical usefulness of new empirical methods to both academics and practitioners in the field.

In both the SNF Project and this Special Issue we place particular emphasis on two components. First, we explore ‘big data’—large amounts of legal information difficult or impossible to investigate using traditional means of analysis. Some of the contributions in this Issue provide a taste of the scale of the research to come. Charlotin, for instance, analyzes fragmentation and judicial dialogue through over 75,000 cross-citations by international tribunals. Other contributors use new datasets to challenge conventional views of how hundreds of preferential trade agreements (PTAs) relate to the WTO (Elsig et al.) or to map the professional roles of thousands of arbitrators and counsels in investment arbitration (Langford et al.). The second component relates to two novel techniques uniquely suited to make sense of ‘big data’: (i) text-as-data approaches that draw from computational linguistics to assess and compare texts;³ and (ii) network analysis, which investigates connections (ties) between actors or documents (nodes).⁴

We are thrilled to introduce seven contributions to this Special Issue that, each in their own fashion, showcase the data-driven future of international economic law. Four contributions relate to international trade (Derlén et al.; Elsig et al.; Pelc et al.; and Morin et al.); two contributions address international investment (Broude et al.;

research on compliance). Of course, it is important to note that during this time there has also been a rise in the use of empirical legal research more generally.

3 Justin Grimmer and Brandon M. Stewart, ‘Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts’, 21(3) *Political Analysis* 267 (2013).

4 Matthew O. Jackson, *Social and Economic Networks* (Princeton: Princeton University Press, 2008); Mark E. J. Newman, *Networks: An Introduction* (Oxford: Oxford University Press, 2010).

Langford et al.); one article covers both trade and investment jurisprudence and its interaction with other international tribunals (Charlotin). In terms of methods, two contributions focus on text-as-data methods using computer programs to identify and compare treaty or jurisprudential texts (Elsig et al.; Pelc et al.); four articles use network analysis techniques (Charlotin; Derlén et al.; Langford et al.; Morin et al.). Other contributions utilize more traditional human coding of data and/or more conventional statistical methods such as regression analysis to explain variation (Broude et al.; Morin et al.).

Our goal with this Special Issue is to both introduce these new data and methodological research opportunities to legal scholars and, hopefully, convince the readership of this *Journal* that the data-driven future of international economic law offers meaningful applications and insights of value also to legal practitioners. In the remainder of this introduction, we will present the main characteristics of the new data-driven frontier in empirical international economic law research. Wherever possible, we will refer to the contributions in this Special Issue and outside research to highlight its promises, challenges, and limitations. Furthermore, a list of databases and an illustrative bibliography can be found in the Annex to help readers navigate this emerging field.

II. EMPIRICAL LEGAL RESEARCH AND INTERNATIONAL ECONOMIC LAW

So far two strands of research dominate empirical investigations in international economic law: one examines the *impact* of treaties or judicial decisions on behavior; the other looks at what *causes* the creation of and/or variation between treaties or tribunal outcomes. The first strand consists of studies that use international law as an independent variable (i.e. the explaining factor). It looks at whether certain economic treaties caused an increase in trade, finance, or foreign investment volumes or whether certain commitments (or decisions) lead to a change in behavior of governments (e.g. cooperation) or other commercial actors (e.g. corruption).⁵ In a second strand scholars have used international law as dependent variable (i.e. the factor to be explained). This type of research attempts to elucidate, for instance, why states sign international agreements in the first place or asks what may explain the seemingly large variation in the complex web of international treaties. More recently, studies have also looked at international dispute settlement decisions as dependent variables seizing on the expansion of international courts and tribunals and with it, the massive growth of enforcement actions brought by states or private actors.⁶

To some extent, the traditional (and narrower) expression of empirical legal studies can be traced to the academic disciplines and methodological approaches brought by scholars trained in disciplines beyond law who pioneered this scholarship in international law. For instance, informed by ideas of economic and political power and influence,

5 See Shaffer and Ginsburg, above n 2, at 1 (suggestion that ‘the conditions under which international law is formed and has effects’ is the most relevant empirical line of investigation).

6 See, e.g. Christina L. Davis, *Why Adjudicate? Enforcing Trade Rules in the WTO* (Princeton, Princeton University Press, 2012); Susan D. Franck, ‘Development and Outcomes of Investment Treaty Arbitration’, 50 *Harvard International Law Journal* 435 (2009); Suha Jubran Ballan, ‘Investment Treaty Arbitration and Institutional Backgrounds: An Empirical Study’, 34 *Wisconsin International Law Journal* 31 (2016).

political scientists and international relations scholars have tested hypotheses to explain the conditions under which power matters. Conversely, informed by ideas of public choice, microeconomics and more recently behavioral economics, economists, sociologists, and psychologists have spearheaded research that looks at how incentives matter for decision-makers in international courts or governments.

With all the benefits that interdisciplinary approaches have brought to the study of international law, including the deployment of new methodologies and the use of data for the interrogation of long-held, and oft-uncontested assumptions, the success of particular approaches is not without problems. For one, empirical research can be blind to its own limitations. For example, some approaches rely on the construction of explanatory variables. These variables are useful and necessary for regression analysis, but may hide the complexity of what is often referred to as the *law in action* (i.e. how rules arise and operate in the real world).⁷

More importantly for our purposes is how mainstream empirical legal scholarship has largely neglected other, more recent research strategies that have emerged in neighboring fields under diverse labels such as ‘digital humanities’,⁸ ‘computational social science’⁹ or the study of ‘big data’¹⁰. We describe these new research strategies here in this Special Issue as ‘data-driven’ means of analysis since they have resulted from unprecedented data availability, the exponential growth of data-crunching power, and the development of new data analytics tools from network analysis to natural language processing and artificial intelligence. These data-driven approaches provide new ways to empirically analyze international law in all its complexity. Rather than only looking at the causes of international law’s emergence or its effect on the world, modern technology allows us to place international law on both sides of the equation. We can assess how thousands of investment treaties differ and investigate the impact of these differences on the outcome of hundreds of investment awards. Vice versa, we can trace how scores of international arbitrations trigger changes in generations of investment treaties. Greater availability of data and more sophisticated means of its analysis thus not only allow us to gain a more accurate picture of the role of law in the world, but also of the inner-workings of international economic law itself. It is this new frontier of empirical legal research that we explore further in this Special Issue.

III. LET THE DATA SPEAK FOR ITSELF: WHY ‘DATA-DRIVEN’ ANALYSIS IS DIFFERENT

Empirical legal research relies by definition on real world observations and their qualitative or quantitative analysis. So what is so dramatically new or different about the new frontier of data-driven research? The answer is that data is becoming ‘big’—and

7 Stewart Macaulay, ‘Wisconsin’s Legal Tradition, Speaking to Benchers before Receiving Brazeau Professorship’, 24(3) *Gargoyl* 6–10 (1994).

8 David M. Berry (ed.), *Understanding Digital Humanities* (Basingstoke: Palgrave Macmillan, 2012), <http://link.springer.com/10.1057/9780230371934> (visited in 5 April 2017).

9 David Lazer et al., ‘Life in the Network: The Coming Age of Computational Social Science’, 323(5915) *Science* 721 (2009).

10 Viktor Mayer-Schönberger and Kenneth Cukier, *Big Data: A Revolution That Will Transform How We Live, Work, and Think*, Reprint ed. (Boston: Eamon Dolan/Mariner Books, 2014).

not just in a quantitative sense. It becomes big because (i) it occupies a more central place in empirical research design; (ii) it allows us to think big looking at *all* available treaties, citations, or adjudicators rather than just subsamples thereof; and (iii) it requires us to make use of big data tools developed in other disciplines for other, related purposes.

A. Data first

To begin with, data is becoming the centerpiece of empirical legal research. It is not only a means to empirically validate a theory, but increasingly a driver of empirical research. In the past, scholars were compelled to start with theory to reduce research questions to a set of more manageable hypotheses that could be validated or falsified. Today, these practical motivations for approaching empirical research questions deductively are less strong. Instead, we can ‘let the data speak for itself’¹¹—inductively identifying patterns and trends in data which then frame our subsequent analysis. Researchers can thus begin their work with the open mind of the explorer.

Data-driven analysis does not, however, mean ‘the end of theory’.¹² Indeed, better and bigger data improves our ability to test, validate, and refine existing frameworks. Theory is also crucial to tell pattern from noise and to turn what would otherwise be ‘fishing for correlations’ into a principled knowledge-advancing exercise. Yet, the greatest potential of data-driven analysis lies in areas where theory is lacking or is under-developed and inductive research can roam freely. Spirling’s work on treaties signed between the US government and Native Americans is a case in point.¹³ Prior research on the impact of an 1871 constitutional change on the content of these treaties, which shifted treaty-making power from the President to Congress, was inconclusive and empirical scholarship patchy. By investigating almost 600 of these treaties through a text-as-data framework, Spirling inductively concluded that they grew harsher over time finding that the USA’s growing bargaining clout rather than institutional change better explained the evolving design of agreements. In contexts like these, where theory alone provides little guidance, inductive, data-driven approaches have their greatest potential.

Just like the natural and social sciences are harnessing the power of data-driven research to explore the unknown and unexpected, legal scholars should use them to uncover the secrets international economic law still holds in store. Inductive, data-driven research has already begun to map the spaghetti bowl of trade agreements,¹⁴ chartered trends in investment treaties¹⁵ or probed the invisible college of International

11 Ibid, at 14.

12 Chris Anderson, *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*, Wired, <https://www.wired.com/2008/06/pb-theory/> (visited 17 April 2017).

13 Arthur Spirling, ‘U.S. Treaty Making with American Indians: Institutional Change and Relative Power, 1784–1911’ 56 *American Journal of Political Science* 84 (2012).

14 Joost Pauwelyn and Wolfgang Alschner, ‘Forget About the WTO: The Network of Relations between Preferential Trade Agreements (PTAs) and “Double PTAs”’, in Andreas Dür and Manfred Elsig (eds), *Trade Cooperation: The Purpose, Design And Effects Of Preferential Trade Agreements* (Cambridge: Cambridge University Press, 2015) 497 <http://ebooks.cambridge.org/ref/id/CBO9781316018453A031> (visited 7 March 2017).

15 See e.g. Wolfgang Alschner and Dmitriy Skougarevskiy, ‘Mapping the Universe of International Investment Agreements’ 19 *Journal of International Economic Law* 561 (2016).

Centre for Settlement of Investment Disputes (ICSID) arbitrators.¹⁶ By putting data first and letting it guide them empirical legal scholars today enjoy new opportunities to discover relevant, but counterintuitive empirical patterns and trends that they did not know existed or would not have expected to find.

B. All the (available) data, not just subsamples

The second major difference with conventional empirical research concerns the scope of the data. In the past, manual coding and analysis forced empirical legal scholars to limit themselves to convenience or random samples or to take other shortcuts even where larger datasets were available.¹⁷ In international investment law, for instance, it has been common practice to refer to provisions in model agreements rather than negotiated bilateral investment treaties (BITs) to describe trends and practices in treaty making more generally.¹⁸ Yet, a host of problems can arise once researchers draw inferences from samples or substitutes. Selection bias in sampling, missing observations, or unwarranted generalizations can produce a skewed or misleading picture of the larger universe in question.

The existence of more data and better techniques for its analysis now allows us to work with the entirety of available data rather than subsets thereof. According to Mayer-Schönberger and Cukier, ‘sampling is an artifact of a period of information scarcity, a product of the natural constraints on interacting with information in an analog era’.¹⁹ Today, in the digital age, data size increasingly ceases to matter as a practical constraint. With the help of modern computers, we can analyze *all* investment agreements, look at *all* cross-references in WTO reports or investigate ties between *all* investment arbitrators. While we will not be able to overcome selection issues entirely, as some treaty texts will remain unpublished and some disputes secret, reliance on all rather than part of the available data allows us to assemble a much more accurate picture of the international economic law universe by avoiding some of the sampling or substitution pitfalls mentioned above. In addition, it enables us to see details and patterns that can only be properly recognized in the aggregate.²⁰ What looked like a tree, rope, snake, and two spears before, suddenly becomes an elephant. The techniques we showcase in this issue—network analysis and text-as-data—are tools uniquely suitable to uncover latent patterns underlying international economic law thereby allowing us to better understand and study the field.

C. Computing rather than reading or counting

The third and final characteristic of data-driven research is a heavy reliance on computers and computer science. This entails two important transformations of the research process. First, we see a gradual transition from legions of research assistants diligently counting or classifying elements of interests to a greater delegation of tasks

16 Sergio Puig, ‘Social Capital in the Arbitration Market’, 25(2) *European Journal of International Law* 387 (2014).

17 See Mayer-Schönberger and Cukier, above n 10, at 12, 23.

18 Chester Brown and Devashish Krishan (eds), *Commentaries on Selected Model Investment Treaties* (Oxford: Oxford University Press, 2013).

19 Mayer-Schönberger and Cukier, above n 10, at 12.

20 Ibid.

to computers. On the one hand, these computers perform the same tasks previously done by humans—just much faster and more cheaply. On the other hand, computers also allow us to perform analyses impossible to conduct by humans, for example, by comparing the textual similarity of hundreds of PTAs with the WTO agreements, as Elsig et al. do in this Issue. Automation and computation thus has the potential to render empirical research both cheaper and more potent. Second, computers more generally play a more important role at all stages of the research. From the collection of data through web scraping, to its analysis using programming languages, and finally, to the dissemination of findings through dedicated websites²¹ and visualizations, computers impact every stage of the research cycle.

Data-driven research is thus likely to gradually transform the life of legal scholars and hopefully also inform legal practice. That is not to say, however, that computers and sophisticated artificial intelligence algorithms will replace human researchers or lawyers. Decisions on what information to collect, which techniques to select, and what conclusions to draw are all choices that remain squarely with us, humans. Moreover, as both Broude et al. and Morin et al. argue in this Issue some tasks cannot be automated or involve ambiguity best handled by human coders. Traditional means of content analysis are thus unlikely to be displaced entirely by automation. Yet, so-called supervised machine-learning techniques in which algorithms are trained on human coded data to then code new data automatically undoubtedly offer new means that will make the life of researchers easier. Data-driven research is thus increasingly shaped, but not dominated by computers.

IV. UNCOVERING LATENT PATTERNS, DEBUNKING PAST MYTHS AND FORECASTING THE FUTURE: THE PROMISES OF DATA- DRIVEN RESEARCH

Now that we know how data-driven work differs from traditional legal empirical scholarship, the next question is: what can we do with all this new data and computational power? Data-driven research offers exciting opportunities for empirical legal research on three different fronts. First, it promises better descriptive and exploratory analyses. Second, new data and analytical techniques present novel ways to investigate causal relationships and open the hitherto largely untouched terrain of legal prediction. Third, data-driven research provides new impetus for theory-building and analytical frameworks that keep pace with the increased complexity of legal data.

A. Exploratory and descriptive research

The availability of large amounts of legal data provides unprecedented opportunities to unveil hidden patterns across treaties or judicial decisions, debunk widely held myths formed through anecdotal evidence, and quantify otherwise abstract debates ranging from international law's fragmentation to the relationship between regionalism and multilateralism in the trade regime.

First, data-driven analysis, especially through text-as-data and network analysis, can be used to make latent patterns in legal data visible. Alschner and Skougarevskiy,

21 An example for a dedicated website accompanying academic publications is www.mappinginvestmenttreaties.com.

for instance, use textual similarity metrics to trace legal consistency, innovation, and diffusion in the investment treaty universe over time.²² They find, for example, that 81% of the Trans-Pacific Partnership (TPP) investment chapter has been copied from prior US PTA investment chapters.²³ Similarly, network analysis can make latent structures visible. Puig's work reveals that ICSID arbitrators form a tight professional network dominated by a small number of elite and repeat arbitrators.²⁴ In this Special Issue, Behn et al. build on Puig's research and investigate the larger network of investment arbitration practitioners identifying new groups of tightly connected individuals and quantifying the degree of double-hatting where actors play the different roles of arbitrator and counsel.

Second, text-as-data and network analysis techniques can be used to debunk widely held truths. Pelc, for instance, uses network analysis of WTO citations to demonstrate how the same countries that publicly insist on the lack of binding precedent in WTO law, in fact, regularly file trade disputes that are primarily aimed at setting a favorable precedent, which they exploit in subsequent cases.²⁵ Alschner relies on automated feature coding to show that, contrary to conventional wisdom, investment arbitration clauses and multi-million dollar claims do not trigger changes in investment treaty design, but that states do tweak their agreements incrementally in response to jurisprudential trends in arbitral case law.²⁶ In this Special Issue, Elsig et al. use textual similarity to show that references to WTO agreements and incorporation of WTO language in PTAs have increased rather than decreased over time and that those countries most forcefully pursuing bilateral and regional agreements are also those that link their PTAs most explicitly to the WTO. This challenges widely held views that PTAs and the WTO are primarily seen as competing fora for norm setting. In a similar fashion, Pauwelyn applied descriptive statistics and network analysis techniques to explain the deep differences between adjudicators in the main international investment and trade dispute settlement bodies. With this Pauwelyn also dismissed the increasingly popular idea of a 'convergence' in these two fields—at least with respect to adjudicators.²⁷

Third, text-as-data and network analysis allows us to quantify otherwise abstract phenomena. Creamer and Godzimirska, for instance, rely on WTO members' statements in DSB meetings relating to the adoption of Panel and Appellate Body reports

22 See Alschner and Skougarevskiy, above n 15; Wolfgang Alschner and Dmitriy Skougarevskiy, 'Rule-Takers or Rule-Makers? A New Look at African Bilateral Investment Treaty Practice' Special Issue on international arbitration involving commercial and investment disputes in Africa TDM, <http://papers.ssrn.com/abstract=2791474> (visited 27 July 2016).

23 Wolfgang Alschner and Dmitriy Skougarevskiy, 'The New Gold Standard? Empirically Situating the Trans-Pacific Partnership in the Investment Treaty Universe' (2016) 17 *Journal of World Investment and Trade* 339 (2016).

24 Puig, above n 16.

25 Krzysztof J. Pelc, 'The Politics of Precedent in International Law: A Social Network Application' (2014) 108 *American Political Science Review* 547 (2014).

26 Wolfgang Alschner, 'The Impact of Investment Arbitration on Investment Treaty Design: Myth Versus Reality' 42 *Yale Journal of International Law* (2017).

27 Joost Pauwelyn, 'The Rule of Law Without the Rule of Lawyers? Why Investment Arbitrators are from Mars, Trade Panelists are from Venus', 109 *American Journal of International Law* 761 (2015) (generally arguing that WTO's success is due to the fact that panelists are 'political' as opposed to investor-state arbitration, which is under fire in part because of 'a-political' arbitrators).

to measure how members perceive the legitimacy of the WTO's judicial organs.²⁸ Stone Sweet and Grisel code arbitral awards, citations and treaty provisions to trace the degree of arbitral lawmaking in investment disputes.²⁹ In this Special Issue, Derlén and Lindholm engage with the conceptual question of how we can best measure the importance of judicial decisions over time. By identifying the place of a decision in a citation network, they show that selected network analysis measures provide a good (or not so good) proxy for the importance of a judicial decision. Finally, Charlotin, also in this Issue, creates the largest dataset to date of citations of international courts and tribunals to put numbers on international law's fragmentation debate by investigating the degree of cross-citations between international economic law tribunals and other international adjudicatory institutions.

B. Causal inference and prediction

Aside from better understanding international economic law, big data and the novel means of its analysis create new opportunities for conducting causal inference studies and pave the way for more accurate predictions of the outcome of legal proceedings or negotiations.

While big data provides more data points and thus better and more accurate assessments, text-as-data or network analysis can also generate new variables for the use in inferential research design. Daku and Pelc's contribution in this Issue, for instance, uses the textual proximity between party submissions in WTO disputes and Appellate Body reports as proxy for parties' influence over judicial rulemaking. Using this method, they find, for example, that WTO members with less legal capacity also have less of an impact on the precedents that shape WTO case law. Broude et al., also in this Issue, create a novel index representing the regulatory space of host states in international investment agreements and use it as a dependent variable to explore drivers of treaty design similarity. The exploitation of more data and better means of analysis thus permits the investigation of relationships and questions that could not have been assessed empirically at the same scale before.

The perhaps most daunting new frontier of data-driven research in international economic law, however, lies in the use of artificial intelligence to look into the future. In the US domestic law context, big data approaches are already deployed to quantify the likelihood of success of claims or to forecast the outcome of a Supreme Court decision.³⁰ We anticipate that similar applications may be on the horizon for trade and investment dispute settlement before long. Artificial intelligence algorithms can also be harnessed to write entire treaties autonomously. Alschner and Skougarevskiy have

28 Cosette D Creamer and Zuzanna Godzimirska, '(De)Legitimation at the WTO Dispute Settlement Mechanism' 49 *Vanderbilt Journal of Transnational Law* 275 (2016)

29 Alec Stone Sweet and Florian Grisel, *The Evolution of International Arbitration: Judicialization, Governance, Legitimacy* (Oxford: Oxford University Press, 2017).

30 Daniel Martin Katz, 'Quantitative Legal Prediction-Or-How I Learned to Stop Worrying and Start Preparing for the Data-Driven Future of the Legal Services Industry' (2013) 62 *Emory Law Journal* 909 (2013); Daniel Martin Katz, Michael James Bommarito and Josh Blackman, 'A General Approach for Predicting the Behavior of the Supreme Court of the United States' (Social Science Research Network 2017) SSRN Scholarly Paper ID 2463244 <https://papers.ssrn.com/abstract=2463244> (visited 24 March 2017).

trained a recurrent neural network on existing investment treaties and built a pipeline that produces treaty text automatically based on user-defined criteria.³¹ This pipeline may help alleviate power asymmetries by proposing a negotiation template that reflects the prior treaty practice of both parties in equal terms. It may also be useful in creating baseline documents for multilateral negotiations by consolidating the texts of thousands of agreements into a single document. Big data and artificial intelligence thus open an exciting new chapter for predictive research and practical applications in international economic law.

C. Shifting theoretical perspectives

Finally, data-driven research also creates demand for new theories as scholars reveal hitherto unknown patterns and trends or offer unique bird's eye perspectives on international economic law phenomena. As big data and new techniques allow us to look both deeper and more widely, new analytical frameworks are needed that keep pace with the complexity of the data. Some of these frameworks, like the contribution by Morin et al. in this Issue, draw from evolutionary biology and complexity theory.³² They acknowledge that the whole of, for example, the global trade governance regime, is larger than its constituent parts, and that the regime evolves in ways similar to other complex adaptive systems we see in nature or other social phenomena, a system that creates its own, nonlinear opportunities for, and balance between, exploration (change or invention of new norms) and exploitation (stability or adoption of existing norms) beyond simply the influence of exogenous pressures and crises.³³ Data-driven empirical analysis thus provides fertile ground to test and develop new theories.

V. CHALLENGES AND LIMITATIONS

A greater role for big data and computational techniques also raises a number of practical challenges and comes with a set of potential dangers and limitations.

A. Challenges

The challenges associated with big data research of international economic law fall into three categories. First, there are challenges associated with getting and sharing data. Second, there is the question of competence: legal scholars often lack the methodological training to fully exploit legal data even where it is available. Finally, even where data and skills are aligned, current research outlets are not well placed to accept or rigorously assess data-driven research.

31 Wolfgang Alschner and Dmitriy Skougarevskiy, 'Can Robots Write Treaties? Using Recurrent Neural Networks to Draft International Investment Agreements' in Floris Bex and Serena Villata (eds), *Legal Knowledge and Information Systems: Jurix 2016* (Amsterdam: IOS Press, 2016) 119–124.

32 J.B. Ruhl, Daniel Martin Katz and Michael J. Bommarito, 'Harnessing Legal Complexity' 355 *Science* 1377 (2017).

33 See also Joost Pauwelyn, 'At the Edge of Chaos? Foreign Investment Law as A Complex Adaptive System, How It Emerged and How It Can Be Reformed', 29 *ICSID Review—Foreign Investment Law Journal* 372 (2014).

1. Databases and data sharing

First, getting good data is often problematic. Sometimes data is only available through commercial providers, which raises licensing and usage issues when data is scraped or subsequently disseminated. More often, data does not exist at all or only in non-machine readable formats, which precludes the deployment of the computational tools described above. Laborious effort is thus required to build datasets for legal analysis from scratch. These data infrastructure considerations raise thorny issues on data versatility, database architecture, and research collaboration.

Building databases is costly. To ensure that public money and researchers' time is spent wisely, databases have to be versatile as opposed to research-question specific. Data versatility has three components: flexibility, sustainability, and compatibility. First, data collection efforts should be flexible enough to help a wider range of scholars in answering their research questions. Rather than omitting information that could later prove helpful, it is important to preserve data to the maximum extent possible³⁴ or to retain links between coding and text³⁵ in order to permit the use and re-use of data in ways that are previously unknown. Second, database efforts have to be sustainable. That means databases have to endure over time and be dynamic rather than static allowing for automatic updates as the information evolves, while tracking versions and changes to the data.³⁶ Finally, data versatility requires compatibility so that it can be exchanged with other researchers and combined with other datasets.

These exigencies of data versatility force researchers to carefully think about questions of database architecture before they embark on their work. What types of information should be collected in the first place? How should information be stored? Ideally, modern databases combine textual data (e.g. treaty texts) with metadata (e.g. treaty signatories) and link both sets of information in a single database.³⁷ It takes foresight and active management to make these databases sustainable. Care has to be taken to secure long-term funding and to allocate data management responsibilities between institutions. Finally, good databases also require at least some degree of standardization in order to make data comparable across datasets. On a spectrum, standardization can span from simply using ISO codes for country names to adhering to elaborate legal data standardization efforts advanced by the legal informatics community.³⁸

34 This can include font type or size in the digitation of texts that may later help tell footnotes apart from main text.

35 Instead of entering coded treaty content information in a spreadsheet, for instance, it should be integrated in the text itself through annotation software to facilitate subsequent machine-learning applications.

36 This is important to ensure comparability across different versions of the same datasets. Backups and version-tracking software also facilitate the restoration of data in case errors occur.

37 These databases can take different forms and can be combined. SQL databases allow for the efficient queries of metadata. XML databases embed metadata directly in texts using tags and are particularly useful for integrating coded information or metadata in documents. For an example of a combination of both, see Alschnner, Wolfgang and Umov, Aleksander, 'Towards an Integrated Database of International Economic Law (IDIEL) Disputes for Text-as-data Analysis', CTEI Working Paper 2016-08, available at: http://graduateinstitute.ch/files/live/sites/iheid/files/sites/ctei/shared/CTEI/working_papers/CTEI-2016-08.pdf

38 E.g. <http://www.legalxml.org>. See also, Sergio Puigand Enric G. Torrents, 'The Case for Linking World Law Data' (15 December 2013). Stanford Public Law Working Paper No. 2371795. Available at SSRN: <https://ssrn.com/abstract=2371795>

The last data challenge relates to data sharing and collaboration. Should databases be open to other researchers? On the one hand, there is a legitimate interest among researchers who have invested time and effort in the creation of a database to be its first users. Furthermore, young databases often have nascent problems that should be fixed before they are opened to others. On the other hand, these justifications grow less persuasive over time. Hence, research databases should become accessible for general use after a reasonable period of time. Otherwise, researchers are incentivized to reinvent the wheel and duplicate already existing infrastructural efforts instead of dedicating scarce research resources to novel work. Ultimately, however, larger collaborative efforts will be needed to build the infrastructure required to conduct large-scale empirical research on international economic law. Interdisciplinary projects bringing together lawyers and computer scientists are particularly important in this respect.³⁹

In an Annex to this Introduction, we provide a non-exhaustive overview of existing databases on international economic law to facilitate further research and, we hope, broader cooperation between entities collecting data.

2. Skills and competencies

A second challenge for data-driven research has to do with the skills and competencies it requires. The surge of legal data and processing power only occurred over the last decade. Today's machine and deep learning capabilities are of an even more recent vintage. Hence, it should come as no surprise that even those researchers trained in empirical methods, including at the PhD level, have not had significant exposure to web-scraping, programming, or machine learning. Even today, and in spite of the growing legal tech industry, law schools with few exceptions do not train their students in computational tools for legal analysis. As a result, few legal scholars have the skills and competencies necessary to fully exploit the legal data treasures now available. But the community is growing. Established researchers can reach out to computer scientists or computational social science scholars in order to access the skill set required. Young legal scholars, in particular, however, should consider acquiring at least a basic data science competency in order to recognize new research opportunities, converse more easily with computer or social scientists, and be able to conduct basic data analysis or web-scraping autonomously.

To be sure, the most sophisticated computational tools are of little use without knowing what questions to ask. That is why legal scholars cannot look to political scientists, economists, or computational scholars working in isolation for empirical answers to pressing legal debates. In the age of increased automation and computation, legal expertise remains crucial. Legal scholars are in the best position to identify relevant queries or think of useful legal applications.

39 An example of such efforts includes the Swiss Network for International Studies (SNIS) project 'Diffusion of International Law: A Textual Analysis of International Investment Agreements', which brings together lawyers, political scientists, economists, and computational linguists. See http://www.snis.ch/project_diffusion-international-law-textual-analysis-international-investment-agreements. See Wolfgang Alschner et al., 'Building a Corpus of Multi-Lingual and Multi-Format International Investment Agreements' in Floris Bex and Serena Villata (eds), *Legal Knowledge and Information Systems: JURIX 2016* (Amsterdam: IOS Press, 2016).

Through their intimate knowledge of the law they have the capacity to contribute in meaningful ways to empirical research without turning themselves into ‘quants’, in particular, by assisting scholars with more technical capacity in the creation of adequate variables or measurements that reflect the complexity of the law in action. Legal data analysis thus needs the involvement of competent scholars with legal training and should not be left to other disciplines operating in silos.⁴⁰

3. Outlets and audiences for computational legal research

Even where researchers join efforts and expertise to conduct data-driven empirical research, they are likely to find it challenging to find proper outlets for their research. First of all, since most legal scholarship is argumentative and not empirical, it may not be intuitively clear to editors and reviewers why empirical scholarship is worth publishing when it confirms rather than challenges widely held views or when it is descriptive rather than normative. Yet, knowing something to be true is different from arguing or believing it to be true. While counterintuitive findings with important normative implications remain incredibly relevant, journals should not shy away from publishing empirical research where it incrementally adds to or confirms our understanding of the world.

Second, computational legal scholarship faces another, more unique hurdle in the sense that tools often have to be perfect to convince in the eyes of many lawyers. Most cutting-edge computational tools rely on probabilistic models to classify data. Just like *google translate* gets most but not all its translations right, machine learning algorithms that map awards based on a human-coded training set will not achieve perfect accuracy. Nor will rule-based coding extract all the features a diligent human may have spotted. This is the trade-off we need to accept in return for the efficiency gains achieved. Some people tend to be more comfortable with the errors of judgments involved in employing human coders than the errors of statistics in relying on computers. In both cases, however, we need to get used to the fact that information is often imperfect, understand the limitations and assumptions on which research findings are based and work on ways to improve accuracy. As Anne van Aaken put it ‘[b]eing shortsighted is usually better than being blind’.⁴¹ The same holds true for computational analysis. These tools do not have to be perfect, as long as they are good enough to be useful (and they will get better over time).

Third, both outlets and consumers of data-driven research have to learn to evaluate its quality. Editorial teams need to acquire the expertise to tell high from low-quality data-driven empirics. Moreover, legal journals also need to start thinking about putting in place conventions that mandate publication of underlying data to allow other researchers to replicate or challenge results. This not only enhances

40 See, Jason Webb Yackee, ‘Do states bargain over investor-state dispute settlement? Or, Toward Greater Collaboration in the Study of Bilateral Investment Treaties’, 12 Santa Clara Journal of International Law 277 (2014).

41 Anne van Aaken, ‘Opportunities for and Limits to an Economic Analysis of International Law’, 3(1) Transnational Corporations Review 27 (2011), at 43.

quality control but also spurs new research as scholars cannot keep datasets to themselves but need to grant access to wider communities.

B. Limitations

Data-driven research also comes with a set of limitations, two of which are particularly dangerous. The first is the risk to conduct theory-less research that mistakes noise in the data for actual patterns. The second concerns data biases that can lead to incorrect conclusions.

1. Pattern or noise? The risk of theory-less research

The main danger of data-driven research is that we let data not only speak but also think for itself. Whenever large sets of data are being analyzed, some variation will emerge. In the best of cases, this variation turns out to be an important but hitherto unknown pattern in the data that spurs new research and inspires new theories. In the worst of cases, random noise is mistaken as a meaningful relationship that is then used to make ambitious claims about how international economic law operates. So how can we tell the difference between pattern and noise?

The answer is theory, good statistical analysis, and context-specific qualitative triangulation. While we can let the data speak for itself we need to interpret whatever it tells us through the lens of our knowledge of the world and the law. If some novel pattern emerges from our data that challenges our prior beliefs, it could mean that we are onto something new and exciting, but, more likely, it simply means that we attribute too much importance to what is only random noise. Inductive, data-driven research is no substitute to critical reflection, regression analysis, interviews or in-depth case studies, and theorization, but just a complement to or another path to put to use these other, more traditional methods. Hence, researchers need to be careful not to take variation in data at face value, but should rather consider it as an opportunity to dig deeper using the entire toolkit of qualitative and quantitative methods with the same analytical rigor as they would employ in deductive empirical scholarship.

2. Skewed data and incorrect conclusions

A plight of all empirical legal scholarship is missing or skewed data. The omission of data points can bias measurements and variables leading to faulty conclusions. This risk is particularly pronounced when working with network analysis or text-as-data techniques. Imagine you draw up a list of your friends and your friends' friends. The resulting network will have you at its center. Now, if we were to take you out of the network, the shape of the network would shift dramatically and we would no longer understand that what all these different people have in common is their friendship with you. Therefore, in a real-life network, it is crucial to have a complete dataset since even a few missing data points can dramatically alter our conclusions.

Text-as-data is also susceptible to skewed data, but in a different manner. Many text-as-data techniques are language-specific. Hence, we can meaningfully compare the textual similarity of two English BITs or two Spanish BITs but not one English and one Spanish BIT. While there are ways around that problem, it creates incentives to work with documents in the most widely used language, which in the international

economic law context is likely to be English.⁴² This then risks neglecting non-English documents. These limitations are not insurmountable obstacles to high-quality research, but important considerations to keep in mind when collecting data, selecting methods, and drawing conclusions.

VI. THE TIME IS RIPE

In light of its promises and, notwithstanding its challenges and limitations, we believe the time is ripe to seriously embark on this new data-driven frontier in international economic law scholarship. The availability of data on treaty texts and drafts, as well as international decisions and proceedings has increased manifold. The sheer volume of treaties and international decisions is now exponentially higher than the universe available just 10 or 20 years ago. Text-as-data and network analysis allows the processing of larger amounts of data than traditional approaches and permits scholars to explore new and unexpected patterns.

The highly decentralized and often contentious nature of international economic law makes the use of these techniques even more appealing. Overlapping layers of treaties, tribunals, cases, and international organizations form a complex and quickly evolving web hard to penetrate using traditional analyses. This same complexity makes the benefits, but also the weaknesses of international economic law harder to grasp at a time in which challenges to its institutions, norms, and intellectual underpinnings are multiplying. As a result, thinking about the production, enforcement, and contestation of international rules through the lens of complex systems and making sense of it with the help of network analysis or text-as-data study is today more pertinent than ever.

Last but not least, data-driven analysis of international economic law does not only provide exciting new opportunities for scholars, but also has the capacity to make life easier for legal practitioners in all corners of the world. Be it through open-access interactive websites, legal tech startups, *pro bono* legal data research through transnational law clinics⁴³ or simply practice-oriented scholars, this new frontier has the potential to impact the practice of law as much as the study of it.

For all these reasons, we believe that the time is ripe for data-driven empirical analysis of international economic law and we hope that this Special Issue will help introduce this emerging field and its growing cohort of computer-savvy scholars to a wider range of legal scholars and practitioners.

Annex: Databases on International Economic Law

42 On English as international law's lingua franca, see Anthea Roberts, *Is International Law International?* (Oxford: Oxford University Press, 2017, forthcoming).

43 For a prominent example in the international economic law field, see www.tradelab.org and its global network of legal clinics and publicly shared research output.