Do legal remedies promote investment? New evidence from a natural experiment in the investment treaty network*

Cree Jones†

December 15, 2017

Abstract

Many developing countries are considering curtailing legal remedies available to investors in bilateral investment treaties (BITs). As with any trade or investment policy, there are two primary challenges to evaluating how this policy change will affect investment: (1) investment policy is correlated with unobserved drivers of investment, and (2) studying a change in policy at the provision level requires provision-level information on a large number of BITs. I identify a natural experiment to address the first challenge: an arbitration decision that unexpectedly endowed some investors with new and stronger legal remedies through an unanticipated application of the "most favored nation" treatment. I also introduce a new comprehensive database, created by me in partnership with the United Nations Conference on Trade and Development (UNCTAD), that contains provision-level information for over 2,500 BITs. Using this database and natural experiment, I present robust evidence that stronger legal remedies in a BIT do not lead to more investment. I also present suggestive evidence that stronger legal remedies imposed on a host economy may lead to a decline in investment as host economies react to the increase in the implicit cost of hosting foreign capital caused by its increased exposure to arbitration and tighter constraints on its ability to regulate foreign capital.

*I am very grateful to my advisors Steve Levitt, Jonathan Dingel, Felix Tintelnot, Thomas Miles and Thomas Ginsburg for their many conversations, constructive feedback, and direction and encouragement throughout my PhD. I thank John Romalis, Chad Syverson, Stephane Bonhomme, Derek Neal and Rodrigo Adao for comments and suggestions that substantially improved the paper. I also thank seminar participants at the University of Chicago Economics Trade Working Group and the University of Chicago Economics Applied Micro Lunch for helpful comments. I gratefully acknowledge the International Investment Agreements (IIA) Section at the United Nations Conference on Trade and Development (UNCTAD) and thank them for the opportunity to work with them to create the UNCTAD IIA Database, which made this paper possible. All errors are my own.

†University of Chicago. Email: creejones@uchicago.edu. Website: www.creelanejones.com.

1 Introduction

A bilateral investment treaty (BIT) is a treaty signed by two countries that grants certain legal protections to investors from one signatory when investing in the other signatory. Common protections include protection from expropriation and a guarantee of the free transfer of assets. If a host country violates the treaty, a BIT grants investors legal remedies that often include the right to seek treaty or contract enforcement in an arbitration tribunal outside the jurisdiction of the host country.

More than 3,200 BITs have been signed since 1959, with 2,321 BITs currently in force.[1] There is mixed empirical evidence on whether BITs attract investment. There is also mounting concern among many scholars and policy makers that legal remedies in BITs impose high costs on host economies. These costs are realized through eroded domestic sovereignty caused by disputes being removed from domestic courts. Costs are also realized through higher exposure to arbitration proceedings, resulting always in attorneys fees and often in the payment of damages. Sixty-two new arbitration cases were filed in 2016 bringing the total number of cases filed to more than 760.[2]

In response to an increasing exposure to arbitration proceedings, many host countries are considering limiting legal remedies in investment treaties. Some countries have recently taken decisive action. For example, India, the tenth largest recipient of FDI in 2015, was respondent to 12 arbitration disputes in 2015.[3] In July of 2016 India sent notice to 57 countries of its unilateral termination of BITs with those countries and requested an opportunity to renegotiate those agreements. Based on the current draft of the India Model BIT, new BITs will likely include much weaker legal remedies.[4]

There is little existing empirical work on how weaker legal remedies in a BIT may affect investment decisions of MNCs. There are three primary challenges to credibly answering this question. The first is that one would need data on the precise content of each investment treaty, both in terms of legal remedies and other substantive protections. Having these data would enable a researcher to compare outcomes of BITS with weak versus strong legal remedies while controlling for the strength of the rest of the protections in the agreement. The second challenge is endogeneity: two countries are more likely to sign a strong BIT if the investment opportunities between the two countries are high and the protections in the treaty are relevant. The third challenge is isolating the effect of the legal remedies from the effect of the remainder of the BIT. Even if one could credibly argue that an event study centered around the entry into force of BITs was causal evidence of the effect of a BIT on attracting investment, it would be difficult to isolate the effect of legal remedies on FDI from the effect of the other protections in the BIT that are introduced concurrently with the legal remedies.

To address the data challenge I have worked in partnership with the United Nations Conference on Trade and Development (UNCTAD), managing a global team of more than 550 law student at 42 universities in 22 countries, to create a comprehensive database on the content of all BITs with publicly available text (approximately 2,500 BITs). To address the endogeneity and isolation challenges I have identified a natural experiment arising from an arbitration decision (Maffezini v. Spain) issued in January of 2000. The arbitration decision was the first to allow an investor to invoke the most favored nation (MFN) protection in its treaty to gain access to a better legal remedy in another active BIT signed by

[1] UNCTAD (2016a)
[2] UNCTAD (2016b)
[3] UNCTAD (2016b)
[4] India Model BIT (2016)

the host country. This natural experiment enables me to both eliminate the endogeneity concern as well as isolate the effect of legal remedies on FDI from the effect of the rest of the BIT protections.

Using my new database, in tandem with this unexpected arbitration decision, I am able to identify a treatment and control group of origin-destination country pairs and implement a series of difference-in-differences regressions to study how bilateral FDI responds to an unexpected improvement in the legal remedies of a BIT. I find robust evidence that stronger legal remedies introduced by Maffezini v. Spain do not cause an increase in foreign investment. Only one of the 32 treatment coefficient estimates presented in this paper is positive, and this estimate is approximately zero. I also find suggestive evidence that stronger legal remedies imposed on a host economy may actually lead to a decline in foreign investment: My preferred model specification and treatment measure produce a statistically significant coefficient estimate of -0.57, indicating that a one standard deviation increase in treatment will cause, on average, a 14.25 percent reduction in FDI. This coefficient estimate is sensitive to subjective decisions I have made when developing continuous measures of treatment: estimates calculated using other continuous measures built from different decisions also produce negative coefficient estimates, but the size of the estimate is dampened for some measures and the standard errors are amplified for others. This negative result may be explained by the following mechanism: imposing new legal remedies on a host economy may increase the implicit cost of accepting foreign capital by increasing the host economy's risk of arbitration and constraining its ability to regulate foreign capital. This price increase lowers the marginal benefit to the host economy of attracting or accepting foreign capital, leading to a decrease in its demand for foreign capital.

This project builds on an active literature on whether or not BITs affect FDI flows. There are many papers that use variations of gravity models of investment to run regressions of FDI on an indicator variable for an active BIT and various controls. The findings of these papers, i.e. whether there is a positive, negative or no correlation between BITs and FDI, vary depending on the country pairs included in the sample and the control variables added to the model specification. Most of these papers are able to say whether the correlation they find is the causal effect of a BIT. See UNCTAD (2014) for a comprehensive summary.

There are a handful of papers that attempt causal inference. For example, Busse et al. (2010) use the number of BITs signed by each country in a pair as an instrument for whether a BIT is active between the country pair. The authors find that BITs increase FDI. However, the instrumental variable they use is likely correlated with unobservables: the number of BITs signed by each country may be correlated with things like unobserved investment opportunities, the availability of investment capital in an origin economy, the production technology in an origin economy, etc. It is also possible that the exclusion restriction is not satisfied, since the signing of other BITs may increase bilateral FDI through channels outside of a BIT (i.e. new investment induced by the other BITs may introduce new production technologies, create a more skilled labor force, etc.). Berger et al. (2012) is the only paper in the literature that goes beyond coding BITs as a single indicator variable. That paper considers the content of a BIT in a single dimension: whether or not arbitration is included in the treaty. The authors find a positive correlation between FDI flows and the inclusion of an arbitration provision. They stop short of making a causal argument.

This paper makes four contributions to this literature. First, this paper introduces a new and comprehensive dataset on the content of over 2,500 BITs. Second, this paper identifies a new causal identification

strategy: using an unanticipated interpretation of a treaty provision to study whether or not a particular class of protections (better legal remedies) in a BIT cause an increase in FDI. Third, this new dataset and identification strategy are combined to provide a credibly identified answer (stronger legal remedies in a BIT do not cause an increase in foreign investment) to a variation of the literature's primary research question. Finally, this paper presents suggestive evidence of a new result: imposing stronger legal remedies on a host economy may actually lead to a decline in investment.

This paper also contributes to a broader literature that considers whether and to what extent the rule of law plays a role in facilitating growth in a developing economy.[5] Domestic legal institutions in a developing economy may be biased or unreliable. Legal remedies in a BIT create an outside option for investors to resolve legal disputes with the host government. Understanding how protected investors respond to changes to these legal remedies can help academics and policy makers understand how better legal institutions may facilitate investment and long run growth in a developing economy.

The paper proceeds as follows: Section 2 introduces the UNCTAD International Investment Agreements (IIA) Database, which is necessary to implement the identification strategy described in section 3. Section 3 describes the Maffezini v. Spain decision and how it acts like a natural experiment to modify the legal remedies that may be invoked under a BIT. Section 4 presents the empirical findings of the paper. Section 5 concludes.

2 The UNCTAD IIA Database

The UNCTAD IIA Database (the Database) contains information on the content of more than 2,500 BITs across 121 treaty provisions and their variations. This provision-level detail is necessary to study how investors may respond to changes to these provisions. In my role as lead consultant to the IIA Section at UNCTAD, I designed and executed the data collection process and managed the efforts of more than 550 law students at 42 universities in 22 countries. This section begins with a description of the methodology used to create the Database. The section concludes with a broad description of the content of the Database and a more detailed description of the legal remedies and their summary statistics for BITs that were active at the time of the Maffezini v. Spain decision.

2.1 Methodology

The method used to create the Database consists of two parts. The first is the method used to code the content of a BIT. The second is the method used to scale the project to cover the universe of BITs. The first part was developed internally by UNCTAD and the second part was developed by me, with technical and administrative support from UNCTAD.

The method for cataloging the content of a BIT is recorded in a 93 page "IIA Mapping Guide". This document describes, in detail, the 121 different provisions that are coded as part of the project, complete with coding instructions and treaty text examples for each provision. Each BIT is coded independently by two readers, who then discuss and resolve disagreements in their coding.

This coding method is time intensive. During the fall of 2012, I, at the request of UNCTAD, developed a framework to scale the coding of BITs through the UNCTAD IIA University Mapping Project

[5] See, for example, Beck and Levine (2005) and Rubin (2005).

(the Project). From January 2013 through December 2016 I worked as lead consultant to manage the implementation of this Project. It is estimated that more than 20,000 hours were spent on the creation of the Database, which was completed in January of 2017.

The Project was completed in six segments. Each segment was organized during an academic term, with the first segment being held in the winter of 2013 and the sixth segment being held in the autumn of 2016. Six sets of students, organized in teams by university, participated in the different segments. Student teams were each led by a professor at their home institution with a specialization in international investment law. University professors were recruited and trained by me prior to the beginning of each segment. During each segment, all students completed a training exercise in which they each mapped the same two BITs. These training BITs were also mapped by UNCTAD. I then compared each student mapping to the UNCTAD mapping to gauge accuracy of student work and provide students with feedback as part of their training.

After training, students completed three rounds of mapping. During each round I assigned each student a partner and a mapping assignment. Each mapping assignment included four unique BITs and one common BIT (the control BIT already mapped by UNCTAD). Students first completed the mapping assignment independently. Students then compared their own mapping with their partner mapping and worked together to resolve any discrepancies. I then compared a partnership mapping of the control BIT with the UNCTAD mapping. I included a mapping assignment in the Database only if a partnership achieved 95 percent accuracy on their mapping of the control BIT (i.e. matched the UNCTAD mapping of the control BIT on at least 95 percent of the provisions). Student accuracy developed over the course of each round, with approximately 50 percent of partnerships achieving 95 percent accuracy in their first mapping round and approximately 80 percent achieving 95 percent accuracy in their final mapping round. Mapping assignments that did not achieve 95 percent accuracy were mapped by a new partnership in a subsequent round. The average accuracy rate of student work included in the Database is 98.2 percent.

2.2 Content

BITs are typically structured to perform five major functions: (1) define the entities and/or assets that will be protected (definitions), (2) articulate investor protections (protections), (3) set forth Host State obligations (obligations) (4) establish how long the protections and obligations will last (duration), and (5) specify the legal remedies, i.e. how the treaty may be enforced (legal remedies). While the structure is uniform across BITs, there is variation in the composition of protections in each of these categories. This section provides an overview of the treaty provisions that are relevant to the identification strategy used in Section 3: (1) most favored nation (MFN) treatment (one standard of treatment under the investor protections category), and (2) legal remedies.[6]

The MFN protection guarantees that a host government will treat protected investors at least as well as any other foreign investor. The MFN doctrine was adopted from trade agreements and historically applied to commercial policies (and not to procedural protections like legal remedies) governing investment.[7] Of the 1,268 BITs in the Database that were active in 2000, 1,245 (98.2 percent) include MFN protection.

The legal remedies in a BIT confer rights on investors that enable them to enforce the treaty through international arbitration. These legal remedies can be boiled down to seven provisions: three baseline

[6] A summary of the remaining provisions is provided in Appendix A.

[7] The historical context of the MFN doctrine prior to Maffezini v. Spain is discussed in more detail in Section 3.

provisions, three limitations and one qualification. Each are considered in turn.

The first baseline provision specifies what can give rise to a claim (scope). The Database documents three variations of this provision. The first is the most broad, allowing an investor to bring a claim against a host-country for any dispute related to an investment (treaty + any). The second is more narrow, allowing an investor to bring a claim only for violations of the treaty or for other specified reasons (i.e. a dispute arising under a contract between the investor and the host country) (treaty + other). The third is the most narrow, only allowing claims for disputes arising under violations of the treaty itself (treaty only).

The second baseline provision specifies the court or tribunals before whom a claim can be brought (forum). These options typically include the International Centre for Settlement of Investment Disputes (ICSID) housed at the World Bank and/or the United Nations Commission on International Trade Law (UNCITRAL) housed at the United Nations. Other forum options may also be included, such as the Arab Investment Council or the International Chamber of Commerce (other forums). A small number of BITs list domestic courts of the host economy as the only forum available to arbitrate a dispute (domestic courts only). A handful of BITs don't list any forum options (none).

The third baseline provision specifies how different forums may interact with one another, i.e. if an investor begins a claim in a particular forum they may or may not forfeit their right to bring a later claim arising under the same breach of the treaty in a different forum (forum interactions). The Database documents five variations of interactions between forums. First, the treaty may explicitly preserve the right of an investor to file an arbitration proceeding even if a claim has already been brought before a domestic court (arbitration right preserved). Second, the treaty may make no reference to these interactions (none). Third, the treaty may impose a no U-turn condition, meaning that once an investor has chosen arbitration, the claim cannot be brought before a domestic court (no U-turn). Fourth, the treaty may impose a fork-in-the-road condition, meaning that an investor must choose at the outset to bring their claim either before a domestic court or before an arbitration tribunal (fork-in-the-road). Note that this variation is more restrictive than the no U-turn variation, since the no U-turn variation still allows an investor to move a dispute to arbitration after starting a domestic proceeding. Fifth, the treaty may require that a dispute first be litigated in the domestic court as a pre-condition to bringing a claim in another forum (local remedy first). The enforcement of this local remedy first requirement is the point at issue in the Maffezini v. Spain decision.

There are three limitations that a BIT may impose on the scope of legal remedies. The first limitation may only allow legal remedies for violations of a subset of the protections in the BIT (limit: provision). One common iteration of this limitation only allows claims that arise under expropriation and excludes all other possible claims. The second limitation may exclude certain policy areas from legal remedies (i.e. if the excluded policy area is public health, an investor may not bring a claim against an action taken by the host government if that action was taken to protect the public health) (limit: policy). The third limitation prevents legal remedies for a claim arising under a tax policy imposed by the host country (limit: tax).

There is also a qualification that may severely limit an investor's access to legal remedies. This qualification is when a host country does not consent to arbitration in the treaty and instead requires an investor to get permission from the host country after a breach has occurred and before a claim can be brought in arbitration (consent withheld).

There are 1,485 BITs that are active at the time of the Maffezini v. Spain decision. Of these BITs, 1,268 (85.4 percent) are included in the Database. The remaining 217 BITs either do not have a publicly available text or are only available in a less accessible language, i.e. a language other than English, Spanish, French, German, Portuguese, Russian, Italian or Arabic. Of these 1,268 BITs, 1,163 BITs (91.6 percent) include legal remedies for investors. The following table summarizes the number and share of the 1,268 BITs in the Database that were active in 2000 that include a particular variation of each of the legal remedy provisions discussed in this section.

Table 1: Legal Remedies in Active BITs in 2000

Provision	BIT count	BIT share
scope: BIT + any	898	70.76
scope: BIT + other	43	3.39
scope: BIT only	222	17.49
scope: no legal remedies	105	8.28
forum: ICSID + UNCITRAL	540	42.55
forum: ICSID or UNCITRAL	575	45.31
forum: other only	25	1.97
forum: domestic courts only	14	1.10
forum: none	9	0.71
forum: no legal remedies	105	8.28
interact: arbitration right preserved	42	3.31
interact: none	805	63.44
interact: no-U-turn	35	2.67
interact: fork in the road	218	17.18
interact: local remedy first	63	4.96
interact: no legal remedies	105	8.28
limit: provision	105	8.28
limit: policy	21	1.65
limit: tax	27	2.13
consent withheld	76	5.99

These data identify the precise nature of the legal remedies in each of these BITs. Using this information, I can now identify which origin-destination country pairs are endowed with new and stronger legal remedies through the Maffezini v. Spain decision. The method for identifying these treated pairs as well as their level of treatment is described in detail in the following section.

3 Maffezini v. Spain

In 1997 an Argentine investor, Emilio Agustin Maffezini, filed an arbitration claim against Spain under the Argentina-Spain BIT signed in 1991. According to the BIT, Maffezini was required to first fully litigate his claim in Spanish Courts before a claim could be brought before an arbitration tribunal (local

remedy first).[8] Maffezini cited two facts:

1. Spain had signed a BIT with Chile that did not include the local remedy first condition.

2. The Argentina-Spain BIT included MFN protection.

Maffezini then argued that the MFN protection in the Argentina-Spain BIT allowed him to invoke the better legal remedy in the Chile-Spain BIT to avoid litigating first in Spanish courts. In its 2000 decision, the panel of three arbitrators unanimously agreed with Maffezini, thus allowing the claim to move forward.[9]

As discussed briefly in the previous section, the MFN doctrine originated in trade agreements and historically applied to commercial policies like tariffs and market access.[10] Prior to Maffezini v. Spain it was generally understood by policy makers that MFN treatment in the context of investment treaties was limited in scope to similar commercial policies like taxes, subsidies and regulatory takings and did not extend to legal remedies like access to arbitration.[11] After Maffezini v. Spain, investors now had legal precedent to invoke any legal remedy in any active investment treaty signed by the host state, rather than relying exclusively on the legal remedies in the treaty with the investor's home country. Thus, protected investors now have access to (weakly) better legal remedies.

Maffezini v. Spain presents an opportunity to study whether or not foreign investment increases in response to an unexpected endowment of stronger legal remedies. I can identify a treatment group: the set of origin and destination pairs endowed with at least one new legal remedy by Maffezini v. Spain; and a control group: the set of origin and destination pairs with an active BIT that have no change to their legal remedies. I can also construct various measures of treatment for each dyad. These treatment measures include a simple indicator variable, a normalized count variable that counts the number of provisions that adopt a stronger variation, and other more nuanced (and subjective) measures of treatment. I can use these treatment measures to run difference-in-difference regressions, discussed in detail in Section 4. The remainder of this section first describes how to identify the treatment and control groups. The section concludes by introducing various methods of measuring treatment.

3.1 Identifying treatment and control groups

Treatment by Maffezini v. Spain is determined by two conditions. The first is whether a BIT between an origin and destination pair includes an MFN provision. The second is whether the host economy has at least one other active BIT that include one or more stronger legal remedies relative to the BIT governing the country pair. Satisfying the MFN condition can be observed directly in the UNCTAD IIA Database. Satisfying the better legal remedy condition can be observed by constructing a treaty network for each host economy as described below.

First, it is straight forward to construct an ordinal ranking of the different variations of the legal remedy provisions. Recall, for example, that there are three variations of the scope provision: (1) treaty + any, (2) treaty + other, and (3) treaty only. Treaty + any is more comprehensive than treaty + other,

[8] Argentina Spain BIT (1991)
[9] Maffezini v. Spain (2000)
[10] See, for example, Hornbeck (1909), discussing MFN provisions in 18th and 19th century trade agreements.
[11] Vesel (2007)

which is more comprehensive than treaty only. As a result I can ordinally rank these variations using rankings A, B and C, where A is the most favorable variation for an investor and C is the least favorable.

Next, I construct a treaty network for each country when acting as a host economy. Consider a simple case in which a country has signed four investment treaties. Table 2 documents the rank of the variation of each legal remedy provision in each of the four treaties (black text). Table 2 also documents the rank of the best variation of each legal remedy across all active treaties signed by the host economy (red text). Note that an investor may now invoke this best variation through its MFN protection after the Maffezini v. Spain decision.

Table 2: Identifying Treatment by M. v. Spain

	Treaty 1	Treaty 2	Treaty 3	Treaty 4
Scope (X_1)	A	B	B	C
	A	A	A	A
Forum (X_2)	C	C	C	B
	B	B	B	B
Forum Interaction (X_3)	D	C	C	B
	B	B	B	B
Limit on provisions (L_1)	A	A	A	A
	A	A	A	A
Limit on policy (L_2)	B	B	B	B
	B	B	B	B
Limit on tax (L_3)	B	B	B	A
	A	A	A	A
Consent withheld (Q_1)	A	B	A	A
	A	A	A	A

Conditional on each of the four BITs including a MFN clause, the ordinal rank of each provision is uniform across the treaties after Maffezini v. Spain, where all provisions now take on the highest ordinal rank in the set of treaties. As a result, the set of legal remedies is now the same for all four treaties. Note that all four BITs may be treated by Maffezini v. Spain, even if a single BIT is most favorable across a majority of provisions. This is because this most protective treaty may still incorporate a stronger remedy from a relatively weaker treaty. For example, in Table 2, Treaty 4 has the strongest ordinal rank across the destination network for all but one of the provisions (scope). After Maffezini v. Spain, the scope has increased from an ordinal rank of C to A, thus expanding the legal remedies available in Treaty 4.

Using this methodology I can divide the 2,536 origin-destination dyads governed by the 1,268 BITs in the Database that were active in 2000 into treatment and a control groups. This method results in a control group of 342 dyads (13.5 percent) and a treatment group of 2,194 dyads (86.5 percent). Within the control group, 44 dyads avoid treatment by not including MFN in their active BIT, 296 dyads avoid treatment by not realizing a change in legal remedies, and 2 dyads avoid treatment by failing both conditions.

3.2 Estimating treatment

This paper develops eight different measures of treatment. The first codes treatment as an indicator variable. The second simply counts how many of the seven provisions are treated and divides by seven so that treatment is bounded by 0 and 1. The remaining six measures take into account the magnitude of the change in each provision, assign scores to these changes, and impose a functional form to create a continuous treatment measure bounded by 0 and 1. The remainder of this section describes the precise method used to calculate one of these six measures. The remaining five are described in Appendix B.

Each continuous measure of treatment is calculated in three steps. First I develop a method to measure the strength of the original set of legal remedies in each BIT. Second, I use this same method to measure the strength of legal remedies after treatment. Third, I construct a treatment variable by taking the difference of the new and the old strength measures.

Table 3: Estimating Treatment by M. v. Spain

	Treaty 1	Treaty 2	Treaty 3	Treaty 4
Scope (X_1)	1	0.9	0.9	0.8
	1	1	1	1
Forum (X_2)	0.8	0.8	0.8	0.9
	0.9	0.9	0.9	0.9
Forum Interaction (X_3)	0.7	0.8	0.8	0.9
	0.9	0.9	0.9	0.9
Limit on provisions (L_1)	1	1	1	1
	1	1	1	1
Limit on policy (L_2)	0.9	0.9	0.9	0.9
	0.9	0.9	0.9	0.9
Limit on tax (L_3)	0.9	0.9	0.9	1
	1	1	1	1
Consent withheld (Q_1)	1	0.5	1	1
	1	1	1	1
REM (strength)	0.675	0.3375	0.675	0.78
	0.84	0.84	0.84	0.84
Treatment (ind.)	1	1	1	1
Treatment (count)	$\frac{2}{7} \approx 0.286$	$\frac{5}{7} \approx 0.714$	$\frac{4}{7} \approx 0.571$	$\frac{1}{7} \approx 0.143$
Treatment (cont. 1)	0.165	0.5025	0.165	0.06

To measure the strength of the original legal remedies in each BIT I start with the ordinal rank system developed in the previous section. For the three baseline provisions (X_1, X_2, and X_3), I replace the ordinal rank with a numerical score, starting at 1 and decreasing the score by 0.1 for each step down in the ranking.[12] For the three limitations (L_1, L_2, and L_3), I assign a score of 1 if the limitation is not

[12]Note that the methodology is similar for the other 5 continuous measures of treatment, but with different variations in the size of the step function and in the functional form of the measure. The purpose of these additional measures of treatment is to provide a robustness check to the subjective nature of the decisions made when creating a continuous

included and a score of 0.9 if the limitation is included. For the consent qualification (Q_1), I assign a score of 1 if consent to arbitration is unconditional and a score of 0.5 if consent is conditioned on the approval of a host economy.[13] The overall strength of legal remedies is calculated using a simple average of the strength of the three baseline provisions, scaled by the limitations and the consent qualification. Note that the baseline provisions are scaled by the limitations and qualification provisions, since including these provisions limit access to the baseline provisions.

$$REM = Q_1 \left[\frac{1}{3}(X_1 + X_2 + X_3) \right] L_1 \times L_2 \times L_3$$

The measure is bounded by 0 and 1 by construction with strong legal remedies being close to 1 and weak legal remedies being close to 0. Applying this method to the same set of four treaties in the previous section yields the results in Table 3. Note that, as before, the strength of each provision is uniform across the treaties after Maffezini v. Spain, where all provisions now take on the highest value across the set of treaties. As a result, the strength of the legal remedies is now the same for all four treaties.

Using the original (REM_o) and modified (REM_m) strength of legal remedies we can now create a measure of treatment by Maffezini v. Spain:

$$TREAT = REM_m - REM_o$$

The treatment measures for our hypothetical treaties above are 0.165, 0.5025, 0.165 and 0.06 respectively.

Applying this method to the set of 2,536 origin-destination dyads governed by the 1,268 mapped active BITs enables me to measure the original and modified strength of legal remedies as well as the measure of treatment for each dyad. The distribution of each measure is plotted in Figure 1.

Prior to Maffezini v. Spain, there is more variation in the strength of legal remedies across the investment treaty network. After Maffezini v. Spain, legal remedies are almost equalized across the network through the application of MFN. Note that, although the strength of legal remedies is much more homogeneous after Maffezini v. Spain, the distribution of our continuous measure of treatment is not. This distribution is right skewed with a large number of dyads having very little exposure to treatment. The distribution also has a fat right tail with 328 dyads (12.9 percent) realizing an absolute gain of at least 0.4 (1.6 standard deviations) in the strength of their legal remedies.

---

measure of treatment.

[13] Note that the size of this penalty reflects the fact that this qualification will severely limit an investor's access to the broader set of legal remedies.

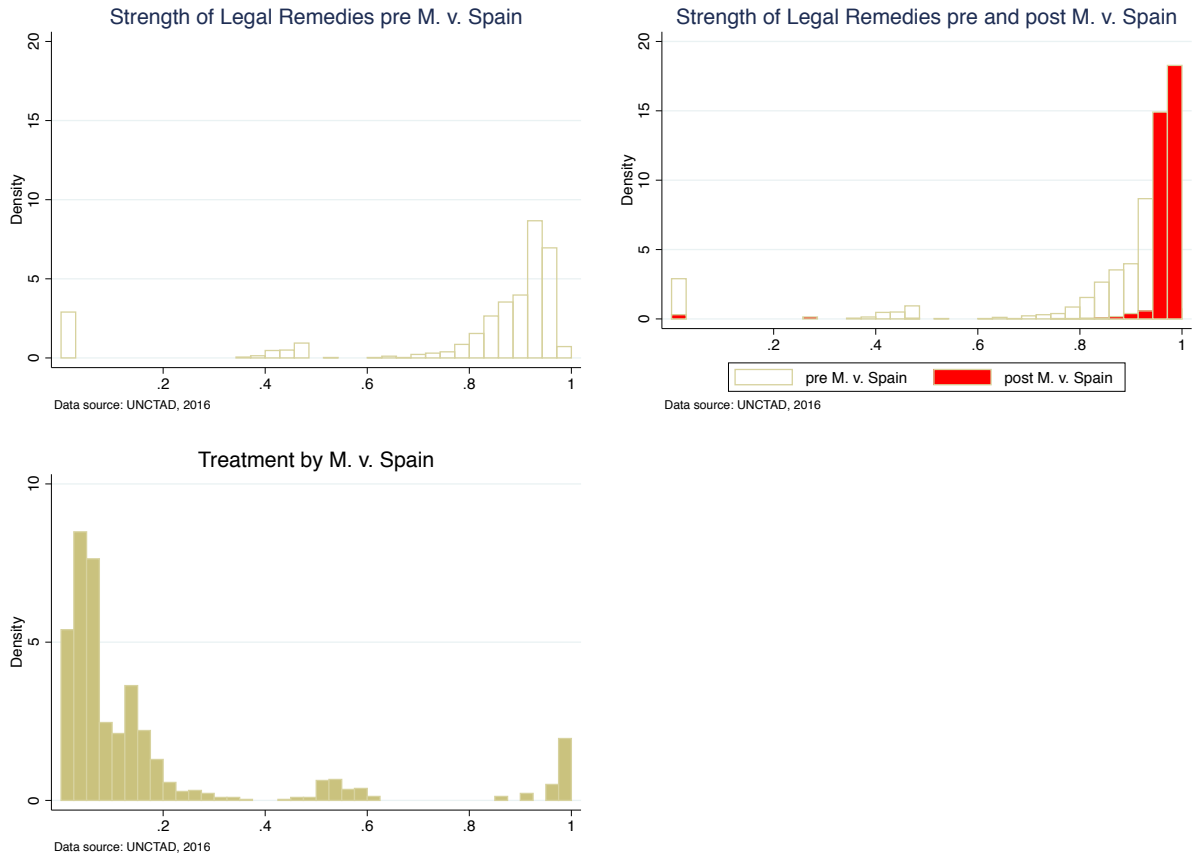

Figure 1: Distribution of Treatment

4 Empirical Results

With treatment and control groups precisely defined and treatment measures in hand, I can now implement a series of difference-in-differences regressions to study how foreign investors and host economies respond to an unexpected increase in the strength of legal remedies in a BIT. This section begins with a short description of the data and an introduction to the difference-in-differences model specification. The section then presents the empirical findings of the paper. The section concludes by presenting additional robustness checks to support the primary findings of the paper.

4.1 The Data

The dataset is a panel from 1990 to 2014 for 2,527 origin-destination dyads that have an active BIT in 2000 that is included in the UNCTAD IIA Database.[14] The outcome variable of interest is bilateral FDI flows, available in a private database maintained by UNCTAD. The treatment variable is treatment by Maffezini v. Spain. This section will use eight measures of treatment, the three measures discussed in the previous section and five additional measures described in Appendix B. Standard control variables from the trade literature (GDP, distance between countries, common border, common language, prior

[14] Due to data limitations in other data sources, 9 of the 2,536 dyads are not included in the panel.

colonial affiliation, common currency, the existence of a free trade agreement (FTA), General Agreement on Tariffs and Trade (GATT) membership, common market, and common customs union) are available in a public database maintained by CEPII.[15] I also use seven different measures of the original strength of each active BIT as an additional control. The first measure is an indicator that equals 1 if a BIT is active between a dyad in a particular year. The remaining six measures are constructed from the UNCTAD IIA Database using a method similar to the method used to construct the continuous measures of treatment.[16]

4.2 The Model

This section will use the following four variations of a difference-in-differences model specification:

$$FDI_{ijt} = \exp\Big[\beta_0 + \beta_1 \times (Treat_{ij} \times Post_t) + \beta_2 \times Treat_{ij} + \beta_3 \times Post_t\Big] \times \epsilon_{ijt} \tag{1}$$

$$FDI_{ijt} = \exp\Big[\beta_0 + \beta_1 \times (Treat_{ij} \times Post_t) + \beta_2 \times Treat_{ij} + \beta_3 \times Post_t + \alpha' X^1_{ijt}\Big] \times \epsilon_{ijt} \tag{2}$$

$$\begin{aligned} FDI_{ijt} = \exp\Big[&\beta_0 + \beta_1 \times (Treat_{ij} \times Post_t) + \beta_2 \times Treat_{ij} + \beta_3 \times Post_t + \alpha' X^1_{ijt} \\ &+ \gamma_i + \gamma_j + \delta_t\Big] \times \epsilon_{ijt} \end{aligned} \tag{3}$$

$$\begin{aligned} FDI_{ijt} = \exp\Big[&\beta_0 + \beta_1 \times (Treat_{ij} \times Post_t) + \beta_2 \times Treat_{ij} + \beta_3 \times Post_t + \alpha' X^2_{ijt} \\ &+ \delta_{it} + \delta_{jt}\Big] \times \epsilon_{ijt} \end{aligned} \tag{4}$$

$$FDI_{ijt} = \exp\Big[\beta_0 + \beta_1 \times (Treat_{ij} \times Post_t) + \alpha' X^3_{ijt} + \gamma_{ij} + \delta_{it} + \delta_{jt}\Big] \times \epsilon_{ijt} \tag{5}$$

The first specification is a simple difference-in-differences model with no control variables, where $Treat_{ij}$ is our measure of treatment for each dyad and $Post_t$ is an indicator variable for whether or not the observation occurs after 1999 (i.e. after the Maffezini v. Spain decision issued in January of 2000) (simple DiD). This model serves as a baseline model to estimate the effect of treatment without controlling for other covariates or fixed effects. The second specification adds in controls for whether or not a BIT is active and the standard controls from the trade literature, all captured in the X^1_{ijt} vector (simple gravity). The standard controls include: GDP in the origin and destination countries, distance between countries, and indicator variables for common border, common language, prior colonial affiliation, common currency, the existence of a free trade agreement (FTA), whether both countries in a dyad are members of the General Agreement on Tariffs and Trade (GATT), common market, and common customs union. The third specification adds in controls for origin, destination and year fixed effects (captured by γ_i, γ_j and δ_t respectively) (simple FE). The fourth specification includes orig-year and destination-year fixed effects (in lieu of origin, destination, and year fixed effects) captured by δ_{it} and δ_{jt}, respectively (moderate FE). Note that in the fourth specification the vector of control variables is X^2_{ijt} which is similar to X^1_{ijt} but now excludes origin and destination GDP, since GDP is soaked up by the fixed effects. The fifth specification includes origin-destination, origin-year and destination-year fixed effects, captured by γ_{ij}, δ_{it} and δ_{jt} respectively (full FE). In this specification the vector of control variables is X^3_{ijt} which, similar to X^2_{ijt}, is a truncated version of X^1_{ijt} that excludes controls that are soaked up by the fixed effects in the fifth specification. These controls include whether or not a BIT is active between the dyad in a particular year, the existence of a free trade agreement (FTA), whether both countries in a dyad are members of the

[15] CEPII (2016)

[16] This methodology is discussed in more detail in Appendix B.

General Agreement on Tariffs and Trade (GATT), common market, and common customs union. They are referred to in the regression tables below as measures of "market integration." Note that the *Treat* and *Post* variables are also soaked up by the fixed effects in the fifth model specification.

The first three models are incorrectly specified because they fail to account for multi-lateral resistance terms, i.e. how a host and source economies' other bilateral relationships impact investment flows within the dyad. Including more comprehensive fixed effects, as in specifications 4 and 5, correct for this misspecification and including these fixed effects is a common best practice in the trade literature.[17] The first three model specifications are only used to demonstrate the evolution of the coefficient estimate as I add in controls and fixed effects. Note that specification 5 (full FE) is the most comprehensive model and is less likely to suffer from omitted variable bias relative to specification 4 (moderate FE). However, as discussed in Baldwin and Taglioni (2006), one potential drawback of the full FE specification is that including origin-destination fixed effects will eliminate information in the cross-section variation, so all identification will now come from time variation in the variables. In this particular case, all treated dyads (88 percent of all dyads) will have time variation in their treatment and so the elimination of cross-section variation is less of a concern.

Seventy-six percent of the observations in our panel dataset report zero bilateral FDI flows. These zero observations pose two challenges. The first is that running a regression on a log-linearization of the above model specifications will drop a lot of observations and perhaps bias our estimates. The second challenge is verifying whether or not our treatment and control groups have parallel trends prior to treatment, i.e. whether or not the standard identification assumption in difference-in-differences is satisfied. To get around the first challenge I can implement the regression using Poisson-Pseudo-Maximum-Likelihood estimation, a standard best practice in structural gravity models.[18]

To address the second challenge I can eliminate some of the zero observations by summing bilateral FDI flows across three year increments (1991-1993, 1994-1996, etc.) and then trimming the dataset to only include dyads with a history of positive FDI prior to 2000. This trimmed dataset includes 762 of the original 2,527 dyads (30 percent); of these 762 dyads, 671 (88 percent) are in the treatment group and 91 (12 percent) are in the control group. Approximately 64 percent of the observations still report zero bilateral FDI, but enough of the zeros are eliminated to consider the parallel trends condition. Using this trimmed dataset I check for parallel trends on the log of FDI for six different snapshots of the distribution for the treatment and control groups: the 50th, 60th, 65th, 75th, 90th and 95th percentiles.[19] These distribution snapshots appear in Figure 2.

[17] See Baldwin and Taglioni (2006) for a careful discussion of correct model specifications in gravity models.

[18] See Silva & Tenreyro (2006). Note that the fixed effects estimation in our fourth and fifth specification is implemented using a new structural gravity ppml command introduced in Zylkin (2017).

[19] Note that the 50th, 60th and 65th percentile values of bilateral FDI flows in one or both of the distributions in 1993 is zero, resulting in the missing data points in the first three graphs.

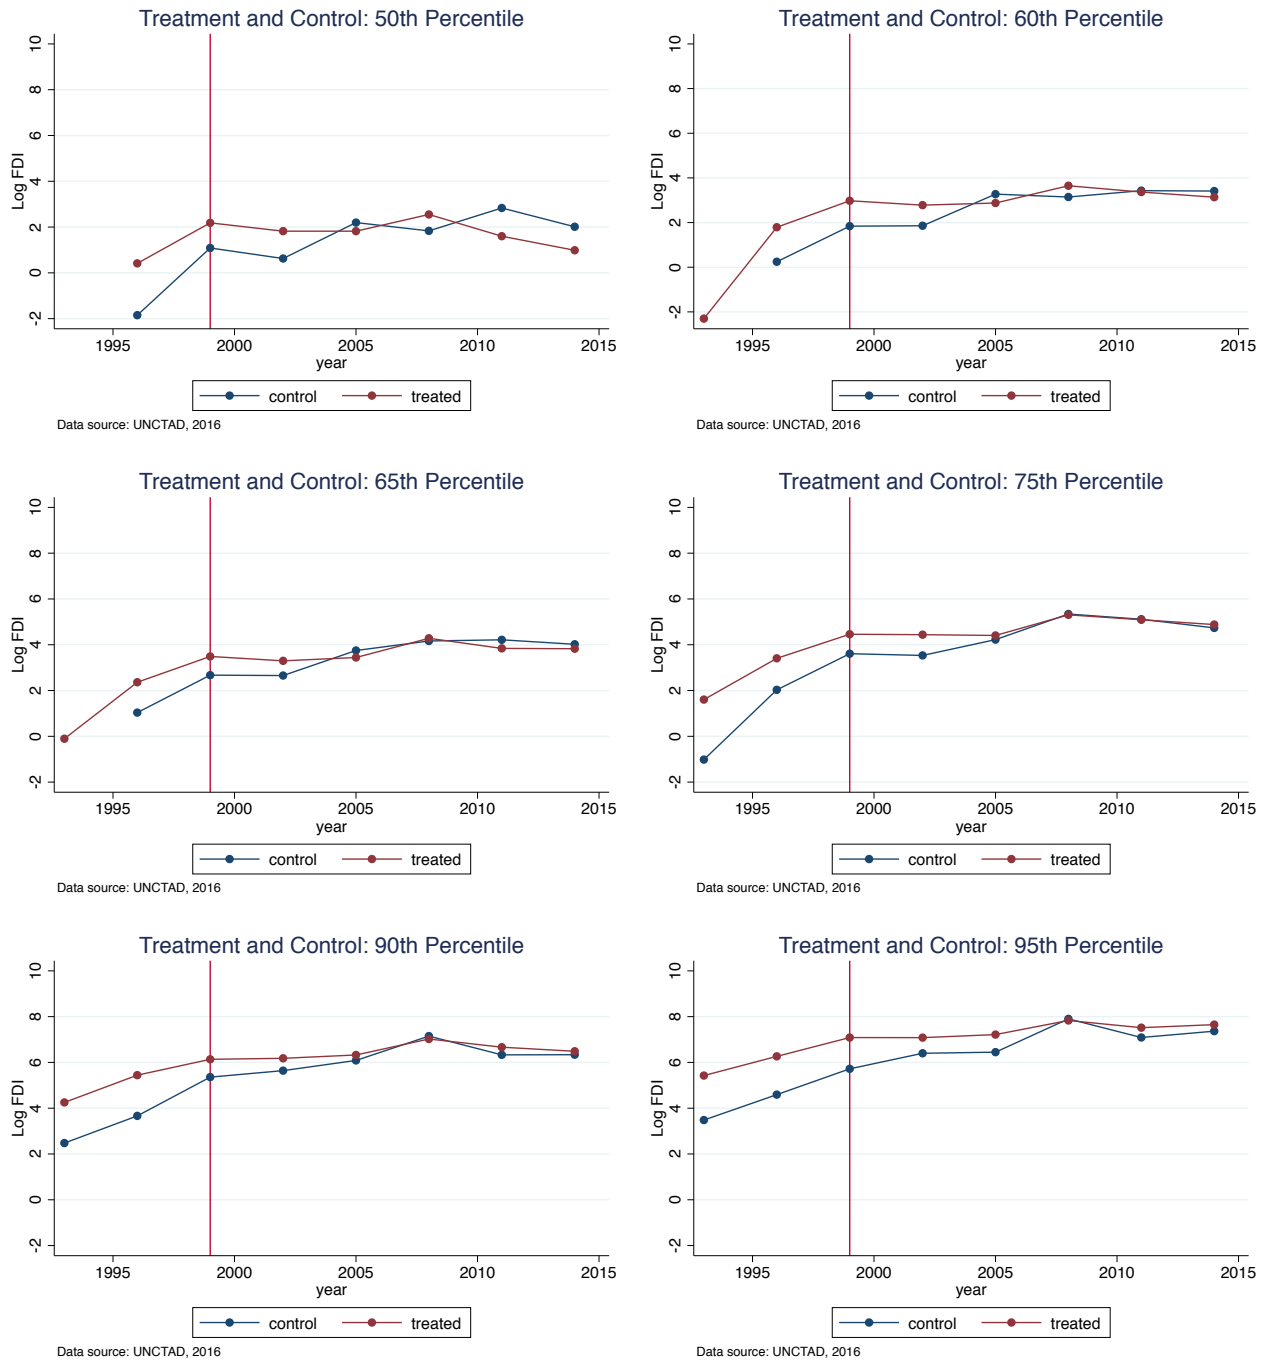

Figure 2: Verifying Parallel Trends

These snapshots of the trends in the distributions provide approximate evidence of parallel trends prior to treatment and a common trend pattern across the distributions: for each of the percentiles, the treatment group always starts with higher FDI, realizes a decline in FDI after treatment, and FDI approximately converges to the same amount of FDI for the control group after treatment. After verifying parallel trends I am now ready to implement the difference-in-differences regressions.

4.3 The Findings

The first set of regressions are implemented using all five model specifications to document the evolution of the coefficient estimates as controls and fixed effects are added to the model. These regressions use the simplest measure of treatment (coded as an indicator variable), annual panel data, and all dyads in the treatment and control groups. The coefficient estimates are reported in Table 4.

The first three model specifications produce negative and statistically significant coefficient estimates for the treatment variable ($TREAT \times POST$). Recall that model specifications 1, 2 and 3 are incorrectly specified, since they fail to control for the multi-lateral resistance terms. In model specifications 4 and 5, adding higher level fixed effects to control for multilateral resistence terms reduces the magnitude of the coefficient estimates, which are no longer statistically different than zero.

Table 4: Comparing Model Specifications: Treatment as an indicator, Annual Data, All Dyads

	simple DiD	simple gravity	simple FE	moderate FE	full FE
treat × post	-0.73**	-0.63**	-0.70***	-0.10	-0.095
	(0.23)	(0.23)	(0.21)	(0.22)	(0.20)
treat	1.26**	0.55	0.86*	0.27	
	(0.44)	(0.42)	(0.37)	(0.36)	
post	2.38***	0.85***	1.70***		
	(0.20)	(0.22)	(0.51)		
BIT		0.87***	0.18	-0.41	-0.27^{+}
		(0.23)	(0.21)	(0.26)	(0.16)
standard gravity		x	x	x	
market integration		x	x	x	x
simple FE			x		
moderate FE				x	
full FE					x
Observations	62425	60441	50658	34141	27369
R^2	0.006	0.108	0.424	0.670	0.850

Standard errors in parentheses. All standard errors are clustered at the origin-destination level.

Note that standard gravity covariates in the moderate FE model exclude origin and destination GDP.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The next set of regressions use the correctly specified models (moderate FE and full FE) to compare the initial coefficient estimates in these models to coefficient estimates where treatment is allowed to vary across the treatment group. This is done implementing regressions across three measures of treatment: (1) indicator, (2) the normalized count of the number of treated provisions(count), and (3) the first continuous measure of treatment developed in Section 3 (Cont. 1). These estimates are reported in Table 5. Allowing for treatment to vary across the treatment group increases the size of the coefficient estimates in both models. The estimate in the moderate FE model using the count measure of treatment is statistically

significant at the 90 percent level. This shift in the size and significance of the estimates indicates that accounting for the amount of treatment is important when estimating the effect of treatment.

Table 5: Comparing Treatment Measures: Moderate and Full FE Models, Annual Data, All Dyads

	Moderate FE Indicator	Moderate FE Count	Moderate FE Cont. 1	Full FE Indicator	Full FE Count	Full FE Cont. 1
treat × post	-0.10	-0.58^{+}	-0.21	-0.095	-0.49	-0.33
	(0.22)	(0.33)	(0.30)	(0.20)	(0.31)	(0.25)
treat	0.27	0.39	0.41^{+}			
	(0.36)	(0.33)	(0.24)			
BIT	-0.41	-0.44^{+}	-0.44^{+}	-0.27^{+}	-0.31^{+}	-0.31^{+}
	(0.26)	(0.26)	(0.26)	(0.16)	(0.16)	(0.16)
standard gravity	x	x	x			
market integration	x	x	x	x	x	x
moderate FE	x	x	x			
full FE				x	x	x
Observations	34141	34141	34141	27369	27369	27369
R^2	0.670	0.668	0.670	0.850	0.850	0.850

Standard errors in parentheses. All standard errors are clustered at the origin-destination level.

Note that standard gravity covariates in the moderate FE model exclude origin and destination GDP.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

There are two additional considerations when estimating the coefficients. The first is the lumpy nature of investment: foreign firms may not choose to make investments every year, and this fluctuation in firm investment decisions may affect the estimates. The second consideration is whether the estimates of the effect of treatment are different if dyads that have no prior history of FDI are excluded from the regression. Since the preliminary estimates are suggesting a possible negative effect of treatment, including dyads with no prior FDI will, by construction, dampen the effect of treatment, since it is not possible for FDI flows to decline if there is no prior FDI.

To take these considerations into account I run another set of regressions using the moderate and full FE model specifications and the continuous measure of treatment (Cont. 1) using four different variations of the panel dataset: (1) annual data and all dyads, (2) annual data and dyads with prior FDI, (3) three-year data and all dyads, and (4) three-year data and dyads with prior FDI. These coefficient estimates are reported in Table 6. As expected, removing dyads with no prior history of FDI increases the size of the coefficient estimate. Note that summing the data over three year increments to smooth out investment also increases the size of the coefficient estimates. None of the coefficient estimates in the moderate FE model are statistically significant. However, estimates in the full FE model using three-year data for all dyads and dyads with prior FDI as well as using the annual data for dyads with prior FDI, are all statistically significant. The largest effect is estimated when I sum over three-year increments and only include dyads with prior FDI.

Table 6: Comparing Estimates: Annual v. 3 Year Data and All Dyads v. Dyads with prior FDI

	Mod FE 1Y ALL	Mod FE 1Y Prior	Mod FE 3Y ALL	Mod FE 3Y Prior	Full FE 1Y ALL	Full FE 1Y Prior	Full FE 3Y ALL	Full FE 3Y Prior
treat x post	-0.21	-0.41	-0.29	-0.43	-0.33	-0.46^{+}	-0.49^{+}	-0.57*
	(0.30)	(0.26)	(0.34)	(0.30)	(0.25)	(0.24)	(0.28)	(0.27)
treat	0.41^{+}	0.28	0.46^{+}	0.33				
	(0.24)	(0.25)	(0.25)	(0.26)				
BIT	-0.44^{+}	-0.50*	-0.42^{+}	-0.42^{+}	-0.31^{+}	-0.52**	-0.048	-0.24
	(0.26)	(0.23)	(0.25)	(0.25)	(0.16)	(0.16)	(0.19)	(0.20)
standard gravity	x	x	x	x				
market integration	x	x	x	x	x	x	x	x
moderate FE	x	x	x	x				
full FE					x	x	x	x
Observations	34141	14546	12652	5190	27369	14546	9754	5182
R^2	0.670	0.776	0.707	0.806	0.850	0.888	0.916	0.926

Standard errors in parentheses. All standard errors are clustered at the origin-destination level.

Note that standard gravity covariates in the moderate FE model exclude origin and destination GDP.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

One final consideration is whether these coefficient estimates are sensitive to the subjective decisions I made when constructing my continuous measure of treatment. I check this by running a final set of regressions on my full FE model specification, using three-year aggregated data and only including dyads with prior FDI.[20] Regressions are implemented on all eight measures of treatment: treatment as an indicator, treatment as a normalized count, the continuous measure of treatment I have used thus far and the five other continuous measures of treatment developed in Appendix B. In these regressions I use a control for BIT strength that corresponds with the measure of treatment. I also run another set of regressions that include separate controls for the original strength of each of the five sets of BIT provisions: definitions (DEF), investor protection (PRO), host state obligations (OBL), treaty duration (DUR), and legal remedies (REM). These coefficient estimates are reported in Table 7.

The method used to construct the five new continuous measures of treatment is discussed in Appendix B. Note that measures 1 and 3 are the most conservative measures and are constructed using smaller step sizes in adjustments between different variations of a treaty provision relative to measures 2 and 4. As a result, the treatment distributions of these measures have more clumping towards zero.[21] Measures 5 and 6 are more crude measures, using a more rudimentary count measure of treatment. These measures also produce a treatment distribution that is more spread out. The coefficient estimates in Table 7 indicate that when the size of the step increases and treatment is amplified for dyads in the center (as is the case for measures 2 and 4), the estimate of the effect of treatment is dampened. When a more rudimentary approach to treatment is used (as is the case for measures 5 and 6), the treatment effect is still large but the standard errors have increased.

[20] This combination of model, time increment and sample produces the largest negative coefficient estimate, and so I use this as a baseline to compare estimates using the different measures of treatment.

[21] These distributions are included in Appendix B.

Table 7: Comparing Measures of Treatment: 3Y Data, Dyads with prior FDI

	ind	count	cont. 1	cont. 2	cont. 3	cont. 4	cont. 5	cont. 6	cont. 1	cont. 2	cont. 3	cont. 4	cont. 5/6
treat x post	0.034	-0.87*	-0.57*	-0.34	-0.62*	-0.33	-0.61	-0.61	-0.55*	-0.28	-0.24	-0.64	-0.64
	(0.22)	(0.34)	(0.27)	(0.25)	(0.28)	(0.25)	(0.40)	(0.40)	(0.27)	(0.26)	(0.25)	(0.42)	(0.42)
BIT	-0.19	-0.25	-0.30	-0.32	-0.33	-0.32	-0.36	-0.36					
	(0.20)	(0.20)	(0.24)	(0.27)	(0.27)	(0.28)	(0.31)	(0.30)					
DEF									1.22	1.57	0.93	0.84	0.84
									(3.11)	(1.86)	(1.70)	(1.54)	(1.54)
PRO									-4.00+	-4.27+	-4.84*	-3.03+	-3.03+
									(2.41)	(2.41)	(2.11)	(1.74)	(1.74)
OBL									0.25	0.32	-0.33	-1.65	-1.65
									(1.16)	(1.11)	(1.14)	(1.64)	(1.64)
DUR									0.19	0.25	0.97	0.84	0.84
									(3.37)	(1.73)	(1.53)	(1.09)	(1.09)
REM									1.93**	1.59**	1.69***	1.98+	1.98+
									(0.71)	(0.59)	(0.51)	(1.18)	(1.18)
market integration	x	x	x	x	x	x	x	x	x	x	x	x	x
full FE	x	x	x	x	x	x	x	x	x	x	x	x	x
Observations	5182	5182	5182	5182	5182	5182	5182	5182	5182	5182	5182	5182	5182
R^2	0.925	0.926	0.926	0.925	0.926	0.925	0.925	0.925	0.926	0.926	0.926	0.926	0.926

Standard errors in parentheses. All standard errors are clustered at the origin-destination level.

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Finally, to check whether or not the coefficient estimates are capturing a pre-trend decline in FDI in treated dyads, I can run the full FE regression of bilateral FDI on a series of treatment variables interacted with year dummies and then plot the year coefficients of the interacted terms. Two coefficient event steadies are presented below, estimated using the full FE model specification, the primary continuous measure of treatment and three-year aggregate data. The first event study includes all dyads and the second event study only includes dyads with prior FDI.

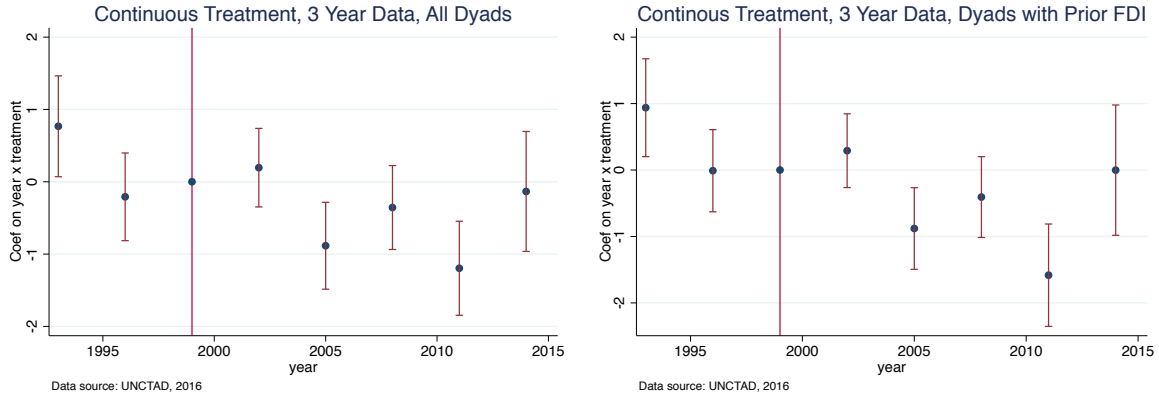

Figure 3: Coefficients Event Study

These event studies collectively suggest that treatment by the Maffezini v. Spain decision does cause a decline in investment, although the effect of the treatment appears to be delayed. This lag in the

effect may be caused by a delay in a host country's understanding of how the Maffezini v. Spain decision modified their risk of arbitration. It may also indicate that policy tools available to host economies to counteract the effect of the decision were not immediately available (i.e. host governments may have already signed contracts with foreign investors and the cost of terminating the contracts exceeded the cost in terms of new arbitration risk that would arise from honoring the contract).

Essentially all coefficient estimates in this section are negative. Together these estimates provide robust evidence that endowing foreign investors with stronger legal remedies does not cause an increase in FDI. At the same time, many of the coefficient estimates from the preferred model specifications using more sophisticated measures of treatment are large, negative, and statistically significant. However, these coefficient estimates are sensitive to subjective decisions made in the construction of treatment measures and to whether or not I exclude dyads with no prior FDI. Together these estimates provide suggestive evidence that imposing stronger legal remedies on a host economy may lead to a decline in foreign investment as host economies react to the increase in the implicit cost of hosting foreign capital caused by an increased exposure to arbitration and tighter constraints on its ability to regulate foreign capital.

To give some idea of the magnitude of the negative effect consider the coefficient estimate produced by the first and third continuous treatment measures in the model that includes separate controls for the different provision categories in a BIT (-0.55 and -0.24, respectively). Note that one standard deviation in the treatment distribution is approximately 0.25. Taking the estimates at face value implies that exposure to treatment equal to one standard deviation in the treatment distribution, will, on average, cause a decline in foreign investment ranging from 6 to 14.25 percent with 95 percent confidence intervals ranging from 4.1 to 27.5 percent (under measure 1) and -6.25 to 18.25 percent (under measure 3).

5 Conclusion

This paper considers whether including stronger legal remedies in an investment treaty causes an increase in foreign investment. To answer this question I identify the Maffezini v. Spain decision as a natural experiment that endowed some investors with stronger legal remedies through an unanticipated application of the MFN protection. I also introduce the UNCTAD IIA Database, built by me in partnership with UNCTAD, on the provision level detail of more than 2,500 BITs. Using this database and natural experiment I find robust evidence that stronger legal remedies in a BIT do not lead to more investment. I also find suggestive evidence that stronger legal remedies imposed on a host economy may actually lead to a decline in investment as host economies respond to an increase in the implicit cost of receiving foreign investment. These results suggest that India's current policy trajectory (and other host economies considering similar policy actions) to negotiate new BITs with weaker legal remedies will likely not have a negative affect on its ability to attract new foreign investment. The evidence also suggests that negotiating new BITs with weaker legal remedies may actual lead to more foreign investment as host economies become more likely to accept and pursue foreign investment in response to a reduction in the implicit cost of receiving FDI.

References

[1] Argentina Spain BIT (1991). http://investmentpolicyhub.unctad.org/Download/TreatyFile/119. Accessed 1 December 2017.

[2] Baldwin, R. and D. Taglioni (2007). "Trade Effects of the Euro: A comparison of estimators", *Journal of Economic Integration*, Vol. 22, No. 4, pp. 780-818.

[3] Beck, T. and R. Levine (2005). "Legal Institutions and Financial Development", *Handbook of New Institutional Economics*, The Netherlands: Springer, pp. 251-280.

[4] Berger, A., M. Busse, P. Nunnenkamp and M. Roy (2013). "Do Trade and Investment Agreements Lead to More FDI? Accounting for Key Provisions Inside the Black Box", *International Economics and Economic Policy*, Vol. 10, No. 2, pp. 247-275.

[5] Busse, M., J. Koniger and P. Nunnenkamp (2010). "FDI Promotion through Bilateral Investment Treaties: More than a BIT?", *Review of World Economics*, Vol. 146, No. 1, pp. 147-177.

[6] Buthe, T. and H. Milner (2009). "Bilateral Investment Treaties and Foreign Direct Investment: A Political Analysis", *The Effect of Treaties on Foreign Direct Investment: Bilateral Investment Treaties, Double Taxation Treaties, and Investment Flows*, New York: Oxford University Press, pp. 171-224.

[7] CEPII (2016). CEPII Gravity Database. http://www.cepii.fr/cepii/en/bdd_modele/bdd.asp. Accessed 5 February 2017.

[8] Hornbeck, S. (1909). "The Most-Favored-Nation Clause", *American Journal of International Law*, Vol. 3, No. 4, pp. 797-827.

[9] India Model BIT (2016). http://investmentpolicyhub.unctad.org/Download/TreatyFile/3560. Accessed 21 November 2017.

[10] Maffezini v. Spain (2000). https://www.italaw.com/sites/default/files/case-documents/ita0479.pdf. Accessed 1 December 2017.

[11] Rubin, P. (2005). "Legal Systems as Frameworks for Market Exchanges", *Handbook of New Institutional Economics*, The Netherlands: Springer, pp. 205-228.

[12] Silva, S. and S. Tenreyro (2006). "The Log of Gravity", *The Review of Economics and Statistics*, Vol. 88 No. 4, pp. 641-658.

[13] Vesel, S. (2007) "Clearing a Path through a Tangled Jurisprudence: Most-Favored-Nation in Bilateral Investment Treaties," *Yale Journal of International Law*, Vol. 32, No. 1, pp. 125-190.

[14] UNCTAD (2014). "The Impact of International Investment Agreements on Foreign Direct Investment: An Overview of Empirical Studies 1998-2014", *IIA Issues Note*, New York and Geneva: United Nations.

[15] UNCTAD (2016a). UNCTAD International Investment Agreements Database. investmentpolicyhub.unctad.org. Accessed 9 December 2016.

[16] UNCTAD (2016b). UNCTAD Investor State Dispute Settlement Database. investmentpolicyhub.unctad.org. Accessed 9 December 2016.

[17] Zylkin, T. (2017). "PPML_PANEL_SG: Stata module to estimate 'structural gravity' models via Poisson PML". http://EconPapers.repec.org/RePEc:boc:bocode:-s458249. Accessed 27 July 2017.

A BIT Provisions and Summary Statistics

This Appendix provides a summary of BIT provisions as well as summary statistics of the share of BITs that include a particular variation of a provision.

A.1 Defining investor/investment

All BITs include a definition of who and/or what is protected by the BIT. This definition typically includes a definition of investment and/or a definition of investor.

The broadest formulation of the definition of investment is stated with no qualifications. The definition may be limited in two ways. The first is by excluding specific types of assets from the definition. Examples of excluded assets include portfolio investments (i.e. shares in a company), commercial transactions and non-business real-estate investment. The second way an investment definition may be limited is by including criteria that an investment must satisfy in order to qualify as an investment under the definition in the treaty. These criteria may include a commitment of capital and the assumption of risk or may only include investments that are made in accordance with the laws of the host country.

The definition of investor typically includes a definition for natural persons and a definition for legal entities. The broadest variation of the natural person definition will include both permanent residents and persons with dual nationality. The broadest variation of the legal entity definition will not include criteria the entity must satisfy in order to qualify as a defined investor. Limiting criteria may include the existence of a business activity and/or a requirement of ownership and control over the investment.

The Definitions section of a BIT my be further limited by excluding particular policy actions from the scope of the treaty. Common scope exclusions include taxation, subsidies and grants and government procurement. These limitations mean that a country will be able to pursue these policy actions (i.e. taxation) without risk of being accused of violating the BIT (i.e. no expropriation of an investment).

Table A1: Definitions in Active BITs in 2000

Provision	BIT count	BIT share
Investment Defined	1,239	97.71
Limit 1: excluding portfolio	2	0.16
Limit 2: excluding other assets	26	2.05
Limit 3: characteristics of investment	6	0.47
Limit 4: in accordance w/ state laws	768	61.99
Limit 5: closed list of investment	36	2.84
Investor Defined	1,204	94.95
Natural Person	164	12.93
NP Limit 1: Includes Permanent Residents	127	10.02
NP Limit 2: Excludes Double Nationality	52	4.10
LE Limit 1:Business activity requirement	157	12.38
LE Limit 2: Ownership and control defined	150	11.83
Limit 1: excludes taxation	86	6.79
Limit 2: excludes subsidies/grants	30	2.37
Limit 3: excludes gov. procurement	21	1.66
Limit 4: excludes other subject matter	71	5.60

A.2 Investor Protections

The primary protections in a BIT are (1) protection from discrimination and (2) protection from expropriation.

A BIT may protect investors from discrimination relative to domestic investors through a "National Treatment" (NT) clause. This clause guarantees that investors protected by the BIT will be treated at least as well by the host government as domestic investors. Similarly, a BIT may protect investors from discrimination relative to other foreign investors through a "Most-Favored-Nation" (MFN) clause. This clause guarantees that protected investors will be treated at least as well by the host government as it treats any other foreign investor. As discussed in this paper, the MFN protection has turned out to be an important provision in the BIT network.

The second BIT protection is protection from expropriation, which protects investors from having their assets taken by the host government without due process of law and without prompt and just compensation. There are two variations of the expropriation protection. The first includes a reference to indirect expropriation. The second does not. A reference to indirect expropriation extends the protection to include protection from actions by the host government that may not involve the actual seizure of assets, but still result in depriving the investor of either the value or the control of their assets (i.e. requiring that a majority holding in the asset be transferred to a domestic investor). The expropriation protection may also be limited by including a carve-out for general regulatory measures. This means that general regulatory measures implemented by the host government that may undermine the value of an investment do not qualify as an act of expropriation.

Table A2: Investor Protections in Active BITs in 2000

Provision	BIT count	BIT share
Non-Discrimination: NT	1,001	78.94
Non-Discrimination: MFN	1,244	98.11
Expropriation	1,266	99.84
Limit 1: indirect expropriation defined	33	2.60
Limit 2: carve-out for regulatory measures	73	5.76

A.3 Host State Obligations

Host State obligations in a BIT may include: (1) compensation for losses, (2) prohibition on performance requirements, (3) allowing the employees of the investor to enter and work in the country, (4) allowing investors autonomy over their choice of senior managers and/or board members, and (5) allowing the free transfer of assets.

Compensation for losses are paid out by the host government in specified circumstances such as armed conflict, civil unrest or natural disasters. This has two common variations. The first is a relative right to compensation. The second is an absolute right to compensation. The relative right to compensation may be relative to either domestic investors (NT), other foreign investors (MFN), or both. An absolute right to compensation, when included in this provision, is typically for losses directly caused by an action of the host government.

The second obligation is to not impose performance requirements on an investor. Performance re-

quirements often involve different restrictive measures designed to ensure that the foreign investment benefits the host economy in a particular way. For example, a host government may require a foreign investor to use a certain amount of domestic inputs when producing a product or require that a certain share of output be sold in the domestic market. BITs often include a host state obligation to not impose particular performance requirements. There are two variations of this obligation. The first is the incorporation of the provisions of the WTO Agreement on Trade-Related Investment Measures (TRIMs) which lists various performance requirements that are prohibited. The second is an ad hoc listing of various prohibited performance requirements.

The third obligation is to allow the employees of the investor to enter and work in the country. This obligation deals with visa issues and is typically limited to not exceed what current domestic legislation would allow with respect to the entry and work opportunities of foreign employees.

The fourth obligation is allowing investors to choose their senior managers and/or board members. This obligation is typically absolute (i.e. without reference to current domestic legislation). Thus the host state is obligated to honor senior management and board leader appointments by the investor and must help facilitate the visa process for these individuals.

The fifth obligation is allowing investors to freely transfer their assets without undue delays. This obligation may be subject to two qualifications. The first is a balance of payments (BOP) exception. The second is other listed exceptions. The BOP exception allows countries to deviate from the free transfers obligation if it faces a serious balance of payments difficulty. Other listed exceptions may include cases where an investor has filed for bankruptcy or owes money as a result of a judicial or arbitral award.

Table A3: Host State Obligations in Active BITs in 2000

Provision	BIT count	BIT share
Compensation for losses	1,165	91.88
Relative right: No reference	31	2.44
Relative right: NT only	30	2.37
Relative right: MFN only	258	20.35
Relative right: NT and MFN	846	66.72
Absolute right	338	26.66
Prohibition on performance requirements	71	5.60
Variation 1: No reference	17	1.36
Variation 2: TRIMs	4	0.32
Variation 3: ad hoc list	50	3.94
Entry and sojourn of personnel	464	36.59
Senior managers and board members	86	6.78
Free transfer of assets	1,260	99.37
Qualification 1: BOP	98	7.73
Qualification 2: other	93	7.34

A.4 Duration

The duration of a BIT depends on the following provisions: (1) the actual duration of the BIT, (2) whether a BIT automatically renews, (3) the method for terminating the BIT, and (4) the amount of time the BIT protections and obligations survive after an event of termination. The term of time that the BIT protections are in place may also be modified by (5) limitations to the scope of the BIT.

The duration of a BIT is the length of time the BIT is in force. The period is typically 10 or 15 years but the duration of some BITs may be as long as 20 years. A handful of BITs have an indefinite duration.

Many BITs include an automatic renewal clause that comes into effect after the initial duration of the treaty concludes. There are two common variations of the automatic renewal clause. The first is automatic renewal of an indefinite duration. The second is automatic renewal of a fixed duration. Renewal of indefinite duration is subject to the right of either party to terminate the treaty as specified in the termination clause. Renewal of a fixed duration typically mirrors the initial duration of the treaty (i.e. 10 or 15 years).

Most BITs specify when and how a BIT may be terminated. This typically includes a time component (i.e. after the BIT has been in force for a certain number of years) and a notice component (i.e. by written notice to the other party). The termination provision will also specify the length of the notice period (i.e. one year's written notice).

A survival clause may also be included to extend certain protections and obligations of the treaty for a certain amount of time after an event of termination. A survival clause will typically specify which protections and obligations are extended and for how long.

Additional provisions that can modify the duration of the treaty are: (1) limiting the temporal scope of the investments covered and (2) limiting the temporal scope of the disputes covered. Under the first limitation, BIT protections and obligations only cover investments that happen after the BIT enters into force, cover all investments regardless of whether the investment occurred prior to entry into force, or the BIT does not specify either case. Under the second limitation, the signatories may prevent legal claims arising out of actions or disputes that occur prior to the entry into force of the BIT.

Table A4: Duration Provisions in Active BITs in 2000

Provision	BIT count	BIT share
Duration: <10 years	59	4.73
Duration: 10 years	925	72.95
Duration: 15 years	202	15.93
Duration: 20+ years	50	4.01
Duration: indefinite	33	2.64
Automatic renewal	1,237	97.48
Variation 1: indefinite	781	61.58
Variation 2: fixed term <10	195	15.63
Variation 2: fixed term 10+	261	20.91
Termination	1,242	97.95
Survival Clause: None	33	2.64
Survival Clause <10 years	48	3.85
Survival Clause 10-14 years	677	53.44
Survival Clause 15 years	288	22.73
Survival Clause 20 years	222	17.52
Temporal scope limitation 1: investments post EIF only	90	7.09
Temporal scope limitation 1: investments both pre & post EIF	995	78.41
Temporal scope limitation 1: no specification	161	12.69
Temporal scope limitation 2: disputes post EIF	350	27.58

B Appendix B

Six different methods of calculating a continuous measure of treatment are developed in this paper. All involve, first, ordinally ranking the different variations of each protection, assigning a score to each variation, and then using these scores and a simple function to calculate a measure of the strength of legal remedies before and after treatment. The treatment measure is just the difference of these two measures. This appendix describes the six different methods used to assign scores to each variation and to then combine those scores into a single strength measure. This appendix also describes how these different methods can be applied to other provision categories in a BIT to develop measures of both treaty strength broadly, and provision category strength individually. These measures are used as additional controls as part of the empirical analysis in Section 3.

B.1 Methods 1-4: step rules

The first four methods assign scores to each variation using different step rules:

- Method 1: 0.1/0.9 - The most favorable variant is assigned a score of 1. Each subsequent variant is decreased by 0.1. Limitations and exceptions are scaled by 0.9. Consent withheld is scaled by 0.5.
- Method 2: 0.2/0.8 - The most favorable variant is assigned a score of 1. Each subsequent variant is decreased by 0.2. Limitations and exceptions are scaled by 0.8. Consent withheld is scaled by 0.2.
- Method 3: 0.5/0.9 - The best variant is assigned a score of 1. Each subsequent variant is decreased by a factor of 0.5 divided by the number of variants less 1. For example, if a provision has two variations, the best variant is scored as 1 and the second is scored as 0.5. If a provision has three variations, the best variant is scored as 1, the second is scored as 0.75 and the third is scored as 0.5, etc. Limitations and exceptions are scaled by 0.9. Consent withheld is scaled by 0.5.
- Method 4: 0.5/0.8 - The same as Method 3 except limitations and exceptions are scaled by 0.8 and arbitration consent withheld is scaled by 0.2.

Under methods 1-4 a final score for legal remedies (REM) is calculated by taking a simple average of the core provisions scaled by the score of the limitations and exceptions. Note that these methods for measuring the strength of legal remedies can also be applied to measure the strength of the other four categories of BIT provisions: definitions (DEF), host state obligations (OBL), investor protections (PRO), and duration (DUR). I can then use these measures (in addition to the original legal remedies (REM) measure), in lieu of a simple indicator variable for an active BIT, as additional controls for the strength of the original BIT in our difference-in-differences regressions.[22] It is also possible to take a simple average of these category scores to come up with a single measure of original treaty strength:

$$BIT = \frac{1}{5} * \left(DEF + OBL + PRO + DUR + REM \right) \tag{6}$$

[22] This approach is implemented in Table 7

B.2 Methods 5 and 6: simple sums

The fifth and sixth method simply count how many of the seven provisions in the legal remedies category are the most favorable variant. This measure is normalized by dividing the count by seven. As was the case for the first four methods, this same count method can be used to calculate the strength of the other provision categories.

Methods 5 and 6 only differ in their measure of overall BIT strength. Similar to the first four methods, method 5 calculates an overall strength score by taking a simple average of the category scores. Method 6 is the most rudimentary of the continuous measures, using a normalized count of the scores across all of the provisions in the BIT. Note that, in contrast to the either five methods, weighting across categories is not equal in method 6 and will depend on the number of provisions in each category.

B.3 Maffezini v. Spain and the BIT strength distribution

The same approach used in Section 4 to visualize the change in the distribution of legal remedies after Maffezini v. Spain can also be used to visualize the change in the distribution of the strength of BITs broadly. This is done by estimating the original strength of each BIT and then estimating the modified strength by using the modified strength of the legal remedies section. The distributions of treaty strength prior to treatment and after treatment are presented in Figures B1 and B2, respectively. The distributions of exposure to treatment for each method are presented in Figure B3.

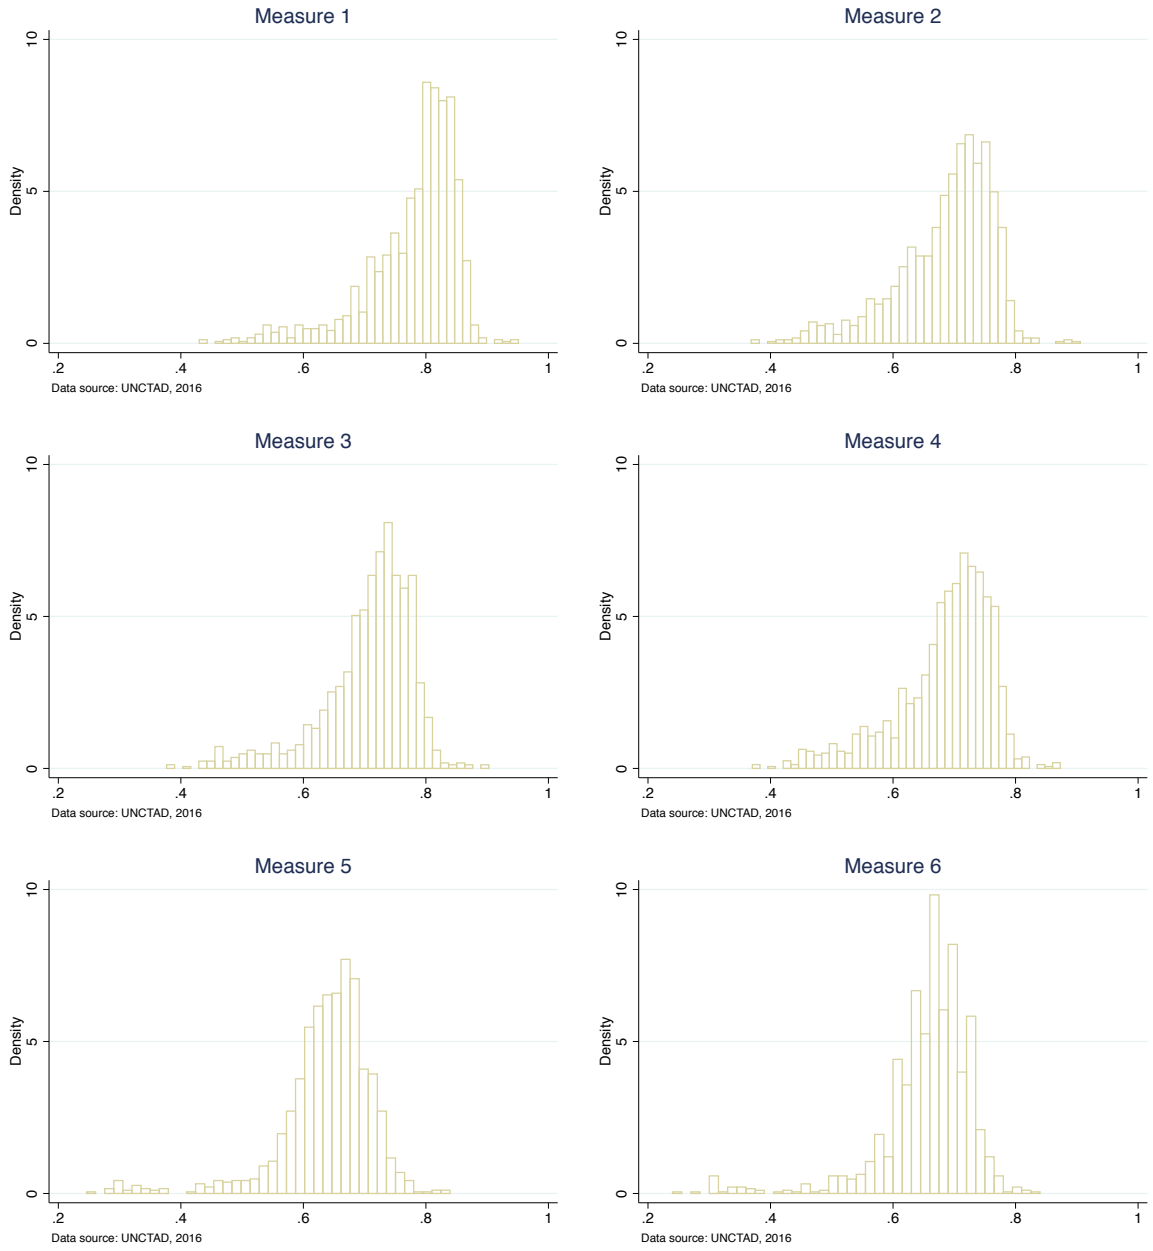

Figure B1: BIT Strength Distribution Prior to M. v. Spain

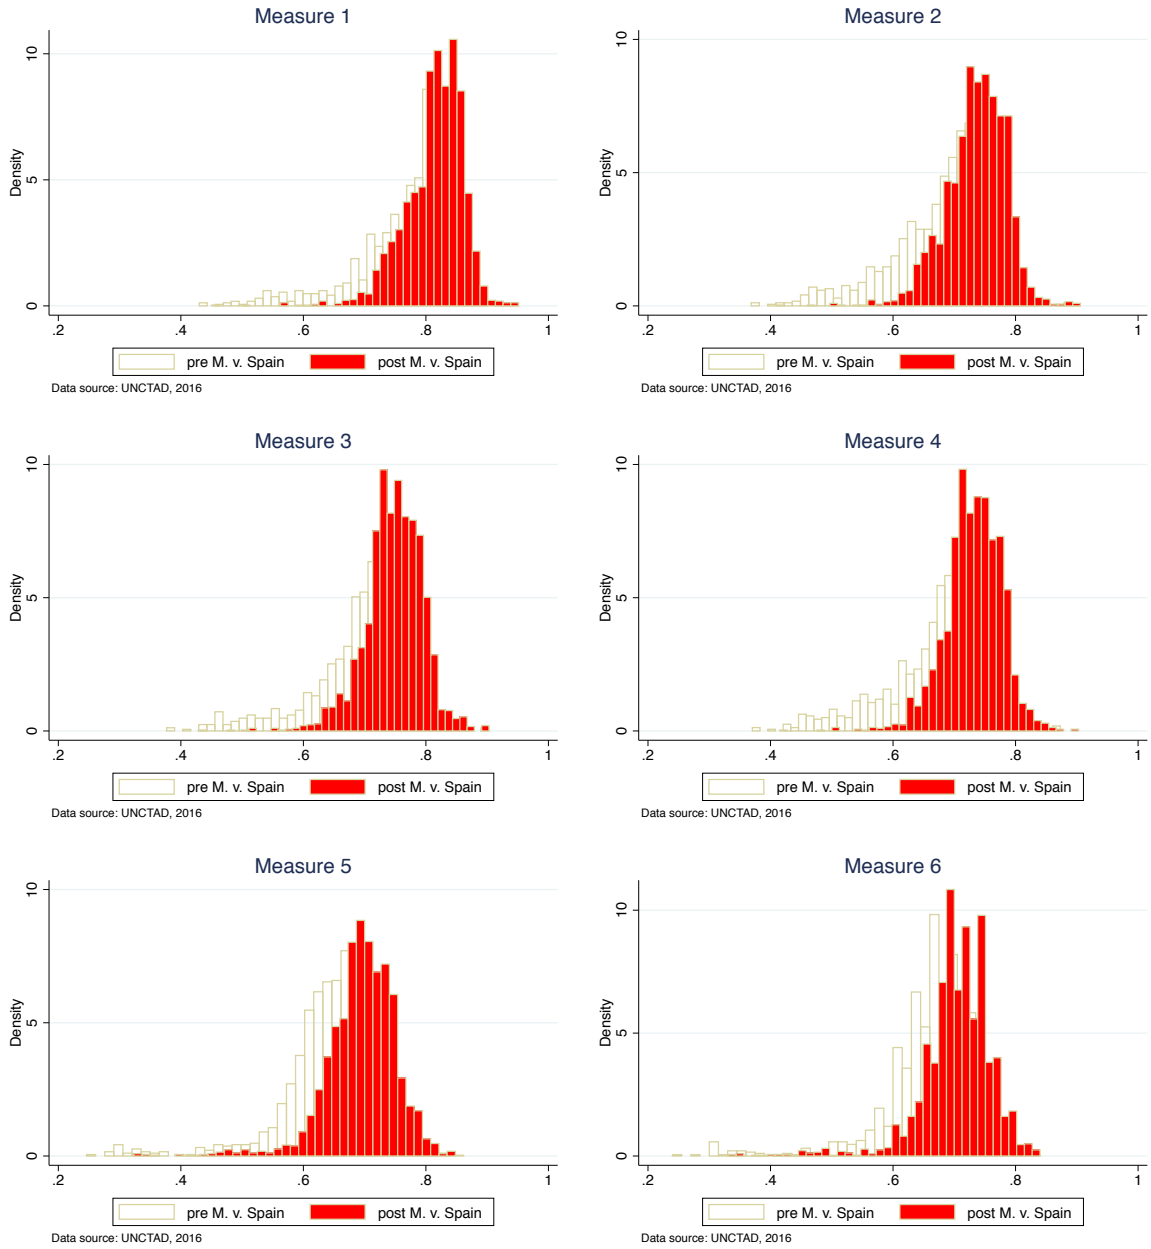

Figure B2: Comparing BIT Strength Pre and Post M. v. Spain

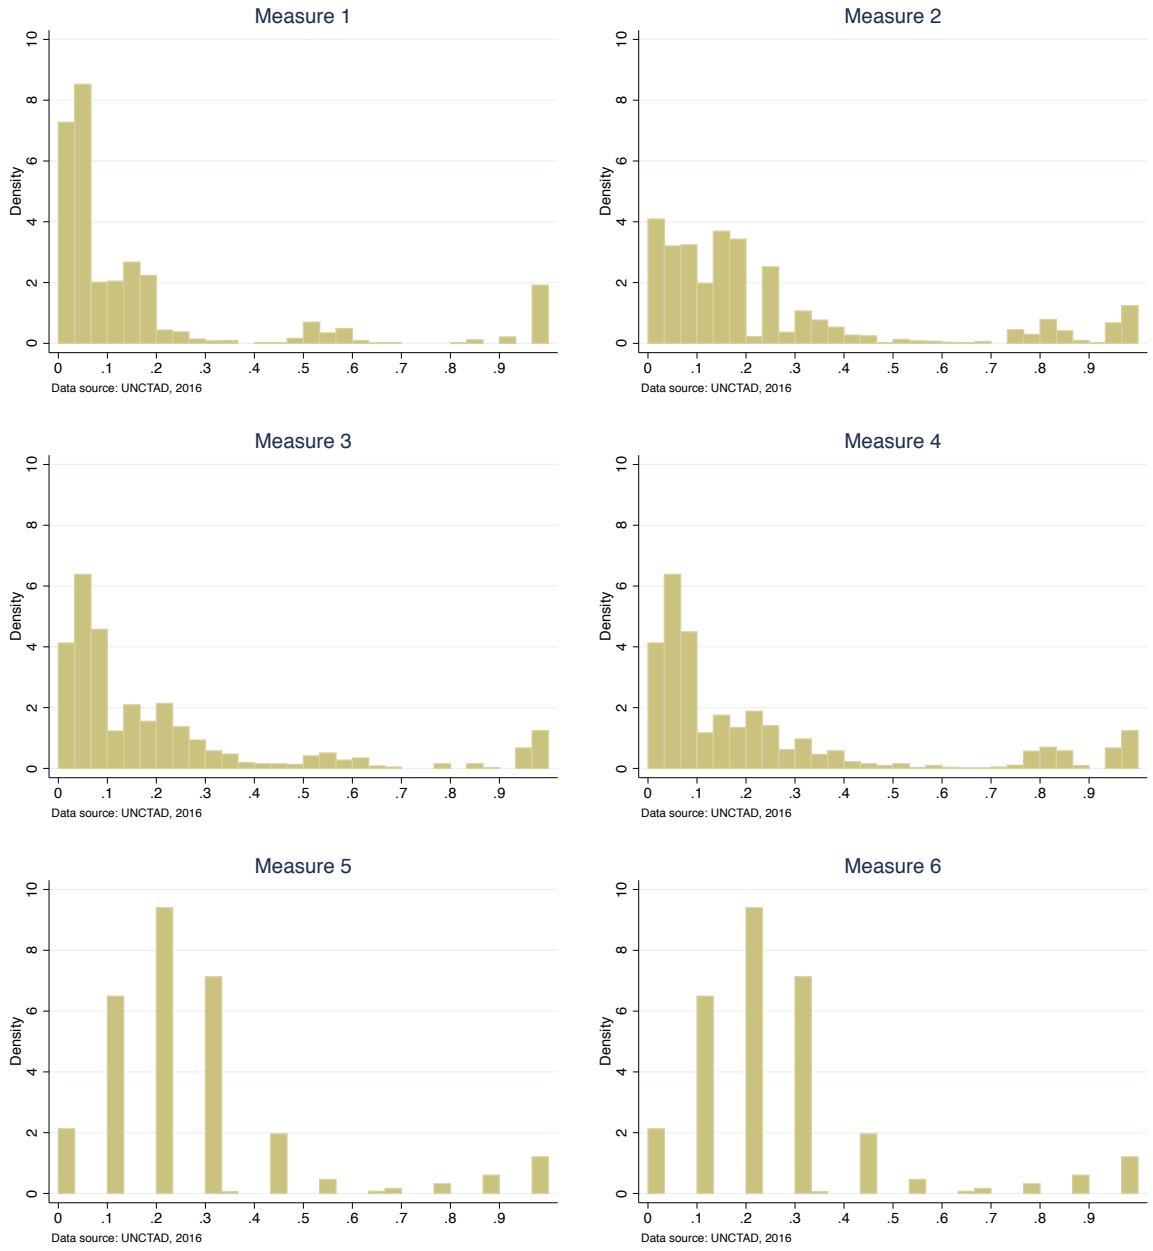

Figure B3: Treatment Distributions