

NBER WORKING PAPER SERIES

LEARNING BY RULING AND TRADE DISPUTES

Giovanni Maggi
Robert W. Staiger

Working Paper 23774
<http://www.nber.org/papers/w23774>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2017

We thank Kala Krishna and Andres Rodriguez-Clare, our discussants at the 2016 Princeton IES Summer Workshop, Nuno Limao, our discussant at the 2016 NBER ITI Summer Institute, and Gene Grossman, our discussant at the 2016 Philadelphia Fed Trade conference, for very helpful and detailed comments on an earlier draft. We are also grateful to David Atkin, Chad Bown, Petros C. Mavroidis, Doug Staiger and participants in seminars at the Board of Governors of the Federal Reserve, CREI, FGV-Rio, FGV-Sao Paulo, Georgetown, Insper, Stanford, Syracuse and the University of Geneva for useful comments and discussions. Junhui Zeng provided outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23774.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Giovanni Maggi and Robert W. Staiger. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning by Ruling and Trade Disputes
Giovanni Maggi and Robert W. Staiger
NBER Working Paper No. 23774
September 2017
JEL No. D02,F13,K33

ABSTRACT

Over the WTO years, the frequency of disputes and court rulings has trended downwards. Such trends are sometimes interpreted as symptoms of a dispute resolution system in decline. In this paper we propose a theory that can explain these trends as a result of judicial learning; thus according to our theory such trends represent good news, not bad news. We then offer evidence that the predictions of our model are consistent with WTO trade dispute data, and we take a first step towards estimating the strength and scope of court learning.

Giovanni Maggi
Department of Economics
Yale University
37 Hillhouse Avenue Rm 27
New Haven, CT 06511
and FGV/EPGE-Rio
and also NBER
giovanni.maggi@yale.edu

Robert W. Staiger
Department of Economics
Dartmouth College
319 Silsby
Hanover, NH 03755
and NBER
rstaiger@dartmouth.edu

1. Introduction

Since the inception of the World Trade Organization (WTO) in 1995, there have been roughly 500 WTO trade disputes. The WTO is endowed with a sophisticated court system, the Dispute Settlement Body (DSB), which adjudicates disputes if governments fail to reach settlement. There is considerable variation in the outcome of these disputes: sometimes governments settle early, sometimes they “fight it out” to a DSB ruling. The stakes of disputes also vary widely across cases: sometimes stakes are small, but sometimes they involve large volumes of trade, thus it is important to understand what determines the initiation of disputes and their outcomes.

There are some interesting patterns in the initiation and resolution of disputes over time. Plot 1 shows the raw numbers of disputes and DSB rulings over the WTO years: the plot suggests a declining trend both in the frequency of disputes and in the frequency of DSB rulings, although in a more pronounced way for disputes than for rulings. The impression from Plot 1 is that countries “fight” less as the institution ages. If anything, Plot 1 understates this trend because the number of WTO members has increased since its founding, and Plot 1 does not control for this. A simple way to control for the expanding WTO membership is to check whether country *dyads* fight less as they “age,” that is, the longer the pair of countries have both been WTO members. Plot 2 shows that indeed country dyads initiate fewer disputes and end up less frequently in court as they age.

One could interpret these declining trends in different ways. One possible interpretation is that the WTO institution is becoming less effective over time, or that governments are losing confidence in it.¹ While this is certainly a possibility, in this paper we propose a theory that can explain the declining trends in disputes and rulings as a result of institutional – or “judicial” – learning. According to our theory, these trends represent good news, not bad news.

Judicial learning may occur, for example, because the court learns through past rulings to use and interpret data and to make more effective and timely use of rigorous economic reasoning in arriving at its rulings. Or it may occur because the court learns to better interpret the legal nuances of the WTO contract and thereby learns to better translate the contract into the intent of the contracting parties. Or judicial learning might take the form of a reduction in the verification cost in a setting of costly state verification, or simply the cost of issuing a ruling

¹In the context of the General Agreement on Tariffs and Trade (GATT), which was the precursor to the WTO, Hudec (1993, pp. 11-15) associates a declining use of the GATT dispute settlement system with a decline in the effectiveness of the system as perceived by the member governments.

even in a complete-information setting. We refer to these possible forms of judicial learning as “learning by ruling,” and we describe the impacts of such learning as raising the “quality” (or “efficiency”) of the court. All these possibilities then amount to a court whose quality rises as it learns from past rulings, where court quality naturally translates into a higher expected joint payoff for governments under a court ruling. Our theory is built around this basic structure.

The importance of judicial learning has been emphasized by many legal scholars, although typically in the context of domestic legal systems, not international institutions. For example, an interesting informal and personal account of this importance can be found in former Supreme Court Justice John Paul Stevens’ discussion of learning on the job (Stevens, 2006). After describing the differences of opinion between Justices Holmes and Brandeis in an important legal ruling that would determine the basis for regulatory takings under U.S. law, Justice Stevens concludes:

“I suspect that Justices Holmes and Brandeis would also agree that learning on the job is essential to the process of judging. At the very least, I know that learning on the bench has been one of the most important and rewarding aspects of my own experience over the last thirty-five years.” (p. 1567)

There are good reasons to believe that judicial learning may also be a phenomenon of first-order importance in the WTO. The WTO is a relatively young institution, and the adjudication of trade disputes is a complex task; so it is reasonable to think that, especially in the early stages of the institution, there is significant learning by the actors involved in the WTO’s judicial system. These actors include the Appellate Body, the Dispute Settlement panels, and quite possibly also the WTO’s Secretariat, a group of experts that plays a key role in the dispute settlement process.²

The broad idea behind our theory is simple: in the presence of learning by ruling, going to court today generates future payoff gains for governments; and to the extent that the marginal gains from court learning are diminishing, the investment value of triggering a court ruling

²As we discuss in more detail later in the paper, the Appellate Body is a standing judicial body, so in this case judges may learn directly from their own experience. The WTO Secretariat has considerable “institutional memory,” so similar statements apply. But also the Dispute Settlement panel, which is a rotating body, may learn from reading previous panel reports, as panel reports are public. Another form of learning that is probably important as well is that governments, as they litigate repeatedly over time, may learn how the court operates and adjudicates cases, thus they may learn to predict more accurately the outcome of a ruling. As we discuss below, intuitively the implications of this type of learning by governments should be largely similar to those of learning-by-ruling, but a rigorous examination would require a non-trivial extension of our model.

decreases over time, thus leading to a decline in the frequency of rulings and disputes. We examine this idea through a formal model, which helps elucidate the conditions under which the above prediction holds.

The core assumptions of our model are the following: (i) in each period, a Home government makes a policy choice and a Foreign government decides whether or not to dispute this choice; (ii) if a dispute is initiated, governments (whose objective functions may include political-economy concerns) bargain “in the shadow of the law,” subject to negotiation costs; (iii) if governments disagree, the court intervenes and issues an imperfect ruling (with the objective of maximizing the governments’ joint payoff); (iv) the efficiency of the court rises as court experience grows, but at a diminishing rate; and (v) governments are “large players” that interact repeatedly, so they internalize the benefits of court learning.

Most existing models of “bargaining in the shadow of the law” explain equilibrium court intervention as bargaining failure due to incomplete information (or overconfidence about the ruling). Our model, by contrast, generates equilibrium court intervention for a different reason: due to learning by ruling, going to court today may imply payoff gains for tomorrow. Importantly, such payoff gains arise also if governments do not go to court tomorrow, because court learning improves the disagreement point for tomorrow’s bargain (an off-equilibrium effect), and arise even if governments do not have a dispute tomorrow, because court learning improves the outcome of the would-be bargain (an “off-off-equilibrium” effect).

We focus first on a static setting, and show that in such a setting there can never be a DSB ruling in equilibrium, but there can be a dispute, and a dispute is more likely when the quality of the DSB is lower. In a dynamic setting, on the other hand, the presence of court learning can give rise to rulings in equilibrium. When we examine how the likelihood of current disputes and rulings depends on court experience (cumulative rulings), we find that the relation is decreasing if governments are patient enough; and even if governments are impatient, these results hold when the stock of cumulative rulings is large enough. The role played by government patience is due to the fact that an increase in court experience has both a “dynamic effect” and a “static effect” that push in opposite directions. For example, in the case of rulings, the dynamic effect of an increase in court experience makes a ruling less likely because the future payoff gain from going to court is diminishing as the court walks down its learning curve;³ but the static effect

³As will become clear, there are two mechanisms that contribute to the diminishing returns from court learning. One is that learning occurs at a diminishing rate, and the other one is more subtle: conditional on there being no dispute, an increase in court quality induces the Home government to select a more efficient

goes in the opposite way because an increase in DSB quality reduces the inefficiency of going to court today; and the dynamic effect dominates when the discount factor is sufficiently high.

Our basic two-country model focuses on the case in which court learning is general in scope, in the sense that a ruling today raises the quality of the court tomorrow regardless of which country is the defendant tomorrow. But it is possible that the scope of learning might be narrower, in that learning could be defendant-specific or complainant-specific or even “directed-dyad-specific” (applying only to future disputes in which the same disputants play the same role in the dispute); and it is also possible that the scope of learning could be broader, in the sense that the effects of learning might spill over to disputes involving third parties. We consider these possibilities in the context of a many-country extension of our basic model. We show in this extended setting that our main results continue to hold, but now the pattern of the impacts of court experience on the likelihood of current rulings is also informative about the scope and kind of court learning that is present.

Returning to the interpretive question we raised at the outset, our model suggests that the frequency of DSB use is not a reliable measure of the effectiveness of the institution. According to our theory, a declining trend in DSB disputes or rulings does not imply that the quality of the institution declines over time, in fact it is a symptom of beneficial learning. However, this is a statement about the *change* in the frequency of DSB use over time. In our model, a lower *level* of this frequency may well be a symptom of lower court quality: if the quality of the court (for given stock of cumulative rulings) is higher, the likelihood of a ruling is higher, because of the static effect described above.

Finally, we explore the empirical content of our theory using WTO trade dispute data. We focus on a key prediction of the model, namely that the likelihood of current disputes and rulings should tend to decrease with the stock of cumulative past rulings.⁴ Our empirical investigation has a dual objective. First, we want to ask whether the above prediction is consistent with the data. And second, to the extent that the answer to this question is affirmative, we want to gauge the empirical importance of learning by ruling and assess its scope and form.

Unlike the existing empirical work on learning by doing for firms where direct measures

policy, and the associated efficiency gain is diminishing due to the concavity of the Pareto frontier.

⁴Our focus on learning ignores another dynamic mechanism that has a similar flavor to learning but has distinct implications and may also be present in the WTO, namely legal precedent. While our formal model abstracts from the issue of legal precedent, we consider this issue (along with several other alternative mechanisms) in the empirical part of the paper, where we argue that in principle the effect of legal precedent could explain some – but not all – of the patterns we observe in the data.

of productivity growth are available (see, for example, Irwin and Klenow, 1994, Clerides et al, 1998, Bernard and Jensen, 1999, Benkard, 2000, Thornton and Thompson, 2001, Kellogg, 2011 and Levitt et al, 2013), we cannot observe directly the quality of the court, so we cannot estimate directly the relationship between court experience and court quality; but we can use the predictions of our model to indirectly gauge the importance of learning by ruling. In particular, our model suggests that a stronger (negative) effect of cumulative rulings on the likelihood of current rulings and disputes signals the presence of stronger court learning. Furthermore, our model suggests a way to gauge empirically the scope of court learning. In a world with many countries and many issue areas, the scope of court learning might be general, or specific to the disputant countries, or specific to the disputed issue area. By exploring how the likelihood of current rulings and disputes is affected by different measures of court experience, e.g. disputant-specific, issue–area-specific, or general-scope experience, we attempt to gauge the relevant domains of court learning.

Our empirical findings are broadly consistent with our model, and interestingly, we find evidence consistent with issue–area specific learning and with some forms of disputant-specific learning (in particular, complainant-specific and directed-dyad-specific), while we find only weak evidence of general-scope learning. We then discuss alternative explanations for the correlations we find in the data and argue that, while some of these explanations may well play a role, they cannot fully account for the patterns found in our data. In this light, we interpret our empirical findings as providing promising support for the proposition that court learning is an important phenomenon for understanding the pattern of WTO dispute resolution.

To our knowledge this is the first paper that explores the implications of judicial learning for trade disputes, or more generally for international institutions. A related model is Maggi and Staiger (2011), but that paper does not consider learning and does not allow for bargaining or settlement, and focuses on questions of institutional design such as the desirability of legal precedent, while here we focus mostly on how learning affects the initiation and outcome of trade disputes. In Maggi and Staiger (2015) we do allow governments to settle or fight it out in court, but the model is static, and focuses on how the outcome of trade disputes is affected by the form of the contract (property vs liability rules).

By contrast, there is a large literature on the broader implications of judicial learning, but this literature is mostly informal and does not focus on international institutions. A few recent papers have developed formal models of judicial learning (see for example Baker and Mezzetti,

2012, and Beim, 2017), but the structure and focus of these models is very different from ours. In the literature on trade agreements, other models that generate trade disputes in equilibrium are Park (2011), Beshkar (2016) and Staiger and Sykes (2017), but these papers do not focus on the determinants of dispute outcomes (with the partial exception of Beshkar). Our model is also related to the law-and-economics literature on bargaining in the shadow of the law (e.g. Bebchuck, 1984, Reinganum and Wilde, 1986); these models however are typically static, do not focus on court learning, and are not concerned with international institutions.

On the empirical side, there are papers that examine various determinants of the initiation and outcome of trade disputes, including Busch (2000), Busch and Reinhardt (2000, 2006), Guzman and Simmons (2002, 2005), Bown (2005), Davis and Bermeo (2009), Bown and Reynolds (2014), Conconi et al. (2017), Kuenzel (2017) and Maggi and Staiger (forthcoming). But none of these papers are concerned with the dynamics of court learning.

The rest of the paper is organized as follows. Section 2 presents our benchmark static model. Section 3 develops our dynamic model with learning by ruling. Section 4 develops a multi-country version of our model which allows for various forms of learning spillovers. Section 5 examines the empirical content of our theory through WTO dispute data. Section 6 offers concluding remarks. Finally, an Appendix presents a simple parametrized model that illustrates how the reduced-form assumptions we make in the main model can be “micro founded.”

2. The static model

We consider a partial equilibrium setting of trade between two countries, postponing until a later section the extension to many countries. In the industry under consideration, Home is the importing country and Foreign the exporting country. Home can choose an import barrier T , while the Foreign government is passive in this industry. For concreteness we will interpret T as a tariff, but our analysis is valid also for a more general (possibly multidimensional) policy.

The Home government’s objective function is $\omega(T, s)$, where s is a (possibly multidimensional) “state of the world” that is ex-ante uncertain. Similarly, the Foreign government’s payoff is denoted $\omega^*(T, s)$. We let the government joint payoff be $\Omega(T, s) \equiv \omega(T, s) + \omega^*(T, s)$. The state of the world s could include political-economy shocks, demand/supply parameters, or the severity of a market failure (e.g. a domestic externality) that calls for a corrective policy.⁵ We

⁵A special case of the structure above, which we consider in the Appendix, is a standard competitive partial-equilibrium setting where the Home government maximizes a politically-adjusted welfare function that attaches

assume that $\omega(T, s)$ is strictly concave in T , and we denote the Home government’s unilaterally optimal policy conditional on s by $T^N(s) \equiv \arg \max_T \omega(T, s)$. While we do not impose any special structure directly on $\omega^*(T, s)$, below we will impose some structure on the government Pareto frontier.

In what follows we describe a simple setting where, depending on the state of the world, governments may or may not get into a dispute over the setting of T ; and if they do get into a dispute, they may settle their dispute with negotiations or allow the dispute to proceed to a court ruling. Negotiations are modeled as a Nash bargain with a disagreement point given by the court ruling. We assume that if (and only if) governments engage in negotiations, they can use efficient transfers as means of compensation.

Key to our analysis in this setting is the government Pareto frontier in payoff space. Consider first the case where governments do not negotiate and hence transfers are not available. In this case, for any s the Pareto frontier is traced out by varying the Home import barrier T and recording the impacts on ω and ω^* . We assume that this frontier is strictly concave, as shown in Figure 1. This assumption is satisfied in most models of trade policy, for example in competitive settings where T is a tariff and governments maximize politically-adjusted welfare functions (as described in note 5 and provided political economy weights are not too large). We label this the “no-transfer frontier.” Note from Figure 1 that this frontier has a peak at the unilaterally optimal policy $T^N(s)$ (point N) and has slope equal to -1 at the joint-payoff-maximizing (“first best”) policy $T^{fb}(s) \equiv \arg \max_T \Omega(T, s)$ (point FB). Consider next the case in which negotiations occur and hence efficient international transfers are available. In this case, for any s the Pareto frontier is linear with slope -1 and tangent to the no-transfer frontier at the FB point as depicted in Figure 1. We label this the “negotiation frontier.”⁶ For future reference, note also that the iso-joint-payoff curves (along which Ω is constant) are lines with slope -1 .

Also key to our analysis is the disagreement point for government negotiations, that is, the expected payoff for each government if the governments fail to reach agreement in their

an extra weight to import-competing producers. In that special setting, the state of the world s will simply be the Home political-economy-shock parameter.

⁶Our assumption that governments have access to efficient transfers when they negotiate to settle a dispute simplifies the model and makes our points more transparent, but our main results would hold under the more realistic assumption that the transfers used for this purpose are costly. Relative to Figure 1, the only change this would imply is that the negotiation frontier would be concave (assuming a convex cost of transfers) but would still lie above the no-transfer frontier except for a tangency at the FB point.

negotiations and trigger a court ruling. We do not formalize explicitly the court’s information and decision-making process, nor the specific form of the contract that governments sign ex-ante. Rather, we specify in reduced-form fashion the governments’ expected payoffs from triggering a court ruling. We denote these expected payoffs as $\omega^R(s; \lambda)$ and $\omega^{*R}(s; \lambda)$, where λ is a parameter reflecting the imperfections/inefficiencies of the court system (or equivalently, the inverse of court “quality”). To the extent that the court ruling is uncertain from the governments’ point of view when they negotiate, ω^R and ω^{*R} should be interpreted as expectations over the possible rulings. We take the parameter λ to be between zero and some finite bound $\bar{\lambda}$: it can be viewed as capturing the inaccuracy of the court’s information and/or the costs of delays associated with court rulings (costs of litigation).

The most direct interpretation of a “ruling” in this setting is that the court prescribes a specific value of T , but since we do not explicitly model the nature of the ruling, other forms of ruling are also compatible with the model. For example the court could specify bounds on the policy T (e.g. a tariff binding), or it could require Home to compensate Foreign for higher trade barriers. Regardless of the specific form of the ruling, what will matter for our results is the governments’ expected joint payoff from going to court and how it is impacted by the court quality; for this reason, our reduced-form approach to modeling this impact makes the logic of our theory more transparent. As we make clear below, the core restriction of our static model is that the governments’ expected joint payoff from going to court is higher if the court is more efficient; this core assumption seems compatible with a wide variety of institutional/contractual settings and types of rulings.

In terms of concrete settings that are captured by our model, one simple environment where the quality of the court shapes the governments’ disagreement payoffs is one where the role of the court is to “complete” an incomplete contract, for example along the lines of Maggi and Staiger (2011). This is an environment where the state of the world s is observed by governments but imperfectly verifiable by the court, so governments cannot write a complete contingent contract, and the court is endowed with the authority to “fill the gaps” of the contract ex-post. In the Appendix we present a more structured model that is a special case of the environment just described.⁷ Another contractual/institutional environment that would fit our model equally

⁷In the Appendix model, the state of the world s is a political-economy parameter, and the policy T is a tariff, but the trade agreement does not specify the tariff at all. If invoked, the court “fills the gap” in this incomplete contract, by trying to estimate the first-best tariff $T^{fb}(s)$, but does so with noise. This noise represents the key imperfection in the court system (corresponding to the λ parameter described in the text), and is the key

well is one where governments can write a contingent contract but s is imperfectly verifiable, so the court – if invoked – applies the contract based on its imperfect estimate of s . In both of these environments, the quality of the court’s information affects the payoffs that governments can expect if they go to court.

Finally, while government negotiations have the advantage of enabling governments to compensate each other through transfers, we assume that such negotiations are subject to transaction costs. The reason we introduce transaction costs is that, if there were none, the disagreement point for government negotiations would be irrelevant to their joint payoff, and hence any institutional or contractual arrangement, as well as court learning, would be irrelevant. In particular, we assume an “iceberg” negotiation cost: a fraction $1 - \kappa$ of the bargaining surplus “melts” away, so governments can only move part-way toward the Pareto frontier when they negotiate. Formally, if $\omega^B(s; \lambda)$ and $\omega^{*B}(s; \lambda)$ are respectively Home’s and Foreign’s bargaining payoffs absent negotiation costs, then the net payoffs from the bargain are given by $\omega^{B_{net}}(s; \lambda) = \omega^R(s; \lambda) + \kappa[\omega^B(s; \lambda) - \omega^R(s; \lambda)]$ and $\omega^{*B_{net}}(s; \lambda) = \omega^{*R}(s; \lambda) + \kappa[\omega^{*B}(s; \lambda) - \omega^{*R}(s; \lambda)]$, or rewriting slightly:

$$\begin{aligned}\omega^{B_{net}}(s; \lambda) &= \kappa\omega^B(s; \lambda) + (1 - \kappa)\omega^R(s; \lambda), \\ \omega^{*B_{net}}(s; \lambda) &= \kappa\omega^{*B}(s; \lambda) + (1 - \kappa)\omega^{*R}(s; \lambda),\end{aligned}\tag{2.1}$$

where $\kappa \in (0, 1)$. For future reference, we let Ω^{fb} denote the first-best joint payoff and Ω^R denote the joint disagreement payoff. Noting that the joint bargaining payoff absent negotiation costs is Ω^{fb} , we can write the joint *net* bargaining payoff as

$$\Omega^{B_{net}}(s; \lambda) = \kappa\Omega^{fb}(s) + (1 - \kappa)\Omega^R(s; \lambda).$$

We now outline the timing of the static game: (1) After the state of the world s is realized and observed by governments, Home chooses T ; (2) Foreign acquiesces or initiates a dispute; (3) If a dispute is initiated, governments Nash bargain over T and a transfer; (4) If governments disagree, the court intervenes and issues a ruling (which is automatically enforced).

Notice that in our model engaging in a “dispute” is synonymous with “bargaining in the shadow of the law.” In the context of the WTO, the first step of a trade dispute is indeed that governments engage in consultations and negotiation (in fact this step is mandatory according to WTO rules). However it is important to note that in practice governments may negotiate

determinant of the governments’ disagreement point.

and settle outside the institutional framework, or in other words through informal (rather than formal) negotiations. Our model can be interpreted as applying to both formal and informal negotiations.

We will conduct much of the analysis at a graphical level, since this will aid intuition. Let us start by identifying the disagreement point and the bargaining payoffs for the stage-3 negotiations. In Figure 1, which recall depicts the government Pareto frontier for a given state s , we use the label R to indicate the expected payoffs from triggering a ruling, (ω^R, ω^{*R}) , the label B to indicate the bargaining payoffs absent negotiation costs, (ω^B, ω^{*B}) , and the label B_{net} to indicate the net bargaining payoffs $(\omega^{B_{net}}, \omega^{*B_{net}})$ defined in (2.1). We omit the arguments s and λ from the labels in this figure, as this keeps the notation simple and should not cause confusion. Clearly, the B point lies on the negotiation frontier and Northeast of point R , with the exact location determined by bargaining powers, and the B_{net} point is a linear combination of point R and point B , with the exact location determined by the negotiation cost κ .

Moving backwards in our game, we can now ask: Knowing that the dispute subgame would yield payoffs B_{net} , what policy does Home choose at stage 1? The answer to this question follows from two observations. First, if the B_{net} point is below the no-transfer frontier, Home chooses the policy T that maximizes its payoff while keeping Foreign indifferent between complaining and not. Graphically, this corresponds to the vertical projection of B_{net} onto the no-transfer frontier, which we label B^0 in Figure 1. In this case, there is no dispute in equilibrium.⁸ And second, if the B_{net} point is above the no-transfer frontier, Home will trigger a dispute by choosing a level of T that induces Foreign to complain, the governments will settle, and the equilibrium payoffs are given by point B_{net} . Intuitively, other things equal, a dispute is more likely to occur when the negotiation cost κ is lower. And in either case, it is clear that there is never a ruling in equilibrium in this static model. Summarizing:

Remark 1. *In the static setting, there may or may not be a dispute, but there is never a ruling in equilibrium.*

An important ingredient for understanding the implications of learning by ruling is to understand how changes in the court imperfection λ affect the governments' equilibrium joint payoff

⁸To be precise, if the B_{net} point lies to the left of the N point (the unilateral optimum), then the equilibrium payoff point B^0 is not the vertical projection of B_{net} onto the no-transfer frontier, but rather the N point itself. But this does not change anything of substance in our analysis. We also note that we are implicitly assuming that Foreign does not complain in case of indifference.

in this static setting. Let us start by considering the impact of λ on disagreement payoffs in the dispute subgame. Graphically, in Figure 2 we label R^λ the locus traced by the disagreement point R as λ varies (as before, we omit the argument s from the label to keep the notation simple). We impose some structure on the R^λ locus. Our core assumption is that decreasing λ (i.e. increasing court efficiency) leads to a higher expected joint disagreement payoff for each s : graphically, this means that if λ decreases the R point moves to a higher iso- Ω line. We also assume that the R point coincides with the first-best point (FB) if $\lambda = 0$ (see Figure 2); that is, a perfect court implements the first best outcome. The implicit assumption behind these restrictions is that the court attempts to maximize the governments' joint payoff, so if the court has better information and/or litigation costs are lower then the governments' joint payoff from going to court is higher; and absent any informational imperfection or litigation costs, going to court will deliver the maximum feasible joint payoff.⁹ We record this in:

Assumption 1: For all states s , the governments' joint payoff from going to court ($\Omega^R(s; \lambda)$) is decreasing in λ , with $\Omega^R(s; \lambda) = \Omega^{fb}(s)$ if $\lambda = 0$.

Note that Assumption 1 allows for the possibility that increasing court efficiency (decreasing λ) might hurt one of the governments, while increasing the “size of the pie,” thus the R^λ locus may be upward or downward sloping in payoff space.¹⁰ In Figure 2, to fix ideas we draw the R^λ locus as downward sloping.

Next we consider the impact of λ on the net bargaining payoffs in the dispute subgame. We label B_{net}^λ the locus traced by the B_{net} point as λ varies. Since, as noted above, the B_{net} point is a linear combination of the R point and the B point, the B_{net}^λ locus inherits two properties

⁹The assumption that the court attempts to maximize the governments' joint payoff seems a natural one in this setting. The idea is that governments design the institution at some ex-ante stage and endow the court with a certain objective function. As long as efficient international transfers are available at this ex-ante stage, which according to our model is the case whenever governments negotiate, and which in reality seems especially likely for ex-ante negotiations of the kind we are considering here (as distinct from the ex-post negotiations which are the topic of note 6), it is natural to suppose that this objective function is the governments' joint payoff. We also note that our results are qualitatively robust to small deviations from this assumption (i.e. if the court's objective diverges slightly from the governments' joint payoff).

¹⁰While an upward sloping R^λ locus (implying both governments gain from greater court efficiency) might seem the most natural case, the R^λ locus can also be downward sloping in standard economic environments, and is in fact downward sloping for example in the special model we present in the Appendix. There, a decrease in court efficiency takes the form of a noisier ruling, and this benefits the exporting country because its payoff is convex in the tariff. Nevertheless, in more general settings it is certainly plausible that court inefficiency might hurt both countries. This is more likely to happen if court inefficiency takes the form of costs/delay in the rulings, or if the policy T is multidimensional (so that getting the policy “wrong” may well hurt both countries), or if governments are risk-averse.

from the R^λ locus: (i) the B_{net} point coincides with the FB point if $\lambda = 0$, and (ii) it moves to lower iso- Ω lines as λ increases (that is, $\Omega^{B_{net}}$ decreases with λ). Moreover, since the R^λ locus hits the FB point with slope different than -1 , this must be true also for the B_{net}^λ locus, thus a third property is that (iii) when λ is small enough the B_{net}^λ locus lies below the no-transfer frontier (and it may or may not cross the frontier as λ increases). These three properties will play an important role in the analysis below.

We need to impose a regularity condition on the B_{net}^λ locus. As with the R^λ locus, we allow the B_{net}^λ locus to slope upward or downward, but we assume that it does not bend backwards. Note that this is the case if and only if Foreign's net bargaining payoff $\omega^{*B_{net}}$ is monotonic in λ , so we can state:

Assumption 2: For all states s , the B_{net}^λ locus is not backward-bending, or equivalently, $\omega^{*B_{net}}$ is monotonic in λ .

To clarify the conditions under which Assumption 2 is satisfied, we derive a more explicit expression for $\omega^{*B_{net}}$. Using the Nash bargaining assumption, and letting σ^* denote Foreign's bargaining power, we have $\omega^{*B_{net}} = \omega^{*R} + \kappa\sigma^*(\Omega^{fb} - \Omega^R)$, hence (assuming ω^R and ω^{*R} are differentiable in λ):

$$\frac{d\omega^{*B_{net}}}{d\lambda} = (1 - \kappa\sigma^*)\frac{d\omega^{*R}}{d\lambda} - \kappa\sigma^*\frac{d\omega^R}{d\lambda}. \quad (2.2)$$

Thus, Assumption 2 is satisfied as long as $\frac{\frac{d\omega^R}{d\lambda}}{\frac{d\omega^{*R}}{d\lambda}} \neq \frac{1-\kappa\sigma^*}{\kappa\sigma^*}$. To interpret this condition, suppose first that $\frac{d\omega^R}{d\lambda}$ and $\frac{d\omega^{*R}}{d\lambda}$ have opposite signs, so that court quality affects governments in opposite ways (this is the case for example in the special model presented in the Appendix; see also note 10): in this case, clearly Assumption 2 is automatically satisfied. Consider next the remaining possibility that $\frac{d\omega^R}{d\lambda}$ and $\frac{d\omega^{*R}}{d\lambda}$ are both negative (increasing court quality improves both disagreement payoffs): then condition (2.2) – and hence Assumption 2 – is satisfied as long as negotiation costs are sufficiently large (i.e. κ sufficiently small), because in this case the slope of the B_{net}^λ curve is sufficiently close to that of the R^λ curve.¹¹

We can now back up to stage 1 and examine how λ affects the equilibrium joint payoff in the full game. We will argue that, under Assumptions 1 and 2, the equilibrium joint payoff is decreasing in λ . To see this recall that, if the B_{net} point is below the no-transfer frontier (so

¹¹In addition, while we have stated Assumption 2 as an assumption that holds for each s , as we point out later our main results in the dynamic model of the next section will hold even if the assumption is violated for some s as long as it is satisfied on average across s .

that there is no dispute), the equilibrium point is the vertical projection of the B_{net} point onto the no-transfer frontier, and if the B_{net} point is above the no-transfer frontier (so that there is a dispute), the equilibrium point is the B_{net} point itself. It follows that, as λ increases from zero, the equilibrium point starts from the FB point and moves along the outer envelope of the no-transfer frontier and the B_{net}^λ locus (the red locus in Figure 2). Figure 2 focuses on the case in which the B_{net}^λ locus is downward sloping, but a similar argument applies if it is upward sloping. Focus first on a range of λ such that the equilibrium point is on the no-transfer frontier: as λ increases, the equilibrium point moves away from the FB point along the frontier, and hence the equilibrium joint payoff decreases. Next focus on a range of λ such that the equilibrium point is on the B_{net}^λ locus: recalling from the analysis above that $\Omega^{B_{net}}$ decreases with λ , it follows again that increasing λ will decrease the joint payoff. We can thus state:

Remark 2. *In the static setting, the equilibrium joint payoff is decreasing in λ .*

Note that an increase in court quality increases the equilibrium joint payoff through off-equilibrium effects, because there is no ruling in equilibrium: if there is a dispute and governments settle, higher court quality exerts its beneficial effect by improving the disagreement point; and if there is no dispute, higher court quality improves the would-be negotiation outcome, thus inducing Home to choose a more efficient policy (an “off-off-equilibrium” effect).

Two other implications of the model are also noteworthy, and can be understood in terms of Figure 2. First, note that for λ sufficiently low, further reductions in λ must hurt one of the governments, because the equilibrium point slides along the no-transfer frontier. Interestingly, this is the case even if reductions in λ do not affect the countries’ relative bargaining positions, that is, even if the R^λ is increasing with slope equal to one. And second, if the B_{net} point is left of the N point when λ is high, so that the equilibrium point is N (see footnote 8), then decreasing λ initially has no effect, and it starts affecting the equilibrium outcome only later on. That is, the threat of using the court may impose no discipline on the Home government’s policy choice if λ is above some threshold, and in this case court learning will have no effect until that threshold is crossed.

If we impose a bit more structure, the model also yields a prediction on how λ affects the likelihood of a dispute. Suppose the B_{net}^λ locus crosses the no-transfer frontier at most once (this will be the case if the B_{net}^λ locus is convex, linear or not too concave). Then, recalling that the B_{net}^λ locus must lie below the no-transfer frontier if λ is sufficiently close to zero, it follows

that it can only cross the no-transfer frontier from below. Thus, for a given state s , there is a low interval of λ such that there is no dispute, and there may be a high interval of λ such that there is a dispute (the latter interval may be empty). This implies that, for any ex-ante distribution of s , the likelihood of a dispute is increasing in λ . We summarize with:

Remark 3. *In the static setting, the likelihood of a dispute is increasing in λ , provided the B_{net}^λ locus crosses the no-transfer frontier at most once.*

Thus, at a broad level, the static model suggests that disputes tend to be less likely when the court system is more efficient.

3. Learning By Ruling

As we observed in the Introduction, the WTO is a relatively young international institution characterized by a fairly sophisticated judicial system. The adjudication process that this judicial system is designed to conduct is complex and subtle, and there is little doubt that the actors involved in this system have much to learn along many dimensions, especially in the early stages of the institution. In this section we extend the static model of the previous section to a dynamic setting, and explore the implications of judicial learning for the dynamics of disputes and rulings. We begin with a brief discussion of the nature of judicial learning.

3.1. Judicial Learning

One could consider different types of judicial learning. A first possibility is that the court can learn from its past experience. This is the notion that we refer to as *learning by ruling*. There are several mechanisms by which a court can learn from experience. One is that the court may become more accurate in conducting investigations and figuring out the economic and political costs/benefits of trade policies (and of domestic policies that have impacts on trade). This may involve learning to use and interpret data, or to choose the right experts, or just learning to use rigorous economic reasoning as such reasoning relates to the particular legal issues at hand. We can think of this as “methodological” learning, or in other words, “learning by doing” in investigating and adjudicating. But one can think also of a “factual” type of learning by the court: for example, by repeatedly studying the policies of a certain country (say, China) or in a certain issue area (say, health and safety), the court may gain knowledge about persistent

aspects of that country’s policy environment or of that issue area (the “state of the world”). Our model is a better fit for methodological learning than for factual learning, because we will assume for tractability that the state of the world s is *iid* over time, but our main insights intuitively extend also to the latter type of learning.

Another kind of court learning that our model can accommodate is a reduction in the costs associated with court rulings, which we refer to loosely as “litigation costs.” For example, as the court system gains experience, the delays associated with court investigations and rulings may become shorter, and this may lead to significant efficiency gains for governments. Similarly, court learning can lead to lower costs of conducting investigations and verifying the “state of the world.”

In the case of a standing judicial body such as the WTO’s Appellate Body, it may be the judges who learn directly from their own experience. But also in the case of a rotating body such as the WTO’s Dispute Settlement panels, today’s panel may learn from reading panel reports from previous cases, since such reports are publicly available. And finally, in the WTO there is another important standing body, namely the Secretariat, which is a group of experts that plays a supporting role in the adjudication process. To the extent that the Secretariat learns how to more effectively aid in the adjudication of WTO cases over time, this too can be thought of as part of court learning.

What kinds of learning fall outside our formal model? One type of learning that is probably quite relevant in the WTO is *governments’ learning about the court*. For example, it is possible that as governments litigate repeatedly in court, they learn how the court operates and adjudicates cases, and therefore they learn to better predict the outcome of a ruling. Intuitively, some of the implications of this type of learning should be similar to those of learning by ruling, because both types of learning imply that governments reap future gains by going to court today. However there may also be subtle differences in implications, because governments’ learning about the court does not *per se* increase the quality of court decisions. For this reason a formal analysis of this type of learning would be interesting in its own right, but we leave this extension for future research.

Another type of learning that falls outside our formal model is *governments’ learning about each other*. While it is certainly plausible that such learning may go on in the context of WTO disputes, we will present evidence in our empirical section that this is not the kind of learning that can explain the dynamic patterns in our data.

Finally, a dynamic mechanism that has a similar flavor but is quite distinct from learning is *legal precedent*.¹² We will discuss the likely implications of legal precedent in the empirical section, and argue that this mechanism could explain some, but not all, of the dynamic patterns observed in our data.

3.2. The two-period setting

We now extend our static model to two periods, $t = 1, 2$. In each period, the same game as described in the static setting takes place. The state of the world s_t is *iid*, so learning by ruling will be the only source of dynamics. The governments' common discount factor is $\delta \in (0, \infty)$.¹³

We assume a veil of ignorance: before the period-2 state of the world is realized, each government is equally likely to be the importer or the exporter, and hence is equally likely to be the complainant or the defendant in a dispute. The essence of the veil of ignorance is that in the future each government may find itself on either side of a trade dispute, that of complainant or that of defendant.

We model learning by ruling in a similar fashion as in the typical models of learning by doing for firms, where increasing a firm's current output increases its future efficiency: we assume that adjudicating one more case today increases the efficiency of the court tomorrow. More specifically, if there has been a ruling at $t = 1$, the court imperfection (λ) at $t = 2$ is lower. This bare-bones two period model will allow us to make a couple of key points, but later in this section we consider a slightly richer version of the model to examine how the current likelihood of disputes and rulings depends on cumulative rulings.

We start with a key observation: in contrast with the static setting, where no rulings can occur in equilibrium, the presence of learning by ruling can give rise to equilibrium rulings, because going to court today generates future payoff gains.

To establish this, we proceed by backward induction. At $t = 2$ the outcome is the same as in the static setting analyzed above, and hence there are no rulings in equilibrium. But the situation is different at $t = 1$, because there is an investment value in going to court due to the learning effect. Recall that in the static setting, under Assumptions 1 and 2, a decrease in λ leads to a higher equilibrium joint payoff. Thus, given the veil of ignorance, going to court at

¹²In Maggi and Staiger (2011) we explore the implications of legal precedent for trade disputes, in a setting without any learning.

¹³Since we have only two periods, it is natural to allow δ to be higher than one, as the second period can be thought of as condensing a potentially long future.

$t = 1$ implies a common future payoff gain, which we label Δ .

It is worth emphasizing that, much as in the static model, in our two-period model governments benefit from increasing future court efficiency through off-equilibrium mechanisms, because at $t = 2$ there is no court activity in equilibrium: making the court more efficient improves the disagreement point in case of dispute, and even if no dispute takes place, improving the would-be negotiation outcome leads to a more efficient policy choice by Home. But if we had a richer model with more than one period ahead of $t = 1$, the payoff gain Δ would include also a direct effect of increasing court efficiency in case a ruling occurs in equilibrium.

Next focus on the equilibrium outcome at $t = 1$. We suppress the time index on period-1 variables, as well as the arguments s and λ , since this should not cause confusion. At $t = 1$, the disagreement payoffs are $(\omega^R + \delta\Delta, \omega^{*R} + \delta\Delta)$. Graphically, in Figure 3 we label the corresponding payoff point $R + \delta\Delta$. Owing to the veil of ignorance, this point lies somewhere on the 45° line emanating from point R , and in general may be below or above the negotiation frontier. If point $R + \delta\Delta$ is above the negotiation frontier, then a dispute will end in ruling; and going backwards, in this case Home chooses a policy T that triggers a complaint by Foreign.

Thus it is possible that a ruling will occur in equilibrium. This will be the case if the learning effect (and hence the future gain from going to court, $\delta\Delta$) is strong relative to the loss in joint payoff that governments incur *today* if they disagree and go to court (or graphically, the distance along a 45° line between the R point and the negotiation frontier).

Finally, what is the outcome at $t = 1$ if point $R + \delta\Delta$ is below the negotiation frontier, so there is no ruling? In this case, the net bargaining payoffs at $t = 1$ given disagreement point $R + \delta\Delta$ are given by:

$$\begin{aligned}\omega_{\Delta}^{B_{net}} &= \kappa\omega^B + (1 - \kappa)(\omega^R + \delta\Delta) \\ \omega_{\Delta}^{*B_{net}} &= \kappa\omega^{*B} + (1 - \kappa)(\omega^{*R} + \delta\Delta).\end{aligned}$$

Graphically, we label B_{Δ}^{net} the point corresponding to payoffs $(\omega_{\Delta}^{B_{net}}, \omega_{\Delta}^{*B_{net}})$. Because of negotiation costs (κ), point B_{Δ}^{net} lies somewhere between the $R + \delta\Delta$ point and the negotiation frontier. It is then easy to argue that the outcome is a dispute with settlement if B_{Δ}^{net} is above the no-transfer frontier, and no dispute if B_{Δ}^{net} is below the no-transfer frontier.

3.3. Impact of court experience on current outcomes

How do past rulings affect the likelihood of current rulings and disputes? This question cannot be examined in the two-period scenario considered thus far, because rulings can occur only at $t = 1$, where there is no “past”; but a slight enrichment of the model allows us to address this question in a meaningful way.

To this end, we continue to assume two periods, $t = 1, 2$, but we now suppose there is an initial stock of rulings x , inherited from a “past” period $t = 0$. To examine how past rulings affect current outcomes, we focus on the equilibrium outcome at $t = 1$ conditional on x .¹⁴ Learning by ruling in this setting is represented by a decreasing function $\lambda(x)$.

We assume that learning occurs at a diminishing rate when evaluated in terms of its effect on net bargaining payoffs:

Assumption 3: Increases in court experience have weakly diminishing marginal impact on net bargaining payoffs: $|\frac{d\omega^{Bnet}(s, \lambda(x))}{dx}|$ and $|\frac{d\omega^{*Bnet}(s, \lambda(x))}{dx}|$ are weakly decreasing in x (for any s).

The role of this assumption will soon become transparent, but to help interpret it, note that the units of measure of court quality can be defined in an arbitrary (ordinal) way, so for a notion of diminishing rate of learning to be invariant to the units of measure, it needs to be stated in payoff terms. We focus on the impact of court experience on net bargaining payoffs because this will lead to sharp predictions in our setting.¹⁵

We are now ready to study how an increase in x affects the likelihood of rulings and disputes at $t = 1$. We evaluate these likelihoods as viewed from the beginning of period $t = 1$, that is before the state s_1 is realized. We first focus on the likelihood of a ruling.

Recall that a ruling occurs at $t = 1$ if and only if the $R + \delta\Delta$ point is above the negotiation frontier, so we can write $\Pr(\text{ruling}) = \Pr(g < \delta\Delta)$, where in terms of Figure 3 g is the distance from the R point to the negotiation frontier along a 45° line, and

$$\Delta = E_s[\Omega_{t=2}(s; \lambda(x+1)) - \Omega_{t=2}(s; \lambda(x))],$$

¹⁴It would be conceptually easy to endogenize the occurrence of a ruling at $t = 0$, but this would not add much to the question of how cumulative rulings affect current outcomes, because in a three-period setting this question is meaningful only from the perspective of the central period ($t = 1$), since at $t = 0$ there is no past and at $t = 2$ there cannot be rulings.

¹⁵Recall that we allow court quality λ – and hence court experience x – to have opposite impacts on the two governments’ payoffs, and for this reason we state Assumption 3 in terms of “diminishing marginal impact,” rather than concavity in x : we need each payoff function to be concave (convex) in x if it is increasing (decreasing) in x .

where $\Omega_{t=2}(s; \lambda)$ is the equilibrium joint payoff at $t = 2$ in state s given λ .

First note that g decreases with x , because as x increases and $\lambda(x)$ decreases, the R point moves to a higher iso- Ω line and hence moves closer to the negotiation frontier. This is the *static* effect of increasing court experience on the likelihood of a ruling: holding Δ fixed, increasing x decreases today's inefficiency from going to court, thus increasing the likelihood of a ruling.

Now focus on the *dynamic* effect of increasing court experience, that is the effect of x on Δ . Clearly, Δ is decreasing in x provided the expected joint payoff $E_s \Omega_{t=2}$ is concave in x . Recall from the static analysis that, as x increases, and hence λ decreases, the equilibrium payoff point moves along the outer envelope of the no-transfer frontier (no-dispute range) and the B_{net}^λ locus (dispute range). It is easy to show that, given Assumption 3, the equilibrium joint payoff $\Omega_{t=2}$ is strictly concave in x over the no-dispute range for any state s .¹⁶ Intuitively, this is so for two reasons: first, learning occurs at a diminishing rate; and second, conditional on there being no dispute, an increase in court quality induces the Home government to select a more efficient policy, and the associated efficiency gain is diminishing due to the concavity of the Pareto frontier.

Next note that the expected joint payoff $E_s \Omega_{t=2}$ will be globally (and strictly) concave in x if the probability of an equilibrium dispute at $t = 2$ is sufficiently small.¹⁷ Given that in reality, the average frequency with which trade disputes arise between two given governments is quite low (and that this seems true even accounting for informal disputes that occur outside the formal WTO procedures), it seems empirically relevant to restrict attention to constellations of model parameters implying a small probability of equilibrium disputes. We therefore impose:

Assumption 4: The probability of an equilibrium dispute at $t = 2$ is sufficiently small.

Under Assumptions 3 and 4, the expected future joint payoff $E_s \Omega_{t=2}(s; \lambda(x))$ is strictly concave in x , and hence increasing x reduces Δ , the future gain from going to court (dynamic effect). Recall that an increase in x also reduces g , the current inefficiency from going to court (static effect), but clearly, if governments care enough about the future (i.e. if δ is sufficiently large),

¹⁶To see this, note that in the no-dispute range the equilibrium point is the vertical projection of the B^{net} point onto the no-transfer frontier. Letting $\omega = F(\omega^*)$ denote the no-transfer frontier, the joint payoff is given by $\omega^{*Bnet} + F(\omega^{*Bnet})$. It is easy to show that this is concave in x , given that x has diminishing marginal impact on ω^{*Bnet} .

¹⁷In our model this is the case, for example, if both governments benefit from a more efficient court and the negotiation cost is high; in that case, for each state s the R point lies southwest of the FB point and the B_{net} point is close to the R point, so the probability of equilibrium disputes at $t = 2$ will be small (and may well be zero).

the dynamic effect dominates the static effect. Since this is true for any period-1 state of the world, it follows that the probability of a ruling decreases with x . Notice also that, even if δ is small, the probability of a ruling is decreasing in x for x sufficiently large,¹⁸ because when learning vanishes Δ goes to zero, while g does not go to zero. Thus we can state:

Proposition 1. *The probability of a ruling at $t = 1$ is (i) globally decreasing in x if δ is high enough and (ii) decreasing in x regardless of δ for x sufficiently large.*

Proposition 1 states that if δ is sufficiently high or x is sufficiently large, the frequency of rulings should decrease with x . But even if δ and x are small, the model does not necessarily predict that the frequency of rulings will increase with x . To see this, note that in our analysis above we have taken x as given, but if x were determined endogenously the following observation would immediately apply: if δ is sufficiently small then $(\delta\Delta - g)|_{x=0} < 0$ for all s , so rulings would never get started in equilibrium. In this light it is not obvious whether our model could ever predict that the likelihood of rulings increases with x and rulings occur in equilibrium.

Next we consider the impact of cumulative rulings on the probability of a *dispute* at $t = 1$. Recall that there is a dispute at $t = 1$ if either (i) there is a ruling, which happens if the disagreement point $R + \delta\Delta$ is above the negotiation frontier, or (ii) there is a dispute with settlement, which happens if the disagreement point $R + \delta\Delta$ is below the negotiation frontier and the net bargaining payoff point B_{Δ}^{net} is above the no-transfer frontier.

It is easy to argue that if δ is sufficiently high the likelihood of a dispute is globally decreasing in x . Intuitively, when δ is high, the dynamic effect of an increase in x (that is the decrease in Δ) dominates the static effect, so we can focus on the former. Since decreasing Δ worsens the disagreement point $R + \delta\Delta$, it also worsens the net bargaining payoffs (due to negotiation costs), thus making a dispute less appealing.¹⁹

Furthermore, under the regularity condition that the B_{net}^{λ} locus crosses the no-transfer frontier at most once, the probability of a dispute at $t = 1$ is decreasing in x regardless of δ for x sufficiently large. To see this intuitively, suppose governments are at the margin of indifference between disputing and not, so that the B_{Δ}^{net} point is on the no-transfer frontier.

¹⁸A more precise statement is that there exist values x_1 and x_2 (with $x_1 < x_2$) such that the probability of a ruling is strictly decreasing for $x \in (x_1, x_2)$ and equal to zero for $x \in (x_2, \infty)$.

¹⁹More formally, suppose governments are at the margin of indifference between disputing and not. This means that the B_{Δ}^{net} point is on the no-transfer frontier. If δ is high enough, we can focus on the dynamic effect (decrease in Δ): clearly, a decrease in Δ pushes the B_{Δ}^{net} point below the no-transfer frontier, where there is no dispute.

If x is very large then Δ is close to zero, so the B_{Δ}^{net} point is approximately the same as the B^{net} point. Increasing λ has a static and a dynamic effect; both are small, but they go in the same direction. Recalling Remark 3, the static effect is that a reduction in λ pushes down the likelihood of a dispute. And the dynamic effect is that reducing Δ makes the B^{net} point dip below the no-transfer frontier, thus reinforcing the conclusion. Hence we have:

Proposition 2. *The probability of a dispute at $t = 1$ is (i) globally decreasing in x if δ is high enough and (ii) decreasing in x regardless of δ for x sufficiently large, provided the B_{net}^{λ} locus crosses the no-transfer frontier at most once.*

Before moving on, it is worth emphasizing an important implication of the model: the frequency of disputes and rulings is not a reliable indicator of the effectiveness of the institution. Consider, for example, the frequency of rulings: according to our theory, a declining frequency of rulings does not imply that the institution is getting worse over time, in fact it is a symptom of beneficial learning by the institution. But note that this statement concerns the *change* in ruling frequency over time. A higher *level* of the ruling frequency, on the other hand, is associated with higher court efficiency according to our model. To make this point in the simplest way, suppose we shift down the schedule $\lambda(x)$ in such a way that Δ is not affected (or more generally, in such a way that the static effect of the change in $\lambda(x)$ dominates its dynamic effect): then the disagreement point is more likely to be above the negotiation frontier, so other things equal the probability of a ruling will be higher.

Our final point of this section is that the model does not yield sharp predictions regarding the conditional likelihood of settlement:

Remark 4. *At $t = 1$ the likelihood of settlement conditional on a dispute may go up or down with x , even if δ is high.*

The intuition for this result is that the effect of an increase in x on the ruling margin (that is, when the $R + \delta\Delta$ point is on the negotiation frontier) may be stronger or weaker than the effect of x on the dispute margin (that is, when the B_{Δ}^{net} point is on the no-transfer frontier), depending on the probability distribution of s , and for this reason the ratio between the probability of a ruling and the probability of a dispute can either increase or decrease.²⁰

²⁰A more formal proof is the following. Consider a realization $s = s'$ such that $R + \delta\Delta$ is just above the

According to Remark 4, it would be a mistake to look for evidence of court learning by examining how the conditional likelihood of rulings is impacted by cumulative rulings. Rather, according to our theory, court learning effects should show up most strongly in the impacts of cumulative rulings on the unconditional likelihood of a ruling and of a dispute. This serves as an important guide for the empirical work that we present later in the paper.

4. The scope of learning

Thus far we have assumed that issuing a ruling today increases the court’s future efficiency regardless of which country is the defendant in the future. But one could consider more narrow forms of learning. For example, learning might be specific to the directionality of the dispute (“directed-dyad specific”), meaning that a ruling where country i is the complainant and j the defendant increases the court’s future efficiency only for disputes where again i is the complainant and j the defendant, but not if roles are reversed. This could be the case if for example the court learns about features of the political economy of the industry in country j that competes with imports from country i . At the same time, we have restricted our attention to a two-country world, but in a many-country world one could consider broader forms of learning. For example, learning might be “defendant specific,” meaning that a ruling where country i is the complainant and j the defendant increases the court’s efficiency for any future dispute where again j is the defendant, regardless of whether the complainant in the future dispute is i or some third country.

In this section we extend our analysis to allow for many countries and a range of learning possibilities that include the “undirected-dyad specific” form of learning considered in the previous sections, as well as the alternative forms of learning just described, plus a number of other possibilities. This extension serves not only to explore the robustness of our main results, but also to enrich the model for the purposes of our empirical exploration.

We consider the simplest possible multi-country extension of our two-country partial equilibrium model: suppose there are $N \geq 2$ countries, with each of the $\frac{N!}{2(N-2)!}$ dyads trading two non-numeraire goods (one in each direction) which are separable from each other and from all

negotiation frontier. As x increases, Δ decreases and hence for $s = s'$ the outcome switches from ruling to settlement. So if the probability mass of s is concentrated around s' , then $\Pr(\text{settlement})/\Pr(\text{dispute})$ goes up. On the other hand, suppose there is zero probability mass around s' . Then a small-enough increase in x does not affect $\Pr(\text{ruling})$, while it decreases $\Pr(\text{dispute})$, thus $\Pr(\text{ruling})/\Pr(\text{dispute})$ goes up and hence $\Pr(\text{settlement})/\Pr(\text{dispute})$ goes down.

other non-numeraire goods that countries trade, and with an outside good that enters utility quasi-linearly. The payoff to a given government j is then the sum of $2(N - 1)$ (separable) product-specific payoff terms: for each trading partner, say country i , there is a payoff term associated with the product that country j imports from country i , with state variable s_{ij} , and a payoff term associated with the product that country j exports to country i , with state variable s_{ji} . As before, the state variables are ex ante uncertain. Finally, we continue to assume that there can be at most one dispute in any period, and we continue to assume that analogues of Assumptions 1-4 apply in this extended setting.

We denote by λ_{ij} the court's imperfection in ruling on disputes brought by country i against country j . We assume that $\lambda_{ij} = \lambda(X_{ij})$, where $\lambda(\cdot)$ is a decreasing and convex function, and X_{ij} is a composite experience variable that takes the form

$$X_{ij} = \beta_1 x_{ij} + \beta_2 x_{ji} + \beta_3 x_{i(nj)} + \beta_4 x_{(ni)j} + \beta_5 x_{other}. \quad (4.1)$$

Here, x_{ij} is the number of past rulings where i was the complainant and j was the defendant, x_{ji} the number of past rulings where j was the complainant and i the defendant, $x_{i(nj)}$ the number of past rulings where i was the complainant and the defendant was not j , $x_{(ni)j}$ the number of past rulings where j was the defendant and the complainant was not i , and x_{other} is the number of remaining past disputes.

All β 's are assumed weakly positive. Moreover, it is natural to assume that β_1 is at least as large as each of the other β 's, because it is plausible that direct experience is at least as relevant as indirect experience; and by a similar argument, it is natural to suppose that β_2 , β_3 and β_4 are at least as large as β_5 .

Our formulation of court learning includes several interesting possibilities. At one extreme, learning could be purely general, in the sense that prior experience improves the court's efficiency in future disputes regardless of the identities of the disputants or the roles they play. This case of pure *general learning* corresponds to the case where all the β 's are equal and strictly positive. At the other extreme, court learning could be highly specific, so that prior experience is applicable only to future disputes in which the same governments play the same roles: this *directed-dyad specific learning* corresponds to the case where $\beta_1 > 0$ and all other β 's are zero. And in between these two extremes are the cases of *undirected-dyad specific learning* ($\beta_1 = \beta_2 > 0$ and all other β 's are zero), where prior experience is applicable to future disputes between the same governments regardless of the roles they play; *defendant-specific learning*

($\beta_1 = \beta_4 > 0$ and all other β 's are zero), where prior experience is applicable to future disputes that involve the defendant again in the role of a defendant, regardless of who the complainant is; and *complainant-specific learning* (only $\beta_1 = \beta_3 > 0$ and all other β 's are zero), where prior experience is applicable to future disputes that involve the complainant again in the role of a complainant, regardless of who the defendant is. And of course these possibilities are not mutually exclusive: for example, the case in which there is general learning as well as directed-dyad specific learning would correspond to $\beta_1 \geq \beta_2 = \beta_3 = \beta_4 = \beta_5 \geq 0$ (with the difference between β_1 and the other β 's interpreted as the directed-dyad specific component).²¹

Key to the predictions of the model are the future payoff changes implied by a ruling at $t = 1$. Let us denote these future payoff changes for the period-1 complainant (i) and for the period-1 defendant (j) respectively as Δ^i and Δ^j . It is easy to show that, if an increase in court quality improves the disagreement payoff for both the complainant and the defendant, then the joint future payoff change $\Delta^i + \Delta^j$ must be positive. However, if court quality has opposing impacts on the defendant and the complainant, as is possible in our model, $\Delta^i + \Delta^j$ in principle could be negative. In the Appendix we show that, if at $t = 2$ the probability of a potential dispute between governments i and j is large enough relative to the probability of a potential dispute between some other country pair – that is, if there is a large enough probability of “re-matching” between governments – then $\Delta^i + \Delta^j$ is guaranteed to be positive, and hence a ruling can occur in equilibrium at $t = 1$. We note that this condition is much stronger than we need, and the same result could be obtained in a variety of other ways, for example by assuming enough symmetry (so there is enough of a “veil of ignorance,” ensuring that Δ^i is not very different from Δ^j), or taking the opposite approach, by assuming enough heterogeneity across country dyads, so that $\Delta^i + \Delta^j$ is positive at least some of the time (note that we need $\Delta^i + \Delta^j$ to be positive only in some states of the world, not always).²²

We can now focus on the impact of past rulings on the likelihood of current rulings and

²¹Note an implicit restriction in our model: we are assuming that past disputes between third countries (i.e. countries other than i and j) have the same relevance as past disputes where today’s complainant (i) played the role of defendant, or past disputes where today’s defendant (j) played the role of complainant. We could write down a more general model where these learning effects are allowed to be different, but this would substantially complicate the notation and exposition without much gain in insight.

²²It is worth noting that, to the extent that there are learning spillovers that can benefit third countries in the future (e.g. β_3 or β_5 are strictly positive) there is potential for under-utilization of the court system, in the sense that $\Delta^i + \Delta^j$ may not be large enough to generate a ruling in equilibrium even though this would be desirable from a global standpoint. While this is an interesting possibility that merits further examination, it goes beyond the scope of this paper, which is mainly positive rather than normative.

disputes: how does an increase in x_m (for $m = ij, ji, i(nj), (ni)j, other$) affect the probability of a ruling between (complainant) i and (defendant) j at $t = 1$?

Just as in the previous section, increasing x_m has a static effect and a dynamic effect. The static effect is that increasing any of the x_m 's increases (at least weakly) current court quality, implying that the disagreement point gets closer to the negotiation frontier; other things equal, this static effect pushes up the probability of a ruling. The dynamic effect is that the increase in x_m affects $\Delta^i + \Delta^j$. In the Appendix we show that, under the assumption made above that the probability of "re-matching" is high (and the analogs of Assumptions 1-4), $\Delta^i + \Delta^j$ is weakly decreasing in each of the x_m 's. Thus, as in our two-country setting, the dynamic effect works in the opposite way as the static effect, pushing down the probability of a ruling. It follows that, if δ is large enough, the dynamic effect outweighs the static effect, and the probability of a ruling is weakly decreasing in each x_m . Furthermore it is easy to see that, regardless of δ , if the composite experience variable X_{ij} is large enough then the probability of a ruling is decreasing in each x_m , just as in our two-country setting, because the probability of a ruling eventually must hit zero as X_{ij} becomes large.

Finally, we ask whether also the likelihood of a dispute is decreasing in each x_m . The answer is yes if β_1 is sufficiently close to β_2 , so that learning is essentially undirected-dyad specific, and δ is sufficiently large: intuitively, a setting with $\beta_1 = \beta_2$ is analogous to our two-country setting of the previous section. But in general, court experience can have a positive or negative impact on the likelihood of a dispute in this multi-country setting.²³

5. Empirical Evidence

We now provide an initial assessment of the empirical content of our theory, examining patterns in WTO dispute behavior and focusing on the theory's main prediction: if there is court/DSB learning, the likelihood of current disputes and rulings should tend to decrease with the stock of cumulative past rulings, although this prediction applies more strongly for rulings than for

²³To see why, suppose for instance that only β_1 is positive (directed-dyad specific learning). Suppose further that the B_{Δ}^{net} point is at the dispute margin, i.e. on the no-transfer frontier, and suppose δ is large, so that we can ignore the static effect and focus on the dynamic effect. And finally, suppose that an increase in DSB efficiency is beneficial to the defendant but hurts the complainant, as is possible (see note 10). Consider an increase in x_{ij} or x_{ji} : as we argued above, in this case $\Delta^i < 0$, $\Delta^j > 0$ and $\Delta^i + \Delta^j > 0$, and furthermore, increasing x_{ij} or x_{ji} reduces $\Delta^i + \Delta^j$. It can also be shown that increasing x_{ij} or x_{ji} reduces both Δ^i and Δ^j in absolute value. This implies that the $R + \delta\Delta$ point moves Northwest with slope steeper than -1 , and so does the B_{Δ}^{net} point. This could lead the B_{Δ}^{net} point to dip below the no-transfer frontier or to rise above it, thus the impact on the likelihood of a dispute is ambiguous.

disputes, especially in our multi-country extension. Our empirical investigation has a dual objective. First, we want to ask whether our theory’s central prediction is consistent with the data. And second, to the extent that the answer to this question is affirmative, we want to gauge the empirical importance of learning by ruling and assess its scope and form.

Our dataset consists of 388 WTO disputes initiated between 1995 and 2009 as contained in the WTO Dispute Settlement Database. Our Data Appendix describes the steps taken in constructing this dataset. We note at the outset that, while our data on the frequency of DSB rulings is quite reliable, we face a potential limitation when it comes to data on the frequency of disputes, because a dispute can either end in a DSB ruling or it can end in settlement; and as we observed in section 2, settlement in our model can be interpreted either as a deal struck within the formal WTO dispute process or as a deal struck outside this process. Unfortunately, we only have data on settlements that occur within the formal WTO dispute process, thus we face a potentially important sample selection problem when measuring the frequency of disputes. Nevertheless, with this caveat in mind, we will examine how past rulings affect the current frequency of both rulings and disputes.

Recall that our multi-country model of section 4 allows for a rich set of possibilities regarding the scope of court learning, including the five possibilities of *directed-dyad*-specific learning, *undirected-dyad*-specific learning, *complainant*-specific learning, *defendant*-specific learning and *general-scope* learning. But our model considers only one sector or issue area. For empirical purposes, it seems compelling to allow for one more dimension of learning, namely, learning may or may not be specific to the disputed issue area. To operationalize the notion of “issue area” in a simple way, we assume that an issue area is embodied in a GATT/WTO Article. If learning can be specific to the GATT/WTO Article ruled upon by the court, then we have five additional possibilities: court learning could be *article specific*, *directed-dyad-and-article* specific, *undirected-dyad-and-article* specific, *complainant-and-article* specific, and *defendant-and-article* specific. And of course, combinations of these different dimensions of learning might be present but in different degrees.

Below we investigate empirically all of these different potential domains of learning, but for simplicity we organize our investigation into two sections. We first treat dyads as undirected, and focus on learning that might be general, undirected-dyad specific, article specific or undirected-dyad-and-article specific. We then turn to the possibility of directed-dyad-specific learning effects, and allow for all the potential domains of learning described above.

To facilitate the distinction between undirected and directed dyads, we use \overleftrightarrow{ij} to index undirected dyads, and \overrightarrow{ij} to index directed dyads where country i is the complainant and country j the defendant. And we use k to index GATT/WTO Articles and t to index years. For undirected dyads, we then define the following variables: $D_{\overleftrightarrow{ij},k,t}$ is the number of disputes initiated by country-dyad \overleftrightarrow{ij} on article k in year t ; $R_{\overleftrightarrow{ij},k,t}$ is the number of country-dyad- \overleftrightarrow{ij} disputes on article k that led to an adopted panel ruling in year t ; and $CR_{\overleftrightarrow{ij},k,t}$ is the cumulative number of country-dyad- \overleftrightarrow{ij} disputes on article k that led to an adopted panel ruling prior to year t . In what follows, we refer to $R_{\overleftrightarrow{ij},k,t}$ simply as the number of “rulings” for dyad \overleftrightarrow{ij} on article k in year t , and similarly for the variable $CR_{\overleftrightarrow{ij},k,t}$. The analogous variables for directed dyad \overrightarrow{ij} are $D_{\overrightarrow{ij},k,t}$, $R_{\overrightarrow{ij},k,t}$ and $CR_{\overrightarrow{ij},k,t}$.

Notice that our convention is to date disputes by the year in which they are formally initiated (through a “request for consultation,” the official start of formal WTO dispute settlement proceedings), and to date DSB rulings by the year in which the DSB panel report containing the ruling is formally “adopted” (approved) by the WTO membership. The latter dating convention reflects our belief that the entire panel process – investigation, preliminary and final reports, and appeals – that leads up to final adoption of DSB rulings is a potentially important source of DSB learning.

5.1. Undirected Dyads

We begin our investigation of undirected dyads with some simple plots, focusing on rulings. The analogous plots for disputes are somewhat weaker but reflect broadly similar patterns.

In Plot 3 we depict on the vertical axis $R_{.,k,t}$, the number of rulings on article k (for any dyad) in year t , and on the horizontal axis we depict $CR_{.,k,t}$, the cumulative number of rulings on article k (for any dyad) prior to year t . The appearance of a negative relationship in Plot 3 is consistent with the presence of article-specific DSB learning according to our model.

In Plot 4 we depict on the vertical axis $R_{\overleftrightarrow{ij},.,t}$, the number of rulings for dyad \overleftrightarrow{ij} (on any article) in year t , and on the horizontal axis we depict $CR_{\overleftrightarrow{ij},.,t}$, the cumulative number of rulings for dyad \overleftrightarrow{ij} (on any article) prior to year t . The appearance of a negative relationship in Plot 4 is consistent with the presence of undirected-dyad-specific DSB learning according to our model, though this relationship seems weaker than the relationship in Plot 3.

Finally, in Plot 5 we depict on the vertical axis $R_{\overleftrightarrow{ij},.,t}$, the number of rulings for dyad \overleftrightarrow{ij} (on any article) in year t , and on the horizontal axis we depict $CR_{.,.,t}$, the cumulative number

of rulings – on any article and for any dyad – prior to year t . Unlike for Plots 3 and 4, Plot 5 shows no discernible relationship between current rulings and cumulative past rulings, and hence no suggestion of general DSB learning according to our model.

We next turn to some basic regressions, in order to probe the visual impressions suggested by Plots 3-5. Below we present results from two regressions, one for disputes and one for rulings, under both logit and OLS estimation. We focus our discussion in the text on the logit results, but we point out where our logit results diverge from the OLS results and emphasize only those findings that are common to both. We estimate the dispute regression with a panel spanning the 15 years 1995-2009 and consisting of observations on each of the 126 undirected country dyads that initiated at least one WTO dispute during this period and each of the 241 GATT/WTO Articles that were disputed at least once during this period. For the ruling regression, we restrict the sample to the 55 undirected country dyads that generated at least one WTO adopted panel ruling report as a result of a dispute initiated during this period and to the 140 GATT/WTO Articles that were ruled upon at least once in an adopted panel report as a result of a dispute initiated during this period.²⁴

The dependent variable in the dispute logit regression is $DLogit_{ij,k,t}^{\leftrightarrow}$, defined as 1 if $D_{ij,k,t}^{\leftrightarrow} \geq 1$ and 0 otherwise. The dependent variable in the ruling logit regression is $RLogit_{ij,k,t}^{\leftrightarrow}$, defined as 1 if $R_{ij,k,t}^{\leftrightarrow} \geq 1$ and 0 otherwise.²⁵ For both the dispute and ruling regressions, the key independent variables of interest are four measures of cumulative past rulings, which we denote by $CR_{ij,k,t}^{\leftrightarrow}$, $CR_{n(ij),k,t}$, $CR_{ij,nk,t}$ and $CR_{n(ij),nk,t}$ where a subscript nz denotes “not z ” for index z .²⁶ The variable $CR_{n(ij),k,t}$, defined as the cumulative number of rulings for dyads other

²⁴Our panel is unbalanced, due to WTO accessions that occurred between the WTO’s inception in 1995 and the end of our sample period in 2009: as a result of these accessions, the number of undirected dyads for the dispute regression rises from 110 in 1995 to 126 in 2009, while the number of undirected dyads for the ruling regression rises from 50 in 1995 to 55 in 2009. For our purposes here it seems reasonable to treat accessions as exogenous, and under this assumption the unbalanced nature of our panel raises no special econometric issues (see, e.g., Wooldridge, 2010, pp 828-830). Nevertheless, to check that our unbalanced panel is not impacting our results, we have also re-estimated all of the (undirected- and directed-dyad) regressions we present below on the restricted sample of dyads between countries that were original members of the WTO (i.e., members beginning in 1995), and find that our results are unchanged. A similar issue arises with the growing membership of the EU over our sample period. To check that our results are also not impacted by this issue, we re-estimated all of the (undirected- and directed-dyad) regressions we present below on the restricted sample that excludes the disputes that involved an EU-joiner (before it joined the EU) during our sample period, and again find that none of our results are materially impacted.

²⁵The dependent variables for the OLS undirected-dyad dispute and ruling regressions are, respectively, $D_{ij,k,t}^{\leftrightarrow}$ and $R_{ij,k,t}^{\leftrightarrow}$. We have also re-estimated all of our OLS (undirected- and directed-dyad) regressions with the 0/1 dependent variables used in our logit regressions, and find that the results are not materially impacted.

²⁶With regard to the ruling regressions, recall from Remark 4 and our discussion following this Remark that

than \overleftrightarrow{ij} on article k prior to year t , captures *article-specific* court experience. The variable $CR_{\overleftrightarrow{ij},nk,t}$, defined as the cumulative number of rulings for dyad \overleftrightarrow{ij} on articles other than k prior to year t , captures *undirected-dyad-specific* court experience. The variable $CR_{n(\overleftrightarrow{ij}),nk,t}$, defined as the cumulative number of rulings for dyads other than \overleftrightarrow{ij} on articles other than k prior to year t , captures *general* court experience. Finally, the variable $CR_{\overleftrightarrow{ij},k,t}$ is the cumulative number of rulings for dyad \overleftrightarrow{ij} on article k prior to year t ; this variable is meant to capture the narrowest form of court experience that is specific to both the disputants involved and the article that they are disputing. The top half of Table 1 provides summary statistics for each of the variables used in the undirected dyad regressions.

The results of the undirected dyad regressions are presented in columns 1 and 2 of Table 2 (with the corresponding OLS results contained in Columns 1 and 2 of Table 3). Each regression includes a quadratic time trend, as well as (undirected)-dyad- and article- fixed effects to control for unobserved heterogeneity in the disputes and rulings behavior at the level of the dyad (the countries in dyad \overleftrightarrow{ij} may have a particularly litigious relationship) and the level of the article (article k may be particularly susceptible to litigation).²⁷

Importantly, we do not include an $\overleftrightarrow{ij}k$ fixed effect, and therefore do not control for unobserved heterogeneity at the level of the dyad *and* article (the countries in dyad \overleftrightarrow{ij} might have a particularly litigious relationship over article k), for two reasons. First, and most obviously, including such a fixed effect and relying only on within- $\overleftrightarrow{ij}k$ variation over time to estimate the regression coefficients would diminish our ability to assess the impact of those variables that exhibit little within- $\overleftrightarrow{ij}k$ variation over time.²⁸ And second, for the ruling regressions the right-hand-side variable $CR_{\overleftrightarrow{ij},k,t}$ is the sum of lagged values of the dependent variable, and inclusion of an $\overleftrightarrow{ij}k$ fixed effect would introduce an incidental parameters problem and lead

our theory points to looking for evidence of court learning by examining the impacts of cumulative rulings on the unconditional likelihood of a ruling; hence we do not control for selection into rulings when estimating our ruling regressions.

²⁷We have also experimented with the inclusion of further controls, including variables that capture the tendency of richer (OECD) countries to be claimants in WTO disputes involving intellectual property rights (TRIPS articles) and to be respondents in WTO disputes involving subsidies (SCM articles) and technical barriers (SPS/TBT articles), as well as even more specific controls (such as disputes that involve obligations specific to China's accession agreement to the WTO) and also more general controls (such as the bilateral real trade volume between countries i and j in year t , and measures of exchange rate overvaluation as a time-varying indicator of a country's incentive to initiate WTO disputes over the policies of its trading partners). Our results are robust to the inclusion of these additional controls.

²⁸Relatedly, we choose to include a quadratic time trend rather than year fixed effects because the inclusion of year fixed effects would interfere with our ability to assess the importance of our general learning variable (which exhibits little within-year variation over the cross-section of $\overleftrightarrow{ij}k$).

to biased and inconsistent estimates for our relatively short panel.²⁹ An implication is that if there *is* important unobserved heterogeneity at the dyad-and-article level, our estimates of the coefficient on $CR_{\overleftrightarrow{ij},k,t}$ will be biased upward, a bias that works against finding evidence of the most narrow form of learning.

Focusing first on the ruling logit regression in column 2 of Table 2, the estimated coefficients on $CR_{n(\overleftrightarrow{ij}),k,t}$ and $CR_{\overleftrightarrow{ij},nk,t}$ are negative and strongly significant, confirming the visual impressions of Plots 3 and 4 and suggesting the presence of article-specific and dyad-specific DSB learning. And while the coefficient estimate on $CR_{n(\overleftrightarrow{ij}),nk,t}$ in column 2 of Table 2 is negative and significant, the corresponding OLS coefficient estimate in column 2 of Table 3 is insignificantly different from zero, suggesting overall only weak evidence of general-scope learning, in line with the visual impression of Plot 5.³⁰ Finally, notice that the point estimate of the coefficient on $CR_{\overleftrightarrow{ij},k,t}$, our narrowest measure of DSB experience, is positive (and strongly significant according to the OLS estimates in column 2 of Table 3). This may reflect the upward bias in this coefficient that would be expected if there is unobserved heterogeneity at the dyad-and-article level. Below we offer more evidence consistent with this interpretation.

Turning to the dispute regression, column 1 of Table 2 presents the coefficient estimates from the $DLogit_{\overleftrightarrow{ij},k,t}$ regression. The results are broadly similar to those of the ruling regressions. The coefficient estimates on $CR_{n(\overleftrightarrow{ij}),k,t}$ and $CR_{\overleftrightarrow{ij},nk,t}$ are negative and significant, while the coefficient estimate on $CR_{n(\overleftrightarrow{ij}),nk,t}$ is negative and significant in the logit estimation but is insignificantly different from zero under OLS (column 1 of Table 3). Thus, as with the ruling regressions, the dispute regressions are suggestive of article-specific and dyad-specific DSB learning and there is only weak evidence of general-scope learning. And now the coefficient on $CR_{\overleftrightarrow{ij},k,t}$, our narrowest measure of DSB experience, is positive and strongly significant.

As a partial check on the interpretation that our failure to find a negative coefficient on $CR_{\overleftrightarrow{ij},k,t}$ reflects the presence of unobserved heterogeneity at the dyad-and-article level, we

²⁹Letting T denote the length of the panel, the issue that arises for our ruling regressions if an $\overleftrightarrow{ij}k$ fixed effect is included is that for T fixed and relatively small, the estimates of the slope parameter on $CR_{\overleftrightarrow{ij},k,t}$ will be biased and inconsistent even as the \overleftrightarrow{ij} and k dimensions of the panel become large. This is because the number of $\overleftrightarrow{ij}k$ fixed effects to be estimated grows proportionately with the \overleftrightarrow{ij} and k dimensions of the panel, and only the “within” dimension of the data (with T observations) can be used to estimate the slope parameter on $CR_{\overleftrightarrow{ij},k,t}$; and the presence of a lagged endogenous variable ensures that this regressor will be correlated with the error term unless $T \rightarrow \infty$. See Wooldridge (2010) for a textbook treatment of the incidental parameter problem and possible approaches to addressing it.

³⁰For the OLS results, we report standard errors clustered by dyad, but clustering by dyad and article makes no material difference to the results we emphasize.

next present estimates of the dispute regressions (logit and OLS) with an \overleftrightarrow{ij} k fixed effect. Recall that inclusion of this fixed effect will diminish our ability to assess the impact of those variables that exhibit little within- \overleftrightarrow{ij} k variation over time, but should address the upward bias in the estimated coefficient on $CR_{\overleftrightarrow{ij},k,t}$ induced by any unobserved heterogeneity at the dyad-and-article level; and for the dispute regressions, the inclusion of an \overleftrightarrow{ij} k fixed effect does not mechanically lead to biased or inconsistent estimates as would be the case for the rulings regressions where $CR_{\overleftrightarrow{ij},k,t}$ constitutes a lagged dependent variable. The results are contained in columns 1 (logit) and 2 (OLS) of Table 4. When an \overleftrightarrow{ij} k fixed effect is included in the dispute regressions the coefficient on $CR_{\overleftrightarrow{ij},k,t}$ turns strongly and significantly negative, consistent with our interpretation and with the presence of DSB learning even at the dyad-and-article level.

5.2. Directed Dyads

We next turn to our analysis based on directed dyads. As we discussed in section 4, DSB learning might be specific to the defendant country (which is under the magnifying glass of the DSB), or to the complainant country (e.g. because the DSB learns about the political-economic impacts of trade barriers on this country's exporters), or even to the directed dyad itself (e.g. by adjudicating disputes brought by China against the US, the DSB may learn about sectors where China exports to the US). Our directed dyad regressions can provide evidence on the possible importance of these finer dimensions of DSB learning.

As with the undirected dyads, we estimate two regressions, one for disputes and one for rulings, and we report both logit and OLS results but emphasize the logit results in our discussion in the text. For the dispute regressions, our panel (spanning the 15 years 1995-2009) consists of observations on each of the 156 directed country dyads that initiated at least one WTO dispute during this period and each of the 241 GATT/WTO Articles that were disputed at least once during this period. The dependent variable in the dispute logit regression is $DLogit_{\overleftrightarrow{ij},k,t}$ (defined as 1 if $D_{\overleftrightarrow{ij},k,t} \geq 1$ and 0 otherwise). And for the ruling regression, we restrict the sample to the 73 directed country dyads that generated at least one WTO adopted panel ruling report as a result of a dispute initiated during this period and to the 140 GATT/WTO Articles that were ruled upon at least once in an adopted panel report as a result of a dispute initiated during this period. The dependent variable in the ruling logit regression is $RLogit_{\overleftrightarrow{ij},k,t}$ (defined as 1 if $R_{\overleftrightarrow{ij},k,t} \geq 1$ and 0 otherwise).³¹

³¹The dependent variables for the OLS directed-dyad dispute and ruling regressions are $D_{\overleftrightarrow{ij},k,t}$ and $R_{\overleftrightarrow{ij},k,t}$.

For both the dispute and ruling regressions, the key independent variables of interest are now the 10 measures of court experience denoted by $CR_{ij,k,t}^{\rightarrow}$, $CR_{(ni)j,k,t}^{\rightarrow}$, $CR_{i(nj),k,t}^{\rightarrow}$, $CR_{ji,k,t}^{\rightarrow}$, $CR_{other,k,t}^{\rightarrow}$, $CR_{ij,nk,t}^{\rightarrow}$, $CR_{(ni)j,nk,t}^{\rightarrow}$, $CR_{i(nj),nk,t}^{\rightarrow}$, $CR_{ji,nk,t}^{\rightarrow}$ and $CR_{other,nk,t}^{\rightarrow}$. The meaning of these variables can be understood as follows. Consider first the five k -specific variables:

- (1) the variable $CR_{ij,k,t}^{\rightarrow}$ captures directed-dyad-and-article specific court experience;
- (2) the variable $CR_{ji,k,t}^{\rightarrow}$ captures court experience that is specific to the “reverse” directed dyad, i.e. where j complains against i , and to the article (and thus, together with $CR_{ij,k,t}^{\rightarrow}$, captures undirected-dyad-and-article specific court experience);
- (3) the variable $CR_{i(nj),k,t}^{\rightarrow}$ captures complainant-and-article specific court experience;
- (4) the variable $CR_{(ni)j,k,t}^{\rightarrow}$ captures defendant-and-article specific court experience;
- (5) the variable $CR_{other,k,t}^{\rightarrow}$ captures article specific (but not disputant specific) court experience.

The second group of five variables is analogous, except that cumulative rulings are aggregated over all non- k articles. And the interpretation of these variables is also analogous, except that they capture non-article-specific court experience: for example, $CR_{ij,nk,t}^{\rightarrow}$ captures directed-dyad-specific experience, $CR_{i(nj),nk,t}^{\rightarrow}$ captures complainant-specific experience, and $CR_{other,nk,t}^{\rightarrow}$ captures general experience. The bottom half of Table 1 provides summary statistics for the variables in the directed dyad regressions.

The results of the directed-dyad logit regressions are presented in columns 3 and 4 of Table 2 (with the corresponding OLS results contained in Columns 3 and 4 of Table 3). Similarly to the undirected-dyad regressions, in both of our directed-dyad regressions we also include a quadratic time trend and article- and (directed)-dyad- fixed effects.

Focusing first on the ruling regression in column 4 of Table 2, the estimated coefficient on $CR_{other,k,t}^{\rightarrow}$ is negative and strongly significant, which suggests the presence of article-specific learning. The estimated coefficient on $CR_{ij,nk,t}^{\rightarrow}$ is negative and strongly significant, suggesting the presence of directed-dyad-specific learning (and the estimated coefficient on $CR_{ji,nk,t}^{\rightarrow}$ is negative and significant for the logit but insignificantly different from zero with a positive point estimate for OLS, suggesting at best weak evidence for *undirected*-dyad-specific learning). The estimated coefficient on $CR_{i(nj),nk,t}^{\rightarrow}$ is also negative and strongly significant, suggesting complainant-specific learning. And the estimated coefficient on $CR_{other,nk,t}^{\rightarrow}$, while negative and

See also notes 24 and 25.

significant in the logit specification, is insignificantly different from zero for OLS, suggesting at best only weak evidence of general-scope learning. Finally, the estimated coefficient on $CR_{\vec{ij},k,t}$ is positive and strongly significant, possibly reflecting as we indicated earlier an upward bias in the estimated coefficient on $CR_{\overleftarrow{ij},k,t}$ from the presence of unobserved heterogeneity at the dyad-and-article level (and the estimated coefficient on $CR_{\overleftarrow{ji},k,t}$ is negative and significant in the logit specification but insignificantly different from zero for OLS).

Turning to the dispute regression results in column 3 of Table 2, the results are broadly consistent with the ruling regressions of column 4. In particular, the estimated coefficients on $CR_{\overrightarrow{other},k,t}$, $CR_{\overrightarrow{ij},nk,t}$ and $CR_{\overrightarrow{i(nj)},nk,t}$ are each negative and strongly significant, suggesting the presence of article-specific, directed-dyad-specific and complainant-specific learning. And there is no evidence of general learning from the dispute regression (the estimated coefficient on $CR_{\overrightarrow{other},nk,t}$ is statistically insignificant), reinforcing the caution with which we interpreted the coefficient on this variable in the ruling regression. And as with the ruling regression, the estimated coefficient on $CR_{\vec{ij},k,t}$, our most narrow measure of DSB experience, is positive and strongly significant. Again to check our interpretation that this positive coefficient reflects the presence of unobserved heterogeneity at the dyad-and-article level and an upward bias in the estimated coefficient on $CR_{\vec{ij},k,t}$, we present estimates of the directed-dyad dispute regression with an $\vec{ij}k$ fixed effect in columns 3 (logit) and 4 (OLS) of Table 4. When the $\vec{ij}k$ fixed effect is included in the directed-dyad dispute regressions the coefficient on $CR_{\vec{ij},k,t}$ turns strongly and significantly negative, consistent with our interpretation above and with the presence of DSB learning at our most narrow level.³²

The one difference relative to the ruling regression results in column 4 of Table 2 is that in the dispute regression results in column 3 of Table 2 the estimated coefficient on $CR_{\overrightarrow{(ni)j},k,t}$ has switched from negative but insignificantly different from zero to positive and strongly significant. Recalling that our model (particularly in its multi-country extension) yields more ambiguous predictions about the impacts of experience variables such as $CR_{\overrightarrow{(ni)j},k,t}$ on the frequency of disputes than it does for rulings, it is possible that the positive coefficient on $CR_{\overrightarrow{(ni)j},k,t}$ in the

³²Note that when an $\vec{ij}k$ fixed effect is included in the directed-dyad dispute regressions in columns 3 and 4 of Table 4 the estimated coefficient on the article-specific learning term $CR_{\overrightarrow{other},k,t}$ loses its significance (and in fact turns slightly positive). This likely reflects the loss of effective variation used to estimate the regression coefficient on this variable in the presence of an $\vec{ij}k$ fixed effect. And the positive and significant coefficient on $CR_{n(\vec{ij}),k,t}$ in the undirected dyad logit estimate in column 1 of Table 4 can be similarly understood from the perspective of the directed dyad logit in column 3 as reflecting the loss of significance of the coefficient on $CR_{\overrightarrow{other},k,t}$ together with the positive and significant coefficient on $CR_{\overrightarrow{(ni)j},k,t}$, which we discuss next.

dispute regression of column 3 is a manifestation of this ambiguity. An alternative interpretation is that this reflects a “bandwagon effect” that falls outside our model, whereby other potential complainants follow up with claims of their own once a ruling against defendant-country j on article k has been issued and adopted.³³

Overall, the results of our regressions reveal several important points. First, we find evidence consistent with article-specific and disputant-specific court learning, with the latter taking the form of directed-dyad-specific and complainant-specific learning, while we do not find much evidence of general-scope learning.³⁴ It is also notable that the coefficient on the linear time trend is positive in all of our regressions. The fact that controlling for our measures of court experience (the CR variables) helps wipe out the negative effect of calendar time suggests that court learning can indeed help explain the raw declining trend in disputes and rulings that was evidenced in Plots 1 and 2, as we hypothesized at the outset. And finally, there is evidence consistent with a possible “bandwagon” effect, and so a more complete empirical account of the pattern of WTO disputes and rulings may require an extended model that captures these effects in addition to the effects of court learning on the dynamics of dispute resolution.³⁵

³³It is also interesting to note that, while the logit coefficient on $CR_{\vec{j},k,t}$ in column 3 of Table 2 is negative but insignificant, the OLS coefficient in column 3 of Table 3 is positive and significant, providing some weak evidence for a possible “tit-for-tat” effect (e.g., if the US files an article- k complaint today against China, in the future China is more likely to file an article- k complaint against the US) that is outside our model. Indeed, there is some anecdotal evidence of such tit-for-tat behavior in the practice of WTO disputes (see for example the article by Jennifer Freedman in Bloomberg Business, 2012, and the discussion in Davis, 2012).

³⁴While the WTO was created in 1995, it included both the set of pre-existing GATT Articles from 1947 and also a set of new WTO Articles. In this light, one might conjecture that court learning effects in the WTO era would be stronger for WTO than for GATT articles. When we estimate the regressions in Tables 2 and 3 allowing for separate learning effects for WTO versus GATT articles, we find that the learning effects are statistically indistinguishable across the two sets of articles with one exception: in our directed dyad logit ruling regression, the estimated coefficient on $CR_{\vec{i(nj)},k,t}$, which captures complainant-and-article specific court experience, is negative and strongly significant for WTO articles but insignificantly different from zero for GATT articles, and the hypothesis that the two coefficients are the same is strongly rejected. This provides some evidence that court learning effects in the WTO era may indeed be stronger for WTO than for GATT articles, though our OLS estimates show no statistically significant difference across any of the learning coefficients so we interpret this evidence as at best weak and only suggestive.

³⁵Various stories about a bandwagon effect seem plausible, but the details of court remedies (e.g., how complete they are, whether they apply effectively to 3rd parties) would matter, and as a result it is not obvious whether rulings for or rather against the defendant would be more likely to stimulate follow-up disputes by other claimants. Similar subtleties arise with tit-for-tat effects (see note 33). This points to the value of modeling such effects before going further in investigating their empirical content, a task we leave to future research.

5.3. Alternative Interpretations

Thus far we have interpreted our empirical findings as reflecting the effects of DSB learning, and of DSB learning that embodies a particular scope and form. An important question is whether there are alternative interpretations of these empirical findings. In this section we consider the plausibility of the key alternatives.

We begin with the most narrow version of this question: Can we be sure that, when viewed through the lens of our model, our empirical findings admit only the interpretation we have given them? Put differently, while we do not claim to have structurally estimated the key learning parameters (the β 's) of our model, can the model be used to infer from our empirical findings which of the β 's are positive and which are zero? We argue now that the answer is "Yes." To this end, we return to our multi-country, continuous-policy model of section 4. That model focuses on a single issue area, but the key points can be extended to a setting with multiple issue areas if government payoffs are separable in issue areas. Recall from expression (4.1) that there are five non-negative parameters $(\beta_1, \beta_2, \beta_3, \beta_4, \beta_5)$ describing the nature and scope of court learning, with five corresponding experience variables x_m . Suppose data can be used to estimate the derivatives of the likelihood of rulings and disputes with respect to the x_m 's. We can interpret our regressions as estimating these derivatives: in particular and as we have reported above, according to our empirical findings the likelihood of a ruling where i is the complainant and j the defendant is decreasing in x_{ij} (directed-dyad specific court experience) and in $x_{i(nj)}$ (complainant-specific court experience), while it is essentially independent of the other x_m 's. It can be shown that, according to the model, this implies that β_1 and β_3 are positive while the other β 's are zero.³⁶ It is an extension of this logic to a setting with multiple issues/articles that underlies our statements above that the data is consistent with directed-dyad-specific, complainant-specific and article-specific learning.

We next ask whether there are alternative interpretations of our empirical findings based on alternative models. One plausible candidate is that there is learning going on, but that it takes the form of governments learning about each other. To consider this alternative interpretation,

³⁶Two important qualifications are needed here. First, the statement above is valid under the natural restriction we have assumed on the β 's, namely $\beta_1 \geq \beta_2, \beta_3, \beta_4 \geq \beta_5$, and ignoring non-generic possibilities (and in particular, the knife-edge case where the static effect of x_m on the likelihood of a ruling exactly offsets the dynamic effect through Δ). Second, in the text we discuss only what can be inferred about the β 's from the ruling regression, and not from the dispute regression. The reason is that, as we point out in section 4, the model predictions regarding the likelihood of disputes are somewhat more ambiguous; thus, in the absence of further restrictions it is not clear that one can make inferences on the β s using the dispute regressions.

we have re-run the regressions in Tables 2 and 3 replacing the cumulative-stock-of-ruling CR variables on the right-hand side with analogous “ CS ” variables that measure the cumulative stock of formal consultations (facilitated by the WTO secretariat and held in private between the disputing parties) that *settle* prior to panel formation. If governments learn about each other during these consultations and if this has an important impact on the frequency of subsequent disputes and rulings along similar lines to the DSB learning in our model, we would expect this to show up in negative and significant coefficients on the CS variables pertaining to the dyad of the consulting parties (that is, on the $CS_{ij,k,t}^{\leftarrow}$ and $CS_{ij,nk,t}^{\leftarrow}$ variables in the undirected dyad regressions, and on the $CS_{ij,k,t}^{\rightarrow}$, $CS_{ij,nk,t}^{\rightarrow}$, $CS_{ji,k,t}^{\rightarrow}$ and $CS_{ji,nk,t}^{\rightarrow}$ variables in the directed dyad regressions). In fact, we fail to find any robust evidence for such coefficient estimates.³⁷

A second plausible candidate that could provide an alternative interpretation of *some* of our findings is the impact of legal precedent. Under this interpretation, court rulings help to complete the incomplete WTO contract (as for example in Maggi and Staiger, 2011), so as the stock of rulings accumulate, there are fewer and fewer contingencies that are left uncovered by the contract, thus the frequency of rulings may naturally decrease. More specifically, suppose that a given WTO article k is initially incomplete and is silent about a set M_0 of contingencies, out of the total set of contingencies M . Suppose further that, in each period of time, one contingency is randomly selected out of the set M , and if this contingency is not covered by the contract, the court may be called upon to specify the contractual obligations for this contingency. Then, in this simple scenario, as the stock of rulings accumulates the probability of new rulings goes down. Admittedly, the legal precedent interpretation may well explain our empirical findings with regard to effects we attribute to article-specific learning.³⁸ But importantly, this explanation does not seem compelling as an alternative to our DSB learning story when it comes to our findings regarding defendant- and complainant-specific effects. Hence, we view legal precedent as plausibly being part, but *only* part, of the explanation of our empirical

³⁷One might alternatively conjecture that learning about each other reduces the governments’ negotiation costs in the future (e.g., by eliminating within-dyad persistent private information). This conjecture could be captured within our model in a reduced form way with the assumption that κ rises when governments learn about each other. But it is not hard to show in the context of our model that an increase in κ increases the settlement rate (by increasing the likelihood of disputes and reducing the likelihood of rulings). Contrary to this prediction, the trend in the settlement rate in WTO disputes has been flat or slightly negative.

³⁸An interesting possibility to distinguish between these two interpretations of our findings regarding article-specific effects might be to investigate whether these effects are also present in the early GATT era, when legal precedent was by all accounts not operative (see, for example, the discussion of the views of GATT/WTO legal scholars on this point in Maggi and Staiger, 2011). We view this as a promising avenue for further research.

findings above.

A third candidate is the presence of government learning about the court. It is useful to distinguish between two types of learning within this category. A first possibility is that, by observing how the court operates, governments may learn to better predict the outcome of rulings (a possibility we mentioned previously in section 3). This might be the case, for example, if governments learn about the court’s preferences and possible biases. Our model assumes that the court’s objective is given by the governments’ joint surplus and is common knowledge, but different court objectives are certainly possible in the real world. Intuitively, this type of learning might explain our findings about the impacts of cumulative rulings on the likelihood of current rulings and disputes. However one can view this type of learning as falling into a broader notion of “institutional learning,” and so we view this interpretation as broadly complementary – rather than competing – with our interpretation of judicial learning.

Another possibility is that governments might learn about the court’s inefficiency (λ), and in particular one might hypothesize that, as a result of past rulings, governments have become more pessimistic about the quality of court rulings. This is essentially the “bad news” story we mentioned in the Introduction. One issue that makes this candidate interpretation unappealing is that if one is willing to assume systematically biased prior beliefs, virtually anything can be explained. But even putting this issue aside, while it is possible that a formal version of this story could deliver predictions that match the main features of the data, this is not automatically the case. To see this, consider the simplest two-country version of this story, where governments initially think the court learning curve is $\lambda(x)$, and then they receive bad news that leads them to believe $\lambda(x)$ is higher than previously thought. And let us suppose for simplicity that $\lambda(x)$ shifts up in a way that preserves the initial Δ (or Δ doesn’t change enough to outweigh the static effect). Then the probability of a ruling will go down, but the probability of a dispute will go up (recalling Remark 3); and the second implication is inconsistent with our data (as is the additional implication that the settlement rate will rise – see note 37).³⁹

³⁹For both types of government learning about the court that we have discussed above, if one assumes that a government learns more by participating as a third party in a dispute than by observing the dispute as an outsider, then there is a simple way to gauge whether such learning is important: a government’s experience as third-party participant in disputes that end in ruling should be a stronger predictor of the likelihood of current disputes and rulings than a government’s experience as outside observer. To assess this possibility, we have re-estimated each of our directed-dyad regressions breaking out the cumulative rulings where at least one of the disputing parties i or j was a third-party participant from our “cumulative ruling other” variables $CR_{other,k,t}^{\rightarrow}$ and $CR_{other,nk,t}^{\rightarrow}$. We find that the third-party experience variables do not have a stronger effect than the “other” experience variables; thus by this metric there is no evidence in our data that government learning

It is also possible that the increasing complexity of WTO disputes combined with a fixed resource constraint faced by the WTO court could account in a mechanical way for the overall declines in the numbers of WTO disputes and rulings that are depicted in our Plots 1 and 2.⁴⁰ But our regressions are picking up something more: a resource-constraint story would not predict that declines in disputes and rulings would be systematically related to our *CR* variables, but rather this story would naturally show up along with other factors in our time trends; and even if one were to argue that our *CR* variables somehow reflect this explanation, it would presumably be the general *CR* effect, not a disputant-specific or article-specific *CR* effect, where this explanation would be reflected, and we don't find strong general effects.⁴¹

Finally, our 1995-2009 sample period witnessed unprecedented growth in the numbers of preferential trade agreements (PTAs) to which WTO members belong, and it has been observed that such PTA membership may itself reduce the reliance of countries on the WTO to resolve their trade disputes (this argument is made most forcefully by Mavroidis and Sapir, 2015), offering an alternative PTA-driven interpretation of the decline in WTO disputes and rulings over this period. To explore this alternative interpretation, we have re-estimated the regressions in Tables 2, 3 and 4 with a PTA-partner variable included on the right-hand side. Our results indicate the PTA partnership does indeed reduce the likelihood of WTO disputes that end in rulings (though we find no evidence that PTA partnership reduces the likelihood of WTO disputes); but importantly, inclusion of the PTA partnership variable on the right-hand side of our regressions leaves our findings relating to the cumulative rulings variables unchanged. Hence, while the growth in PTA partnership appears to be a part of the explanation of the decline in WTO disputes and rulings, accounting for it does not diminish the importance of DSB learning as an explanation according to our cumulative ruling variables.

about the court is important.

⁴⁰For example, the average number of stages in a WTO dispute at which a ruling is issued (e.g., panel, appeal, compliance panel) is growing over time. This is one dimension on which the demands placed on the WTO court for handling a given WTO dispute has increased. Plots 1 and 2 indicate the numbers of WTO disputes that are initiated, and the numbers that make it to the ruling stage; but these numbers do not reflect directly the proliferation of ruling stages, and it is possible that this proliferation could interact with a WTO resource constraint to indirectly and mechanically account for the declining numbers depicted in the plots.

⁴¹As a further check on whether a "WTO capacity constraint" story could be an alternative explanation for our results, we have re-estimated all of our regressions with the inclusion of an additional right-hand-side variable that reflects the median resolution time of all disputes over article k initiated in year t , and we find that inclusion of this variable in our regressions has no material impact on our findings.

6. Conclusion

Over the two decades that the WTO has been in existence, the frequency of WTO disputes and court rulings has trended downwards. Such trends are sometimes interpreted as symptoms of a dispute resolution system in decline. In this paper we have proposed a theory that can explain these trends as a result of judicial learning. And according to our theory, such trends represent good news, not bad news. We have also confronted the theory with data from WTO disputes. We interpret our empirical findings as supporting the proposition that court learning is an important phenomenon for understanding the pattern of WTO dispute resolution.

Beyond providing support for the theory, our empirical findings shed some light on the scope and form that learning by ruling may take in the WTO. As interpreted through the lens of our model, we have found robust evidence in the pattern of WTO disputes and rulings that is consistent with article-specific learning and with some forms of disputant-specific learning, but only weak evidence of general-scope learning. And we have argued that our learning-by-ruling model is better able to account for these patterns than simple alternative models. Still, we have only focused empirically on the most central prediction of our theory, and have therefore only scratched the surface of exploring the potential role of court learning in accounting for the dynamics of dispute resolution. Based on the promising results from our initial empirical exploration, a deeper empirical analysis of the impacts of court learning on the dynamics of disputes and rulings seems an important task for future research.

The theory itself can also be extended in interesting ways. For example, we have abstracted from the possibility of free-rider issues in the context of court learning, and the fact that we find only weak evidence of general-scope learning suggests that at least the most extreme free-rider possibilities may not arise in practice. Moreover, one might expect free-rider issues to become more severe over time as the WTO membership has expanded, and yet according to our estimates the time trends in the frequency of WTO disputes and rulings are positive, suggesting at most a more modest role for free-rider effects. But incorporating free-rider issues into our model could nevertheless yield interesting further implications, including predictions about how the frequency of disputes and rulings depend on the probability of future interaction (persistence of matches), the size of countries (bigger countries internalize more the benefits of learning) and the total number of countries in the agreement.

7. Appendix

A. An example model

Here we present a parameterized model that is a special case of the model developed in the main text, with the purpose of illustrating how the reduced-form assumptions we make in that model can be justified in a more “structural” way.

We assume the Home government chooses a tariff T to maximize a weighted welfare function which allows for political economy considerations. In particular, Home’s payoff is

$$\omega(T, s) = CS(T) + R(T) + s \cdot PS(T),$$

where CS is consumer surplus, PS is producer surplus and R is tariff revenue, and where $s \geq 1$ is a parameter that captures the political importance of the domestic industry (in the spirit of Baldwin, 1987, and Grossman and Helpman, 1994). In this parameterized model, the “state of the world” is simply the political parameter s . The Foreign government is passive in this industry (it has no policy of its own) and its payoff is given by national welfare, which in this setting is just the sum of consumer and producer surplus:

$$\omega^*(T) = CS^*(T) + PS^*(T).$$

We assume the demand and supply functions are linear in both countries. Note that $\omega^*(T)$ is decreasing and convex in T ; intuitively, this is because increasing T reduces trade volume, and hence reduces the impact on foreign welfare of further increases in T . On the other hand, note that $\omega(T, s)$ is concave in T provided s is not too high: the reason is that $CS(T) + R(T)$ is concave but $PS(T)$ is convex.

It is easy to verify that the no-transfer government Pareto frontier is concave if $\omega_{TT}(s) + \omega_{TT}^* < 0$, and this is the case if s does not exceed a threshold level.⁴² In what follows we assume this condition is satisfied.

Recall our assumption that the state of the world s is ex-ante uncertain, and its realization is observed by governments but is not verifiable, so governments cannot write a complete contingent contract. Here we further suppose that the contract does not specify the policy T at all (discretion), but the court is endowed with the authority to “fill the gap” of this contract

⁴²Clearly, this threshold level is lower than the threshold level of s below which ω is concave in T , but higher than one (because if $s = 1$ both governments maximize welfare).

ex-post. More specifically, denoting by T^{fb} the “first best” level of T that maximizes joint surplus in state s (and where for notational simplicity we suppress the dependence of T^{fb} on s), the DSB can observe a noisy signal of T^{fb} , given by $T^{dsb} = T^{fb} + \varepsilon$, where ε is a white noise with mean zero and variance λ .⁴³ If invoked, the DSB issues a ruling to maximize the governments’ expected joint payoff conditional on its noisy information. Given the above assumptions, the DSB ruling will prescribe the tariff level T^{dsb} . In this example we abstract from litigation costs, so the court-inefficiency parameter λ captures simply the noise in the DSB signal. Finally, we assume that governments have symmetric bargaining powers when negotiating at stage 3.

We can now characterize the equilibrium outcome of the static model. We focus first on the dispute subgame (stage 3). Given that the no-transfer frontier is concave, the disagreement point for the negotiation is below this frontier as a result of the uncertainty in the DSB ruling; moreover, it lies Southeast of the FB point, because the uncertainty in the ruling hurts the importer (whose payoff is concave in T) and benefits the exporter (whose payoff is convex in T). Given that payoffs are quadratic and bargaining powers are symmetric, it is direct to verify that the expected disagreement payoffs are given by

$$\begin{aligned}\omega^R(s; \lambda) &= \omega^{fb}(s) + \lambda \cdot \omega_{TT}(s), \\ \omega^{*R}(s; \lambda) &= \omega^{*fb}(s) + \lambda \cdot \omega_{TT}^*,\end{aligned}$$

where $\omega^{fb}(s) \equiv \omega(T^{fb}(s), s)$ and $\omega^{*fb}(s) \equiv \omega^*(T^{fb}(s))$. Recall from the discussion above that $\omega_{TT}(s) < 0$, $\omega_{TT}^* > 0$ and $\omega_{TT}(s) + \omega_{TT}^* < 0$ for all s . Thus, increasing the DSB noise λ worsens Home’s disagreement payoff, improves Foreign’s disagreement payoff, and worsens the joint disagreement payoff.

Note that: (i) $\Omega^R(s; \lambda) = \Omega^{fb}(s) + \lambda[\omega_{TT}(s) + \omega_{TT}^*]$ is decreasing in λ for any s , and $\Omega^R(s; 0) = \Omega^{fb}(s)$, thus our reduced-form Assumption 1 is satisfied; (ii) $\omega^{*R}(s; \lambda)$ is clearly monotonic in λ for all s , thus our reduced-form Assumption 2 is also satisfied. Thus Remark 2 holds in this example model.

We now focus on the occurrence of disputes in the static setting. It is easy to derive the net bargaining payoffs as

$$\begin{aligned}\omega^{Bnet}(s; \lambda) &= \omega^{fb}(s) + \lambda \cdot \left[\left(1 - \frac{\kappa}{2}\right) \omega_{TT}(s) - \frac{\kappa}{2} \omega_{TT}^* \right], \\ \omega^{*Bnet}(s; \lambda) &= \omega^{*fb}(s) + \lambda \cdot \left[\left(1 - \frac{\kappa}{2}\right) \omega_{TT}^* - \frac{\kappa}{2} \omega_{TT}(s) \right].\end{aligned}$$

⁴³We could assume that the court observes a noisy signal of s rather than a noisy signal of T^{fb} , at the cost of a slightly more complicated analysis.

Clearly, these payoffs are linear in λ for any s , hence the B_{net}^λ locus is linear and can cross the no-transfer frontier at most once. It follows that Remark 3 holds in this example model.

Finally consider the dynamic setting. Here we do not need to impose any additional structure relative to our more general model. Note that, while Assumption 4 needs to be imposed just as in the more general model, Assumption 3 is automatically satisfied here, because net bargaining payoffs are linear in λ and λ is convex in x . As discussed in the main text, Assumption 4 directs attention to what seems to be the empirically relevant case given that the frequency of disputes is empirically very low.

B. Derivation of $\Delta^i + \Delta^j$ in the multi-country setting

Consider first the static version of our extended setting. A key observation is that, for a given country pair (i, j) , the equilibrium joint payoff is decreasing in λ_{ij} and λ_{ji} . This follows because Remark 2 and the discussion preceding it apply also to our extended setting.

Next consider the dynamic setting. Suppose that at $t = 1$ a dispute occurs in which country i is the exporter/complainant and country j is the importer/defendant, and consider the future impacts of a ruling in this dispute. At $t = 2$ there are six possibilities that we need to consider for the future impacts on countries i and j : (i) with probability P^{ij} , there will be the potential for a dispute with country i the complainant and country j the defendant, in which case the relevant court experience X_{ij} increases by an amount β_1 ; (ii) with probability P^{ji} there will be the potential for a dispute with j the complainant and i the defendant, in which case X_{ji} increases by an amount β_2 ; (iii) with probability P^{io} there will be the potential for a dispute with i the complainant and some third country o the defendant, in which case X_{io} increases by an amount β_3 ; (iv) with probability P^{oj} there will be the potential for a dispute with j the defendant and o the complainant, in which case X_{oj} increases by an amount β_4 ; (v) with probability P^{oi} there will be the potential for a dispute with o the complainant and i the defendant, in which case X_{oi} increases by an amount β_5 ; and (vi) with probability P^{jo} there will be the potential for a dispute with j the complainant and o the defendant, in which case X_{jo} increases by an amount β_5 .

In light of the above considerations, we can write the joint future payoff gain for today's disputants, $\Delta^i + \Delta^j$, as

$$\begin{aligned}
\Delta^i + \Delta^j = & \\
& P^{ij} \{ [\tilde{\omega}_i^{ij}(X_{ij} + \beta_1) + \tilde{\omega}_j^{ij}(X_{ij} + \beta_1)] - [\tilde{\omega}_i^{ij}(X_{ij}) + \tilde{\omega}_j^{ij}(X_{ij})] \} \\
& + P^{ji} \{ [\tilde{\omega}_i^{ji}(X_{ji} + \beta_2) + \tilde{\omega}_j^{ji}(X_{ji} + \beta_2)] - [\tilde{\omega}_i^{ji}(X_{ji}) + \tilde{\omega}_j^{ji}(X_{ji})] \} \\
& + \sum_{o \neq i, j} P^{io} [\tilde{\omega}_i^{io}(X_{io} + \beta_3) - \tilde{\omega}_i^{io}(X_{io})] + \sum_{o \neq i, j} P^{oj} [\tilde{\omega}_j^{oj}(X_{oj} + \beta_4) - \tilde{\omega}_j^{oj}(X_{oj})] \\
& + \sum_{o \neq i, j} P^{oi} [\tilde{\omega}_i^{oi}(X_{oi} + \beta_5) - \tilde{\omega}_i^{oi}(X_{oi})] + \sum_{o \neq i, j} P^{jo} [\tilde{\omega}_j^{jo}(X_{jo} + \beta_5) - \tilde{\omega}_j^{jo}(X_{jo})],
\end{aligned}$$

where $\tilde{\omega}_r^{rs}(\cdot)$ and $\tilde{\omega}_r^{sr}(\cdot)$ denote the expected equilibrium (sub-)payoffs for country r when country r is respectively the period-2 complainant or period-2 defendant in a dispute with country s , and where we use the shorthands $\tilde{\omega}_r^{rs}(X_{rs}) \equiv \tilde{\omega}_r^{rs}(\lambda(X_{rs}))$ and $\tilde{\omega}_r^{sr}(X_{sr}) \equiv \tilde{\omega}_r^{sr}(\lambda(X_{sr}))$.

As the expression for $\Delta^i + \Delta^j$ makes clear, if both complainant and defendant benefit from higher court quality, then $\Delta^i + \Delta^j$ must be positive, because every term in the expression for $\Delta^i + \Delta^j$ is then positive. But consider the case where court quality has opposing impacts on the defendant and the complainant: then the $P^{ij}\{\cdot\}$ and $P^{ji}\{\cdot\}$ terms in the expression above are positive, but of the four additional terms, up to two could be negative.⁴⁴ Thus, in principle $\Delta^i + \Delta^j$ could be negative. As mentioned in the main text, this ambiguity can be resolved by assuming that $P^{ij} + P^{ji}$ is sufficiently close to one for each dyad ij . This ensures that the four “good” effects dominate the two “bad” effects, and hence $\Delta^i + \Delta^j$ is positive.

Next note that, under the assumption that $P^{ij} + P^{ji}$ is sufficiently close to one (and the analogs of Assumptions 1-4), $\Delta^i + \Delta^j$ is decreasing in each of the x_m 's. This follows immediately from the fact that the $P^{ij}\{\cdot\}$ and $P^{ji}\{\cdot\}$ terms dominate the remaining terms in the expression for $\Delta^i + \Delta^j$, and that the expected joint payoff $\tilde{\omega}_i^{ij} + \tilde{\omega}_j^{ij}$ is concave in X_{ij} .

In light of the above arguments, it is clear that if δ is large enough the probability of a ruling is decreasing in each x_m . Furthermore, it is also clear that, regardless of δ , if X_{ij} is large enough the probability of a ruling is decreasing in each x_m , because the probability of a ruling eventually must hit zero as X_{ij} gets large.⁴⁵

⁴⁴This follows because these four additional terms are associated with cases where either i or j is either a defendant or a complainant with a third party. If complainants are hurt by an increase in DSB efficiency, then the two defendant terms must be positive, because the joint defendant-and-complainant payoff increases with the increase in DSB efficiency. And if defendants are hurt by an increase in DSB efficiency, then the two complainant terms must be positive, for the same reason.

⁴⁵A more formal statement of this result is that there exist critical values x_m^1 and x_m^2 (with $x_m^1 < x_m^2$) such that the probability of a ruling is strictly decreasing for $x_m \in (x_m^1, x_m^2)$ and equal to zero for $x_m \in (x_m^2, \infty)$.

Finally, turning to the impact of x_m on the likelihood of a dispute, we can show that this impact in general is ambiguous. Suppose for instance that only β_1 is positive (directed-dyad specific learning). Suppose further that the B_{Δ}^{net} point is at the dispute margin, i.e. on the no-transfer frontier, and suppose δ is large, so that we can ignore the static effect and focus on the dynamic effect. And finally, suppose that an increase in DSB efficiency is beneficial to the defendant but hurts the complainant, as is possible (see note 10). Consider an increase in x_{ij} or x_{ji} : as we argued above, in this case $\Delta_{ij}^i < 0$, $\Delta_{ij}^j > 0$ and $\Delta_{ij}^i + \Delta_{ij}^j > 0$, and furthermore, increasing x_{ij} or x_{ji} reduces $\Delta_{ij}^i + \Delta_{ij}^j$. It can also be shown that increasing x_{ij} or x_{ji} reduces both Δ_{ij}^i and Δ_{ij}^j in absolute value. This implies that the $R + \delta\Delta$ point moves Northwest with slope steeper than -1 , and so does the B_{Δ}^{net} point. This could lead the B_{Δ}^{net} point to dip below the no-transfer frontier or to rise above it, thus the impact on the likelihood of a dispute is ambiguous.

8. Data Appendix

The data used in this paper comes from the WTO Dispute Settlement Database (see Horn, Johannesson and Mavroidis, 2011 for a description). This data set is maintained by the World Bank, and its current coverage includes each of the 426 documented WTO disputes between 1995 and August 2011.⁴⁶ We exclude from our analysis the 24 disputes that were initiated after January 1 2010 (to avoid truncation of dispute outcomes in the dataset); and we exclude as well the 8 cases where the issue returns in a later dispute (which we include) or is simply handled in another dispute (which we include).⁴⁷ And finally, we drop the 6 multi-complainant cases in this dataset that were each treated as a single dispute by the WTO (i.e., each of the claimants against the common respondent was listed under the same WTO dispute number), on the grounds that these cases reflect especially tight links across the claimants that would likely impact dispute behavior through channels about which our model is silent.⁴⁸ After this set of exclusions we are left with 348 WTO disputes.

⁴⁶Each dispute is associated with a unique DSnumber, which is the official case number recorded in WTO documents. Thus this data set includes disputes from DS1 to DS426.

⁴⁷The 8 excluded cases are DS3, DS16, DS52, DS101, DS106, DS185, DS228, DS271, which respectively return or are handled in DS41, DS27, DS65, DS132, DS126, DS187, DS230, DS270.

⁴⁸The 6 excluded cases are DS27, DS35, DS58, DS158, DS217, DS234.

References

- Baker, Scott and Claudio Mezzetti (2012), "A Theory of Rational Jurisprudence," *Journal of Political Economy* 120 (3): 513 - 551.
- Baldwin, Richard (1987), "Politically Realistic Objective Functions and Trade Policy," *Economics Letters* 24: 287-90.
- Bebchuck, Lucian A. (1984), "Litigation and Settlement under Imperfect Information," *RAND Journal of Economics* 15(3): Autumn, 404-415.
- Beim, D. (2017), "Learning in the Judicial Hierarchy", *Journal of Politics* 79(2).
- Benkard, C. Lanier (2000), "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review* 90:1034–1054.
- Bernard, Andrew B. and J. Bradford Jensen (1999), "Exceptional Exporter Performance: Cause, Effect or Both?" *Journal of International Economics* 47 (1):1–25.
- Beshkar, Mostafa (2016), "Arbitration and Renegotiation in Trade Agreements," *Journal of Law, Economics and Organization* 32.
- Bown, Chad (2005), "Participation in WTO Dispute Settlement: Complainants, Interested Parties and Free Riders," *World Bank Economic Review* 19 (2)
- Bown, Chad. P., and Kara. M. Reynolds (2015), "Trade Flows and Trade Disputes," *Review of International Organizations* 10.
- Busch, Mark L. (2000), "Democracy, Consultation, and the Paneling of Disputes under GATT," *Journal of Conflict Resolution* 44 (4).
- Busch, M. L. and E. Reinhardt (2000), "Bargaining in the Shadow of the Law: Early Settlement in GATT/WTO Disputes." *Fordham International Law Journal* 24 (1): 158-172.
- Busch, Mark L. and Eric Reinhardt (2006), "Three's a Crowd: Third Parties and WTO Dispute Settlement," *World Politics* 58: 446-77.

- Clerides, Sofronis K., Lach, Saul and James R. Tybout (1998), "Is Learning by Exporting Important? Micro-Dynamic Evidence from Colombia, Mexico, and Morocco," *Quarterly Journal of Economics*, August: 903-947.
- Conconi, Paola, David R. DeRemer, Georg Kirchsteiger and Lorenzo Trimarchi (2017), "Suspiciously Timed Trade Disputes", *Journal of International Economics* 105: 57-76.
- Davis, Christina (2012), *Why Adjudicate? Enforcing Trade Rules in the WTO*. Princeton University Press.
- Davis, Christina L. and Sarah Blodgett Bermeo (2009), "Who files? Developing Country Participation in GATT/WTO Adjudication," *Journal of Politics* 71 (3).
- Freedman, Jennifer (2012), "China Floods the WTO with Tit-for-Tat," *Bloomberg*, June 7.
- Grossman, Gene M. and Elhanan Helpman (1994), "Protection for Sale," *American Economic Review* 84(4): 833-850.
- Guzman, Andrew and Beth A. Simmons (2002), "To Settle or Empanel? An Empirical Analysis of Litigation and Settlement at the World Trade Organization," *Journal of Legal Studies*, January, S205-S235.
- Guzman, Andrew T. and Beth A. Simmons (2005), "Power Plays and Capacity Constraints: The Selection of Defendants in World Trade Organization Disputes," *Journal of Legal Studies* 34 (2): 557-598.
- Horn, Henrik, Louise Johannesson and Petros C. Mavroidis (2011), "The WTO Dispute Settlement System 1995-2006: Some Descriptive Statistics," *Journal of World Trade* 45(6).
- Hudec, Robert E. (1993), *Enforcing International Trade Law: The Evolution of the Modern GATT Legal System*, Butterworth Legal Publishers: USA.
- Irwin, Douglas A. and Peter J. Klenow (1994), "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy* 102 (6): 1200-1227.
- Kellogg, Ryan (2011), Learning by Drilling: Interfirm Learning and Relationship Persistence in the Texas Oilpatch," *Quarterly Journal of Economics* 126:1961-2004.

- Kuenzel, David J. (2017), “WTO Dispute Determinants”, *European Economic Review* 91: 157-179.
- Levitt, Steven D., List, John A. and Chad Syverson (2013), “Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant,” *Journal of Political Economy* 121(4): 643-681.
- Maggi, Giovanni and Robert W. Staiger (2011), “The Role of Dispute Settlement Procedures in International Trade Agreements,” *Quarterly Journal of Economics* 126: 475–515.
- Maggi, Giovanni and Robert W. Staiger (2015), “Optimal Design of Trade Agreements in the Presence of Renegotiation,” *American Economic Journal: Microeconomics* 7(1): 109–143.
- Maggi, Giovanni and Robert W. Staiger (forthcoming), “Trade Disputes and Settlement”, *International Economic Review*.
- Mavroidis, Petros C. and Andre Sapir (2015), “Dial PTAs for Peace: The Influence of Preferential Trade Agreements on Litigation between Trading Partners,” *Journal of World Trade* 49(3): 351-374.
- Park, Jee-Hyeong (2011), “Enforcing International Trade Agreements with Imperfect Private Monitoring,” *Review of Economic Studies*.
- Reinganum, Jennifer F. and Louis L. Wilde (1986), “Settlement, litigation, and the allocation of litigation costs,” *RAND Journal of Economics* 17(4): Winter, 557-566.
- Staiger, Robert W. and Alan O. Sykes (2017), “How Important can the Non-Violation Clause be for the GATT/WTO?,” *American Economic Journal: Microeconomics* 9(2): 149-187.
- Stevens, John Paul (2006), “Learning on the Job,” *Fordham Law Review* 74: 1561-7.
- Thornton, Rebecca A. and Peter Thompson (2001), “Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding,” *American Economic Review* 9: 1350–1368.
- Wooldridge, Jeffrey M (2010), *Econometric Analysis of Cross Section and Panel Data*, Second Edition; The MIT Press, Cambridge Mass.

Figure 1: Equilibrium payoffs in the static setting




Figure 2: Impact of court quality in the static setting


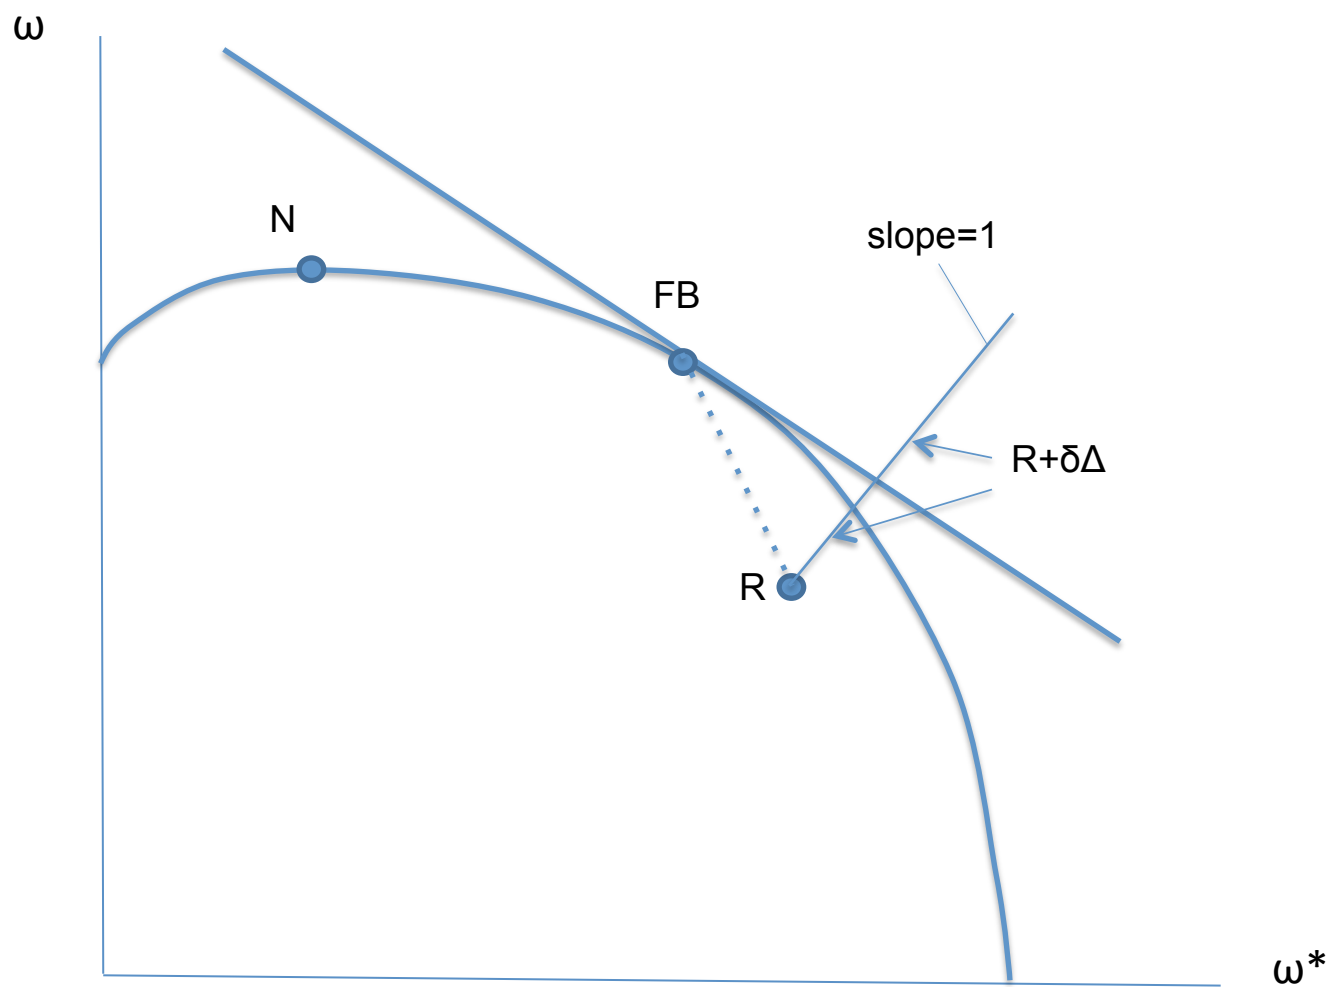
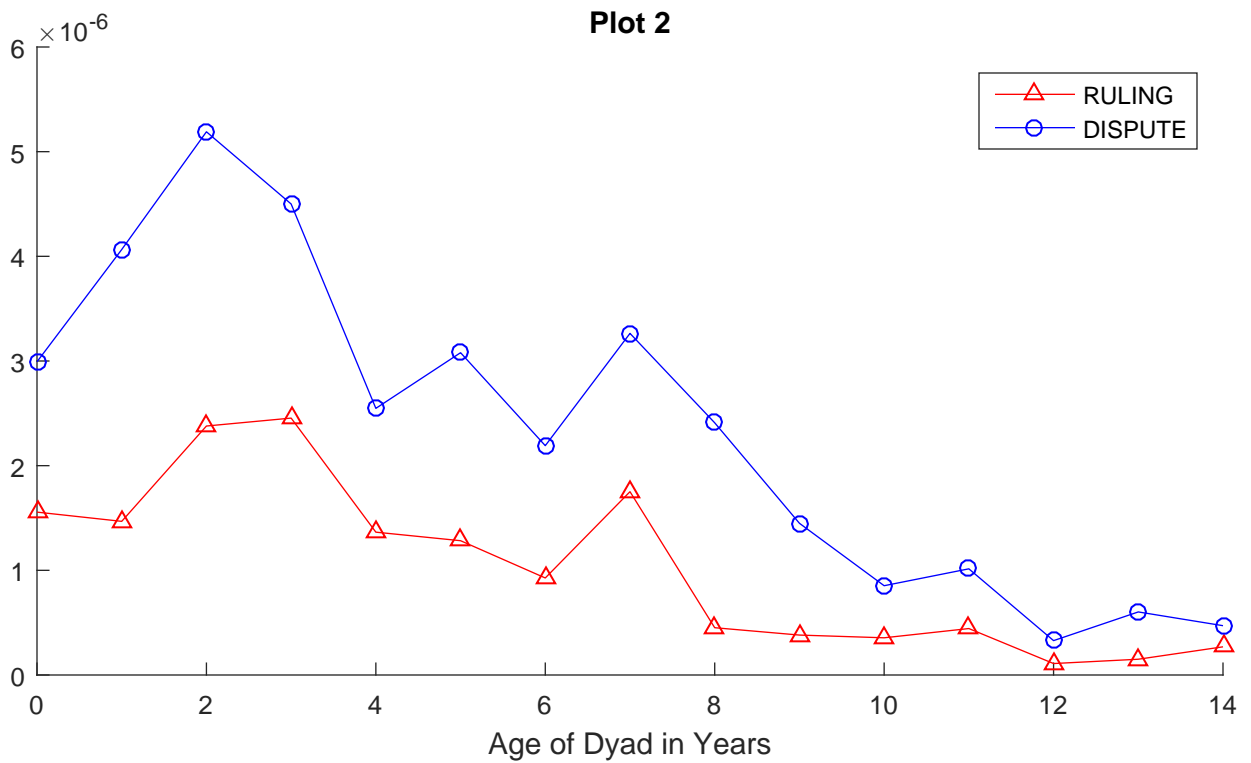
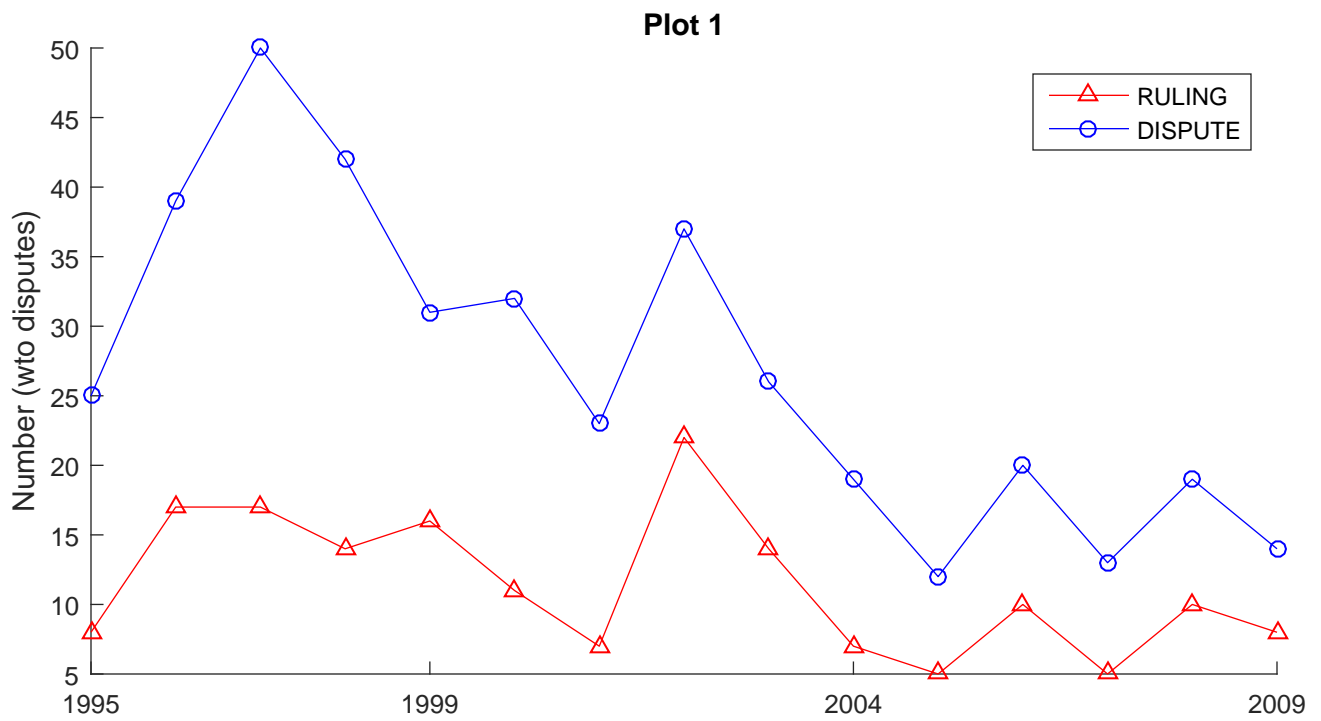


Figure 3: Two-period setting





Note: The vertical axis records the ratio X/Z , where X is the average number of disputes or rulings involving dyads of a given age, and Z is the average trade volume of dyads in this age group.


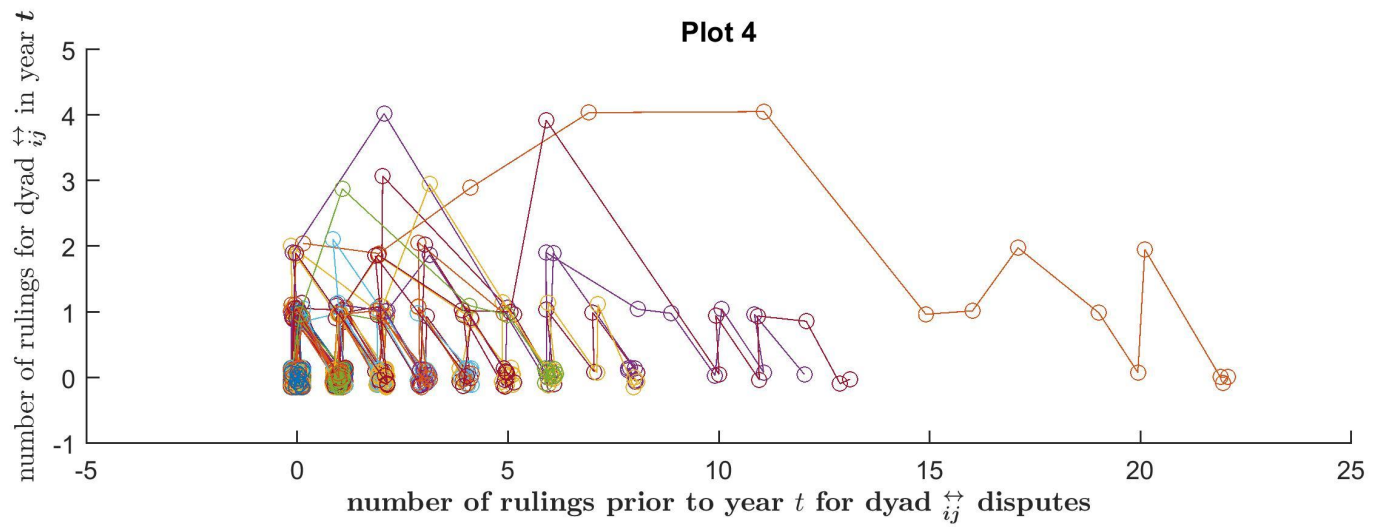
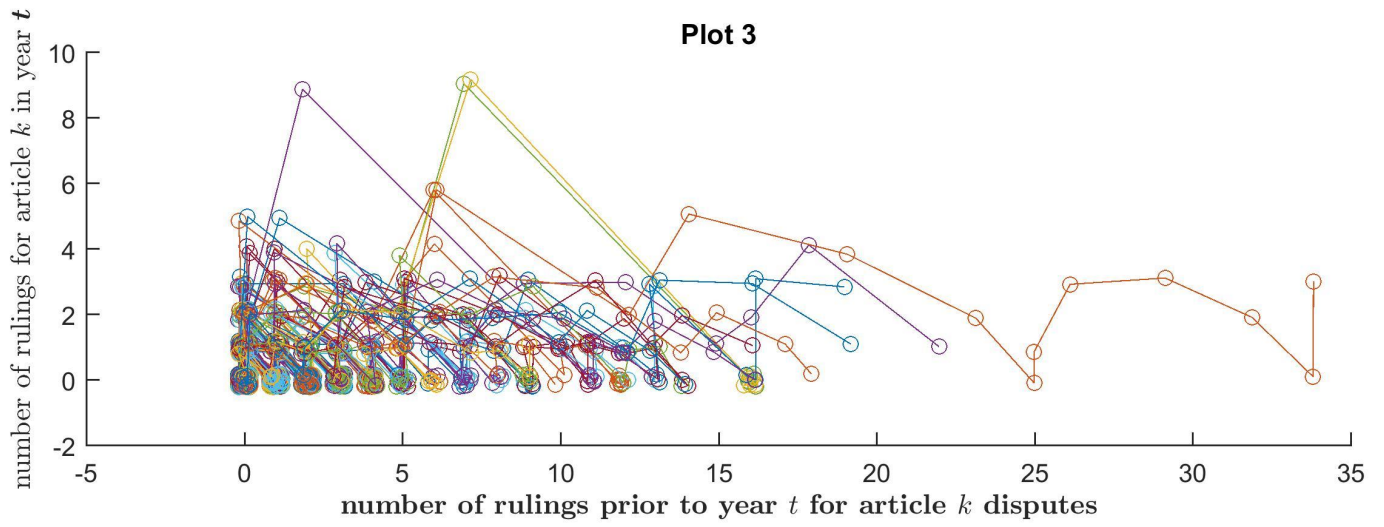


Table 1: Summary Statistics

VARIABLES		Dispute Regression					Ruling Regression				
		N	mean	sd	min	max	N	mean	sd	min	max
Undirected Dyads	Dependent Variable (logit)	439,584	0.00587	0.0764	0	1	112,560	0.00565	0.0750	0	1
	Dependent Variable (ols)	439,584	0.00635	0.0863	0	5	112,560	0.00589	0.0797	0	3
	$CR_{\overrightarrow{ij},k,t}$	439,584	0.00885	0.111	0	5	112,560	0.0346	0.218	0	5
	$CR_{n(\overrightarrow{ij}),k,t}$	439,584	1.101	2.925	0	34	112,560	1.861	3.589	0	34
	$CR_{\overrightarrow{ij},nk,t}$	439,584	2.123	6.933	0	78	112,560	4.803	9.772	0	78
	$CR_{n(\overrightarrow{ij}),nk,t}$	439,584	264.3	221.0	0	631	112,560	258.7	218.2	0	631
Directed Dyads	Dependent Variable (logit)	545,142	0.00486	0.0696	0	1	149,520	0.00431	0.0655	0	1
	Dependent Variable (ols)	545,142	0.00512	0.0755	0	5	149,520	0.00443	0.0684	0	3
	$CR_{\overrightarrow{ij},k,t}$	545,142	0.00713	0.0970	0	5	149,520	0.0260	0.184	0	5
	$CR_{(\overrightarrow{ni})j,k,t}$	545,142	0.00567	0.0879	0	5	149,520	0.0187	0.161	0	5
	$CR_{\overrightarrow{i(nj)},k,t}$	545,142	0.0920	0.602	0	13	149,520	0.246	0.993	0	13
	$CR_{\overrightarrow{ji},k,t}$	545,142	0.0661	0.392	0	10	149,520	0.143	0.572	0	10
	$CR_{\overrightarrow{other},k,t}$	545,142	0.938	2.583	0	34	149,520	1.460	2.990	0	34
	$CR_{\overrightarrow{ij},nk,t}$	545,142	1.712	5.197	0	66	149,520	3.615	7.069	0	66
	$CR_{(\overrightarrow{ni})j,nk,t}$	545,142	1.361	5.104	0	66	149,520	2.598	6.999	0	66
	$CR_{\overrightarrow{i(nj)},nk,t}$	545,142	22.07	51.90	0	265	149,520	34.17	63.52	0	264
	$CR_{\overrightarrow{ji},nk,t}$	545,142	15.87	26.93	0	117	149,520	19.93	29.39	0	117
	$CR_{\overrightarrow{other},nk,t}$	545,142	225.0	194.4	0	631	149,520	202.9	177.4	0	631

Table 2: Logit

VARIABLES	Undirected Dyad		VARIABLES	Directed Dyad	
	(1) <i>DLogit</i> \overrightarrow{ij},k,t	(2) <i>RLogit</i> \overrightarrow{ij},k,t		(3) <i>DLogit</i> \overrightarrow{ij},k,t	(4) <i>RLogit</i> \overrightarrow{ij},k,t
$CR_{\overrightarrow{ij},k,t}$	0.186*** (0.0646)	0.119 (0.0913)	$CR_{\overrightarrow{ij},k,t}$	0.209*** (0.0786)	0.235** (0.0999)
$CR_{n(\overrightarrow{ij}),k,t}$	-0.0127* (0.00761)	-0.160*** (0.0168)	$CR_{\overrightarrow{(ni)j},k,t}$	0.117*** (0.0186)	-0.0479 (0.0299)
$CR_{\overrightarrow{ij},nk,t}$	-0.0272*** (0.00256)	-0.0644*** (0.00540)	$CR_{\overrightarrow{i(nj)},k,t}$	0.0293 (0.0277)	-0.0799 (0.0502)
$CR_{n(\overrightarrow{ij}),nk,t}$	-0.00190** (0.000841)	-0.00663*** (0.00199)	$CR_{\overrightarrow{ji},k,t}$	-0.0128 (0.0941)	-0.416** (0.184)
			$CR_{\overrightarrow{other},k,t}$	-0.0374*** (0.00844)	-0.196*** (0.0193)
			$CR_{\overrightarrow{ij},nk,t}$	-0.0509*** (0.00410)	-0.102*** (0.00775)
			$CR_{\overrightarrow{(ni)j},nk,t}$	0.000918 (0.00109)	0.00122 (0.00253)
			$CR_{\overrightarrow{i(nj)},nk,t}$	-0.0107*** (0.00188)	-0.0212*** (0.00448)
			$CR_{\overrightarrow{ji},nk,t}$	-0.000515 (0.00362)	-0.0248*** (0.00912)
			$CR_{\overrightarrow{other},nk,t}$	-0.00130 (0.000834)	-0.00691*** (0.00203)
t	0.198*** (0.0260)	1.277*** (0.111)	t	0.235*** (0.0258)	1.409*** (0.117)
t^2	-0.00726*** (0.00225)	-0.0313*** (0.00416)	t^2	-0.0104*** (0.00225)	-0.0373*** (0.00433)
Constant	-9.201*** (1.076)	-10.90*** (0.723)	Constant	-9.701*** (1.077)	-12.87*** (0.792)
Observations	439,584	112,560	Observations	545,142	149,520
Pseudo R ²	0.277	0.229	Pseudo R ²	0.242	0.218
\overrightarrow{ij} FE	Y	Y	\overrightarrow{ij} FE	Y	Y
kFE	Y	Y	kFE	Y	Y

Standard errors in parentheses

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

\overrightarrow{ij} FE are undirected dyad fixed effects. kFE are article fixed effects.

Table 3: OLS

VARIABLES	Undirected Dyad		VARIABLES	Directed Dyad	
	(1) $D_{\vec{ij},k,t}$	(2) $R_{\vec{ij},k,t}$		(3) $D_{\vec{ij},k,t}$	(4) $R_{\vec{ij},k,t}$
$CR_{\vec{ij},k,t}$	0.0345*** (0.00734)	0.0138*** (0.00470)	$CR_{\vec{ij},k,t}$	0.0150*** (0.00525)	0.0119** (0.00549)
$CR_{n(\vec{ij}),k,t}$	-0.00140*** (0.000474)	-0.000849** (0.000354)	$CR_{(\vec{ni})j,k,t}$	0.00288** (0.00127)	0.000884 (0.000797)
$CR_{\vec{ij},nk,t}$	-0.00135*** (0.000362)	-0.000525*** (0.000126)	$CR_{i(\vec{nj}),k,t}$	0.00100 (0.00139)	0.000407 (0.000927)
$CR_{n(\vec{ij}),nk,t}$	-1.09e-05 (1.32e-05)	-8.13e-06 (1.76e-05)	$CR_{\vec{ji},k,t}$	0.0128** (0.00493)	-0.00122 (0.00250)
			$CR_{\vec{other},k,t}$	-0.00160*** (0.000374)	-0.000957*** (0.000254)
			$CR_{\vec{ij},nk,t}$	-0.00117*** (0.000188)	-0.000713*** (0.000187)
			$CR_{(\vec{ni})j,nk,t}$	-5.47e-06 (1.28e-05)	6.72e-06 (2.23e-05)
			$CR_{i(\vec{nj}),nk,t}$	-5.33e-05** (2.13e-05)	-5.12e-05* (2.82e-05)
			$CR_{\vec{ji},nk,t}$	-0.000199 (0.000261)	8.92e-05 (0.000103)
			$CR_{\vec{other},nk,t}$	-6.42e-06 (1.01e-05)	-3.11e-06 (1.55e-05)
t	0.00130** (0.000603)	0.00317*** (0.000699)	t	0.00112** (0.000435)	0.00249*** (0.000521)
t^2	-1.84e-05 (3.17e-05)	-0.000106 (6.77e-05)	t^2	-1.90e-05 (2.44e-05)	-8.47e-05* (5.00e-05)
Constant	-0.00806** (0.00329)	0.00311 (0.00377)	Constant	-0.00820*** (0.00256)	-6.12e-06 (0.00289)
Observations	439,584	112,560	Observations	545,142	149,520
R ²	0.044	0.021	R ²	0.029	0.017
\vec{ij} FE	Y	Y	\vec{ij} FE	Y	Y
kFE	Y	Y	kFE	Y	Y
CE	\vec{ij}	\vec{ij}	CE	\vec{ij}	\vec{ij}

Standard errors in parentheses

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

\vec{ij} FE (\vec{ij} FE) are undirected (directed) dyad fixed effects. kFE are article fixed effects.

CE are standard errors clustered by undirected or directed dyads

Table 4: Dispute Regressions with Dyad-and-Article Fixed Effects

VARIABLES	Undirected Dyad		VARIABLES	Directed Dyad	
	(1)	(2)		(3)	(4)
	LOGIT	OLS		LOGIT	OLS
	$DLogit_{\vec{ij},k,t}$	$D_{\vec{ij},k,t}$		$DLogit_{\vec{ij},k,t}$	$D_{\vec{ij},k,t}$
$CR_{\vec{ij},k,t}$	-0.932*** (0.0853)	-0.148*** (0.0128)	$CR_{\vec{ij},k,t}$	-1.262*** (0.102)	-0.152*** (0.0115)
$CR_{n(\vec{ij}),k,t}$	0.0236*** (0.00835)	3.07e-06 (0.000268)	$CR_{(\vec{ni})j,k,t}$	0.0582** (0.0276)	0.00207* (0.00113)
$CR_{\vec{ij},nk,t}$	-0.0134*** (0.00254)	-0.000593* (0.000304)	$CR_{i(\vec{nj}),k,t}$	-0.0101 (0.0441)	-0.00392** (0.00189)
$CR_{n(\vec{ij}),nk,t}$	-0.00219*** (0.000845)	-1.68e-05 (1.38e-05)	$CR_{\vec{ji},k,t}$	0.182 (0.133)	0.00414 (0.00773)
			$CR_{\vec{other},k,t}$	0.00778 (0.0108)	5.52e-05 (0.000210)
			$CR_{\vec{ij},nk,t}$	-0.0276*** (0.00403)	-0.000476*** (0.000144)
			$CR_{(\vec{ni})j,nk,t}$	0.00177 (0.00117)	-2.09e-06 (1.28e-05)
			$CR_{i(\vec{nj}),nk,t}$	-0.0101*** (0.00198)	-3.27e-05* (1.96e-05)
			$CR_{\vec{ji},nk,t}$	-0.00204 (0.00371)	-0.000163 (0.000245)
			$CR_{\vec{other},nk,t}$	-0.00178** (0.000846)	-1.33e-05 (1.05e-05)
t	0.204*** (0.0262)	0.00130** (0.000603)	t	0.246*** (0.0262)	0.00112** (0.000435)
t^2	-0.00743*** (0.00225)	-1.84e-05 (3.16e-05)	t^2	-0.0106*** (0.00226)	-1.90e-05 (2.44e-05)
Constant		0.00428** (0.00181)	Constant		0.00332** (0.00136)
Observations	26,253	439,584	Observations	29,193	545,142
(Pseudo) R ²	0.0392	0.024	(Pseudo) R ²	0.0516	0.023
\vec{ij} kFE	Y	Y	\vec{ij} kFE	Y	Y
CE	N	\vec{ij}	CE	N	\vec{ij}

Standard errors in parentheses

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

\vec{ij} kFE (\vec{ij} kFE) are undirected (directed) dyad-and-article fixed effects.

CE are standard errors clustered by undirected (\vec{ij}) or directed (\vec{ij}) dyads