

LEARNING BY RULING AND TRADE DISPUTES*

Giovanni Maggi[†]

Robert W. Staiger[‡]

April 2020

Abstract

We explore the implications of judicial learning for trade disputes through a model where both the initiation of disputes and the occurrence of rulings are endogenous, governments bargain “in the shadow of the law,” and the efficiency of the court increases with experience. Judicial learning can explain litigation on the equilibrium path, since going to court today implies future payoff gains for the governments. Our model predicts that where learning is present the likelihood of both disputes and rulings should tend to decrease with court experience. Using detailed data on WTO disputes, we find evidence consistent with significant judicial learning at the WTO, but this learning appears to be article-specific and disputant-specific, rather than general, in scope.

*We thank Arnaud Costinot (the Editor) and two anonymous referees for very helpful comments and suggestions. We have also benefited from the detailed comments of Kala Krishna, Andres Rodriguez-Clare, and Gene Grossman, our discussants at various conferences, and from discussions with David Atkin, Chad Bown, Petros C. Mavroidis, Doug Staiger and participants at various seminars. Junhui Zeng and Winston Chen provided outstanding research assistance.

[†]Department of Economics, Yale University; FGV EPGE; and NBER.

[‡]Department of Economics, Dartmouth College; and NBER.

1. Introduction

The World Trade Organization (WTO) is endowed with a sophisticated court system contained within its Dispute Settlement Body (DSB). The adjudication of trade disputes is a complex task, however, and the WTO's DSB is a relatively young judicial system, having come into being on January 1 1995. In such an environment, it seems likely that a court would learn valuable lessons from its early experience, and that this learning would contribute importantly to the evolution of its performance over time.

Judged by trends in the frequency of new disputes and DSB rulings, the evolution of the performance of the WTO's court system displays some interesting patterns. Plot 1 shows the raw annual numbers of disputes initiated and DSB rulings issued over the WTO years 1995-2009.¹ 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 have engaged in fewer "fights" at the WTO as the institution has aged. If anything, Plot 1 understates this trend because the number of WTO members has increased over the last 20 years, 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 older dyads initiate fewer disputes and end up less frequently in court.

These declining trends could be interpreted in different ways. One interpretation is that they signal bad news for the WTO as an institution, in the sense that the DSB is becoming less effective – or that governments at any rate are losing confidence in it – and therefore less relevant over time.² While this is certainly a possibility, in this paper we will propose a theory that can explain the declining trends in disputes and rulings as a result of court learning. According to our theory, these trends represent good news, not bad news. And guided by the core predictions of our theory, we show that the data on WTO dispute settlement is consistent

¹This is the time period covered by the WTO Dispute Settlement Database (see Horn, Johannesson and Mavroidis, 2001, for a description), whose detailed information on WTO disputes we utilize for the empirical analysis in this paper.

²This is an interpretation adopted by a number of legal scholars. For example, Pauwelyn and Zhang (2018) document the declining trend in numbers of new WTO disputes and rulings (p. 17), and suggest that these trends may be driven by potential claimants who are discouraged from filing new cases due to the increasing delays that are being experienced in the process of adjudicating increasingly complex WTO disputes (p. 14). Similarly, in 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.

with this interpretation.

In domestic legal settings, the importance of court learning has been widely emphasized, both by legal practitioners and by the law and economics literature. 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 how early differences of opinion and accumulated experience on the bench shaped the opinions of Justices Holmes and Brandeis in legal rulings 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).

The law and economics literature has also emphasized the importance of legal precedent and its impact on the evolution of legal rules in the context of common law systems, stimulated by Posner's (1973) claim that in the presence of precedent "common law will effect order from chaos" and generate a process that converges to efficient legal rules. One can think of the underlying process described by Posner as a "collective" form of court learning.

In the context of the WTO's DSB, such learning is no doubt also present. At the level of the Appellate Body (which is a standing body of seven judges) where an important function is to interpret and clarify the terms of the treaty "in the light of its object and purpose," judges learn to better interpret WTO treaty provisions from their own accumulated case-load experience as well as from required exchanges over differences of opinion with each other (Zdouc, 2008, p. 271). Judges may also learn from their ruling experience to use and interpret data and to make more effective and timely use of rigorous economic reasoning in arriving at their rulings. Similar forms of court learning can occur also at the level of the lower courts' rotating Panel judges, though here the primary source of learning relies on the accumulation of arguments developed in previous DSB decisions (Zdouc, 2008, pp. 272-73). This type of court learning is therefore similar to the collective form of learning that occurs under legal precedent in domestic common-law systems and that led Posner (1973) to make his famous claim, as discussed above.³

³We should point out that, in comparison to domestic common-law systems, the application of legal precedent in the WTO is generally thought to be more fluid (see, for example, Jackson, 2006). It should also be noted that an entire literature in law and economics has arisen to challenge and assess Posner's claim (see, for example, Baker and Mezzetti, 2012, who emphasize the possibility of convergence to inefficient legal rules when judicial resources are scarce; Parameswaran, 2018, who shows that endogenizing the case-arrival process in a model à la Baker-Mezzetti introduces a further reason why legal rules may not converge to efficiency; Daughety and

And 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.

We refer to these possible forms of learning interchangeably as “judicial learning” or “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.

A key difference between a government-to-government court system like the WTO-DSB and a domestic legal setting is that the number of potential disputants in the WTO system (the member governments) is small and therefore the impact of current rulings on the future quality of the court is at least in part *internalized* by the potential disputants. This “large players” setting means in turn that the full implications of judicial learning for a court system such as the WTO-DSB cannot be understood in settings where the potential disputants are atomistic, as has been assumed in the law and economics literature that studies court learning in the domestic legal context. Moreover, the internalization of court learning may have implications not only for the decision of whether to go to court, but also for the choices that determine whether there is a dispute in the first place. The law and economics literature does not speak to this issue, because that literature typically treats the dispute margin as exogenous. In contrast, we explore the implications of judicial learning through a simple dynamic model that endogenizes the occurrence of both disputes and rulings in a large-players setting.⁴

The main ingredients of our model are as follows. In each period, a Home government makes a policy choice and a Foreign government decides whether or not to dispute this choice. If a dis-

Reinganum, 1999, who explore the possibility of informational cascades that might preclude convergence to efficient legal rules; and Talley, 1999, for a skeptical appraisal of the likelihood that such precedential cascades would interfere with the attainment of efficient legal rules in practice; and see also Hadfield, 1992, 2011, and Beim, 2017). For our purposes, to the extent that this collective form of court learning is relevant for the WTO, what matters is whether it leads *toward* efficiency, not that it necessarily attains efficiency, and on this point the law-and-economics literature is in broad agreement under a variety of conditions (see for example, Gennaioli and Shleifer, 2007, who establish that legal evolution under a common law system of precedent is on average beneficial even if judges are motivated by personal agendas).

⁴Pelc (2014) presents evidence that governments often seek DSB rulings not for their immediate commercial value but for their future precedential value. As we discussed above, the evolution of legal precedent can be interpreted as a “collective” form of court learning, thus when viewed in this perspective, Pelc’s evidence is consistent with our large-players story.

pute is initiated, governments bargain “in the shadow of the law,” subject to negotiation costs. In case of disagreement, the court intervenes and issues a ruling with the objective of maximizing the governments’ joint payoff. Finally, the quality of court rulings increases with experience (cumulative rulings), implying a higher joint disagreement payoff for the governments.

Our first main result is to demonstrate that the presence of judicial learning can explain litigation in equilibrium. In most existing models of “bargaining in the shadow of the law,” litigation is explained by informational frictions, either in the form of private information or of inconsistent prior beliefs about the court 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 the governments tomorrow. Importantly, such payoff gains arise even 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 policy that Home chooses to avoid a dispute (an “off-off-equilibrium” effect).

To demonstrate how the prospect of judicial learning can give rise to litigation, we begin with a static benchmark setting, and show that in such a setting disputes never go to court in equilibrium. In spite of the lack of rulings, however, the governments’ joint payoff increases with court quality, due to the off- and off-off-equilibrium effects described above. Notably, this joint payoff increase is assured even though we find that, if court efficiency is above a threshold level, further increases in court efficiency must hurt one government and benefit the other. We also find that, while they will never go to court, disputes *can* occur in equilibrium in the static setting, and the likelihood of disputes tends to be higher if the court is less efficient. With these results established in the benchmark static setting, we then turn to a dynamic setting and show that the presence of court learning can give rise to rulings in equilibrium. Furthermore, the probability of rulings is higher – other things equal – if governments are more patient.

Our second main result is to show that, when judicial learning is the driving force of equilibrium disputes and rulings, the likelihood of current disputes and rulings must decrease with court experience (cumulative rulings) if governments are sufficiently patient; and even if the governments’ patience is low, the likelihood of disputes and rulings decreases with court experience in an average sense. 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. Focusing on the case of rulings, the dynamic effect of an increase in court

experience makes a ruling less likely as long as the future joint returns from court learning are diminishing, but the static effect pushes in the opposite direction because an increase in DSB quality reduces the inefficiency of going to court today. As we demonstrate, the static effect can overwhelm the dynamic effect for specific levels of court experience if governments are not very patient, but on average the dynamic effect must dominate. And regarding whether the future joint returns from court learning are diminishing, our model highlights two forces that push in this direction. One is not surprising, and arises if learning itself occurs at a diminishing rate. But even if learning occurs at a constant rate, returns from court learning may still be diminishing, and the reason is linked to the off-off-equilibrium effect of court quality that we mentioned above: conditional on a dispute *not* being initiated, 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.

These results therefore suggest that a declining frequency of disputes and rulings over time can be a symptom of the beneficial learning that is associated with a rise in cumulative rulings, and hence good news for the WTO. But we also show that the *settlement rate* (likelihood of settlement conditional on a dispute) may go up or down with cumulative rulings even if governments are patient. Hence, when viewed through the lens of our model, one should look at the impact of cumulative rulings on the unconditional likelihood of rulings and disputes, not the settlement rate, when seeking evidence of court learning. This finding helps guide our empirical analysis.

Our basic model focuses on two countries, and therefore cannot make a distinction between disputant-specific learning (applying only to future disputes that involve the same disputants) and general-scope learning (applying to future disputes regardless of which countries are involved). But with many countries, learning could be disputant-specific, or purely general, or in-between these two extremes. We consider these possibilities in a many-country extension of our model, with the aim of making our framework more flexible in anticipation of our empirical analysis. 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 disputes and rulings is also informative about the *scope* of court learning.

Finally, we look for evidence of learning by ruling using WTO trade dispute data. Unlike the existing empirical work on learning by doing for firms, where direct measures of productivity 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. We focus on the key prediction of our model regarding the likelihood of rulings, because data on rulings are more reliable than data on disputes. Our model suggests that a stronger (negative) effect of cumulative rulings on the likelihood of current rulings 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 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, following a similar empirical approach to Kellogg (2011) who looks for evidence of various forms of learning by doing in the context of drilling activity in the oil and gas industry.

Our empirical findings suggest that judicial learning in the WTO is significant. We find evidence consistent with issue-area- and disputant-specific learning, while we find only weak evidence of general-scope learning. Moreover, we find that calendar time does not have a negative impact on the frequency of rulings once we control for our measures of court experience, suggesting that court learning can help explain the raw declining trends in disputes and rulings as evidenced by Plots 1 and 2 and provide a “good news” interpretation of these trends. We also discuss alternative explanations for the correlations we find in the data.

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, while here we focus 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). 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.⁵

By contrast, there is a vast law-and-economics literature on the implications of legal precedent for the evolution of common law, as we have noted. But its emphasis on whether a system of common-law precedent can deliver efficient law under conditions that typify domestic legal settings makes the focus of this literature very different. Among other things, the difference in emphasis accounts for the fact that this literature treats the potential disputants as atomistic, and therefore it is not directly relevant for the large-players setting of WTO disputes.

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 (2018). 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 contains several proofs not included in the body of the paper and 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. The Home government can choose an import barrier T , while the Foreign government has no export policy 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.

⁵Our model also shares some features with the theoretical literature on learning by doing for firms (which includes e.g. Spence, 1981, Cabral and Riordan, 1994, and Besanko et al., 2010), a common feature being that incurring some extra cost today produces efficiency benefits in the future. But the nature of the future benefits is very different across the two settings, and there is no analog in the firm learning-by-doing setting to the off- and off-off-equilibrium effects described above.

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 denote the joint government payoff by $\Omega(T, s) \equiv \omega(T, s) + \omega^*(T, s)$. 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.⁶

We assume that the Home government’s payoff $\omega(T, s)$ 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 the Foreign government’s payoff $\omega^*(T, s)$, below we will impose some structure on the government Pareto frontier.

In what follows we describe a simple setting where a choice of $T^N(s)$ by the Home government would be internationally inefficient, and where the Home and Foreign governments have signed a trade agreement to address this inefficiency and have agreed to abide by a dispute settlement procedure under which the Foreign government may lodge a complaint with the court against the Home government’s choice of T . Under this arrangement, governments may or may not get into a dispute ex post 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. We model these ex-post negotiations as a Nash bargain, with a disagreement point given by the expected court ruling and Home and Foreign bargaining powers given respectively by σ and $(1 - \sigma)$ with $\sigma \in [0, 1]$. We assume that if (and only if) governments engage in negotiations, they can use efficient transfers as a means of compensation.

Key to our analysis in this setting is the government ex post (i.e., conditional on s) Pareto frontier in payoff space. Consider first the case where 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. We will refer to this as the “no-transfer frontier.” Note from Figure 1 that the no-transfer frontier has

⁶There are three reasons we include uncertainty in the model. First, below we will introduce a trade agreement with a court to help settle disputes, and as will become clear, for the court to play a non-trivial role the trade agreement must not be a fully contingent contract, and so there must be some uncertainty in the future state of the world. Second, as will also become clear below, uncertainty in s allows for variation in equilibrium outcomes and non-degenerate probabilities of disputes and court rulings. Third, in the two-period model developed in the next section we will assume a veil of ignorance, and this too requires uncertainty about the future. We also note that our results do not depend on the distribution of s , so we do not need to introduce explicit notation for such a distribution.

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).⁷ Consider next the case in which ex post 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 refer to this as the “negotiation frontier.”⁸

Also key to our analysis is the disagreement point for government ex post negotiations, that is, the expected payoffs from going to court. We do not formalize explicitly the court’s information and decision-making process, nor the specific form of the contract over T 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 $\omega^R(s; \lambda)$ for the Home government and $\omega^{*R}(s; \lambda)$ for the Foreign government. The parameter λ captures the imperfections/inefficiencies of the court system (or equivalently, the inverse of court “quality”); these may include the imperfections of the court’s information and/or the costs of litigation. 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 assume that ω^R and ω^{*R} are differentiable in λ , and take λ to be between zero (perfect court) and some finite bound. What matters for our results is the governments’ expected joint payoff from going to court and how it is impacted by court quality; for this reason, our approach of working directly with the reduced-form expected payoffs $\omega^R(s; \lambda)$ and $\omega^{*R}(s; \lambda)$ makes the logic of our theory more transparent. As we make clear below, the core restriction of our static model is simply that the governments’ expected joint payoff from going to court is higher if the court is more efficient.

With our reduced form approach we can keep the model compatible with a variety of different forms of ex-ante contracts and court mandates, ranging from vague agreements that endow the court with “interpretive” authority, to agreements with gaps that endow the court with “gap-filling” authority, to contracts that specify contingent policy rules where the relevant state of

⁷To be more precise, the no-transfer frontier does not include the upward sloping part of the curve drawn in Figure 1, because these points are Pareto-dominated by the N point. For this reason this part of the curve is dashed.

⁸Our 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.

the world is imperfectly observed by the court.⁹ The most direct interpretation of a “ruling” in this setting is that the court prescribes a specific value of T , but other forms of ruling are also compatible with the model; for example, the court might specify bounds on the policy T (e.g. a tariff binding) and let the importing country choose a policy within these bounds. Notice that we have specified the governments’ expected payoffs from triggering a court ruling as independent of Home’s chosen import barrier T . This means that the court ruling amounts to a policy (or policy bounds) for the Home government that is independent of the initially chosen Home import barrier that triggered the dispute. We have in mind that regardless of how egregious the initial Home choice of T might be perceived by the court to be, the role of the court ruling is simply to ensure that the Home government adopts the appropriate policy “going forward” as determined by the court, a role that is consistent with the nature of DSB rulings in the WTO.¹⁰

Finally, we assume that negotiations are subject to transaction costs. Note that, if there were no transaction costs, the disagreement point would be irrelevant to the governments’ joint payoff, and hence any institutional or contractual arrangement, as well as court learning, would be irrelevant. For simplicity 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.¹¹ Formally, if $\omega^B(s; \lambda)$ and $\omega^{*B}(s; \lambda)$ are the governments’

⁹For example, the role of the court may be to “complete” an incomplete contract along the lines of Maggi and Staiger (2011). This is an environment where governments cannot write a complete contingent contract, and the court is endowed with the authority to interpret or fill the gaps of the contract ex-post. Another contractual/institutional environment is the one described in Maggi and Staiger (2018), 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, and both environments are captured by our reduced-form approach. What our approach does not allow are contracts that are unambiguous and specify perfectly verifiable rules, and therefore leave the court with a pure enforcement role and no possibility of errors. Such contracts, however, are clearly unrealistic in the context of the WTO.

¹⁰In particular, at no time in the history of the GATT/WTO has there been any formal remedy for harm done prior to the adjudication of a violation (see, for example, Ossa, Staiger and Sykes, 2020). Notice also that if the court is for example interpreting a vague provision of the ex-ante contract, then the appropriate policy as determined by the court is its best guess at what the contract demands in the current state s (e.g., Is the safeguard action taken by the Home government consistent with what the DSB interprets is stipulated in the WTO contract for state s , or not?). The important feature for our purposes is that the court ruling implies a policy to be adopted by the Home government (or a bound on that policy) that is independent of the Home government’s initial policy.

¹¹We have also considered the case of fixed bargaining costs, which seems a natural alternative to the assumption of iceberg negotiation costs. Our main qualitative results can be shown to hold also with fixed bargaining costs. See also note 18.

bargaining payoffs absent negotiation costs, the net payoffs from the bargain are

$$\omega^{B_{net}}(s; \lambda) = \omega^R(s; \lambda) + \kappa[\omega^B(s; \lambda) - \omega^R(s; \lambda)], \quad (2.1)$$

$$\omega^{*B_{net}}(s; \lambda) = \omega^{*R}(s; \lambda) + \kappa[\omega^{*B}(s; \lambda) - \omega^{*R}(s; \lambda)]. \quad (2.2)$$

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).¹²

In Appendix A we present a more structured model that is a special case of the environment just described. That model focuses on a standard competitive partial-equilibrium setting where the Home government maximizes a politically-adjusted welfare function that attaches an extra weight to import-competing producers. There, the Home political-economy parameter is ex-ante uncertain and the policy T is a tariff. If invoked, the court estimates the first-best tariff as a function of the political-economy parameter, but does so with noise. This noise represents the key imperfection in the court system, corresponding to the λ parameter in the reduced-form model we present here.

We will conduct much of the analysis at a graphical level, since this will aid intuition. We start by fixing a level of court quality $\lambda > 0$ and a state s and identifying the disagreement point and the equilibrium bargaining payoffs for the government negotiation (stage 3). In Figure 1, point R indicates the expected payoffs from triggering a court ruling, (ω^R, ω^{*R}) , point B indicates the bargaining payoffs absent negotiation costs, (ω^B, ω^{*B}) , and point B_{net} indicates the net bargaining payoffs $(\omega^{B_{net}}, \omega^{*B_{net}})$. We omit the arguments s and λ from the labels in this figure, as this should not cause confusion. We depict R as residing below the no-transfer frontier in Figure 1; this implicitly assumes that the court ruling does not dictate that transfers take place.¹³ Point B is the projection of point R onto the negotiation frontier, with this

¹²Notice 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 outside the institutional framework, through informal negotiations. Our model can be interpreted as applying to both formal and informal negotiations.

¹³This conforms with WTO rulings in the typical (“violation” claim) dispute, where compensation plays no

projection reflecting the bargaining power parameter σ . And finally, the point B_{net} is a linear combination of point R and point B , with the exact location determined by the negotiation cost parameter 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 and what action will Foreign take at stage 2? The answer follows from two observations. First, if the B_{net} point is above the no-transfer frontier as in the higher of the two B_{net} points pictured in Figure 1 – which will be the case if κ is sufficiently close to 1 (i.e. negotiation costs are sufficiently small) – 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 the point B_{net} . In this case there is a dispute but no ruling. And second, if the B_{net} point is below the no-transfer frontier as in the lower of the two B_{net} points pictured in Figure 1 – which will be the case if κ is sufficiently close to 0 – 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 not even a dispute, let alone a ruling.¹⁴ It is intuitive, then, and straightforward to show, that for given s and λ a dispute will occur in the static model if and only if κ is above some threshold level $\hat{\kappa}(s, \lambda)$; and it is clear that there is never a ruling in equilibrium. Summarizing:

Proposition 1. *In the static setting, there is a dispute if and only if $\kappa > \hat{\kappa}(s, \lambda)$, and there is never a ruling in equilibrium.*

We next consider the impact of court quality λ on the governments’ equilibrium payoffs. This will be a key ingredient for the dynamic analysis of the next section.

role in the ruling of the DSB. A role for compensation in DSB rulings arises, if at all, under two circumstances: first, if the defending government refuses to abide by an initial DSB ruling against it, then the DSB may authorize the claimant to engage in reciprocal tariff actions of its own in order to achieve an appropriate level of compensation for the continuing violation; and second, in the case of “non-violation” claims, an option to compensate is part of the initial DSB ruling. Our model could be extended to accommodate these additional circumstances by, for example, allowing the Home government to choose between a low level of protection and a high level of protection coupled with a transfer to the Foreign government, with the implication that in Figure 1 the point R could now lie above the no-transfer frontier. This would in turn alter some of our results regarding the likelihood of disputes, but would leave all of our results regarding the likelihood of rulings unaltered.

¹⁴To 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.

Our core assumption is that decreasing λ leads to a higher joint disagreement payoff Ω^R for each state s . This assumption can be justified as long as governments have access to side transfers at the stage when they write their ex-ante contract, because the goal of their contract should then be to maximize their joint surplus. If the ex-ante contract is vague or has gaps and the role of the court is to interpret the contract or fill the gaps, it is then natural to suppose that the court attempts to complete the contract as the governments would have done ex ante and therefore attempts to maximize the governments' joint payoff with its rulings, and thus a higher-quality court is better at achieving this objective.¹⁵ And if the ex-ante contract is a contingent contract, then the contract will specify the first-best policy $T^{fb}(s)$ and if the court is invoked it will simply apply the contract and prescribe its best estimate of the first-best policy, and therefore again a higher-quality court is better at achieving joint-surplus-maximizing rulings.

To keep our analysis focused on the main points, we also make two technical assumptions whose specific roles will become clear below: first, a perfect court implements the first best outcome, that is $\Omega^R(s; 0) = \Omega^{fb}(s)$; and second, Foreign's disagreement payoff is at no point insensitive to court quality, that is $\frac{d\omega^{*R}}{d\lambda} \neq 0$ for all λ . We summarize these restrictions in:

Assumption 1: (i) $\Omega^R(s; \lambda)$ is strictly decreasing in λ ; (ii) $\Omega^R(s; 0) = \Omega^{fb}(s)$; (iii) $\frac{d\omega^{*R}}{d\lambda} \neq 0$ for all λ .

It is worth pointing out that Assumption 1 allows for the possibility that increasing court efficiency (decreasing λ) might decrease the disagreement payoff of one of the governments, while increasing the joint disagreement payoff. While it might seem natural that both governments should gain from greater court efficiency, it is straightforward to show that in some standard settings this is not the case, and so the generality provided by Assumption 1 in this regard is relevant.¹⁶

Recalling that $\omega^R(s; \lambda)$ and $\omega^{*R}(s; \lambda)$ are represented for a fixed λ and s as the point R in

¹⁵The assumption that the court seeks to complete the contract as the governments would have done ex ante is broadly in line with the rules set out by the Vienna Convention (and adhered to by the WTO). In WTO (2005), the Appellate Body states: "We recall that Article 31(1) of the Vienna Convention stipulates that: 'A treaty shall be interpreted in good faith in accordance with the ordinary meaning to be given to the terms of the treaty in their context and in the light of its object and purpose.'... Importantly, the ordinary meaning of a treaty term must be seen *in the light of the intention of the parties...*" (emphasis added). The same is true in domestic legal settings, where, for example, Posner (2005, p.8) writes: "Gap filling and disambiguating are both 'interpretive' in the sense that they are efforts to determine how the parties would have resolved the issue that has arisen had they foreseen it when they negotiated their contract."

¹⁶For example, in the special model we present in Appendix A, 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.

Figure 1, in Figure 2 we depict the locus traced by the disagreement point R as λ varies, which we label \mathcal{R}^λ (as before, we omit the argument s from the label to keep the notation simple). Note that Assumption 1 allows the \mathcal{R}^λ locus to be upward or downward sloping; Figure 2 depicts the case in which the \mathcal{R}^λ locus is downward sloping, but the analysis is similar if this locus is upward sloping. Also note that Assumption 1(ii) implies that the \mathcal{R}^λ locus emanates from point FB , while Assumption 1(iii) implies that the \mathcal{R}^λ locus cannot be backward bending. The significance of this last property will become clear below.

We now turn to the impact of court quality (λ) on the net bargaining payoffs in the dispute subgame $\omega^{B_{net}}(s; \lambda)$ and $\omega^{*B_{net}}(s; \lambda)$ (which for fixed λ and s we represented by the B_{net} point in Figure 1). In Figure 2 we label $\mathcal{B}_{net}^\lambda$ the locus traced by the B_{net} point as λ varies. Since the B_{net} point is a linear combination of the R point and the B point in Figure 1, it is immediate that the $\mathcal{B}_{net}^\lambda$ locus in Figure 2 inherits two properties from the \mathcal{R}^λ locus: (a) it coincides with the FB point for $\lambda = 0$, and (b) the joint payoff $\Omega^{B_{net}}(s; \lambda)$ is strictly decreasing in λ . Moreover, this second property implies that the $\mathcal{B}_{net}^\lambda$ locus hits the FB point with slope different than -1 , and thus a third property of the $\mathcal{B}_{net}^\lambda$ locus is that (c) when λ is small enough, the $\mathcal{B}_{net}^\lambda$ locus lies below the no-transfer frontier (and it may or may not cross the frontier as λ increases).

Finally, we ask whether the $\mathcal{B}_{net}^\lambda$ locus also inherits the property of the \mathcal{R}^λ locus that it cannot be backward bending. As we show in Appendix B, this is always the case if the \mathcal{R}^λ locus is downward sloping (as is the case for example in the special model presented in Appendix A); and if the \mathcal{R}^λ locus is upward sloping, this is the case provided that either the negotiation cost or Home's bargaining power is high enough. In what follows we assume that the $\mathcal{B}_{net}^\lambda$ locus is not backward bending, and we comment below on the qualifications to our results that would arise if this condition were not satisfied (see note 17).

We can now back up to stage 1 and examine how λ affects the equilibrium joint payoff in the full game, which we denote by $\Omega^e(s; \lambda)$. We will argue that $\Omega^e(s; \lambda)$ is decreasing in λ . To see this recall that, if the B_{net} point is below the no-transfer frontier (so 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 $\mathcal{B}_{net}^\lambda$ locus, with the resulting locus depicted by the red curve in Figure 2. To confirm that $\Omega^e(s; \lambda)$ is decreasing in λ , 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 $\Omega^e(s; \lambda)$ decreases.¹⁷ Next focus on a range of λ such that the equilibrium point is on the $\mathcal{B}_{net}^\lambda$ locus. Recalling that $\Omega^{B_{net}}$ decreases with λ , it follows again that increasing λ will decrease $\Omega^e(s; \lambda)$.¹⁸ We can thus state:

Proposition 2. *In the static setting, the equilibrium joint payoff is increasing in court quality, even though there is no ruling in equilibrium. This results from two effects: (i) if there is a dispute and governments settle, higher court quality improves the disagreement point, and this leads to a better bargaining outcome (an “off-equilibrium” effect); (ii) if there is no dispute, higher court quality improves the would-be bargaining outcome, thus inducing Home to choose a more efficient policy (an “off-off-equilibrium” effect).*

A further implication of the model is also noteworthy, and can be understood in terms of Figure 2. 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 \mathcal{R}^λ locus is increasing with slope equal to one. We record this implication in:

Remark 1. *If court quality is above some threshold level, further improvements in court quality must hurt one of the governments (and benefit the other).*

Finally, the model also yields a prediction on how court quality affects the likelihood of a dispute, under one additional regularity condition. Suppose the $\mathcal{B}_{net}^\lambda$ locus crosses the no-transfer frontier at most once.¹⁹ Then, recalling that the $\mathcal{B}_{net}^\lambda$ 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

¹⁷This is where our assumption that the $\mathcal{B}_{net}^\lambda$ locus is not backward bending comes in: if this locus were backward bending, this could induce “oscillating” movements of the equilibrium point – first moving away from the FB point along the frontier, and then moving back toward the FB along the frontier – over some range of λ as λ increases. It is easy to show, however, that even if $\mathcal{B}_{net}^\lambda$ locus were backward bending, all of our results would still hold provided λ is below some critical threshold.

¹⁸We note that, if the negotiation cost were modeled as a fixed cost rather than an iceberg cost, the only change in results would be that the equilibrium joint payoff is *weakly* increasing in court quality, and in particular, it is constant in the dispute region and strictly decreasing in no-dispute region. All our remaining results would not be affected.

¹⁹Note that this will be the case if the $\mathcal{B}_{net}^\lambda$ locus is convex, linear or not too concave. For example, in the Appendix A model the $\mathcal{B}_{net}^\lambda$ locus is linear.

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 distribution of s , the likelihood of a dispute is increasing in λ . We summarize with:

Proposition 3. *In the static setting, the likelihood of a dispute decreases in court quality, provided the $\mathcal{B}_{net}^\lambda$ locus crosses the no-transfer frontier at most once.*

To summarize the main points of this section, in the static setting disputes never go to court in equilibrium (Proposition 1), but nevertheless the governments' joint payoff increases with court quality, due to the off- and off-off- equilibrium effects identified in Proposition 2. Disputes, on the other hand, can occur in equilibrium and tend to be more likely when the court is less efficient (Proposition 3). These static results will play an important role in the analysis of the dynamic setting, which we develop in the next section.

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 in the form of the DSB. 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 introduce a notion of learning by ruling into the static model with the assumption that the court quality parameter λ is a weakly decreasing function of the stock of prior court rulings (court experience).

We have already described the kinds of court learning that can be captured by our model. But what kinds of learning fall outside our model? One type of learning that may be relevant in the WTO is *governments' learning about the court*. For example, it is possible that as governments litigate repeatedly, they learn how the court operates and adjudicates cases, and therefore they learn to better predict the outcome of a ruling. Another type of learning that may be relevant is *governments' learning about each other*. In principle this type of learning could reduce asymmetric information between bargainers. Our modeling approach abstracts

from these learning possibilities, but we will come back to them in the empirical section, when we discuss alternative mechanisms that might explain our empirical findings.

We consider 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 is *iid*, so learning by ruling will be the only source of dynamics. The governments' common discount factor is $\delta \in (0, \infty)$.²⁰

In order to examine how past rulings affect current outcomes, we suppose that at $t = 1$ there is an initial stock of rulings (court experience) x inherited from a "past" period $t = 0$. We model learning by ruling in a fashion similar to the typical models of learning by doing for firms, where a firm's efficiency increases with the firms' cumulative output: we assume that the efficiency of the court increases with cumulative rulings. Formally, as we noted above, we model learning by ruling as a weakly decreasing function $\lambda(x)$.

We let $\lim_{x \rightarrow \infty} \lambda(x) \equiv \bar{\lambda} \geq 0$ (this limit clearly exists given that $\lambda(x)$ is decreasing and bounded below by zero). To simplify our formal arguments, we assume that $\bar{\lambda}$ is reached for some finite level of x , that is, there exists a large enough \bar{x} such that $\lambda(\bar{x}) = \bar{\lambda}$ and learning is exhausted.²¹

Finally, 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 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 and generating a ruling delivers 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 at $t = 2$. But the situation is different at $t = 1$, because going to court and generating a ruling delivers payoff gains at $t = 2$ due to the learning effect. We next analyze under what conditions a ruling occurs in equilibrium at $t = 1$; we will not use a time index, as this should not cause confusion.

²⁰The reason we allow δ to be higher than one is that we think of period 2 as being longer than period 1, capturing the fact that the future may matter more than the present. Of course, a richer specification of the model would allow for many future periods (in which case it would be compelling to have $\delta < 1$), but this would make the model much less tractable.

²¹Note that we allow $\bar{\lambda}$ to be zero or strictly positive; that is, once learning is exhausted the court may or may not be able to implement the first-best outcome.

Recall that in the static setting, under Assumption 1 the equilibrium joint payoff is decreasing in λ . Thus, given the veil of ignorance, going to court at $t = 1$ implies a common future payoff gain for each government, which we denote by Δ . Since a ruling at $t = 1$ changes the value of λ at $t = 2$ from $\lambda(x)$ to $\lambda(x + 1)$, we can write Δ as half of the joint future payoff gain:

$$\Delta = \frac{1}{2}E_s[\Omega^e(s; \lambda(x + 1)) - \Omega^e(s; \lambda(x))], \quad (3.1)$$

It is worth emphasizing that, much as in the static model, here governments benefit from higher 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 in a richer model with more than one period ahead of $t = 1$, the payoff gain Δ would include as well the direct effect of increasing court efficiency in case a ruling occurs in equilibrium.

At $t = 1$, the disagreement payoffs for Home and Foreign are now respectively given by $\omega^R + \delta\Delta$ and $\omega^{*R} + \delta\Delta$ (we omit the arguments s and x for simplicity), that is, the expected payoff from a court ruling in the current period plus the discounted future payoff gain. Graphically, in Figure 3 we label as $R + \delta\Delta$ the corresponding disagreement payoff point. This point lies somewhere on the 45° line emanating from point R (which represents the point (ω^R, ω^{*R}) in Figure 3), and in general may be below or above the negotiation frontier. If point $R + \delta\Delta$ is above the negotiation frontier (which is the case illustrated in Figure 3a), then a dispute will end in ruling; and going backwards, in this case Home chooses a policy T that triggers a complaint by Foreign.²²

Note from Figure 3 that a ruling occurs if and only if the distance between the R point and the negotiation frontier along a 45° line is less than $\delta\Delta$, or equivalently, if

$$\Omega^{fb}(s) - \Omega^R(s; \lambda(x)) < 2\delta\Delta. \quad (3.2)$$

Given s and x , (3.2) defines the ruling condition for our dynamic model. This condition is satisfied if δ is above some threshold, therefore for any distribution of s the probability of a

²²It is now easy to understand the role played in our model by the veil of ignorance. If we removed the veil of ignorance, the main change would be that the probability of a ruling could be lower, other things equal. The reason lies in Remark 1: while an increase in court quality increases the equilibrium joint payoff, when λ is sufficiently small it must *hurt* one of the governments. When this is the case, it would take a higher δ for the disagreement point at $t = 1$ to be above the negotiation frontier.

ruling (as viewed from the beginning of period $t = 1$, before s is realized) is increasing in δ . We may therefore state:

Proposition 4. *Given x , the probability of a ruling at $t = 1$ is increasing in δ (and is equal to zero if $\delta = 0$).*

This result makes clear that in our dynamic model rulings can happen in equilibrium, and the likelihood of rulings is higher when governments care more about the future.

Next we examine how court experience x affects the likelihood of a ruling at $t = 1$. We assume that learning occurs at a diminishing rate. Given that the units of measure of court quality can be defined in an arbitrary (ordinal) way, for a notion of diminishing rate of learning to be invariant to the units of measure it needs to be stated in payoff terms. In what follows, we assume that learning occurs at a diminishing rate when evaluated in terms of its effect on disagreement payoffs:²³

Assumption 2: Increases in court experience have weakly diminishing marginal impacts on disagreement payoffs: $|\frac{d\omega^R(s,\lambda(x))}{dx}|$ and $|\frac{d\omega^{*R}(s,\lambda(x))}{dx}|$ are weakly decreasing in x (for any s).

Let us now examine how x affects the condition for equilibrium rulings (3.2). First note that the left-hand side of this condition, $\Omega^{fb}(s) - \Omega^R(s; \lambda(x))$, decreases with x – because the joint disagreement payoff Ω^R is increasing in court quality and hence in x – and this makes a ruling more likely. This is the *static* effect of 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 court experience, that is the effect of x on Δ and hence on the right-hand side of (3.2). From (3.1) it is clear that Δ is decreasing in x if $E_s\Omega^e$ is concave in x . Recall from the static analysis and Figure 2 that, as x increases, and hence λ decreases, the equilibrium payoff point moves along the outer envelope of the B_{net}^λ locus (dispute range) and the no-transfer frontier (no-dispute range). It is easy to show that, given Assumption 2, the equilibrium joint payoff Ω^e is strictly concave in x over the no-dispute range for any

²³Recall 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 2 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 .

s .²⁴ 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 no-transfer frontier.

This last point bears emphasis. The concavity of the no-transfer Pareto frontier is a key force behind the diminishing impact of court experience x on the future joint payoff, and this in turn is related to the off-off-equilibrium impacts of court quality identified in Proposition 2: 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 frontier.²⁵ Indeed, even if Assumption 2 is not satisfied, Ω^e may be concave in x due to the concavity of the frontier.

Next note that, in light of the strict concavity of Ω^e in x over the no-dispute range for any s , the expected joint payoff $E_s \Omega^e$ is guaranteed to be concave in x if the probability of an equilibrium dispute at $t = 2$ is sufficiently small.²⁶ Given that in reality the average frequency with which trade disputes arise between two given governments is quite low (and that this is true even accounting for informal disputes that occur outside the formal WTO procedures), it seems reasonable to restrict attention to parameter constellations that imply a small probability of equilibrium disputes. We therefore impose:

Assumption 3: The probability of an equilibrium dispute at $t = 2$ is sufficiently small.

Under Assumptions 2 and 3, the expected future joint payoff $E_s \Omega^e(s; \lambda(x))$ is concave in x , and hence an increase in x reduces Δ (dynamic effect): this dynamic effect makes the ruling

²⁴To see this, first note that Assumption 2 ensures that x has diminishing marginal impact on ω^{*Bnet} , since by (2.2) this is linear in ω^{*R} and Ω^{fb} , and the latter does not depend on x . Next recall from Figure 1 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, it then follows that the joint payoff in the no-dispute range 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} .

²⁵As described in note 24, letting $\omega = F(\omega^*)$ denote the no-transfer frontier, the joint payoff in the no-dispute range is given by $\omega^{*Bnet} + F(\omega^{*Bnet})$, and this is more likely to be concave in x if the no-transfer frontier is more concave.

²⁶In our model the probability of an equilibrium dispute at $t = 2$ will be small if, for example, both governments benefit from a more efficient court and the negotiation cost is high; in that case, and with reference to Figure 1, 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 $t = 2$ – that is, the probability that the B_{net} point is above the no-transfer frontier – will be small (and may be zero).

condition in (3.2) less likely to be satisfied, and therefore makes a ruling less likely.²⁷

Recalling now that an increase in x also reduces $\Omega^{fb}(s) - \Omega^R(s; \lambda(x))$, the current inefficiency from going to court (static effect), and through this static effect makes a ruling more likely, it is apparent that the static effect pushes one way and the dynamic effect pushes the other. But clearly, if δ is sufficiently large the dynamic effect dominates the static effect. Since this is true for any s , it follows that if δ is sufficiently large then for any distribution of s the probability of a ruling decreases with x .

Next notice that, regardless of δ , the likelihood of a ruling must reach zero as x reaches \bar{x} , and therefore the marginal impact of x on the likelihood of a ruling on average must be negative. More formally, let $\pi_R(x)$ denote the probability of a ruling at $t = 1$ as a function of x , and suppose $\pi_R(0) > 0$. Then clearly $\pi_R(\bar{x}) = 0 < \pi_R(0)$, and hence we can write $\frac{1}{\bar{x}} \int_0^{\bar{x}} \pi'_R(x) dx < 0$. The next proposition summarizes:

Proposition 5. (i) *If δ is sufficiently high, the probability of a ruling at $t = 1$ is decreasing in x for all x .* (ii) *Regardless of δ , increases in x reduce the probability of a ruling at $t = 1$ on average.*

As Proposition 5 indicates, in our model there is a clear tendency for the probability of a ruling to decrease with court experience: if δ is sufficiently high, the probability of a ruling is decreasing in x for any x ; and regardless of δ , this inverse relationship must hold on average provided only that the potential for learning is eventually exhausted. This latter result relies only on the basic feature of the model that equilibrium rulings can occur only if there is potential for further judicial learning.²⁸

Next we consider the impact of court experience x on the probability of a *dispute* at $t = 1$. Let us first ask: what is the outcome at $t = 1$ if the disagreement point $R + \delta\Delta$ is below the negotiation frontier, as in Figure 3b, so that there is no ruling? In this case, the net

²⁷While Assumptions 2 and 3 suffice for our result, we note that our model abstracts from a factor that is arguably empirically relevant and would contribute toward the concavity of $E_s\Omega^e$ in x . As observed in note 8, if we introduced the assumption of a convex cost of transfers in the context of dispute settlement negotiations, then the negotiation frontier would be concave, and this would be a force toward concavity of $E_s\Omega^e$ in x .

²⁸As a matter of empirical interpretation, this latter result suggests that, in a world with many disputants, on average the probability of rulings should tend to decrease with court experience, in the following sense. Consider a world with many country dyads or issue areas, where court learning may be specific to the country dyad and/or the issue area, and suppose the level of court experience x is heterogeneous across these country dyads and issue areas. Then according to this result, the impact of x on the probability of rulings on average should be negative, regardless of the discount factors of the disputants.

bargaining payoffs at $t = 1$, which we denote $(\omega^{B_{net}^\Delta}, \omega^{*B_{net}^\Delta})$, can be written as follows (omitting the arguments s and x for simplicity):

$$\begin{aligned}\omega^{B_{net}^\Delta} &= \kappa\omega^B + (1 - \kappa)(\omega^R + \delta\Delta) \\ \omega^{*B_{net}^\Delta} &= \kappa\omega^{*B} + (1 - \kappa)(\omega^{*R} + \delta\Delta).\end{aligned}\tag{3.3}$$

Graphically, in Figure 3b we label B_{net}^Δ the point corresponding to payoffs $(\omega^{B_{net}^\Delta}, \omega^{*B_{net}^\Delta})$. Because of the negotiation cost (κ), point B_{net}^Δ lies somewhere between the disagreement point $R + \delta\Delta$ and the negotiation frontier. It is then easy to see that the outcome is a dispute with settlement if B_{net}^Δ is above the no-transfer frontier (as depicted by the higher B_{net}^Δ point displayed in Figure 3b), and the outcome is no dispute if B_{net}^Δ is below the no-transfer frontier (as depicted by the lower B_{net}^Δ point displayed in Figure 3b).

Armed with the above observation, we can establish that a similar result as Proposition 5(i) holds also for disputes. As was the case for rulings, increasing x has two effects on the likelihood of disputes at $t = 1$, a static effect (through ω^R , ω^{*R} , ω^B and ω^{*B}) and a dynamic effect (a reduction in Δ). If δ is sufficiently high, the dynamic effect dominates the static one, so we can focus on the former. How does a reduction in Δ affect the likelihood of a dispute at $t = 1$? It is clear from (3.3) that, holding all else equal, decreasing Δ reduces $\omega^{B_{net}^\Delta}$ and $\omega^{*B_{net}^\Delta}$ by an equal amount, or in graphical terms and using Figure 3b, the B_{net}^Δ point shifts inwards along a 45 degree line. This implies that, for any given s , the B_{net}^Δ point can cross the no-transfer frontier only from above, not from below; and as a result the equilibrium outcome can only switch from dispute to no-dispute, but not vice-versa. We can then conclude that, for any distribution of s , increasing x reduces the likelihood of a dispute at $t = 1$.

Finally, a similar result as Proposition 5(ii) holds also for disputes, under the regularity condition that the $\mathcal{B}_{net}^\lambda$ locus crosses the no-transfer frontier at most once. To show this we will argue, as we did above for the case of rulings, that the likelihood of a dispute for $x = 0$ is higher than for $x = \bar{x}$. For this purpose it is useful to return to Figure 2. First note that the B_{net}^Δ point (not pictured in Figure 2) must be (weakly) above the $\mathcal{B}_{net}^\lambda$ locus. Next note that, if we increase x from zero to \bar{x} , we are reducing Δ to zero (dynamic effect) and reducing λ to its minimum level $\bar{\lambda} \geq 0$ (static effect). How do these two changes affect the position of the B_{net}^Δ point? Reducing Δ to zero shifts the B_{net}^Δ point inward along a 45⁰ line until it hits the $\mathcal{B}_{net}^\lambda$ locus; and reducing λ causes the B_{net}^Δ point to travel along the $\mathcal{B}_{net}^\lambda$ locus toward the FB point. It is then easy to see that, if the $\mathcal{B}_{net}^\lambda$ locus crosses the no-transfer frontier at most once, the

B_{net}^Δ point can only cross the no-transfer frontier from above, not from below. And since this is true for any given s , we can conclude that increasing x from zero to \bar{x} reduces the likelihood of a dispute at $t = 1$. Hence we have:

Proposition 6. (i) If δ is sufficiently high, the probability of a dispute at $t = 1$ is decreasing in x for all x . (ii) Regardless of δ , increases in x reduce the probability of a dispute at $t = 1$ on average, provided the $\mathcal{B}_{net}^\lambda$ locus crosses the no-transfer frontier at most once.

Together Propositions 5 and 6 suggest an important possibility: a declining frequency of disputes and rulings over time can be a symptom of the beneficial court learning that is associated with a rise in cumulative rulings. As we discussed in the Introduction, our model therefore provides a “good news” interpretation of recent trends in WTO disputes and rulings.

It also bears emphasis that, while a declining trend in the frequency of rulings is not a sign of declining court efficiency according to our model, a lower *level* of the frequency of rulings *is* associated with lower court efficiency. To see this point in the simplest way, suppose we shift up 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 with reference to Figure 3, for any given s the disagreement point $R + \delta\Delta$ can only cross the negotiation frontier from above, not from below, so the probability of a ruling decreases.²⁹

We close this section by emphasizing one measure that according to our model is *not* a reliable metric for gauging the performance of the court, namely, the *settlement rate* (likelihood of settlement conditional on a dispute), which may go up or down with cumulative rulings even if governments are patient. We can state:

Remark 2. At $t = 1$ the likelihood of settlement conditional on a dispute may go up or down with x , even if δ is high.

The reason for this ambiguity is that the effect of an increase in x on the probability of a ruling may be stronger or weaker than its effect on the probability of a dispute, depending on the exact distribution of s , and thus the ratio between the two probabilities can increase or decrease

²⁹Perhaps interestingly, the same result does not hold for the frequency of disputes: if we shift up the $\lambda(x)$ schedule in a way that keeps Δ constant, the frequency of disputes can go up or down. This can be easily gleaned from inspection of Figure 3.

with x .³⁰

According to Remark 2, 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 considered a simple setting with two countries and a single issue area, where there cannot be a meaningful distinction between general-scope court learning and more narrow forms of court learning. But in a world with many countries, court learning could vary in scope. For example, learning could be “general,” in the sense that issuing a ruling today increases the court’s future efficiency regardless of which countries and issues are involved in future rulings; or it could be “dyad specific,” in the sense that a ruling issued for a dispute between countries i and j increases the court’s future efficiency only for disputes that involve the same pair of countries. In this section we consider a many-countries extension of the model, where we can distinguish between general-scope learning and dyad-specific learning. The main purpose of this extension is to enrich the model in view of our empirical exploration.³¹

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

³⁰More formally, consider a realization $s = s'$ such that $R + \delta\Delta$ is just above the negotiation frontier. As x increases, Δ decreases and hence for $s = s'$ the outcome switches from ruling to settlement; and if the probability mass of s is concentrated around s' , then $\Pr(\text{settlement})/\Pr(\text{dispute})$ goes up. Alternatively, 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.

³¹The scope of court learning could also be limited by the directionality of disputes. For example, learning might be “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. This could be the case if for example the court learns about features of the industry in country j that competes with imports from country i . Or as another example, learning could 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 our working paper (Maggi and Staiger, 2019), we allow for this richer set of possibilities where the role of a disputant (defendant or complainant) matters for the scope of learning. We come back to this point in the empirical section below, where we also discuss the possibility that learning could be specific to the disputed issue area.

other non-numeraire goods that countries trade, and with an outside good that enters utility quasi-linearly. The payoff to a given government i is then the sum of $N - 1$ (separable) payoff terms, one for each of country i 's trading partners, and each of these payoff terms depends on a state variable s_{ij} that is ex-ante uncertain.

We assume that, in each period, a single pair of countries is selected randomly from the population of countries to play the game detailed in the previous section. We continue to assume a veil of ignorance: conditional on a pair of countries having a dispute at $t = 2$, there is an equal probability that each country will be the defendant. Finally, we assume that analogues of Assumptions 1-3 apply in this extended setting.

We denote by λ_{ij} the court's imperfection when ruling on a dispute between countries i and j . We assume that $\lambda_{ij} = \lambda(X_{ij})$, where $\lambda(\cdot)$ is a weakly decreasing function, and X_{ij} is a composite experience variable that takes the form

$$X_{ij} = \beta_1 x_{ij} + \beta_2 x_{other}, \quad (4.1)$$

where x_{ij} is the number of past rulings for disputes between countries i and j , and x_{other} is the number of past rulings for all other disputes. Both β_1 and β_2 are assumed weakly positive. Moreover, it is natural to assume $\beta_1 \geq \beta_2$, because it is plausible that direct experience is at least as relevant as indirect experience.³²

Our formulation of court learning includes a few 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 purely general learning corresponds to the case where $\beta_1 = \beta_2 > 0$. At the other extreme, court learning could be purely dyad-specific, so that prior experience is applicable only to future disputes in which the same pair of governments is involved: this corresponds to the case where $\beta_1 > 0$ and $\beta_2 = 0$. And in between these two extremes are intermediate cases where both general-scope learning and dyad-specific learning are present ($\beta_1 > \beta_2 > 0$); here, β_2 is interpreted as the general component of learning, and $\beta_1 - \beta_2$ as the dyad-specific component.³³

³²Notice that with the learning spillovers to third countries implied by the presence of general learning, there is potential for under-utilization of the court system when this form of learning is present. We leave an investigation of the normative and positive implications of this observation to future work.

³³Note an implicit restriction in this formulation: we are assuming that past disputes between third countries (i.e. countries other than i and j) have the same relevance as past disputes between a third country and one of the two countries (i, j) . In our working paper we relax this restriction (see also note 31).

Consider first the static version of our extended setting. The key observation is that, for a given country pair (i, j) , the equilibrium joint payoff is decreasing in λ_{ij} . This follows because Proposition 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 between country i and country j , and let Δ_i and Δ_j denote respectively the future payoff changes for country i and country j if these countries go to court. From the perspective of country i , at $t = 2$ there are two possibilities that need to be considered: (i) with probability P^{ij} , country i will be “matched” with country j again, in which case the relevant court experience X_{ij} increases by an amount β_1 ; (ii) with probability P^{io} country i will be “matched” with third country o , in which case X_{io} increases by an amount β_2 . Note that, since only one pair of countries plays the game in each period, $P^{ij} + \sum_{o \neq i, j} P^{io} < 1$.

We can thus write the future payoff gain for country i from going to court with country j at $t = 1$:

$$\Delta_i = P^{ij}[\tilde{\omega}^{ij}(X_{ij} + \beta_1) - \tilde{\omega}^{ij}(X_{ij})] + \sum_{o \neq i, j} P^{io}[\tilde{\omega}^{io}(X_{io} + \beta_2) - \tilde{\omega}^{io}(X_{io})],$$

where $\tilde{\omega}^{ik}(\cdot)$ denotes the expected equilibrium payoff for country i if it is matched with country k in period 2, and where we use the shorthand $\tilde{\omega}^{ik}(X_{rs}) \equiv \tilde{\omega}^{ik}(\lambda(X_{rs}))$. An analogous expression holds for Δ^j .

Clearly, under our assumptions both Δ_i and Δ_j are positive.³⁴ This immediately implies that Proposition 4 extends to this more general setting: intuitively, given that Δ_i and Δ_j are positive, the more governments care about the future, the more likely that they go to court at $t = 1$.

Next focus on the impact of past rulings on the likelihood of current rulings: How do changes in x_{ij} and x_{other} affect the probability of a ruling between i and j at $t = 1$? We now argue that Proposition 5 extends in a natural way to this setting.

First, it is clear that an increase in either x_{ij} or x_{other} must decrease the likelihood of a ruling in an average sense as described in the previous section, provided that learning eventually stops for some finite value of X_{ij} . More precisely, for any fixed x_{ij} there is a high enough value of x_{other} at which learning stops, and hence the probability of a ruling reaches zero, and the same

³⁴Recall that we have assumed a veil of ignorance for each pair of countries, so the only reason why Δ_i and Δ_j might be different from each other is that country i and j have different probabilities of engaging in a dispute with a third country at $t = 2$.

is true when these variables are flipped. Thus, applying the same logic as in the two-country model, we can state that the impact of both experience variables on the likelihood of a ruling on average must be negative.

Furthermore, if δ is sufficiently high, the probability of a ruling is weakly decreasing in each of the experience variables x_{ij} and x_{other} for all values of x_{ij} and x_{other} . Notice that, just as in the previous section, increasing either experience variable has a static effect and a dynamic effect. The static effect is that increasing either experience variable increases (at least weakly) the current efficiency of the court, 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 increasing each experience variable weakly reduces Δ_i and Δ_j . 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 then follows that, if δ is large enough, the dynamic effect outweighs the static effect, and the probability of a ruling is weakly decreasing in x_{ij} and x_{other} .

Finally, applying a similar logic as in the previous section, it is easy to show that also the predictions of the model regarding the likelihood of disputes, as summarized in Proposition 6, extend in a natural way to this many-country setting.

5. Empirical Evidence

We now investigate the empirical content of our theory with the help of data from WTO trade disputes. Our central objective is to assess the empirical importance of learning by ruling. The main challenge is that, unlike the existing empirical work on learning by doing for firms, where direct measures of productivity are available, we cannot observe directly the quality of the court, so we cannot estimate directly the relationship between court experience and court quality. We address this challenge by using the predictions of our model to indirectly gauge the importance of learning by ruling. In particular, our model suggests that if court learning is important we should find a strong negative impact of cumulative rulings on the likelihood of current rulings and disputes; and that the scope of court learning can be gauged by examining how the likelihood of current rulings and disputes is affected by different measures of court experience, e.g. disputant-specific or general-scope experience. We will also consider the plausibility of some key alternative interpretations for our empirical findings.

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.

In our empirical investigation we will focus on the determinants of rulings, while omitting here the description of a parallel investigation of the determinants of disputes.³⁵ The reason is that, while our data on the frequency of DSB rulings is quite reliable, the data on the frequency of disputes suffers from a potentially important limitation. This is 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, data only exist on settled disputes that occur within the formal WTO dispute process, and so we have no data on disputes between WTO members that occur outside the WTO process and are settled informally, never making it to the WTO. Thus there is a potentially important sample selection issue when measuring the frequency of disputes that does not arise with our analysis of rulings.³⁶

Recall that our multi-country model of section 4 allows for the possibilities of *dyad*-specific learning and *general-scope* learning (as well as a range of intermediate possibilities). But our model considers only one sector or issue area. For empirical purposes, it seems compelling to allow for an additional 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.³⁷ If learning can be specific to the GATT/WTO Article ruled upon by the court, then we have two additional possibilities beyond dyad-specific and general-scope learning: court learning could also be *article specific* or it could be *dyad-and-article* specific. And of course, combinations of these different dimensions of learning might be present but in different degrees. Below we investigate empirically each of these different potential domains of court learning, assessing the evidence that learning might be general, dyad specific, article specific or dyad-and-article specific.

³⁵In our working paper Maggi and Staiger (2019) we describe parallel empirical results relating to disputes.

³⁶It is important to note that we examine the unconditional likelihood of rulings, rather than the likelihood of a ruling conditional on a dispute. Recall that our model yields predictions about the unconditional likelihood of rulings, while as Remark 2 highlights, the same is not true for the likelihood of rulings conditional on a dispute.

³⁷The initiation of a WTO dispute involves a description by the complainant country of the specific policy commitments that are the subject of the dispute, and these policy commitments are enumerated in the various GATT and WTO Articles (e.g., Article VI on antidumping measures, Article XI on quantitative restrictions). Hence our use of the GATT/WTO Articles named in the dispute as representing the dispute issue areas seems the natural approach.

In what follows we use $R_{ij,k,t}$ to denote the number of country-dyad- ij disputes on article k that led to an adopted panel ruling in year t ; and we use $CR_{ij,k,t}$ to denote the cumulative number of country-dyad- ij disputes on article k that led to an adopted panel ruling prior to year t . We will refer to $R_{ij,k,t}$ simply as the number of “rulings” for dyad ij on article k in year t , and similarly for the variable $CR_{ij,k,t}$. Notice that our convention is to date DSB rulings by the year in which the DSB panel report containing the ruling is formally “adopted” (approved) by the WTO membership. This 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 court learning.

5.1. Main Results

We now present some simple regressions, following an approach similar to Kellogg (2011), who looks for evidence of various forms of learning by doing in the context of drilling activity in the oil and gas industry (e.g., producer-specific, rig-specific and producer-rig-specific learning). We exploit a key prediction of our theory and look for evidence of various forms of judicial learning by regressing the current frequency of WTO rulings on various cumulative court experience variables and controls.

Any attempt to identify the effect of cumulative past rulings on the current likelihood of a ruling must confront two basic issues. First, there may be spurious positive correlation due to unobserved serially-correlated shocks or cross-sectional heterogeneity. Importantly, this would introduce a bias *against* the prediction of our model, but at any rate we will attempt to minimize such bias by including fixed effects and other controls. Second, cumulative rulings are positively correlated with calendar time, so if there are unobserved determinants of rulings that decline over time, this will generate a spurious negative correlation that we might erroneously attribute to learning. We will address this issue by controlling for a (potentially non-linear) time trend.

We present results from our regressions under both logit and OLS estimation. We focus our discussion in the text on the logit results, but we point out where our logit and OLS results diverge and emphasize only those findings that are common to both. We estimate our regressions with a panel spanning the 15 years 1995-2009 and consisting of observations on each of the 55 country dyads that generated at least one WTO adopted Panel ruling report as a result of a dispute initiated during this period and 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.³⁸

The dependent variable in the logit regression is $RLogit_{ij,k,t}$, defined as 1 if $R_{ij,k,t} \geq 1$ and 0 otherwise.³⁹ The key independent variables of interest are four measures of cumulative past rulings, which we denote by $CR_{ij,k,t}$, $CR_{n(ij),k,t}$, $CR_{ij,nk,t}$ and $CR_{n(ij),nk,t}$ where a subscript nz denotes “not z ” for index z .⁴⁰ The variable $CR_{ij,k,t}$ is the cumulative number of rulings for dyad ij on article k prior to year t ; this variable captures *dyad-and-article* specific court experience, and reflects the narrowest form of court experience that is specific to both the disputants involved and the article that they are disputing. The variable $CR_{n(ij),k,t}$, defined as the cumulative number of rulings for dyads other than ij on article k prior to year t , captures *article-specific* court experience. The variable $CR_{ij,nk,t}$, defined as the cumulative number of rulings for dyad ij on articles other than k prior to year t , captures *dyad-specific* court experience. And finally, the variable $CR_{n(ij),nk,t}$, defined as the cumulative number of rulings for dyads other than ij on articles other than k prior to year t , captures *general-scope* court experience. Table 1 provides summary statistics for each of the variables used in the regressions (including both summary statistics for variables used in our ruling regressions and for variables used in the dispute regressions that, as we explain below, we present as a check on our ruling regressions).

Each regression includes a quadratic time trend, as well as dyad- and article- fixed effects to control for unobserved heterogeneity in the rulings behavior at the level of the dyad (the countries in dyad ij may have a particularly litigious relationship) and the level of the article

³⁸Our 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 dyads in our regressions rise 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 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 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.

³⁹The dependent variable for the OLS regression is $R_{ij,k,t}$. We have also re-estimated our OLS regression with the 0/1 dependent variable used in our logit regression, and find that the results are not materially impacted. And analogously, making use of the count variable $R_{ij,k,t}$ and re-estimating our logit regression as a Poisson or negative binomial regression makes no material difference to our results.

⁴⁰As we noted at the outset of this section, Remark 2 and our discussion following this Remark guide us to look 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.

(article k may be particularly susceptible to litigation).⁴¹ Importantly, we do not include a dyad-by-article (ij, k) fixed effect, and therefore do not control for unobserved heterogeneity at the level of the dyad *and* article (the countries in dyad 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- (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- (ij, k) variation over time.⁴² And second, the right-hand-side variable $CR_{ij,k,t}$ is the sum of lagged values of the dependent variable, and inclusion of an (ij, k) fixed effect would introduce an incidental parameters problem and lead to biased and inconsistent estimates for our relatively short panel.⁴³ An implication is that, if there *is* important unobserved heterogeneity at the dyad-and-article level, our estimates of the coefficient on $CR_{ij,k,t}$ will be biased upward, a bias that works against finding evidence of the most narrow form of learning. We will return to this point below.

The results of the logit regression are presented in column 1 of Table 2, with the corresponding OLS results contained in column 2 of Table 2.⁴⁴ The estimated coefficients on both $CR_{n(ij),k,t}$

⁴¹We 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.

⁴²Relatedly, we choose to include a quadratic (and also a cubic, as we discuss below) 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-scope learning variable. Note that our general-scope learning variable does exhibit some within-year variation over the cross-section of (ij, k) (because rulings occur on specific dates within each year, as do dispute initiations, and hence these events may be associated with different magnitudes of cumulative rulings even when they involve the same disputants and the same article and occur in the same year), so identification of a general-scope learning effect in the presence of year fixed effects is possible; but this variation alone may be insufficient to detect the effect. And as we report below, even without year fixed effects we find only weak evidence of a general-scope learning effect.

⁴³Letting T denote the length of the panel, the issue that arises if an (ij, k) fixed effect is included in regressions of Table 2 is that for T fixed and relatively small, the estimates of the slope parameter on $CR_{ij,k,t}$ will be biased and inconsistent even as the ij and k dimensions of the panel become large. This is because the number of (ij, k) fixed effects to be estimated grows proportionately with the 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 $CR_{ij,k,t}$; 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.

⁴⁴For 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.

and $CR_{ij,nk,t}$ are negative and strongly significant in both the logit and OLS regressions. On the other hand, the coefficient estimate on $CR_{n(ij),nk,t}$ is negative and significant in the logit regression but insignificantly different from zero in the OLS regression. Taken together, these coefficient estimates suggest the presence of dyad-specific and article-specific court learning, but only weak evidence of general-scope learning.⁴⁵ The average marginal effects for $CR_{n(ij),k,t}$ and $CR_{ij,nk,t}$ are -0.0008 and -0.0003 , respectively, which are sizeable in light of the mean value of the dependent variable $RLogit_{ij,k,t}$ (0.0056); for example, according to these estimates an additional ruling on article k in a dispute between i and j is all else equal predicted to lower the probability of a future dispute between these countries on article k by 15%. Finally, notice that the point estimate of the coefficient on $CR_{ij,k,t}$, our narrowest measure of DSB experience, is positive (and strongly significant according to the OLS estimates in column 2). One would expect that if court learning occurs anywhere, it should be found here. However, our failure to find a negative coefficient on $CR_{ij,k,t}$ may reflect the upward bias in this coefficient that would be expected when estimating the regressions in Table 2 if there is unobserved heterogeneity at the dyad-and-article level.

As a partial check on the interpretation that our estimated coefficient on $CR_{ij,k,t}$ reflects the presence of unobserved heterogeneity at the dyad-and-article level, we next present estimates of *dispute* regressions (logit and OLS), first without an (ij, k) fixed effect under the same specifications as in our ruling regressions of Table 2, and then with an (ij, k) fixed effect included. Recall that inclusion of this fixed effect will diminish our ability to assess the impact of those variables that exhibit little within- (ij, k) variation over time, but it should address the upward bias in the estimated coefficient on $CR_{ij,k,t}$ induced by any unobserved heterogeneity at the dyad-and-article level. The reason we now focus on disputes, rather than rulings, as the dependent variable, is that for the dispute regressions the inclusion of an (ij, k) fixed effect does not mechanically lead to biased or inconsistent estimates due to an incidental parameters problem as would be the case if we were to introduce an (ij, k) fixed effect into the rulings regressions, where $CR_{ij,k,t}$ constitutes a lagged dependent variable. This makes the dispute regressions better suited for the purpose at hand, despite the potential sample-selection issues with the dispute data that we discussed above.

⁴⁵When we add a cubic term to the time trend, the evidence of article-specific and dyad-specific learning across both logit and OLS specifications remains, and there is now also evidence of general-scope learning across both logit and OLS specifications; but since the evidence for general-scope learning hinges on the inclusion of the cubic time trend, we interpret our results as overall providing only weak evidence of general-scope learning, as we report in the text.

The results are contained in Table 3. Defining $D_{ij,k,t}$ as the number of disputes initiated by country-dyad ij on article k in year t , and defining $DLogit_{ij,k,t}$ as 1 if $D_{ij,k,t} \geq 1$ and 0 otherwise, columns 1 and 2 present the logit and OLS dispute regressions, respectively, without an (ij, k) fixed effect under the same specifications as in our ruling regressions of Table 2, while columns 3 and 4 present these regressions with an (ij, k) fixed effect included.⁴⁶ Without the (ij, k) fixed effect, the coefficient on $CR_{ij,k,t}$ is positive and strongly significant, much like the results for our ruling regressions. But when an (ij, k) fixed effect is included in the dispute regressions, the coefficient on $CR_{ij,k,t}$ turns strongly and significantly negative, consistent with our interpretation that there is unobserved heterogeneity at the dyad-and-article level, and suggesting that the upward bias in the estimated coefficient on $CR_{ij,k,t}$ that would be expected from such unobserved heterogeneity may be masking evidence of the presence of court learning at the dyad-and-article level in our ruling regressions of Table 2.

In our working paper Maggi and Staiger (2019), we consider in both our theory and our empirical investigation richer possibilities for the scope of learning. There we allow for the possibility of *directed-dyad-specific* learning, that is, 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. This could be the case if for example the court learns about features of the industry in country j that competes with imports from country i . We also allow for the possibility of *defendant-specific* learning, 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, and similarly for the possibility of *claimant-specific* learning. When we allow for these richer learning possibilities, we continue to find strong evidence of article-specific learning and only weak evidence of general-scope learning, while we find strong evidence of two finer forms of disputant-specific learning, namely directed-dyad-specific learning and complainant-specific learning.

To summarize, our empirical findings are consistent with the presence of significant article-specific and disputant-specific court learning, while we find little evidence of general-scope learning. It is also worth noting that the coefficient on the linear time trend is positive in all of our regressions. The fact that calendar time does not have a negative impact on the frequency

⁴⁶The drop in the number of observations for the logit regression in column 3 of Table 3 reflects the issue of quasi-separation that arises when the (ij, k) fixed effect is included.

of disputes and rulings once we control for our measures of court experience (the *CR* variables) suggests that court learning can help explain the raw declining trend in WTO disputes and rulings displayed by Plots 1 and 2 and discussed in the Introduction.

5.2. Alternative Interpretations

Thus far we have interpreted our empirical findings as reflecting the effects of court learning, and of court learning that embodies a particular scope and form. And we have suggested that court learning may be responsible for the declining trends observed in the numbers of WTO disputes and rulings. An important question is whether there are alternative interpretations of these empirical findings. In this section we consider the plausibility of some key alternatives.

A first possibility is that the relationships we estimate between cumulative rulings and the likelihood of current rulings may reflect the presence of *government learning about the court*. In particular, by observing how the court operates, governments may learn to better predict the outcome of rulings. If governments' beliefs about the outcome of rulings are initially asymmetric, they may converge as a result of this kind of learning, and as a consequence the frequency of equilibrium rulings may decline. This might be the case, for example, if governments learn about the court's preferences and possible biases, or about the applicable legal precedent (Priest and Klein, 1984). 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.

If we assume that a government learns about the court more effectively when participating as a third party in a dispute that goes to a ruling than when observing the same dispute as an outsider, there is a simple way to gauge whether this type of learning is important: a government's experience as third-party observer of rulings should be a stronger predictor of the likelihood of current rulings than a government's experience as outside observer of rulings.⁴⁷ To consider this possibility, we have re-estimated each of our regressions breaking up our experience

⁴⁷More specifically, suppose that disputants go to court when they have inconsistent and overly-optimistic priors about the court's expected ruling. And, to make things sharper, suppose that when a dispute arises and the court is invoked, the countries that join the dispute as third-party participants learn about the court and their priors converge toward the truth, while the "outsider" countries learn nothing. Then we would expect that countries i and j are less likely to go to court the greater the cumulative number of times that i and/or j have been third-party dispute participant (because their priors are less likely to be far from the truth and hence less likely to be inconsistent), while the cumulative number of times that i and/or j have been outside observers to rulings should not have an impact on their likelihood of going to court.

variables $CR_{n(ij),k,t}$ and $CR_{n(ij),nk,t}$ into two components: the number of past rulings where at least one of the disputing parties i or j was a third-party participant, and the remaining ones (where both i and j were “outsiders”). We find that the third-party experience variables do not have a stronger effect than the outsider experience variables; thus by this metric we fail to find evidence of government learning about the court in our data.

A second possibility is that, by repeatedly engaging in disputes, *governments may learn about each other*. In principle this type of learning could explain a negative impact of cumulative past rulings on the current likelihood of rulings, because reducing or eliminating private information may reduce the probability of bargaining failures and hence decrease the propensity to go to court.

Our data however do not support the government-learning-about-each-other story, for two reasons. First, an immediate implication of this story would be that governments’ learning about each other should increase the settlement rate.⁴⁸ Contrary to this prediction, the trend in the settlement rate in WTO disputes has been flat or slightly negative (see for example, Pauwelyn and Zhang, 2018). Second, we have re-run our regressions 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 negative impact on the frequency of subsequent rulings, we would expect this to show up as negative and significant coefficients on the CS variables pertaining to the dyad of the consulting parties (that is, on the $CS_{ij,k,t}$ and $CS_{ij,nk,t}$ variables). In fact, we fail to find any robust evidence for such coefficient estimates.⁴⁹

Having discussed alternative explanations for the impact of cumulative rulings, we now

⁴⁸By contrast, recall from Remark 2 that our model does not imply that the settlement rate should increase with cumulative rulings.

⁴⁹Notice that, in contrast to our theory, neither of these alternative explanations requires that countries be “large players” who internalize the impacts on court quality of their decision to litigate. In this light, one might wonder why we do not pay special attention in our empirical investigation to the most frequent litigants, on the conjecture that the strongest evidence for our theory should be found within this subsample of countries. But this conjecture does not follow from our theory, because as we have indicated the future benefits of a higher-quality court are not limited to better rulings, but also arise from better settlements (off-equilibrium effects) and even better policy choices made to avoid disputes entirely (off-off-equilibrium effects). In principle, then, even a government who is observed to seek a single ruling and then never engages in another dispute may have been incentivized to seek the ruling in order to enjoy the future benefits of a better court. All that is required is that the government expects that there will be future interactions with other governments.

discuss two alternative explanations for the declining trends in the numbers of WTO disputes and rulings (as evidenced in Plots 1 and 2) that do not go through the channel of cumulative rulings.

A first possibility, suggested by Pauwelyn and Zhang (2018, see note 2), is that the increasing delays experienced in the process of adjudicating increasingly complex WTO disputes are contributing to the declining trend in WTO disputes and rulings. But our regressions are picking up something more: the Pauwelyn-Zhang story would not predict that declines in disputes and rulings would be systematically related to our *CR* variables, but rather this story would more naturally be picked up by 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 as we noted we do not find strong general-learning effects.

Second, it is possible that the decline in disputes and rulings may reflect the impacts of the growing membership in preferential trade agreements (PTAs), on the grounds that (i) it has been observed (e.g., Mavroidis and Sapir, 2015) that PTA membership may reduce the reliance of countries on the WTO to resolve their trade disputes and (ii) our 1995-2009 sample period witnessed unprecedented growth in the numbers of PTAs to which WTO members belong. The point made just above for the Pauwelyn-Zhang story applies here too: while the growth in PTA membership can potentially explain the downward trends in rulings and dispute, it does not explain the *CR* effects we find in the data. But in addition to this conceptual point, our data allow us to investigate whether PTA membership contributes to explain the variation in the likelihood of rulings, and through what channel. We have re-estimated all of our regressions controlling for a PTA-partner variable (constructed using the NSF-Kellogg Institute Data Base on Economic Integration Agreements for September 2015) on the right hand side. Our results indicate that the PTA partnership does indeed reduce the likelihood of WTO rulings, but inclusion of the PTA partnership variable on the right-hand side of our regressions leaves our findings relating to both the cumulative rulings variables and the positive coefficient on the linear time trend unchanged.

Finally, one may ask whether our empirical findings admit only the interpretation we have given them *when viewed through the lens of our model*. 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 to this question is “Yes.” To this end, we return to our multi-country 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 two non-negative parameters (β_1, β_2) describing the nature and scope of court learning, with two corresponding experience variables, x_{ij} and x_{other} . Suppose data can be used to estimate the derivatives of the likelihood of rulings with respect to x_{ij} and x_{other} . 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 in a dispute between i and j is decreasing in x_{ij} (measuring dyad-specific experience) while it is essentially independent of x_{other} (measuring general-scope experience). It can be shown that, according to the model, this implies that β_1 is positive while β_2 is zero. It is an extension of this logic to a setting with multiple issue areas that underlies our statements above that the data is consistent with disputant-specific and issue-specific court learning.

6. Conclusion

The importance of judicial learning has been widely emphasized in the context of domestic court systems, but little attention has been paid to judicial learning in the context of international institutions such as the WTO. Judicial learning in international institutions is likely to have distinct implications from those in domestic court systems because the number of potential disputants is small and thus the impact of current rulings on the future efficiency of the court is at least in part internalized by the disputants.

We have explored the implications of learning by ruling through a two-country model where both the initiation of disputes and the occurrence of rulings are endogenous, governments bargain “in the shadow of the law,” and the efficiency of the court increases with its experience. We have shown that the presence of judicial learning can explain litigation on the equilibrium path, since going to court today may imply payoff gains for the governments tomorrow, and that the probability of rulings is higher when governments are more patient. The future payoff gains from increasing court quality occur even if governments do not go to court tomorrow, because a more efficient court implies a better disagreement point for tomorrow’s bargain (an off-equilibrium effect), and even if governments do not have a dispute tomorrow, because a more efficient court induces governments to choose better policies in order to avoid disputes (an off-off-equilibrium effect). These off- and off-off-equilibrium effects shape the implications of learning

by ruling in important ways. Our model predicts that the likelihood of both disputes and rulings should tend to decrease with court experience (cumulative rulings). This prediction is sharper if governments are sufficiently patient, but holds in an average sense also if the governments' patience is low. And we have shown how this prediction extends to a many-country setting. Underlying these results is the interplay between a "static" effect of increasing court experience, which pushes toward a higher likelihood of disputes and rulings, and a "dynamic" effect, which pushes in the opposite direction.

We have also confronted the theory with data from WTO disputes. Our empirical findings suggest that judicial learning is an important phenomenon in the WTO. More specifically, our evidence is consistent with article-specific and disputant-specific learning, but we find only weak evidence of general-scope learning. And we have argued that our empirical findings suggest that court learning may be responsible for the declining trends observed in the numbers of WTO disputes and rulings.

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

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

where CS is consumer surplus, PS is producer surplus and R is tariff revenue, and where $\gamma \geq 1$ is a parameter that captures the political importance of the domestic industry (in the spirit of Baldwin, 1987, and Grossman and Helpman, 1994). 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:

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

We assume the demand and supply functions are linear in both countries. Note that ω^* 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 ω is concave in T provided γ 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 $\omega_{TT} + \omega_{TT}^* < 0$, and this is the case if γ does not exceed a threshold level.⁵⁰ In what follows we assume this condition is satisfied.

Let $T^{fb}(\gamma)$ denote the “first best” level of T , which maximizes the government joint payoff $\omega + \omega^*$ given γ . We assume that γ is ex-ante uncertain. Its realization is perfectly observed by governments but imperfectly observed by the court.

We can consider two equivalent forms of contract. One possibility is that the contract specifies a contingent tariff schedule $T^{fb}(\gamma)$; in this case, if the court is invoked, it estimates $T^{fb}(\gamma)$ and imposes it on the importing government. The second possibility is that the contract

⁵⁰Clearly, this threshold level is lower than the threshold level of γ below which ω is concave in T , but higher than one (because if $\gamma = 1$ both governments maximize welfare).

takes a discretionary form in that it does not specify the policy T at all, but endows the court with the authority to “fill the gap” in this contract ex-post; in this case, if the court is invoked, it maximizes the governments’ expected joint payoff given its noisy information. In either case, we assume that the court observes a noisy signal of $T^{fb}(\gamma)$ given by $T^{court} = T^{fb}(\gamma) + \varepsilon$, where ε is a white noise with mean zero and variance λ , and prescribes the tariff level T^{court} .⁵¹

In this example we abstract from litigation costs, so the court-inefficiency parameter λ captures simply the noise in the court 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 court 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

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

where ω^{fb} and ω^{*fb} are respectively the domestic and foreign payoffs given the first best tariff. Recall from the discussion above that $\omega_{TT} < 0$, $\omega_{TT}^* > 0$ and $\omega_{TT} + \omega_{TT}^* < 0$ for all γ . Thus, increasing the court noise λ worsens Home’s disagreement payoff, improves Foreign’s disagreement payoff, and worsens the joint disagreement payoff.

Note that: (i) $\Omega^R = \Omega^{fb} + \lambda(\omega_{TT} + \omega_{TT}^*)$ is strictly decreasing in λ , (ii) if $\lambda = 0$ then Ω^R coincides with Ω^{fb} , and (iii) by the expression for ω^{*R} above, $\frac{d\omega^{*R}}{d\lambda} \neq 0$ for all λ ; thus our reduced-form Assumption 1 is satisfied. Next note that the net bargaining payoffs are

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

⁵¹We could have assumed that the court observes a noisy signal of γ rather than a noisy signal of $T^{fb}(\gamma)$, at the cost of a slightly more complicated analysis. Recall that in the reduced-form model in the main text we assumed that the court observes a noisy signal of the state of the world s , thus the structural tariff model presented here can be mapped into the reduced-form model if we think of $T^{fb}(\gamma)$ as the state of the world.

Clearly, these payoffs are linear in λ for any γ , hence the B_{net}^λ locus is linear, it cannot be backward bending and can cross the no-transfer frontier at most once. It follows that Propositions 1, 2 and 3 hold 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 3 needs to be imposed just as in the more general model, Assumption 2 is automatically satisfied here, because disagreement payoffs are linear in λ , and λ is convex in x . As discussed in the main text, Assumption 3 directs attention to what seems to be the empirically relevant case given that the frequency of disputes is empirically very low. Under this assumption, all the results of the dynamic setting hold in this example model.

B. Conditions under which B_{net}^λ is not backward bending

Clearly the B_{net}^λ locus is not backward bending if $\frac{d\omega^{*B_{net}}}{d\lambda} \neq 0$ for all λ . To check under what conditions this is satisfied, we use

$$\omega^{*B_{net}} = \omega^{*R} + \kappa(1 - \sigma)(\Omega^{fb} - \Omega^R)$$

to derive

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

With this we may conclude that $\frac{d\omega^{*B_{net}}}{d\lambda} \neq 0$ for all λ if and only if

$$\frac{d\omega^R/d\lambda}{d\omega^{*R}/d\lambda} \neq \frac{1 - \kappa(1 - \sigma)}{\kappa(1 - \sigma)} \text{ for all } \lambda. \quad (7.1)$$

To evaluate this condition, first note that the left-hand side is the slope of the \mathcal{R}^λ locus, which can be either positive or negative but which is finite by Assumption 1(iii). Next note that the right-hand side of this condition is non-negative, and goes to infinity if $\kappa(1 - \sigma)$ goes to zero. Hence, the condition is automatically satisfied when the \mathcal{R}^λ locus is downward sloping (as in the special model described in Appendix A). And if the \mathcal{R}^λ locus is upward sloping, the condition will still be satisfied provided that either the negotiation cost is large enough or Home's bargaining power is high enough.

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 2011.⁵² 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).⁵³ 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.⁵⁴ After this set of exclusions we are left with 348 WTO disputes.

⁵²Each 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.

⁵³The 8 excluded cases are DS3, DS16, DS52, DS101, DS106, DS185, DS228, DS271, which repectively return or are handled in DS41, DS27, DS65, DS132, DS126, DS187, DS230, DS270.

⁵⁴The 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.
- Besanko, D. et al., (2010), Learning-by-Doing, Organizational Forgetting, and Industry Dynamics, *Econometrica*, 78(2): 543-508.
- 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.

- Cabral, Louis M. and Michael H. Riordan (1994), "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica* 62(5): 1115-1140.
- 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.
- Daughety, Andrew F. Daughety and Jennifer F. Reinganum (1999), "Stampede to Judgment: Persuasive Influence and Herding Behavior by Courts," *American Law and Economics Review* 1(1/2) Fall 1999: 158-189.
- Davis, Christina L. and Sarah Blodgett Bermeo (2009), "Who files? Developing Country Participation in GATT/WTO Adjudication," *Journal of Politics* 71 (3).
- Gennaioli, Nicola and Andrei Shleifer (2007), "The Evolution of Common Law," *Journal of Political Economy*, 115(1): 43-68.
- 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.
- Hadfield, Gillian K. (1992), "Bias in the Evolution of Legal Rules," *Georgetown Law Journal* 80: 583: 583-616.
- Hadfield, Gillian K. (2011), "The dynamic quality of law: The role of judicial incentives and legal human capital in the adaptation of law," *Journal of Economic Behavior & Organization* 79: 80-94.


- 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 (2018), “Trade Disputes and Settlement”, *International Economic Review* 59(1): 19-50.
- Maggi, Giovanni and Robert W. Staiger (2019), “Learning by Ruling and Trade Disputes,” working paper version, February 2019.
- 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.
- Ossa, Ralph, Staiger, Robert W. and Alan O. Sykes (2020), “Disputes in International Investment and Trade,” mimeo.

- Parameswaran, Giri (2018), “Endogenous cases and the evolution of the common law,” *RAND Journal of Economics* 49(4) Winter:791–818.
- Park, Jee-Hyeong (2011), “Enforcing International Trade Agreements with Imperfect Private Monitoring,” *Review of Economic Studies*.
- Pauwelyn, Joost and Weiwei Zhang (2018), “Busier than Ever? A Data-Driven Assessment and Forecast of WTO Caseload,” *Journal of International Economic Law* 21(3), September: 461-487.
- Pelc, Krzysztof J. (2014), “The Politics of Precedent in International Law: A Social Network Application,” *The American Political Science Review* 108(3): 547-564
- Posner, Richard A. (1973), *The Economic Analysis of Law*. Boston: Little, Brown.
- Posner, Richard A (2005), “The Law and Economics of Contract Interpretation,” *Texas Law Review*, 83, 1581-1614.
- Priest, George L. and Benjamin Klein (1984), “The Selection of Disputes for Litigation,” *The Journal of Legal Studies*, Vol. XIII, January.
- 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.
- Spence, A., (1981), The Learning Curve and Competition, *Bell Journal of Economics*, 12(1): 49-70
- 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.
- Talley, Eric (1999), “Precedential Cascades: An Appraisal,” *Southern California Law Review* 74: 87-138.
- 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.

WTO (2005), *European Communities - Customs Classification of Frozen Boneless Chicken Cuts - Report of the Appellate Body*, WT/DS269/AB/R and WT/DS286/AB/R.

Zdouc, Werner (2008), “Features of the Appellate Body that have Defined its Performance,” Chapter 25 in Janow, Merit E., Donaldson, Victoria and Alan Yanovich (eds.) *The WTO: Governance, Dispute Settlement, and Developing Countries*, Juris Publishing Inc.: New York.



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.

Figure 1: Equilibrium payoffs in the static setting




Figure 2: Impact of court quality in the static setting




Figure 3: Two-Period Setting

Figure 3a: When a ruling occurs in equilibrium




Figure 3b: When a ruling does not occur in equilibrium




Table 1: Summary Statistics

VARIABLES	Ruling Regression					Dispute Regression				
	N	mean	sd	min	max	N	mean	sd	min	max
Dependent Variable (logit)	112,560	0.00565	0.0750	0	1	439,584	0.00587	0.0764	0	1
Dependent Variable (ols)	112,560	0.00589	0.0797	0	3	439,584	0.00635	0.0863	0	5
CR _{ij,k,t}	112,560	0.0346	0.218	0	5	439,584	0.00885	0.111	0	5
CR _{n(ij),k,t}	112,560	1.861	3.589	0	34	439,584	1.101	2.925	0	34
CR _{ij,nk,t}	112,560	4.803	9.772	0	78	439,584	2.123	6.933	0	78
CR _{n(ij),nk,t}	112,560	258.7	218.2	0	631	439,584	264.3	221.0	0	631

Table 2: Main Results

VARIABLES	(1) Logit RLogit _{ij,k,t}	(2) OLS R _{ij,k,t}
CR _{ij,k,t}	0.119 (0.0913)	0.0138*** (0.00470)
CR _{n(ij),k,t}	-0.160*** (0.0168)	-0.000849** (0.000354)
CR _{ij,nk,t}	-0.0644*** (0.00540)	-0.000525*** (0.000126)
CR _{n(ij),nk,t}	-0.00663*** (0.00199)	-8.13e-06 (1.76e-05)
<i>t</i>	1.277*** (0.111)	0.00317*** (0.000699)
<i>t</i> ²	-0.0313*** (0.00416)	-0.000106 (6.77e-05)
Constant	-10.90*** (0.723)	0.00311 (0.00377)
Observations	112,560	112,560
R-squared		0.021
Pseudo R-squared	0.229	
ijFE	Y	Y
kFE	Y	Y
CE	N	Y

Standard errors in parentheses

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

ijFE are dyad fixed effects. kFE are article fixed effects.

CE are standard errors clustered by dyads.

Table 3: Dispute Regressions with and without Dyad-and-Article Fixed Effects

VARIABLES	Without (ij,k) Fixed Effects		With (ij,k) Fixed Effects	
	(1) Logit DLogit _{ij,k,t}	(2) OLS D _{ij,k,t}	(1) Logit DLogit _{ij,k,t}	(2) OLS D _{ij,k,t}
CR _{ij,k,t}	0.186*** (0.0646)	0.0345*** (0.00734)	-0.932*** (0.0853)	-0.148*** (0.0128)
CR _{n(ij),k,t}	-0.0127* (0.00761)	-0.00140*** (0.000474)	0.0236*** (0.00835)	3.07e-06 (0.000268)
CR _{ij,nk,t}	-0.0272*** (0.00256)	-0.00135*** (0.000362)	-0.0134*** (0.00254)	-0.000593* (0.000304)
CR _{n(ij),nk,t}	-0.00190** (0.000841)	-1.09e-05 (1.32e-05)	-0.00219*** (0.000845)	-1.68e-05 (1.38e-05)
<i>t</i>	0.198*** (0.0260)	0.00130** (0.000603)	0.204*** (0.0262)	0.00130** (0.000603)
<i>t</i> ²	-0.00726*** (0.00225)	-1.84e-05 (3.17e-05)	-0.00743*** (0.00225)	-1.84e-05 (3.16e-05)
Constant	-9.201*** (1.076)	-0.00806** (0.00329)		0.00428** (0.00181)
Observations	439,584	439,584	26,253	439,584
(Pseudo) R-squared	0.277	0.044	0.0392	0.024
(ij,k)FE	N	N	Y	Y
CE	N	Y	N	Y

Standard errors in parentheses

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

(ij,k)FE are dyad-and-article fixed effects.

CE are standard errors clustered by dyads.