

LEARNING BY RULING: A DYNAMIC MODEL OF TRADE DISPUTES* (Preliminary Draft)

Giovanni Maggi[†] Robert W. Staiger[‡]

February 14, 2016

Abstract

Over the WTO years, the frequency of disputes and court rulings has trended downwards. This trend has sometimes been interpreted as a symptom that the WTO institution is becoming less effective over time. In this paper we propose a theory that can explain the observed trends as a result of judicial learning, thus according to our theory such trends may represent good news, not bad news. We then explore whether the predictions of our model are consistent with WTO trade dispute data, and we take a first step towards estimating the strength and scope of court learning.

*We thank participants in seminars at Georgetown, Syracuse, CREI, FGV-Rio, FGV-Sao Paulo and Inesper for very helpful comments and discussions. Junhui Zeng provided outstanding research assistance.

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

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

1. Introduction

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

There are some interesting patterns in the initiation and resolution of disputes over time. Plot 1 shows how the raw numbers of disputes and DSB rulings change over the WTO years: the plot suggests a declining trend both in the frequency of disputes and in the frequency of DSB rulings, although in a more pronounced way for disputes than for rulings, and with some spikes. The visual impression from Plot 1 is that countries “fight” less as the institution ages. If anything, Plot 1 understates this trend because the number of WTO members has increased substantially 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. Plot 2 shows that indeed country dyads engage in fewer disputes and end up less frequently in court as they age.

One can interpret these declining trends in different ways, and some commentators have suggested that this might be because the WTO institution is becoming less effective over time, or governments are losing confidence in it. While this is a legitimate hypothesis, in this paper we will propose a theory that can explain the declining trends in disputes and rulings as a result of institutional learning, thus according to our theory such trends may represent good news, not bad news.

We believe that judicial learning in the WTO is a phenomenon of first-order importance. The WTO is a relatively young institution, and the adjudication of trade disputes is a complex task, so it is reasonable to think that there is significant learning by the actors involved in the WTO’s judicial system, especially in the early stages of the institution. These actors include the Dispute Settlement panels, the Appellate Body and quite possibly also the WTO’s

Secretariat, a group of experts that plays a key role in the dispute settlement process.¹ Another form of institutional learning that is probably important is that governments, as they litigate repeatedly over time, may learn how the court operates and adjudicates cases, thus they may learn to predict more accurately the outcome of a ruling.

In our formal model we will consider only a very simple form of institutional learning, which we brand “learning by ruling”: we will assume that, as the stock of cumulative rulings increases, the court becomes more accurate in evaluating the economic and political costs/benefits of trade policies. This may be, for example, because the court learns to better infer these cost/benefit trade-offs from the legal nuances of the WTO contract, and thereby learns to better translate the contract into the intent of the contracting parties; or it may be because the court learns to use and interpret data and to make more effective use of rigorous economic reasoning in arriving at its rulings.

The importance of judicial learning in practice has been emphasized by many legal scholars and political scientists, although typically in the context of domestic legal systems, not international institutions. For an interesting informal and personal account, see John Paul Stevens (2005).

Before proceeding, we note that in the WTO there may be another dynamic mechanism that has a similar flavor as learning but has distinct implications, namely *legal precedent*. Our model will focus only on learning and abstract from legal precedent, but we will come back to this notion in the empirical part of the paper, because in principle the effect of legal precedent could explain some of the patterns we observe in the data.

The key ingredients of our model are the following: (i) if a dispute is initiated, governments (whose objective functions may include political-economy concerns) bargain “in the shadow of the law,” subject to negotiation costs; (ii) if invoked, the court issues a ruling (with the objective of maximizing the governments’ joint payoff) based on a noisy signal of the state of the world; (iii) the court becomes more accurate as the stock of cumulative rulings increases, but at a diminishing rate; and (iv) governments repeatedly engage in disputes, so they internalize the benefits of court learning.

We note that most existing models of “bargaining in the shadow of the law” explain equilibrium court intervention as bargaining failure due to incomplete information (or overconfidence

¹The Appellate Body is a standing judicial body, so in this case judges may learn directly from their own experience. But also the Dispute Settlement panel, which is a rotating body, may learn from reading previous panel reports, as panel reports are public.

about the ruling). Our model, in contrast, generates equilibrium court intervention for two different reasons: first, due to learning-by-ruling (and the fact that disputants may interact repeatedly in court), going to court today may imply future payoff gains; and second, if the policy is all-or-nothing in nature and international transfers are costly, two attributes that we argue feature prominently in many trade disputes, an additional reason for equilibrium court intervention is the non-convexity of the bargaining set.

We next preview our main theoretical results. We start by focusing on the case of continuous policy (e.g. a tariff). In this case, we show that in a static setting there can never be a DSB ruling in equilibrium, but there can be a dispute, and a dispute is more likely when the DSB is less accurate. In a dynamic setting, on the other hand, the presence of court learning can give rise to rulings in equilibrium. Other things equal, an equilibrium ruling is more likely when the importing country faces stronger political pressures from import-competing producers. This is due to the fact that, when the weight attached by a government on producer surplus is higher, its objective function is less concave in the tariff, and hence the Pareto frontier becomes less concave, therefore the static inefficiency from going to court (due to the DSB noise) is smaller.

When we examine how the likelihood of current rulings depends on court experience (cumulative rulings), we find that this relation is decreasing, at least if governments are patient enough; and even if governments are impatient, this result holds when the stock of cumulative rulings is large enough. The role played by government patience is due to the fact that an increase in court experience has a dynamic effect that makes a current ruling less likely (because the future payoff gain from going to court is diminishing as the court walks down its learning curve), but also a static effect that goes in the opposite way (a decrease in DSB noise reduces the inefficiency of going to court today), and the dynamic effect dominates when the discount factor is sufficiently high.

Our basic model focuses on the case in which court learning is general in scope, in the sense that a ruling today makes the court more accurate tomorrow regardless of which country is the defendant tomorrow, but we also consider the case where the scope of learning is narrower, and in particular where learning is defendant-specific or complainant-specific, and show that the main results described above continue to hold.

Turning to the likelihood of current disputes, if learning is general in scope and governments are patient enough, the likelihood of current disputes decreases with court experience. However, unlike for rulings, if governments are impatient the prediction may not hold even if the stock of

experience is large, and moreover, the prediction may not hold if learning is defendant-specific or complainant-specific. Thus the effect of court experience on disputes is a bit more ambiguous than its effect on rulings.

Coming back to a question we raised at the outset, the model suggests that the frequency of DSB use is not a good measure of the effectiveness of the institution. A declining trend in DSB disputes or rulings does not imply that the quality of the institution declines over time, in fact it is a symptom of beneficial learning. However, this is a statement about the *change* in the frequency of DSB use over time. According to our theory a lower *level* of this frequency may well be a symptom of lower court accuracy: if the level of court accuracy (for given stock of cumulative rulings) is higher, the disagreement point is more likely to be above the Pareto frontier, so the likelihood of a ruling is higher.

We then focus on the case in which the trade policy is all-or-nothing in nature and international transfers are costly. In this case, rulings can occur in equilibrium even in a static setting, and the reason lies in the non-convexity of the bargaining set: the uncertainty in the ruling may help governments share the (expected) surplus more efficiently than by using costly transfers. We show that in this case the likelihood of rulings and disputes decreases with court experience, even if governments are impatient. The reason this prediction holds regardless of the discount factor, unlike for a continuous policy, is that the static effect of an increase in DSB accuracy is to *decrease* the likelihood of a ruling, since this implies a less uncertain ruling and hence the surplus-sharing appeal of going to court is diminished. Thus the static effect and the dynamic effect of an increase in court experience go in the same direction. Furthermore, we note that when the policy is binary, introducing learning implies a “pivoting” of the time path of rulings and disputes, with more rulings/disputes occurring early on and fewer occurring later on.

Finally, we explore the empirical content of our theory using WTO trade dispute data. We focus on a key prediction of the model, namely that the likelihood of current disputes and rulings should tend to decrease with the stock of cumulative past rulings. Our empirical investigation has a dual objective. First, we want to test whether the above prediction is consistent with the data. And second, to the extent that the answer to the previous question is affirmative, we want to gauge the empirical importance of learning-by-ruling. Note that, unlike the existing empirical work on learning-by-doing for firms, we cannot observe directly the productivity/accuracy of the court, so we cannot estimate directly the relationship between court experience and court accuracy; but we can use the predictions of our model to indirectly infer the importance of

learning-by-ruling: the model suggests that, the stronger the effect of cumulative past rulings on the likelihood of rulings and disputes, the more important learning-by-ruling is likely to be.

While our model assumes only two countries and one sector, in a world with many countries and many issue areas the scope of court learning might be general, or specific to the disputing countries, or 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. To fix ideas, suppose that, by ruling on a dispute brought by country i against country j on article k , the court may learn about (i) disputes in general (“general-scope” learning), or (ii) disputes where the defendant is country i (“defendant-specific” learning), or (iii) disputes where the complainant is country j (“complainant-specific” learning), or (iv) disputes that i brings against j (“directed-dyad-specific” learning), or (v) disputes on article k (“article-specific” learning). Or learning could be even more narrow: for example the court might learn about disputes that are brought against country j on article k . And further, these different dimensions of learning might all be present but in different degrees. Our data allows us to disentangle these different dimensions of learning, thus we can attempt to gauge not only the strength but also the scope of court learning.

Our findings are broadly consistent with the model, and interestingly, we find evidence consistent with article-specific learning and with disputant-specific learning (in particular, complainant-specific and directed-dyad-specific), while we only find weak evidence of general-scope learning. We also find evidence of mechanisms that are outside our model, and in particular of “tit for tat” behavior and of a “bandwagon” effect in dispute initiation, suggesting possible directions of extension of our model.

To our knowledge this is the first paper that explores the implications of judicial learning for trade disputes, or more generally, the implications of institutional learning for international relations. A related model is Maggi and Staiger (2011), but that paper does not consider learning and does not allow for bargaining or settlement, and focuses on questions of institutional design such as the desirability of legal precedent, while here we focus mostly on how learning affects the 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 form of the contract (property vs liability contract) affects the outcome of trade disputes.

There is a fairly large literature on the implications of judicial learning, but this literature is mostly informal and does not focus on international institutions. A few recent papers have

developed formal models of judicial learning (see for example Baker and Mezzetti, 2012, and Beim, 2014), but the structure and focus of these models is very different from ours.

In the literature on trade agreements, other models that generate trade disputes in equilibrium are Beshkar (2016), Staiger and Sykes (2013) and Park (2011). These papers however do not focus on the determinants of trade dispute outcomes (with the partial exception of Beshkar).

On the empirical side, there are several papers that examine the determinants of the initiation and outcome of trade disputes. Guzman and Simmons (2002) examine whether disputes involving continuous policies are more likely to end in settlement than disputes involving all-or-nothing policies (more on this below). Other papers that examine various determinants of the initiation and outcome of trade disputes include Busch (2000), Busch and Reinhardt (2000, 2006), Guzman and Simmons (2005), Bown (2005), Davis and Bermeo (2009), Kuenzel (2015) and Conconi et al. (2015).²

Finally, our model is related to the law-and-economics literature that focuses 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.

The rest of the paper is organized as follows. In Section 2 we present our benchmark static model, focusing on the case of a continuous policy. In Section 3 we develop our dynamic model with learning-by-ruling. In Section 4 we consider the case of binary policy. In Section 5 we examine the empirical content of our theory through WTO dispute data. In Section 6 we offer some concluding remarks.

2. The static model

We consider a two-country partial equilibrium setting. In the industry under consideration, Home is the importing country and Foreign the exporting country. Home can choose an import barrier T , while the Foreign government is passive in this industry.

In this basic model we focus on the case in which the policy is continuous. For concreteness

²In our above-mentioned paper, Maggi and Staiger (2015), we also examine empirically whether trade disputes are more likely to end in settlement when the contract form is a liability rule or a property rule. See also Bown and Reynolds (2014) for a descriptive exploration of the kind of ‘trade’ that countries fight about (e.g. nature of the products, trade volumes, prices).

we assume T is a tariff. Later we will consider the case of a binary policy.

We assume the Home government maximizes a weighted welfare function which allows for political economy considerations. In particular, Home's payoff is

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

where CS is consumer surplus, PS is producer surplus and R is tariff revenue, and θ captures the political importance of the group of domestic producers. This government objective function is in the spirit of Baldwin (1987) and Grossman and Helpman (1994).

For simplicity we abstract from political economy considerations in the Foreign country, and assume the Foreign government maximizes national welfare, which in this setting is just the sum of consumers and producer surplus:

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

Allowing for an extra weight attached to producer surplus in this country would not affect our main results.

To simplify the exposition we assume that the demand and supply functions are linear in both countries. For future reference, note that $\omega^*(T)$ is decreasing and convex in T . Intuitively, the reason is that increasing the tariff T reduces trade volume, and hence reduces the incidence of further increases in T . On the other hand, note that $\omega(T, \theta)$ is concave in T provided θ is not too high: the reason is that $CS(T) + R(T)$ is concave but $PS(T)$ is convex. If θ is so high that $\omega(T, \theta)$ is convex, the unilaterally optimal tariff is prohibitive; we assume that θ does not exceed such threshold.

We let the joint-payoff maximizing tariff be

$$T^{fb}(\theta) \equiv \arg \max_T [\omega(T, \theta) + \omega^*(T)]$$

We will refer to this as the “first best” policy, and note that it is increasing in θ . Also, we let the unilaterally optimal tariff be $T^N(\theta) \equiv \arg \max_T \omega(T, \theta)$. Clearly, $T^N(\theta)$ is increasing in θ and weakly higher than $T^{fb}(\theta)$.

We now characterize the government Pareto frontier, that is the locus of feasible government payoffs (ω, ω^*) . In the absence of government-to-government transfers, the frontier is concave for any θ . Note that the frontier has a peak at the Nash tariff T^N (point N) and has slope equal to -1 at the first best tariff T^{fb} (point FB), as depicted in Figure 1. We label this the

“no-transfer frontier”. Note that the no-transfer frontier becomes less concave as θ increases (since this increases the weight of the convex component in Home’s payoff function, PS), and becomes linear as θ reaches the “prohibitive” threshold.

Now suppose that (costless) international transfers can be used. Then clearly the Pareto frontier is linear with slope -1 and tangent to the no-transfer frontier at the FB point (see Figure 1). In our basic model we assume that, if governments engage in negotiations, they can use efficient transfers, thus we label this the “negotiation frontier.”

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

We now describe the informational structure, the institutional environment and the role of the court (DSB). We consider the simplest possible environment in which the DSB plays an active role. In particular, the role of the court will be to “complete” an incomplete contract.

The political parameter θ (“state of the world”) is ex-ante uncertain, and distributed according to some distribution with support $(\theta_{\min}, \theta_{\max})$. The realization of θ is observed by governments but is not verifiable (not observed by the court), so governments cannot write a complete contingent contract. To simplify the model, we take the incompleteness of the contract to an extreme, assuming it does not specify the policy T at all (discretion). However, the court is endowed with the authority to “fill the gap” of this contract ex-post. This institutional setting can be interpreted more broadly as one in which the contract imposes vague obligations and the court can interpret the court ex-post.³

More specifically, the DSB can observe a noisy signal of T^{fb} , given by $T^{dsb} = T^{fb} + \varepsilon$, where

³In earlier work (Maggi and Staiger, 2011) we have examined the optimal design of the role of the DSB when a complete contingent contract cannot be written, and argued that under some conditions the optimal institution entails a contract that leaves discretion on trade policy and a court that plays a “gap filling” role ex-post. Another institution that may be optimal in that setting is one where the contract is vague and the court plays an “interpretation” role as ex-post. As we argue in that paper, these two institutional forms have similar features and implications in many respects. An alternative possibility in that model is to let the court play only an enforcement role (non-activist court); this can be optimal if the accuracy of the court’s information is low. Thus an implicit assumption in the present model is that the accuracy of the court’s information is not too low, so that the optimal institution entails an activist court.

ε is a white noise with mean zero and variance σ^2 .⁴ If invoked, the DSB issues a (perfectly enforceable) ruling to maximize the governments' expected joint payoff conditional on its noisy information.⁵ Given our assumptions, the DSB ruling will prescribe the tariff level T^{dsb} .⁶ We also assume that, if governments go to court, they incur a symmetric litigation cost C^L .

Governments can avoid court intervention by bargaining. In particular, after observing the realization of θ , they can bargain over the policy T and a transfer, with a disagreement point given by the court ruling. That is, governments bargain “in the shadow of the law.”

Throughout the paper, we say that there is a “dispute” if governments engage in bargaining. 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 and settle outside the institutional framework, or in other words through informal (rather than formal) negotiations. Our model can be interpreted as applying to both formal and informal negotiations.

We can now describe the full timing of the game:

1. After θ is realized, Home chooses T ;
2. Foreign acquiesces or initiates a dispute;
3. If a dispute is initiated, governments negotiate over policy T and a transfer;
4. If the governments disagree, they each incur the litigation cost C^L , the DSB is invoked and a ruling is triggered

We assume the simplest possible symmetric bargaining protocol: each government gets to make a take-or-leave offer with probability 1/2.

Government negotiations are subject to transaction costs. In particular, we assume an “iceberg” negotiation cost: a fraction $1 - \kappa$ of the bargaining surplus “melts” away. Formally, if

⁴We could assume that the court observes a noisy signal of θ rather than a noisy signal of T^{fb} , at the cost of a slightly more complicated analysis. Note also that the assumption that ε is independent of θ is somewhat restrictive, because it implies that the DSB can always make mistakes in either direction, so if the true T^{fb} is close to 0 (free trade), it may prescribe a negative level of T (import subsidy). Again, this feature could be made more realistic at the cost of a more complicated analysis.

⁵The assumption that the DSB attempts to maximize the governments' joint payoff seems a natural one in this setting. The idea is that governments design the institution at some ex-ante stage and endow the court with a certain objective function. Given that international transfers are available, it is natural to suppose that such objective function is the governments' joint payoff.

⁶Alternatively we could allow the DSB to impose a maximum level of T , e.g. a tariff cap.

ω_i^{ND} is government i 's net disagreement payoff (i.e. net of litigation costs) and ω_i^B its bargaining payoff absent negotiation costs, its payoff gain from the bargain is $\kappa (\omega_i^B - \omega_i^{ND})$ with $\kappa \in (0, 1)$.

The reason we need transaction costs in our model is plain: if there were no transaction costs, the disagreement point for government negotiations would be irrelevant, and hence any institutional or contractual arrangement would be irrelevant. Later on we will consider another type of transaction cost, which is the presence of transfer costs.

Finally, we assume a veil of ignorance: from an ex-ante perspective (before θ is realized) each government is equally likely to be the importer or the exporter, so is equally likely to be the claimant 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 claimant or that of defendant. This assumption will play an important role in the dynamic setting analyzed below.

We are now ready to characterize the equilibrium outcome of the static model. We focus on subgame perfect equilibria of the game described above. We go by backward induction and start with the dispute subgame (stage 3).

Given that the Pareto frontier is concave, the disagreement point for the negotiation is below the frontier, and in particular it is Southeast of the FB point. This is 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, it is direct to verify that the expected disagreement payoffs are given by

$$\begin{aligned}\omega^D &= \omega(T^{FB}, \theta) + \omega_{TT} \cdot \sigma^2, \\ \omega^{*D} &= \omega^*(T^{FB}) + \omega_{TT}^* \cdot \sigma^2\end{aligned}$$

(with $\omega_{TT} < 0$, $\omega_{TT}^* > 0$). These expressions make clear that increasing the DSB noise σ worsens Home's threat point and improves Foreign's threat point.

Taking into account litigation costs, we can write the net disagreement payoffs as

$$\begin{aligned}\omega^{ND} &= \omega^D - C^L, \\ \omega^{*ND} &= \omega^{*D} - C^L.\end{aligned}$$

These are the payoffs that governments get if they disagree, since in this case a ruling is triggered and governments pay litigation costs. This payoff pair $(\omega^{ND}, \omega^{*ND})$ is labeled ND in Figure 1.

Since point ND is always below the negotiation frontier, it is clear that there can never be a ruling in the static setting.

Next we examine when a dispute occurs in the static setting. Given our bargaining protocol and the iceberg negotiation cost, it is easy to derive the net bargaining payoffs in case a dispute is initiated. Letting Ω denote joint payoff, and omitting the argument θ , we can write net bargaining payoffs as

$$\begin{aligned}\omega^{Bnet} &= \omega^{ND} + \kappa \cdot \frac{\Omega^{FB} - \Omega^{ND}}{2} \\ &= \omega^{FB} + \sigma^2 \cdot \left[\frac{\kappa}{2} \omega_{TT}^* + \left(1 - \frac{\kappa}{2}\right) \omega_{TT} \right] - (1 - \kappa) C^L, \\ \omega^{*Bnet} &= \omega^{*ND} + \kappa \cdot \frac{\Omega^{FB} - \Omega^{ND}}{2} \\ &= \omega^{*FB} + \sigma^2 \cdot \left[\left(1 - \frac{\kappa}{2}\right) \omega_{TT}^* + \frac{\kappa}{2} \omega_{TT} \right] - (1 - \kappa) C^L.\end{aligned}$$

Graphically, the net bargaining payoff point (labeled $Bnet$ in Figure 1) is somewhere between the ND point and its 45^0 projection onto the negotiation frontier.

Moving backwards to stages 1 and 2, it is easy to argue that: (i) if the $Bnet$ point is below the no-transfer frontier, Home chooses a tariff T such that Foreign is indifferent between complaining and not, so there is no dispute.⁷ Graphically, the outcome in this case is the vertical projection of $Bnet$ onto the no-transfer frontier, that is point B^0 in Figure 1. (ii) If point $Bnet$ is above the no-transfer frontier, Home will trigger a dispute, by choosing a level of T that Foreign will complain about, and the equilibrium payoffs are given by point $Bnet$.

Next we observe that, given θ , a dispute occurs if σ is above some threshold level. To see this note that, if σ increases, the $Bnet$ point moves Southeast in a linear way, starting from a point below the no-transfer frontier, and since this frontier is concave, it can only cross the frontier from the left. This is true for any θ , so we can conclude that the probability of a dispute increases with σ .

Proposition 1. *In the static setting: (i) there is never a DSB ruling; (ii) the likelihood of a dispute is increasing in the DSB noise σ .*

The economic intuition for the effect of σ on the likelihood of a dispute is that, when σ is high, the disagreement point is bad and skewed in favor of the exporter, so the marginal benefit of using side payments – and hence of initiating a dispute – is high.

⁷We are implicitly assuming that Foreign does not complain in case of indifference.

We can characterize how the governments' payoffs change with σ . This can be done easily in a graphical way. Figure 2 depicts how the equilibrium outcome varies with σ for a given θ . As σ increases from zero, initially there is no dispute and the equilibrium payoff point (which is the vertical projection of $Bnet$ onto the no-transfer frontier) travels down along the no-transfer frontier, and after σ crosses a threshold, the outcome is a dispute and the payoff point moves Southeast linearly with slope steeper than -1. It is then clear that increasing σ leads to (i) an increase in the exporter's equilibrium payoff; (ii) a decrease in the importer's equilibrium payoff, and (iii) a decrease in the joint equilibrium payoff. Observe also that the joint equilibrium payoff is weakly concave in σ . The following remark highlights the impact of σ on the joint payoff, since this will play an important role in the dynamic analysis to follow.

Remark 1. *In the static setting, the equilibrium joint payoff is decreasing and weakly concave in the DSB noise σ .*

3. Learning By Ruling

The WTO is a relatively young international institution characterized by a fairly sophisticated judicial system. The adjudication process that such judicial system is designed to conduct is a complex job, and there is little doubt that the actors involved in this system have much to learn in many dimensions, especially in the early stages of the institution. We believe this kind of “institutional learning” is a phenomenon of first-order importance, and we are interested in exploring its implications for the dynamics of disputes and rulings.

One could consider different types of institutional learning. A first possibility is that the court can learn from its past experience. This is the notion that we refer to as *learning by ruling*. There are several mechanisms by which a court can learn from experience. One is that the court may become more accurate in conducting investigations and figuring out the economic and political costs/benefits of trade policies (and of domestic policies that have impacts on trade). This may involve learning to use and interpret data, or to choose the right experts, or just learning to use rigorous economic reasoning. We can think of this as “methodological” learning, or in other words, “learning by doing” in investigating and adjudicating.

But we can also think of a “factual” type of learning by the court: for example, by repeatedly studying the policies of a certain country (say, China) or in a certain issue area (say, health and safety), the court may gain knowledge about persistent aspects of that country's policy

environment or of that issue area (the “state of the world”). In our model we will consider only methodological learning, by assuming that the court’s information becomes less noisy (σ goes down) as rulings accumulate, and we shut down any factual learning by assuming that the state of the world is *iid* over time.

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

Another type of institutional learning that is probably quite relevant in reality is *governments’ learning about the court*. By this we mean that, as governments litigate repeatedly in court, they learn how the court operates and adjudicates cases, and therefore they learn to better predict the outcome of a ruling. Intuitively, some of the implications of this type of learning should be similar as those of learning by ruling, to the extent that both types of learning increase the predictability of future rulings. However there may also be subtle differences in implications, due to the fact that governments’ learning about the court does not *per se* increase the quality of court decisions. It would be interesting to study more formally this latter type of learning, but for the purposes of this paper we will focus more narrowly on the case of learning by ruling.

Finally, a dynamic mechanism that has a similar flavor but is quite distinct from learning is *legal precedent*. While it would be interesting to explore the implications of legal precedent, this is beyond the scope of this paper.⁸ We will come back to the notion of legal precedent, however, in the empirical section, because in principle this mechanism could explain some of the dynamic patterns observed in our data.

3.1. The two-period setting

We start by considering two periods, $t = 1, 2$. In each period, the same game as described in the static setting takes place. The state of the world θ is *iid*, so learning-by-ruling will be the

⁸In Maggi and Staiger (2011) we explore the implications of legal precedent for trade disputes, in a setting without any learning.

only source of dynamics. The governments' common discount factor is $\delta \in (0, \infty)$.⁹

We model learning-by-ruling in a similar fashion as in the typical models of learning-by-doing for firms, where increasing a firm's current output increases its future productivity: we assume that adjudicating one more case today increases the accuracy of the court tomorrow. More specifically, if there has been a ruling at $t = 1$, the DSB noise (σ) at $t = 2$ is lower.

This bare-bones two period model will allow us to make a couple of key points, but later in this section we consider a slightly richer version of the model to examine how the current likelihood of disputes and rulings depends on cumulative rulings.

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

Going by backward induction, at $t = 2$ the outcome is the same as in the static setting analyzed above, and hence there can be no rulings, but the situation is different at $t = 1$, because there is an investment value in going to court due to the learning effect. Recall from Remark 1 that, in the static setting, decreasing σ increases the equilibrium joint payoff. Thus, given the veil of ignorance, going to court at $t = 1$ implies a common future payoff gain, which we label Δ .

It is worth emphasizing that increasing future court accuracy benefits governments through an indirect off-equilibrium mechanism, because at $t = 2$ there is no court activity in equilibrium. Making the court more accurate improves the disagreement point in case of dispute at $t = 2$, and the disagreement point matters because of the negotiation cost (κ). Moreover, if no dispute takes place at $t = 2$, improving the would-be dispute outcome leads to a more efficient policy choice by Home (an off-off-equilibrium effect).¹⁰

In what follows we characterize the equilibrium outcome at $t = 1$. We do not use a time index, as this should not cause confusion.

At $t = 1$, the disagreement payoffs are $(\omega^{ND} + \delta\Delta, \omega^{*ND} + \delta\Delta)$. Graphically, we label the corresponding payoff point $ND + \delta\Delta$. In terms of Figure 1, this point lies somewhere on the 45° line emanating from point ND , and in general may be below or above the negotiation frontier. If point $ND + \delta\Delta$ is above the negotiation frontier, then a dispute will end in ruling; and going

⁹Since we have only two periods, it is natural to allow δ to be higher than one, as the second period can be thought of as condensing a potentially long future.

¹⁰If we had a richer model with more than one period ahead of $t = 1$, the payoff gain Δ would include also a direct effect of increasing court accuracy in case a ruling occurs in equilibrium.

backwards, Home chooses a policy T that triggers a complaint by Foreign.

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

At a broad level, it is worth highlighting that the possibility of equilibrium rulings depends not only on the presence of learning by ruling, but also on the presence of “large” players who interact repeatedly in court, as well as on the implicit assumption that governments internalize the future gain from going to court. In our basic model, since there are only two countries, governments fully internalize this gain, but in a multi-country setting there would be international externalities from using the court: when a pair of countries goes to court today, there will be a benefit for other countries that may use the court in the future, to the extent that the scope of court learning is not narrowly confined to today’s disputant countries, and as a consequence, there is potential for a free-rider problem in the use of the judicial system. We will come back to the question of possible learning spillovers later in the paper.

We can characterize how the occurrence of a ruling at $t = 1$ depends on the realization of the political economy shock θ at $t = 1$. Intuitively, a higher realization of θ implies a less concave no-transfer frontier, and hence the ND point is closer to the negotiation frontier at $t = 1$; or equivalently, as θ increases the joint payoff at ND is closer to the first best joint payoff. On the other hand, Δ is independent of the realization of θ at $t = 1$, because θ is *iid* over time and Δ is the *expected* future payoff gain from going to court at $t = 1$. Thus, as θ increases, the $ND + \delta\Delta$ point can only cross the negotiation frontier from below, and as a consequence a ruling occurs if θ is above a certain threshold. Formally, the difference in joint payoff between the first best and the ND point is $\Omega^{FB} - \Omega^{ND}$, and the impact of θ on this difference is $\frac{\partial}{\partial\theta}(\Omega^{FB} - \Omega^{ND}) = \omega_{\theta}^{FB} - (\omega_{\theta}^{FB} + \sigma^2\omega_{TT\theta}) = -\sigma^2\omega_{TT\theta} < 0$, where the sign follows from the fact that ω_{TT} increases with θ .¹¹ This leads to

Remark 2. *There is a threshold level $\hat{\theta} \leq \theta_{\max}$ such that a ruling occurs at $t = 1$ for $\theta \in (\hat{\theta}, \theta_{\max})$.*

Thus the model predicts that rulings should be more likely to occur, other things equal, when

¹¹We note here that the impact of θ on the occurrence of a dispute is more ambiguous, because θ affects not only the ND point but also the position of the no-transfer frontier.

the importing government faces strong political pressures from domestic import-competing producers.

What is the outcome at $t = 1$ if point $ND + \delta\Delta$ is below the negotiation frontier, so there is no ruling? In this case, let B_{Δ}^{net} denote the net bargaining payoffs at $t = 1$ given disagreement point $ND + \delta\Delta$.¹² Point B_{Δ}^{net} lies between the $ND + \delta\Delta$ point and its 45⁰ projection onto the negotiation frontier (with the distance determined by the negotiation cost). Then it is easy to argue that the outcome will be a dispute with settlement if B_{Δ}^{net} is above the no-transfer frontier, and no dispute if B_{Δ}^{net} is below the no-transfer frontier.

Note that, if learning is more important, in the sense that a ruling at $t = 1$ reduces σ by a bigger amount, not only the likelihood of a ruling but also the likelihood of a dispute is higher at $t = 1$. That a ruling is more likely when learning is more important is obvious, given our discussion thus far. The reason why also a dispute becomes more likely can be understood graphically as follows. Suppose the B_{Δ}^{net} point is on the no-transfer frontier, so governments are at the margin between having a dispute and not: then, if Δ increases, the B_{Δ}^{net} point will move above the no-transfer frontier and there will be a dispute.

3.2. Impact of past rulings on current outcomes

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

Continue assuming two periods, $t = 1, 2$, but now suppose there is an initial stock of rulings x , inherited from a “past” period $t = 0$. To examine how past rulings affect current outcomes, we can focus on the equilibrium outcome at $t = 1$ conditional on x .¹³

Learning by ruling in this setting is represented by a decreasing function $\sigma(x)$, which we assume convex with $\lim_{x \rightarrow \infty} \sigma(x) > 0$. The interpretation of these assumptions is that learning proceeds at a diminishing rate and there is a “baseline” level of noise that cannot be removed by learning.

¹²More explicitly, the B_{Δ}^{net} payoffs are the same as the B^{net} payoffs derived above in the static setting, except that the disagreement payoffs are $(\omega^D - C^L + \delta\Delta, \omega^{*D} - C^L + \delta\Delta)$ instead of $(\omega^D - C^L, \omega^{*D} - C^L)$.

¹³It would be conceptually easy to endogenize the occurrence of a ruling at $t = 0$, but this would not add much to the question of how cumulative rulings affect current outcomes, because in a three-period setting this question is meaningful only from the perspective of the central period ($t = 1$), since at $t = 0$ there is no past and at $t = 2$ there cannot be rulings.

We are now ready to study how an increase in x affects the likelihood of rulings and disputes at $t = 1$. We first focus on the likelihood of a ruling.

Recall that a ruling occurs at $t = 1$ if and only if the $ND + \delta\Delta$ point is above the negotiation frontier, so we can write

$$\Pr(\text{ruling}) = \Pr(g < \delta\Delta),$$

where g is the distance between the ND point and the negotiation frontier along a 45° line, and

$$\Delta = E_\theta[\Omega^{t=2}(\sigma(x+1), \theta) - \Omega^{t=2}(\sigma(x), \theta)],$$

where $\Omega^{t=2}(x, \theta)$ is the equilibrium joint payoff at $t = 2$ given x and θ .

First note that g decreases with x , because as σ decreases, the ND point gets closer to the negotiation frontier.

We next argue that Δ also decreases with x . Ignoring the integer nature of x , we approximate $\partial\Delta/\partial x$ as

$$\frac{\partial\Delta}{\partial x} = \frac{\partial E_\theta[\Omega^{t=2}(x+1, \theta) - \Omega^{t=2}(x, \theta)]}{\partial x} \approx E_\theta \Omega_{xx}^{t=2}(\cdot) = E_\theta[\Omega_{\sigma\sigma}^{t=2}(\cdot)(\sigma'(x))^2 + \Omega_{\sigma}^{t=2}(\cdot)\sigma''(x)].$$

Remark 1 above implies $\Omega_{\sigma}^{t=2} < 0$ and $\Omega_{\sigma\sigma}^{t=2} \leq 0$, and since the learning curve is convex, then $\frac{\partial\Delta}{\partial x} < 0$.¹⁴

Thus increasing x reduces the future gain from going to court (Δ) but decreases today's inefficiency from going to court (g). Clearly, if governments care enough about the future (i.e. δ is sufficiently large), the dynamic effect dominates the static effect and the probability of a ruling decreases with x . But notice also that, even if δ is small, the probability of a ruling is decreasing in x for x sufficiently large, because when learning vanishes Δ goes to zero, while g does not go to zero. We can thus state:

Proposition 2. *At $t = 1$ the likelihood of a ruling is decreasing in x for x sufficiently large, and it is globally decreasing in x if the discount factor δ is high enough.*

Proposition 2 suggests that if δ is sufficiently high or x is sufficiently large, the frequency of rulings should decline. But note that even if δ and x are small the model does not necessarily predict that the frequency of rulings will increase. In our analysis above we have taken x as given, but if x were determined endogenously, the following observation would immediately

¹⁴It is easy to extend the argument to take into account the integer nature of x .

apply: if δ is sufficiently small then $(\delta\Delta - g)|_{x=0} < 0$ for all θ , so rulings would never get started in equilibrium. It is not obvious whether in this model it is possible that the likelihood of rulings increases with x **and** rulings occur in equilibrium.

Before moving on, it is worth emphasizing an important implication of the model: the frequency of DSB use is not a good measure of the effectiveness of the institution. According to our theory, a declining frequency of rulings does not imply that the institution is getting worse over time, in fact it is a symptom of beneficial learning by the institution. But note that this statement concerns the *change* in ruling frequency over time. A higher *level* of the ruling frequency, on the other hand, is associated with higher court efficiency according to our model: if we decrease the noise σ for a given level of x , the disagreement point is more likely to be above the Pareto frontier, so other things equal it increases the probability of a ruling.

Next we consider the impact of cumulative rulings on the probability of a dispute at $t = 1$.

It is easy to argue that if δ is sufficiently high the likelihood of a dispute is globally decreasing in x . Intuitively, suppose governments are at the margin between disputing and not; this means that the B_{Δ}^{net} point defined above is on the no-transfer frontier. Increasing x has two effects on B_{Δ}^{net} : a “static” effect, that is a shifts in the ND point (towards Northwest) and a “dynamic” effect (a decrease in Δ). If δ is high enough, the dynamic effect dominates, so the B_{Δ}^{net} point dips below the no-transfer frontier, where there is no dispute. We can thus state:

Proposition 3. *At $t = 1$ the likelihood of a dispute is globally decreasing in x if δ is high enough.*

It is interesting to note that, unlike the case of rulings, the likelihood of a dispute may not be decreasing in x for x very large. Suppose δ is small, so the effect of x on the ND point dominates the effect on Δ , and suppose that κ is small (high negotiation cost). Consider the margin of indifference between dispute and no-dispute, where the B_{Δ}^{net} point is on the no-transfer frontier: if σ decreases, it is not hard to show that the B_{Δ}^{net} point will move above the no-transfer frontier, where there is a dispute, even if x is large.¹⁵ Note that this logic does not

¹⁵To see why, suppose $\kappa = 0$, so that the B_{Δ}^{net} point coincides with the disagreement point $ND + \delta\Delta$. Let us show that as σ changes the ND point moves with a slope that is steeper than the no-transfer frontier. It suffices to show that the slope of the no-transfer frontier at the end point $T = T^{\min}$ is flatter than the slope of the $ND(\sigma)$ line. The slope of the $ND(\sigma)$ line is $\frac{\omega_{TT}}{\omega_T^*}$, while the slope of the no-transfer frontier at T^{\min} is $\frac{\omega_T(T^{\min})}{\omega_T^*(T^{\min})}$. Note that $\omega_T(T^{\min}) = \omega_T(T^N) - \omega_{TT} \cdot (T^N - T^{\min}) = -\omega_{TT} \cdot (T^N - T^{\min})$, and $\omega_T^*(T^{\min}) = \omega_T^*(T^{prohib}) - \omega_{TT}^* \cdot (T^{prohib} - T^{\min}) = -\omega_{TT}^* \cdot (T^{prohib} - T^{\min})$ [where we are using the facts that (i) $\omega_T(T^N) = 0$;

apply to the likelihood of a ruling, because if x is large governments cannot possibly be at the margin between ruling and no ruling.

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

Remark 3. *At $t = 1$ the likelihood of settlement conditional on a dispute may go up or down with x , even if δ is high.*

The intuition for this result is that the effect of an increase in x on the ruling margin (which occurs when the $ND + \delta\Delta$ point is on the negotiation frontier) may be stronger or weaker than the effect of x on the dispute margin (which occurs when the B_{Δ}^{net} point is on the no-transfer frontier), depending on the probability distribution of θ , and for this reason the ratio $\Pr(\text{ruling})/\Pr(\text{dispute})$ can either increase or decrease.¹⁶

Remark 3 will be relevant for our empirical work, because it will motivate us to look at how cumulative rulings impact the unconditional likelihood of a ruling and of a dispute, as opposed to the conditional likelihood of settlement.

3.3. Scope of learning

Thus far we have assumed that issuing a ruling today increases the court's future accuracy regardless of which country is the defendant in the future. But one could consider more narrow forms of learning. For example, learning might be "directed-dyad specific," meaning that a ruling where country i is the claimant and j the defendant increases the court's future accuracy only for disputes where again i is the claimant and j the defendant, but not if roles are reversed. This could be the case if learning is specific to the defendant country (e.g. because the court learns about the political economy of Home's import-competing sector) or to the complainant country (e.g. because the court learns about the political economy of Foreign's export sector).

(ii) $\omega_T^*(T^{prohib}) = 0$ and (iii) for a linear function f , we can write $f(x_1) = f(x_0) + f' \cdot (x_1 - x_0)$. We can conclude that $\frac{\omega_T(T^{\min})}{\omega_T^*(T^{\min})} = \frac{\omega_{TT} \cdot (T^N - T^{\min})}{\omega_{TT}^* \cdot (T^{prohib} - T^{\min})}$, and since $T^N \leq T^{prohib}$, the claim follows.

¹⁶A more formal proof is the following. Consider a realization $\theta = \theta'$ such that $ND + \delta\Delta$ is just above the negotiation frontier. As x increases, Δ decreases and hence for $\theta = \theta'$ the outcome switches from ruling to settlement. So if the probability mass of θ is concentrated around θ' , then $\Pr(\text{settlement})/\Pr(\text{dispute})$ goes up. On the other hand, suppose there is zero probability mass for a small neighborhood around θ' . 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.

In this section we will extend our analysis by allowing for a whole range of possibilities that include the general form of learning considered in the previous sections and the case of directed-dyad specific learning just described.

Formally, it is convenient to label θ_{ij} the political-economy shock if country j is the importer and country i the exporter,¹⁷ and σ_{ij} the noise in the DSB signal of θ_{ij} . Thus σ_{ij} is interpreted as the court's accuracy in ruling on disputes brought by country i against country j . We assume that σ_{ij} is a function of the composite experience variable $X_{ij} = \beta_1 x_{ij} + \beta_2 x_{ji}$, where x_{ij} is the number of past ruling where j was the defendant and x_{ji} is the number of past ruling where i was the defendant. We thus write the court's learning curve as $\sigma_{ij} = \sigma(X_{ij}) = \sigma(\beta_1 x_{ij} + \beta_2 x_{ji})$, where $\sigma(\cdot)$ is decreasing and convex. This formulation includes the special cases of general learning and directed-dyad specific learning: learning is general if $\beta_1 = \beta_2 > 0$, and learning is directed-dyad specific if $\beta_1 > \beta_2 = 0$. More generally, it is reasonable to assume $\beta_1 \geq \beta_2 \geq 0$: for a dispute where j is the defendant, the court learns weakly more from past rulings where the defendant was j than from past rulings where the defendant was i .

We start by examining how the likelihood of current rulings depends on past rulings. We will show that, if δ is sufficiently high, the likelihood of a ruling where country j is the defendant is decreasing in both x_{ij} and x_{ji} , with the effect of x_{ij} weakly stronger (when starting from equal values of x_{ij} and x_{ji}).

The first observation is that, in our static setting, increasing court accuracy increases the defendant's payoff, decreases the claimant's payoff and increases the joint payoff (recall Remark 1 and the discussion preceding it).

Next consider the dynamic setting. Suppose that at $t = 1$ country j is the importer/defendant, and consider the future impacts of a ruling. Given the veil of ignorance, at $t = 2$ there are two possibilities: (i) with probability 1/2 country j is again the defendant, in which case the relevant court experience X_{ij} increases by an amount β_1 ; (ii) with probability 1/2 country i is the defendant, in which case the relevant court experience X_{ji} increases by an amount $\beta_2 \leq \beta_1$.

In light of the above considerations, what are the future payoff changes implied by a ruling at $t = 1$? Let us denote such future payoff changes for the claimant (i) and for the defendant

¹⁷The reason we use a double index, rather than using simply θ_j to denote the political-economy shock in importing country j will become clear later.

(j) respectively as Δ_i and Δ_j . We can write

$$\begin{aligned}\Delta_i &= \frac{1}{2}E_\theta[\omega^i(\sigma(X_{ij} + \beta_1), \theta_{ij}) - \omega^i(\sigma(X_{ij}), \theta_{ij})] + \frac{1}{2}E_\theta[\omega^j(\sigma(X_{ji} + \beta_2), \theta_{ji}) - \omega^j(\sigma(X_{ji}), \theta_{ji})] \text{ and} \\ \Delta_j &= \frac{1}{2}E_\theta[\omega^j(\sigma(X_{ij} + \beta_1), \theta_{ij}) - \omega^j(\sigma(X_{ij}), \theta_{ij})] + \frac{1}{2}E_\theta[\omega^i(\sigma(X_{ji} + \beta_2), \theta_{ji}) - \omega^i(\sigma(X_{ji}), \theta_{ji})]\end{aligned}$$

where $\omega^i(\sigma(\cdot), \cdot)$ and $\omega^j(\sigma(\cdot), \cdot)$ denote respectively the claimant's and defendant's equilibrium payoffs when the DSB noise is σ .

We know from the discussion above that, in a given period, a decrease in the DSB noise benefits the defendant more than it hurts the claimant. It is then straightforward to show that (i) Δ_i and Δ_j may be positive or negative depending on the values of β_1 and β_2 , and (ii) $\Delta_i + \Delta_j > 0$ regardless of the values of β_1 and β_2 .

Focus first on how a ruling affects the future payoff of today's claimant, that is Δ_i . To see why the sign of Δ_i is ambiguous, note that if learning is general ($\beta_1 = \beta_2 > 0$) then $\Delta_i > 0$, as we argued in the previous section, but if learning is directed-dyad specific ($\beta_1 > \beta_2 = 0$) then $\Delta_i < 0$, because in this case the expression above yields $\Delta_i = \frac{1}{2}E_\theta[\omega^i(\sigma(X_{ij} + \beta_1), \theta_{ij}) - \omega^i(\sigma(X_{ij}), \theta_{ij})]$, which is negative because in a given period a decrease in σ hurts the claimant.

Next focus on how a ruling affects the future payoff of today's defendant, that is Δ_j . The reason why Δ_j is ambiguous is that the initial values of X_{ij} and X_{ji} may be different, so the slope of the learning curve may be different depending on who is the defendant at $t = 2$. Suppose X_{ij} is large, so that the learning for cases where j is defendant is close to exhausted; then the impact of a ruling on j 's future payoff will be close to zero if j is again the defendant at $t = 2$, and negative if j is the claimant. On the other hand, it is easy to see that if $X_{ij} = X_{ji}$, or if learning is directed-dyad specific ($\beta_1 > \beta_2 = 0$), or if learning is general ($\beta_1 = \beta_2 > 0$), then $\Delta_j > 0$.

Finally, it is intuitive that $\Delta_i + \Delta_j > 0$ regardless of β_1 and β_2 . This is because, whichever country is the defendant, a ruling today will lead to an increase in the relevant stock of court experience and hence to an increase in DSB accuracy, and we know that this increases the countries' joint payoff.

Having established that a ruling at $t = 1$ increases future joint payoff, it is easy to see that a ruling at $t = 1$ will occur if the learning effect is strong enough. Denoting $\mathbf{\Delta}$ the vector (Δ_i, Δ_j) , and recalling with reference with Figure 2 that the disagreement point at $t = 1$ is given by $ND + \delta\mathbf{\Delta}$, this point lies Northwest of point ND at an angle steeper than -1 . Thus if the joint gain from learning $\Delta_i + \Delta_j$ is large enough, $ND + \delta\mathbf{\Delta}$ lies above the negotiation

frontier and hence a ruling occurs at $t = 1$.

How does an increase in x_{ij} or x_{ji} affect the probability of a ruling at $t = 1$? Just as in the previous section, there is a tradeoff between a static effect and a dynamic effect. In particular, increasing either x_{ij} or x_{ji} increases (at least weakly) the experience variables X_{ij} and X_{ji} , and hence: (i) this decreases (at least weakly) the current DSB noise, whichever country is the defendant, implying that the ND point in Figure 2 moves Northwest and gets closer to the negotiation frontier; and (ii) by an analogous argument as in the previous section, this reduces the joint future gain from a ruling, $\Delta_i + \Delta_j$, thus making it less likely that the point $ND + \delta\Delta$ lies above the negotiation frontier. Other things equal, the first effect pushes up the probability of a ruling, while the second effect pushes down the probability of a ruling. We can then conclude that if δ is large enough the probability of a ruling is decreasing in x_{ij} and x_{ji} .

Furthermore, using the assumption $0 \leq \beta_2 \leq \beta_1$ one can easily show that, when starting from equal values of x_{ij} and x_{ji} , the impact of an increase in x_{ij} on the probability of a ruling where j is the defendant is (weakly) stronger than its impact on the probability of a ruling where j is the claimant.

A natural question is whether also the likelihood of a dispute is decreasing in x_{ij} and x_{ji} . The answer is yes if learning is sufficiently general (i.e. β_1 is sufficiently close to β_2) and governments are sufficiently patient, but we do not have a more general result. We argued in the previous section that, if learning is fully general ($\beta_1 = \beta_2$) and δ is high enough, the likelihood of a dispute is decreasing in x . With a continuity argument one can show that if β_1 is sufficiently close to β_2 and δ is high enough, the likelihood of a dispute is decreasing in x_{ij} and x_{ji} . But if β_1 is not close to β_2 this result is no longer guaranteed. To see why, suppose $\beta_1 > \beta_2 = 0$ (directed-dyad specific learning). With reference to Figure 2, suppose the B_{Δ}^{net} point is right at the dispute margin, i.e. on the no-transfer frontier, and consider an increase in x_{ij} or x_{ji} . Suppose δ is large, so that we can ignore the static effect and focus on the dynamic effect. As we argued above, in this case $\Delta_i < 0$, $\Delta_j > 0$ and $\Delta_i + \Delta_j > 0$, and furthermore, increasing x_{ij} or x_{ji} reduces $\Delta_i + \Delta_j$. It can also be shown that increasing x_{ij} or x_{ji} reduces both Δ_i and Δ_j in absolute value. This implies that the $ND + \delta\Delta$ point moves Southwest with slope steeper than -1, and so does the B_{Δ}^{net} point. This could lead the B_{Δ}^{net} point to dip below the no-transfer frontier or to rise above it, thus the impact on the likelihood of a dispute is ambiguous.

4. Binary policy

A second reason for equilibrium rulings, which is static in nature, arises if policy is discrete and transfers are costly. This will provide a “baseline” frequency of rulings which persists even absent learning (or if learning is exhausted), and will deliver a “pivoting” effect on the time path of rulings/disputes as learning becomes more important. As we have argued elsewhere (see Maggi and Staiger, 2014, 2015), costly transfers are a highly relevant case in the context of trade disputes, where compensation almost never takes the form of cash payments, and many trade disputes involve features of regulatory regimes that are arguably discrete in nature.

As before, we consider a single industry where the importing government (H) makes a policy choice and the exporting government (F) can initiate a dispute. But here we consider discrete policies, and to focus sharply on the essential ingredients, we now assume that H makes a binary policy choice $T \in \{FT, P\}$, where FT denotes “Free Trade” and P denotes “Protection.” And we continue to assume the availability of a transfer instrument b which is positive if H makes a transfer to F and negative if F makes a transfer to H , but we now assume that the transfer carries with it a dead-weight-loss $c(b)$ borne by the government making the payment. For tractability, we assume $c(b) = c \cdot |b|$, with $c \in (0, 1)$.

H 's payoff is given by $\omega = v(T) - b - c^+(b)$, where $v(T)$ is the importing government's payoff associated with the policy choice T (and could amount to a politically weighted sum of producer surplus, consumer surplus and tariff revenue as in the previous section), and where $c^+(b) = c(b)$ if $b > 0$ and 0 otherwise. Similarly, F 's payoff is given by $\omega^* = v^*(T) + b - c^-(b)$, where $v^*(T)$ is the exporting government's payoff associated with the (Home) policy choice T , and where $c^-(b) = c(b)$ if $b < 0$ and 0 otherwise.

In this binary policy setting, it is useful to define the importing government's gain from protection: $\gamma \equiv v(P) - v(FT) \geq 0$. Similarly, we may define the exporting government's loss from protection: $\gamma^* \equiv v^*(FT) - v^*(P) \geq 0$. We take the “state of the world” $(\gamma, \gamma^*) \equiv \gamma$ to be uncertain ex-ante, distributed according to $F(\gamma)$, and we assume that F is symmetric: $F(\gamma, \gamma^*) = F(\gamma^*, \gamma)$. And we allow that the joint loss from protection $\gamma^* - \gamma$ can be positive or negative (due, say, to political economy forces or the existence of various potential domestic market failures). We continue to refer to the policy that maximizes joint surplus as the first-best policy. When $\gamma^* - \gamma > 0$ the first-best policy is $T = FT$, and when $\gamma^* - \gamma < 0$ the first-best policy is $T = P$. As we did in the previous section, below we use T^{fb} to denote this first-best

policy choice.

As Figure 3 illustrates, for any given state of the world and with our assumptions on the transfer b , the negotiation frontier is given by the outer envelope of two concave sub-frontiers in (ω, ω^*) space, with one sub-frontier emanating from the point $(v(P), v^*(P))$ labeled P in Figure 3, and the other sub-frontier emanating from the point $(v(FT), v^*(FT))$ labeled FT . Figure 3 illustrates three particular state realizations that together span the possible shapes of the negotiation frontier. In the top left panel, FT is the joint surplus maximizing policy choice, and the joint gains from FT relative to P are sufficiently large that the FT sub-frontier everywhere dominates the P sub-frontier; and the negotiation frontier is therefore concave. It is easy to show that this corresponds to state realizations satisfying $\frac{\gamma}{\gamma^*} < \frac{1}{1+c}$. In the top right panel, P is the joint surplus maximizing policy choice, and the joint gains from P relative to FT are sufficiently large that the P sub-frontier everywhere dominates the FT sub-frontier; and the negotiation frontier is again concave. This corresponds to state realizations satisfying $\frac{\gamma}{\gamma^*} > 1 + c$. Finally, the bottom panel illustrates the case where $\frac{1}{1+c} < \frac{\gamma}{\gamma^*} < 1 + c$ and the joint surplus from P and FT are sufficiently similar that neither sub-frontier dominates everywhere; here the negotiation frontier has a region of convexity.

The realized γ is observed by both governments but is not verifiable, and the policy T is not specified in the contract. As before, we allow the court/DSB, if invoked, to “fill the gap” in the contract. As in the previous section, we assume that each government incurs the litigation cost C^L when the DSB is invoked; and if invoked, we assume that the DSB observes T^{fb} imperfectly.

In the binary policy context, we model the signal technology in a particularly simple way: we assume that the DSB makes a mistake with probability $q \in (0, \frac{1}{2})$.¹⁸ So for example, if $T^{fb} = P$ the DSB ruling is P with probability $1 - q$ and FT with probability q . This assumption is restrictive, because q is independent of γ , but it helps to keep the analysis tractable and transparent. We will later discuss where our results might be qualified with other signal technologies.

Finally, we focus on the case of “general” learning, as in the basic continuous-policy model. We let x again denote the cumulative stock of rulings, and assume that q is decreasing and convex in x . Hence $q(x)$ with $q_x < 0$ and $q_{xx} > 0$ is the DSB learning curve.

The timing is as before: (0) γ is realized and observed by the governments; (1) The importer chooses T ; (2) The exporter acquiesces or initiates a dispute; (3) If a dispute is initiated, the

¹⁸When $\gamma = \gamma^*$ there can be no “mistake,” but we define $q \equiv 1/2$ for this state realization.

governments negotiate over the policy T and a transfer b , subject to the negotiation cost κ ;
(4) If the governments disagree, they each incur the litigation cost C^L , the DSB is invoked and a ruling is triggered. We maintain the bargaining protocol from the previous section, namely, that with equal probability each government gets to make a take-or-leave offer.

We begin by analyzing the static benchmark, and consider first the subgame where a dispute has been initiated, and governments negotiate over the policy T and a transfer b , subject to the negotiation cost κ . When will governments settle at the negotiation stage (stage 3), and when will their negotiations end in disagreement and trigger a DSB ruling?

Recall from the previous section that a ruling will occur if and only if the expected net disagreement point ND lies above the negotiation frontier. In terms of Figure 3, if litigation costs were zero the expected disagreement point would lie on the line segment connecting the points P and FT ; and with positive litigation costs the point ND then lies somewhere below this line segment. It then follows immediately by inspection of the top left and top right panels of Figure 3 that a ruling will never be triggered for state realizations satisfying $\frac{\gamma}{\gamma^*} < \frac{1}{1+c}$ or $\frac{\gamma}{\gamma^*} > 1+c$.

To see when a ruling will be triggered, we may therefore focus on the range of state realizations satisfying $\frac{1}{1+c} < \frac{\gamma}{\gamma^*} < 1+c$, where the negotiation frontier is illustrated in the bottom panel of Figure 3. Figure 4 replicates the negotiation frontier for a state realization in this range, and illustrates the case where the ND point lies above the negotiation frontier and a ruling therefore occurs. To characterize the conditions where rulings occur, we need expressions for the expected net disagreement payoffs and the negotiation frontier over this range of state realizations. The expected net disagreement payoffs for the negotiation are easily derived. For $\gamma > \gamma^*$ within this range, the expected net disagreement payoffs are:

$$\omega^{ND} = v(P) - q\gamma - C^L \text{ and } \omega^{*ND} = v^*(P) + q\gamma^* - C^L. \quad (4.1)$$

For $\gamma < \gamma^*$ within this range, the expected net disagreement payoffs are:

$$\omega^{ND} = v(FT) + q\gamma - C^L \text{ and } \omega^{*ND} = v^*(FT) - q\gamma^* - C^L. \quad (4.2)$$

It is straightforward to derive that, for $1 \leq \frac{\gamma}{\gamma^*} \leq 1+c$, if a dispute occurs, then a ruling is triggered iff

$$(1+c)\gamma^* - \gamma > \frac{(2+c)C^L}{q} \text{ and } (1+c)\gamma - \gamma^* > \frac{(2+c)C^L}{(1-q)} \quad (4.3)$$

while for $\frac{1}{1+c} < \frac{\gamma}{\gamma^*} \leq 1$, if a dispute occurs, then a ruling is triggered iff

$$(1+c)\gamma^* - \gamma > \frac{(2+c)C^L}{(1-q)} \text{ and } (1+c)\gamma - \gamma^* > \frac{(2+c)C^L}{q}. \quad (4.4)$$

The “ruling region” is the set of states characterized by (4.3) and (4.4). Notice that as DSB noise decreases (i.e., as q falls from $\frac{1}{2}$ toward 0), the ruling region shrinks. Hence, increasing DSB accuracy makes a ruling less likely, and $\Pr(\text{ruling}) \rightarrow 0$ as DSB noise vanishes. Also, an increase in C^L clearly decreases the likelihood of a ruling. Finally, notice that as the ruling region is contained in the cone $\gamma/\gamma^* \in (\frac{1}{1+c}, 1+c)$, it follows that DSB rulings can be triggered only if “efficiency stakes” are relatively small.

We now move backwards to stage 2 and examine when a dispute occurs. We let ω^D, ω^{*D} denote the expected equilibrium payoffs in the dispute subgame (net of bargaining and litigation costs). Does F initiate a dispute?

If H has chosen FT , F does not complain. If H has chosen P , F complains iff $\omega^{*D} > v^*(P)$. Backing up, H chooses P if either $\omega^{*D} < v^*(P)$ (so F will acquiesce) or $\omega^{*D} > v^*(P)$ and $\omega^D > v(FT)$ (so F will complain but H is better off in a dispute than under FT and no transfer). Otherwise H will choose FT .

Collecting these points, we have the following: (i) if $\omega^{*D} > v^*(P)$ and $\omega^D < v(FT)$, H chooses FT and there is no dispute; (ii) if $\omega^{*D} < v^*(P)$, H chooses P and there is no dispute; (iii) if $\omega^{*D} > v^*(P)$ and $\omega^D > v(FT)$, H chooses P and there is a dispute. Thus, there is dispute iff $\omega^D > v(FT)$ and $\omega^{*D} > v^*(P)$.

Now let’s consider the state realizations for which a dispute arises. First note that, if γ is in the ruling region, clearly there is a dispute, because in this case $\omega^{ND} > v(FT)$ and $\omega^{*ND} > v^*(P)$. So we can focus on the remaining regions, that is, the case where the ND point is below the negotiation frontier. In this case, there is a dispute iff $\omega^{Bnet} > v(FT)$ and $\omega^{*Bnet} > v^*(P)$; and here governments will settle. Our remaining task is therefore to characterize the state realizations for which settlement occurs, which we refer to as the “settlement region.”

Given our bargaining protocol, the gross expected payoff from the bargain for H is $(\omega^{TOL} + \omega^{ND})/2$, where ω^{TOL} is H ’s gross payoff from the bargain when H makes the take-or-leave offer; and similarly, the gross expected payoff from the bargain for F is $(\omega^{*TOL} + \omega^{*ND})/2$. Taking into account negotiation costs, the net bargaining payoffs are easily shown to be $\omega^{Bnet} = \frac{\kappa}{2}\omega^{TOL} + (1 - \frac{\kappa}{2})\omega^{ND}$ and $\omega^{*Bnet} = \frac{\kappa}{2}\omega^{*TOL} + (1 - \frac{\kappa}{2})\omega^{*ND}$. Given $\gamma > \gamma^*$, the settlement region

is defined by the conditions

$$\frac{\kappa}{2}\omega^{TOL} + (1 - \frac{\kappa}{2})\omega^{ND} > v(FT) \text{ and } \frac{\kappa}{2}\omega^{*TOL} + (1 - \frac{\kappa}{2})\omega^{*ND} > v^*(P),$$

whenever the ND point is below the negotiation frontier (otherwise the take-or-leave payoffs are not defined, and we are in the ruling region). The dispute region of course is the sum of the settlement region and the ruling region. Finally, given the symmetry of the model, in the octant $\gamma < \gamma^*$ the dispute region is the mirror image of the one in the region $\gamma > \gamma^*$.

As a special case, note that if $\kappa = 0$ the dispute region boils down to the region where $\gamma > \frac{C^L}{q}$ and $\gamma^* > \frac{C^L}{q}$. In the opposite benchmark case of no negotiation costs, $\kappa = 1$, it is easy to show that there is always a dispute. For general κ , it can be shown that decreasing q or increasing C^L shrinks the dispute region. This can be established easily by a graphical argument based on the negotiation frontiers in Figure 3. Fix a state γ and consider a decrease in q . There is a dispute iff the B^{net} point is in the quadrant defined by the P point and the FT point. By graphical inspection of Figure 3, as q falls, the B^{net} point can only move from inside this quadrant to outside it, not vice-versa. This is true both if the frontier has a convex portion (as in the bottom panel of Figure 3) and if it does not. Since this is true for any γ , it follows that a reduction in q decreases the likelihood of dispute. A similar graphical argument can be used to check that an increase in C^L reduced the likelihood of dispute.

We may now state:

Proposition 4. *Suppose the policy T is binary and transfers are costly. Then, even in the static setting the probability of a ruling is strictly positive, provided C^L is sufficiently small.*

Hence, when policies are discrete and transfers are costly, disputes can result in rulings (provided C^L is sufficiently small) independent of any learning effect. With continuous policies, no ruling is possible, whether or not transfers are costly; and it is easily shown that with discrete policies as in the current section, if transfers were costless rulings would again not be possible. It is the combination of these two features – discrete policies and costly transfers – that is capable of generating rulings in the static model.

The intuition for the finding reported in Proposition 4 is clearest in the case where the efficiency stakes of the dispute are very low and $\gamma^* - \gamma$ is essentially zero (though the stakes for each party in the dispute, γ and γ^* , could each still be very large). In this case joint surplus is hardly effected by the particular policy choice, and so there is essentially no cost to joint

surplus of a DSB ruling that is mistaken; and here, settling the dispute with the correct policy choice but with a costly transfer will be unattractive to the disputing parties relative to the expected joint surplus that comes from allowing the DSB to issue a (possibly mistaken) ruling and avoiding the use of the costly transfer.

Finally, we may also state:

Remark 4. *If DSB accuracy increases (q falls toward 0), both the frequency of disputes and the frequency of rulings goes down, but the likelihood that a dispute will end in a ruling may go up or down.*

This result highlights that, in the case of binary policy, the static effect of an increase in DSB accuracy is to depress the frequency of disputes and rulings, by increasing the predictability of the ruling. In what follows we will refer to this as the “predictability” effect.

We now turn to the dynamic setting. As before, we let x denote the cumulative number of rulings that have occurred prior to period 1, and $x + I$ the cumulative number of rulings that have occurred prior to period 2, where I is an indicator variable that takes the value 1 if there is a ruling in period 1 and takes the value 0 otherwise. At $t = 2$, we have the static outcome characterized in the previous sub-section given DSB noise $q(x)$: that is, x affects $\Pr(\text{Dispute})$ and $\Pr(\text{Ruling})$ at $t = 2$ only through the static predictability effect (which by Remark 4 is negative). We want to evaluate this effect at $t = 1$, where there is a static and a dynamic effect.

We consider first the probability of a ruling in period 1. At $t = 1$, $\Pr(\text{Ruling}) = \Pr(g < \delta\Delta)$, where as before g is the distance between the ND point and the negotiation frontier along the 45° line, and where $\Delta = E_\gamma[\Omega^{t=2}(x + 1, \gamma) - \Omega^{t=2}(x, \gamma)]$ is the gain in future joint payoff associated with increased DSB accuracy (the same across H and F given the veil of ignorance). Increasing x increases g (the static predictability effect). How does x affect Δ ? As we argued previously, this effect is negative if $\Omega^{t=2}$ is concave in x .

A complication is that, as x rises and $q(x)$ goes down, at the dispute margin there is a jump up in the joint payoff as we go from the no-dispute region (where the outcome is either at the P point or the FT point in Figure 3) to the dispute region (where the expected outcome will in general not correspond to either of these two points). This is a source of convexity in $\Omega^{t=2}$ with respect to x , and depending on the probability distribution of γ it could swamp the other effects. To mute this effect in what follows we will assume that C^L is sufficiently small. Focusing for the moment on $C^L = 0$, it is then easy to see that $\Omega^{t=2}$ is continuous in x , and

in fact the no-dispute region vanishes. Notice also from our derivation of (4.3) and (4.4) that $\Omega^{t=2}$ is continuous at the ruling margin for any C^L . Hence, when C^L is small, we need consider only two effects: (i) how x affects $\Omega^{t=2}$ inside ruling region; and (ii) how x affects $\Omega^{t=2}$ inside the settlement region. We will show that both effects are concave.

Let us first consider state realizations inside the ruling region. We focus on realizations such that $\gamma^* - \gamma > 0$ (the same argument applies for states such that $\gamma^* - \gamma < 0$). The expected joint payoff conditional on γ is $(1 - q(x))\Omega(P) + q(x)\Omega(FT)$. This is clearly concave in x for any γ , and hence concave in expectation.

Next consider state realizations inside the settlement region. We again focus on realizations such that $\gamma^* - \gamma > 0$. The disagreement payoffs are $\omega^D = (1 - q(x))v(P) + q(x)v(FT)$ and $\omega^{*D} = (1 - q(x))v^*(P) + q(x)v^*(FT)$. The gross bargaining payoffs are $\omega^B = \frac{1}{2}(\omega^{TOL} + \omega^D)$ and $\omega^{*B} = \frac{1}{2}(\omega^{*TOL} + \omega^{*D})$, where ω^{TOL} is H 's payoff if it makes a take-or-leave offer. It is straightforward to show that ω^{TOL} and ω^{*TOL} are linear in q . Finally, the net bargaining joint payoff is $\Omega^D + (1 - \kappa)(\Omega^B - \Omega^D)$, which of course preserves linearity in q . We can conclude again that expected joint payoff is concave in x . With this we may conclude that, provided C^L is sufficiently small, the likelihood of a ruling is decreasing in x for any δ .

Next we examine how $\Pr(\text{Dispute})$ at $t = 1$ depends on x . Recall that a dispute happens if the B^{net} point $(\omega^{Bn}, \omega^{*Bn})$ is outside the quadrant defined by the P and FT points in Figure 3, where now the B^{net} point is based on the disagreement point $ND + \delta\Delta$, that is, $\omega^{Bn} = \frac{\kappa}{2}\omega^{TOL} + (1 - \frac{\kappa}{2})[\omega^{ND} + \delta\Delta]$ and $\omega^{*Bn} = \frac{\kappa}{2}\omega^{*TOL} + (1 - \frac{\kappa}{2})[\omega^{*ND} + \delta\Delta]$. We have established just above that Δ is decreasing in x . And as we observed in the previous subsection, as x increases and $q(x)$ falls, the ND point can only move from inside the quadrant to outside it, a point that is clear from inspection of Figure 3. And the reduction in Δ associated with the increase in x only strengthens this effect, ensuring that the B^{net} point can only move from inside this quadrant to outside it, not vice-versa. So we can conclude that, provided C^L is sufficiently small, the likelihood of a dispute is decreasing in x for any δ .

Finally, it is straightforward to construct examples where the likelihood of settlement conditional on a dispute goes up or down with x . We can now state:

Proposition 5. *At $t = 1$ the likelihood of disputes and rulings is globally decreasing in x , provided C^L is sufficiently small (but the likelihood of settlement conditional on a dispute can go up or down).*

Now consider the forward-looking impact of learning. Consider $t = 1$. Comparing the case of learning with the benchmark case of no learning, we move from $\Delta = 0$ to $\Delta > 0$, and this clearly increases the likelihood of period-1 disputes and rulings. We may thus state:

Proposition 6. *The presence of learning-by-ruling increases the likelihood of period-1 disputes and rulings (but does not necessarily increase the likelihood that disputes end in a ruling in period 1).*

Note that introducing learning generates a “pivoting” of the time path of rulings/disputes: there are more early on and fewer later on.

Considering together the findings of this and the previous section, our model suggests that (i) disputes and rulings should be more frequent for binary than continuous policies (two reasons to go to court instead of one), and (ii) the negative effect of x on the likelihood of disputes and rulings should be stronger for binary than continuous policies, because with binary policies the static effect and dynamic effect of a change in x go in the same direction, whereas for continuous policies these two effects counteract each other; as a consequence, with continuous policy the effect of x is guaranteed to be negative only for high δ , whereas with binary policies it is negative for any δ .

5. Empirical Evidence

To explore the empirical content of our theory, we examine patterns in WTO dispute behavior and focus on a broad prediction of the model: if there is DSB learning, the likelihood of current disputes and rulings should tend to decrease with the stock of cumulative past rulings. We have shown that this prediction applies both to general learning (for rulings and disputes) and to directed-dyad learning (for rulings, though not necessarily for disputes), and that it holds when policies are continuous if governments are sufficiently patient and when policies are discrete if litigation costs are not too high.

Our model only has two countries and a single sector, but in a multi-country reality with multiple policy areas, DSB learning could be general or specific in scope, and in the latter case it could be specific to disputant countries (e.g. learning about China, or about the trade interaction between the US and India) and/or to an issue area (e.g. learning about agricultural policy issues). We will attempt to investigate empirically these different potential domains of

learning. As a first pass, we distinguish among three levels: country dyad, WTO article (under the assumption that an “issue area” is embodied in a GATT/WTO article), and general.

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

Our dataset consists of 388 WTO disputes initiated between 1995 and 2009 as contained in the WTO Dispute Settlement Database and described in Maggi and Staiger (2015).¹⁹ Letting $i = 1, \dots, I$ index country dyads that have had at least one WTO dispute in our sample period, $k = 1, \dots, K$ index GATT/WTO Articles that were disputed at least once in our sample period, and $t = 1, \dots, T$ index years in our sample period, we define the following variables, each of length IKT :

- $D_{i,k,t} \equiv$ number of disputes initiated by country-dyad i on article k in year t ,
- $R_{i,k,t} \equiv$ number of disputes by country-dyad i on article k that led to an adopted panel ruling in year t ,
- $CR_{i,k,t} \equiv$ cumulative number of disputes by country-dyad i on article k that led to an adopted panel ruling prior to year t .

Notice that our convention is to date disputes by the year in which they are formally initiated (through a “request for consultation,” the official start of formal WTO dispute settlement

¹⁹Specifically, starting with the 426 WTO disputes initiated between 1995 and August 2011 covered in the WTO Dispute Settlement Database, we follow Maggi and Staiger (2015) and drop the 24 disputes in this data set initiated after January 1 2010 (to avoid truncation of dispute outcomes in the dataset). And like Maggi and Staiger, we drop as well 8 cases where the issue formally returns in a later dispute (which we include) or is simply handled formally in another dispute (which we include). And finally, we drop the 6 multi-complainant cases in this dataset – the 5 described in Maggi and Staiger, 2015, plus the additional multi-complainant case DS35 which was dropped by Maggi and Staiger on other grounds – 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.

proceedings), and to date DSB rulings by the year in which the DSB panel report containing the ruling is formally “adopted” (approved) by the WTO membership. The latter dating convention reflects our belief that the entire panel process – investigation, preliminary and final reports, and appeals – that leads up to final adoption of DSB rulings is a potentially important source of DSB learning. Our primary goal is to explore the possibility that $D_{i,k,t}$ and $R_{i,k,t}$ might vary with $CR_{i,k,t}$ in ways that could reflect the impacts of DSB learning.

5.1. Some simple plots

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

In Plot 3 we depict on the vertical axis $R_{.,k,t}$, the number of disputes over article k (between the countries of any dyad) that ended in an adopted panel ruling in year t , and on the horizontal axis we depict $CR_{.,k,t}$, the cumulative number of disputes over article k (between the countries of any dyad) that ended in an adopted panel ruling prior to year t . The appearance of a negative relationship in Plot 3 is consistent with the impacts of article-specific DSB learning according to our model.²⁰

In Plot 4 we depict on the vertical axis $R_{i.,t}$, the number of disputes between the countries of dyad i (on any article) that ended in an adopted panel ruling in year t , and on the horizontal axis we depict $CR_{i.,t}$, the cumulative number of disputes between the countries of dyad i (on any article) that ended in an adopted panel ruling prior to year t . The appearance of a negative relationship in Plot 4 is consistent with the impacts of dyad-specific DSB learning according to our model, though this relationship seems weaker than the negative relationship in Plot 3.

Finally, in Plot 5 we depict on the vertical axis $R_{i.,t}$, the number of disputes between the countries of dyad i (on any article) that ended in an adopted panel ruling in year t , and on the horizontal axis we depict $CR_{.,t}$, the cumulative number of disputes – on any article and by any dyad – that ended in an adopted panel ruling prior to year t . Unlike for Plots 3 and 4, Plot 5 shows no discernible relationship between current rulings and cumulative past rulings, and hence no suggestion of general DSB learning according to our model.

²⁰The extreme outlier in Plot 3 corresponds to GATT Article III, the National Treatment clause. There are a number of sub-articles associated with the National Treatment clause that have been the topic of WTO disputes, and it is likely that disaggregating further the GATT/WTO articles will reveal that this outlier is an artifact of aggregation over sub-articles. This is the subject of some of our ongoing work.

5.2. Logit regressions

We next turn to some basic logit regressions, in order to probe the visual impressions suggested by Plots 3-5. We will consider both *undirected dyads* (i.e., dyads defined by the identities of the two countries, without regard to which role – claimant or defendant – they play in the dispute) and *directed dyads* that distinguish between a dyad composed of countries l and m depending on whether country l is the claimant and m the defendant or vice-versa. To facilitate this distinction, we now introduce the notation $D_{ij}^{\leftrightarrow,k,t}$ in place of $D_{i,k,t}$ to refer to the undirected dyad version of this variable, and similarly we now use $R_{ij}^{\leftrightarrow,k,t}$ and $CR_{ij}^{\leftrightarrow,k,t}$ to refer respectively to $R_{i,k,t}$ and $CR_{i,k,t}$. We will then use $D_{ij}^{\rightarrow,k,t}$, $R_{ij}^{\rightarrow,k,t}$ and $CR_{ij}^{\rightarrow,k,t}$ to represent the directed dyad versions of these variables, where country i is the claimant and country j is the defendant. In what follows we present our results from logit estimation, but similar results are obtained under OLS.

Focusing first on undirected dyads, we run two logit regressions, one for disputes and one for rulings. In the dispute regression, the dependent variable is $DLogit_{ij}^{\leftrightarrow,k,t}$ (defined as 1 if $D_{ij}^{\leftrightarrow,k,t} \geq 1$ and 0 otherwise) and the independent variables are $CR_{ij}^{\leftrightarrow,k,t}$, $CR_{n(ij),k,t}$, $CR_{ij}^{\leftrightarrow,nk,t}$ and $CR_{n(ij),nk,t}$, where a subscript nz denotes “not z ” for index z . And in the ruling regression, the dependent variable is $RLogit_{ij}^{\leftrightarrow,k,t}$ (defined as 1 if $R_{ij}^{\leftrightarrow,k,t} \geq 1$ and 0 otherwise) and the independent variables are the same as above. The results of these regressions are presented in columns 1 and 2 of Table 1. In both regressions we also include a quadratic time trend and article- and (undirected)-dyad- fixed effects.²¹

The variable $CR_{n(ij),k,t}$ is the cumulative number of disputes on article k by dyads other than \overleftrightarrow{ij} that ended in a ruling prior to year t . This variable is meant to capture article-specific learning: if there is DSB learning about GATT/WTO Articles that is transferable across pairs of disputants, then we would expect according to our theory that the estimated coefficient on $CR_{n(ij),k,t}$ would be negative. The variable $CR_{ij}^{\leftrightarrow,nk,t}$ is the cumulative number of disputes by dyad \overleftrightarrow{ij} on articles other than k that ended in a ruling prior to year t . This variable is meant to capture dyad-specific learning: if there is DSB learning about the pair of disputants that is

²¹For the dispute regression our sample spans the period 1995-2009, and it includes country dyads that initiated at least one WTO dispute during this period and GATT/WTO Articles that were disputed at least once during this period. For the ruling regression, we restrict the sample to country dyads that generated at least one WTO adopted ruling report as a result of a dispute initiated during this period and to 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.

transferable across GATT/WTO Articles, then we would expect that the estimated coefficient on $CR_{ij, nk, t}^{\leftrightarrow}$ would be negative. And the variable $CR_{n(ij), nk, t}^{\leftrightarrow}$ is the cumulative number of disputes by dyads other than \overleftrightarrow{ij} on articles other than article k that ended in a ruling prior to year t . This variable is meant to capture general learning: if there is DSB learning that is transferable across dyads and GATT/WTO Articles, then we would expect that the estimated coefficient on $CR_{n(ij), nk, t}^{\leftrightarrow}$ would be negative.

Finally, the variable $CR_{ij, k, t}^{\leftrightarrow}$ is the cumulative number of disputes on article k by dyad \overleftrightarrow{ij} that ended in a ruling prior to year t . This variable is meant to capture learning that is specific to both the disputants involved and the article that they are disputing, and we would expect its coefficient to be negative if any DSB learning is present; however, it is likely that this narrowest measure of DSB experience may conflate other forces that are outside our model, such as “unfinished business” between a given pair of disputants on a particular article over which a dispute between them has arisen.²² Hence, while we include this variable in our estimating equations, we do not emphasize it as a reliable indicator of DSB learning the way we do with the other variables.²³

Focusing first on the ruling regression in column 2, the estimated coefficients on $CR_{n(ij), k, t}^{\leftrightarrow}$ and $CR_{ij, nk, t}^{\leftrightarrow}$ are negative and strongly significant, as is the estimated coefficient on $CR_{n(ij), nk, t}^{\leftrightarrow}$. The coefficient estimates on $CR_{n(ij), k, t}^{\leftrightarrow}$ and $CR_{ij, nk, t}^{\leftrightarrow}$ confirm the visual impressions of Plots 3 and 4: there is evidence consistent with article-specific and dyad-specific DSB learning. And contrary to the visual impression left by Plot 5, the coefficient estimate on $CR_{n(ij), nk, t}^{\leftrightarrow}$ in column 2 also indicates evidence of general DSB learning. As we will see, however, the evidence for general learning weakens substantially when we consider directed dyad logits. Finally, notice that the coefficient on $CR_{ij, k, t}^{\leftrightarrow}$, our narrowest measure of DSB experience, is positive but not significantly different from zero, with a relatively large standard error. As we have discussed above, we attribute this finding to the fact that $CR_{ij, k, t}^{\leftrightarrow}$ very likely conflates possible learning effects with other forces that are outside our model.

Turning to the dispute regression, column 1 presents the coefficient estimates from the $DLogit_{ij, k, t}^{\leftrightarrow}$ regression. Here the results are somewhat ambiguous: the estimated coefficient on

²²As we describe in note 19, we have excluded from our sample of WTO disputes those cases where the disputed issue formally returns in a later dispute which we include, but there is still plenty of room for “unfinished business” to be present in our sample of disputes at a less formal level.

²³Moreover, the variation in $CR_{ij, k, t}^{\leftrightarrow}$ is likely to be rather small, and as we discuss below the comparatively high standard error that we find strengthens this suspicion.

$CR_{ij,nk,t}^{\leftrightarrow}$ is still negative and strongly significant, consistent with dyad-specific DSB learning; but the estimated coefficients on $CR_{n(\overleftarrow{ij}),k,t}$ and $CR_{n(\overrightarrow{ij}),nk,t}$, while also still negative, lose some of their significance, suggesting only weaker evidence consistent with article-specific and general DSB learning. And now the coefficient on $CR_{ij,k,t}^{\leftrightarrow}$, our narrowest measure of DSB experience, is positive and strongly significant. We will come back to the interpretation of these findings below, after discussing the results of the dispute regression with directed dyads.

We next turn to our analysis based on *directed* dyads. There are at least two reasons to consider directed dyads. The first reason is that, as we discussed in section 3.3, DSB learning might be specific to the defendant country (which is under the magnifying glass of the DSB), or to the complainant country (e.g. because the DSB learns about the political-economic impacts of trade barriers on this country’s exporters), or even to the directed dyad itself (e.g. by adjudicating disputes brought by China against the US, the DSB may learn about sectors where China exports to the US). And a second reason for focusing on directed dyads is that this may allow us to filter out possible “tit-for-tat” effects, meaning that for example if the US files today against China, in the future China is more likely to file against the US.²⁴

As with the undirected dyads, we run two logit regressions for the directed dyads, one for disputes and one for rulings. In the dispute regression, the dependent variable is $DLogit_{ij,k,t}^{\rightarrow}$ (defined as 1 if $D_{ij,k,t}^{\rightarrow} \geq 1$ and 0 otherwise), and the independent variables are $CR_{ij,k,t}^{\rightarrow}$, $CR_{(ni)j,k,t}^{\rightarrow}$, $CR_{i(nj),k,t}^{\rightarrow}$, $CR_{j^i,k,t}^{\rightarrow}$, $CR_{other,k,t}^{\rightarrow}$, $CR_{ij,nk,t}^{\rightarrow}$, $CR_{(ni)j,nk,t}^{\rightarrow}$, $CR_{i(nj),nk,t}^{\rightarrow}$, $CR_{j^i,nk,t}^{\rightarrow}$ and $CR_{other,nk,t}^{\rightarrow}$. And in the ruling regression, the dependent variable is $RLogit_{ij,k,t}^{\rightarrow}$ (defined as 1 if $R_{ij,k,t}^{\rightarrow} \geq 1$ and 0 otherwise) and the independent variables are the same as above. The results of these regressions are presented in columns 3 and 4 of Table 1. Similarly to the undirected-dyad regressions, in both of our directed-dyad regressions we also include a quadratic time trend and article- and (directed)-dyad- fixed effects.²⁵

We now explain the meaning and interpretation of the independent variables listed above. There are two groups of five CR variables, the first group specific to article k and the second to

²⁴Indeed, there is a fair amount of anecdotal evidence of such tit-for-tat behavior in the practice of WTO disputes (see for example the article by Jennifer Freedman in Bloomberg Business, 2012, and the discussion in Davis, 2012).

²⁵In analogy with our sample for the undirected dyad regressions (see note 21): for our directed dyad dispute regressions the sample spans the period 1995-2009 and includes *directed* dyads that initiated at least one WTO dispute during this period and GATT/WTO Articles that were disputed at least once during this period; and for the ruling regression, we restrict the sample to *directed* dyads that generated at least one WTO adopted ruling report as a result of a dispute initiated during this period and to 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.

all other articles (not k). Consider first the k -specific variables: these five variables correspond to five subsets of the universe of directed dyads:

(1) the same directed dyad as in the dependent variable (\overrightarrow{ij}); the corresponding variable $CR_{\overrightarrow{ij},k,t}$ is meant to capture learning that is directed-dyad-and-article specific;

(2) directed dyads where j is the defendant but the complainant is not i ; the corresponding variable $CR_{\overrightarrow{(ni)j},k,t}$ is meant to capture learning that is defendant-and-article specific;

(3) directed dyads where i is the complainant but the defendant is not j ; the corresponding variable $CR_{\overrightarrow{i(nj)},k,t}$ is meant to capture learning that is complainant-and-article specific;

(4) the “reverse” directed dyad, i.e. where j complains against i ; the corresponding variable $CR_{\overrightarrow{ji},k,t}$ might in principle capture learning that is reverse-dyad-and-article specific, but it might also realistically pick up tit-for-tat effects, which would go in the opposite direction.

(5) all the remaining directed dyads; the corresponding variable ($CR_{\overrightarrow{other},k,t}$) is meant to capture learning that is article specific (but not disputant specific).

The second group of five variables is analogous, except that cumulative rulings are aggregated over all non- k articles. And the interpretation of these variables is also analogous, except that they capture non-article-specific effects: for example, the $CR_{\overrightarrow{ij},nk,t}$ variable is meant to capture *directed*-dyad-specific DSB learning, the $CR_{\overrightarrow{i(nj)},nk,t}$ variable is meant to capture claimant-specific DSB learning, and $CR_{\overrightarrow{other},nk,t}$ is meant to capture general learning.

Focusing again first on the ruling regression in column 4, the estimated coefficient on $CR_{\overrightarrow{other},k,t}$ is negative and strongly significant, which is consistent with the presence of article-specific DSB learning (of the remaining four k -specific variables, the estimated coefficient on the defendant-and-article-specific variable $CR_{\overrightarrow{(ni)j},k,t}$ is negative but only marginally significant while none of the others are significantly different from zero). The estimated coefficient on $CR_{\overrightarrow{ij},nk,t}$ is negative and strongly significant, consistent with directed-dyad-specific DSB learning. And the estimated coefficient on $CR_{\overrightarrow{i(nj)},nk,t}$ is negative and strongly significant, consistent with claimant-specific DSB learning. But the estimated coefficient on $CR_{\overrightarrow{other},nk,t}$, while negative, is only marginally significant, suggesting only weak evidence of general DSB learning. And in comparison with the more aggregated measure of general DSB experience $CR_{n(\overleftarrow{ij}),nk,t}$ in the undirected dyad ruling regression of column 2, the strongly negative coefficient on that variable can be understood as reflecting the coefficient on $CR_{\overrightarrow{i(nj)},nk,t}$ in the directed dyad ruling regression of column 4 (and therefore can be said to reflect claimant-specific DSB learning). Finally, the estimated coefficient on $CR_{\overrightarrow{ji},nk,t}$ is positive and strongly significant, suggesting as

we indicated earlier the possible importance of a tit-for-tat dynamic that is outside our model.

Turning to the dispute regression results in column 3, the results are broadly consistent with the ruling regressions of column 4. In particular, the estimated coefficients on $CR_{other,k,t}^{\rightarrow}$, $CR_{ij,nk,t}^{\rightarrow}$ and $CR_{i(nj),nk,t}^{\rightarrow}$ are each negative and strongly significant, consistent with the presence of article-specific, defendant-and-article-specific and claimant-specific DSB learning. And there is no evidence of general DSB learning from the dispute regression (the estimated coefficient on $CR_{other,nk,t}^{\rightarrow}$ is now statistically insignificant), reinforcing the caution with which we interpreted the marginally significant coefficient on this variable in the ruling regression. And as with the ruling regression, the estimated coefficient on $CR_{ji,nk,t}^{\rightarrow}$ is positive and strongly significant, possibly reflecting tit-for-tat effects.

The one difference with respect to the ruling regression results in column 4 is that in the dispute regression results of column 3 the estimated coefficient on $CR_{(ni)j,k,t}^{\rightarrow}$ has switched from negative and marginally significant to positive and strongly significant. Recalling that our model yields more ambiguous predictions about the impacts of non-general learning variables such as $CR_{(ni)j,k,t}^{\rightarrow}$ on the frequency of disputes than it does for rulings, it is possible that the positive coefficient on $CR_{(ni)j,k,t}^{\rightarrow}$ in the dispute regression of column 3 is a manifestation of this ambiguity. A further possible interpretation is that this reflects a “bandwagon effect” that falls outside our model, whereby other potential claimants follow up with claims of their own once a ruling on an article k claim against defendant-country j has been issued and adopted. In any event, notice that, comparing the directed dyad dispute regression results in column 3 with the undirected dyad dispute regression results in column 1, the marginal significance of the negative coefficient on the article-specific variable $CR_{n(\overrightarrow{ij}),k,t}$ in the undirected dyad dispute regression can be understood as reflecting the positive and strongly significant coefficient on the more disaggregated $CR_{(ni)j,k,t}^{\rightarrow}$ in the directed dyad dispute regression.

It is notable that the coefficient of the linear time trend is positive in all of our regressions. The fact that controlling for our measures of court experience (the CR variables) wipes out the negative effect of calendar time suggests that court learning can indeed explain the raw declining trend in disputes and rulings that was evidenced in Plots 1 and 2, as we hypothesized at the outset.

Thus far we have interpreted our empirical findings as reflecting the effects of DSB learning that occur over the panel process of investigation, preliminary and final reports, and appeals leading up to final adoption of DSB rulings. An important question is whether there are

alternative interpretations of these empirical findings. One plausible candidate is that there is learning going on, but that it actually takes the form of governments learning about each other. To consider this alternative interpretation, we have re-run the regressions in Table 1 (in both logit and OLS) replacing the cumulative-stock-of-ruling CR variables on the right-hand side with analogous “ CS ” variables that measure the cumulative stock of formal consultations (facilitated by the WTO secretariat and held in private between the disputing parties) that *settle* prior to panel formation. If governments learn about each other during these consultations and if this has an important impact on the frequency of subsequent disputes and rulings along similar lines to the DSB learning in our model, we would expect this to show up in negative and significant coefficients on the CS variables pertaining to the dyad of the consulting parties (that is, on the $CS_{ij,k,t}^{\leftarrow}$ and $CS_{ij,nk,t}^{\leftarrow}$ variables in the undirected dyad regressions, and on the $CS_{ij,k,t}^{\rightarrow}$, $CS_{ij,nk,t}^{\rightarrow}$, $CS_{ji,k,t}^{\leftarrow}$ and $CS_{ji,nk,t}^{\leftarrow}$ variables in the directed dyad regressions). In fact, we fail to find any robust evidence for such coefficient estimates.

A second plausible candidate that could provide an alternative interpretation of at least *some* of our findings is the impact of legal precedent. Under this interpretation, court rulings help to complete the incomplete WTO contract (as for example in Maggi and Staiger, 2011), so as the stock of rulings accumulate, there are fewer and fewer contingencies that are left uncovered by the contract, thus the frequency of rulings may naturally decrease. More specifically, suppose that a given WTO article k is initially incomplete and is silent about a set M_0 of contingencies, out of the total set of contingencies M . Suppose further that, in each period of time, one contingency is randomly selected out of the set M , and if this contingency is not covered by the contract, the court may be called upon to specify the contractual obligations for this contingency. Then, in this simple scenario, as the stock of rulings accumulates the probability of new rulings goes down. Admittedly, the legal precedent interpretation may well explain part or even all of our empirical findings with regard to effects we attribute to article-specific learning.²⁶ But importantly, this explanation does not seem compelling as an alternative to our DSB learning story when it comes to our findings regarding defendant- and claimant-specific effects. Hence, we view legal precedent as plausibly being part, but *only* part, of the explanation of our empirical findings above.

²⁶An interesting possibility to distinguish between these two interpretations of our findings regarding article-specific effects might be to investigate whether these effects are also present in the early GATT era, when legal precedent was by all accounts not operative (see, for example, the discussion of the views of GATT/WTO legal scholars on this point in Maggi and Staiger, 2011). We view this as a promising avenue for further research.

Overall, the results of the undirected and directed dispute and ruling regressions reveal several important points. First, there appears to be evidence consistent with court learning, and in particular with article-specific, directed-dyad-specific and claimant-specific learning. Interestingly, we find only weak evidence of general-scope learning. Second, focusing on the pre-panel consultation stage of the WTO dispute settlement process, we do not find evidence in support of the possibility that governments learn about each other. And finally, our empirical investigation has uncovered evidence consistent with possible tit-for-tat and “bandwagon” effects, and so a more complete empirical account of the pattern of WTO disputes and rulings may require an extended model that captures these effects in addition to the effects of court learning on the dynamics of dispute resolution.²⁷

6. Conclusion

[TBA]


²⁷While various stories about a bandwagon effect seem plausible, the details of court remedies (e.g., how complete they are, whether they apply effectively to 3rd parties) would matter, and as a result it is not immediately obvious whether rulings for or rather against the defendant would be more likely to stimulate follow-up disputes by other claimants. Similar subtleties may arise with tit-for-tat effects. This points to the value of carefully modeling such effects before attempting to go further in investigating their empirical content, a task we leave to future research.

References


- Baker, Scott and Claudio Mezzetti (2012), "A Theory of Rational Jurisprudence," *Journal of Political Economy*, vol. 120(3), pp. 513 - 551.
- Baldwin, Richard (1987), "Politically Realistic Objective Functions and Trade Policy," *Economics Letters*, vol. 24, pp. 287-90.
- Bebchuck, Lucian A. (1984), "Litigation and Settlement under Imperfect Information," *RAND Journal of Economics* 15(3): Autumn, 404-415.
- Beim, D. (2014), "Learning in the Judicial Hierarchy", mimeo.
- 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, no. 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, no. 4 (2000).
- 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.
- Conconi, Paola, David R. DeRemer, Georg Kirchsteiger and Lorenzo Trimarchi (2015), "Suspiciously Timed Trade Disputes", mimeo.
- Davis, Christina (2012), *Why Adjudicate? Enforcing Trade Rules in the WTO*. Princeton University Press.
- Davis, Christina L. and Sarah Blodgett Bermeo (2009), "Who files? Developing Country Participation in GATT/WTO Adjudication," *Journal of Politics*, 2009, 71 (3), 1033–1049.

- Freedman, Jennifer (2012), "China Floods the WTO with Tit-for-Tat," *Bloomberg Business*, June 7.
- Grossman, Gene M. and Elhanan Helpman (1994), "Protection for Sale," *American Economic Review*, 84(4), pp. 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*, 2005, 34 (2), 557-598.
- Horn, Henrik, Louise Johannesson and Petros C. Mavroidis (2011), "The WTO Dispute Settlement System 1995-2006: Some Descriptive Statistics," *Journal of World Trade*, 45(6).
- Kuenzel, David J. (2015), "WTO Dispute Determinants", mimeo.
- Maggi, Giovanni and Robert W. Staiger (2011), "The Role of Dispute Settlement Procedures in International Trade Agreements," *Quarterly Journal of Economics*, February.
- Maggi, Giovanni and Robert W. Staiger (2014), "Optimal Design of Trade Agreements in the Presence of Renegotiation," *AEJ-Micro*.
- Maggi, Giovanni and Robert W. Staiger (2015), "Trade Disputes and Settlement", mimeo.
- Park, Jee-Hyeong (2011), "Enforcing International Trade Agreements with Imperfect Private Monitoring," *Review of Economic Studies*.
- Reinganum, Jennifer F. and Louis L. Wilde (1986), "Settlement, litigation, and the allocation of litigation costs," *RAND Journal of Economics* 17(4): Winter, 557-566.
- Staiger, Robert W. and Alan O. Sykes (2013), "How Important can the Non-Violation Clause be for the GATT/WTO?," mimeo.
- Stevens, John Paul (2005), "Learning on the Job," *Fordham Law Review*, 74, 1561-7.

Plot 1



Plot 2



Note: the ratio X/Z , where X is the number of disputes or rulings involving diads of age i, and Z is the total trade volume of diads of age i.

Figure 1: Continuous policy (static setting)

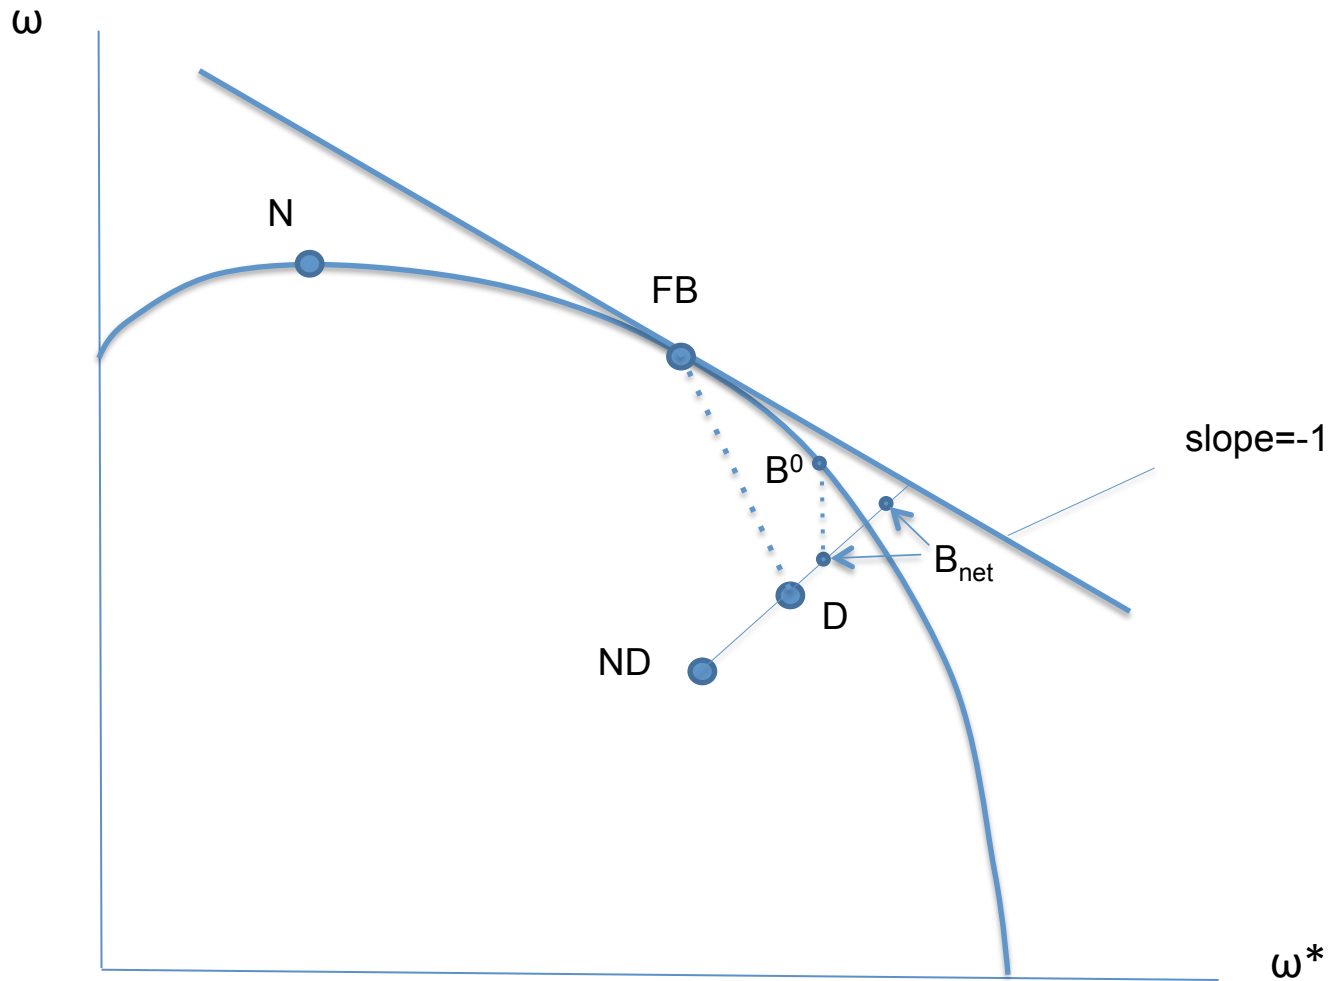


Figure 2: Continuous policy (static setting)




Figure 3: Binary policy, Pareto frontier





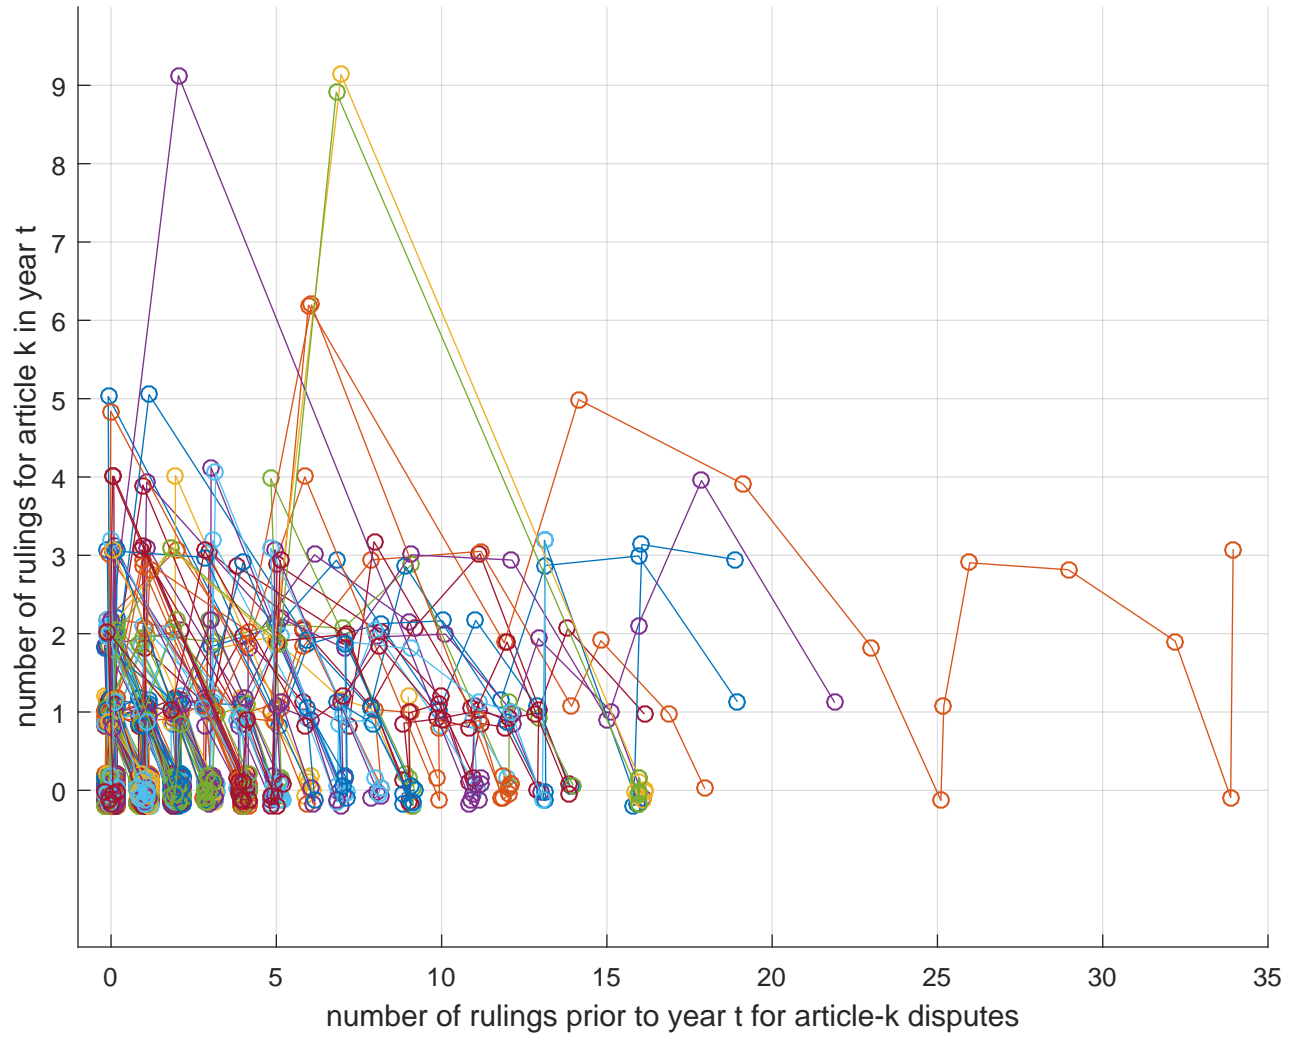



Figure 4: Binary policy: when negotiation ends in ruling



Plot 3



Plot 4



Plot 5

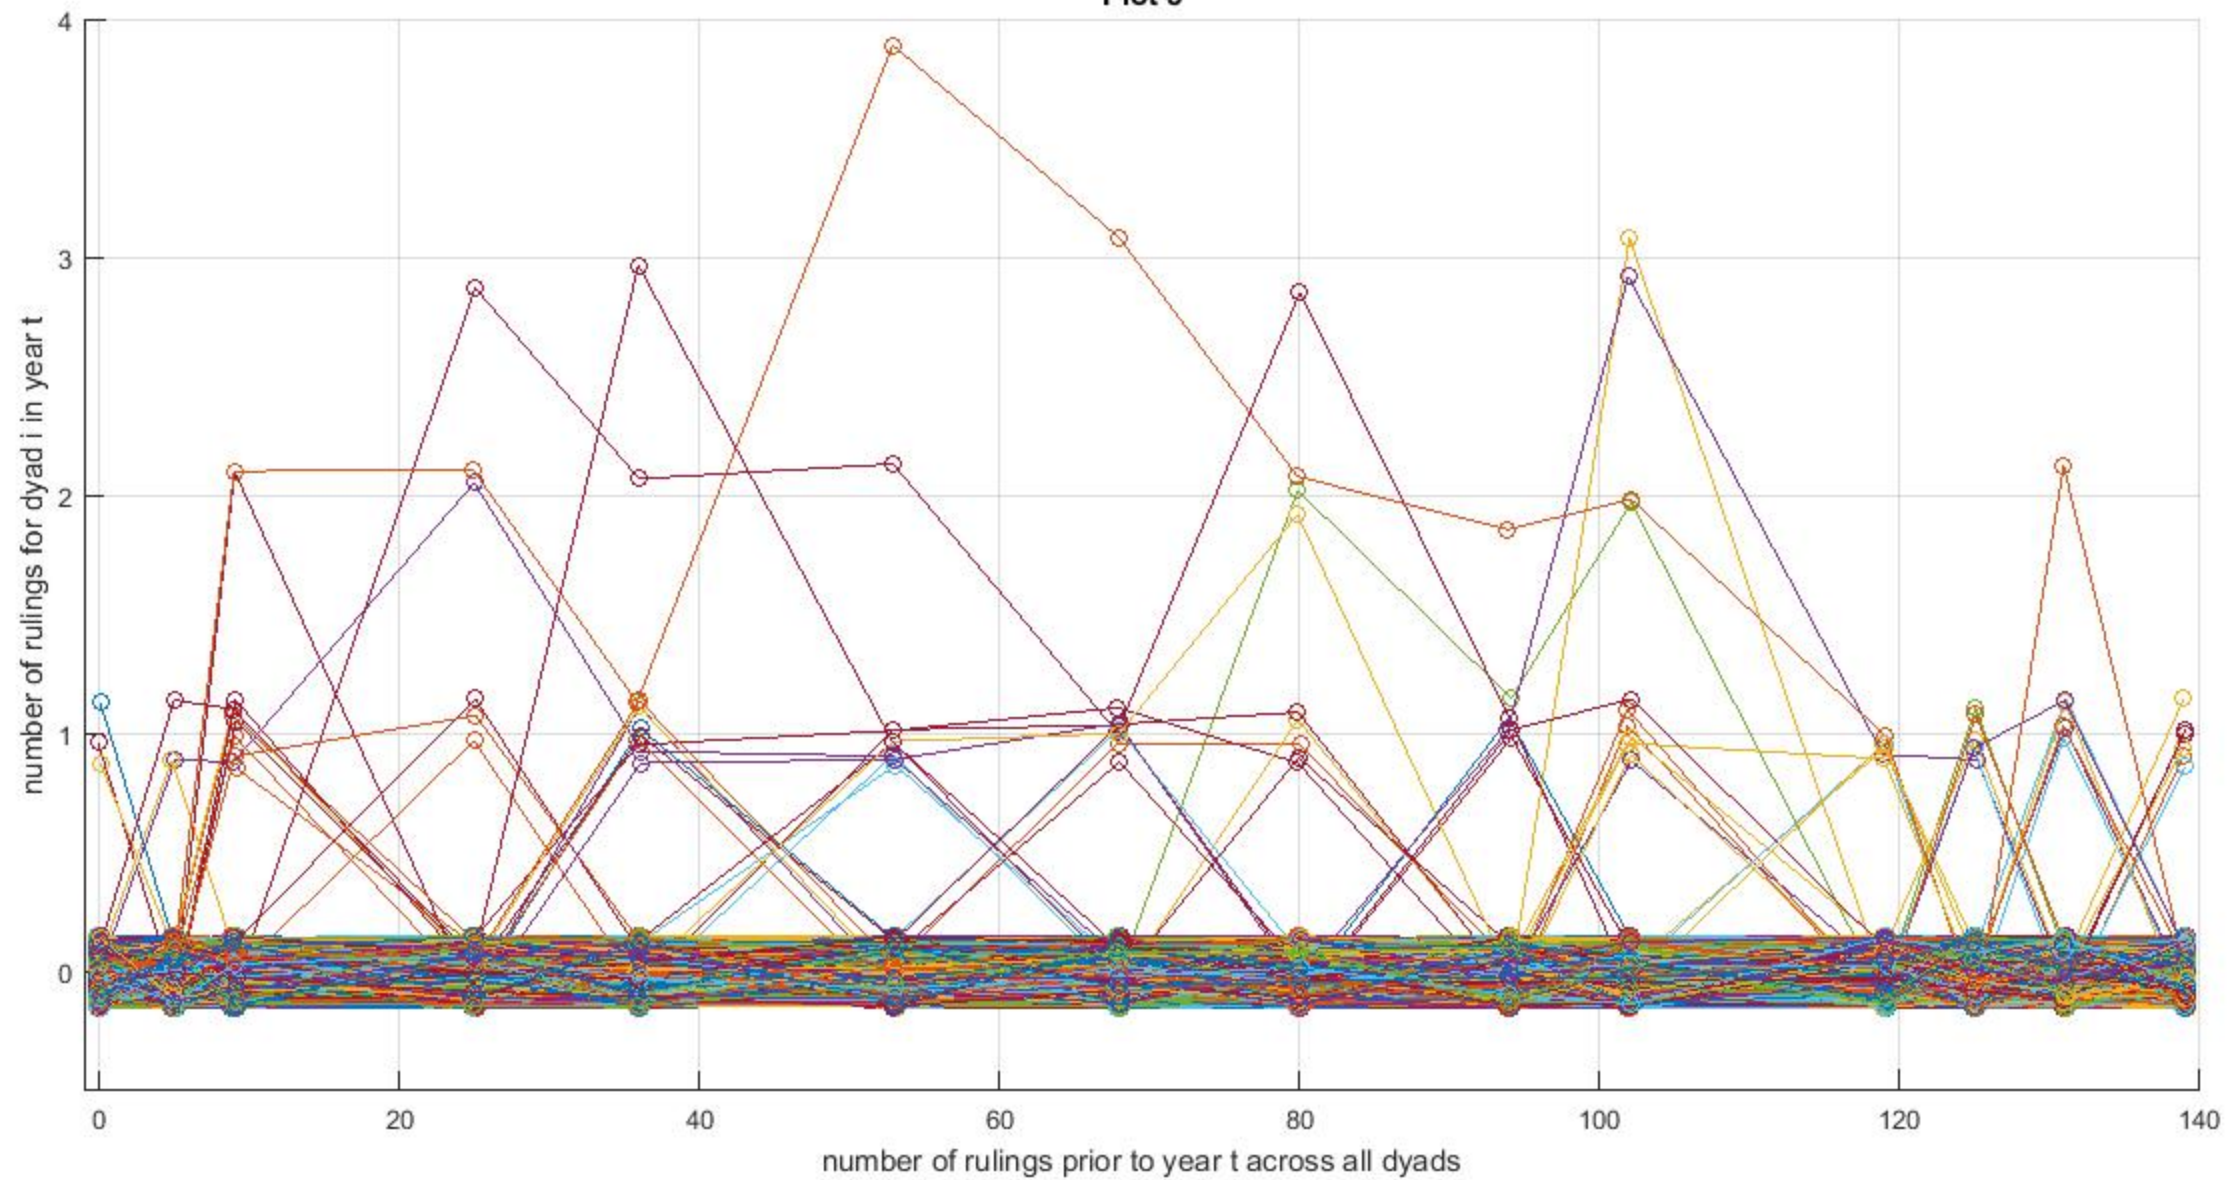


Table 1

VARIABLES	Undirected Dyad		VARIABLES	Directed Dyad	
	(1) DLogit_ \vec{ij} kt	(2) RLogit_ \vec{ij} kt		(3) DLogit_ \vec{ij} kt	(4) RLogit_ \vec{ij} kt
CR_ \vec{ij} kt	0.186*** (0.0646)	0.119 (0.0913)	CR_ \vec{ij} kt	0.144 (0.168)	-0.422 (0.271)
CR_n(\vec{ij})kt	-0.0127* (0.00761)	-0.160*** (0.0168)	CR_ $\overline{(ni)}$ jkt	0.115*** (0.0186)	-0.0528* (0.0309)
CR_ \vec{ij} nkt	-0.0272*** (0.00256)	-0.0644*** (0.00540)	CR_ $\overline{i(nj)}$ kt	0.0397 (0.0277)	-0.0529 (0.0509)
CR_n(\vec{ij})nkt	-0.00190** (0.000841)	-0.00663*** (0.00199)	CR_ \vec{ji} kt	0.0310 (0.186)	0.464 (0.288)
			CR_ \overline{other} kt	-0.0371*** (0.00833)	-0.196*** (0.0192)
			CR_ \vec{ij} nkt	-0.266*** (0.0174)	-0.803*** (0.0464)
			CR_ $\overline{(ni)}$ jnkt	0.000216 (0.00108)	-0.00366 (0.00269)
			CR_ $\overline{i(nj)}$ nkt	-0.00493*** (0.00184)	-0.0148*** (0.00568)
			CR_ \vec{ji} nkt	0.218*** (0.0172)	0.693*** (0.0449)
			CR_ \overline{other} nkt	-0.00125 (0.000838)	-0.00434* (0.00224)
t	0.198*** (0.0260)	1.277*** (0.111)	t	0.270*** (0.0264)	1.832*** (0.144)
t ²	-0.00726*** (0.00225)	-0.0313*** (0.00416)	t ²	-0.0113*** (0.00227)	-0.0530*** (0.00501)
Constant	-9.201*** (1.076)	-10.90*** (0.723)	Constant	-9.650*** (1.077)	-14.54*** (0.895)
Observations	439,584	112,560	Observations	545,142	149,520
Pseudo R ²	0.277	0.229	Pseudo R ²	0.249	0.270
\vec{ij} FE	Y	Y	\vec{ij} FE	Y	Y
kFE	Y	Y	kFE	Y	Y

Standard errors in parentheses

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