

The Price of Decentralized Enforcement in the WTO

Leslie Johns[†]

Krzysztof J. Pelc[‡]

EARLY DRAFT: November 2014

Abstract

Like many international organizations, the WTO allows states to litigate potential rule violations through its dispute settlement system. Yet many policies that appear to violate WTO rules remain unchallenged, even when they have a significant economic impact. Why is this? We argue that the likelihood that a trade policy is challenged is linked to the policy's distributional impact. When a policy's impact is concentrated on a single country, that country has a stronger incentive to challenge it than when the policy's impact is spread across many WTO members. States disproportionately favor disputes with concentrated results, even though this results in aggregate under-enforcement. The resulting selection effect has a paradoxical implication: violations with more concentrated effects, which are less likely to be challenged in any period, should also fare better in litigation, conditional on being filed. Conversely, violations with concentrated effects, that are challenged the fastest, should fare the worst before judges. The evidence, which considers all WTO disputes from 1995 to 2013, bears out these beliefs.

[†]Department of Political Science, University of California, Los Angeles, ljohns@polisci.ucla.edu

[‡]Department of Political Science, McGill University, kj.pelc@mcgill.ca

1 Introduction

Like many international organizations, the World Trade Organization (WTO) relies upon member-states to enforce its rules. While the WTO oversees the adjudication of trade disputes, it lacks the central authority to flag or challenge possible rule violations on behalf of its members. Instead, individual members must monitor each other's trade policies and challenge suspected violations themselves. We argue that such decentralized enforcement carries an observable cost.

WTO enforcement can be highly effective, and members often challenge suspicious trade policies as soon as they appear. For example, the United States blocked entry to Canadian trucks carrying cattle and swine in 1998, citing health concerns. Canada held that this policy, which affected only Canadian trucks, violated the WTO's Sanitary and Phytosanitary Agreement. Canada did not wait long before acting: 15 days after the policy began, it filed a dispute and requested expedited consultations with the US.¹

Yet WTO enforcement is not always as effective. Trade violations can go unchallenged for years, if not decades. For example, the US Agriculture Improvement and Reform Act of 1996 (FAIR Act) subsidized US producers of corn in an export-contingent manner, seemingly in contravention of the Subsidies and Countervailing Measures Agreement. Because corn is such a widely produced commodity, and because subsidies depress world prices, hurting all corn-producing countries, the FAIR Act harmed a huge number countries. When Canada finally challenged US corn subsidies, however, it had been 4,025 days since their implementation—over 11 years! The broad multilateral effect of the US policy was apparent in WTO proceedings: after the case was filed by Canada, 15 WTO members (including the EU, and thus representing a total of 40 states) joined as third parties in the consultations.² Nevertheless, the US subsidies were not challenged for over a decade. Why did Canada sue the US in 1998 over a relatively limited trade restriction on cattle and swine, while ignoring massive US corn subsidies? We argue that Canada's decision can be explained in part by how these trade policies affected *other* states.

Our broader claim is that the effectiveness of decentralized enforcement—in the WTO and other international organizations—depends on the distributional impact of rule violations. Countries that bear the cost of decentralized enforcement often provide a public good at a private cost.

¹See DS144: *United States—Certain Measures Affecting the Import of Cattle, Swine and Grain from Canada*.

²See DS357: *United States—Subsidies and Other Domestic Support for Corn and Other Agricultural Products*.

Initiating a WTO dispute is costly: filing requires considerable legal capacity (Busch, Reinhardt and Shaffer, 2009) and political considerations can loom even larger. Accusing a trade partner of a violation inevitably antagonizes it, putting diplomacy at risk. For example, Japan’s foreign affairs ministry has repeatedly prevented the Japanese trade ministry from filing WTO complaints against China because diplomats fear that a WTO case would exacerbate Japan’s ongoing foreign policy conflicts with China.³ WTO disputes are also financially costly. Even the US Trade Representative—one of the most frequent WTO litigants—faces real budget constraints that limit its ability to file cases (Brutger, 2014). By contrast, the benefits from a legal challenge are often widespread. Trade rules require that any solution that results from dispute settlement be Most-Favored Nation (MFN) compliant, meaning that any concessions obtained through dispute settlement must be extended to the membership as a whole.⁴ Complainants can sometimes extract private benefits in settlements, yet the public benefits of litigation are so large that trade scholars routinely refer to WTO dispute settlement as a public good (Bown, 2005). This is usually presented as a positive aspect of the regime, yet we show that it has significant drawbacks.

How diffuse a trade policy’s impact is affects enforcement. Trade violations vary in the number of countries they affect. Some violations, like the illegitimate use of an antidumping duty, have concentrated effects that may affect only one or two countries. Other violations, like raising a tariff beyond the bound limit, have a more diffuse impact, harming all the countries that export this product to the violator country. In extreme cases, policies like US corn subsidies have such a diffuse effect that they can distort the entire world market for a good. Yet the existence of policies with a highly concentrated effect, like the US’s 1998 restrictions on Canadian trucks, mean that enforcement *can* sometimes be a private good. In such situations, countries fully internalize the benefits of their enforcement effort, making enforcement decision-making highly efficient. However, policies with a diffuse effect, like US corn subsidies, lead to under-enforcement and inefficient decision-making: individual countries consider only their own benefit from litigation and have strong incentives to free-ride on the enforcement efforts of others.

This mechanism has implications for the timing of WTO disputes. While a country with more at stake should challenge possible violations more quickly, increasing the number of countries

³Author interviews with officials from Japan’s Ministry of Economy, Trade and Industry in Tokyo on November 19, 2013. Notes on file with authors.

⁴See the WTO Dispute Settlement Understanding, Article 3.5.

with trade at stake can delay enforcement. All things equal, the more affected states there are, the more diffuse the impact of a trade policy. And because countries are less likely to challenge a trade policy with a more diffuse effect in any given period, they should wait longer, on average, to file cases against policies with more diffuse effects. Trade policies that affect only one or two countries, like antidumping duties, should be challenged more swiftly than policies that affect the entire membership, like health and safety standards. These delays can be quite significant: in Canada’s case, a concentrated policy was challenged in 15 days, while a diffuse trade policy remained in place for over 11 years. Moreover, such delays can be interpreted as the odds of a challenge in any given period. In other words, the more diffuse a violation’s effects, the less likely it is to be challenged at any given point.

Our argument also suggests that in addition to under-enforcement and excessive delays, the price of decentralized enforcement is that the “wrong cases” get filed. While litigation always carries uncertain outcomes, some cases have higher legal merit than others. Because litigation is costly, countries carefully weigh the legal merit of their claims before filing. Given how diffuse trade policies generate buck-passing behavior, an affected country must believe that a case has especially high legal merit to justify filing it. But when a trade policy has a concentrated impact, an affected country is less concerned by free-riding and is willing to file cases with relatively weak legal merit. This selection effect entails that conditional on being filed, disputes that affect fewer countries are worse cases, on average, than disputes that affect many countries. The implication is that decentralized enforcement generates inefficient outcomes: instead of spending political and financial resources on cases with broad economic effects and strong legal claims, decentralized enforcement induces states to focus on cases with narrow economic effects and relatively weak legal claims.

As is true with most domestic and international legal systems, empirical analysis of WTO dispute settlement is plagued by the dual problems of observability and endogeneity. We cannot observe all potential disputes—we only observe those disputes that trigger litigation. And litigation itself is endogenous—complainants strategically choose which cases to file. These dual challenges generate selection effects that plague valid statistical inference. Rather than being constrained by these selection effects, we use them to our advantage by leveraging the selection process inherent in our model’s equilibrium behavior. Namely, we examine the timing of observable disputes to

demonstrate that complainants are hesitant to challenge trade policies with diffuse effects, and we examine the outcomes of legal rulings to demonstrate that complainants are only willing to challenge more diffuse policies when they have stronger legal claims.

We test our arguments on all WTO disputes from 1995 to 2013. In doing so, we rely on data from Bown and Reynolds (2014), which we complement with an original dataset that codes the ruling direction on every one of the 4,484 individual claims ever brought by complainants at the WTO. The empirical evidence supports our theory: there is a price to be paid for decentralized enforcement. Cases where many stand to gain from a challenge take longer to file. And those disputes that do not face a collective action problem because they affect a single country are weaker cases, and fare less well once in court.

One implication of our argument, which we do not test in this paper, pertains to *how* countries choose to violate their trade commitments. If violations that affect a greater number of trading partners are less likely to be swiftly challenged, then a government that wants to violate its trade obligations to protect domestic constituents should avoid trade policies with concentrated effects. A government is best off when it can spread the pain around, as it were. We return to this issue and discuss other implications of our argument in the conclusion of this paper.

2 Decentralized Enforcement in International Organizations

International organizations (IOs) play a vital role in enforcing international law. As noted by legalization scholars, states are sometimes willing to delegate authority to international organizations by granting them “authority to implement, interpret and apply the rules; to resolve disputes; and (possibly) to make further rules” (Abbott et al., 2000, 401). As attested to by a vast literature, these organizations vary a great deal in their design.

Many IOs—including most international economic agreements—rely upon decentralized enforcement. For example, bilateral investment treaties (BITs) allow contract enforcement through the International Centre for the Settlement of Investment Disputes (ICSID) and other arbitral bodies (Allee and Peinhardt, 2010, 2011). All of these bodies rely upon investors to challenge and litigate possible violations. ICSID and other arbitral bodies cannot themselves sue countries for violating BITs. Similarly, the WTO—one of the most highly institutionalized IOs—relies on its

individual members to identify possible violations, construct legal arguments, and file disputes.

In contrast, other regimes feature centralized enforcement, in which the international organization itself has authority to monitor state behavior and formally challenge possible violations. The International Criminal Court (ICC), which features a prosecutorial model, is an example. While cases can be referred to the Court's attention in various ways, the ICC's Office of the Prosecutor is ultimately responsible for investigations and prosecutions. The ICC's member-states cannot file lawsuits—the ICC itself chooses which cases to pursue, constructs legal arguments, and ultimately rules on a defendant's guilt.

Finally, some international organizations, like the European Union (EU), feature elements of both centralized and decentralized enforcement. The EU allows both its member-states and the European Commission, the EU's executive body, to file suit at the European Court of Justice. Even when a case is filed by an EU member-state, the EU can play a role in constructing legal arguments via its Advocates-General (AGs). These AGs are individuals who are employed by the EU—not its member-states—to advise ECJ judges (Carrubba and Gabel, forthcoming). AGs write advisory opinions on EU law while “acting with complete impartiality and independence” from the individual litigants.⁵

Scholars usually associate greater institutionalization with stronger and more powerful institutions, like the European Union (Abbott and Snidal, 1998). Yet this relationship breaks down in practice. Even relatively weak institutions—which lack the authority to issue binding legal rulings—can have some degree of centralization. For example, each of the nine major multilateral human rights treaties has created an administrative body to monitor the behavior of its members.⁶ Each of the bodies has formal procedures for challenging possible violations of its associated treaty.⁷ These bodies all rely partly on decentralization: all of the nine treaties allow individuals to file complaints, and most allow member-states to do so as well. However, six of these nine bodies also

⁵See the Treaty on the Functioning of the European Union, Article 252.

⁶These treaties are the International Covenant on Civil and Political Rights (ICCPR), International Covenant on Economic, Social and Cultural Rights (ICESCR), International Convention on the Elimination of All Forms of Racial Discrimination (ICERD), Convention on the Elimination of All Forms of Discrimination against Women (CEDAW), Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment (CAT), Convention on the Rights of the Child (CRC), International Convention on the Protection of the Rights of All Migrant Workers and Members of Their Families (ICMW), Convention on the Rights of Persons with Disabilities (CRPD), and the International Convention for the Protection of All Persons from Enforced Disappearance (CPED).

⁷These dispute procedures were created by optional protocols. State can ratify the human right treaties without necessarily joining these optional protocols.

have centralized enforcement, allowing the treaty-based body to make its own complaints against treaty members.⁸ In sum, relatively weak institutions—like the UN’s human rights treaty-based bodies—are sometimes granted the authority to identify, investigate, and challenge possible violations, while more powerful institutions—like the WTO—lack such authority. What, then, explains the design of enforcement?

Two major factors can be said to drive the choice of enforcement design. First, states are traditionally wary of delegating authority to an IO. Centralized enforcement, in this sense, represents an extreme form of delegation. Sovereign states would rather not let unelected bureaucrats in international agencies examine their behavior and point out instances of noncompliance. In fact, this is seen as the one major reason that has kept the US from joining the International Criminal Court, where the ICC prosecutor can file cases independently. The US has cooperated with the ICC informally—by providing monetary rewards for arrest warrants and even surrendering an individual to the Court—yet it is unwilling to delegate to the ICC the power to decide whom to prosecute (Kaye, 2013). It has even pressured allies to sign non-surrender agreements that would limit the ICC’s ability to prosecute US nationals (Kelley, 2007), thus undercutting the centralized enforcement aspect of the institution from the US’ standpoint.

Second, decentralized enforcement is cost-effective, particularly when it is difficult to monitor whether states are complying with their obligations. In their study of US Congressional delegation, McCubbins and Schwartz (1984) famously distinguished between “police patrols” and “fire-alarms”. Police patrols are akin to centralized enforcement, in that they require a government agency or international organization to seek out non-compliant behavior. Fire alarms, like decentralized enforcement, rely upon individuals who are directly involved in a dispute to trigger the enforcement process. The WTO does not take on enforcement itself for the same reason that McCubbins and Schwartz (1984) argued Congress was progressively shifting to “fire alarms” in its oversight of the executive branch: going after violations in the trade regime in a centralized fashion would require an impractically large bureaucracy.

Yet decentralized enforcement also comes at a major cost: it risks generating collective action problems. Decentralized enforcement imposes a private cost on the member-state that

⁸See the ICESCR Optional Protocol Art. 11, CEDAW Optional Protocol Art. 8, CAT Art. 20, CRC Optional Protocol on Communication Procedures Art. 13(1), CRPD Optional Protocol Art. 6, CED Art. 33.

challenges a possible violation. Yet enforcement generates a public good if remedying, punishing, or deterring a legal violation benefits multiple states. Individual states that can benefit from enforcement therefore have incentive to free-ride on the enforcement effort of others (Olson, 1965).

To understand the impact of collective action problems on enforcement, consider the EU's enforcement of environmental regulation. Most EU law is driven by decentralized enforcement: private actors file cases against governments in national courts, and these national courts can then refer these matters to the European Court of Justice. In fact, the conventional wisdom among EU scholars is that European integration has been the product of a European court that has successfully enlisted national courts in its integration project (Burley and Mattli, 1993; Kumm, 2005). Yet EU environmental regulation largely relies on the European Commission for enforcement. Kelemen (2009, 49) describes how the European Commission has brought the supermajority of cases enforcing environmental laws, even as individuals can also raise environmental issues through national courts. Kelemen's explanation for this pattern echoes our own: "most environmental regulations concern matters of diffuse public interest. For this reason, coupled with the fact that litigation is costly, private parties often lack the individual incentive to commence legal action to secure enforcement." When the benefits are overly diffuse, the odds that any one actor will bring a case are reduced. As a result, the European Commission plays a larger role in enforcing environmental law than it does in other issue-areas.

Similar collective action dynamics are clearly at work in the WTO. The WTO requires that all of its members accord one another Most Favored-Nation (MFN) treatment. This is true not only of trade measures, but also of solutions to legal disputes.⁹ When a defendant removes an illegal trade barrier on a good, not only the complainant, but all countries that export that good should benefit. Of course, reality sometimes falls short of this ideal. Countries can sometimes skirt MFN treatment by writing narrowly-tailored policies or negotiating discriminatory settlements that benefit only the complainant (Davis, 2003; Johns and Pelc, 2014; Kucik and Pelc, forthcoming). However, even in such "discriminatory settlements"—which might appear to be generating purely private benefits—enforcement imposes costs on rule-violators and the expectation of these costs deters future violations. This deterrence effect benefits all states, regardless of whether they are

⁹See Article 3.5 of the Dispute Settlement Understanding: "All solutions to matters formally raised under the consultation and dispute settlement provisions of the covered agreements, including arbitration awards, shall be consistent with those agreements".

involved in an individual dispute.

Balancing the benefits and costs of decentralized enforcement suggests that the optimal design of an IO's enforcement system should depend critically on the nature of the cooperation problem that the IO solves. When the benefits of cooperation are private, there is no collective action problem. In such situations, decentralized enforcement can be optimal because it allows states to avoid sovereignty costs, reduces the cost of IO enforcement, and avoids agency problems. Yet when the benefits of cooperation are public, collective action problems mean that decentralized enforcement comes with a high price. As we go on to show, it not only delays enforcement, but leads to enforcement of the “wrong cases”.

The WTO provides the ideal laboratory for testing our argument. It is widely held to be a highly successful court, in part because of the existence of the Appellate Body, and the richness of its jurisprudence. The result is that countries are able to assess with some precision the legal merit of a case they are considering, and predict the odds of success of their legal claims in court. This ability to gauge legal merit is a precondition of our theory.¹⁰ More importantly, however, the trade regime allows us to measure each country's stake in every dispute: because disputes are over identifiable products, we can collect data on how much each country trades in the disputed commodity, and thus how much they care about any given dispute. In other words, considering the WTO allows us to assess how diffuse the trade effects are in every case with great precision. Finally, and by contrast to e.g. investment, the WTO features the same countries repeatedly making filing decisions. This comes at a considerable analytical advantage, since it allows us to keep complainant characteristics constant. We thus look to the WTO, which is routinely considered among the most successful international courts in global governance, and ask: does decentralized enforcement come at a cost?

3 Theory

We present an infinite-horizon game with discrete time periods ($t = 1, 2, \dots$). If no state ever files a case, then the game continues forever. However, if at least one state files the case, then the

¹⁰In the realm of investment, by comparison, outcomes are less transparent and less predictable. As opposed to trade, each party can appoint one arbitrator, and outcomes often hinge on the third arbitrator. There is no appeal process or higher court equivalent. Decisions are often not published. Finally, the multiplicity of different legal forums—countries can go to UNCITRAL instead of ICSID—mean that outcomes are even less predictable.

dispute goes to the WTO and the strategic interactions in our model end. This is therefore not an infinitely-repeated game since some strategy profiles allow the game to end.

3.1 Primitives and Structure

In our model, a group of n states has been harmed by a policy that restricts the trade of their goods or services. We let the parameter $\tau_i > 0$ denote player i 's trade stake, which represents the degree to which state i 's trade has been disrupted. Conditional on reaching period t , each state i learns its type for that period, α_{it} , which represents the domestic pressure on state i to challenge the trade policy. This pressure is the state's private information, and it is independently and identically distributed across players and time according to distribution function F over the interval $[\alpha_L, \alpha_H]$.¹¹ After each state learns its true type, all states must simultaneously decide whether to file a dispute at the WTO. If a state i does not file, it receives the payoff $-\alpha_{it}\tau_i$ for period t . If no state files, then the game progresses to period $t + 1$. If at least one state files, then the dispute goes to the WTO.

Because our focus is on the decision to file a case, we model the WTO dispute settlement procedure in reduced form. When the dispute goes to the WTO, all players benefit from having the case resolved. While the complainant may lose a panel ruling, it may also win a ruling or get the respondent to accept a settlement in which the policy is either partly or completely removed. Under the most-favored nation principle, all affected states—including those states that were not involved in litigation—benefit from such outcomes. We let parameter $r > 0$ represent the quality of the case, and payoff $r\tau_i$ represent state i 's expected per period payoff from having the dispute go to the WTO.

We allow the state that has filed the case to receive additional private benefits from dispute settlement. This can include discriminatory settlements, in which a portion of the respondent's concessions are targeted to benefit the complainant rather than the members as a whole, as well as other indirect benefits of litigation that only the complainant receives. We let parameter $b > 0$ represent the expected magnitude of private benefits, and payoff $b\tau_i$ represent state i 's expected per period private benefit from dispute settlement. Finally, we let parameter $k > 0$ represent the

¹¹We assume that $0 \leq \alpha_L < \frac{k}{\tau_i} < \alpha_H$. This ensures an interior equilibrium in which players are sometimes willing to file and sometimes unwilling to do so.

one-period cost of litigation for a state that files the dispute. Since WTO dispute settlement suffers from the well-known free-riding problem, we assume that the cost of litigation k is large relative to the complainant's potential stream of private benefits.¹² We assume that all states have a common discount factor, $\delta \in (0, 1)$. Finally, we let V_i denote player i 's continuation value when no state has previously filed the case; this is the state's *ex ante* expected utility from playing the game prior to learning its true type.¹³

So conditional on reaching period t —because no one has previously filed the case—each player i 's expected utility functions are:

$$\begin{aligned} EU_{it}(\text{file}|\alpha_{it}, \tau_i) &= \frac{\delta}{1-\delta} (r+b) \tau_i - k \\ EU_{it}(\text{don't file}|\alpha_{it}, \tau_i) &= -\alpha_{it} \tau_i + (1-\rho_{-i}) \frac{\delta}{1-\delta} r \tau_i + \rho_{-i} \delta V_i \end{aligned}$$

where ρ_{-i} is the probability that no other state files the case, and $1-\rho_{-i}$ is the probability that at least one other state files the case.

3.2 Equilibrium Behavior

We solve the game for the weak perfect Bayesian equilibrium.¹⁴ We first show that such an equilibrium exists and describe a property of the equilibrium:

Proposition 1. *When players are sufficiently impatient ($\delta > 0$ is small), there exists a unique weak perfect Bayesian equilibrium. In the equilibrium, each player adopts a cutpoint strategy: conditional on reaching period t , high types will file the case and low types will not file.*

Payoffs in our game are streams over an infinite number of periods, so we must constrain the value of the discount factor (δ) to identify a reasonable equilibrium. To understand why this is necessary, assume the opposite: suppose that players are extremely patient and δ is close to one. Then the stream of private benefits from litigation $\left(\frac{\delta}{1-\delta}b\right)$ will dramatically outweigh the short-term cost of litigation (k) and all players will want to immediately file the case. Given the

¹²If the opposite were true— k were small relative to $\frac{\delta}{1-\delta}b\tau_i$ —then no free-riding problem would exist. We would expect to see an abundance of enforcement via litigation.

¹³Because α -parameters are independently distributed, V_i is not a function of time.

¹⁴This solution concept requires that strategies are sequentially rational and beliefs are consistent with Bayes' Rule where possible. Since each player's type is independent across time and decisions in any period t are simultaneous, we do not need to specify off-the-equilibrium-path beliefs.

significant delay in filing a WTO dispute, we find this equilibrium behavior implausible. The more plausible equilibrium is one in which the short-term cost of litigation deters states from filing a case, which means that states place relatively low weight on future payoffs.

Because players adopt cutpoint strategies, in which high types file and low types do not, we can examine the impact of various factors on the likelihood that a case is filed by examining changes in the cutpoint value, which is the type of player that is indifferent between filing and not filing.

3.3 Comparative Statics

When a state has relatively little trade at stake in a dispute, both the potential benefit of filing the case and the cost of the policy remaining unchallenged are small. Both factors reduce the incentive of a player to bear the short-term cost of litigation. Not surprisingly, a state is more likely to challenge a trade policy when that policy has a larger impact on the state.

Proposition 2. *Each player is more likely to file the case if its own trade stake increases.*

However, a state's incentives are different when a policy has a larger impact on the trade stake of another state. As shown in the expected utility functions above, an individual player does not directly care about the trade stake of other players. However, the trade stakes of other players indirectly affect a player's decision-making by changing its beliefs about how other players will behave. Players want to receive the long-term expected benefit of the policy being challenged ($\frac{\delta}{1-\delta}r$), but do not want to bear the short-term cost of litigation (k). So a state is less likely to file the case itself when another player becomes more likely to file, and this occurs when the trade stake of another player increases. This indirect effect means that a state is less likely to file a case when the trade stake of another player increases.

Proposition 3. *Each player is less likely to file a case if the trade stake of another player increases.*

One additional factor that affects individual decision-making is the overall quality of the case, r . Since all states receive an infinite stream of r -payoffs when the case is filed—regardless of whether they themselves filed it—the expected utility from both filing and not filing both increase when a case is of higher-quality. However, the payoffs change at different rates. When a player files the case itself, it knows that its future stream of payoffs is increasing as r increases. However,

if a player does not file the case, an increase in r only increases his own payoff if someone else files the case. This means that increasing the quality of a case makes filing the case more attractive relative to not filing the case.

Proposition 4. *As the quality of the case increases, each player is more likely to file the case in any given period.*

While Propositions 2 and 3 are suggestive, neither adequately addresses the impact of diffusion because increasing the trade stake of an individual player corresponds to increasing the total impact of the trade policy on all players, $\tau = \sum_i \tau_i$. To isolate the role of diffusion, we must control for the total impact of the trade policy in both our theoretical model and our empirical tests. We therefore impose an additional assumption for the next two results: we assume that each individual player's trade stake is an equal share of the total trade stake, $\tau_i = \frac{\tau}{n}$. When there are few players (small n), the overall impact of the trade policy is concentrated. However, as the number of players increases, the total impact of the trade policy is spread across more players, making it less concentrated. This approach allows us to control for the total economic impact of the policy because we can hold parameter τ fixed when taking comparative statics.

When the impact of a trade policy is spread across more players, each player's individual trade stake decreases. This diffusion exacerbates the collective action problem that is inherent in the strategic interaction. Each individual player is more tempted to free-ride on the efforts of other players, and therefore is less likely to file the case itself.

Proposition 5. *As the number of players increases, each player is less likely to file the case in any given period.*

However, this individual-level effect does not necessarily extend to the collective outcome—whether *someone* files a case. At the individual-level, diffusion makes each player less likely to file. Yet diffusion also increases the number of actors that could potentially file the case. Which effect is dominant—the individual versus the collective—depends on the cost of litigation (k). Consider what happens if we increase the number of states affected by a trade policy from n to $n+1$. Because this change spreads the economic impact of the trade policy amongst more players, it decreases the marginal benefit of filing the case for the original n states. If the cost of litigation is relatively small, then the probability that the new state files the case is relatively high, offsetting the decrease

in the probability that the original n states will file. However, as the cost of litigation grows larger, the new state becomes less likely to file the case, meaning that negative impact of diffusion on the original n players outweighs the effect of increasing the number of states that can file.

Proposition 6. *When the cost of litigation is large, as the number of players increases, the overall probability that the case is filed by at least one state decreases.*

This behavior is illustrated in Figure 1. Panel (a) shows that as the number of affected states increase, the probability that an individual state files the case is decreasing. Panel (b) shows that the same relationship holds for the collective outcome: as a trade policy becomes more diffuse, the probability that *someone* files the case decreases.

[Insert Figure 1 here.]

3.4 Empirical Implications

Under an ideal research design, we could identify all possible trade violations, observe which policies are challenged in WTO dispute settlement, and ask whether diffuseness matters. However, we can only observe cases that are actually filed. Yet our model holds two major empirical implications—concerning the delay in enforcing trade obligations and the legal merit of a case—that can be tested by exploiting data on these observed legal challenges.

First, our model results above are stated in terms of the likelihood that a dispute is filed in a given period. Yet the fact that we have created an infinite-horizon model means that we can make meaningful inferences about duration; namely, how long we expect it to take countries to challenge a violation at the WTO. We refer to this as “enforcement delay” in our empirical analysis. Any exogenous variable that decreases the likelihood that the case is filed in a given period also increases the expected enforcement delay. Because violations with more diffuse effects are less likely to be challenged in a given period, our model implies that conditional on a case being filed, the enforcement delay will be longer.

Hypothesis 1. *In observable WTO disputes, enforcement delay—the time between when a trade policy is implemented and when it is challenged at the WTO—will be longer for trade policies with more diffuse effects.*

Second, our model demonstrates that diffusion and case quality have opposite effects on filing decisions. While diffusion reduces the likelihood that a case is filed, case quality increases this likelihood. Every state must balance the expected benefit from litigation against its expected cost. Suppose that a given player is indifferent between filing and not filing the dispute. If diffusion increases, then the expected benefit of filing decreases. If we wish to offset this effect to ensure that the player remains indifferent, we must increase the quality of the case. This implies that conditional on being filed, a case that challenges a policy with diffuse effects should be of higher quality, on average, than a case that challenges a policy with concentrated effects. While we cannot observe player's *ex ante* beliefs about the quality of potential cases, we can observe the outcome of WTO rulings. If the perceptions of players are correct—that is, if higher quality cases are more likely to generate pro-complainant rulings—then cases that challenge diffuse policies should be more likely to generate legal victories for the complainant than cases that challenge more concentrated policies.

Hypothesis 2. *In observable WTO disputes, complainants are more likely to prevail at trial when they challenge policies with a diffuse (rather than a concentrated) effect.*

Note that both of these conditional hypotheses are driven by the selection effects in our model's equilibrium behavior. Because we cannot observe disputes that are not filed at the WTO, we cannot test our arguments directly by examining the impact of diffusion and case quality on the likelihood that a case is filed. However, the selection effects that drive equilibrium behavior have implications for those challenges that are observed, allowing us to test our theory indirectly. By constructing conditional hypotheses, we can leverage the endogeneity inherent in international disputes, rather than being constrained by it.

4 Empirics

Testing our two hypotheses requires two distinct datasets. Our first dataset is built at the dispute-country level, and includes data on the trade stake of every WTO member in every WTO dispute since 1995. Our second dataset is constructed at the dispute level, and considers the number of claims the complainant won in each WTO dispute. Next, we describe the data and the estimations that we use to test our two main claims, which together address the question: what are the

consequences of decentralized enforcement at the WTO?

4.1 How Does the Diffusion of Benefits Impact Enforcement Delays?

To test our first hypothesis, we need to measure the ENFORCEMENT DELAY—the time it took for a complainant to file each WTO dispute. The recently coded data in Bown and Reynolds (2014) includes the precise implementation date for the trade policy underlying every WTO dispute, stretching back into the GATT era. This is the date on which the measure being challenged—be it an antidumping duty, a tariff increase, new labeling standards, or an embargo—first went into effect in the country at issue. We compare this implementation date to the date on which each WTO dispute was filed. This allows us to measure, in days, exactly how much time elapsed between the start of a trade policy and its challenge.

We test our theoretical argument using three measures of diffusion on the right-hand side of the equation. The first two measures—NUMBER OF COUNTRIES AFFECTED and DISPUTED TRADE FLOWS HHI—capture economic diffusion by examining trade flows. The third measure—MULTILATERAL VIOLATION—captures diffusion through a legal angle, by examining whether the type of discrimination being challenged is targeted at a few trading partners or whether it affects the entire membership.

Our first measure of diffusion—NUMBER OF COUNTRIES AFFECTED—is the most straightforward. For each dispute, we identify which products are affected by the trade policy that is being challenged. We then count the number of countries with any trade at stake in the year that the complainant initiated the dispute, or the two preceding years. While some challenged policies concern products that are traded by a great number of countries, like corn, others are exported by a handful of countries, like commercial ships. In our sample, the count variable ranges from 1 to 128 countries. Given that this count is so widely distributed, NUMBER OF COUNTRIES AFFECTED denotes the log of the number of countries affected by the trade policy at dispute.

While our first measure of diffusion captures how many countries have trade at stake, it tells us little about how trade is distributed across these countries. It could be, for instance, that a dispute with 128 interested parties only has one country that exports non-trivial amounts of the disputed commodity. In such a case, as per Olson (1965), the collective action problem is diminished: we would expect the one country that has disproportionately more at stake to front

the costs of enforcement for everyone. In sum, beyond the number of countries with something at stake, the distribution of trade among them matters. Our second measure of diffusion—DISPUTED TRADE FLOWS HHI—accordingly measures the distribution of a trade policy’s economic impact. Namely, we construct a Herfindahl-Hirschman index (HHi) measure of trade flows into the defendant country. The HHi measure is usually used to summarize the market structure of an industry in a single figure. It varies from 0 to 1, with 0 indicating a perfectly competitive market with a large number of small actors, and 1 indicating a monopoly with a single actor.¹⁵ If only two countries exported into the defendant country, and if they exported the same amount, the HHi measure would be 0.5. In our usable sample, the HHi measure varies from 0.04 to 1. The greater (smaller) this variable, the more concentrated (diffuse) trade in the disputed product is. According to our theory, we would expect that the greater DISPUTED TRADE FLOWS HHI for a dispute, the faster it will be filed.

Both the NUMBER OF COUNTRIES AFFECTED and DISPUTED TRADE FLOWS HHI variables rely on bilateral trade flow data from the World Integrated Trade Service, which is maintained by the World Bank. This inevitably leaves out non-merchandise disputes for which we cannot quantify the amount of trade. For example, when a group of WTO members successfully challenged Section 301, a piece of legislation that the US used to coerce other countries to amend their policies, there was no underlying traded product, even though the legislation had widespread consequences on the trade regime as a whole.¹⁶ We code such cases as missing for the purpose of our first two diffusion measures, but our third diffusion measure allows us to analyze both merchandise and non-merchandise disputes.

At their core, trade disputes involve some form of discrimination. Broadly speaking, the WTO requires that its members meet two different treatment standards: national treatment and most-favored nation (MFN) treatment. National treatment requires that each WTO member treat foreign imports no less favorably than the comparable domestic good. When a policy violates the national treatment standard, the entire membership is concerned: all foreign exporters are hurt,

¹⁵The HHi measure is calculated as

$$HHi = \sum_{i=1}^n f_i^2$$

where f_i is the amount of trade from country i to the defendant country, and n is the number of countries with non-zero trade to the defendant country.

¹⁶See DS152: *United States—Section 301 Trade Act*.

and thus all stand to benefit from a legal challenge of the policy. In contrast, most-favored nation treatment keeps WTO members from discriminating among different partners and favoring some over others. By definition, trade policies that violate the MFN standard only harm a subset of WTO members. Trade policies that violate the MFN standard therefore have a relatively concentrated impact, while policies that violate the national treatment standard have a more diffuse impact.

Our third measure of diffusion—`MULTILATERAL VIOLATION`—captures such variation in how a trade policy is applied. This dichotomous variable comes from Bown and Reynolds (2014) and indicates whether a trade policy affects a large number of countries, or only a small subset. About 49% of the disputes in our sample are coded as concerning multilateral violations. For example, most safeguards are applied to imports from all countries and are thus multilateral measures affecting all trade partners equally. However, some safeguards are targeted at only a small subset of countries, making them non-multilateral. Bown and Reynolds include such distinctions in their coding.¹⁷ Despite these precautions, the resulting variable, `MULTILATERAL VIOLATION`, remains imperfect: the number of countries affected by violations is a continuous concept, and reducing it to a dichotomous measure conceals some of this variation. Yet it remains a useful way of thinking about the effect of violations. In combination with our other two measures of diffusion, it provides a fuller sense of whether the trade policy that is being challenged has a diffuse or concentrated impact. Even without running our survival model estimates, the descriptives for `MULTILATERAL VIOLATION` are telling: on average, multilateral cases in our sample take over 80% longer to be filed.

Additionally, we include some controls. We account for the total size of the trade at stake in every dispute, and for every individual country’s stake, since the dataset is at the dispute-country level. Both of these variables are coded as logs of the trade flows concerned. We also add a variable indicating the year in which a dispute was initiated, to control for any trend in time that would affect the swiftness with which violations are challenged.

Using these variables, we begin by estimating a Cox proportional hazards model, shown in Table 1. Recall that the data are at the dispute-country level of observation: they include information not only about the country that eventually filed, but also about all the countries that

¹⁷Bown and Reynolds (2014) code all disputes according to whether they are over “global” or “partial” legal violations. We use the term “multilateral” to mean what they intend by the term “global”.

did not. We thus right censor any countries that did not challenge the violation, to account for how, had the violation not been challenged by the eventual complainant, these countries may still have done so. Related to this point, some violations are challenged by more than one party, and often not simultaneously. In the aforementioned case when Canada filed against US subsidies on corn, Brazil followed suit 7 months later, challenging the same agricultural measures, but also adding one additional legal claim. As is most often the case in disputes over the same set of measures by the same defendant, the Dispute Settlement Body chose to create a common panel, which would eventually produce a single ruling assessing both complainants' claims. We account for the common characteristics of such violations by clustering standard errors on the common violation.

Recall that we are limited in our assessment of economic diffusion to those disputes where we can observe trade flows in the disputed product. This leaves out a large number of disputes that concern non-merchandise issues. Yet these still vary along our variable capturing the legal aspect of diffusion, `MULTILATERAL VIOLATION`. We thus begin our analysis with a univariate Cox model that only considers the relation between `MULTILATERAL VIOLATION` and the likelihood the measure is challenged in any given period, shown in Column 1. This exploits our maximal sample.

In the second column of Table 1, we add our economic measures, including our two economic measures of diffusion, `DISPUTED TRADE FLOWS HHI` and `NUMBER OF COUNTRIES AFFECTED`. In the third column, we also adjust the Cox estimation for shared frailty of the defendant country. Shared-frailty survival models are used to model within-group correlation. It might be that some countries have some unobserved quality that makes them more likely to be swiftly challenged, apart from the amount of trade at stake, which we already account for. Accounting for shared frailty is akin to accounting for the panel aspect of time-series data. The salience of a country within the trade regime, its perceived reputation, or other latent qualities affecting the likelihood that a given country is challenged would thus be captured by this adjustment, which adds a respondent-specific frailty parameter to the equation. Looking at descriptive statistics suggests that there is considerable variation in the time it takes different countries to be challenged. Some, like Japan, are swiftly taken to task (in an average of 143 days), while others, like Australia, seem to take much longer (2,871 days). The two superpowers fall somewhere around the membership median in this respect (1,654 days for the US and 1,155 days for the EU). It is worth noting that such descriptives are prone to selection—they do not tell us about the violations that have yet to be challenged, and

that have likely been around for longest—and should thus be interpreted with care. Nonetheless, they are enough to suggest the utility of a frailty parameter.

[Insert Table 1 here.]

The results, shown in Table 1, offer support for the theory. We convert hazard ratios into coefficients to make the results easier to read. A negative coefficient represents a decrease in the hazard function, meaning that a challenge is less likely in any given period. The type of legal discrimination at issue has an unambiguous effect: throughout our different estimations, MULTILATERAL VIOLATION is associated with a decreased likelihood of a legal challenge in any given period. The effect is highly substantively significant. In the first model, the rate of legal challenge decreases by 45.3% when the legal violation at issue is of a multilateral nature.

The indicators of economic diffusion behave as expected, too. The greater the number of countries with a stake in challenging the violation, the less likely such a challenge becomes in any given period. As expected, our HHi indicator carries the opposite sign: the higher the HHi, meaning the more concentrated trade flows are across exporters, the shorter the expected delay before a policy is challenged. Both indicators thus support the same belief about the concentration of benefits: the more countries stand to gain from a challenge, the less likely it becomes at any given point, and the longer it takes, on average.

Our control variables present one unexpected result. The total stake of the violation is consistently negatively related to the hazard rate, meaning that all else equal, violations over more trade flows are less likely of being filed in any given instance. This relationship is highly robust, and also holds in a univariate estimation. The relationship is all the more striking since our economic diffusion variables already account for some measure of the collective action problem at the heart of our theory. The total trade stake variable is thus capturing additional collective action effects, or the daunting effect of taking on large trade flows. One thing is certain, however: under decentralized enforcement, more serious violations, as proxied for by the total underlying trade flows, are less likely of being challenged in any given period. Our dispute-country level indicator, which is highly significant, behaves in a more intuitive way: the more an individual country has at stake, the greater the hazard rate for that country, meaning that it becomes more likely to become a complainant in a dispute over the underlying violation. In sum, enforcement choices

appear individually, if not socially, rational. Our control for initiation year, by comparison, shows no significant effect.

We are aware of the likely (negative) correlation between our two measures of economic diffusion. The greater the number of countries with something at stake, the smaller the HHi index is likely to be, and this is reflected in the negative bivariate correlation between the two variables of -44.6%. To ensure that this is not affecting our results, in columns 4-6 of Table 1, we rerun our main estimation with clustered standard errors on the common dispute including a single one of our diffusion indicators at a time. The results remain: the more concentrated the effect of a violation, measured by any of our three economic and legal variables, the less likely a challenge in any given period becomes.

4.2 Does the Diffusion of Benefits from Litigation Affect Dispute Outcomes?

Now that we have found support for our beliefs about which violations are challenged under decentralized enforcement, we move on to test the implications of this selection process on the legal merit of disputes. The simplest way of assessing legal merit is to look at the outcome of rulings. On average, we should expect disputes with greater legal merit to win more claims. It is well known that WTO disputes display a pro-complainant bias: most rulings find some violation at play. Indeed, this fact is consistent with our theory: international trade rules are sufficiently consistent that countries are able to observe legal merit, and choose to challenge those violations that have the highest likelihood of success before judges. It follows that countries select meritorious cases on average, since they would rather not expend political and financial resources pursuing a dispute they might later lose. Yet this pro-complainant bias conceals quite a bit of variation, since complainants file a number of claims in a given dispute. We construct an original dataset to assess the success of each case in a continuous fashion, by looking at the number of claims won.

To do so, we code the direction of every claim in every WTO dispute from 1995 to 2013.¹⁸ This is a considerable coding exercise, as complainants have brought 4,484 such claims over the WTO's history. Of these, a minority are actually ruled on,¹⁹ though some claims receive more

¹⁸A "claim" is an alleged violation of a given Article or sub-article of the WTO texts. These are taken directly from the complainants' request for consultations. There is thus no room for ambiguity in the coding.

¹⁹For instance, all claims alleging any type of national treatment discrimination will contain a claim under GATT III, yet the panel rarely actually rules on GATT III, and more often delivers more specific rulings under a specific agreements, such as e.g. the Agreement on Sanitary and Phytosanitary (SPS) Measures.

than one finding. All told, panels have delivered 1,429 findings on 820 individual claims. We first collapse these findings at the claim level, and then collapse claims at the dispute level, to obtain a variable indicating the number of claims won by the complainant.

We do the same for the Appellate level. These days, most rulings are appealed, and the Appellate Body frequently overturns panel rulings. The AB is a standing body, as opposed to panels, where judges are chosen *ad hoc*, and the AB's pronouncements are thought to have greater authority and be more attentive to the WTO's jurisprudence than panel rulings. As a result, when assessing the legal merit of a dispute by looking at the direction of rulings, we are interested in rulings "net of appeal". The resulting variable CLAIMS WON NET OF APPEAL corresponds to the claims won at the panel, as modified (or not) by the AB, in the case of an appeal. In the absence of an appeal, this measure is simply the number of claims won at the panel stage. CLAIMS WON NET OF APPEAL is our main dependent variable in this second set of tests, and we express it both as a number and as a percentage.

Our expectation is that owing to the selection process, and the collective action problem therein, the higher "threshold" for challenges of multilateral violations, or violations with more diffuse economic consequences, is likely to positively affect the legal merit of those claims that *are* filed on. Disputes with less concentrated enforcement benefits, in other words, should be better cases. On the right hand side of the equation, our main explanatory variables are thus as described above in Section 4.1: two economic measures of diffusion, DISPUTED TRADE FLOWS HHI and NUMBER OF COUNTRIES AFFECTED, and one legal measure, MULTILATERAL VIOLATION, indicating the nature of the violation. We also control for the total amount of trade at stake, since judges may be swayed by the importance of the case in their verdicts. Finally, we include a year indicator to account for any trends in time: with the accumulation of jurisprudence over time, it might be that the rate of pro-complainant findings is affected, though we remain agnostic about the direction of this effect. A simple descriptive statistic offers some early support for our beliefs: the percentage of claims ruled in favor of the complainant in multilateral disputes is 74%, compared to 65% for non-multilateral cases, and the difference is highly significant.

Our estimation of the number of claims ruled on in favor of the complainant takes the form of a Heckman probit selection model. Selection models benefit from simplicity, and thus we keep our estimations as parsimonious as possible. The purpose is to account for the selection of cases

that reach the ruling stage. We know that a little more than half of all disputes never make it to a ruling; this risks biasing our results. We thus begin by estimating the odds of a ruling in a first stage equation, and use those estimates in our second stage, outcome equation.

To identify our model, we rely on a measure of the number of third parties in the room. There is a rich literature testifying to how the presence of third parties decreases the odds of settlement, and increases the odds of litigation, a belief that has been repeatedly empirically validated. Yet the presence of third parties in the room should not have an effect on the direction of the ruling. Indeed, the bivariate correlation between our outcome variable, CLAIMS WON NET OF APPEAL, and NUMBER OF THIRD PARTIES is -0.03, suggesting that there is little to no relationship between the two, as theory would suggest. The other variable in the first stage equation is TOTAL TRADE AT STAKE, which is the log of the defendant’s total imports of the disputed product. The expected direction of this variable is also unambiguous: we would expect that the more there is at stake, the less likely a settlement, and the more likely a ruling becomes. Anecdotal evidence suggests that in very large stakes, defendants cannot allow themselves to concede for domestic reasons without the “help” of an unfavorable ruling. We also include this variable in the second stage, as politically savvy judges may be swayed by the size of the stakes in a case.

In the second column, we include a few additional control variables. We add standard controls for economic power—COMPLAINANT GDP (LOGGED) and DEFENDANT GDP (LOGGED), as these may play a role both in the outcome of litigation. In addition, as another proxy for power, we include a SUPERPOWER DEFENDANT dummy, coded as 1 if the defendant is either the United States or the European Union, as recent findings have suggested that politically savvy judges try and render compliance easier for the regime’s two superpowers (Brutger and Morse, 2013), even once market power is controlled for. Given how our dependent variable is a “net” number of claims won, we also include an indicator of whether the dispute was appealed, in case there is any difference in the rate of pro-complainant findings among panel judges vs. AB judges. As always, we cluster our robust standard errors on the common dispute, to allow for the fact that some disputes are filed over the same violation.

The number of claims won, the dependent variable in models (1) and (2) in Table 2, is our preferred measure, since it offers the best sense of the magnitude of the win. Indeed, recent work emphasizes how complainants work hard to win a maximum number of claims, and how defendants

work equally hard to be shown in violation of as little legal claims as they can, even once they know it is clear their policies are non-compliant Busch and Pelc (2010). However, since the total number of rulings varies across disputes, we rerun our regression using a different operationalization of our dependent variable, as a percentage of claims won by the complainant. Another advantage of this operationalization is that it aids in interpretation. The right-hand side of model (3) remains the same as in model (2), and only the operationalization of the dependent variable changes.

[Insert Table 2 here.]

Our beliefs about the relationship between the concentration of benefits and legal merit are broadly supported in our three selection models, shown in Table 2. The greater the number of countries with something at stake, the more successful the case, on average. Conversely, the more concentrated trade across those countries, the worse the prospects of the case. And disputes over multilateral cases fare better on average, except when looking at our model (3), which estimates the percentage of claims won. There, although it remains positive, the coefficient for MULTILATERAL VIOLATION loses significance. Yet even in model (3), the three concentration variables, taken together, remain jointly significant. In short, disputes where the benefits of enforcement are more concentrated appear to be worse cases.

These effects are substantively important. The third model, with the dependent variable measured as proportion of claims won, offers the easiest interpretation. An average case with concentrated benefits results in less than half (45%) of claims being found in favor of the complainant. Compare this to an average case with diffuse benefits, which results in 73% of claims ruled in favor of the complainant.²⁰ This represents a 62% jump in the success rate of disputes, according to their economic concentration. Cases with diffuse consequences are better cases, on average.

The total amount of trade at stake has a highly significant impact on the odds of litigation, in accordance with intuition. It also has a positive, though less significant, impact on the success rate of cases, though this effect becomes insignificant when using the Percent of Claims Won dependent variable. The market power and superpower controls appear to explain little of the outcomes of litigation. Importantly, and as expected, the variable that identifies the model,

²⁰We run these calculations by relying only on the two economic concentration variables, since MULTILATERAL VIOLATION falls short of significance in Model (3). We take a dispute with concentrated (diffuse) benefits to be on that ranks one standard deviation higher (lower) on HHi, and one standard deviation lower (higher) on the Number of Countries Affected.

NUMBER OF THIRD PARTIES, bears a strong positive relationship with the odds of litigation. Finally, our reliance on the selection model is itself confirmed. The Wald test allows us to reject the hypothesis that both our equations are independent, and as theory would suggest, supports the use of a selection stage. Rulings are not randomly assigned, and our first stage equation allows us to correct for this selection.

5 Conclusion

In recent years, the effort in the legal community to improve enforcement in trade has been directed at increasing the legal capacity of countries, especially that of developing countries (Busch, Reinhardt and Shaffer, 2009; Bown, 2005; Bown and Hoekman, 2005). What our findings suggest is that such doubling-down on decentralized enforcement might be missing the point: if the objective is to target those violations that distort the world economy the most, increasing legal capacity may not be a sufficient answer.

This paper highlights how decentralized enforcement, where countries' policies are only called into question if one government formally challenges another, comes at a high cost. When we compare the record of enforcement at the WTO against a hypothesized ideal where the most egregious violations are the ones being challenged, we come up against two shortcomings. The first is that those cases that affect the greater number—as measured by the number of countries that have any trade at stake; the distribution of trade flows across those countries; and the legal nature of the violation at issue—are also the least likely of being pursued in any given period. The second finding is that this selection of cases has repercussions on the legal merit of the disputes that *do* arise. When countries challenge disputes that affect them in a narrow way, they often do so at the expense of legal merit, as compared with those cases that affect many. The result is that the “wrong cases” are being filed. Taken together, these findings represent the cost of decentralized enforcement.

There is good reason to believe that this cost will continue to rise in the future. The WTO is progressively shifting towards a regulatory focus, reflecting the changing nature of violations in the regime. Indeed, protectionism in the 21st century takes the form not of trade remedies like antidumping, which are usually targeted at a specific state, but of *standards*: these include

measures such as labeling requirements, health and safety standards, and environment regulation (Kono, 2006). What standards have in common is that their effects are highly diffuse, affecting all trading partners alike. Drawing on our findings, all things equal, we might expect that under-enforcement will become an even greater problem as the type of protection facing the trade regime shifts over to standards.

One striking implication of our findings, which we leave for further research, is that we are correct about the drivers of enforcement delays, then countries have a clear incentive in how they design trade protection. To avoid enforcement, countries should “spread the pain around” as much as possible when providing import relief to domestic industries. The more countries are affected by the same violation, the less likely a challenge becomes, and the more likely delays in enforcement are.

In general, the lesson of our theory is that institutional design should match the nature of the underlying cooperation problem. When enforcement generates a private good, it ought to be paired with a decentralized design. Yet when it generates a public good, as the literature unanimously agrees is the case in the trade regime, enforcement gains from being more centralized. The WTO already has institutional features that could be built on to fill in where decentralized enforcement falls short. Trade Policy Reviews (TRPs), in particular, are the only means through which the WTO Secretariat points out shortcomings in the trade policies of its member states. Yet TPRs are highly infrequent and largely neglected. Our findings suggest that such policy reviews may be playing an undervalued role within the trade regime, and may gain from being more institutionalized.

We hold no illusions over the feasibility of changing the WTO’s institutional design. Decentralized enforcement has been hardwired into the institution since the creation of the GATT in 1947, and for good reason. Rather, our point is that relying exclusively on decentralized enforcement leaves the institution vulnerable to collective action problems. What this article suggests is that the price being paid is a high one.

Appendix

Proof of Proposition 1. Conditional on reaching period t , player i 's expected utility functions in period t are:

$$\begin{aligned} EU_{it}(\text{file}|\alpha_{it}, \tau_i) &= \frac{\delta}{1-\delta} (r+b) \tau_i - k \\ EU_{it}(\text{don't file}|\alpha_{it}, \tau_i) &= -\alpha_{it} \tau_i + (1-\rho_{-i}) \frac{\delta}{1-\delta} r \tau_i + \rho_{-i} \delta V_i \end{aligned}$$

So player i has incentive to file iff:

$$\begin{aligned} \frac{\delta}{1-\delta} (r+b) \tau_i - k &\geq -\alpha_{it} \tau_i + (1-\rho_{-i}) \frac{\delta}{1-\delta} r \tau_i + \rho_{-i} \delta V_i \\ \Leftrightarrow \alpha_{it} &\geq \frac{k}{\tau_i} - \frac{\delta}{1-\delta} b - \rho_{-i} \frac{\delta}{1-\delta} r + \frac{\delta \rho_{-i}}{\tau_i} V_i \equiv \bar{\alpha}_i \end{aligned}$$

In equilibrium:

$$\rho_i = \Pr(\alpha_{it} < \bar{\alpha}_i) = F(\bar{\alpha}_i) \quad \text{and} \quad \rho_{-i} = \prod_{j \neq i} \rho_j = \frac{\prod_k F(\bar{\alpha}_k)}{F(\bar{\alpha}_i)} \quad \text{and} \quad \rho = \prod_k \rho_k = \prod_k F(\bar{\alpha}_k)$$

Player i 's continuation value—her *ex ante* expected utility from playing the game—is therefore:

$$\begin{aligned} V_i &= \int_{\alpha_L}^{\bar{\alpha}_i} \left[-\alpha \tau_i + (1-\rho_{-i}) \frac{\delta}{1-\delta} r \tau_i + \rho_{-i} \delta V_i \right] f(\alpha) d\alpha + \int_{\bar{\alpha}_i}^{\alpha_H} \left[\frac{\delta}{1-\delta} (r+b) \tau_i - k \right] f(\alpha) d\alpha \\ &= \rho_i \left[(1-\rho_{-i}) \frac{\delta}{1-\delta} r \tau_i + \rho_{-i} \delta V_i \right] + (1-\rho_i) \left[\frac{\delta}{1-\delta} (r+b) \tau_i - k \right] - \tau_i \int_{\alpha_L}^{\bar{\alpha}_i} \alpha f(\alpha) d\alpha \\ &= \frac{1}{1-\delta \rho} \left[(1-\rho) \frac{\delta}{1-\delta} r \tau_i - (1-\rho_i) \left(k - \frac{\delta}{1-\delta} b \tau_i \right) - \tau_i \int_{\alpha_L}^{\bar{\alpha}_i} \alpha f(\alpha) d\alpha \right] \end{aligned}$$

This means that:

$$\begin{aligned} \bar{\alpha}_i &= \frac{k}{\tau_i} - \frac{\delta}{1-\delta} b - \rho_{-i} \frac{\delta}{1-\delta} r \\ &\quad + \frac{\delta \rho_{-i}}{\tau_i (1-\delta \rho)} \left[(1-\rho) \frac{\delta}{1-\delta} r \tau_i - (1-\rho_i) \left(k - \frac{\delta}{1-\delta} b \tau_i \right) - \tau_i \int_{\alpha_L}^{\bar{\alpha}_i} \alpha f(\alpha) d\alpha \right] \end{aligned}$$

And the cutpoint $\bar{\alpha}_i$ is implicitly defined by:

$$\Psi^i \equiv \bar{\alpha}_i(1 - \delta\rho) - (1 - \delta\rho_{-i}) \left(\frac{k}{\tau_i} - \frac{\delta}{1 - \delta} b \right) + \delta\rho_{-i}r + \delta\rho_{-i} \int_{\alpha_L}^{\bar{\alpha}_i} \alpha f(\alpha) d\alpha = 0$$

where: $\frac{\partial \Psi^i}{\partial \bar{\alpha}_i} = \frac{\partial \Psi^i}{\partial \bar{\alpha}_i} = \bar{\alpha}_i(-\delta\rho_{-i}f(\bar{\alpha}_i)) + (1 - \delta\rho) + \delta\rho_{-i}\bar{\alpha}_i f(\bar{\alpha}_i) = 1 - \delta\rho > 0$

and: $\lim_{\delta \rightarrow 0} \Psi^i = \bar{\alpha}_i - \frac{k}{\tau_i} \Rightarrow \lim_{\delta \rightarrow 0} \bar{\alpha}_i = \frac{k}{\tau_i}$

These results ensure that player i has an interior cutpoint, $\bar{\alpha}_i \in (\alpha_L, \alpha_H)$. Since this argument holds for an arbitrary player i , there exists a Bayesian Nash equilibrium characterized by the vector of equilibrium cut points, $\bar{\alpha} = (\bar{\alpha}_1, \bar{\alpha}_2, \dots, \bar{\alpha}_n)$, which is implicitly defined by the system of equations $(\Psi^1, \Psi^2, \dots, \Psi^n)$. \square

The following results are useful in proving the comparative statics of the model.

Lemma 1. *In equilibrium, $\bar{\alpha}_i < \frac{k}{\tau_i}$.*

Proof of Lemma 1. By the calculations in the proof of Proposition 1, $\Psi_{\bar{\alpha}_i}^i > 0$. Note that:

$$\Psi^i \left(\bar{\alpha}_i = \frac{k}{\tau_i} \right) = \delta\rho_{-i}(1 - \rho_i) \frac{k}{\tau_i} + (1 - \delta\rho_{-i}) \frac{\delta}{1 - \delta} b + \delta\rho_{-i}r + \delta\rho_{-i} \int_{\alpha_L}^{\frac{k}{\tau_i}} \alpha f(\alpha) d\alpha > 0$$

So the equilibrium value that solves $\Psi^i(\bar{\alpha}_i) = 0$ must be less than $\frac{k}{\tau_i}$. \square

Some partial derivatives:

$$\begin{aligned} \frac{\partial \Psi^i}{\partial \bar{\alpha}_i} &= 1 - \delta\rho > 0 \Rightarrow \lim_{\delta \rightarrow 0} \frac{\partial \Psi^i}{\partial \bar{\alpha}_i} = 1 \\ \frac{\partial \Psi^i}{\partial \bar{\alpha}_j} &= \delta f(\bar{\alpha}_j) \frac{\rho}{\rho_i \rho_j} \left[\frac{k}{\tau_i} - \frac{\delta}{1 - \delta} b + r + \int_{\alpha_L}^{\bar{\alpha}_i} \alpha f(\alpha) d\alpha - \bar{\alpha}_i \rho_i \right] \Rightarrow \lim_{\delta \rightarrow 0} \frac{\partial \Psi^i}{\partial \bar{\alpha}_j} = 0 \\ \frac{\partial \Psi^i}{\partial \tau_i} &= (1 - \delta\rho_{-i}) \frac{k}{\tau_i^2} > 0 \Rightarrow \lim_{\delta \rightarrow 0} \frac{\partial \Psi^i}{\partial \tau_i} = \frac{k}{\tau_i^2} > 0 \\ \frac{\partial \Psi^i}{\partial \tau_j} &= 0 \\ \frac{\partial \Psi^i}{\partial r} &= \delta\rho_{-i} > 0 \end{aligned}$$

The Jacobian for the n -player game is:

$$J_n = \begin{bmatrix} \Psi_{\bar{\alpha}_1}^1 & \dots & \Psi_{\bar{\alpha}_n}^1 \\ \dots & \dots & \dots \\ \Psi_{\bar{\alpha}_1}^n & \dots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix} \Rightarrow \lim_{\delta \rightarrow 0} J_n = I \Rightarrow \lim_{\delta \rightarrow 0} |J_n| = 1 > 0$$

where I is the identity matrix.

Proof of Proposition 2. Because the indexing of players is arbitrary, we solve for the impact of τ_1 on $\bar{\alpha}_1$. By the implicit function theorem,

$$\frac{\partial \bar{\alpha}_1}{\partial \tau_1} = \frac{-|B|}{|J|} \quad \text{where } B = \begin{bmatrix} \Psi_{\tau_1}^1 & \Psi_{\bar{\alpha}_2}^1 & \cdots & \Psi_{\bar{\alpha}_n}^1 \\ \Psi_{\tau_1}^2 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_n}^2 \\ \cdots & \cdots & \cdots & \cdots \\ \Psi_{\tau_1}^n & \Psi_{\bar{\alpha}_2}^n & \cdots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix}$$

The transpose of B is: $B^T = \begin{bmatrix} \Psi_{\tau_1}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_n}^n \\ \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_2}^n \\ \cdots & \cdots & \cdots & \cdots \\ \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \cdots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix} = \begin{bmatrix} \Psi_{\tau_1}^1 & 0 & \cdots & 0 \\ \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_2}^n \\ \cdots & \cdots & \cdots & \cdots \\ \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \cdots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix}$

Let B_{ij}^T denote the submatrix formed by deleting the i -th row and j -th column of matrix B^T . Then:

$$\lim_{\delta \rightarrow 0} |B| = \lim_{\delta \rightarrow 0} |B^T| = \lim_{\delta \rightarrow 0} \Psi_{\tau_1}^1 |B_{11}^T| = \frac{k}{\tau_i^2} |I| = \frac{k}{\tau_i^2} > 0 \Rightarrow \frac{\partial \bar{\alpha}_1}{\partial \tau_1} < 0 \quad \text{for small } \delta$$

A lower value of $\bar{\alpha}_i$ means that player i is more likely to file (be type $\alpha_{it} \geq \bar{\alpha}_i$). \square

Proof of Proposition 3. Because the indexing of players is arbitrary, we solve for the impact of τ_n on $\bar{\alpha}_1$. By the implicit function theorem,

$$\frac{\partial \bar{\alpha}_1}{\partial \tau_n} = \frac{-|C|}{|J|} \quad \text{where } C = \begin{bmatrix} \Psi_{\tau_n}^1 & \Psi_{\bar{\alpha}_2}^1 & \cdots & \Psi_{\bar{\alpha}_n}^1 \\ \Psi_{\tau_n}^2 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_n}^2 \\ \cdots & \cdots & \cdots & \cdots \\ \Psi_{\tau_n}^n & \Psi_{\bar{\alpha}_2}^n & \cdots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix} \quad \text{and } C^T = \begin{bmatrix} \Psi_{\tau_n}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_n}^n \\ \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_2}^n \\ \cdots & \cdots & \cdots & \cdots \\ \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \cdots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix}$$

We can rearrange the matrix using $(n-1)$ column switches and then $(n-2)$ row switches to get:

$$D \equiv \begin{bmatrix} \Psi_{\tau_n}^n & \Psi_{\tau_n}^1 & \Psi_{\tau_n}^2 & \cdots & \Psi_{\tau_n}^{n-1} \\ \Psi_{\bar{\alpha}_n}^n & \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \cdots & \Psi_{\bar{\alpha}_n}^{n-1} \\ \Psi_{\bar{\alpha}_2}^n & \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_2}^{n-1} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \Psi_{\bar{\alpha}_{n-1}}^n & \Psi_{\bar{\alpha}_{n-1}}^1 & \Psi_{\bar{\alpha}_{n-1}}^2 & \cdots & \Psi_{\bar{\alpha}_{n-1}}^{n-1} \end{bmatrix} = \begin{bmatrix} \Psi_{\tau_n}^n & 0 & 0 & \cdots & 0 \\ \Psi_{\bar{\alpha}_n}^n & \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \cdots & \Psi_{\bar{\alpha}_n}^{n-1} \\ \Psi_{\bar{\alpha}_2}^n & \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_2}^{n-1} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \Psi_{\bar{\alpha}_{n-1}}^n & \Psi_{\bar{\alpha}_{n-1}}^1 & \Psi_{\bar{\alpha}_{n-1}}^2 & \cdots & \Psi_{\bar{\alpha}_{n-1}}^{n-1} \end{bmatrix}$$

$$\text{Note that: } |C| = |C^T| = (-1)^{2n-3} |D| = -|D| = -\Psi_{\tau_n}^n |D_{11}|$$

$$\text{Define: } E \equiv D_{11} = \begin{bmatrix} \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \cdots & \Psi_{\bar{\alpha}_n}^{n-1} \\ \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \cdots & \Psi_{\bar{\alpha}_2}^{n-1} \\ \cdots & \cdots & \cdots & \cdots \\ \Psi_{\bar{\alpha}_{n-1}}^1 & \Psi_{\bar{\alpha}_{n-1}}^2 & \cdots & \Psi_{\bar{\alpha}_{n-1}}^{n-1} \end{bmatrix}$$

$$\begin{aligned}
\text{Then: } |E| &= \sum_{k=1}^{n-1} (-1)^{k+1} \Psi_{\bar{\alpha}_n}^k |E_{1k}| > 0 \\
\Leftrightarrow \phi &\equiv \sum_{k=1}^{n-1} (-1)^{k+1} f(\bar{\alpha}_n) \frac{\rho}{\rho_k \rho_n} \left[\frac{k}{\tau_k} - \frac{\delta}{1-\delta} b + r + \int_{\alpha_L}^{\bar{\alpha}_k} \alpha f(\alpha) d\alpha - \bar{\alpha}_k \rho_k \right] |E_{1k}| > 0
\end{aligned}$$

Note that: $\lim_{\delta \rightarrow 0} E_{11} = I \Rightarrow \lim_{\delta \rightarrow 0} |E_{11}| = 1$. For $k = 2, 3, \dots, n-1$, calculating $|E_{1k}|$ requires that we remove the k -th column of E . This removes $\Psi_{\bar{\alpha}_k}^k$ from the k -th row of E . Since all other entries in the k -th row of E approach 0 as δ approaches 0, $\lim_{\delta \rightarrow 0} |E_{1k}| = 0$. So:

$$\lim_{\delta \rightarrow 0} \phi = f(\bar{\alpha}_n) \frac{\rho}{\rho_1 \rho_n} \left[\frac{k}{\tau_1} - \bar{\alpha}_1 \rho_1 + r + \int_{\alpha_L}^{\bar{\alpha}_1} \alpha f(\alpha) d\alpha \right] > 0 \quad \text{by Lemma 1}$$

$$\text{which implies: } \lim_{\delta \rightarrow 0} |D_{11}| > 0 \Rightarrow \lim_{\delta \rightarrow 0} |C| < 0 \Rightarrow \frac{\partial \bar{\alpha}_1}{\partial \tau_n} > 0$$

A higher value of $\bar{\alpha}_i$ means that player i is less likely to file (be type $\alpha_{it} \geq \bar{\alpha}_i$). \square

Proof of Proposition 4. Because the indexing of players is arbitrary, we solve for the impact of r on $\bar{\alpha}_1$. By the implicit function theorem,

$$\frac{\partial \bar{\alpha}_1}{\partial r} = \frac{-|G|}{|J|} \quad \text{where } G = \begin{bmatrix} \Psi_r^1 & \Psi_{\bar{\alpha}_2}^1 & \dots & \Psi_{\bar{\alpha}_n}^1 \\ \Psi_r^2 & \Psi_{\bar{\alpha}_2}^2 & \dots & \Psi_{\bar{\alpha}_n}^2 \\ \dots & \dots & \dots & \dots \\ \Psi_r^n & \Psi_{\bar{\alpha}_2}^n & \dots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix} \quad \text{and } G^T = \begin{bmatrix} \Psi_r^1 & \Psi_r^2 & \dots & \Psi_r^n \\ \Psi_{\bar{\alpha}_2}^1 & \Psi_{\bar{\alpha}_2}^2 & \dots & \Psi_{\bar{\alpha}_2}^n \\ \dots & \dots & \dots & \dots \\ \Psi_{\bar{\alpha}_n}^1 & \Psi_{\bar{\alpha}_n}^2 & \dots & \Psi_{\bar{\alpha}_n}^n \end{bmatrix}$$

Note that:

$$\begin{aligned}
|G| = |G^T| &= \sum_{k=1}^n (-1)^{k+1} \Psi_r^k |G_{1k}^T| = \sum_{k=1}^n (-1)^{k+1} \delta \rho_{-k} |G_{1k}^T| > 0 \\
\Leftrightarrow \lambda &\equiv \sum_{k=1}^n (-1)^{k+1} \rho_{-k} |G_{1k}^T| > 0
\end{aligned}$$

Note that: $\lim_{\delta \rightarrow 0} G_{11}^T = I \Rightarrow \lim_{\delta \rightarrow 0} |G_{11}^T| = 1$. For $k = 2, 3, \dots, n$, calculating $|G_{1k}^T|$ requires that we remove the k -th column of G^T . This removes $\Psi_{\bar{\alpha}_k}^k$ from the k -th row of G^T . Since all other entries in the k -th row of G^T approach 0 as δ approaches 0, $\lim_{\delta \rightarrow 0} |G_{1k}^T| = 0$. So:

$$\lim_{\delta \rightarrow 0} \lambda = \rho_{-1} > 0 \Rightarrow \lim_{\delta \rightarrow 0} |G| > 0 \Rightarrow \frac{\partial \bar{\alpha}_1}{\partial r} < 0 \quad \text{for small } \delta$$

A lower value of $\bar{\alpha}_i$ means that player i is more likely to file (be type $\alpha_{it} \geq \bar{\alpha}_i$). \square

Proof of Proposition 5. If $\tau_i = \frac{\tau}{n}$, then there is a single cutpoint strategy used by all players in the

n -player game that is implicitly defined by:

$$\begin{aligned}\Psi^n &\equiv \bar{\alpha}_n (1 - \delta \rho_n^n) - (1 - \delta \rho_n^{n-1}) \left(\frac{kn}{\tau} - \frac{\delta}{1 - \delta} b \right) + \delta \rho_n^{n-1} \left(r + \int_{\alpha_L}^{\bar{\alpha}_i} \alpha f(\alpha) d\alpha \right) = 0 \\ &\Rightarrow \lim_{\delta \rightarrow 0} \Psi^n = \bar{\alpha}_n - \frac{kn}{\tau} = 0 \Rightarrow \lim_{\delta \rightarrow 0} \bar{\alpha}_n = \frac{kn}{\tau}\end{aligned}$$

Similarly, the cutpoint for the $(n + 1)$ -player game is implicitly defined by:

$$\begin{aligned}\Psi^{n+1} &= \bar{\alpha}_{n+1} (1 - \delta \rho_{n+1}^{n+1}) - (1 - \delta \rho_{n+1}^n) \left(\frac{k(n+1)}{\tau} - \frac{\delta}{1 - \delta} b \right) \\ &\quad + \delta \rho_{n+1}^n \left(r + \int_{\alpha_L}^{\bar{\alpha}_{n+1}} \alpha f(\alpha) d\alpha \right) = 0 \\ &\Rightarrow \lim_{\delta \rightarrow 0} \Psi^{n+1} = \bar{\alpha}_{n+1} - \frac{k(n+1)}{\tau} = 0 \Rightarrow \lim_{\delta \rightarrow 0} \bar{\alpha}_{n+1} = \frac{k(n+1)}{\tau} > \lim_{\delta \rightarrow 0} \bar{\alpha}_n\end{aligned}$$

As the number of players increases, cutpoint $\bar{\alpha}_n$ increases, which means that each player is less likely to file (be type $\alpha_{it} \geq \bar{\alpha}_n$). \square

Proof of Proposition 6. Conditional on reaching a period t , the probability that at least one state files the case in period t when there are n players is: $1 - \rho = 1 - F(\bar{\alpha}_n)^n$. This probability is decreasing in n iff $F(\bar{\alpha}_n)^n < F(\bar{\alpha}_{n+1})^{n+1}$. By the derivations in the proof of Proposition 5, $\lim_{\delta \rightarrow 0} F(\bar{\alpha}_n)^n = F\left(\frac{kn}{\tau}\right)^n$ and $\lim_{\delta \rightarrow 0} F(\bar{\alpha}_{n+1})^{n+1} = F\left(\frac{k(n+1)}{\tau}\right)^{n+1}$. Define:

$$\Psi \equiv F\left(\frac{kn}{\tau}\right)^n - F\left(\frac{k(n+1)}{\tau}\right)^{n+1}$$

And note that:

$$\Psi < 0 \Leftrightarrow 1 < \left[\frac{F\left(\frac{kn}{\tau} + \frac{k}{\tau}\right)}{F\left(\frac{kn}{\tau}\right)} \right]^n F\left(\frac{k(n+1)}{\tau}\right)$$

This holds as $F\left(\frac{k(n+1)}{\tau}\right) \rightarrow 1$, or, equivalently, as k grows large. So when k is sufficiently large, more players decrease the probability that a case is filed. The precise critical value of k will depend on the shape of the distribution function F . \square

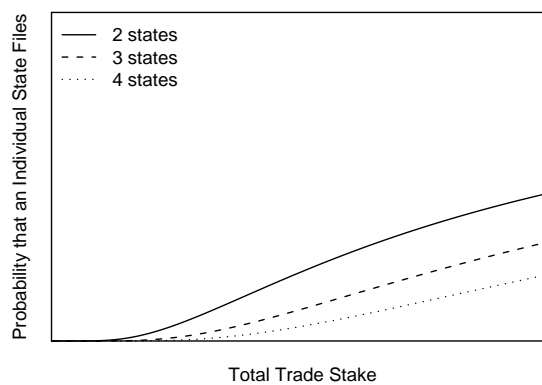
References

- Abbott, Kenneth W. and Duncan Snidal. 1998. "Why States Act through Formal International Organizations." *Journal of Conflict Resolution* 42:3–32.
- Abbott, Kenneth W., Robert O. Keohane, Andrew Moravcsik, Anne-Marie Slaughter and Duncan Snidal. 2000. "The Concept of Legalization." *International Organization* 54:401–419.
- Allee, Todd and Clint Peinhardt. 2010. "Delegating Differences: Bilateral Investment Treaties and Patterns of Dispute Resolution Design." *International Studies Quarterly* 54:1–26.
- Allee, Todd and Clint Peinhardt. 2011. "Contingent Credibility: The Reputational Effects of Investment Treaty Disputes on Foreign Direct Investment." *International Organization* 65:401–432.
- Bown, Chad P. 2005. "Participation in WTO Dispute Settlement: Complainants Interested Parties, and Free Riders." *World Bank Economic Review* 19:287–310.
- Bown, Chad P and Bernard M Hoekman. 2005. "WTO Dispute Settlement and the Missing Developing Country Cases: Engaging the Private Sector." *Journal of International Economic Law* 8(4):861–890.
- Bown, Chad P. and Kara M. Reynolds. 2014. "Trade Flows and Trade Disputes." CEPR Discussion Paper No. DP10164.
- Brutger, Ryan. 2014. "Screening for Success: The Effect of Firm Signaling on WTO Case Selection." Princeton University Working Paper.
- Brutger, Ryan and Julia Morse. 2013. "Balancing Law and Politics: Judicial Incentives in WTO Dispute Settlement." Princeton University Working Paper.
- Burley, Anne-Marie and Walter Mattli. 1993. "Europe Before the Court: A Political Theory of Legal Integration." *International Organization* 47(01):41–76.
- Busch, Marc L., Eric Reinhardt and Gregory Shaffer. 2009. "Does Legal Capacity Matter? Explaining Dispute Initiation and Antidumping Actions in the WTO." *World Trade Review* 8:559–577.
- Busch, Marc L. and Krzysztof J. Pelc. 2010. "The Politics of Judicial Economy at the World Trade Organization." *International Organization* 64:257–279.
- Carrubba, Clifford and Matthew Gabel. forthcoming. *International Courts and the Performance of International Agreements*. Cambridge University Press.
- Davis, Christina. 2003. *Food Fights Over Free Trade: How International Institutions and Issue Linkage Promote Agricultural Trade Liberalization*. Princeton University Press.
- Johns, Leslie and Krzysztof J. Pelc. 2014. "Fear of Crowds in WTO Disputes: Why Don't More Countries Participate?" UCLA and McGill University Working Paper.
- Kaye, David. 2013. "Stealth Multilateralism: U.S. Foreign Policy Without Treaties—or the Senate." *Foreign Affairs* 92:113–124.
- Kelemen, R. Daniel. 2009. *The Rules of Federalism: Institutions and Regulatory Politics in the EU and Beyond*. Harvard University Press.


- Kelley, Judith. 2007. "Who Keeps International Commitments and Why? The International Criminal Court and Bilateral Nonsurrender Agreements." *American Political Science Review* 101(03):573–589.
- Kono, Daniel Y. 2006. "Optimal Obfuscation: Democracy and Trade Policy Transparency." *American Political Science Review* 100(03):369–384.
- Kucik, Jeffrey and Krzysztof J. Pelc. forthcoming. "Measuring the Cost of Privacy: A Look at the Distributional Effects of Private Bargaining." *British Journal of Political Science* xx:xx–xx.
- Kumm, Mattias. 2005. "The Jurisprudence of Constitutional Conflict: Constitutional Supremacy in Europe Before and After the Constitutional Treaty." *European Law Journal* 11(3):262–307.
- McCubbins, Mathew D. and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *American Journal of Political Science* pp. 165–179.
- Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press.

Figure 1: Impact of Diffusion on Filing Outcomes

(a) Individual Behavior



(b) Collective Outcome



These figures were created in R from simulations using a gamma distribution.

Figure 2: Marginal Effect of Number of Countries Affected on Probability of Challenge

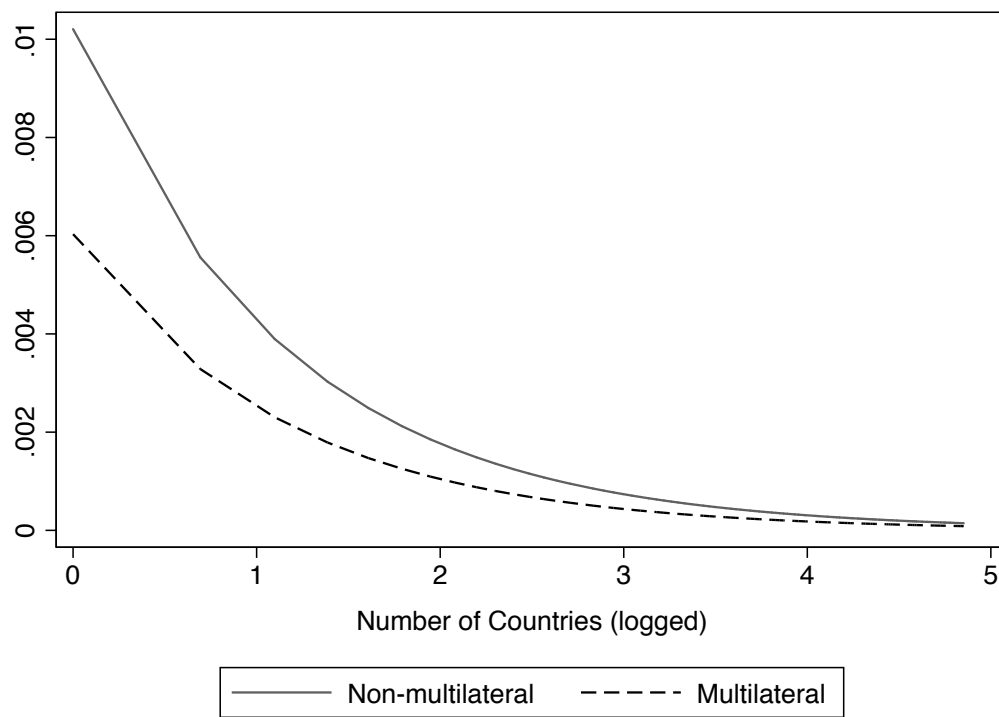


Figure 3: Marginal Effect of Disputed Trade Flows HHi on Probability of Challenge

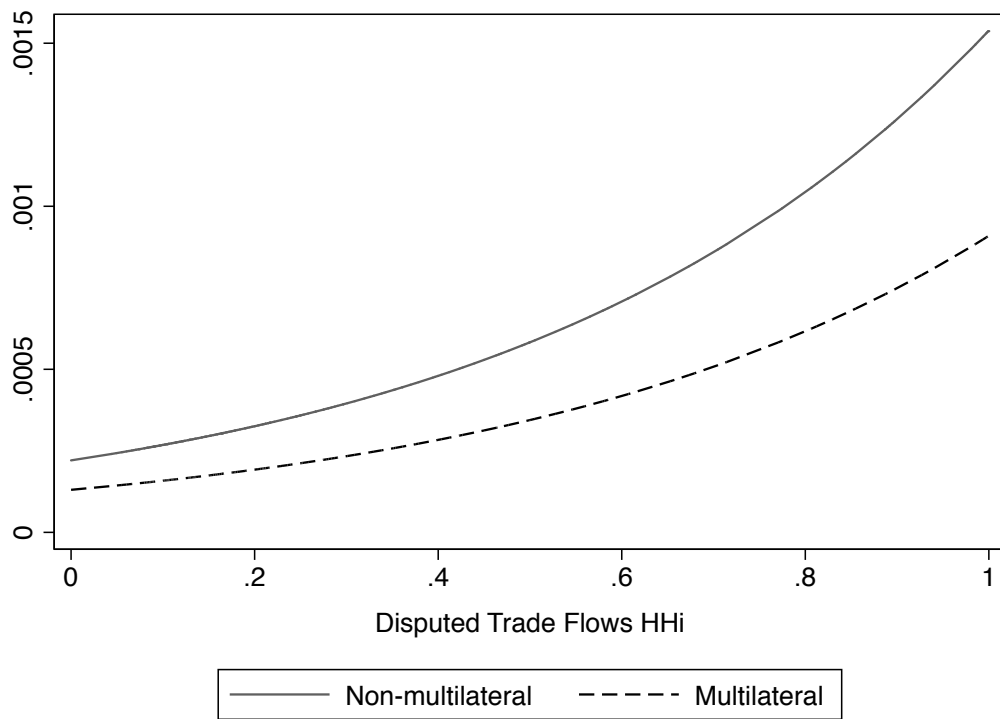


Table 1: Concentration of Benefits and Likelihood of Legal Challenges

	(1)	(2)	(3)	(4)	(5)	(6)
Multilateral Violation	-0.60*** (0.18)	-0.34** (0.17)	-0.53*** (0.17)	-0.56** (0.22)		
Number of Countries Affected		-1.12*** (0.13)	-0.88*** (0.13)		-1.33*** (0.09)	
Disputed Trade Flows HHi		1.16** (0.48)	1.94*** (0.44)			2.65*** (0.51)
Total Trade at Stake (logged)		-0.06*** (0.02)	-0.08*** (0.02)	-0.16*** (0.02)	-0.04** (0.02)	-0.16*** (0.02)
Country Exports at Stake (logged)		0.38*** (0.03)	0.44*** (0.03)	0.33*** (0.03)	0.38*** (0.03)	0.36*** (0.03)
Initiation Year		-0.02 (0.02)	-0.00 (0.02)	0.01 (0.02)	-0.02 (0.02)	0.02 (0.02)
Number of Observations	28966	28774	28774	28873	28774	28873
Number of Disputes	360	254	254	254	254	254

Columns (1) and (2) show Cox regression estimates with robust standard errors in parenthesis clustered on common dispute. Column (3) shows Cox Frailty Model estimates adjusted for shared frailty of a given respondent, standard errors in * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 2: Concentration of Benefits and Legal Merit

	(1)	(2)	(3)
CLAIMS WON NET OF APPEAL (2nd Stage eq.)			
Number of Countries Affected (logged)	0.09** (0.04)	0.09*** (0.03)	0.05*** (0.02)
Disputed Trade Flows HHi	-0.49* (0.27)	-0.50* (0.28)	-0.26** (0.13)
Multilateral Violation	0.30* (0.18)	0.26* (0.15)	0.03 (0.07)
Total Trade at Stake (logged)	0.06* (0.03)	0.05* (0.03)	0.01 (0.01)
Dispute Year	0.02 (0.02)	0.02 (0.01)	-0.00 (0.01)
Defendant GDP (logged)		-0.12 (0.08)	0.02 (0.04)
Complainant GDP (logged)		0.05 (0.04)	-0.01 (0.01)
US/EU Defendant		0.41 (0.34)	-0.20 (0.14)
Appealed		-0.16 (0.16)	0.02 (0.08)
DISPUTE PRODUCES RULING (1st Stage eq.)			
Number of Third Parties	0.13*** (0.03)	0.15*** (0.03)	0.22*** (0.04)
Total Trade at Stake (logged)	0.05*** (0.01)	0.05*** (0.01)	0.05*** (0.02)
N 1st stage / 2nd stage	338 / 176	332 / 176	332 / 176
Wald Test ($\rho=0$)	40.19	32.44	4.99

Heckman selection model with maximum likelihood (ML) estimates. First stage estimates likelihood of a ruling. Second stage in models (1) and (2) estimates number of claims in favor of complainant, and percentage of claims in favor of complainant in model (3). Robust standard errors clustered on the common dispute. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$