BREACH, REMEDIES AND DISPUTE SETTLEMENT IN TRADE AGREEMENTS

By

Giovanni Maggi and Robert W. Staiger

October 2009

COWLES FOUNDATION DISCUSSION PAPER NO. 1735

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.econ.yale.edu/

Breach, Remedies and Dispute Settlement in Trade Agreements*

Giovanni Maggi
Yale University and NBER

Robert W. Staiger
Stanford University and NBER

October 2009

Abstract

We provide a simple but novel model of trade agreements that highlights the role of transaction costs, renegotiation and dispute settlement. The model allows us to characterize the appropriate remedy for breach and whether the agreement should be structured as a system of "property rights" or "liability rules." We then study how the optimal rules depend on the underlying economic and contracting environment. Our model also delivers predictions about the outcome of trade disputes, and in particular about the propensity of countries to settle early versus "fighting it out."

*We thank Kyle Bagwell, Chad Bown, Gene Grossman, Petros C. Mavroidis, TN Srinivasan, Alan Sykes, participants in seminars at FGV- Rio, FGV-SP, PUC-Rio, Stanford and the World Trade Organization, and participants in the 2009 NBER Summer Institute and the conferences "The New Political Economy of Trade" at EUI (Florence), "The Economics, Law and Politics of the GATT-WTO" at Yale and "The Festschrift in Honor of Alan V. Deardorff" at The University of Michigan for very helpful comments.

1. Introduction

When governments make international commitments, what should be the legal remedy for breach under their agreement? Should the entitlements assigned by the agreement take the form of "property rights," whereby an entitlement can only be removed from its holder through a voluntary transaction, or of "liability rules," whereby an entitlement can be removed from its holder for the payment of legally determined damages? And what do these institutional rules imply for the propensity of countries to settle early versus "going to court"?

In this paper we address these questions as they arise in the context of international trade agreements. While our analysis applies to trade agreements generally, we pay particular attention to the rules of the World Trade Organization (WTO) and the General Agreement on Tariffs and Trade (GATT), its predecessor. This is a natural institution on which to focus, given its prominence in the world trading system.

Our analysis features the important possibility that governments may settle their differences without recourse to a ruling by the court, and hence bargain "in the shadow of the law."[1] Indeed, Busch and Reinhardt (2006) report that two thirds of WTO disputes are resolved before a final ruling of the Dispute Settlement Body (DSB). Here the legal rules and remedies do not directly determine the outcome, but they do impact the outcome indirectly by shaping what governments can expect if their attempts at settlement fail. In this case the critical role played by the DSB lies precisely in defining the disagreement point provided by the legal system. And in the presence of transaction costs, these legal rules and remedies can have important efficiency consequences even when, as in the GATT/WTO, bargaining and settlement is the dominant outcome.

These themes are virtually unexplored in the existing economics literature on trade agreements, in part because the modeling of disputes in that literature does not accommodate the possibility of settlement in a meaningful way and is therefore ill-equipped to provide a platform for formal analysis of the questions raised above. By contrast, in the law and economics literature analogous questions and issues have been extensively studied in a domestic context as they relate to the actions of private agents. There are two related literatures. A fundamental question in the literature concerned with domestic contracts (see, for example, Schwartz, 1979,

[1]This phrase appeared in the title of a paper by Mnookin and Kornhauser (1979). Like us, those authors were "...concerned primarily with the impact of the legal system on negotiations and bargaining that occur *outside* the courtroom." (p. 950, emphasis in the original).

Ulen, 1984, and Shavell, 2006) is when contracting parties would want specific performance as a remedy for contract breach and when they would instead prefer damage payments. There is also a vast literature (the seminal contributions are Calabresi and Melamed, 1972, and Kaplow and Shavell, 1996) that is concerned with the related question of when property rules are preferred to liability rules in the design of domestic law. A major goal of our paper is to initiate the formal analysis of these themes in the context of international trade agreements.[2]

Looking across the GATT/WTO, one can see evidence of both liability rules and property rules at work. For instance, there is broad agreement that in the early years GATT operated in effect as a system of liability rules, where breach remedies were similar to compensation for escape clause actions and amounted to damage payments for the purpose of "rebalancing" (see Jackson, 1969, p. 147, Schwartz and Sykes, 2002, and Lawrence, 2003, p. 29). However, it has been observed that with the advent of the WTO the remedy for breach has moved away from a system of damage payments and toward specific performance or "compliance" (see Jackson, 1997, Charnovitz, 2003, and Pelc, 2009), resulting in an institution where some entitlements (e.g. the prohibitions against quantitative restrictions, export and to some extent domestic subsidies, and discriminatory trade tariffs) are now generally protected by property rules while others (e.g., negotiated tariff bindings) are still protected by liability rules through applicable escape clauses (and Pelc, 2009, argues that these escape clauses are themselves evolving toward property rules as well). This interpretation of the evolution of the GATT/WTO is itself a matter of debate among legal scholars (see Hippler Bello, 1996 and Schwartz and Sykes, 2002 who argue that the WTO can still best be viewed as a system of liability rules), but it is clear that the choices that determine whether the GATT/WTO behaves more like a system of property rules or rather liability rules are a key feature of institutional design. Our paper aims at contributing toward an understanding of the forces that explain these choices.

In the domestic context that is the focus of the law and economics literatures, transaction costs are typically associated with private information or other bargaining frictions. Such frictions are surely present as well in international bargaining, but in the international context there is an additional feature that is particularly salient: there generally does not exist an efficient government-to-government transfer mechanism that can be used to make damage payments, either to settle disputes or to compensate for the exercise of escape clauses. In the

[2] Lawrence (2003) provides a lucid treatment of these themes as they relate to international trade agreements at an informal level, and Srinivasan (2007) contains a related discussion.

GATT/WTO, the typical means by which one government achieves compensation for the harm caused by another government's actions is through "counter-retaliation," that is, by raising its own tariffs above previously negotiated levels. Such compensation mechanisms entail important inefficiencies (deadweight loss) that, while plausibly absent in the domestic private agent context, introduce a novel transaction cost in the international context.

A major point of departure of our model is precisely this difference between the domestic private agent setting and the international government-to-government setting. In particular, we consider a setting where governments, operating in the presence of ex-ante uncertainty about the joint benefits of free trade (which could be positive or negative due to the possible presence of political economy factors), contract over trade policy and define a mandate for the DSB in the event that contract disputes should arise ex post, once uncertainty has been resolved and trade policies are chosen. We assume that transfers between governments are costly and that the marginal cost of transfers is (weakly) increasing in the magnitude of the transfer.

We consider agreements that specify a baseline commitment to free trade but allow the importing government to escape (breach) this commitment by compensating the exporter with a certain amount of damages. If the level of damages is set either at zero or at a level so high that the importing government would never choose to breach, then this amounts to a property rule in which either the entitlement to protect is assigned to the importing government or the entitlement to free trade is assigned to the exporting government. If instead damages are set at an intermediate level, then this is a liability rule. Our formalization of damages for breach can be interpreted in either of two ways. The first is that breach damages are specified explicitly in the contract in the form of an escape clause; the second is that the contract specifies a rigid free trade commitment, but the DSB may be given a mandate to require the payment of damages in case of breach. As we have indicated above and discuss further below, both of these interpretations are relevant for the GATT/WTO.

We assume that the DSB, if invoked, initiates an investigation and observes a noisy signal of the joint benefits of free trade, and issues a ruling based on this imperfect information. At the time when the DSB can be invoked, the governments are uncertain about the outcome of the DSB investigation and hence about the DSB ruling, and subject to this uncertainty must decide whether to invoke the DSB or to settle. And if a DSB ruling is reached, the governments can renegotiate the ruling; this is a further possibility for renegotiation and settlement, in addition to the possibility before any DSB ruling which we already mentioned.

These two features of our model – the possibility that governments negotiate an early settlement under *uncertainty* about the DSB ruling, and the possibility of a *later settlement* in which the DSB ruling is renegotiated – constitute a second important point of departure from the law-and-economics literature discussed above. Allowing governments to be uncertain about the DSB ruling is important because, as will become clear, this allows for the possibility that the governments may *not* settle early; and as a consequence, the model generates a variety of predictions regarding when governments settle early or a dispute arises in equilibrium, and how the disputes are resolved. And allowing for the possibility of later settlement generates predictions about the circumstances under which a DSB ruling is renegotiated in equilibrium.[3]

We start with a benchmark scenario in which the DSB receives no information ex post, and we consider the impact of ex-ante uncertainty over the joint benefits of free trade. We find that a property rule, which either demands strict performance or permits complete discretion in the choice of trade policy, is optimal if ex-ante uncertainty is sufficiently low, whereas a liability rule tends to be optimal when ex-ante uncertainty is high. This finding suggests that, as uncertainty over the joint benefits of free trade falls, the optimal institutional arrangement should tend to move away from liability rules toward property rules; and that liability rules should be more prevalent than property rules in issue areas characterized by a higher degree of uncertainty over the joint benefits of free trade, and vice-versa.

Next we turn to the more general case in which the DSB observes a noisy signal ex post about the joint benefits of free trade. In this case, if ex-ante uncertainty about the joint benefits of free trade is small, a property rule is again optimal, but with the assignment of entitlements contingent on the signal received by the DSB; and if ex-ante uncertainty is sufficiently large, a liability rule again tends to be optimal, but with the DSB reducing the level of damages when it receives a signal that the joint benefits from free trade are smaller or negative. We also establish that, for a given degree of ex-ante uncertainty, if the noise in the DSB signal is itself sufficiently small, a property rule is optimal, with the assignment of entitlements contingent on the signal. This finding suggests that, as the accuracy of DSB rulings increases, the optimal institutional arrangement should tend to move away from liability rules toward property rules.

More broadly, if one accepts that the accuracy of DSB rulings has increased over time, or that the degree of ex-ante uncertainty about the joint benefits of free trade has fallen over time,

[3] As we discuss later in the paper, subsequent to DSB rulings it is not uncommon for WTO disputes ultimately to be settled through "mutually agreed solutions" between the disputing parties rather than implementation of the ruling, corresponding roughly to a renegotiation of the DSB ruling in our formal analysis.

then our model predicts a gradual shift from liability rules to property rules. As we indicated above, a number of legal scholars maintain that this shift can be seen in the GATT/WTO.

We also find that, in the circumstances where a liability rule is optimal, it is never optimal to set damages high enough to make the exporter "whole." This runs counter to the presumption established by the "efficient breach" perspective of the law-and-economics literature, according to which damages should be set at a level that makes the injured party whole so that breach occurs whenever it is optimal to "buy out" the injured party. In fact, this presumption is often extended informally to the context of the WTO, where it is sometimes suggested that the principle of reciprocity which guides much of the GATT/WTO remedy system falls short because it does *not* make the injured party whole (see, for example, Charnovitz, 2002, and Lawrence, 2003). Our finding points out an important caveat in this reasoning: simply put, in the WTO context, the damages paid for breach often take the form of counter-retaliation on the part of the injured party, and this is an inefficient means of compensation that, from an ex-ante perspective, should not be permitted to an extent that makes the injured party whole.

In addition, we find that the damages for breach should be responsive to both the harm caused to the exporter *and* the benefit garnered by the importer. We relate this finding to the WTO Agreement on Safeguards, and suggest that it may be helpful in interpreting the WTO rules on compensation for escape clause actions.

Our model generates interesting insights also with regard to the role of transaction costs in determining the optimal rules. We find that a property rule tends to be preferable to a liability rule when the cost of transfers is high. We also examine the impact of frictions in ex-post bargaining, by comparing the case of frictionless ex-post bargaining with the extreme opposite case in which ex-post bargaining is not feasible, and find that the introduction of frictions in bargaining may favor property rules over liability rules. These results contrast with the findings in the law-and-economics literature that liability rules tend to be preferable to property rules when transaction costs are high (Calabresi and Melamed, 1972, and Kaplow and Shavell, 1996).

Next we consider when disputes arise in equilibrium and how disputes are resolved. We find that early settlement of disputes is more likely when, at the time of contracting, there is more uncertainty about the future joint benefits from free trade; and we find that early settlement occurs when the joint benefits of free trade turn out to be either very high or very low. We also find that, conditional on the DSB being invoked, the ruling is implemented when the DSB receives information that the joint benefits of free trade are either very high or very low, and consequently sets a very high/low level of damages; but the DSB ruling is renegotiated if the realized signal, and hence the level of damages set by the DSB, lies in an intermediate range. A corollary is that renegotiation of DSB rulings need not reflect a "bad" (inaccurate) ruling.

Finally, we consider the possibility that the cost of granting transfers differs across the two countries, and we interpret the high-cost country as a developing country. We find that there is a tendency for developed countries to impose more protection as a result of early settlements than developing countries. In addition we find that there is a pro-trade (anti-trade) selection bias in DSB rulings when a developed country (developing country) is the respondent.

Beyond the literature we have already mentioned, there are a number of additional papers that are related to ours. Like us, Beshkar (2008a,b) considers the possibility of efficient breach with non-verifiable political pressures and costly transfers, but his model differs from ours in a number of important ways: most significantly, he does not allow for the possibility of renegotiation and settlement, which as we have emphasized is central to our analysis. Similarly, Howse and Staiger (2005) investigate whether the GATT/WTO reciprocity rule might be interpreted as facilitating efficient breach, but they do not consider the possibility of settlement either. Our paper is also related to Maggi and Staiger (2008). But that paper abstracts from issues of costly transfers and settlement to highlight instead contract vagueness and interpretation.[4]

The rest of the paper proceeds as follows. The next section lays out the basic model. Section 3 examines the nature of the optimal rules. Section 4 focuses on the outcome of disputes. Section 5 considers two extensions: a more general informational environment where uncertainty can be multi-dimensional, and a more general class of contracts that allows not only for a "stick" associated with protection, but also for a "carrot" associated with free trade. Section 6 concludes. All proofs are contained in the Appendix.

2. The Basic Model

We focus on a single industry in which the Home country is the importer and the Foreign country is the exporter. The government of the importing country chooses a binary level of

[4] Bagwell and Staiger (2005), Martin and Vergote (2008) and Bagwell (2009) consider models with privately observed political pressures, but they do not consider the role of a court (DSB) and focus instead on self-enforcement issues (from which we abstract). Park (2009) does consider the role of the DSB in a setting with privately observed political pressures, but the DSB is formalized as a device that automatically turns private signals into public signals, without any filing decisions by governments, and hence his model cannot make a distinction between renegotiation/settlement and the triggering of DSB rulings. See also Ethier (2001) for an early attempt to formalize the role of the DSB and its impact on the resolution of disputes.

trade policy intervention for the industry, which we denote by $T \in \{FT, P\}$: "Free Trade" or "Protection."[5] We assume that the exporting government is passive in this industry.

At the time that the Home government makes its trade policy choice, a transfer may also be exchanged between the governments, but at a cost. Here we seek to capture the feature that cash transfers between governments are seldom used as a means of settling trade disputes, while indirect (non-cash) transfers, such as tariff adjustments in other sectors or even non-trade policy adjustments, are more easily available.[6] To allow for this possibility in a tractable way, we let b denote a (positive or negative) transfer from Home to Foreign, and we let $c(b)$ denote the deadweight loss associated with the transfer level b. We assume that $c(b)$ is non-negative, (weakly) convex and smooth everywhere, with the natural features that $c(0) = 0$ and $c(b) > 0$ for $b \neq 0$. Finally, it is convenient to assume that the deadweight loss $c(b)$ is borne by the country that makes the transfer: thus, the loss is borne by Home (Foreign) if $b > 0$ $(b < 0)$.

The importing government's payoff is given by

$$\omega(T, b) = v(T) - b - c(b)I, \tag{2.1}$$

where $v(T)$ is the importing government's valuation of the domestic surplus associated with policy T in the sector under consideration, and I is an indicator function that is equal to one if Home makes the transfer (i.e., if $b > 0$) and equal to zero otherwise. We have in mind that $v(T)$ corresponds to a weighted sum of producer surplus, consumer surplus and revenue from trade policy intervention, with the weights possibly reflecting political economy concerns (as in, e.g., Baldwin, 1987, and Grossman and Helpman, 1994). As we noted above, the exporting government is passive in this industry; its payoff is therefore

$$\omega^*(T, b) = v^*(T) + b - c(b)I^*, \tag{2.2}$$

where $v^*(T)$ is the exporting government's valuation of the foreign surplus associated with policy T, and I^* is an indicator function that is equal to one if Foreign makes the transfer (i.e., if $b < 0$) and equal to zero otherwise.

[5] Our assumption of a binary policy instrument helps to keep our analysis tractable, and captures reasonably well a variety of non-tariff policy choices, such as regulatory regimes or product standards, that are discrete in nature. Many of the trade disputes in the GATT/WTO focus on these kinds of policy issues.

[6] The resolution of GATT/WTO disputes has, with one exception, never involved cash transfers (the one exception to date is the *US-Copyright* case; see WTO, 2007, pp. 283-286). However, in the context of a trade dispute countries do sometimes achieve the indirect payment of compensation through the WTO "self-help" method of counter-retaliation in other sectors. And WTO disputes that are settled by a "mutually agreed solution" under Article 3.6 of the WTO Dispute Settlement Understanding (either before or after a DSB ruling) may involve a variety of indirect transfer mechanisms.

Using (2.1) and (2.2), the joint payoff of the two governments is denoted as Ω and given by

$$\Omega(T, b) = v(T) + v^*(T) - c(b). \tag{2.3}$$

We assume that Home always gains from protection, and we denote this gain as

$$\gamma \equiv v(P) - v(FT) > 0.$$

This gain may be interpreted as arising from some combination of terms-of-trade and political considerations. On the other hand, we assume that Foreign always loses from protection, and we denote this loss as

$$\gamma^* \equiv v^*(FT) - v^*(P) > 0.$$

The joint (positive or negative) gain from protection is then $\Gamma \equiv \gamma - \gamma^*$.

In this simple economic environment, the "first best" (joint-surplus maximizing) outcome is easily described: if $\Gamma > 0$ (or $\gamma > \gamma^*$), the first best is $T = P$ and $b = 0$, and if $\Gamma < 0$ (or $\gamma < \gamma^*$), the first best is $T = FT$ and $b = 0$. Notice that b always equals zero under the first best, because transfers are costly to execute. For future use, we denote by Ω_{FB} the first-best joint payoff level.

We assume that governments are ex-ante uncertain about Γ, but they observe Γ ex post. If Γ were verifiable (i.e. observed ex post by the court/DSB), of course the governments could write a complete contingent contract, and the problem would be uninteresting. We are interested instead in an imperfect-contracting scenario, where such a complete contingent contract cannot be written. We consider the simplest scenario of this kind that allows us to make the relevant points. We assume that γ^* is ex-ante known to all, so that all the uncertainty in Γ originates from γ, and that γ is not verifiable. In section 5 we extend our results to the case in which γ^* is also uncertain (and not verifiable). There is also a further motivation – besides simplicity – for considering the case in which γ^* is known ex ante. This is the case that is most favorable to the standard argument for a liability rule, according to which efficiency can be induced if the exporter is made whole with a damage payment of γ^* in the event of breach. We will show that, even in this most-favorable case, the standard argument for a liability rule must be qualified in our setting along a number of important dimensions.

We denote by $h(\gamma)$ the ex-ante distribution of γ, which we assume to be common knowledge (to the governments as well as the DSB). Unless otherwise noted, we assume that the support of γ is bounded. Finally, to make things interesting, we assume that the value $\gamma = \gamma^*$ is in the interior of the support of γ, so that the first-best is P in some states (when $\gamma > \gamma^*$, and hence $\Gamma > 0$) and FT in some states (when $\gamma < \gamma^*$, and hence $\Gamma < 0$).

The fact that governments cannot write a complete contingent contract does not *necessarily* imply inefficiencies. If international transfers were costless (no deadweight loss), then governments could always achieve the first-best by engaging in ex-post (i.e., after observing γ) negotiations over policies and (costless) transfers.[7] If international transfers are costly, on the other hand, the first best cannot be achieved in general, but ex-ante joint surplus may be enhanced by writing a contract ex ante (before γ is realized), and defining a role for the DSB in the event that contract disputes arise ex post. We look for the contract/DSB combination that maximizes ex-ante joint surplus.[8]

We may now describe the timing of events. The game is as follows:

stage 0. Governments write the contract and define the role of the DSB.

stage 1. The state of the world is realized and observed by the governments.

stage 2. The importer proposes a T and a b that can differ from the terms of the contract. The exporter either accepts the proposal or files with the DSB.

stage 3. If invoked, the DSB issues a ruling.

stage 4. The importer can propose a deviation from the ruling. The exporter either accepts the proposal (so that the DSB ruling is not enforced) or demands enforcement of the ruling.

stage 5. Trades occur and payoffs are realized.

Note that we allow ex-post renegotiation of the initial contract (in stage 2) as well as renegotiation of the DSB ruling (in stage 4); and we assume that the importer makes take-or-leave offers. Opportunities for renegotiation are central to our analysis, and as we indicated

[7] We abstract from issues of enforcement here and simply assume that bargaining outcomes between the two governments are enforced.

[8] There are three ways to justify this emphasis on the maximization of the governments' ex-ante joint surplus. One possibility is to allow for costless ex-ante transfers, i.e., transfers at the time the institution is created. This justification is not in contradiction with our assumption of costly ex-post transfers, if it is interpreted as reflecting the notion that the cost of transfers can be substantially eliminated in an ex-ante setting such as a GATT/WTO negotiating round where many issues are on the table at once (see, for example, the discussion in Hoekman and Kostecki, 1995, Ch. 3). A second possibility would be to keep the single-sector model and introduce a veil of ignorance, so that ex-ante there is uncertainty over which of the two governments will be the importer and which the exporter. And a third possibility would be to introduce a second mirror-image sector.

in the Introduction and describe further below, they are an important feature of the dispute resolution process for trade agreements such as the GATT/WTO. By contrast, the assumption of take-or-leave offers makes our analysis easier, but it is not critical for our results.

2.1. The feasible contracts and the role of the DSB

We next describe the feasible contracts and the role of the DSB. We consider a family of *menu contracts* that allow the importer to choose between (i) setting FT and (ii) setting P and compensating the exporter with a payment b^D. Or, using slightly different terminology, these contracts specify a baseline commitment (FT) but allow the importer to escape/breach this commitment by paying a certain amount of *damages*.

Note that this family of contracts includes two contracts of particular interest: (i) a rigid $\{FT\}$ contract, which corresponds to the case in which b^D is set at a prohibitively high (or infinite) level; we often refer to this contract as one that requires strict *performance* under all circumstances, or in short, a "performance contract," and (ii) discretion over trade policy, which corresponds to the case in which $b^D = 0$. Thus, as the level of damages b^D goes from zero to prohibitive, the menu contract spans all the interesting possibilities, ranging from discretion to a contract that stipulates (non-prohibitive) damages to a strict performance contract.

Observe as well that setting damages at either zero or a prohibitive level amounts to establishing a *property rule*, in which as a legal matter the right to protect is granted to the importer (when damages are set to zero) or the right to free trade is granted to the exporter (when damages are set at a prohibitive level). And setting damages strictly between zero and the prohibitive level amounts to establishing a *liability rule*. In what follows we also draw links between our results and the law-and-economics literature that is concerned with the choice between property rules and liability rules.

We initially consider a benchmark scenario in which the DSB does not receive any information ex post (section 3.1), so that the governments have no uncertainty about the DSB ruling at the ex-post negotiation stage, and later (section 3.2) we consider the case in which the DSB can observe a noisy signal of γ (which we denote $\hat{\gamma}$), so that governments are uncertain about the DSB ruling at the ex-post negotiation stage. In this latter case, we consider a wider class of contracts, where b^D can be contingent on $\hat{\gamma}$: given the b^D schedule specified by the contract, if the DSB is invoked, it estimates the damages due to the exporter based on its information.

There are two interpretations of the optimization problem we have just outlined. The first, more direct interpretation is that governments design a *contract* that specifies a baseline commitment to free trade but includes an explicit *escape clause*. Some WTO contracts/clauses take this form, for example negotiated tariff commitments and the associated GATT Article XIX Escape Clause and/or Article XXVIII renegotiation provisions.[9] Given this interpretation, we may ask what is the appropriate level of *damages* that should be included in the contract: the answer here is relevant for the design of explicit escape clause provisions. Under this interpretation a DSB ruling simply enforces contract performance (i.e. ensures that the importing government either selects FT or pays the contractually specified damages).

A second interpretation is that governments design an *institution* consisting of two parts: (i) a rigid $\{FT\}$ contract with no contractually specified means of escape; and (ii) a mandate for the DSB, which instructs the DSB to enforce a certain *remedy* (the payment b^D) that the importing government must pay the exporting government in case of *breach*. If the breach remedy is prohibitive in all states of the world, we say that the DSB is instructed to enforce contract performance. At the opposite extreme, if $b^D = 0$, so that the DSB permits breach at zero cost, the outcome is the same as with full discretion, just as under the previous interpretation. In the WTO, many contractual commitments are best described as rigid (e.g., the ban on export subsidies). And we may then ask what is the appropriate remedy for breach that should be made available by the DSB: the answer is relevant for the mandate/design of the DSB.[10]

Our analysis applies equally well under either of these interpretations, i.e., whether the breach remedy is specified in the contract or rather in the DSB mandate. (In a richer model, both could coexist, with general breach possibilities determined by the mandate of the DSB and more specific breach possibilities made available in the contract). In either case, the level of the breach remedy is important for the same reason: it serves to define the disagreement point provided by the legal system should ex-post negotiations between the governments fail.

[9] The non-violation nullification-or-impairment clause of the GATT can also be interpreted along the lines of an escape clause, as it permits countries to in effect breach their negotiated market access commitments with unanticipated changes in domestic policies and pay damages to injured parties as a remedy.

[10] In reality, the role of DSB investigations in a setting such as the GATT/WTO is typically twofold: first, to establish whether or not breach exists; and second, to determine damages if this becomes necessary as part of a remedy. In the model we develop here, the first role is trivially fulfilled, because we have assumed that the contractually obligated policy is clearly specified and the policy choice is publicly observable. This assumption allows us to concentrate on the second DSB role listed above (i.e., the determination of damages), which is the focus of our analysis. Of course, a typical GATT/WTO dispute involves multiple claims of breach which the DSB must evaluate, and so in practice the breach determination itself helps to determine the level of damages, and the two roles of the DSB are intertwined.

2.2. The ex-post Pareto frontier

We complete our description of the basic model by describing the ex-post Pareto frontier and how it varies with the realized state (γ). The ex-post frontier plays a key role in the analysis, since it describes the set of feasible payoffs for the negotiations in both stages 2 and 4.

To describe the ex-post frontier, we partition the possible realizations of γ into four intervals (or "regions"): Region I ($\Gamma << 0$), where the efficiency gains from FT are large; Region II ($\Gamma < 0$), where the efficiency gains from FT are relatively small; Region III ($\Gamma > 0$), where the efficiency gains from P are relatively small; and Region IV ($\Gamma >> 0$), where the efficiency gains from P are large. Figure 1 depicts the ex-post Pareto frontier for a representative γ realization in each of Regions I through IV. For each region, the Pareto frontier corresponds to the outer envelope of two concave sub-frontiers, one passing through point P (and associated with $T = P$ and various levels of the transfer b), the other passing through point FT (and associated with $T = FT$ and various levels of b): the concavity of each sub-frontier reflects the convexity of $c(b)$. Recalling our assumption that the value $\gamma = \gamma^*$ is in the interior of the support of γ, it follows that Regions II and III are non-empty. By contrast, Regions I and/or IV are relevant only if the support of γ is sufficiently large.

The top left panel of Figure 1 depicts the ex-post frontier for Region I. Here, the efficiency gains from FT are large, and as a consequence, achieving the frontier always requires $T = FT$; moreover, the frontier is concave (in the relevant part of the payoff space), reflecting the mounting inefficiency associated with greater transfers as we move away (in either direction) from point FT (where $\omega = \omega(FT, 0)$ and $\omega^* = \omega^*(FT, 0)$). The bottom right panel of Figure 1 depicts the ex-post frontier for Region IV. Here the efficiency gains from P are large, and so achieving the frontier always requires $T = P$; and again, in Region IV the frontier is concave, reflecting the mounting inefficiency associated with greater transfers as we move away from point P (where $\omega = \omega(P, 0)$ and $\omega^* = \omega^*(P, 0)$).

Now consider the top right panel of Figure 1, which depicts the ex-post frontier for Region II. Here, FT entails a higher joint payoff than P, but by a relatively small amount, and as a consequence neither of the policies FT or P Pareto-dominates the other; and the frontier is piece-wise concave but globally non-concave, because *both* the policy T and the transfer b change as we move along the frontier. The lower left panel of Figure 1 depicts the ex-post frontier for Region III. Here the situation is qualitatively similar to Region II, except that now it is P that entails a higher joint payoff (by a relatively small amount).

As we have noted, Regions I and IV are relevant only if the support of γ is sufficiently large, and as should now be apparent, the bargaining environment in Regions I and IV is very different from that in Regions II and III; as a consequence, the degree of uncertainty over γ turns out to be a key determinant of the optimal breach remedy. Also, as can be seen by inspection of Figure 1, making the support of γ larger has qualitatively similar effects as making the cost of transfers smaller (holding everything else constant); hence, the cost of transfers is also pivotal in determining the optimal breach remedy.

For now it suffices to observe that, for any realized γ, the outcome of the negotiations is determined by the features of the relevant ex-post frontier and the position of the disagreement point relative to the frontier. And as we have observed, the disagreement point is shaped by the level of damages b^D.

3. The Optimal Rules

We now turn to a complete analysis of the game. We begin with the benchmark scenario in which the DSB receives no information ex post, so that governments have no uncertainty about the DSB ruling at the stage of ex-post negotiations. We then consider noisy DSB investigations.

3.1. A benchmark scenario: the DSB receives no information ex post

In this section we suppose that the DSB receives no information ex post, and thus b^D must be noncontingent. For a given b^D, we solve the game by backward induction, and then determine the value of b^D that maximizes ex-ante joint surplus.

The backward induction analysis is simplified by observing that, in this benchmark scenario, the outcome of the stage-4 bargain must be the same as the outcome of the stage-2 bargain. Intuitively, since bargaining is efficient, the outcome of the stage-4 subgame is on the Pareto frontier for any state γ – call this point Q_4. Now consider what happens at stage 2 when governments negotiate after observing the state. At this stage governments anticipate perfectly the DSB decision, and so the threat point of this negotiation is simply point Q_4. Since Q_4 is on the Pareto frontier, there is no possible Pareto improvement that governments can achieve at stage 2 over the threat point, and hence the equilibrium outcome of the stage-2 bargain is the same as that of the stage-4 bargain. (For this reason, in this section we often refer to "the" bargain without specifying the stage in which the bargain occurs.)

In light of this observation, to determine the optimal b^D we just need to derive the equilibrium joint surplus of the stage-4 bargain as a function of b^D, take the expectation of this joint surplus over all states γ, and optimize b^D. Here we sketch the intuition behind our results, relegating proofs to the Appendix.

Given that FT is the first-best policy in Regions I and II while P is the first-best policy in Regions III and IV, it might be expected that efficiency considerations would push toward a high value of b^D when the state falls in Regions I or II, and toward a low value of b^D when it falls in Regions III or IV. And given this expectation, it seems natural that the optimal level of b^D would then be somewhere in the middle (i.e., a liability rule), since there is uncertainty over whether the realized state lies in Regions I/II or Regions III/IV. But things are more complicated, in part because as we have noted the ex-post frontier can be non-concave (Regions II and III), and in part because the importer has two distinct choices under disagreement (FT with no transfer or P with transfer b^D), only one of which is "active" under the circumstances (i.e., the choice preferred by the importer given the realized γ).

In what follows, we use b to denote the *equilibrium* transfer, as opposed to the contractually-specified damages b^D. Let us examine how the outcome of the bargain depends on b^D, focusing first on Regions II and III, where Γ is not far from zero. The top left panel of Figure 2 depicts how the bargaining outcome varies with b^D in Region II. The first step is to determine the disagreement point as a function of b^D. For small b^D (between the point labelled P and point R), if negotiations failed the importer would choose to set $T = P$ and pay damages b^D (rather than FT with no transfer); this is a disagreement point on the frontier, and so the negotiation yields $T = P$ and $b = b^D$. For intermediate b^D (between the points R and J), if negotiations failed the importer would still choose to set $T = P$ and pay damages b^D; but this is now a disagreement point inside the frontier, and so the negotiations lead to a choice of $T = FT$ and a transfer from the exporter ($b < 0$), and the DSB ruling is therefore renegotiated.[11] Finally, for b^D beyond the "prohibitive" level at J, if negotiations failed the importer would choose to set $T = FT$ and pay zero damages; this is a disagreement point on the frontier, and so the negotiations yield $T = FT$ and $b = 0$. An analogous interpretation applies for Region III as

[11] The equilibrium transfer from the exporter ($b < 0$) is implicitly defined by the negative solution to the equation $b - c(b) = b^D - \gamma^*$. The particular level of b to which the governments renegotiate reflects our assumption that the importing government makes a take-or-leave offer to the exporting government. Alternative bargaining assumptions, such as Nash bargaining, would alter the level of b in a straightforward manner, but would not change our basic results.

depicted in the bottom left panel of Figure 2.

Let $b^{reneg}(\gamma)$ and $b^{prohib}(\gamma)$ denote the levels of b^D associated respectively with points R and J for the realized state. The top right and bottom right panels of Figure 2 depict how joint surplus Ω varies with b^D in Regions II and III, respectively. Notice that in each region Ω is non-monotonic in b^D, that $d\Omega/db^D = 0$ at $b^D = 0$ (reflecting the assumption $c'(0) = 0$), and that $b^{prohib}(\gamma) < \gamma^*$ for all states in Regions II and III.

These pictures suggest a key observation. If the support of γ around γ^* is sufficiently small, so that Γ can never be very far from zero, then only Regions II and III are relevant, and the expected joint surplus is maximized by a property rule which either permits discretion ($b^D = 0$) or requires strict performance in all states ($b^D \geq \bar{b}^{prohib} \equiv \max_\gamma b^{prohib}(\gamma)$); in other words, a liability rule that induces breach under some circumstances ($b^D \in (0, \bar{b}^{prohib})$) is never optimal. To see this more directly, it is helpful to focus on the case where $\gamma = \gamma^*$, which marks the border between Regions II and III: in this case, as Figure 3 clearly indicates, a liability rule in which $b^D \in (0, \bar{b}^{prohib})$ can never be optimal. Intuitively, when $\gamma = \gamma^*$ joint surplus is not sensitive to which policy is chosen, and so the goal is simply to avoid a costly transfer $b \neq 0$, which only the property rules $b^D = 0$ and $b^D \geq \bar{b}^{prohib}$ can accomplish. As long as the support of γ around γ^* is sufficiently small, the sensitivity of joint surplus to the chosen policy can never be large enough to overturn this logic. Thus it is intuitive that a liability rule cannot be optimal in the case of small support of γ; this intuition is confirmed by Proposition 1(i) below.

We turn next to the case of large uncertainty, in which Regions I and IV now also become relevant. For Region I, the top left panel of Figure 4 depicts how the bargaining outcome varies with b^D. For this region, the bargaining outcome always entails $T = FT$; but as b^D rises from zero, the outcome moves from left to right along the bold portion of the frontier up to the prohibitive level of b^D corresponding to point J. For b^D between zero and this prohibitive level, if negotiations failed the importer would choose to set $T = P$ and pay damages b^D; this is a point inside the frontier, and so the negotiation leads to a choice of $T = FT$ and a transfer from the exporter, and the DSB ruling is therefore renegotiated. For b^D beyond this prohibitive level, if negotiations failed the importer would choose to set $T = FT$ and pay zero damages; this is a point on the frontier, and so the negotiation yields $T = FT$ and $b = 0$. The top right panel of Figure 4 depicts how Ω varies with b^D for a given state in Region I. Notice that Ω is increasing in b^D, that $d\Omega/db^D > 0$ at $b^D = 0$ (reflecting the fact that $b < 0$ at $b^D = 0$), and that $b^{prohib}(\gamma) < \gamma^*$ for any state in Region I.

The bottom panels of Figure 4 depict the same information for Region IV. In this region, as the bottom left panel indicates, the outcome always entails $T = P$; but as b^D rises from zero, the outcome moves from left to right along the bold portion of the frontier up to a prohibitive level of b^D corresponding to point J. For b^D between zero and this prohibitive level, if negotiations failed the importer would set $T = P$ and pay damages b^D; this is a point on the frontier, and so negotiations implement the DSB ruling $T = P$ and $b = b^D$. For b^D at or beyond this prohibitive level, if negotiations failed the importer would set $T = FT$ and pay zero damages; this is a point inside the frontier, and so the negotiations lead to $T = P$ and $b < b^D$, and the DSB ruling is renegotiated. The bottom right panel depicts how Ω varies with b^D for a given state in Region IV. Notice that Ω is (weakly) decreasing in b^D, that $d\Omega/db^D = 0$ at $b^D = 0$, and that $b^{prohib}(\gamma) > \gamma^*$ for any state in Region IV.

Together, Figures 2 and 4 suggest a second key observation: if uncertainty about γ is large so that Regions I through IV are all relevant, then a liability rule is optimal. To see this, first note that discretion ($b^D = 0$) cannot be optimal. This is because Ω may be raised for states in Region I by increasing b^D slightly above zero – which reduces the equilibrium transfer b that the importer can extract in exchange for FT – while joint surplus in Regions II-IV are unaffected (to the first order) by this maneuver. Next note that a prohibitive level of b^D cannot be optimal either: by inspection of Figures 2-4, decreasing b^D from a prohibitive level to a level slightly below γ^* strictly improves joint surplus in Region IV (by reducing the equilibrium transfer b in Region IV that the importer must pay to adopt P) while not affecting joint surplus in the other regions. And an immediate corollary of this argument is that the optimal level of b^D is strictly lower than γ^*.

With the intuition developed above, we are now ready to state our first proposition:

Proposition 1. *(i) If the support of γ is sufficiently small, a property rule is optimal (specifically, the optimum is $b^D = 0$ if $E\gamma > \gamma^*$ and $b^D \geq \bar{b}^{prohib}$ if $E\gamma < \gamma^*$). (ii) If the support of γ is sufficiently large, the optimum is a liability rule. Moreover, the optimal b^D is lower than the exporter's loss from protection: $0 < b^D < \gamma^* < \bar{b}^{prohib}$.*

We have used the support of γ as a measure of ex-ante uncertainty. If uncertainty about γ is small in the sense of small variance but with a large support, then the optimum will not be *exactly* a property rule, but the result will hold in an approximate sense, so the qualitative insight goes through.

Proposition 1 states that a liability rule is optimal only if uncertainty about γ is large, and even in this case, the optimal level of damages b^D is lower than the level that makes the exporter "whole," i.e. γ^*. This result qualifies the presumption from the law-and-economics literature (e.g., Kaplow and Shavell, 1996) that the efficient level of breach damages is the one that makes the injured party whole; and this qualification arises even under the conditions that are most favorable to this argument, namely that γ^* is known to the DSB. The source of this qualification comes from our assumption of costly ex-post transfers, and so it applies with particular force to international dispute resolution. Specifically, in the context of the WTO where the damages paid for breach often take the form of counter-retaliation on the part of the injured party, the available means of compensation are inefficient, and therefore from an ex-ante perspective they should not be utilized to an extent that makes the injured party whole. This qualification gains special relevance in light of the emphasis placed on reciprocity in the GATT/WTO system of remedies: as noted in the Introduction, it is sometimes suggested that reciprocity falls short as a mechanism for facilitating efficient breach because it does *not* make the injured party whole, but Proposition 1 suggests that this may in fact be a desirable feature of reciprocity.

Moreover, as Proposition 1 indicates, if uncertainty about γ is sufficiently small, *any* liability rule is suboptimal (let alone the specific one with $b^D = \gamma^*$), and instead the optimum is a property rule. Intuitively, if uncertainty is low and the joint benefits of free trade are never very far from zero, the overriding efficiency concern is to avoid costly transfers; getting the correct policy choice in each state of the world is secondary. In these circumstances, assigning a property right to either party provides a disagreement point from which settlement occurs without transfers, and hence a property rule is preferable to a liability rule.[12]

Finally, note that Proposition 1 suggests a pair of empirical predictions. First, if uncertainty about the joint economic/political benefits of free trade decreases over time, then the optimal contractual/institutional arrangement should tend to move away from liability rules and towards property rules. And second, we should tend to observe more liability rules in issue areas where uncertainty about the joint economic/political benefits of free trade is larger; and conversely, the use of property rules should be more frequent in issue areas where this uncertainty is smaller.

We next turn to the role of transaction costs in determining the optimal type of agreement. Recalling that the cost of transfers plays the role of transaction costs in our model, consider

[12] Recall that property rights are sold ex-post in exchange for transfers only if γ takes very low or very high values (i.e. in regions I or IV), whereas a liability rule induces transfers in equilibrium in any region.

an increase in $c(b)$ for all $b \neq 0$ (while preserving the properties of $c(b)$ that we have assumed), fixing the support of γ. It is clear by inspection of Figure 1 that Regions II and III expand, while Regions I and IV contract, and at some point Regions I and IV disappear. Using similar arguments to those presented above, it then follows that when the cost of transfers is sufficiently high a property rule is optimal. Next consider decreasing the cost of transfers: as Figure 1 indicates, Regions I and IV expand, while Regions II and III contract, and as the cost of transfers goes to zero, the probability of being in Regions I and IV must become strictly positive. Again using similar arguments to those presented above, it then follows that if the cost of transfers is small enough, the optimum is a liability rule.

The following proposition states the result:

Proposition 2. *(i) If the cost of transfers is sufficiently high, a property rule is optimal. (ii) If the cost of transfers is sufficiently small, the optimum is a liability rule.*

Proposition 2 implies that a property rule tends to be preferred to a liability rule when transaction costs are high. This result stands in contrast with the finding in the law-and-economics literature that liability rules tend to be preferable to property rules when transaction costs are high (Calabresi and Melamed, 1972, and Kaplow and Shavell, 1996). Our result differs from this earlier finding because of our focus on the cost of transfers, which as we have indicated is an important transaction cost in the international government-to-government setting. To gain further intuition about this difference in results, recall that transaction costs in Calabresi and Melamed (1972) and Kaplow and Shavell (1996) take the form of bargaining frictions (the bargain fails with a certain probability); this type of transaction costs penalizes property rules more than liability rules because property rules induce more bargaining in equilibrium. In our setting, on the other hand, the presence of a transfer cost penalizes a liability rule more than a property rule because a liability rule induces more transfers in equilibrium.

Even more surprisingly in light of the Calabresi and Melamed (1972) and Kaplow and Shavell (1996) finding, we now show that in the presence of a cost of transfers, higher transaction costs can favor property rules even if these transaction costs take the form of frictions in bargaining. Specifically, we compare the case of frictionless ex-post bargaining (which we have just considered) with the opposite extreme in which ex-post bargaining is not feasible.

As a preliminary observation, consider what happens if transfers are costless: in this case, with frictionless ex-post bargaining, liability rules are equivalent to property rules, because the first best is achieved in all cases; while if ex-post bargaining is not feasible, it is easy to show that the unique optimum is a liability rule with $b^D = \gamma^*$ (i.e. the exporter must be made "whole"). Thus we can say that, if transfers are costless, removing the possibility of ex-post bargaining favors liability rules, and in this sense we can "replicate" the finding of the law-and-economics literature. However, we now show that this result is reversed if transfers are costly.

Recall from Proposition 1(ii) that if the support of γ is sufficiently large, then the optimum is a liability rule. But when ex-post bargaining is not feasible, it is easily established that the property rule $b^D = 0$ can become optimal in the large-γ-support case (with no impact in the reported results of Proposition 1(i) for the small-γ-support case). To understand, it is helpful to return to Figure 4 and recall why $b^D = 0$ cannot be optimal in the large-γ-support case when ex-post bargaining is frictionless: as we described above, and as the top right panel in Figure 4 indicates, joint surplus may be raised for state realizations in Region I by increasing b^D slightly above zero and thereby requiring some payment for breach, while joint surplus in Regions II-IV is unaffected (to the first order) by this maneuver. But notice from the top left panel of Figure 4 that if ex-post bargaining were not feasible, the outcome in Region I would entail $T = P$ for any non-prohibitive level of b^D, and increasing b^D slightly above zero would then have *no* (first order) impact on joint surplus for state realizations in Region I (while joint surplus in Regions II-IV would remain unaffected to the first order). Hence, when ex-post bargaining is not feasible, the property rule $b^D = 0$ is a local optimum in the large-γ-support case, and it is easily shown that it is also the global optimum for a region of parameters.

We summarize this discussion with:

Proposition 3. *If ex-post bargaining is not feasible, then: (i) If the support of γ is sufficiently small, a property rule is optimal; (ii) If the support of γ is sufficiently large, the optimum can be either a liability rule $b^D \in (0, \gamma^*)$ or the property rule $b^D = 0$.*

When compared to Proposition 1, Proposition 3 reveals that removing the possibility of ex-post bargaining expands the parameter region where a property rule is optimal, by opening up the possibility that a property rule may be optimal even in the case of large uncertainty. Intuitively, if uncertainty is large and ex-post bargaining is feasible, the reason a liability rule is preferable to a property rule is that the former induces lower transfers in the ex-post bargain. When the possibility of ex-post bargaining is removed, however, this advantage of a liability rule is removed as well, because in this case property rights do not induce any ex-post transfers, whereas a liability rule still does.

Together with Proposition 2, then, our model suggests that higher transaction costs (either in the form of cost of transfers or bargaining frictions) should tend to favor property rules over liability rules. This in turn points to an interesting empirical implication: we should tend to observe more property rules in issue areas where transaction costs are higher; and if transaction costs rise over time the optimal contractual/institutional arrangement should tend to move away from liability rules and towards property rules.[13]

3.2. Noisy DSB investigations

We now turn to the case where the DSB, if invoked, can observe a noisy signal of γ, which we denote by $\hat{\gamma}$. As we mentioned above, here we can consider a wider class of contracts, where b^D can be contingent on $\hat{\gamma}$. And as we indicated, this scenario is also interesting because governments are then uncertain about the DSB ruling at the ex-post negotiation stage, and this has important implications.

We impose a minimum of structure on the signal technology, by requiring that the conditional density of γ given $\hat{\gamma}$, denoted $h(\gamma|\hat{\gamma})$, is log-supermodular. This condition is relatively standard and is satisfied by several common distributions (see Athey, 2002). For future reference, we also let $g(\hat{\gamma}|\gamma)$ denote the conditional density of $\hat{\gamma}$ given γ.

Since at the ex-post negotiation stage governments face uncertainty over what would happen if the exporter invoked the DSB, the backward induction analysis in this scenario is more involved than it was under our benchmark scenario. For tractability here we impose a linear cost of transfers: $c(b) = c \cdot |b|$. The reason this assumption simplifies the analysis is that, as we establish below, the problem of finding the b^D schedule that maximizes the ex-ante joint surplus is equivalent to a simpler problem, namely finding the level of b^D that maximizes the

[13]This suggests an intriguing possibility: if an importing government wishing to increase protection in one industry can find a relatively low-cost way of compensating foreign exporters by offering a tariff reduction in another industry when that tariff is set at a high initial level, and can thereby avoid altogether the more costly counter-retaliation method of compensation, then it could be argued that as negotiated tariff cuts become deeper and this maneuver becomes more difficult, the cost of transfers becomes greater and property rules become more attractive (see Pelc, 2009, note 37 for a related observation). On the other hand, there are other forces that cut the opposite way: for example, the efficiency costs associated with counter-retaliation itself would be lower when the counter-retaliation begins from tariffs that are closer to their efficient levels, and for this reason the cost of transfers could fall as negotiated tariff cuts become deeper. Also, it is not clear how bargaining frictions (for example stemming from private information) have evolved over time or vary across issue areas. On balance, then, it is difficult to say without further structure how transaction costs might change over time or across issues, and so the empirical predictions of our model in this regard remain unclear.

expected joint surplus as viewed from stage 4, when the true γ is unknown, but conditional on observing a signal $\hat{\gamma}$. With a nonlinear cost of transfers, this equivalence need not hold, and the problem is more complex. We leave the analysis of the more general case for future research (but we believe that our qualitative insights continue to hold).

To state this result, we let $\Omega_4(b^D, \gamma)$ denote the joint payoff in the stage-4 subgame for a given b^D and realized γ, and we define the expected joint surplus as viewed from stage 4, when the realized γ is unknown, but conditional on observing a signal $\hat{\gamma}$, as $E[\Omega_4(b^D|\hat{\gamma})] = \int \Omega_4(b^D, \gamma)h(\gamma|\hat{\gamma})d\gamma$. We may now state:

Lemma 1. *If $c(b^D) = c \cdot |b^D|$, then the ex-ante optimal $b^D(\hat{\gamma})$ maximizes $E[\Omega_4(b^D|\hat{\gamma})]$.*

Armed with Lemma 1, we can now characterize the qualitative properties of the optimal damages level b^D. It is convenient to start with the case of large uncertainty. We assume here that $h(\gamma|\hat{\gamma})$ has full support, that is $[0, \infty)$, for any $\hat{\gamma}$. We also assume $\lim_{\hat{\gamma}\to\infty} \Pr(\gamma < \gamma^*|\hat{\gamma}) = 0$ (for any fixed γ^*): in words, the posterior probability that $\gamma < \gamma^*$ approaches zero as the signal realization $\hat{\gamma}$ becomes infinitely large. The following proposition characterizes how the optimal damages level b^D depends on the signal realization $\hat{\gamma}$ and on the (commonly known) exporter's loss from protection γ^*:

Proposition 4. *Assume $h(\gamma|\hat{\gamma})$ has full support for any $\hat{\gamma}$. Then: (i) b^D is (weakly) decreasing in $\hat{\gamma}$, with $\lim_{\hat{\gamma}\to\infty} b^D = 0$; and (ii) b^D is (weakly) increasing in γ^*.*

According to Proposition 4, the level of damages should be higher when, other things equal, the signal of the importer's gain from protection is lower, or the exporter's loss from protection is higher. Moreover, if the former is very high relative to the latter (i.e., as $\hat{\gamma} \to \infty$ for fixed γ^*), the damages for breach that the importer should be required to pay approach zero.

Notice an interesting feature of this result: contrary to the standard logic of efficient breach whereby damages should reflect only the level of harm caused by the breach, Proposition 4 implies that the damages for breach should be responsive to both the level of harm that the breach causes the exporter *and* the level of benefit that the breach is estimated to provide the importer. Intuitively, since it is not optimal to set damages at a level that fully compensates the exporter, making the damages sensitive to the estimated benefit that the importer gains from breach helps to ensure that breach will occur only when it is likely to be efficient.

This feature in turn suggests a possible interpretation of the WTO Agreement on Safeguards. According to the Safeguards Agreement, an importing government may temporarily impose tariffs as a response to injury to its domestic import-competing producers, and need not compensate the impacted foreign exporters under certain conditions which suggest a more direct link between the injury and foreign exports.[14] If these conditions are interpreted as indicating circumstances in which the imposition of tariffs would be an effective way of addressing the injury, then under these conditions the benefits of protection to the importer could reasonably be thought to be high. And if the DSB receives a signal that this is indeed the case, then Proposition 4 would suggest that the level of damages should be low or even approach zero, broadly in line with what the Agreement on Safeguards stipulates in this circumstance.

Next we focus on the case of small uncertainty. We consider two distinct ways in which uncertainty can be small: first we consider the case in which there is small ex-ante uncertainty about γ, just as in the previous section; and second, we consider the case in which the noise in the signal $\hat{\gamma}$ is small, so that *ex-post* uncertainty in γ (i.e. conditional on $\hat{\gamma}$) is small.

Let us start with the case of small ex-ante uncertainty about γ. It is a simple corollary of Proposition 1(i) that, if the support of the marginal distribution of γ is sufficiently small, the optimum is either the property rule $b^D = 0$ or the property rule $b^D \geq \bar{b}^{prohib}$, with the choice between the two contingent on the realized value of the signal $\hat{\gamma}$.

Next we consider the case in which the DSB information is very precise. If the support of $\hat{\gamma}$ conditional on γ is sufficiently small, we can apply a logic similar to that which led to Proposition 1(i) to conclude again that a contingent property rule is optimal, as the following proposition states:

Proposition 5. *If the signal observed by the DSB is sufficiently precise, in the sense that the support of $g(\hat{\gamma}|\gamma)$ around γ is sufficiently small, then the optimum is a property rule (contingent on $\hat{\gamma}$).*

The analysis of this section suggests two broad implications, one concerning the variation of optimal rules across issues, and one concerning the evolution of optimal rules over time. Starting

[14] In particular, the Safeguards Agreement specifies that no compensation need be paid by the importing government for 3 years when reimposing protection in response to injury of its domestic import-competing industry, provided that the injury is associated with an *absolute* increase in imports; whereas if the level of imports has fallen but the level of domestic production has fallen by more, so that injury is associated with an increase in imports only *relative* to domestic production, trade protection can still be reimposed but the importing government must compensate the impacted exporters from the start (in either case, this compensation may take the form of the temporary withdrawal of equivalent concessions by the exporter).

with the cross-sectional implication, our model suggests that the use of property rules should be more frequent (other things equal) in issue areas where the accuracy of the information that the DSB can gather ex post is higher, whereas liability rules should be more frequent in issue areas where the DSB is not likely to be well informed.

Concerning the evolution of the optimal rules over time, our results suggest that, if the accuracy of DSB rulings increases, the optimal institutional arrangement should move away from liability rules and toward property rules.[15] And as we have previously noted, our model suggests a similar evolution if ex-ante uncertainty about the joint benefits of free trade is reduced. Thus, if one accepts that the accuracy of legal rulings has increased from the time of GATT's inception to the creation of the WTO, and/or that the degree of ex-ante uncertainty about the joint benefits of free trade has diminished over this period, then we may ask whether or not the evolution from GATT to the WTO has indeed been in the direction away from liability rules and toward property rules. Here opinions differ among legal scholars, and we do not take a stand on the merits of the different views that have been expressed.[16] Rather, we simply note that the implications of our model suggest plausible circumstances under which such an evolution would be desirable from an institutional-design perspective.

4. The Outcome of Disputes

In this section we consider the implications of our model for the outcome of trade disputes. To this end, we first link more directly the stages of our game with the stages of a WTO dispute.

Broadly speaking, there are three phases to a WTO dispute. In a first phase, the complainant must request consultations with the respondent. If consultations fail to settle the dispute within

[15] Here it is interesting to note that Pelc (2009), who maintains that the GATT/WTO escape clause has shifted away from a liability-rule approach towards a property-rule approach, attributes this shift to an improved ability of the DSB to verify that a legitimate circumstance for escape has arisen, broadly in line with what would be expected based on our analysis.

[16] On the one hand, Jackson (1997, pp. 62-63) expresses the view that, while the early GATT years were ambiguous on this point, "...by the last two decades of the GATT's history..., the GATT contracting parties were treating the results of an adopted panel report as legally binding...," and that the WTO "...clearly establishes a preference for an *obligation to perform* the recommendation..." (emphasis in the original). On the other hand, Hippler Bello (1996) and Schwartz and Sykes (2002) view the changes in the DSB that were introduced with the creation of the WTO differently. According to Schwartz and Sykes, the GATT was devised to operate according to a liability rule that permitted efficient breach, where the penalty for breach in practice took the form of unilateral retaliation, but in the GATT's final years unilateral retaliation became excessive and discouraged efficient breach. The changes in the DSB that were introduced with the creation of the WTO were motivated, according to Schwartz and Sykes, by a need to *reduce* the penalty for breach, thus returning the system to one based squarely on liability rules.

60 days of the request, then the complainant may request that a Panel be established. In a second phase, the Panel gathers information on the dispute and issues a ruling which may be appealed to the Appellate Body, leading to a final ruling. And in a third phase, governments may engage in negotiations over the extent and modalities of compliance with the DSB ruling (with a "compliance panel" available in case of further disagreements).

Below we seek to develop the predictions of our model, and at a broad level match these predictions to the various possible outcomes under WTO-like contracts and dispute settlement procedures. To this end we now offer interpretations of model outcomes in terms of observable outcomes of the WTO dispute settlement procedures.

Let us consider first the interpretation of stages 2 and 3 in our model. Given that the WTO DSB requires that governments "consult" prior to requesting that a formal dispute Panel be formed for the purpose of issuing a ruling, it is natural to think of the consultation phase of the WTO dispute settlement process as being reflected in a stage 2 negotiation. The interpretation of stage 3 of our model seems equally straightforward: it is natural to think of a stage-3 ruling by the DSB as corresponding to the issuance of the Panel/Appellate Body final ruling.

Next, we turn to the interpretation of stage 4, and in particular the difference between the outcome where the DSB ruling is implemented and the outcome where the DSB ruling is renegotiated. In the former case, the DSB ruling defines a disagreement point for the subsequent negotiations which is on the Pareto frontier, and so there is nothing to gain from renegotiating the DSB ruling. In the latter case, the DSB ruling defines a disagreement point that is inside the Pareto frontier, and so in this case renegotiations take place: in particular, the DSB announces a breach payment under which (i) the home country would prefer to choose P and make the DSB-mandated breach payment rather than the alternative of FT with no payment, but (ii) the home country would prefer a third alternative to the two choices under the DSB ruling, namely, a policy of FT combined with a payment from the exporter. In this light, it seems natural to interpret a renegotiation that occurs in stage-4 as corresponding to a settlement in which the appropriate level of compensation is worked out between the disputants prior to the importer agreeing to bring its policies into compliance by adopting FT.[17]

[17] At a broad level, the renegotiation of DSB rulings in our model corresponds roughly to the determination of a "mutually agreed solution" as provided for under Article 3.6 of the WTO Dispute Settlement Understanding (see, for example, the mutually agreed solution to the U.S.-Canada lumber dispute reached on October 12, 2006 and notified under Article 3.6). At a more specific level, a good illustration is provided by the compliance settlement for the U.S.-EU "Banana" dispute in the WTO (see USTR, 2001). In reaching a settlement for this dispute, the EU (respondent) stated on April 11, 2001, when the dispute was settled/resolved, that it would

4.1. Predictions from the basic model

Having described the broad link from our model outcomes to stages of WTO disputes, we now return to the formal analysis of our model. There are three possible model outcomes to consider: (i) *Early Settlement*, which occurs when the importer's offer at stage 2 is accepted; (ii) *DSB is Invoked and the DSB Ruling is Implemented*; and (iii) *DSB is Invoked but the DSB Ruling is Renegotiated*.[18] We focus in this and the next subsection on the case in which the DSB can observe a noisy signal of γ, and we keep the assumptions we made in section 3.2.

A first observation is that, for realizations of γ in Regions I and IV governments settle early, while for realizations of γ in Regions II and III governments go all the way to a DSB ruling. With this observation, we may state:

Remark 1. *Early settlement occurs if γ is very low or very high, while a DSB ruling is triggered for intermediate values of γ.*

The arguments that establish the first part of this Remark are straightforward. Extreme values of γ correspond to Regions I and IV, and in these regions stage-2 uncertainty about the DSB's signal realization (and hence level of damages) does not place the disagreement point above the Pareto frontier, as the top left and bottom right panels of Figure 5 confirm ($E[D]$ denotes the disagreement point): as a result, governments have no reason to seek a ruling. Intuitively, when the joint surplus associated with FT is either very large and positive or very large and negative, the equilibrium policy choice is independent of the level of damages determined by the DSB, and so governments have nothing to gain by seeking a DSB ruling.

More subtle is the reason why the DSB is invoked in equilibrium for intermediate values of γ. First observe that intermediate values of γ correspond to Regions II and III, where the joint surplus associated with FT may be positive or negative but it is moderate in size. For this reason, the Pareto frontier is convex and the equilibrium policy *does* depend on the level of damages determined by the DSB, as the top right and bottom left panels of Figure 5 confirm.

come into compliance with the DSB ruling, but not fully until January 1, 2006. Hence, during this intervening period, the United States (a claimant) – by accepting the EU's non-to-partial compliance over this period – essentially allowed the EU to take some compensation (by being able to unilaterally deviate from its WTO commitment over this period) in exchange for the promise by the EU to fully comply by January 1, 2006. We thank Chad Bown for pointing us to this dispute as a suggestive illustration of our specific model result.

[18] By construction, in our model governments always engage in stage-2 "consultations," and for this reason, we focus on the model's predictions concerning early settlement and renegotiation of DSB rulings. Our model could be extended to consider the issue of whether or not governments initiate consultations; a natural possibility in this regard would be to introduce a cost of consultation. We leave this extension to future research.

The next step is to understand why a convex frontier leads to a DSB ruling in equilibrium. Graphically, given stage-2 uncertainty about the DSB's signal realization (and hence level of damages), the disagreement point is above the stage-2 Pareto frontier, as the top right and bottom left panels of Figure 5 confirm, and hence the importer prefers to trigger a DSB ruling rather than settle early. To gain a more direct intuition for this insight, consider the extreme case in which c is infinite, so that transfers are not feasible. Then the frontier is made of two points, P and FT, and any payoff combination between those two points is not feasible. In this case, invoking the DSB brings about an (expected) payoff combination that lies between points P and FT, due to the random nature of the DSB ruling; and since this is the disagreement point, there is no scope for early settlement. In essence, then, the role of the DSB ruling is analogous to that of a transfer, in that it makes feasible certain intermediate payoff combinations that would not otherwise be feasible.[19]

Remark 1 highlights ex-post conditions under which governments either settle early or pursue a dispute through to the ruling stage. But it is also interesting to examine the ex-ante probability of early settlement versus DSB ruling. To this end, note that $\Pr(Settlement) = 1 - \Pr(Ruling)$. Thus we can focus on the determinants of $\Pr(Settlement)$, as the determinants of $\Pr(Ruling)$ are mirror images of the former.

If we define an increase in uncertainty over γ as a mean-preserving spread of its distribution, and assume for simplicity that $E\gamma = \gamma^*$, a direct implication of the arguments made above is:

Remark 2. *As uncertainty over γ increases and/or the cost of transfers decreases, the probability of early settlement increases.*

The intuition for Remark 2 is similar to Remark 1, and can be understood again with the aid of Figure 5. In particular, as long as there is any uncertainty in the DSB ruling, i.e. b^D depends on $\hat{\gamma}$, there is settlement in equilibrium *if and only if* γ falls in Regions I or IV. But the probability of Regions I and IV combined is higher when γ is more uncertain and/or when the cost of transfers is lower, and the result then follows.

[19] It is worth emphasizing that the prediction that disputes *ever* proceed to a ruling (i.e., "go to court") – and hence the ability to make statements about when early settlement is likely to occur – distinguishes our model from much of the law-and-economics literature concerned with liability rules versus property rules. For example, Kaplow and Shavell (1996) consider the case of a perfectly uninformed and perfectly informed court, but they do not consider the case of an imperfectly informed court (our case of noisy DSB investigations) and so disputes are always settled early in their analysis.

We next turn to consider the probability that a DSB ruling is implemented versus renegotiated. We focus on ex-post conditions under which rulings are implemented or renegotiated:

Remark 3. *Conditional on a DSB ruling being triggered, the ruling is renegotiated when the DSB-assessed damages $b^D(\cdot)$ fall in an intermediate range, while it is implemented when $b^D(\cdot)$ takes very low or very high values.*

The intuition for this Remark is as follows. First, recall from Remark 1 that a ruling is reached in equilibrium only for realized γ in Regions II and III. Second, recall that renegotiation of the DSB ruling occurs when (i) Home would prefer to choose P and make the DSB-mandated breach payment rather than the alternative of FT with no payment, but (ii) Home would prefer a third alternative to the two choices under the DSB ruling, namely, a policy of FT combined with a payment from the exporter. And finally note that, as Figure 6 confirms, for Regions II and III this occurs for intermediate levels of the damages $b^D(\cdot)$. Hence, according to Remark 3, DSB rulings should be renegotiated when the DSB issues a "close" ruling, i.e., a ruling that does not suggest either very high or very low joint surplus associated with the FT policy.

Note also an interesting implication of Remark 3: it may well happen that compliance with the DSB ruling becomes an issue and the ruling is ultimately renegotiated even though the DSB ruling "gets it right" (i.e. $\hat{\gamma}$ is close to γ). In other words, our model indicates that the renegotiation of DSB rulings does not come about because rulings are "bad."

4.2. Differential cost of transfers between developed and less-developed countries

As illustrated by Remark 2 above, an interesting feature of our model is that the cost of international transfers can have important implications for predictions concerning the outcomes of disputes. If we introduce the further assumption that the cost of granting a (positive) transfer is higher for less-developed countries than it is for developed countries, then our model generates predictions for the variation in outcomes that arise when disputes are between two developed countries, between two developing countries, or between a developed and a developing country.[20] Here we consider the model's predictions regarding a dispute between a developed and a less-developed country, assuming that the cost of transfers c is higher for the latter. For this discussion we also assume for simplicity that the distribution of γ is symmetric about γ^*.

[20] Horn and Mavroidis (2008) document the interesting variation in outcomes of WTO disputes depending on the developed/less-developed status of the disputants.

We first observe that, under these assumptions, the model predicts that if the developed country is the respondent (importer), then with relatively high probability we are in Regions II or IV. This is because, as can be confirmed by graphical inspection, if the developed country is the importer then Region II consists of an interval of γ that is larger than Region III, and similarly Region IV consists of an interval of γ that is larger than Region I, owing to the relatively low (high) cost of transfers for the developed (less-developed) country. On the other hand, if the less-developed country is the respondent (importer), then the model predicts that with relatively high probability we are in Regions I or III, for analogous reasons.

These observations carry several implications. First we consider the implications of asymmetric costs of transfer for the observed outcomes of disputes between developed and developing countries under early settlement. We may state:

Remark 4. *Conditional on the developed country being the respondent, early settlements result in a policy of P (with compensation paid by the developed country to the developing country) with probability higher than $1/2$. Conditional on the less-developed country being the respondent, early settlements result in a policy of FT (with compensation paid by the developed country to the developing country) with probability higher than $1/2$.*

Thus, according to Remark 4, there is a tendency for developed countries to end up imposing more protection in equilibrium as a result of early settlements than is the case for less-developed countries. This Remark follows directly from the fact that if the developed country is the respondent, then the probability is higher that the realized Γ is in Region IV than that it is in Region I, and vice versa if the less-developed country is the respondent.

We next consider the outcomes of disputes between developed and developing countries that proceed all the way to a DSB ruling. We may state:

Remark 5. *Conditional on the developed country being the respondent, DSB rulings tend to occur when $\Gamma < 0$, that is, when FT is the first-best policy. Conditional on the less-developed country being the respondent, DSB rulings tend to occur when $\Gamma > 0$, that is, when P is the first-best policy.*

Thus, according to Remark 5 there is a pro-trade (anti-trade) selection bias in rulings when a developed country (less-developed country) is the respondent. This follows directly from the fact that if the developed country is the respondent, then the probability is higher that the realized Γ is in Region II than that it is in Region III, and vice versa if the less-developed country is the respondent.

5. Extensions

In this section we consider two extensions of our basic model. We first extend our analysis to a richer informational environment, where both γ and γ^* are non-verifiable. We then consider a broader class of contracts, where in addition to a "stick" associated with the choice of P (the damages level), the contract can also specify a "carrot" associated with the choice of FT.

5.1. Both γ and γ^* non-verifiable

Here we consider a scenario in which both γ and γ^* are non-verifiable. In particular, we now assume that γ and γ^* are ex-ante uncertain and independent of each other, with their ex-ante (common-knowledge) distributions denoted respectively by $h(\gamma)$ and $h^*(\gamma^*)$. Both governments observe γ and γ^* ex post, while the DSB does not. We also assume that γ and γ^* have the same support; note that this implies that the support of Γ includes the value $\Gamma = 0$.

We focus here on the case in which the DSB receives no information ex post. In this case, most of the analysis in section 3.1 continues to hold, provided our statements about uncertainty in γ are re-interpreted as statements about uncertainty in Γ (so, for example, the case in which the support of γ is large should be re-interpreted as the case in which the support of Γ is large).[21] But there is one important exception: in the case of large support of Γ, it may no longer be that $b^D \geq \bar{b}^{prohib}$ is dominated by some lower level of b^D.

To understand this point, it is first helpful to recall why it is that, when γ^* is known ex ante and the support of Γ is large enough so that Regions I and IV become relevant, it cannot be optimal to set $b^D \geq \bar{b}^{prohib}$. As our arguments leading up to Proposition 1 indicate, this follows because $b^{prohib}(\gamma) > \gamma^*$ for γ in Region IV but $b^{prohib}(\gamma) < \gamma^*$ for γ in the other regions, and so moving from a strict performance rule to a liability rule by decreasing b^D from $\bar{b}^{prohib}$ to a level slightly below γ^* strictly improves joint surplus in Region IV while not affecting joint surplus in the other regions. But when γ^* is also uncertain, this argument can only be made if Region IV is relevant for some γ when the realized γ^* takes its maximum value. And if the support of

[21] Our results for the case in which the DSB receives a noisy signal as reported in Propositions 4 and 5 and the Remarks also generalize to the setting in which both γ and γ^* are ex-ante uncertain, with the one exception being that b^D need not be (weakly) increasing in the signal $\hat{\gamma}^*$ unless the signal is sufficiently precise.

γ^* is wide enough relative to that of γ (which is the case under our assumption that γ and γ^* share the same support), Region IV cannot be relevant when the realized γ^* takes its maximum value; hence this argument does not apply.

As is true when γ^* is known ex ante, in the current setting where γ and γ^* are both uncertain, altering the degree of ex-ante uncertainty and altering the cost of transfers each has a similar effect on the optimal rules. We therefore have the following result:

Proposition 6. *Suppose that the DSB receives no information ex post, and that both γ and γ^* are uncertain. Then: (i) If the support of Γ is sufficiently small, or the cost of transfers is sufficiently large, a property rule is optimal (either $b^D = 0$ or $b^D \geq \bar{b}^{prohib}$). (ii) If the support of Γ is sufficiently large, or the cost of transfers is sufficiently small, the optimum can be either a liability rule ($0 < b^D < \bar{b}^{prohib}$) or the property rule $b^D \geq \bar{b}^{prohib}$.*

Two interesting aspects of Proposition 6 should be highlighted. The first is that the presence of uncertainty about the exporter's loss from protection (γ^*) further weakens the case for a liability rule, since now even in the case of large support of Γ a property rule may be optimal. And the second is that, together with Proposition 1, Proposition 6 and the arguments leading up to it suggest that the case for a liability rule is strongest when uncertainty is large *and* the importer faces asymmetrically large risk (so that most of the uncertainty about Γ is derived from uncertainty about γ, not γ^*).

Finally, in light of Proposition 6 it is interesting to revisit the effects of bargaining frictions in the current setting where both γ and γ^* are uncertain. Recall that, under the assumption that γ^* is known ex ante, a comparison across Propositions 1 and 3 establishes that removing the possibility of ex-post bargaining expands the parameter region where a property rule is optimal. However, if both γ and γ^* are uncertain, the introduction of bargaining frictions may or may not favor property rules. More specifically, when γ and γ^* are uncertain it is still true (as is the case when γ^* is known ex ante) that if ex-post bargaining is not feasible, the property rule $b^D = 0$ can be optimal in the large-Γ-support case. But when γ and γ^* are uncertain, it can also be established that the strict performance rule $b^D \geq \bar{b}^{prohib}$ *cannot* be optimal in the large-Γ-support case when ex-post bargaining is not feasible,[22] which is interesting because, as Proposition 6(ii) indicates, when ex-post bargaining is frictionless, it is the property rule

[22] To see this intuitively, consider Figure 4: as the bottom left panel indicates, when ex-post bargaining is not feasible, setting b^D at its prohibitive level simply forces $T = FT$, and the joint surplus is thereby discretely reduced, rather than discretely increased as Figure 4 indicates is the case when ex-post bargaining is frictionless.

$b^D = 0$ that cannot be optimal and the property rule $b^D \geq \bar{b}^{prohib}$ that *can* be optimal in the large-Γ-support case.

This discussion leads to the following point: when both γ and γ^* are uncertain, removing the possibility of ex-post bargaining *tilts* the use of property rules away from strict performance contracts ($b^D \geq \bar{b}^{prohib}$) and towards discretion ($b^D = 0$). It is not immediately obvious why bargaining frictions should work in a biased way against strict performance contracts relative to discretion. After all, while bargaining frictions reduce the opportunities for ex-post negotiation, this applies whether the ex-post negotiations begin from a position of $T = FT$ (as in strict performance contracts) or $T = P$ (as in discretion). What accounts for this bias is the fact that the elimination of ex-post bargaining opportunities pushes for a lower optimal b^D, both around b^D slightly above 0 and around $b^D = \bar{b}^{prohib}$, and this implies that $b^D = 0$ is more likely to be optimal and $b^D \geq \bar{b}^{prohib}$ is less likely to be optimal when ex-post bargaining is infeasible.

5.2. A More General Class of Contracts

Thus far we have restricted our analysis to a menu contract that gives the importing government a choice between setting P and compensating the exporting government with a payment b^D, or setting FT. In this section we consider a richer menu contract that allows the importer to choose between (P, b^D) and (FT, b^{FT}): that is, the importing government is given a choice between setting P and compensating the exporting government with a payment b^D, or setting FT and making the associated payment b^{FT}. Intuitively, in addition to the "stick" implied by the payment of damages $b^D > 0$ when the importer chooses P, it might be optimal to include a "carrot" implied by $b^{FT} < 0$ when the importer chooses FT. Notice, though, that the carrot, like the stick, is an ex-post transfer and hence costly in our model; and so it is not obvious that ex-ante efficiency would in fact be served by the inclusion of a carrot in the menu contract.

Let us continue to focus on the case (as in section 5.1) where both γ and γ^* are uncertain and the DSB receives no information ex post, so that b^D and b^{FT} must be noncontingent. Intuition for our findings can be developed by returning to Figures 3 and 4. Recall that Figure 3 depicts the case in which $\Gamma = 0$, which marks the border between Regions II and III. By inspection of the right panel of Figure 3, if the support of Γ around zero is sufficiently tight, then even with our more general class of contracts it is still optimal to adopt a property rule which either permits discretion or requires strict performance in all states of the world. But then, introducing a carrot ($b^{FT} < 0$) for FT could never be helpful, because it would simply introduce the equilibrium payment of a costly transfer ($b^{FT} < 0$) which would accompany FT when FT would have been chosen in equilibrium anyway and no transfer would have been paid.

Now consider Figure 4, which depicts the same information for Regions I and IV. As we have observed, when uncertainty over Γ is large, these regions also become relevant. And it again can be verified that when uncertainty over Γ is large the optimal level of b^D is strictly positive. But here, as can be confirmed using the top left panel of Figure 4, offering a carrot ($b^{FT} < 0$) for FT can now be beneficial, because when the realized Γ lies in Region I the DSB ruling is then renegotiated less often and the equilibrium transfer paid by the exporting government to the importing government is smaller as a result.

Noting again that changes in the degree of ex-ante uncertainty and changes in the cost of transfers have similar effects, the following result confirms the intuition developed above:

Proposition 7. *Suppose the DSB receives no information ex post, and consider menu contracts of the type $\{(P, b^D), (FT, b^{FT})\}$: (i) If the support of Γ is sufficiently small, or the cost of transfers is sufficiently large, it is optimal to set $b^{FT} = 0$, and the optimal level of b^D is either $b^D = 0$ or $b^D \geq \bar{b}^{prohib}$; and (ii) If the support of Γ is sufficiently large, or the cost of transfers is sufficiently small, the optimal levels of b^D and $|b^{FT}|$ are strictly positive.*

According to Proposition 7, if uncertainty over Γ is sufficiently small, then there is no gain in expanding the simple menu contract that we analyzed in previous sections to include the possibility of a carrot for FT. Intuitively, introducing a carrot induces ex-post transfers in equilibrium, which are costly, and provides no benefits. On the other hand, using a carrot can help if uncertainty in Γ is large, because in this case costly ex-post transfers occur in equilibrium, and the introduction of a carrot can then reduce the size of these transfers.[23]

It is interesting to consider whether this kind of carrot mechanism is observed in actual trade agreements such as the WTO. On the one hand, when a government agrees to reduce its tariffs as

[23] Observe that the contract class we consider in this section is equivalent to a revelation mechanism $T(\tilde{s}), b(\tilde{s})$, where $\tilde{s}$ is the importer's report of the state of the world $s = (\gamma, \gamma^*)$. This is not the most general class of mechanisms within our game, however. We have focused on contracts whereby only the importer makes a choice of policy/transfer or an announcement. Theoretically we could do better by setting up some kind of revelation game that involves also the exporter. For example, suppose that governments simultaneously announce the value of s to the DSB and, if the reports are different, both governments are hit with steep penalties. Clearly this kind of mechanism can implement the first best, because it is an equilibrium for the governments to reveal the true value of s. But we believe it is reasonable to abstract from this kind of mechanism, because in reality the WTO DSB does not have the power to impose penalties on governments for the policies they choose, let alone for the announcements they make. See Maggi and Staiger (2008) for a discussion of self-enforcement issues in related contexts.

a result of a trade negotiation, it typically considers this to be a concession that is only valuable to it in exchange for similar concessions from other governments. So it is clearly the norm for a government to receive some form of compensation from other governments when it agrees to a policy of free trade. According to this observation, the findings recorded in Proposition 7 could potentially be interpreted as suggesting a novel role played by the compensations for trade liberalization that we observe. But when interpreting the carrot $|b^{FT}|$, it must be remembered that this is an ex-post transfer, which is contractually specified to be executed after the state of the world γ has been observed as an additional (ex post) reward for contract performance. This, of course, rules out transfers that are made as part of an ex-ante negotiation. When put this way, it is less clear that the carrot-for-performance mechanism represented in Proposition 7 can be found along side the "stick" of damages-for-breach in existing trade agreements.

Overall, then, our consideration of this more general class of contracts raises some interesting new questions. Nevertheless, it is worth emphasizing that, while it may well be optimal under some circumstances to include a carrot b^{FT} in the contract, this does not invalidate our results from earlier sections regarding the optimal level of damages (b^D) and the conditions under which either liability rules or property rules would be most desirable.

6. Conclusion

In this paper, we analyze the optimal design of legal remedies for breach in the context of international trade agreements, with a particular focus on the GATT/WTO. Our formal analysis delivers sharp conclusions concerning the appropriate remedy for breach and optimal institutional design in light of features of the underlying economic and contracting environment. And our analysis also delivers novel predictions regarding when disputes arise in equilibrium, and how the disputes are resolved.

In order to preserve tractability and focus on the main points, we have made a number of strong assumptions, and it is important to extend our results to more general settings. For example, we have assumed that there is no cost to initiating disputes, but incorporating such costs into our model could yield some interesting additional predictions concerning the conditions under which disputes arise in equilibrium. Relatedly, introducing asymmetries in the costs of initiating disputes between large and small countries and/or developed and developing countries – possibly coupled with asymmetries in bargaining power – could yield insights regarding how the design of legal remedies for breach could best serve a heterogeneous membership.

We have made the strong assumption that there is no private information possessed by either government, an assumption that has helped to bring the distinctive features of our analysis into sharp relief. Allowing for privately informed governments would introduce an additional transaction cost in the form of a bargaining friction into our analysis; unlike the transfer costs that we have emphasized, such bargaining frictions are not specific to the international setting which is our focus, but they are surely important in real-world trade agreements.

Among the most interesting bargaining frictions from which we have abstracted is the possible hold-up problem that could arise for the government of an importing country under a property rule, when there are many exporting governments that hold an entitlement to its markets. As Schwartz and Sykes (2002) have argued, this consideration may be particularly relevant for the GATT/WTO in light of its nondiscrimination rules, and it weighs in favor of a liability-rule interpretation of GATT/WTO commitments. A formal analysis of this issue within our framework would require extending the model to a multi-country setting. We view this as a particularly important extension that we leave for future work.

And finally, we have assumed that DSB rulings are automatically enforced. This is a strong assumption, since in reality DSB rulings must be self-enforcing. Extending our analysis to a setting of self-enforcing agreements is bound to be a complex task, but we can make one simple point here. Suppose the importing country can in principle choose to deviate from – as distinct from renegotiate – the DSB ruling (e.g. choose policy P and make a payment lower than the DSB-mandated damages b^D), but this deviation can be met with a penalty. What is the optimal size of this penalty? In our model, the answer is simple: this penalty should be prohibitive, i.e., sufficiently high to deter this kind of breach in any state of the world. Thus, as a normative matter, our model suggests that the optimal penalties for breach may be non-prohibitive (i.e., induce breach in some states of the world) when it comes to breaching a contractually specified commitment, but should always be prohibitive when it comes to breaching a DSB ruling.

7. Appendix

Proof of Proposition 1: As we argued in the text, the equilibrium outcome of the stage-2 subgame is the same as that of the stage-4 subgame, so we can refer to "the" negotiation without specifying the stage in which the negotiation occurs. We start by describing how the outcome of the negotiation varies with γ for given b^D (notice that in the text we adopted a different perspective to develop intuition, and described how the outcome of the negotiation varies with b^D for given γ). We need to consider separately two cases: $b^D \leq \gamma^*$ and $b^D > \gamma^*$:

(a) If $b^D \leq \gamma^*$: It is convenient to think of the importer as having two threat points: FT with no transfer, and P with transfer b^D. For a generic value of γ, only one of the two threat points is "active": the one that gives the importer a higher payoff. Given b^D, there are two critical levels of γ, $J(b^D)$ and $R(b^D)$, with $J(b^D) < R(b^D)$, such that: (I) for $\gamma \in [0, J(\cdot)]$, the FT threat point is active, and the outcome is a policy of FT with no compensation paid by either party; (II) for $\gamma \in [J(\cdot), R(\cdot)]$, the (P, b^D) threat point is active, but the DSB ruling is renegotiated and the governments agree on a policy of FT and a transfer $b(b^D - \gamma^*)$ defined implicitly by $b^D - \gamma^* = b - c(b)$; and (III) for $\gamma > R(\cdot)$, the (P, b^D) threat point is active and the DSB ruling is not renegotiated, hence the importer chooses P and compensates the exporter with payment b^D. Note that, as γ crosses the level $J(b^D)$, the level of compensation "jumps" from zero to strictly negative. It can be shown that $J(0) = 0 < R(0) < \gamma^*$, and that J and R are increasing functions with $J(b^D) < R(b^D)$ for $b^D \in [0, \gamma^*)$ and $J(\gamma^*) = R(\gamma^*) > \gamma^*$.

(b) If $b^D > \gamma^*$: In this case, the only change relative to case (a) is that $J(b^D) > R(b^D)$, as can be confirmed by graphical inspection. As a consequence, the equilibrium outcome is as follows: for $\gamma \in [0, R(\cdot)]$, the FT threat point is active, and the equilibrium outcome is a policy of FT with no transfer; for $\gamma \in [R(\cdot), J(\cdot)]$, the FT threat point is active, but the DSB ruling is renegotiated and the equilibrium outcome is P with a transfer of γ^*; and for $\gamma > J(\cdot)$, the (P, b^D) threat point is active and the DSB ruling is not renegotiated, hence the equilibrium outcome is (P, b^D).

The next step is to derive the level of b^D that maximizes the expected joint payoff, which we denote $E\Omega(b^D)$ (with a slight abuse of notation). It is easy to show that $E\Omega(b^D)$ is differentiable in b^D. We have to distinguish between cases (a) and (b):

(a) If $b^D \leq \gamma^*$, we can write the ex ante joint payoff in the following way:

$$\begin{aligned} E\Omega(b^D) &= \int_0^{J(b^D)} V(FT)h(\gamma)d\gamma + \int_{J(b^D)}^{R(b^D)} [V(FT) - c(b(b^D - \gamma^*))]dH(\gamma) \\ &+ \int_{R(b^D)}^{\infty} [V(FT) + \gamma - \gamma^* - c(b^D)]dH(\gamma) \end{aligned}$$

where $V(FT) \equiv v(FT) + v^*(FT)$, and $H(\gamma)$ is the cdf of γ. Differentiating, we get

$$\frac{dE\Omega}{db^D} = -\frac{d}{db^D}c(b(b^D-\gamma^*))\cdot[H(R(b^D))-H(J(b^D))]-c'(b^D)\cdot(1-H(R(b^D))+J'(b^D)\cdot h(J(b^D))\cdot c(b(b^D-\gamma^*)) \tag{7.1}$$

Equation (7.1) can be understood by noting that a small increase in b^D affects the equilibrium outcome through its impact on the active threat point. For $\gamma < J(\cdot)$ a small increase in b^D has no effect because in this case the active threat point is FT, which is independent of b^D; for $\gamma \in (J(\cdot), R(\cdot))$, an increase in b^D leads to a reduction in the compensation that the exporter pays in equilibrium (whose initial level is $b(b^D - \gamma^*)$), and hence leads to cost savings of $-\frac{dc(b(b^D-\gamma^*))}{db^D} > 0$; for $\gamma > R(\cdot)$, an increase in b^D translates directly into cost savings of $c'(b^D)$; and finally, an increase in b^D results in a shift forward of the "jump" point $J(b^D)$, which in turn implies cost savings of $J'(b^D) \cdot c(b(b^D - \gamma^*))$. The cost changes just described are weighted by their respective probabilities (and by the density $h(J(\cdot))$ in the case of the jump point).

(b) If $b^D > \gamma^*$, we can write the ex ante joint payoff as

$$\begin{aligned} E\Omega(b^D) &= \int_0^{R(b^D)} V(FT)h(\gamma)d\gamma + \int_{R(b^D)}^{J(b^D)} [V(FT) + \gamma - \gamma^* - c(\gamma^*)]dH(\gamma) \\ &+ \int_{J(b^D)}^{\infty} [V(FT) + \gamma - \gamma^* - c(b^D)]dH(\gamma) \end{aligned}$$

Differentiating, we get

$$\frac{\partial E\Omega}{\partial b^D} = -c'(b^D)(1 - H(J(b^D)) + J'(b^D)h(J(b^D))[c(b^D) - c(\gamma^*)] \tag{7.2}$$

In this case, an increase in b^D has no effect if $\gamma < J(\cdot)$; it leads to an increase in the equilibrium transfer if $\gamma > J(\cdot)$, resulting in a cost increase of $c'(b^D)$; and it shifts the jump point forward, which implies cost savings of $J'(b^D)[c(b^D) - c(\gamma^*)]$.

We are now ready to consider the two cases of "large" and "small" support of γ. It is convenient to start with the case of large support of γ. It suffices to prove part (ii) of the proposition for the case of full support, i.e. $\gamma \in [0, \infty)$. Let us first show that the optimal b^D

is strictly lower than γ^*. To show this, note first that the value $b^D = \gamma^*$ weakly dominates all higher values of b^D. This can be easily seen from Figure 1: in Region IV, setting $b^D > \gamma^*$ is weakly dominated by setting $b^D = \gamma^*$, and in Regions I-III the joint payoff is constant for $b^D \geq \gamma^*$. Next note from (7.1) that $\frac{\partial E\Omega}{\partial b^D}|_{b^D=\gamma^*} = -c'(\gamma^*)(1 - H(R(\gamma^*)) < 0$ (where we have used the fact that $J(\gamma^*) = R(\gamma^*)$), and hence there is a strict gain from lowering b^D below γ^*.

Next let us consider whether it can be optimal to set $b^D = 0$. Evaluating (7.1) at $b^D = 0$ and recalling that $J(0) = 0 < R(0)$, we have $\frac{dE\Omega}{db^D}|_{b^D=0} = -c'(-\gamma^*)H(0) + \frac{\partial J}{\partial b^D}c(\gamma^*)h(J(0))$. Clearly, this derivative is positive.[24] Part (ii) of the proposition follows immediately.

Let us now consider the case of small support of γ (around γ^*). To start with, let us focus first on the knife-edge case $\gamma = \gamma^*$. Clearly, in this case $b^D = 0$ and $b^D \geq b^{prohib}(\gamma^*)$ both yield the first best outcome, and any other value of b^D is suboptimal. Moreover, note that the only other values of b^D that yield a joint payoff "close to" the maximum are those in a right neighborhood of $b^D = 0$; all other values of b^D yield a joint payoff that is discretely lower than the maximum (including those in a left neighborhood of $b^{prohib}(\gamma^*)$, because there is a jump at $b^{prohib}(\gamma^*)$).

Now consider a small support of γ around γ^*, say $(\gamma^* - \varepsilon_1, \gamma^* + \varepsilon_2)$. Focus first on values of b^D that are strictly positive but close enough to zero: clearly, for such values of b^D we have $\frac{\partial \Omega(b^D,\gamma)}{\partial b^D} < 0$ for all $\gamma \in (\gamma^* - \varepsilon_1, \gamma^* + \varepsilon_2)$, and hence no such value of b^D can be optimal. Next focus on a value of b^D that is not close to zero and that is lower than $b^{prohib}(\gamma^* - \varepsilon_1)$ (i.e. non-prohibitive for all values of γ): such a value of b^D is suboptimal, because by continuity it yields a joint payoff that is discretely lower than the maximum, for each value of γ in its support. Finally consider a value of b^D that is prohibitive for some values of γ but not for others, i.e. $b^D \in (b^{prohib}(\gamma^* - \varepsilon_1), b^{prohib}(\gamma^* + \varepsilon_2))$: such a value of b^D is clearly dominated by a fully prohibitive value, i.e. by $b^D > b^{prohib}(\gamma^* + \varepsilon_2)$, because of the jump that occurs at $b^{prohib}(\gamma)$. This establishes that in the case of small support only a property rule can be optimal. **QED**

Proof of Proposition 2: Consider increasing $c(b)$ for all $b \neq 0$ (while preserving the properties of $c(b)$ that we have assumed), fixing the support of γ. It is clear by inspection of Figure 1 that Regions II and III expand, while Regions I and IV contract, and at some point Regions I and IV disappear. With the cost-of-transfers sufficiently high so that only Regions II and III

[24] There is a small loose end here. The point $(\gamma = 0, b^D = 0)$ is a knife-edge point, because the importer is indifferent between the two threat points. If the indifference is broken in favor of P the term $\frac{\partial J}{\partial b^D}c(\gamma^*)h(J(0))$ will appear, otherwise it will not. But the result goes through in both cases.

are relevant, it then follows that there are no transfers in equilibrium under the property rules $b^D = 0$ and $b^D \geq \bar{b}^{prohib}$, while there are non-zero transfers in equilibrium under a liability rule. But then, the cost of transfers can always be made sufficiently high that their cost outweighs the gain in joint surplus associated with inducing the correct policy choice. More specifically, consider a level b^D that is strictly between 0 and $\bar{b}^{prohib}$. By inspection of Figure 2 it is clear that we can make the cost of transfers high enough that this level of b^D is dominated in terms of joint surplus by $b^D = 0$. This shows that the optimum must be a property rule if the cost of transfers is high enough.

Next consider decreasing the cost of transfers: as Figure 1 indicates, Regions I and IV expand, while Regions II and III contract, and as the cost of transfers goes to zero, the probability of being in either of these two regions approaches zero. Using similar arguments to those presented in the text (with the distinction noted that the probability of being in Regions II and III approaches zero), it then follows that if the cost of transfers is small enough, each of the property rules $b^D = 0$ and $b^D \geq \bar{b}^{prohib}$ can be improved upon, and hence it is optimal to adopt a liability rule. **QED**

Proof of Proposition 3: If the support of Γ is small, it can be shown using a similar argument as in the proof of Proposition 1 that a property rule is optimal. Focusing next on the case of large support, $\frac{dE\Omega}{db^D}$ can be written as:

$$\frac{dE\Omega}{db^D} = -c'(b^D)[1 - H(b + c(b^D))] - (1 + c'(b^D))h(b + c(b^D))(b^D - \gamma^*), \tag{7.3}$$

where we have used the fact that $J(b^D) = b + c(b^D)$. Clearly, for $b^D \geq \gamma^*$ we have $\frac{dE\Omega}{db^D} < 0$,and hence the optimal b^D must be lower than γ^*. Also note that $\frac{dE\Omega}{db^D}|_{b^D=0} = 0$ (since $c'(0) = 0$ and $h(0) = 0$), but $\frac{dE\Omega}{db^D}|_{b^D=\varepsilon} < 0$, where ε is an arbitrarily small but positive value. This establishes that $b^D = 0$ is a local maximum. It can be shown by examples that the globally optimal b^D can be either zero or positive. **QED**

Proof of Lemma 1: Consider an arbitrary schedule $b^D(\hat{\gamma})$. At stage 4 this schedule induces equilibrium payoffs $(\omega_4(\hat{\gamma}, \gamma), \omega_4^*(\hat{\gamma}, \gamma))$. Clearly, all of these payoff pairs lie on the ex-post Pareto frontier given γ, and can be characterized as we have done so above for given γ and a level of $\hat{\gamma}$ (and hence b^D). Moving back to stage 2, consider the expected payoffs conditional on γ if stage 4 is reached (i.e. if the DSB is invoked). We denote these expected payoffs as $(E[\omega_4(\hat{\gamma}, \gamma)|\gamma], E[\omega_4^*(\hat{\gamma}, \gamma)|\gamma])$. In this proof we omit the argument b^D from the payoff functions, as this should not cause confusion.

Let us consider the four possible regions of γ. With our assumption of a linear cost of transfers, the ex-post Pareto frontier in each region is now piece-wise linear: in the four panels of Figure 5 we display the ex-post frontiers for each of the four regions. As before, the shaded portion of each frontier depicts the range of stage-4 bargaining outcomes that are induced by varying $\hat{\gamma}$ – and hence b^D – given a representative realized γ in the respective region. In Regions II and III, as the top right and bottom left panels of Figure 5 indicate, the equilibrium payoff points $(\omega_4(\hat{\gamma},\gamma),\omega_4^*(\hat{\gamma},\gamma))$ constitute a convex locus in (ω,ω^*) space. This implies that the expected payoff point $(E[\omega_4(\hat{\gamma},\gamma)|\gamma],E[\omega_4^*(\hat{\gamma},\gamma)|\gamma])$, which we label by $E[D]$ in each panel, lies outside the Pareto frontier for given γ. As a consequence, there is no settlement at stage 2 for realized γ in these regions, and the DSB is invoked. Thus the equilibrium payoffs at stage 2 are given by $(E[\omega_4(\hat{\gamma},\gamma)|\gamma],E[\omega_4^*(\hat{\gamma},\gamma)|\gamma])$ for realized γ falling in either Region II or Region III. In Regions I and IV, the equilibrium payoff points $(\omega_4(\hat{\gamma},\gamma),\omega_4^*(\hat{\gamma},\gamma))$ lie on a single straight line for given γ, as the top left and bottom right panels of Figure 5 make clear. This implies that the expected payoff point $(E[\omega_4(\hat{\gamma},\gamma)|\gamma],E[\omega_4^*(\hat{\gamma},\gamma)|\gamma])$ lies on the Pareto frontier, as the points labeled $E[D]$ in these two panels indicate. And this implies again that the equilibrium payoffs at stage 2 are given by $(E[\omega_4(\hat{\gamma},\gamma)|\gamma],E[\omega_4^*(\hat{\gamma},\gamma)|\gamma])$.

Now let us consider the optimization problem at stage 0. The objective function is $E_\gamma(E[\omega_4(\hat{\gamma},\gamma)+\omega_4^*(\hat{\gamma},\gamma)|\gamma])$, which we can write as follows:

$$\int\left[\int(\omega_4(\hat{\gamma},\gamma)+\omega_4^*(\hat{\gamma},\gamma))h(\hat{\gamma}|\gamma)d\hat{\gamma}\right]h(\gamma)d\gamma=\int\left[\int(\omega_4(\hat{\gamma},\gamma)+\omega_4^*(\hat{\gamma},\gamma))h(\gamma|\hat{\gamma})d\gamma\right]z(\hat{\gamma})d\hat{\gamma}$$

where $z(\hat{\gamma})$ is the marginal density of $\hat{\gamma}$. Clearly, maximizing the objective boils down to maximizing $\int(\omega_4(\hat{\gamma},\gamma)+\omega_4^*(\hat{\gamma},\gamma))h(\gamma|\hat{\gamma})d\gamma$ for each given $\hat{\gamma}$. **QED**

Proof of Proposition 4: We start by writing the derivative $\frac{dE(\Omega_4|\hat{\gamma})}{db^D}$. Clearly we can focus on values of b^D that are lower than γ^*, and hence we can write $E[\Omega_4(b^D|\hat{\gamma})]$ as:

$$\frac{dE(\Omega_4|\hat{\gamma})}{db^D}=\frac{c}{1+c}J'(b^D)h(J(b^D)|\hat{\gamma})(\gamma^*-b^D)+\frac{c}{1+c}[H(R(b^D,\gamma^*)|\hat{\gamma})-H(J(b^D)|\hat{\gamma})]-c[1-H(R(b^D,\gamma^*)|\hat{\gamma})],$$

where we have used the fact that $c(b(b^D-\gamma^*))=\frac{c}{1+c}\cdot(\gamma^*-b^D)$ given a linear $c(\cdot)$.

We now argue that $\frac{dE(\Omega_4|\hat{\gamma})}{db^D}$ is decreasing in $\hat{\gamma}$ when evaluated at the FOC ($\frac{dE(\Omega_4|\hat{\gamma})}{db^D}=0$). Note that the first and second terms in $\frac{dE(\Omega_4|\hat{\gamma})}{db^D}$ are positive, while the third term is negative. Since $h(\gamma|\hat{\gamma})$ is log-supermodular, as $\hat{\gamma}$ increases, $h(\gamma|\hat{\gamma})$ increases proportionally more for higher values of γ. This implies that as $\hat{\gamma}$ increases the negative term increases proportionally more than the sum of the two positive terms. Coupled with the fact that, when evaluated at the FOC ($\frac{dE(\Omega_4|\hat{\gamma})}{db^D} = 0$), the negative term is equal in magnitude to the sum of the positive terms, this implies that as $\hat{\gamma}$ increases, the negative term increases in magnitude by more than the sum of the positive terms, and hence $\frac{dE(\Omega_4|\hat{\gamma})}{db^D}$ decreases. Given that we were starting from an optimal (interior) level of b^D, it follows that the optimal b^D decreases with $\hat{\gamma}$.

The second part of point (i) is a direct consequence of the assumption $\lim_{\hat{\gamma}\to\infty} \Pr(\gamma < \gamma^*|\hat{\gamma}) = 0$. It is clear that if $\Pr(\gamma < \gamma^*|\hat{\gamma}) = 0$, then $b^D = 0$ is optimal. One can then show that, if we make $\Pr(\gamma < \gamma^*|\hat{\gamma}) = 0$ slightly positive, the optimal b^D is close to (or equal to) 0.

To prove that the optimal b^D is increasing in γ^*, note: (i) $J'(b^D) > 0$; (ii) $R(b^D)$ is increasing in γ^*; and (iii) H is increasing. It follows that $\frac{dE(\Omega_4|\hat{\gamma})}{db^D}$ is increasing in γ^*. **QED**

Proof of Proposition 5: Omitted.

Proof of Proposition 6: The analysis of this case is analogous to the case where only γ is uncertain, except that we need to highlight how $J(\cdot)$ and $R(\cdot)$ depend on γ^*. It is not hard to show that R is increasing in γ^* and J is independent of γ^*. We therefore write $J(b^D)$ and $R(b^D, \gamma^*)$.

Given that both γ and γ^* are uncertain, we do not need to write two separate expressions of $E\Omega$ for the cases $b^D > \gamma^*$ and $b^D < \gamma^*$. Instead, we have a single expression that involves the expectation over γ^*, and we can write:

$$\frac{dE\Omega}{db^D} = J'(b^D)h(J(b^D))[\int_{b^D}^{\infty} c(b(b^D - \gamma^*))dH^*(\gamma^*) + \int_0^{b^D} (c(b^D) - c(\gamma^*))dH^*(\gamma^*)] \tag{7.4}$$

$$- \int_{b^D}^{\infty} [\frac{d}{db^D} c(b(b^D - \gamma^*)) \Pr(J(b^D)) < \gamma < R(b^D, \gamma^*))]dH^*(\gamma^*) - c'(b^D) \Pr(\gamma > \max\{R(b^D, \gamma^*), J(b^D)\})$$

In the case of small support of Γ, it can be shown using a similar argument as in the proof of Proposition 1 that a property rule is optimal. That the same result obtains if the cost of transfers is high enough (holding everything else constant) can be shown with a similar argument as in the proof of Proposition 2.

Let us focus now on the case of large support of Γ. We evaluate $\frac{dE\Omega}{db^D}$ at $b^D = 0$. Recall that $J(0) = 0$ and note that $h(0) = 0$, hence $h(J(0)) = 0$. And recalling that $\frac{d}{db^D} c(b(b^D - \gamma^*)) < 0$, it follows that $\frac{dE\Omega}{db^D}|_{b^D=0} > 0$. This implies that the optimal b^D is strictly positive, as claimed. It can be shown by examples that the optimal b^D can be either prohibitive or non-prohibitive. That the same result obtains if the cost of transfers is low enough (holding everything else constant) can be shown with a similar argument as in the proof of Proposition 2. **QED**

Proof of Proposition 7: It is easy to show that it cannot be optimal to set $b^D < 0$ or $b^{FT} > 0$, so we can focus on the case $b^D \geq 0$, $b^{FT} \leq 0$. To keep the notation more intuitive we think of the choice variables as being the absolute transfer levels, b^D and $|b^{FT}|$.

We start by describing how the outcome of the negotiation varies with γ and γ^* for given b^D and $|b^{FT}|$. We need to consider separately four regions in (γ, γ^*) space. Again we can think of the importer as having two threat points: (P, b^D) and $(FT, |b^{FT}|)$; for a generic value of γ, only one of the two threat points is active. Let us define $\gamma = J(b^D, |b^{FT}|)$ as the level of γ such that the importer is indifferent between the two threat points; as in the case where $b^{FT} \equiv 0$, it can be verified that J does not depend on γ^*. Also, in analogy with the case where $b^{FT} \equiv 0$, let $\gamma = R(\gamma^*, b^D, |b^{FT}|)$ denote the level of γ such that the active threat point is just on the ex-post Pareto frontier; this is the border of the region where governments renegotiate. Notice that (as before) R is increasing in γ^*. Finally, let $\gamma^* = \phi(b^D, |b^{FT}|)$ be the level of γ^* for which the curves $\gamma = J(b^D, |b^{FT}|)$ and $\gamma = R(\gamma^*, b^D, |b^{FT}|)$ intersect.

Consider now the four relevant regions in (γ, γ^*) space, for given b^D and $|b^{FT}|$: (A) In the region where $\gamma < \min\{J(\cdot), R(\gamma^*, \cdot)\}$, the $(FT, |b^{FT}|)$ threat point is active, and there is no renegotiation, so the outcome is a policy of FT with the exporter paying $|b^{FT}|$; (B) In the region where $J(\cdot) < \gamma < R(\gamma^*, \cdot)$, the (P, b^D) threat point is active, but the DSB ruling is renegotiated and the equilibrium outcome is a policy of FT with the exporter making a transfer to the importer, which we denote $b(b^D, \gamma^*)$; (C) In the region where $\gamma > \max\{J(\cdot), R(\gamma^*, \cdot)\}$, the (P, b^D) threat is active and there is no renegotiation, hence the outcome is a policy of P with the importer paying b^D; (D) In the region where $R(\gamma^*, \cdot) < \gamma < J(\cdot)$, the $(FT, |b^{FT}|)$ threat is active, but there is renegotiation and the equilibrium outcome is a policy of P with the exporter making a transfer, which we denote $b(|b^{FT}|, \gamma^*)$.

We next consider how $E\Omega$ depends on b^D and $|b^{FT}|$. Using a similar logic as in the proof of Proposition 6, we can write

$$\begin{aligned}
\frac{\partial E\Omega}{\partial b^D} &= -c'(b^D)\Pr(\gamma > \max\{J(\cdot), R(\gamma^*, \cdot)\}) - \frac{\partial c(b(b^D, \gamma^*))}{\partial b^D}\Pr(J(\cdot) < \gamma < R(\gamma^*, \cdot)) \\
&\quad + \frac{\partial J}{\partial b^D} h(J(\cdot))[\int_0^{\phi(\cdot)} [c(b^D) - c(b(|b^{FT}|, \gamma^*))]dH^*(\gamma^*) - \int_{\phi(\cdot)}^{\infty} [c(b(b^D, \gamma^*)) - c(-|b^{FT}|)]dH^*(\gamma^*)] \\
\frac{\partial E\Omega}{\partial |b^{FT}|} &= c'(-|b^{FT}|)\Pr(\gamma < \min\{J(\cdot), R(\cdot)\}) - \frac{\partial c(b(|b^{FT}|, \gamma^*))}{\partial |b^{FT}|}\Pr(R(\cdot) < \gamma < J(\cdot)) \\
&\quad + \frac{\partial J}{\partial |b^{FT}|} h(J)[\int_0^{\phi(\cdot)} [c(b^D) - c(b(|b^{FT}|, \gamma^*))]dH^*(\gamma^*) - \int_{\phi(\cdot)}^{\infty} [c(b(b^D, \gamma^*)) - c(-|b^{FT}|)]dH^*(\gamma^*)]
\end{aligned}$$

Let us first prove Proposition 7(ii). Focus on the case of large support of Γ. A necessary condition for $|b^{FT}| = 0$ to be optimal is $\left|\frac{\partial E\Omega}{\partial |b^{FT}|}\right|_{|b^{FT}|=0, b^D=\tilde{b}^D} \leq 0$, where $\tilde{b}^D$ is the optimal value of b^D conditional on $|b^{FT}| = 0$. Recall that $\tilde{b}^D \in (0, \bar{b}^{prohib})$ when the support of Γ is large enough. Note that (a) $\frac{\partial J}{\partial b^D} > 0$ and $\frac{\partial J}{\partial |b^{FT}|} > 0$; (b) $[c(b(b^D, \gamma^*)) - c(-|b^{FT}|)] > 0$ and $[c(b^D) - c(b(|b^{FT}|, \gamma^*))] > 0$; and (c) $\frac{\partial}{\partial |b^{FT}|} c(b(|b^{FT}|, \gamma^*)) < 0$. It is then direct to conclude that $\left|\frac{\partial E\Omega}{\partial |b^{FT}|}\right|_{|b^{FT}|=0, b^D=\tilde{b}^D} > 0$, and hence the optimal $|b^{FT}|$ is strictly positive. With a similar argument one can show that the optimal value of b^D is strictly positive.

The claim that the optimal b^D and $|b^{FT}|$ are strictly positive if c is small enough (holding everything else equal, including a finite support for Γ) can be shown with a similar argument as in the proof of Proposition 2.

We can now turn to claim (i) of Proposition 7. We focus on the case of small support of Γ. Consider first the knife-edge case $\Gamma = 0$, or $\gamma = \gamma^*$. Let us characterize how $\Omega(b^D, |b^{FT}|; \cdot)$ depends on b^D and $|b^{FT}|$. Clearly, there are two sets of points $(b^D, |b^{FT}|)$ that are optimal: (i) any pair such that $b^D = 0$ and $|b^{FT}| \leq \gamma_{\min}$ induces $T = P$ and $b = 0$, and hence is optimal; (ii) any pair such that $b^{FT} = 0$ and $b^D \geq \bar{b}^{prohib}$ induces $T = FT$ and $b = 0$, and hence is optimal.

Following a similar logic as in the proof of Proposition 6, we can next ask what other points $(b^D, |b^{FT}|)$ yield a joint payoff "close to" the first best in this knife-edge case. The answer is: (a) those such that $|b^{FT}| \leq \gamma_{\min}$ and $b^D > 0$ is close to zero, and (b) those such that $b^D > b^{prohib}(|b^{FT}|, \gamma_{\max})$ and $|b^{FT}| > 0$ is close to zero. All other pairs $(b^D, |b^{FT}|)$ yield a joint payoff that is discretely lower than the first best.

We are now ready to consider a small support of Γ around 0, say $(-\varepsilon_1, +\varepsilon_2)$. Focus first on pairs $(b^D, |b^{FT}|)$ of the type (a) described just above. For these pairs, $\frac{\partial \Omega}{\partial b^D} < 0$ for all $\Gamma \in (-\varepsilon_1, +\varepsilon_2)$, and hence no such pair can be optimal. Focus next on pairs $(b^D, |b^{FT}|)$ of the type (b) described above; for these pairs, $\frac{\partial \Omega}{\partial |b^{FT}|} < 0$ for all $\Gamma \in (-\varepsilon_1, +\varepsilon_2)$, and hence no such pair can be optimal. All other pairs must be suboptimal too, because by continuity they yield a joint payoff that is discretely lower than the first best for all $\Gamma \in (-\varepsilon_1, +\varepsilon_2)$, whereas we know that we can achieve a joint payoff close to the first best with, for example, $|b^{FT}| = 0$ and $b^D > \bar{b}^{prohib}$. This proves that, if the support of Γ is small enough, it is optimal to set $|b^{FT}| = 0$ and the optimal b^D is either zero or prohibitive.

That the same result obtains if the cost of transfers is high enough (holding everything else constant) can be shown with a similar argument as in the proof of Proposition 2. **QED**

References

Athey, Susan (2002), Monotone Comparative Statics Under Uncertainty," *Quarterly Journal of Economics*," vol. CXVII (1): 187-223.

Bagwell, Kyle and Robert W. Staiger (2005), "Enforcement, Private Political Pressure and the GATT/WTO Escape Clause," *Journal of Legal Studies*, June.

Bagwell, Kyle (2009), "Self-Enforcing Trade Agreements and Private Information," unpublished manuscript, Stanford University, February 15.

Baldwin, Richard, "Politically Realistic Objective Functions and Trade Policy," *Economics Letters*, vol. 24, pp. 287-90, 1987.

Beshkar, Mostafa (2008a),"Trade Skirmishes and Safeguards: A Theory of the WTO Dispute Settlement Process," unpublished manuscript, Purdue University, June.

Beshkar, Mostafa (2008b),"Optimal Remedies in International Trade Agreements," unpublished manuscript, Purdue University, May.

Busch, Mark L. and Eric Reinhardt (2006), "Three's a Crowd: Third Parties and WTO Dispute Settlement," *World Politics* 58, April, pp. 446-77.

Calabresi, Guido and A. Douglas Melamed (1972), "Property Rules, Liability Rules and Ineliability: One View of the Cathedral," *Harvard Law Review* 85(6), pp. 1089-1128.

Charnovitz, Steve (2003), "The WTO's Problematic 'Last Resort' Against Noncompliance," mimeograph, August 14.

Ethier, Wilfred J. (2001), "Punishments and Dispute Settlement in Trade Agreements," PIER Working Paper No. 01-021, May 15.

Grossman, Gene M. and Elhanan Helpman, "Protection for Sale," *American Economic Review*, 84(4), pp. 833-850, 1994.

Hippler Bello, Judith (1996), "The WTO Dispute Settlement Understanding: Less is More," *American Journal of International Law* 90.

Hoekman, Bernard and Michel Kostecki (1995), *The Political Economy of the World Trading System*, Oxford University Press, Oxford and New York.

Horn, Henrik and Petros C. Mavroidis (2008), "The WTO Dispute Settlement System 1995-2006: Some Descriptive Statistics," IFN Working Paper No. 740, Research Institute of Industrial Economics.

Howse, Robert and Robert W. Staiger (2005), "United States–Anti-Dumping Act of 1916 (Original Complaint by the European Communities) – Recourse to arbitration by the United States under 22.6 of the DSU, WT/DS136/ARB, 24 February 2004: A legal and Economic Analysis," *World Trade Review* 4:2, 295-316.

Jackson, John H. (1969), *World Trade and the Law of GATT*. New York: Bobbs-Merrill.

Jackson, John H. (1997), "The WTO Dispute Settlement Understanding–Misunderstandings on the Nature of Legal Obligation," *The American Journal of International Law* 91 (1), January, pp. 60-64.

Kaplow, Louis and Steven Shavell (1996), "Property Rules versus Liability Rules: An Economic Analysis," *Harvard Law Review* 109 (4), pp. 713-790

Lawrence, Robert Z. (2003), *Crimes and Punishments? Retaliation under the WTO* (Washington DC: The Institute for International Economics).

Maggi, Giovanni and Robert W. Staiger (2008), "The Role of Dispute Settlement Procedures in International Trade Agreements," NBER Working Paper No. 14067, June.

Martin, Alberto and Wouter Vergote (2008), "On the Role of Retaliation in Trade Agreements," *Journal of International Economics* 76(1).

Mnookin, Robert H. and Lewis Kornhauser (1979), "Bargaining in the Shadow of the Law: the Case of Divorce," 88 *Yale Law Journal* 950.

Park, Jee-Hyeong (2009), "Enforcing International Trade Agreements with Imperfect Private Monitoring: Private Trigger Strategies and a Possible Role for the WTO," unpublished manuscript, Seoul National University, February.

Pelc, Krzysztof J. (2009), "Seeking Escape: The Use of Escape Clauses in International Trade Agreements," *International Studies Quarterly* 53, pp. 349-368.

Schwartz, Alan (1979), "The Case for Specific Performance," *Yale Law Journal* 89.

Schwartz, Warren F. and Alan O. Sykes (2002), "The Economic Structure of Renegotiation and Dispute Resolution in the World Trade Organization," *Journal of Legal Studies* 31.

Shavell, Steven (2006), "Specific Performance versus Damages for Breach of Contract: An Economic Analysis," *Texas Law Review* 84, March.

Srinivasan, TN (2007), "The Dispute Settlement Mechanism of the WTO: A Brief History and an Evaluation from Economic, Contractarian and Legal Perspectives," *The World Economy*. pp. 1033-1068.

Ulen, Thomas S. (1984), "The Efficiency of Specific Performance: Toward a Unified Theory of Contract Remedy," *Michigan Law Review* 83.

USTR (2001), "U.S. Trade Representative Announces the Lifting of Sanctions on European Products as EU Opens Market to U.S. Banana Distributors," Press Release, July 1.

Figure 1

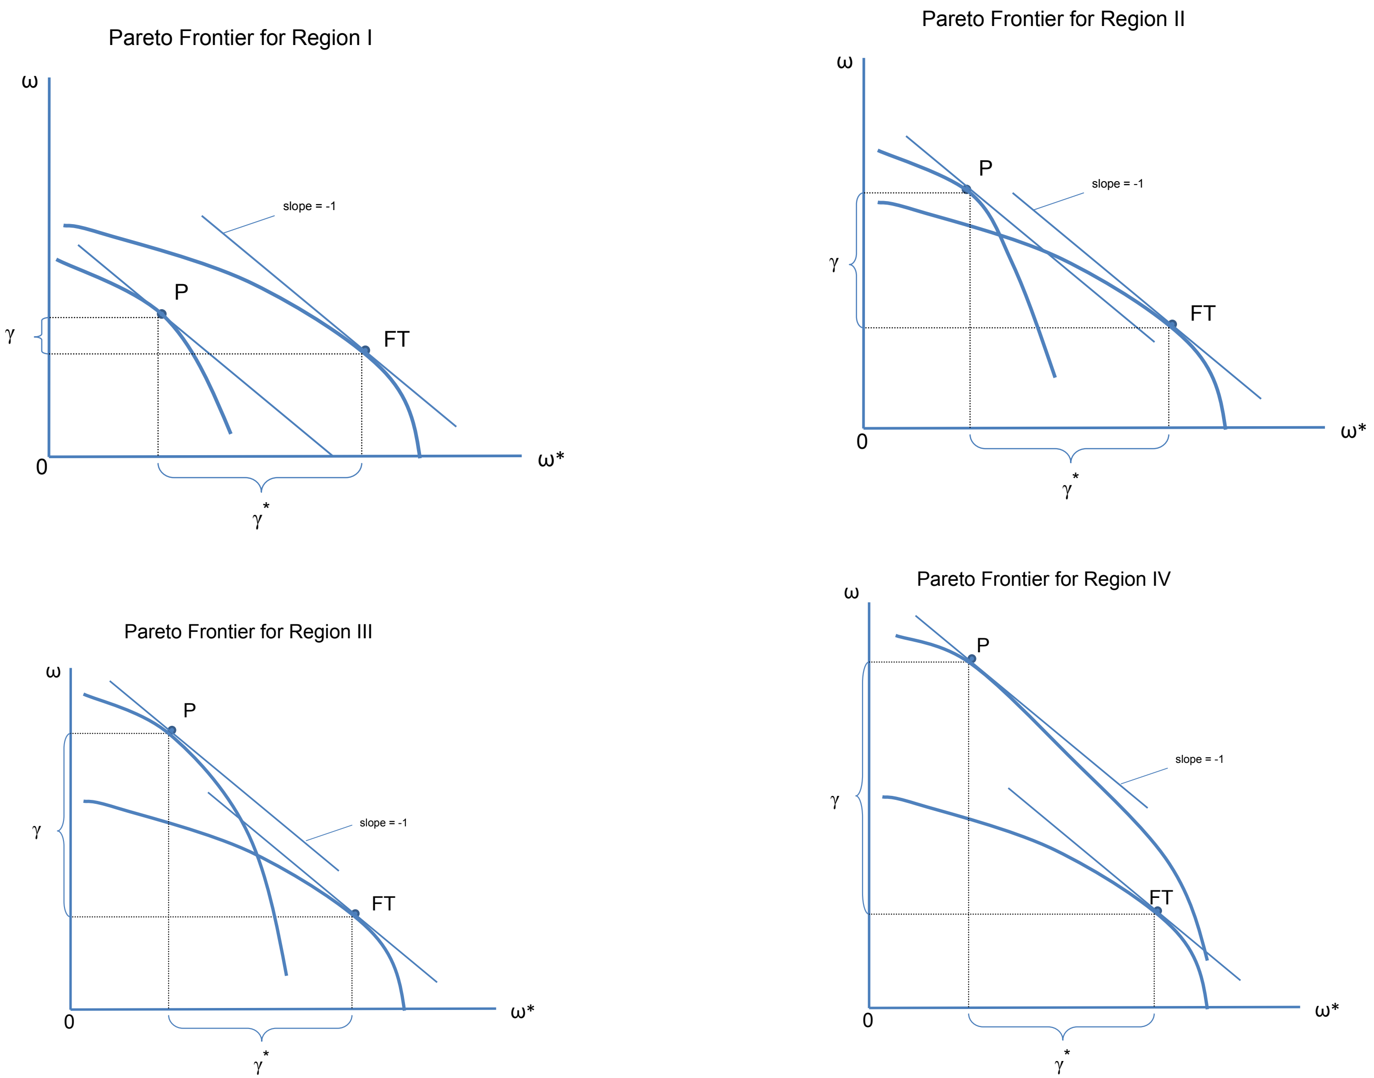

Pareto Frontier for Region I
ω
ω*
0
P
FT
slope = -1
γ
γ*
Pareto Frontier for Region II
ω
ω*
0
P
FT
slope = -1
γ
γ*
Pareto Frontier for Region III
ω
ω*
0
P
FT
slope = -1
γ
γ*
Pareto Frontier for Region IV
ω
ω*
0
P
FT
slope = -1
γ
γ*

Figure 2

Bargaining Outcome in Light of Damages b^D: Region II

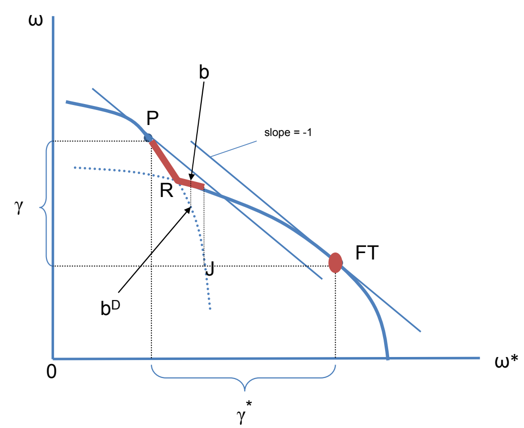

Joint Welfare (Ω) and Damages (b^D): Region II

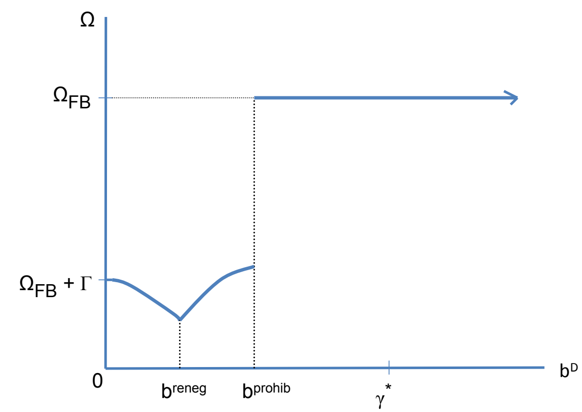

Bargaining Outcome in Light of Damages b^D: Region III

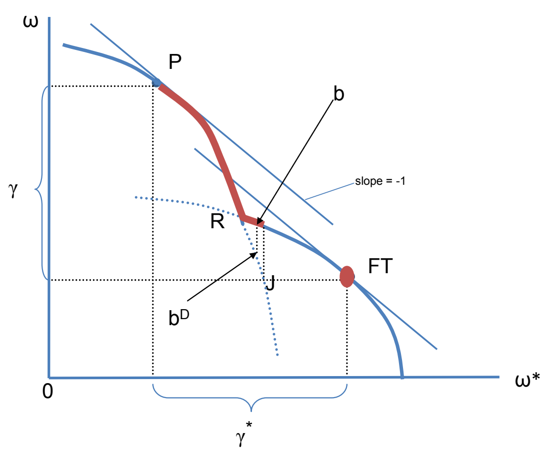

Joint Welfare (Ω) and Damages (b^D): Region III

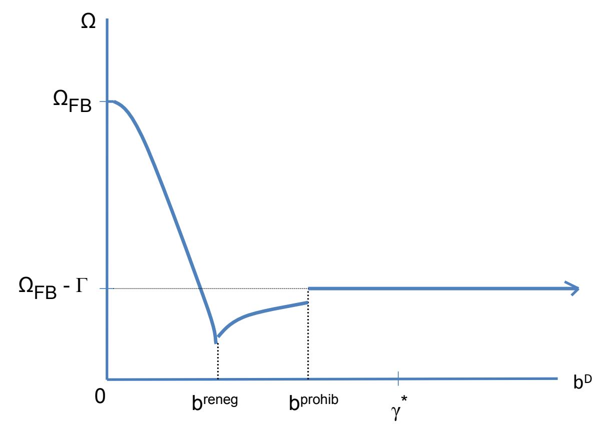

Figure 3

Bargaining Outcome in Light of Damages b^D: Region II/III Border

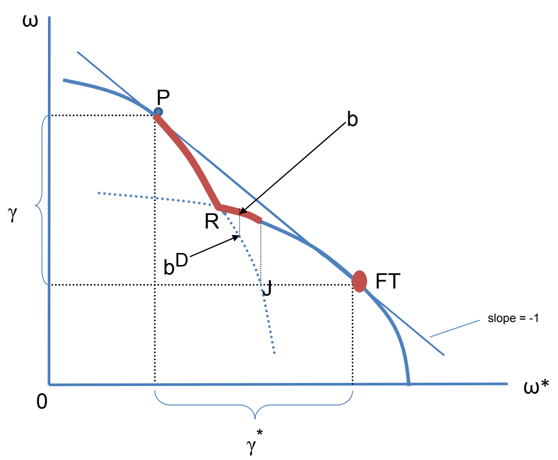

Joint Welfare (Ω) and Damages (b^D): Border of Regions II/III

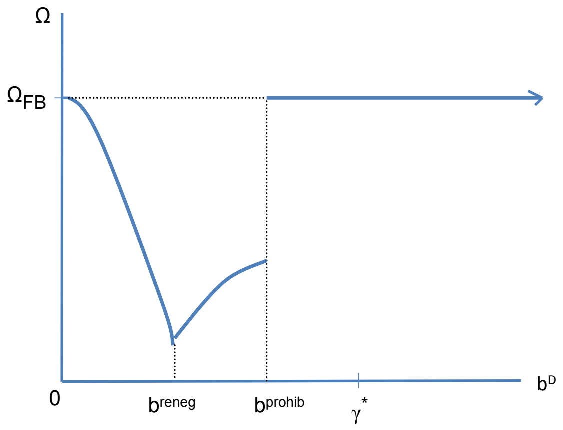

Figure 4

Bargaining Outcome in Light of Damages b^D: Region I

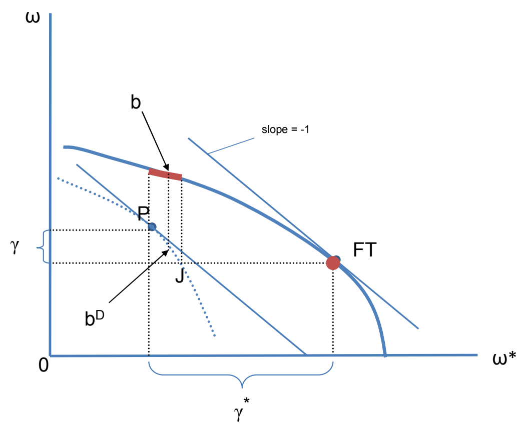

Joint Welfare (Ω) and Damages (b^D): Region I

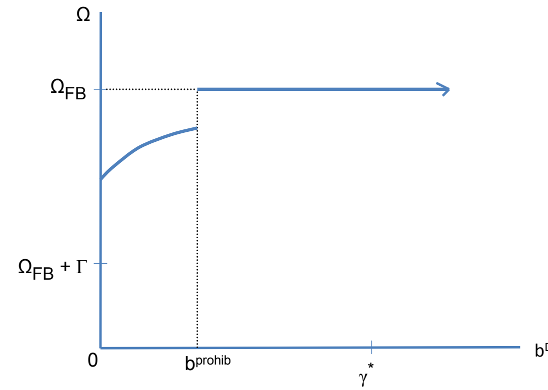

Bargaining Outcome in Light of Damages b^D: Region IV

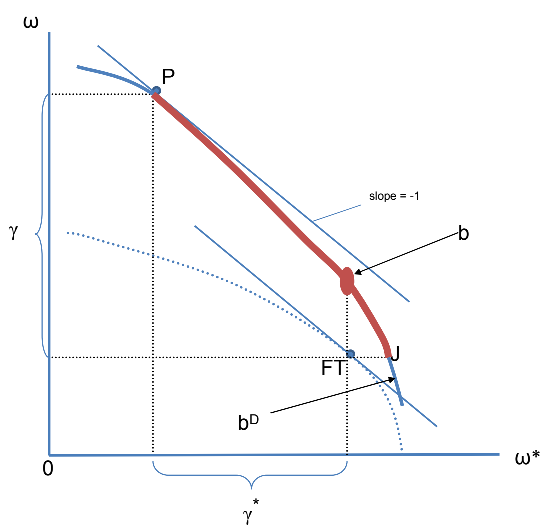

Joint Welfare (Ω) and Damages (b^D): Region IV

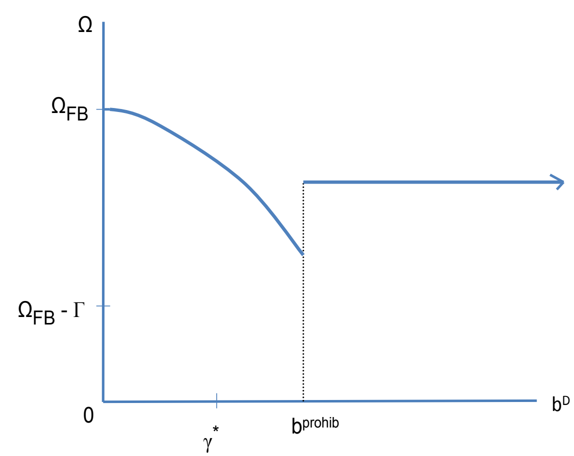

Figure 5

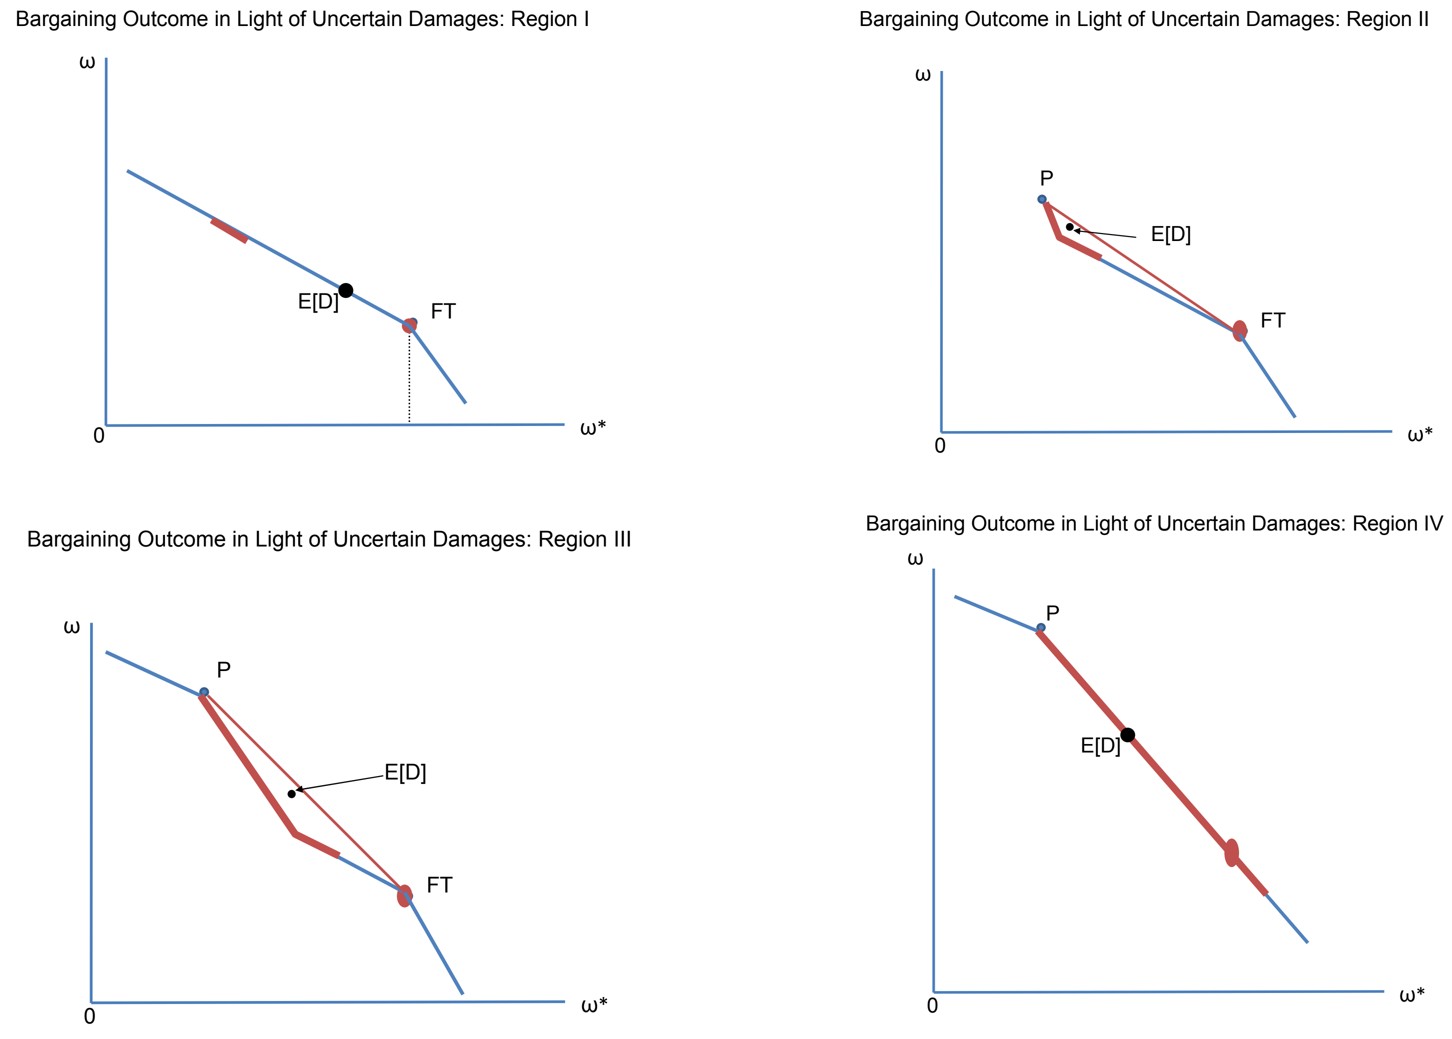

Bargaining Outcome in Light of Uncertain Damages: Region I
ω
E[D]
FT
0
ω*
Bargaining Outcome in Light of Uncertain Damages: Region II
ω
P
E[D]
FT
0
ω*
Bargaining Outcome in Light of Uncertain Damages: Region III
ω
P
E[D]
FT
0
ω*
Bargaining Outcome in Light of Uncertain Damages: Region IV
ω
P
E[D]
0
ω*

Figure 6

Bargaining Outcome in Light of Uncertain Damages: Region II

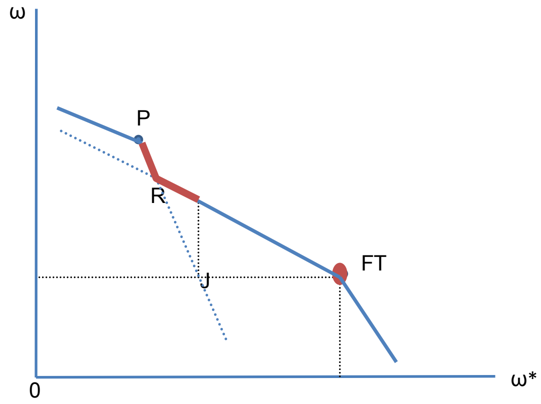

Bargaining Outcome in Light of Uncertain Damages: Region III