2-2015

Optimal Design of Trade Agreements in the Presence of Renegotiation

Giovanni Maggi
Yale University

Robert W. Staiger
Dartmouth College

OPTIMAL DESIGN OF TRADE AGREEMENTS IN THE PRESENCE OF RENEGOTIATION*

Giovanni Maggi
Yale University and NBER

Robert W. Staiger
Stanford University and NBER

September 2011

Abstract

We study the optimal design of trade agreements in a setting where governments can renegotiate the agreement ex-post subject to a key transaction cost, namely that compensation between governments is inefficient. The optimum may be a "property-right" contract or a "liability" contract. A property (liability) rule is optimal if uncertainty over the joint political/economic benefits from free trade is sufficiently small (high). Increasing the bargaining power of the importing country favors property rules over liability rules. And when we introduce a cost of renegotiation, we find that property rules are more likely to be optimal when this cost is higher, a result that reverses a central conclusion of the law-and-economics literature. The model also delivers predictions on the pattern of equilibrium renegotiation. If a property rule is optimal, it is never renegotiated. A liability rule, on the other hand, will be renegotiated for intermediate states of the world. And when renegotiation occurs, it always results in trade liberalization, not protection.

1. Introduction

When governments make international commitments, what is the optimal structure for their contract? This question is at the heart of a growing debate among scholars of international trade agreements. Broadly speaking, these commitments can take one of two possible contractual forms. One type of contract assigns *rights* concerning trade policy to the contract parties. To illustrate, consider a two-country world where the country that imports a good can either practice free trade or engage in protection: this first type of contract either assigns the right to protect to the importing country, or it assigns the right of free trade to the exporting country, and subsequently these rights can be transferred from one government to another only through a voluntary transaction – a renegotiation – between the two governments. In effect, this type of contract assigns *ownership* of rights concerning trade policy, and as a consequence it is commonly referred to in the legal literature as a *property rule.* This contrasts with a second type of contract, where the importer has the option to practice free trade, or to engage in protection and compensate the exporter with a certain amount of *damages.* In the legal literature, this type of contract is referred to as a *liability rule.*[1] As emphasized by many scholars, the choice between these two contract forms is a central issue for the design of a trade agreement (see for instance Jackson, 1997, Schwartz and Sykes, 2002, Lawrence, 2003 and Pauwelyn, 2008).[2]

Looking across the range of real-world trade agreements, there is evidence that both liability rules and property rules are at work. For example, Pauwelyn (2008) argues that property rules provide the "default" approach in both the WTO (applying for example to the WTO prohibitions against quantitative restrictions and export subsidies) and NAFTA, but for certain specific issues a liability-rule approach has instead been taken. Clear examples of liability rules in the GATT/WTO are the provisions for temporary and permanent escapes from negotiated tariff bindings in GATT Articles XIX and XXVIII, respectively, and the rules applying to "actionable" production subsidies in Part III of the WTO Agreement on Subsidies and Countervailing

---

[1]Use of the property- and liability-rule terminology is less common outside of the law-and-economics literature, but the choice between these two types of contract is an important topic also in the economics literature on optimal contract design. In that literature, a liability-rule contract is often referred to as an "option contract," and a property-rule contract is sometimes referred to as a "noncontingent contract," or simply a "property-right" contract (see for example Segal and Whinston, 2002).

[2]More broadly, the choice between liability and property rules relates to the question of the optimal degree of flexibility in a trade agreement. That question has been an important subject of debate in economics, international law and political science, where particular emphasis has been given to the desirability of escape clauses in trade agreements (see for example Bagwell and Staiger, 1990, Sykes, 1991, and Rosendorff and Milner, 2001).

Measures.[3] And in NAFTA (as well as in many bilateral investment treaties), investor protection against expropriation is set up as a liability rule. In addition to this variation across issues, there is evidence that the contract forms featured in a given institution evolve over time. For instance, there is broad agreement that in the early years, GATT operated as a system of liability rules (see Jackson, 1969, p. 147, Schwartz and Sykes, 2002, and Lawrence, 2003, p. 29); however, many legal scholars argue that in more recent times the GATT/WTO has evolved toward a property-rule system (see, for example, Jackson, 1997, Charnovitz, 2003, Pauwelyn, 2008 and Pelc, 2009). A key aim of our paper is to examine the tradeoffs involved in the choice between these two contractual forms.

The intuitive appeal of a liability rule is that it may help mitigate a key limitation of trade agreements, namely, the incompleteness of the contract. If governments could write a complete contingent contract, the efficient outcome could be achieved without any need for a liability rule. But if such a contract is not feasible, for example because some key contingencies (such as the strength of political pressures) are not verifiable, then a liability rule can facilitate the efficient adjustment of trade policy choices without the need to specify those contingencies in the contract. The simple logic is that the damage payment set by a liability rule can help induce the importer to internalize the externalities that it imposes on its trading partner(s) through its trade policy choices. But a liability-rule approach has its limitations. One limitation that stands out in the context of trade agreements is that international lump-sum transfers are generally not available; rather, the payment of damages from one government to another is typically accomplished by "self-help" and takes the form of tariff retaliation, which is inefficient.[4] This transaction cost (inefficient transfers) gives rise to a nontrivial tradeoff between property-rule and liability-rule approaches.[5]

---

[3]Other examples of liability rules in the GATT/WTO can be found in the General Agreement on Trade in Services (Article XXI, which provides for the renegotiation of specific commitments in services trade), and in the Agreement on Trade-Related Aspects of Intellectual Property Rights (Article 31, which sets conditions under which compulsory licenses may be issued). The non-violation nullification-or-impairment clause of the GATT can also be interpreted along the lines of a liability rule, as it permits countries to in effect escape their market access commitments with changes in domestic policies and pay damages to injured parties as a remedy. See Pauwelyn (2008, pp. 134-136) for further discussion.

[4]The empirical magnitude of this inefficiency is arguably of first-order importance, especially if compensation takes the form of tariff retaliation. The inefficiencies associated with the use of tariffs have been quantified by many studies. For example, in one well-cited attempt (Hufbauer and Elliott, 1994), the authors conclude that U.S. consumers pay over six times the average annual compensation of manufacturing workers to preserve each job "saved" by special U.S. import protection.

[5]An additional limitation of a liability-rule approach is that it requires verifiability of the harm inflicted by a country's trade policy on its trading partner(s), which in practice is likely to be very imperfect. This limitation

When evaluating this tradeoff, it is important to allow for the possibility of *renegotiation*, both because of the empirical significance of renegotiation in real-world trade agreements,[6] and because renegotiation changes substantially the nature of the tradeoff between the two contract forms. With renegotiation possible, the ex-ante contract does not determine the policy outcome directly; rather, it defines the disagreement point for the ex-post renegotiation. If there were no transaction costs, the Coase theorem would apply and the ex-ante contract would be irrelevant for the policy outcome. But in the presence of transaction costs, the disagreement point can impact the policy outcome, and the ex-ante contract can then have important efficiency consequences even when, as in the GATT/WTO, there are ample possibilities for renegotiation.

In light of these considerations, it is surprising that virtually all existing models of trade agreements abstract from renegotiation (with the few partial exceptions discussed below). Our paper advances this literature by studying the optimal design of trade agreements in the presence of renegotiation. Our analysis applies to trade agreements generally, but we will pay particular attention to the GATT/WTO, which is a natural institution on which to focus given its prominence in the world trading system.

In recent years there has been considerable research more generally on the optimal design of contracts in the presence of renegotiation, leading examples of which are the papers by Maskin and Moore (1999), Segal and Whinston (2002) and Watson (2007). Our approach broadly follows this literature, by considering an environment with nonverifiable information where the contract is designed ex-ante but can be renegotiated ex post through Nash bargaining. However, we depart from this literature by introducing some new features that are motivated by the international trade context, and we also impose some restrictions to make the model tractable. The main feature we add is that government-to-government transfers involve a deadweight loss, hence utility is nontransferable, whereas the typical models of contracting with renegotiation focus on the case of transferable utility.[7] The main restriction we introduce, on the other hand, is that we focus on a binary policy choice. This buys us tractability, and as we later describe this focus captures many trade-related policies that are discrete in practice.[8] As a consequence

---

is often emphasized in the informal law-and-economics literature, and we discuss it briefly in the Conclusion. However, we abstract from it in our basic model in order to focus more sharply on the role of inefficient transfers.

[6]See for example Hoda (2001) and Busch and Reinhardt (2006) for an account of the many instances in which GATT/WTO members have engaged in renegotiation of their trade policy commitments.

[7]Maskin and Moore (1999) and Segal and Whinston (2002) are able to extend some results to the case of nontransferable utility, but the sharpest findings of this literature are all derived with transferable utility.

[8]A further restriction is that we focus on menu contracts, that is, contracts based on the choice of just one player (in our case the importer). In principle, one could design a more sophisticated mechanism that is based

of the structure we impose, our model delivers sharp results on the pattern of equilibrium renegotiation and on the desirability of property rules versus liability rules.

More specifically, we consider a two-country setting where governments contract over trade policy in the presence of uncertainty about the joint benefits of free trade, which could be positive or negative, due for example to political economy factors. Contracts are perfectly enforceable, but the joint benefits of free trade are not verifiable, so a complete contingent contract cannot be written. The agreement can take the form of a property rule (that either assigns the right of free trade to the exporter or assigns the right to protect to the importer), or of a liability rule (which gives the importer a choice between free trade and protection-cum-compensation). Importantly, we assume that transfers between governments are costly, so that utility is not transferrable; and we allow governments to renegotiate their contract through Nash bargaining once the state of the world is realized.

We start by observing that in our setting it may be optimal to induce renegotiation *in equilibrium*, and indeed the model yields several predictions concerning the pattern and direction of renegotiation. First, if the contract is designed optimally, renegotiation (when it occurs) will result in trade liberalization, not protection. More specifically, according to our model equilibrium renegotiations must take a particular form, in which the exporter agrees to compensate the importer in exchange for trade liberalization, against the importer's (credible) threat to protect and pay damages.[9] Second, while it might be expected that renegotiation would be triggered in extreme states of the world, where the joint benefits of free trade are either very large and positive or very large and negative, we find that in equilibrium renegotiation can only occur for intermediate states of the world. Thus, at a broad level, in our model renegotiation is not an "extraordinary" event, but rather occurs in "ordinary" circumstances.

A third and key prediction concerning the pattern of renegotiation is that, if a property rule is optimal, it is not renegotiated in equilibrium, and hence it entails no equilibrium transfers. We discuss this finding in light of evidence that the use of compensation/tariff-retaliation in

---

on messages sent by both players. But as Segal and Whinston (2002) explain, a (continuous) mechanism that is based on two-sided messages may or may not improve upon a menu contract, depending on the contracting environment. It is therefore an open question whether and to what extent more elaborate mechanisms can improve upon menu contracts in our setting, though as a practical matter it should also be noted that in the context of a policy that is applied on a continuing basis as in our setting, the high frequency with which such message games would have to be played in response to potentially changing states of the world would likely make them exceedingly costly to run. In any case, this is a question that we leave for future research.

[9]While we do not attempt an empirical assessment of our model's predictions, we confirm below that this form of renegotiation does indeed occur in the GATT/WTO.

the GATT/WTO has diminished through time, and we suggest that this diminished role for compensation may be a consequence of the shift from liability to property rules that, in the view of GATT/WTO legal scholars as we have described above, has occurred over time.

We then examine whether the optimal contract takes the form of a property rule or a liability rule. In the standard setting where parties cannot renegotiate the contract, the tradeoff between property rules and liability rules is conceptually simple: property rules involve no ex-post transfers but imply inflexible policy outcomes, while liability rules allow more policy flexibility but entail costly transfers. But the ability to renegotiate the contract ex-post changes the nature of the tradeoff in significant ways: first, property rules no longer necessarily imply inflexible policy outcomes; and second, there is a new consideration that plays an important role, namely how the rule specified in the contract affects the amount of transfers that occur when the contract is renegotiated. Perhaps surprisingly, we show that even though the tradeoff is more subtle, the predictions of the model are sharper when renegotiation is allowed.

We find that a property rule is optimal if uncertainty about the joint benefits of protection is sufficiently low, whereas a liability rule is optimal when this uncertainty is high. Under the interpretation that this uncertainty is caused primarily by political-economy shocks, our finding suggests that the use of liability (property) rules should be more (less) prevalent for issue areas that are more "politicized" and hence prone to political-economy shocks. As we discuss further below, this finding is suggestive of the pattern of liability and property rules observed in the GATT/WTO, and is broadly in line with the emphasis that GATT negotiators placed on uncertainty as they considered the potential benefits of liability rules. Moreover we show that, if a liability rule is optimal, the optimal level of damages falls short of fully compensating the exporter, contrary to the "efficient breach" argument in the law-and-economics literature and in line with features of GATT/WTO remedies based on the principle of reciprocity.

One of the key parameters of the model is the relative bargaining power of the two governments. In our setting, bargaining powers have efficiency consequences, not just distributional consequences, because utility is not transferrable. We find that increasing the bargaining power of the importing country tends to favor property rules over liability rules. This follows from the inefficiency of transfers and the combination of two results we mentioned above: first, if the importer is stronger the equilibrium transfer is larger, because equilibrium renegotiations always entail the exporter compensating the importer; and second, an optimal property rule is never renegotiated in equilibrium and hence entails no equilibrium transfers.

We also consider the impact of renegotiation frictions on the optimal form of contract. We capture renegotiation frictions in a very simple way by introducing a fixed cost of renegotiation. We find that increasing this cost favors property rules over liability rules. As we discuss further below, this result contrasts sharply with a central conclusion of the law-and-economics literature, namely, that bargaining frictions tend to favor liability rules over property rules (Calabresi and Melamed, 1972, and Kaplow and Shavell, 1996).[10] Interestingly, if transfers were costless, the effect of introducing renegotiation frictions in our model would be *reversed*, that is, liability rules would be favored. This explains why our result diverges from the conclusion of the law-and-economics literature: unlike that literature, we focus on a world with costly transfers, which as we argued are a key transaction cost in the international trade context.

These themes are virtually unexplored in the existing economics literature on trade agreements, in part because those models do not accommodate the possibility of renegotiation in a meaningful way.[11] One partial exception is Beshkar (2010b): he considers the optimal design of a trade agreement with privately observed political pressures and costly transfers, but only allows for a limited form of renegotiation and focuses instead on the role of the WTO as a provider of non-binding arbitration for its member governments. Another partial exception is Bagwell and Staiger (1999), who study the properties of a limited form of renegotiation that is restricted to satisfy the GATT/WTO principle of reciprocity, but their focus is also very different from ours. Finally, the possibility of renegotiation is emphasized in the papers by Ludema (2001) and Klimenko et al. (2008), but their focus is on the renegotiation of punishment strategies in repeated-game models of self-enforcing agreements.

By contrast, in the law and economics literature analogous issues have been extensively studied in a domestic context. There are two related literatures. A fundamental question in the literature concerned with domestic contracts (see, for example, Schwartz, 1979, Shavell, 1984 and Ulen, 1984) is when contracting parties would want specific performance as a remedy

---

[10]This is widely seen as a fundamental result in law-and-economics. Wikipedia for example states: "With the opportunity to use either liability or property-based rules to protect entitlements, the academic community soon concluded that the key to figuring out which rule to use turned on the transaction costs. Therefore, if there were low transaction costs, then property rules should be used. If the transaction costs were high, then liability rules should be used." (see the entry "Property Rules, Liability Rules and Inalienability").

[11]For example, a number of papers (such as Bagwell and Staiger, 2005, Martin and Vergote, 2008, Bagwell, 2009, Beshkar, 2010a and Park, forthcoming) consider the optimal design of trade agreements with privately observed political pressures, but none of these papers considers the possibility of renegotiation of the agreement. 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 renegotiation either.

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. But all of this literature maintains the assumption that cash transfers are available (as seems appropriate given the literature's domestic-context focus). By introducing costly transfers, our paper forges a link between the law-and-economics theory of optimal legal rules and the economic theory of trade agreements.

The rest of the paper proceeds as follows. Section 2 lays out the basic model. Section 3 considers a benchmark where no renegotiation is possible. Section 4 characterizes the optimal agreement in the presence of renegotiation. Section 5 considers 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. The Appendix contains proofs not presented in the text.

2. The Model

We focus on a single industry in which the Home country is the importer and the Foreign country is the exporter. We focus on a two-country world because this allows us to make the key points in a more transparent way, but in the concluding section we will briefly discuss the extension to a multi-country setting.

The Home government chooses a binary level of trade policy intervention for the industry, which we denote by $T \in \{FT, P\}$: "Free Trade" or "Protection." The binary policy instrument helps to keep our analysis tractable, and it captures reasonably well a variety of non-tariff policy choices that are discrete in practice, such as trade-related regulatory regimes or product standards, on which many of the trade disputes in the GATT/WTO have focussed. Finally, we assume that the Foreign (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 for providing compensation to trading partners, while indirect (non-cash) transfers, such as tariff adjustments in other sectors or even non-trade policy adjustments, are more easily available.[12] To allow for this possibility in a

---

[12] For example, there are no known cases of cash compensation being provided within the context of the escape clause provisions in GATT Article XIX or the provisions for permanent tariff modifications in Article XXVIII, and the resolution of GATT/WTO disputes has, with two exceptions, never involved cash transfers either (the

tractable way, we let b denote a (positive or negative) transfer from Home to Foreign, with $c(b) \geq 0$ the deadweight loss associated with the transfer level b. The transfer cost $c(b)$ is (weakly) convex and smooth everywhere, with the natural features that $c(0) = 0$ and $c(b) > 0$ for $b \neq 0$. For simplicity, we assume that the Home country bears the deadweight loss $c(b)$, and that the total cost of the transfer inclusive of deadweight loss, $b + c(b)$, is increasing for all b.[13]

The Home government's payoff is given by

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

where $v(T)$ is the Home government's valuation of the domestic surplus associated with policy T in the sector under consideration. 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 the Foreign government is passive in this industry, its payoff is

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

where $v^*(T)$ is the Foreign government's valuation of foreign surplus associated with policy T.

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

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

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

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

---

two exceptions to date are the *US-Copyright* case – see WTO, 2007, pp. 283-286 – and the *Brazil-Cotton* case – see Schnepf, 2010). However, countries do sometimes achieve indirect payment of compensation through various forms of policy adjustments in other sectors (e.g., the GATT/WTO "self-help" method of tariff-retaliation). Our implicit assumption is that these other policies are set efficiently, so that any adjustments to them then entails a deadweight loss.

[13] If the deadweight loss were borne by the Foreign country, none of our qualitative results would change, provided $b - c(b)$ is increasing for all b. Note that both of these assumptions ($b + c(b)$ and $b - c(b)$ increasing for all b) are satisfied if $|c'(b)| < 1$ for all b.

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

Below we will refer to the outcome that maximizes joint surplus as the "first best" outcome. This 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.

We assume that governments are ex-ante uncertain about the joint gains from protection Γ (or equivalently the joint benefits from free trade, $-\Gamma$), but they observe Γ ex post. We also assume that Γ is not verifiable, i.e. not observed ex post by the court/dispute-settlement-body (DSB), so that governments cannot write a complete contingent contract.[14] We consider the simplest environment of this kind that allows us to make the relevant points. We assume that γ^* is known ex-ante, so that all the uncertainty in Γ originates from γ, and that γ is not verifiable.[15] In the Conclusion we briefly discuss the case in which γ^* is also uncertain (and not verifiable). But 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 so-called "efficient breach" argument, according to which efficiency can be induced if Foreign 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). The density $h(\gamma)$ is defined over the positive real line, $\gamma \in [0, \infty)$. We let $\underline{\gamma}$ and $\bar{\gamma}$ denote the bounds of the support of γ, or more formally, $\underline{\gamma} = \inf\{\gamma : h(\gamma) > 0\}$ and $\bar{\gamma} = \sup\{\gamma : h(\gamma) > 0\}$. To make things interesting, we assume that γ^* is strictly positive and 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 transfers were costless (no deadweight loss), then governments could

---

[14]Other papers that also model trade agreements as incomplete contracts include Copeland (1990), Bagwell and Staiger (2001), Horn (2006), Costinot (2008), Horn, Maggi and Staiger (2010) and Maggi and Staiger (2011). We will use the expressions "court" and "DSB" interchangeably. Note that in our model the role of the court/DSB is simply that of an external enforcer of the contract, as in standard models of contracting.

[15]These informational assumptions, namely that uncertainty is one-dimensional and that the uncertain parameter is not verifiable by the court but is observed by both parties, are relatively standard in the literature on mechanism design with renegotiation (see for example Segal and Whinston, 2002). Also, whether the uncertainty over γ reflects underlying uncertainty about $v(FT)$ or $v(P)$ or both is immaterial for our results.

always achieve the first best by engaging in ex-post (i.e., after observing γ) negotiations over policies and (costless) transfers. With costly transfers, 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). We look for the contract that maximizes ex-ante joint surplus.[16]

The contract can be of two different types. The first type of contract is a *property rule*, which either assigns the right of free trade to the exporter (we will sometimes refer to this as a "prohibitive" property rule), or assigns the right to protect to the importer (we refer to this as a "discretionary" property rule). The second type of contract is a *liability rule*, which is a menu contract that gives the Home country a choice between (i) setting FT and (ii) setting P and compensating the Foreign country with a payment b^D. Using slightly different terminology, this type of contract specifies a baseline commitment (FT) but allows Home to escape this commitment by paying a certain amount of *damages*. In section 5 we will consider a more general type of contract that may specify a transfer also for the FT choice; in the basic model we focus on the simpler type of contract because it makes the main insights more transparent.[17]

Note that a prohibitive property rule is outcome-equivalent to an extreme liability rule in which b^D is set prohibitively high (i.e. such that the importer chooses FT in all states of the world), and a discretionary property rule is outcome-equivalent to a liability rule at the other extreme in which $b^D = 0$. Therefore, at a formal level we can focus without loss of generality on the family of liability contracts described above and simply optimize the level of b^D. However, we will call the contracts at the two extremes ($b^D = 0$ and b^D prohibitively high)

---

[16]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, so that the game is overall symmetric; in this case, if the governments focus on the symmetric point of the Pareto frontier, this maximizes the sum of their payoffs.

[17]In our simple model, b^D is noncontingent. But in a richer model where some state-of-the-world variables are verifiable and some are not, b^D could be made contingent on the verifiable variables. This would allow for (partially) contingent liability rules and (partially) contingent property rules. It is useful to keep this in mind because in real-world trade agreements we do observe contingent property and liability rules. Consider, for instance, the WTO. An example of a contingent liability rule is given by the provisions for tariff modifications in GATT Article XXVIII, where the compensation due to the exporting country in case of a tariff increase is contingent on the trade effects of such a tariff change. And an example of a contingent property rule is given by the prohibition against quantitative restrictions, which applies strictly as a default, but does not apply at all under certain "exceptional" contingencies (such as a balance of payments crises, see GATT Article XII).

10

"property rules."[18] We choose to emphasize the property-rule interpretation rather than the extreme-liability-rule interpretation in order to connect with the ongoing debate on the optimal design of trade agreements that we described in the Introduction. It is also important to keep in mind that the difference between a liability rule and a property rule is not just a matter of degree (the level of the damages b^D), but there is an important qualitative difference as well: a liability rule is in essence a "separating" contract, since it induces the importer to choose different policies in different states of the world, while a property rule is in essence a "pooling" contract, since it induces a noncontingent policy choice.[19]

We allow the governments to renegotiate the initial contract after the state of the world γ is realized. More specifically, we assume that the ex-post negotiation is a Nash bargaining game with the initial contract serving as the disagreement point. We let σ (resp. $1 - \sigma$) denote the bargaining power of the Home (resp. Foreign) government. We abstract from underlying issues of enforcement and simply assume that bargaining outcomes between the two governments are enforced (we return to this assumption briefly in the Conclusion). To summarize, the timing of events is as follows: (0) Governments write the contract; (1) γ is realized and observed by the governments; (2) governments can renegotiate the terms of the contract (b and T).

We conclude this section by highlighting an alternative interpretation of the contract-design problem described above. The literal interpretation is that governments write a contract that specifies two options for the importer (choosing FT, or choosing P and compensating the exporter with the payment b^D), and the DSB simply enforces the contract. The alternative interpretation is that governments design an *institution* consisting of two parts: (i) a simple $\{FT\}$ contract with no contractually specified means of escape; and (ii) a mandate for the DSB to implement a certain *remedy for breach* (the payment b^D). Our analysis applies equally well under either of these interpretations (i.e., whether the contract includes an escape provision, or rather a remedy for breach is specified in the DSB mandate), and both of these interpretations are relevant for the GATT/WTO: some WTO clauses take the form of explicit option contracts, for example the escape clause in GATT Article XIX and the provisions for tariff modifications in Article XXVIII; but there are also many contractual commitments for which, when they

---

[18]Kaplow and Shavell (1996) make the same observation and adopt a similar approach.

[19]Arguably, a property rule is also likely to be substantially simpler than a liability rule, because under a liability rule the optimal b^D will inevitably depend on the specifics of the issue area (or the sector) being regulated, so the agreement may have to specify a different b^D for each clause. By contrast, with a property rule, the penalty (zero or prohibitive) need not be issue-specific and could be specified as an agreement-wide default.

apply, there is no escape provision (e.g., the rules governing "actionable" production subsidies, and the ban on export subsidies), and in this case the relevant question is what should be the appropriate remedy applied by the DSB in case of breach. Under either interpretation, 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 fail.[20]

3. The No-Renegotiation Benchmark

Before characterizing the optimal agreement in the presence of renegotiation, it is instructive first to consider the simpler setting where ex-post negotiation is not possible. In this setting governments can be viewed as simply designing a contract ex ante (to maximize ex-ante joint surplus) and then implementing it ex post, and choosing a level of damages b^D amounts to stipulating the actual level of damages that must be paid by Home if it chooses P.

We start by noting that, given b^D, the importer will choose FT if and only if its gain from protection, γ, is below some threshold level $\hat{\gamma}$. The threshold level $\hat{\gamma}$ is the value of γ for which the importer is indifferent between FT and P-plus-damages-b^D, and is given by $\hat{\gamma} = b^D + c(b^D)$. The threshold $\hat{\gamma}$ summarizes the policy "allocation" induced by the contract, and we say that $\hat{\gamma}$ is "implemented" by the level of damages b^D if $\hat{\gamma} = b^D + c(b^D)$.

It is useful to highlight how the notions of property rules and liability rules map into values of $\hat{\gamma}$. For this purpose, we define the "prohibitive" level of damages b^{prohib} as the minimum value of b^D such that the importer chooses FT for all γ in the support $(\underline{\gamma}, \bar{\gamma})$, which is defined implicitly by $b^D + c(b^D) = \bar{\gamma}$. Clearly, then, setting a discretionary property rule ($b^D = 0$) corresponds to setting $\hat{\gamma} = 0$; a prohibitive property rule ($b^D \geq b^{prohib}$) corresponds to $\hat{\gamma} \geq \bar{\gamma}$; and a liability rule ($b^D \in (0, b^{prohib})$) corresponds to a value $\hat{\gamma}$ that is strictly between 0 and $\bar{\gamma}$.

It is helpful to write the optimal contracting problem as choosing the transfer b^D and the policy allocation $\hat{\gamma}$ to maximize the ex-ante joint surplus, subject to the "implementation constraint" $b^D + c(b^D) = \hat{\gamma}$. Letting $E\Omega_\varnothing(b^D, \hat{\gamma})$ denote the ex-ante joint surplus given b^D and $\hat{\gamma}$ when no ex-post negotiation is possible, we can state the optimal contracting problem as

---

[20]There is a semantic distinction in the law-and-economics literature according to which, if a trade agreement is viewed as a piece of international law then the property-rule/liability-rule terminology would be used, whereas if a trade agreement is viewed as a contract then the analogous distinction is between a contract that requires "specific performance" and one that specifies "damages for breach." We use these terms interchangeably.

$$\max_{b^D,\hat{\gamma}} E\Omega_{\varnothing}(b^D, \hat{\gamma}) \;=\; V(FT) + \int_{\hat{\gamma}}^{\infty}(\gamma - \gamma^*)dH(\gamma) - c(b^D)[1 - H(\hat{\gamma})] \qquad (3.1)$$

$$\text{s.t. } b^D + c(b^D) \;=\; \hat{\gamma},$$

where $V(FT) \equiv v(FT) + v^*(FT)$. The expression for $E\Omega_{\varnothing}(b^D, \hat{\gamma})$ in (3.1) is the sum of three terms. The first is the joint surplus under a rigid FT policy and no transfers; the second captures the gains in joint surplus associated with allowing the policy P for $\gamma > \hat{\gamma}$; and the third reflects the deadweight loss associated with the transfer b^D and policy allocation $\hat{\gamma}$.

Note that, if there is no cost of transfers ($c(\cdot) \equiv 0$), the objective in (3.1) is clearly maximized by $\hat{\gamma} = \gamma^*$, the first best allocation; but if $c(\cdot) > 0$, it may be optimal to deviate from this allocation. Next note that, if $c(\cdot) > 0$, implementing the first best allocation implies a deadweight loss, which is given by $c(b^D)[1 - H(\gamma^*)]$; this can be interpreted as the "sorting cost." This cost is incurred for all states higher than γ^*, which explains why it is weighted by $[1 - H(\gamma^*)]$. Thus, in the absence of renegotiation, the tradeoff in choosing $\hat{\gamma}$, and hence the optimal level of damages, can be understood in very simple terms: the choice of $\hat{\gamma}$ hinges on the comparison between the efficiency cost of deviating from the first best allocation and the savings in sorting costs that can be achieved by doing so.

To maximize the objective in (3.1), we can use the implementation constraint to solve for the value of b^D that implements $\hat{\gamma}$, plug this into the objective function and optimize $\hat{\gamma}$. We let $b^D_{\varnothing}(\hat{\gamma})$ denote the value of b^D that implements $\hat{\gamma}$. Differentiating $E\Omega_{\varnothing}$ with respect to $\hat{\gamma}$, and noting that $\frac{dc(b^D_{\varnothing}(\hat{\gamma}))}{d\hat{\gamma}} = \frac{c'(\cdot)}{1+c'(\cdot)}$, we obtain

$$\frac{dE\Omega_{\varnothing}(b^D_{\varnothing}(\hat{\gamma}), \hat{\gamma})}{d\hat{\gamma}} = (\gamma^* - \hat{\gamma}) \cdot h(\hat{\gamma}) + c(\cdot) \cdot h(\hat{\gamma}) - \frac{c'(\cdot)}{1 + c'(\cdot)} \cdot [1 - H(\hat{\gamma})]. \qquad (3.2)$$

The first term of $\frac{dE\Omega_{\varnothing}}{d\hat{\gamma}}$ captures the marginal efficiency gain of increasing $\hat{\gamma}$: this is positive if $\hat{\gamma} < \gamma^*$ and negative otherwise. The second term and third term together capture the marginal savings in sorting costs (positive or negative) from increasing $\hat{\gamma}$: the second term is positive because increasing $\hat{\gamma}$ reduces the range of states for which the importer government chooses to pay the transfer, while the third term is negative because increasing $\hat{\gamma}$ requires an increase in the transfer, which will be paid for all states higher than $\hat{\gamma}$.

At this point one might proceed with a "local" approach, and ask how the objective can be improved starting from the first best allocation $\hat{\gamma} = \gamma^*$: Does improvement require increasing

13

or decreasing $\hat{\gamma}$? Or formally, what is the sign of $\frac{dE\Omega_\varnothing}{d\hat{\gamma}}$ at $\hat{\gamma} = \gamma^*$? Clearly, the sign is positive if and only if sorting costs are saved by increasing $\hat{\gamma}$ slightly from γ^*, but this is ambiguous because, as explained above, the marginal savings in sorting costs are composed of two effects that go in opposite directions. More specifically, it is direct to verify that the sign of $\frac{dE\Omega_\varnothing}{d\hat{\gamma}}$ at $\hat{\gamma} = \gamma^*$ is positive if and only if $\frac{d\ln c}{d\hat{\gamma}} > \frac{h(\hat{\gamma})}{1-H(\hat{\gamma})}$: thus the answer hinges on a comparison between the proportional change in the deadweight loss and the (inverse of) the hazard rate at $\hat{\gamma}$, and a local approach cannot therefore take us very far.

Partly because of the feature we just highlighted, the predictions about the nature of the optimal rules in the no-renegotiation case are somewhat ambiguous. The only sharp prediction obtains under the scenario in which uncertainty about γ is small, in the sense that the support of γ around γ^* is small: in this case, a property rule must be optimal. To see why, consider a liability rule, that is a value of $\hat{\gamma}$ within $(0, \bar{\gamma})$. It is easy to see that such a value of $\hat{\gamma}$ is dominated by $\hat{\gamma} = 0$: the key is to notice from (3.2) that $\frac{dE\Omega_\varnothing}{d\hat{\gamma}} = -\frac{c'(\cdot)}{1+c'(\cdot)} < 0$ for all $\hat{\gamma}$ between 0 and $\underline{\gamma}$ (because $h = 0$ for all these values).[21] This implies that $\hat{\gamma} = 0$ dominates all values of $\hat{\gamma}$ between 0 and $\underline{\gamma}$, and moreover $\hat{\gamma} = 0$ dominates $\hat{\gamma} = \underline{\gamma}$ by a discrete margin, and by continuity will dominate any $\hat{\gamma}$ within the support $(\underline{\gamma}, \bar{\gamma})$ provided that this support is sufficiently small.[22] Intuitively, a liability rule can achieve a contingent, and hence more efficient, policy allocation, but the associated gain is small when the support of γ is small, and it is overwhelmed by the deadweight loss from the transfer.

Let us focus now on the opposite case, in which uncertainty is large. We find that, if the support of γ is sufficiently large, a prohibitive property rule $\hat{\gamma} \geq \bar{\gamma}$ is necessarily suboptimal, but the discretionary property rule cannot be ruled out. To understand the first of these two claims, start by noting that the first two terms in (3.2) collapse to $(\gamma^* - b_\varnothing^D(\hat{\gamma})) \cdot h(\hat{\gamma})$ (using the definition of $b_\varnothing^D(\hat{\gamma})$). Next note that $b_\varnothing^D(\hat{\gamma}) > \gamma^*$ for $\hat{\gamma}$ large enough, and hence $\frac{dE\Omega_\varnothing}{d\hat{\gamma}} < 0$ for $\hat{\gamma}$ in a left neighborhood of $\bar{\gamma}$. This implies that $\hat{\gamma} \geq \bar{\gamma}$ is dominated by setting $\hat{\gamma}$ slightly lower than $\bar{\gamma}$. On the other hand, the discretionary property rule $\hat{\gamma} = 0$ *can* under some conditions be a maximum. To see why, notice that, since $b_\varnothing^D(0) = 0$ and $c'(0) = c(0) = 0$, we have $\frac{dE\Omega_\varnothing}{d\hat{\gamma}}|_{\hat{\gamma}=0} = \gamma^* \cdot h(0)$. If $h(0) = 0$ then $\hat{\gamma} = 0$ is a stationary point, and in this case one can show

---

[21]Recall that $\gamma^* > 0$, and so when the support of γ around γ^* is sufficiently small we must have that $\underline{\gamma} > 0$.

[22]It can be shown that a liability rule is dominated also by the prohibitive property rule $\hat{\gamma} \geq \bar{\gamma}$. This cannot be seen easily from (3.2), but one way to see it intuitively is to focus on the symmetric case where $E\gamma = \gamma^*$: in this case the two property rules yield the same expected joint surplus, and hence if a liability rule is dominated by one property rule it is also dominated by the other.

that if $h'(0)$ is sufficiently small then $\hat{\gamma} = 0$ is a local maximum, and it is straightforward to construct examples in which $\hat{\gamma} = 0$ is a global maximum. Intuitively, when uncertainty is large and renegotiation is not possible, it may be best to allow full discretion (P always), because inducing FT even for just the lowest levels of γ – i.e. those states where FT is most desirable for joint surplus – requires a transfer that will occur in equilibrium for all higher levels of γ – i.e. those states in which the importer will choose P – and the resulting transfer costs may be too high to be worthwhile.

To summarize our analysis thus far, in the no-renegotiation scenario the key tradeoff is conceptually simple. A liability rule can insure against extreme realizations of γ but it entails transfer costs, while a property rule avoids transfer costs but it entails downside risk associated with extreme realizations of γ. However, the predictions about the optimal rules are somewhat ambiguous, at least in the case of large uncertainty. As we will show in the next section, introducing renegotiation complicates the trade-offs, but perhaps surprisingly, leads to sharper predictions about the optimal rules.

4. The Optimal Agreement in the Presence of Renegotiation

We now turn to the central task of characterizing the optimal agreement in the presence of renegotiation. In part we can build on the analysis of the no-renegotiation case in the previous section, because for any level of damages b^D the contract characterized there provides the *disagreement (threat) point* for any renegotiation in the present setting. But to characterize where governments will actually end up for any realization of γ given a level of b^D, we must consider the incentives to renegotiate the initial contract, and this necessitates some new notation and a somewhat more involved analysis.

The first question we need to address is the following: Given a contract that specifies damages b^D, when does renegotiation occur (i.e., for what realizations of γ), and in what direction does it occur (i.e., from P to FT or from FT to P)? As we observed just above, the threat point in the renegotiation is given by the initial contract, which gives the importer the option to choose between $(T = FT, b = 0)$ and $(T = P, b = b^D)$. Clearly, the importer is indifferent between these two options if $\gamma = b^D + c(b^D) \equiv S(b^D)$. In words, $S(b^D)$ is the total cost of the transfer b^D inclusive of deadweight loss, and it is the level of γ at which the threat point "switches": under disagreement, for $\gamma < S(b^D)$ the importer chooses $(T = FT, b = 0)$, while for $\gamma > S(b^D)$ it chooses $(T = P, b = b^D)$. We depict the curve $S(b^D)$ in Figure 1. As

Figure 1 reflects, $S(b^D)$ is increasing and convex and goes through the origin, and the threat point is $(T = FT, b = 0)$ below $S(b^D)$ and $(T = P, b = b^D)$ above it.

Having characterized how the threat point varies with γ for a given b^D, we can now proceed to identify the realizations of γ for which the initial contract will be renegotiated, and determine as well the direction of the renegotiation. Note that the analysis of renegotiation is not straightforward because utility is not transferrable, due to the cost of transfers, and for this reason we cannot simply focus on the governments' ex-post joint surplus to determine whether the contract will be renegotiated.

Let us focus first on the case $\gamma < S(b^D)$, where the threat point is $(T = FT, b = 0)$. Clearly, the governments will renegotiate to the policy P if and only if there exists a transfer b^e such that both governments gain by switching from $(T = FT, b = 0)$ to $(T = P, b = b^e)$, which requires $\gamma > S(b^e)$ (for the importer) and $b^e > \gamma^*$ (for the exporter). The equilibrium b^e will then fall somewhere in the interval $[\gamma^*, S^{-1}(\gamma)]$, depending on bargaining powers.[23] Furthermore, it is clear that there is a Pareto improvement over the threat point if and only if the interval $[\gamma^*, S^{-1}(\gamma)]$ is nonempty, or $\gamma > S(\gamma^*)$. Thus, we can conclude that the contract is renegotiated toward policy P for values of γ such that $S(\gamma^*) < \gamma < S(b^D)$. This condition identifies a region in (γ, b^D) space, which is highlighted in Figure 1 by the vertical shading (and labeled P_R). Notice that $b^e < b^D$ in this region, given that $S(b^e) < \gamma < S(b^D)$ with $S(\cdot)$ increasing.

It is useful at this juncture to observe that it can never be strictly optimal to set $b^D > \gamma^*$. To see why, notice from Figure 1 that setting $b^D > \gamma^*$ induces the same policy allocation as setting $b^D = \gamma^*$ (namely FT for $\gamma < S(\gamma^*)$ and P for $\gamma > S(\gamma^*)$). Therefore, any $b^D > \gamma^*$ is weakly dominated by $b^D = \gamma^*$ because the latter implies a weakly lower expected transfer.[24] A consequence of this observation is that it will never be the case that *in equilibrium* the contract is renegotiated towards P: this follows because renegotiation from FT to P is only possible in the case where $\gamma < S(b^D)$ and the threat point is FT, and as Figure 1 depicts when $\gamma < S(b^D)$ the contract is renegotiated toward P only when $b^D > \gamma^*$, which we have just observed can

---

[23]Note that the function $S(\cdot)$ is invertible, because we assumed that $b + c(b)$ is increasing everywhere.

[24]To see this, fix a level of b^D above γ^*, say $b^D = \tilde{b}^D > \gamma^*$, and replace it with $b^D = \gamma^*$. This decreases the expected equilibrium transfer (weakly) for two reasons: (1) if $\gamma > S(\tilde{b}^D)$, so that the importer chooses $(T = P, b = b^D)$ without renegotiating, the transfer obviously decreases, and (2) if $\gamma \in (\gamma^* + c(\gamma^*), S(\tilde{b}^D))$, so that the contract is renegotiated, the equilibrium transfer b^e is higher than γ^*, as we showed in the text. We also note the reason for the qualifier that $b^D > \gamma^*$ is "weakly" dominated by $b^D = \gamma^*$: if the support of γ around γ^* is small, the expected equilibrium transfer is the same in the two cases, because all states $\gamma > \gamma^* + c(\gamma^*)$ have zero density. But note that, if this is the case, even if $b^D > \gamma^*$ renegotiation from FT to P cannot occur with positive probability. This explains the sentence that follows in the text.

never be strictly optimal. Instead, as we will confirm and discuss further below, the only kind of renegotiation that can occur in equilibrium is from P to FT.

Let us now focus on the case $\gamma > S(b^D)$, where the threat point is $(T = P, b = b^D)$. In this case, the governments will renegotiate toward the policy FT if and only if there exists a (negative) transfer b^e such that both governments gain by switching from $(T = P, b = b^D)$ to $(T = FT, b = b^e)$, which requires $S(b^D) - S(b^e) > \gamma$ (for the importer) and $\gamma^* > b^D - b^e$ (for the exporter). Again using the definition of $S(\cdot)$, the equilibrium b^e will then fall somewhere in the interval $[b^D - \gamma^*, S^{-1}(S(b^D) - \gamma)]$, depending on bargaining powers. Clearly, there exists a Pareto improvement over the threat point if and only if this interval is nonempty, or $\gamma < S(b^D) - S(b^D - \gamma^*) \equiv R(b^D)$, hence we can conclude that the contract is renegotiated toward policy FT when $S(b^D) < \gamma < R(b^D)$. This condition identifies a region in (γ, b^D) space that is highlighted in Figure 1 by the horizontal shading (and labeled FT_R).[25]

Finally we note that the two renegotiation regions highlighted in Figure 1 are themselves independent of the bargaining-power parameter σ. Bargaining powers only affect the exact amount of transfer b^e that will be exchanged inside these regions.

Figure 1 summarizes the analysis thus far, and is of central importance for understanding the results that follow. The two key curves are $S(b^D)$ and $R(b^D)$. The curve $\gamma = S(b^D)$ is the locus of points where the importer's threat point switches between P and FT. If $b^D < \gamma^*$ (which as we argued is the relevant case), this curve also marks the lower boundary of the region where governments renegotiate the contract (in the direction of free trade), while the curve $\gamma = R(b^D)$ marks the upper boundary of the renegotiation region.

Our findings on the pattern of renegotiation are recorded in the following:

Proposition 1. *(i) If $b^D < \gamma^*$, the contract is renegotiated for $\gamma \in (S(b^D), R(b^D))$, in which case the governments agree on FT and the exporter compensates the importer. (ii) If $b^D > \gamma^*$, the contract is renegotiated for $\gamma \in (S(\gamma^*), S(b^D))$, in which case the governments agree on P and the importer compensates the exporter; however, setting $b^D > \gamma^*$ is weakly dominated, and this kind of renegotiation does not happen in equilibrium.*

Proposition 1 implies two interesting predictions regarding the pattern and direction of equilibrium renegotiation. The first prediction is that, as long as damages are set optimally,

---

[25]Note that $R(0) = -S(-\gamma^*) > 0$, so this region is guaranteed to be nonempty, and note also that $R''(b^D) < S''(b^D)$ for all b^D, which ensures that the point of intersection between the R curve and the S curve is unique.

any observed ex-post renegotiation of the ex-ante contract must result in liberalization (from P to FT), not protection (from FT to P). That is, according to Proposition 1, equilibrium renegotiations all take a particular form in which the importer (respondent) agrees to liberalize and the exporter (claimant) agrees to pay something for this. What should not occur in equilibrium according to Proposition 1 is a renegotiation wherein the importer's threat point is FT but the governments agree to a policy of P and a level of damages to the exporter that is less than the contractually-specified level ($b^e < b^D$). Intuitively, if renegotiation took this latter form, it would imply that the contractually-specified damages b^D are suboptimally high, because for the exporter government to agree to such a renegotiation would require $\gamma^* < b^e$, and hence $\gamma^* < b^D$; but this cannot be optimal, as we have explained previously.[26]

The second prediction is that, if $b^D > 0$, renegotiation can occur in equilibrium only for intermediate values of γ. Broadly speaking, then, our model predicts that renegotiation is not an "extraordinary" event that occurs only in extreme states of the world, but instead can occur only for "ordinary" states of the world.[27] This is perhaps surprising, since intuition might suggest that renegotiation should occur in exceptional circumstances where the ex-ante contract turns out to be highly inefficient ex-post. The reason this intuition is not correct in our model is that, in extreme states of the world, the initial contract performs well, in the sense that it induces the importer to make the correct policy choice. Instead, in our model governments may have incentive to renegotiate (for the relevant range of b^D) only if the importer government is relatively close to indifferent between the options placed before it by the initial contract, and this indifference occurs for an intermediate state of the world ($\gamma = S(b^D)$): if the importer

---

[26]Examples of renegotiation that conform to the model's equilibrium predictions, in which the respondent agrees to liberalize and the complainant agrees to pay something for this, are not hard to find. For instance, in an early GATT dispute between India and Pakistan regarding export fees levied by Pakistan on jute sold to India that India claimed violated the MFN obligation, India withdrew its complaint under a settlement (see GATT, 1953) in which Pakistan agreed to eliminate the discriminatory features of its export taxes in exchange for an agreement by India to reduce its (non-discriminatory) export tax on coal. A more recent example is provided by the 2001 compliance settlement for the U.S.-EU "Banana" dispute in the WTO (see USTR, 2001). In this settlement, the EU (respondent) agreed to come into compliance with the DSB ruling, but not fully until 2006. Hence, the U.S. (a claimant), by accepting the EU's non-to-partial compliance over the 2001-2006 period, allowed the EU to take some compensation (by being able to deviate from its WTO commitment over this period) in exchange for the promise of eventual full compliance. Finally, whether or not renegotiations of the second type described in the text – which should not occur in equilibrium according to the model's predictions – are observed in practice is an interesting empirical question, but we are unaware of any evidence from existing empirical studies that is directly relevant in this regard.

[27]Note that this result makes no reference to the probabilities of different realizations of γ. Thus, in principle, intermediate values of γ could be less likely than extreme values of γ, and in this sense the former could be more "exceptional" than the latter. But this can happen only for very non-standard distributions, so we feel justified in stating broadly that in our model renegotiation can occur only for "ordinary" states of the world.

prefers P and is far from this indifference point, the exporter will have to pay a large transfer to convince the importer to switch its policy choice to FT, and this will entail a large dead weight loss. In this sense, the reason the contract may be renegotiated in equilibrium in our model reflects the inefficiencies of the transfers more than the inefficiencies of the policy itself.

Having characterized the pattern and direction of renegotiation, we can now turn to the analysis of the optimal contract. An important step towards characterizing the optimal b^D is to ask the following question: What allocations $\hat{\gamma}$ can be implemented in the presence of renegotiation, and what is the level of damages b^D that implements a given $\hat{\gamma}$?

This question is immediately answered by looking at Figure 1. First note that there exists *no* level of damages b^D that can implement values of $\hat{\gamma}$ outside the interval $[R(0), S(\gamma^*)]$: regardless of b^D, the policy outcome for $\gamma > S(\gamma^*)$ is always P, and the policy outcome for $\gamma < R(0)$ is always FT. This is an important difference relative to the case of no renegotiation: in the absence of renegotiation, any allocation $\hat{\gamma}$ can be implemented by an appropriate choice of b^D. But when renegotiation is feasible, it is impossible to induce FT for values of γ above $S(\gamma^*)$, or P for values of γ below $R(0)$.

Notice also that the range of implementable values of $\hat{\gamma}$ is smaller when the cost of transfer is lower (it can be verified that decreasing $c(\cdot)$ leads to an increase in $R(0)$ and a decrease in $S(\gamma^*)$). This feature is very intuitive in the limiting case where transfers are costless: then the parties will always renegotiate to the efficient outcome regardless of b^D, and hence only the allocation $\hat{\gamma} = \gamma^*$ is implementable. This is a manifestation of the Coase theorem: in the absence of transaction costs the ex-ante contract is irrelevant and the efficient outcome always obtains ex-post. When transfers are quite costly, on the other hand, the renegotiation outcome is very sensitive to the level of damages b^D specified in the initial contract, and hence the range of implementable allocations is wider. Thus *renegotiation limits the scope of implementation, and the more so the lower the transfer cost.* We let $IM_{\hat{\gamma}} \equiv [R(0), S(\gamma^*)]$ denote the set of implementable values of $\hat{\gamma}$.

In spite of the fact that renegotiation imposes bounds on implementation, *renegotiation is beneficial* for ex-ante joint surplus. To see this notice that, for each given b^D and γ, governments renegotiate only if this leads to an ex-post Pareto improvement. Since renegotiation leads to a weak ex-post Pareto improvement for all (b^D, γ), it follows that the ex-ante joint surplus must also be weakly higher. The following lemma summarizes:

Lemma 1. *Renegotiation limits the range of allocations $\hat{\gamma}$ that can be implemented. The*

implementable range of $\hat{\gamma}$ is given by $IM_{\hat{\gamma}} = [R(0), S(\gamma^*)]$. However, renegotiation is (weakly) beneficial for the ex-ante joint surplus.

Lemma 1 indicates that, despite placing limits on what can be implemented, the ability to renegotiate (weakly) enhances the ex-ante joint surplus of the governments. We note that this feature contrasts with other mechanism design settings (e.g., the hold-up and risk-sharing environments studied in Segal and Whinston, 2002) where the ability to renegotiate ex-post can be harmful to ex-ante surplus, for the simple reason that our model abstracts from the sorts of ex-ante investment/risk-aversion issues that are the focus of those settings.[28]

We next ask, what level of b^D is required to implement a given $\hat{\gamma}$? From Figure 1 it is clear that implementing a given $\hat{\gamma}$ in $IM_{\hat{\gamma}}$ requires $b^D(\hat{\gamma}) = R^{-1}(\hat{\gamma})$. Note that $b^D(\hat{\gamma})$ is increasing in the relevant range; and, recalling that the definition of $b_\varnothing^D(\hat{\gamma})$ implies $b_\varnothing^D(\hat{\gamma}) = S^{-1}(\hat{\gamma})$, we note as well that $b^D(\hat{\gamma}) \leq b_\varnothing^D(\hat{\gamma})$ for all $\hat{\gamma} \in IM_{\hat{\gamma}}$: in spite of the fact that renegotiation limits the scope of implementation, it takes a lower level of contractually-specified damages to implement a given $\hat{\gamma}$ than in the absence of renegotiation (for $\hat{\gamma}$ in the implementable set).

Finally, it is important to recall that the level of damages $b^D(\hat{\gamma})$ specified in the contract is not necessarily the transfer that occurs in equilibrium, since the contract may be renegotiated, so Lemma 1 does not tell us the *cost* of implementing $\hat{\gamma}$. For $\hat{\gamma} \in IM_{\hat{\gamma}}$, this cost includes two components: (1) the cost of the transfer $b^e(\cdot)$ made when the contract is renegotiated, which is the case for $\gamma \in (S(b^D(\hat{\gamma})), \hat{\gamma})$, and (2) the cost of the transfer b^D made when the contract is not renegotiated and the importer chooses $(T = P, b = b^D)$, which is the case for $\gamma > \hat{\gamma}$.

Armed with the observations above, we can now write down the optimization problem in the presence of renegotiation. Recalling that we can focus on $b^D \leq \gamma^*$ and that for this range of b^D we have $\hat{\gamma} = R(b^D)$, we can write the problem as follows:

$$\max_{b^D, \hat{\gamma}} E\Omega(b^D, \hat{\gamma}) = V(FT) + \int_{\hat{\gamma}}^{\infty} (\gamma - \gamma^*)dH(\gamma) - c(b^D)[1 - H(\hat{\gamma})] - \int_{S(b^D)}^{\hat{\gamma}} c(b^e(b^D; \gamma))dH(\gamma)$$

$$s.t. \ \hat{\gamma} = R(b^D), \ b^D \leq \gamma^*$$

There are two main differences between this optimization problem and the optimization problem in (3.1) that applies when renegotiation is not possible: first, the expected cost of

---

[28] It is also relevant to observe that efforts to renegotiate ex-post are actively encouraged (and even mandated) in the context of GATT/WTO disputes. Our finding that renegotiation is beneficial is consistent with this stance.

transfers now includes not only the cost of the transfer b^D for states in which the contract is not renegotiated and the importer chooses $(T = P, b = b^D)$, but also the cost of the transfer b^e that is paid by the exporter when renegotiation occurs; and second, the level of b^D required to implement a given $\hat{\gamma}$ is lower than in the case of no renegotiation, as we highlighted above.

A final ingredient for finding the optimal level of contractually-stipulated damages b^D is understanding how the level of b^D affects the transfer b^e paid by the exporter when renegotiation occurs. Intuitively, increasing b^D strengthens the bargaining position of the exporter and hence decreases b^e in absolute size. To see this more formally, note that in the FT_R region b^e solves the generalized Nash bargaining problem:

$$\max_{b} \quad NB(b; b^D) \equiv \left(\omega(FT, b) - \omega(P, b^D)\right)^\sigma \left(\omega^*(FT, b) - \omega^*(P, b^D)\right)^{1-\sigma}.$$

By standard monotone comparative-statics results, $\frac{\partial b^e}{\partial b^D}$ has the same sign as $\frac{\partial^2 NB(b, b^D)}{\partial b \partial b^D}|_{b=b^e}$ which, using the explicit expression $NB(b; b^D) = \left(S(b^D) - S(b) - \gamma\right)^\sigma \left(\gamma^* + b - b^D\right)^{1-\sigma}$, is positive. As $b^e < 0$ in the FT_R region, it follows that $\frac{\partial |b^e|}{\partial b^D} < 0$. We record this finding in

Lemma 2. *For (b^D, γ) in the FT_R region, where governments renegotiate to FT and $b^e < 0$, an increase in b^D leads to a decrease in (the absolute size of) the equilibrium transfer:* $\frac{\partial |b^e|}{\partial b^D} < 0$.

We are now ready to study the optimal level of b^D, and in particular compare property rules with liability rules in the presence of renegotiation. Recall that the discretionary property rule is defined as $b^D = 0$; the prohibitive property rule as $b^D \geq \bar{b}^{prohib}$ (where $\bar{b}^{prohib}$ is determined by $S(\bar{b}^{prohib}) = \bar{\gamma}$); and a liability rule as $b^D \in (0, \bar{b}^{prohib})$.

Before proceeding, however, it is important to emphasize how the introduction of renegotiation changes the tradeoff involved in the choice between liability rules and property rules. Recall that, in the absence of renegotiation, the tradeoff is fairly simple: property rules avoid the cost of transfers but imply rigid policy outcomes, whereas liability rules can introduce policy flexibility but imply some waste associated with the use of transfers. If governments are able to renegotiate the contract, on the other hand, this tradeoff is complicated by the fact that the policy outcome is no longer necessarily rigid under property rules; by the fact that renegotiation imposes a limit on the policy allocations that can be implemented; and perhaps most importantly, by the fact that the level of b^D affects the equilibrium payments that are made when governments renegotiate. But as we now show, in spite of this more complicated tradeoff, the introduction of renegotiation actually sharpens the results of the model.

21

We focus first on the case of small uncertainty. In this case, a property rule must be optimal, and the logic is similar to the case of no renegotiation. Figure 2 depicts the relevant features of the small-uncertainty case. First note from Figure 2 that, if the support of γ around γ^* is sufficiently small, a property rule ($b^D = 0$ or $b^D \geq \bar{b}^{prohib}$) is not renegotiated for any γ, and hence it induces zero transfers in equilibrium. A liability rule may achieve a more efficient policy allocation than a property rule, since the policy can be made contingent on γ, but the associated benefit is small because the support of γ around γ^* is small. On the other hand, the cost of achieving this state-contingency is not small, because implementing a threshold $\hat{\gamma}$ close to γ^* requires a level of damages b^D that is close to $R^{-1}(\gamma^*)$ and hence does not become negligible as the support shrinks.[29] In this case, it is straightforward to establish that the optimum is $b^D = 0$ if $E\gamma > \gamma^*$ and $b^D \geq \bar{b}^{prohib}$ if $E\gamma < \gamma^*$.

Let us focus next on the case where uncertainty is sufficiently large. It is helpful to refer back to Figure 1 for this case. Suppose that $\underline{\gamma} < R(0)$ and $\bar{\gamma} > S(\gamma^*)$. Recalling that the implementable range of $\hat{\gamma}$ is $IM_{\hat{\gamma}} = [R(0), S(\gamma^*)]$, this is the case in which the support includes high-γ states in which the policy outcome is P regardless of the initial contract, and it includes low-γ states in which the policy outcome is FT regardless of the initial contract. In this case, a liability rule must be optimal. To see why, first note that in this case $\bar{b}^{prohib} > \gamma^*$, and recall from Proposition 1 that $b^D > \gamma^*$ can never be optimal, so a prohibitive property rule cannot be optimal. Next consider the discretionary property rule $b^D = 0$. Note that, given $b^D = 0$, for all $\gamma > R(0)$ the contract is not renegotiated and the outcome is $(P, b = 0)$; for these states, increasing b^D slightly from zero entails only a second-order loss, since the marginal cost of the transfer is zero at $b = 0$. But for all $\gamma < R(0)$, given $b^D = 0$ the contract is renegotiated and the exporter pays a sizable transfer b^e, and recall from Lemma 2 that increasing b^D reduces the size of b^e: this is a first-order benefit, and hence increasing b^D slightly from zero improves the objective. We can conclude that if the support of γ is sufficiently large, both property rules are dominated by a liability rule.

The following proposition summarizes:

Proposition 2. *(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*

---

[29] To be more precise, if $\hat{\gamma}$ is close to γ^* then for states above $\hat{\gamma}$ the transfer b^D will be close to $R^{-1}(\gamma^*)$, which is non-negligible; for states below $\hat{\gamma}$ the contract will be renegotiated, and the equilibrium transfer may be lower, but this renegotiated transfer is unrelated to the size of the support of γ and hence does not become small as the support shrinks.

is sufficiently large (on both sides of γ^*), the optimum is a liability rule, and in particular the optimal b^D satisfies $0 < b^D < \gamma^* < \bar{b}^{prohib}$.

As Proposition 2 reflects, the introduction of renegotiation leads to sharper predictions about the optimal rules, despite the fact that the trade-offs involved become more subtle. In particular, when renegotiation is possible, with sufficiently large uncertainty a liability rule dominates both the prohibitive property rule and the discretionary property rule, whereas in the absence of renegotiation we have shown that the discretionary property rule can be optimal even when uncertainty is large. And as we have explained, the reason a liability rule dominates the discretionary property rule in the presence of renegotiation is surprising: introducing a small $b^D > 0$ in the contract leads to a *saving* in transfer costs by strengthening the bargaining position of the exporter, while the *policy* allocation remains unaffected.

We have used the support of γ as a measure of ex-ante uncertainty. If uncertainty about γ is small in the sense that the density of γ is very concentrated around γ^* but the support is large, 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. We note as well that the support of γ need only shrink relative to γ^*, not in absolute size.

Proposition 2 states that a liability rule is optimal if uncertainty about γ is sufficiently large, but in this case the optimal level of b^D is lower than the level that makes the exporter "whole," i.e. γ^*.[30] This result qualifies the argument often made in the law-and-economics literature that the efficient level of 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 verifiable. The source of this qualification comes from our assumption of costly transfers, and so it applies with particular force to international trade agreements. Specifically, in the context of the WTO compensation often takes the form of tariff-retaliation by the injured party, hence it entails inefficiencies, and therefore from an ex-ante perspective it should not be utilized to an extent that fully compensates the injured party for its loss. This qualification gains special relevance in light of the emphasis placed on reciprocity in the GATT/WTO system of remedies: it is sometimes suggested that reciprocity falls short as a mechanism for inducing efficient outcomes because it does *not* make the injured party whole (see, for example, Charnovitz, 2002, Lawrence, 2003 and Pauwelyn, 2008), but Proposition 2 suggests that this may in fact be a desirable feature of reciprocity.

---

[30]A similar result has been shown by Beshkar (2010b).

Moreover, as Proposition 2 indicates, if uncertainty about γ is sufficiently small, *any* liability rule is suboptimal (let alone the specific liability rule with $b^D = \gamma^*$), and instead the optimum is a property rule. Under the interpretation that uncertainty in γ reflects primarily political-economy shocks, Proposition 2 therefore suggests an interesting empirical prediction: we should tend to observe more liability rules for issue areas where political-economy shocks are more intense; and conversely, the use of property rules should be more frequent for issue areas where political-economy shocks are less important. Whether or not this empirical prediction is borne out in observed trade agreements is an open question, but it is tempting to link this prediction to what we observe in the GATT/WTO: as we mentioned in the Introduction, export subsidies and quantitative restrictions (QRs) are prohibited by property rules, while tariffs and production subsidies are regulated through a liability-rule approach. One can argue that tariffs and production subsidies are subject to considerable political-economy shocks. On the other hand, export sectors are typically less subject to political pressures than import-competing sectors;[31] hence, export subsidies are arguably less sensitive to political pressures than tariffs and production subsidies. And while QRs can in principle be the target of strong lobbying pressures, the GATT/WTO has successfully "tariffied" most QRs, thereby channeling political pressures to a large extent away from QRs and into tariffs (and perhaps production subsidies), so QRs are arguably less subject to political economy shocks as well.[32]

We next highlight two interesting predictions of our model that derive from the underlying pattern of equilibrium renegotiation. One relates to a key feature of optimal property rules, while the other relates to the role of bargaining powers for the optimal choice of contract.

First, when a property rule is optimal, it is never renegotiated, and it therefore entails no equilibrium transfers. This can be seen as follows. Consider a prohibitive property rule ($b^D \geq \bar{b}^{prohib}$). By definition, this entails a b^D high enough that for all γ in the support the

---

[31] This claim has a long history, and it reflects the twin observations that it is declining industries rather than expanding industries that typically receive trade policy support and that it is import-competing industries rather than exporting industries that are typically in decline. Empirically, these observations have been extensively documented in the literature. As for why the asymmetric political influence of import-competing relative to export sectors arises, there have been a number of proposed answers, including the free-riding and rent-erosion problems faced by lobbyists in expanding sectors as a result of new entrants (Grossman and Helpman, 1996, and Baldwin and Robert-Nicoud, 2007) and the "identity bias" reflecting asymmetries in the political market (Krueger, 1990, and Fernandez and Rodrik, 1991).

[32] We also note that the degree of uncertainty was a key consideration in the discussions of GATT trade negotiators on the potential benefits of liability rules. As Pauwelyn (2008, p. 137) writes: "...trade negotiators cannot foresee all possible situations, nor can they predict future economic and political developments, both at home and internationally. As a result of this uncertainty, they wanted the flexibility of a liability rule." This kind of consideration seems broadly in line with the message of our Proposition 2.

importer's threat point is FT (or $S(b^D) > \bar{\gamma}$); but we know from Proposition 1 that under an optimal contract there can never be renegotiation from FT to P, hence when a prohibitive property rule is optimal it is never renegotiated. Now consider a discretionary property rule ($b^D = 0$). We have established previously that a necessary condition for this to be optimal is that when $b^D = 0$ there is no renegotiation for any γ (or $R(0) < \underline{\gamma}$). Our claim then follows immediately. We record this in

Proposition 3. *When a property rule is optimal, it is never renegotiated, and therefore entails no equilibrium transfers.*

Intuition might have suggested that the possibility of renegotiation should enhance the performance of property rules, because it insures the contracting parties against the intrinsic rigidity of such rules. Proposition 3 however tells us that this intuition is not really correct. The possibility of renegotiation can enhance the performance of a property rule only if such a rule is suboptimal. If a property rule is optimal, the possibility of renegotiation is immaterial.

In light of Proposition 3, it is relevant to observe that the frequency of renegotiation and compensation in the GATT/WTO has diminished through time.[33] And as we mentioned in the Introduction, in the view of most legal scholars the GATT/WTO began as a liability-rule system but has developed over time into a system of property rules.[34] Proposition 3 links these two observations, and suggests that the observed drop in the use of compensation might be a consequence of this shift in the GATT/WTO from liability to property rules.[35] Finally, we have emphasized the implications of Proposition 3 for changes through time, but we note

---

[33] This feature has been noted by, among others, Hoda (2001), Goldstein and Martin (2002), Pauwelyn (2008), and Pelc (2009). For example, Hoda (2001) notes that during the period 1995-1999 only eight tariff renegotiations took place, as opposed to fifty-six in the period 1980-89. Pelc (2009) focuses on the decline in compensation over the history of the GATT/WTO, and documents this decline in the context of GATT/WTO escape clauses.

[34] Representing this majority view, Jackson (1997, pp. 62-63), argues that the GATT/WTO has evolved from what was in effect a system of liability rules in the early GATT years to a system of property rules under the reforms introduced with the creation of the WTO and embodied in the DSB. On the dissenting view, see Hippler Bello (1996) and Schwartz and Sykes (2002), who view the changes in the DSB that were introduced with the creation of the WTO as serving instead to return the system to one based squarely on liability rules.

[35] It is not obvious what may have caused this shift from liability to property rules. According to our model a reduction in uncertainty could have this effect, but it is not clear that uncertainty has diminished over time. For this reason, here we emphasize only the prediction of the model concerning the co-variation between contract form and frequency of renegotiation, which seems consistent with observations. In the conclusion we suggest an intriguing possibility that arises in a multi-country extension: if renegotiation frictions increase with membership, then an expansion of membership over time can contribute to an explanation of the shift from liability to property rules.

that Proposition 3 suggests an analogous cross-sectional prediction: there should be more renegotiation and compensation in issue areas regulated by liability rules.

We turn now to the role of bargaining powers for the optimal choice of contract. We find that an increase in the importing government's bargaining power σ unambiguously favors property rules over liability rules. The reason is simple. As σ rises, the expected transfer as a result of renegotiations goes up, because by Proposition 1 equilibrium renegotiations always entail a transfer from the exporter to the importer. Next recall that σ has no impact on the equilibrium policy. These observations, combined with the result of Proposition 3 that there are no equilibrium renegotiations when property rules are optimal, lead to the following implication:

Proposition 4. *As σ rises, the optimum can switch from a liability rule to a property rule, but not vice-versa.*

It is interesting to observe that, according to our model, bargaining powers are irrelevant under an optimal property rule (because by Proposition 3 an optimal property rule is never renegotiated), while an importing government with low bargaining power (small σ) will receive a relatively low payoff under a liability rule. Hence our model indicates that, where significant power imbalances exist between countries, moving between a liability rule system and a property rule system will not be distributionally neutral, suggesting in turn that developed and developing countries might naturally have differing preferences with regard to reforms that would move an institution in one direction or the other.[36]

4.1. Costly renegotiation

Thus far we have considered two contracting scenarios: one where renegotiation is frictionless (earlier in this section), and one where renegotiation is not feasible at all (in section 3). We

---

[36]For instance, consider a 2-mirror-image-sectors version of the model along the lines described in note 16, in which the model is overall symmetric across countries except for a country-specific bargaining power parameter (σ for Home and $(1 - \sigma)$ for Foreign with $\sigma < 1/2$), and suppose model parameters are such that property and liability rules are equally efficient. Then under the property rule, Home and Foreign payoffs are equal, but under the liability rule, Home does worse than Foreign. And since the two rules are equally efficient by assumption, a move from the property rule to the liability rule would then require that Foreign compensate Home. In the case of the GATT/WTO, one might expect based on our results that developing countries would take a skeptical view of reforms that had the impact of moving the system away from property rules and toward liability rules. This prediction is broadly in line with bargaining positions taken by member governments over proposed institutional reforms that would have elevated the role of compensation as a "buy out" of one's obligations in the GATT/WTO (see, for example, the discussion in Pelc, 2009, and the references cited therein, as well as the complementary discussion in Pauwelyn, 2008, pp. 90-93).

now examine a range of intermediate scenarios where renegotiation is feasible but is costly. Our main objective is to understand how renegotiation frictions impact the choice between property rules and liability rules. This question is interesting in its own right, but a further motivation for considering this extension comes from the law-and-economics literature. As mentioned in the Introduction, a central result in this literature is that an increase in renegotiation frictions favors liability rules over property rules (Calabresi and Melamed, 1972, and Kaplow and Shavell, 1996), and indeed this result is very intuitive: property rules are rigid in nature, so they work well only if renegotiation is easy.[37] In our setting, as we show next, this result is reversed, and the reason lies in the fact that utility is nontransferable in our model, whereas the above-mentioned law-and-economics models assume transferable utility.

We consider a simple extension of our basic model, wherein governments must pay a fixed (deadweight) cost K if they want to renegotiate the contract. Let α (resp. $(1-\alpha)$) be the share of this cost borne by Home (resp. Foreign). This is a cost that is incurred only if the renegotiation is successful, so it reduces the available surplus but does not affect the disagreement utilities.[38]

To take a first intuitive pass at this question, let us compare the two extreme scenarios: the case in which renegotiation is costless ($K = 0$) and the case in which renegotiation is not feasible ($K = \infty$). The finding in Proposition 3 implies that removing the possibility of renegotiation favors property rules over liability rules. This is because an optimal property rule is never renegotiated, while as we have established above, when renegotiation occurs it increases joint surplus; hence, removing the possibility of renegotiation can only increase the attractiveness of a property rule relative to a liability rule.

The above reasoning suggests that increasing K should favor property rules. However this reasoning is incomplete, because a change in K affects b^e when renegotiation occurs, and this indirect effect can in principle offset the direct intuitive effect. But as we now show, even if this indirect effect works in the "wrong" direction, it can never outweigh the direct effect.

We start by deriving the regions in (b^D, γ) space where the contract is renegotiated, for a given renegotiation cost K. Focus first on the region where $\gamma < S(b^D)$, so that the importer's threat point is FT. For the contract to be renegotiated, there must exist a b^e such that (i) the

---

[37]In fact, it is this result, applied to the GATT/WTO context, that leads Pauwelyn (2008, p. 66) to the following statement: "On balance, one can therefore expect that Calabresi and Melamed's third reason for liability rules (high transaction costs) can find particular application in the international context."

[38]If the cost of renegotiation were a sunk cost, rather than a fixed cost, it could potentially introduce a hold-up issue in the process of renegotiation, since incurring this cost would be akin to an ex-ante investment from the point of view of an individual government. We abstract from this aspect here.

importer is better off, which requires $\gamma > S(b^e) + \alpha K$, and (ii) the exporter is better off, which requires $b^e > \gamma^* + (1-\alpha)K$. Clearly, this is the case if and only if $\gamma > S(\gamma^* + (1-\alpha)K) + \alpha K > S(\gamma^*)$. It follows that the P_R region is as depicted in Figure 3. Intuitively, as K increases, the horizontal line $\gamma = S(\gamma^* + (1-\alpha)K) + \alpha K$ shifts up, thus the P_R region shrinks.

Next focus on the region where $\gamma > S(b^D)$, so that the importer's threat point is P. For the contract to be renegotiated, there must exist a b^e such that (i) the importer is better off, which requires $\gamma < S(b^D) - S(b^e) - \alpha K$, and (ii) the exporter is better off, which requires $b^e > b^D - \gamma^* - (1-\alpha)K$. This is the case if and only if $\gamma < S(b^D) - S(b^D - \gamma^* + (1-\alpha)K) - \alpha K \equiv R(b^D; K)$. The FT_R region is also depicted in Figure 3. Note that as K increases the curve $R(b^D; K)$ shifts down, thus the FT_R region shrinks.

Having derived the renegotiation regions in the presence of the renegotiation cost, the next step is to assess how an increase in K affects the relative performance of property vs. liability rules. The key steps of the argument, spelled out in the Appendix, are two. First we show that, as in the case of costless renegotiation, an optimal property rule is never renegotiated. And second, we then argue that a small increase in K has two first-order effects on the performance of a liability rule: a direct (weakly) negative effect, since it reduces the renegotiation surplus; and an indirect effect through b^e, which can be positive or negative, but even if positive, can never outweigh the negative direct effect.

The following proposition (proved in the Appendix) states the result:

Proposition 5. *As K increases, the optimum may switch from a liability rule to a property rule, but not vice-versa.*

Proposition 5 highlights the impact of renegotiation costs, which is a distinct type of transaction cost from the one that we have focused on more directly – the cost of transfers. It is interesting to observe that, if utility were transferable, renegotiation costs would have the *opposite* effect, that is, they would favor liability rules.[39] This suggests that these different forms of transaction costs *interact* in nontrivial ways, and it points to the importance of taking transfer costs into account when evaluating the effects of bargaining frictions. Moreover, this observation explains why our result is at odds with the conclusion of the law-and-economics

---

[39]To see this, suppose that transfers are costless (i.e., $c(b) \equiv 0$). In this case, with frictionless renegotiation, liability rules are equivalent to property rules, because both achieve the first best; while if renegotiation is costly, the unique optimum is a liability rule with $b^D = \gamma^*$ (i.e. the exporter must be made "whole"). Hence, if transfers are costless, renegotiation costs favor liability rules, in contrast with the case of costly transfers.

literature mentioned above, that renegotiation frictions tend to favor liability rules over property rules. Relative to that literature, our novel finding arises because of our focus on a world with costly transfers, which as we have indicated are an important feature of the international government-to-government contracting environment.[40]

Finally, notice that Propositions 3 and 5 taken together suggest a kind of complementarity between liability rules and renegotiation when transfers are costly. In this environment, lowering the cost of renegotiation makes liability rules more attractive, and the adoption of liability rules makes renegotiation more likely in equilibrium.

5. A More General Class of Contracts

Thus far we have considered a family of menu contracts that admits a natural partition into liability rules and property rules. While a property rule amounts to granting the right of P to the importer (when damages are set to zero) or the right of FT to the exporter (when damages are set at a prohibitive level), a liability rule can be thought of as introducing a "stick" b^D attached to the choice of P (when b^D is strictly positive but set at a non-prohibitive level).

We now ask whether it is desirable in the presence of a liability rule to introduce as well a "carrot" b^{FT} attached to the choice of FT. We accomplish this by considering a more general class of menu contracts of the type $\{(P, b^D), (FT, b^{FT})\}$, where in principle b^D and b^{FT} could be positive or negative. Notice that there is no role for a carrot in a property rule, because by definition a property rule simply assigns property rights without specifying any ex-post transfers.[41] But the introduction of a carrot might be optimal as a *complement* to a stick/liability rule. And if it is optimal, the question then arises whether a property rule can still be optimal when a carrot is available for use with a liability rule.

We find that the performance of a liability rule can always be enhanced by the use of an appropriately chosen carrot $b^{FT} < 0$. The reason for this can be understood as follows. Under a liability rule, there will be some states of the world (some γ's) where the importer

---

[40] It is true that the type of renegotiation friction that we consider here differs from that typically considered by the law-and-economics literature (namely, the presence of private information). Nevertheless, as we have observed above and shown in note 39, the reason for the reversal of the results has to do with non-transferable utility, not the exact nature of the bargaining friction.

[41] For the more restricted class of menu contracts that we considered in earlier sections, we observed that a discretionary (prohibitive) property rule is outcome equivalent to a liability rule with b^D set at zero (b^D set at a prohibitively high level). This outcome equivalence extends naturally to the more general class of menu contracts that we consider here when $b^{FT} = 0$.

would be willing to select FT for a very small transfer from the exporter, but the importer can exploit its bargaining power (i.e., $\sigma > 0$) to extract a sizable transfer b^e from the exporter in order to be induced away from the threat point (P, b^D) to a policy of FT. Introducing a small carrot $b^{FT} < 0$ for FT flips the importer's threat point for these γ's from (P, b^D) to (FT, b^{FT}) and thereby undercuts the ability of the importer to hold out for a bigger transfer from the exporter, ensuring that for these γ's the importer will then select FT and be paid the contractually specified b^{FT}. And from the point of view of ex-ante efficiency, the elimination of large transfers (b^e) that would have to be paid in a few states of the world is worth the addition of a very small transfer (b^{FT}) to be paid in (possibly) many states of the world. The upshot, then, is that introducing a small carrot for "good behavior" can be helpful in the presence of a liability rule, even when that carrot is a costly/inefficient means of transferring surplus between governments; and paradoxically, offering a reward to the importer for FT hurts the importer, by taking away from him the credible threat of choosing P.

As noted above, this raises a second question: Can a property rule still be optimal when a carrot is available for use with a liability rule? Here we find that the answer is "Yes," and in particular that the availability of a carrot for use with a liability rule does not alter our previous finding that a property rule (liability rule) is optimal if uncertainty is sufficiently small (large).

The following proposition (proved in the Appendix) summarizes the discussion thus far:

Proposition 6. *Consider menu contracts of the type $\{(P, b^D), (FT, b^{FT})\}$: (i) If the support of γ is sufficiently small, a property rule is optimal; (ii) If the support of γ is sufficiently large (on both sides of γ^*), it is optimal to use a carrot ($b^{FT} < 0$) together with a liability rule/stick ($b^D > 0$), and in particular the optimal b^D satisfies $0 < b^D < \gamma^* < \bar{b}^{prohib}$.*

Notice that if a liability rule is not optimal in the restricted class of menu contracts considered in previous sections, it might become optimal in the more general class considered here, because the use of a carrot can enhance the value of a liability rule as we have observed. What Proposition 6 shows is that, while the quantitative thresholds may change, it is nevertheless true that the qualitative findings of Proposition 2 extend to this more general class of contracts.

We have focused above on establishing that Proposition 2 extends to the more general class of contracts considered here, but it is straightforward to check that each of our other Propositions extend to this setting as well.[42] In particular, Proposition 1 extends with modification only in

---

[42]Lemma 2 also extends to this setting. In fact, the only result from earlier sections that does not extend to

the critical levels of b^D and γ, while Propositions 3-5 extend without modification.

It is interesting to consider whether the kind of carrot mechanism described in Proposition 6 is observed in actual trade agreements such as the WTO. On the one hand, when a government agrees to reduce its tariffs as 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 6 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. When put this way, it is less clear that the "carrot" device represented in Proposition 6 can be found in existing trade agreements.

Finally, we close our discussion of this more general class of contracts by emphasizing that, while it is optimal under some circumstances to include a carrot b^{FT} in the contract, its inclusion does not alter in any substantive way our results from the previous section, and indeed as we have observed, all of our Propositions extend to this setting as well.

6. Conclusion

We have characterized the optimal design of trade agreements in a model where governments can renegotiate the agreement ex-post subject to a key transaction cost, namely that compensation between governments is inefficient. We have argued that these two features, renegotiation and inefficient government-to-government transfers, figure prominently in the GATT/WTO and other trade agreements. Our model delivers predictions concerning the optimal form of the agreement, the conditions under which the agreement will be renegotiated in equilibrium, and the form that such renegotiation will take. A key question on which we have focused is whether the agreement should be structured as a system of "property rules" or "liability rules." In this respect our paper forges a link between the theory of trade agreements and the law-and-economics theory of optimal legal rules. We have shown that answers to this question diverge from those provided by the law and economics literature once the inefficiency of transfers is

the more general class of contracts considered here is Lemma 1: when a carrot is available, renegotiation does not limit the range of allocations that can be implemented.

introduced, indicating that the inefficiency of government-to-government transfers distinguishes the trade-agreement setting from domestic legal environments in important ways.

To preserve tractability and focus on the main points, we have made a number of strong assumptions. For example, as we mentioned earlier, one important limitation of a liability-rule approach from which our model abstracts and which is often emphasized in the informal law-and-economics literature is that a liability rule requires verifiability of the harm inflicted by a country's trade policy on its trading partner(s), and in practice this can often be difficult to come by. In ongoing work, we allow that the level of harm is not perfectly verifiable, and the DSB (if invoked) can observe a noisy signal of this harm. We find that, in itself, imperfect verifiability of the level of harm makes liability rules relatively less attractive, which is not surprising. But an advantage of studying this noisy-verification setting is that it can generate positive predictions regarding the propensity of governments to settle early versus "going to court" (i.e., invoking the DSB to generate a noisy signal and issue a ruling), and more generally regarding the outcome of trade disputes.

We have also assumed that neither government possesses private information, an assumption that has helped to bring the distinctive features of our analysis into sharp relief. Allowing for private information would introduce an additional transaction cost in the form of a renegotiation friction. Unlike the transfer costs that we have emphasized, private information in bargaining is not specific to the international setting which is our focus (though surely important in real-world trade agreements), so we feel justified in abstracting from this type of friction in our basic model. In our extended model of section 4.1 we have introduced a cost of renegotiation, which captures a form of renegotiation frictions in a simple and tractable way. Still, extending our results to settings with private information would likely yield additional insights.

In our model, contracts between governments can be perfectly enforced. This is a strong assumption, since in reality trade agreements must be self-enforcing, but issues of self-enforcement are logically distinct from the choice between property and liability rules (on this logical distinction see Jackson, 1997, p. 63, and see also Pauwelyn, 2008, pp. 148-197, for an especially detailed discussion of this point), so focusing on a model with enforceable contracts seems to us like a reasonable first step.[43] But saying that these issues are logically distinct does not mean that there is no interaction between them: indeed, our paper leaves open some subtle

---

[43] For papers that model the self-enforcing nature of trade agreements, see for example Bagwell and Staiger (1990), Maggi (1999) and Ederington (2001) in addition to those mentioned in the Introduction.

questions about the interplay between the nature of enforcement mechanisms and the nature of legal rules. We see this interplay as an important topic for future research.

Finally, we have focused on a two-country model. We believe that this is a natural first step to understand the implications of renegotiation for trade agreements, and that such a setting captures some of the fundamental forces governing the tradeoff between property rules and liability rules. Nevertheless, it is important to assess the robustness of our results to a multi-country setting. To make some progress in this direction, we have analyzed a three-country version of our model, where country H can import a good from two exporting countries, say F1 and F2. The presence of competing exporters can in principle affect results, especially if trade barriers are constrained to be non-discriminatory (i.e. satisfy the most-favored-nation (MFN) restriction) because in this case, when renegotiating the agreement, the importer must secure the consensus of both exporters in order to change its trade policy, so renegotiation is intrinsically a multilateral bargaining process. However, we find that our qualitative results extend to this setting with little modification. More specifically, we assume that the importer can choose a single trade policy (P or FT) for both exporters.[44] Exporters can suffer different levels of harm from protection and can have different bargaining powers. We also assume that the deadweight cost of transfers takes the form $C = c(b_1) + c(b_2)$, where b_i is the transfer from the importer to exporter i. In this setting, a contract specifies a level of damages for each exporter, (b_1^D, b_2^D). The contract is defined as a property rule if (b_1^D, b_2^D) is prohibitive or if $b_1^D = b_2^D = 0$, and as a liability rule otherwise.

In the three-country setting just described, we can show that our Propositions 1-5 continue to hold, provided some regularity conditions are satisfied.[45] The only interesting difference that arises compared with our basic model is that, when the contract is renegotiated (toward free trade), it may happen that one of the exporters compensates the importer ($b_i^e < 0$) while the other is compensated by the importer ($b_j^e > 0$). But at least under the regularity conditions mentioned in note 45, this does not affect our results.[46]

---

[44]In this way we are imposing the MFN restriction exogenously. In principle this restriction should be derived endogenously as an optimal rule, and in a complete analysis of the multi-country setting this would be an important objective; but explaining the MFN rule is not the focus of our model, so the interpretation of the exercise we undertake here is that we are optimizing the agreement rules *conditional on* the MFN restriction.

[45]In particular, we can prove our results if the function $c(b)$ is not too asymmetric around $b = 0$. Alternatively, we can prove the results for a general $c(b)$ if the exporting countries are not too asymmetric.

[46]The "competing exporters" setting we describe above is not the only multi-country setting that one could consider. A natural alternative would be that of competing *importers*: for example, a model where a given country exports a good to multiple importers, along the lines of an Armington model. However, the competing-

One may argue that in a multilateral setting, renegotiation frictions (which in our model are captured by the parameter K) are likely to be quite different than in a bilateral setting. But note that this does not affect our qualitative results, because they hold the number of countries fixed (two in our basic model, three in the multilateral extension) and examine the impact of changes in other parameters, for example the degree of uncertainty. Having said this, there is a separate and interesting question that arises in a multi-country setting: How does an expansion of the agreement membership affect the tradeoff between property and liability rules? Some scholars have argued that, if trade policies are constrained to be nondiscriminatory, renegotiation frictions are likely to become more severe as the number of member countries expands, because in case of renegotiation each individual exporter may attempt to hold up the bargain in order to extract rents. This consideration, together with the standard result from the law-and-economics literature that higher renegotiation frictions tend to favor a liability-rule approach, have led these scholars to conclude that the expansion of WTO membership should make a liability-rule approach more attractive.[47] In our model, renegotiation frictions (K) are a black box, so the model is not designed to investigate how these frictions may vary with the number of countries involved; but it is interesting to observe that, if one accepts the premise that renegotiation frictions increase with the number of member countries, our model then leads to the opposite conclusion: as we have emphasized, in the presence of costly transfers it is property rules rather than liability rules that are more likely to be optimal when renegotiation frictions are higher. This in turn suggests an intriguing possibility: the expansion of the GATT/WTO membership may have been a contributing factor in causing the shift from a liability-rule approach to a property-rule approach, which according to many legal scholars has taken place over time. We leave a more formal and complete exploration of this and other extensions to future work.

---

importers setting is conceptually more straightforward, because the renegotiation of a trade policy is then a bilateral bargain between the relevant importer and the exporting country, and hence the analysis would be similar to that of our basic model. The main difference would be that the joint gains from protection for a given exporter-importer pair will depend on the trade policies chosen by the other importers. The results of our basic model should apply to a given exporter-importer pair conditional on the other importers' policies, and therefore should apply conditional on the other importers' *equilibrium* policies as well. Based on this intuition, we conjecture that our main insights would survive also in this alternative multi-country scenario.

[47]On this point, see Schwartz and Sykes (2002) and Pauwelyn (2008, pp. 56-59).

7. Appendix

Proof of Proposition 5: Let $\hat{b}^D$ denote the level of b^D for which the $S(b^D)$ curve intersects the $R(b^D; K)$ curve, and $\tilde{b}^D$ the level of b^D for which the $S(b^D)$ curve intersects the horizontal line $\gamma = S(\gamma^* + (1 - \alpha)K) + \alpha K$, as in Figure 4.

The first step is to establish that, as in the case of costless renegotiation, under an optimal contract there can never be renegotiation from FT to P. With reference to Figure 4, this can be easily established by arguing that it is never optimal to set $b^D > \tilde{b}^D$.

The second step is to argue that, for any K, if a property rule is optimal it is never renegotiated. This step implies that, conditional on a property rule being optimal, a change in K does not affect the performance of that property rule. Consider first the discretionary property rule $b^D = 0$: if the contract is renegotiated for some γ, then an analogous argument as in the case of costless renegotiation establishes that joint surplus can be increased by raising b^D slightly above zero. Consider next a prohibitive property rule $b^D \geq b^{prohib}$: by definition this implies that the importer's threat point is FT for all γ in the support; and we know from the argument just above that under an optimal contract there can never be renegotiation from FT to P.

The third step is to argue that an increase in K weakly worsens the performance of a liability rule. We need to distinguish between two cases: at the initial level of K the contract may or may not be renegotiated for some γ, depending on whether b^D is higher or lower than $\hat{b}^D$. If $\hat{b}^D < b^D < b^{prohib}$, a marginal increase in K will have no effect. If $0 < b^D < \hat{b}^D$, a marginal increase in K will lower the expected joint payoff, as we now show.

Let us fix a value of b^D in $(0, \hat{b}^D)$, and consider the effect of a marginal increase in K. We can write the expected joint payoff as

$$E\Omega = V(FT) + \int_{R(b^D; K)}^{\infty} [\gamma - \gamma^* - c(b^D)]dH(\gamma) - \int_{S(b^D)}^{R(b^D; K)} [c(b^e(b^D, \gamma; K)) - K]dH(\gamma)$$

where the notation $b^e(b^D, \gamma; K)$ highlights the dependence of the renegotiated transfer b^e on K.

It is clear from the expression above that a marginal increase in K has three effects. First, it decreases $R(\cdot)$, the level of γ for which governments are indifferent between renegotiating and not; however this is a second-order effect, because (a) the probability distribution is atomless, and (b) as in the costless-renegotiation case, (it can easily be shown that) Ω is continuous at $\gamma = R(\cdot)$. Second, for γ in the renegotiation interval $(S(\cdot), R(\cdot))$, it has a direct negative effect,

since it is a deadweight cost. And third, for $\gamma \in (S(\cdot), R(\cdot))$, it has an indirect effect through b^e. We now argue that this indirect effect may be positive or negative, but it can never outweigh the direct negative effect.

Recall that, when b^D and γ are in the relevant renegotiation region (FT_R), the joint payoff in this region is $\gamma^* - \gamma - c(b^e) - K$. We will show that any indirect effect of the increase in K through b^e cannot outweigh its direct negative effect. Note that this indirect effect may be positive or negative, and it is more likely to be positive when α is low, so that the exporting country bears a large share of the cost. To see this intuitively, recall first that in the FT_R region the exporting government makes a transfer to the importing government ($b^e < 0$). If the exporter bears most of the renegotiation cost, the transfer b^e will need to go down (in absolute value), in order to satisfy the exporter's individual-rationality constraint, and this is beneficial to joint surplus because it reduces the deadweight loss. However we now argue that even if the indirect effect is positive, it cannot outweigh the direct negative effect.

Focusing on the FT_R region, Home's surplus over its disagreement utility is $S(b^D) - S(b) - \gamma - \alpha K$, and Foreign's surplus is $\gamma^* + b - b^D - (1 - \alpha)K$. The renegotiated transfer b^e is the value of b that maximizes the Nash product

$$NB(b, b^D) = [S(b^D) - S(b) - \gamma - \alpha K]^\sigma \cdot [\gamma^* + b - b^D - (1 - \alpha)K]^{1-\sigma}$$

Recalling that $c'(b) > -1$ for all b by assumption, it suffices to show that $\frac{\partial b^e}{\partial K} < 1$.

It is convenient to work with the logarithm of NB, which is legitimate provided that σ is strictly between zero and one, so that each government gets a strictly positive surplus ($S(b^D) - S(b) - \gamma - \alpha K > 0$ and $\gamma^* + b - b^D - (1 - \alpha)K > 0$). The extension of the proof to the extreme cases $\sigma = 0$ and $\sigma = 1$ is straightforward.

By the implicit function theorem, $\frac{\partial b^e}{\partial K} = -\frac{\partial^2 \ln NB / \partial b \partial K}{\partial^2 \ln NB / \partial b^2}\big|_{b=b^e}$. We will show that $\frac{\partial^2 \ln NB}{\partial b \partial K} < -\frac{\partial^2 \ln NB}{\partial b^2}$ for all b.

Differentiating $\ln NB$, we obtain

$$\frac{\partial \ln NB}{\partial b} = -\frac{\sigma S'(b)}{S(b^D) - S(b) - \gamma - \alpha K} + \frac{1 - \sigma}{\gamma^* + b - b^D - (1 - \alpha)K}$$

$$\frac{\partial^2 \ln NB}{\partial b \partial K} = -\frac{\alpha \sigma S'(b)}{[S(b^D) - S(b) - \gamma - \alpha K]^2} + \frac{(1 - \alpha)(1 - \sigma)}{[\gamma^* + b - b^D - (1 - \alpha)K]^2}$$

$$\frac{\partial^2 \ln NB}{\partial b^2} = -\frac{\sigma S''(b)}{S(b^D) - S(b) - \gamma - \alpha K} - \frac{\sigma S'(b)^2}{[S(b^D) - S(b) - \gamma - \alpha K]^2} - \frac{1 - \sigma}{[\gamma^* + b - b^D - (1 - \alpha)K]^2}$$

Notice that $\frac{\partial^2 \ln NB}{\partial b^2} < 0$, ensuring that the SOC is satisfied. Next note that if α is sufficiently

close to zero $\frac{\partial^2 \ln NB}{\partial b \partial K}$ is positive, and hence $\frac{\partial b^e}{\partial K} > 0$, consistently with the intuitive discussion above.

Recalling that $S'(b) > 0$, it follows that

$$\frac{\partial^2 \ln NB}{\partial b \partial K} < \frac{1 - \sigma}{[\gamma^* + b - b^D - (1 - \alpha)K]^2}.$$

Moreover, recalling that $S''(b) > 0$ and noting that the importer's surplus from the renegotiation $S(b^D) - S(b) - \gamma - \alpha K$ is nonnegative, we obtain

$$-\frac{\partial^2 \ln NB}{\partial b^2} > \frac{1 - \sigma}{[\gamma^* + b - b^D - (1 - \alpha)K]^2}$$

which implies $\frac{\partial^2 \ln NB}{\partial b \partial K} < -\frac{\partial^2 \ln NB}{\partial b^2}$ for all b, and hence $\frac{\partial b^e}{\partial K} < 1$, as claimed.

We have established that an increase in K weakly worsens the performance of a liability rule, and does not affect the performance of an optimal property rule. The claim of Proposition 5 can then be proved as follows. Suppose that at the initial level of K the optimum is a property rule; then an increase in K will make a liability rule (weakly) less attractive than the optimal property rule, hence the optimum will still be a property rule. It remains to argue that, if at the initial level of K the optimum is a liability rule, an increase in K *may* cause a switch to a property rule. This can be established by example. Consider changing K from a prohibitive level (such that renegotiation never occurs) to zero, and focus on a case where uncertainty is large but without renegotiation a discretionary property rule is optimal, along the lines described in section 3. In this case, a liability rule must become optimal as K is dropped to zero, as indicated by Proposition 2. **QED**

Proof of Proposition 6

In this extended family of contracts, we can think of a pair (b^D, b^{FT}) as representing a contract. For each contract and state of the world, $(b^D, b^{FT}; \gamma)$, there will be one of four possible equilibrium outcomes: (i) the importer chooses P without renegotiating; (ii) the importer's threat point is P but the governments renegotiate to policy FT; (iii) the importer chooses FT without renegotiating; (iv) the importer's threat point is FT but the governments renegotiate to policy P. The first step of the analysis is to characterize the mapping from $(b^D, b^{FT}; \gamma)$ to these four possible outcomes. One way to proceed is to build on the graphical apparatus of Figure 1: we continue working within the (b^D, γ) space and think of b^{FT} as a parameter that shifts the key curves in this space.

As we will show later, it can never be optimal to set $b^{FT} > 0$ or $b^D < 0$. Since proving this claim involves a tedious and taxonomic argument, we postpone this argument to a later part of this proof, and here we focus on the intuitive case where $b^{FT} \leq 0$ and $b^D \geq 0$.

Let us start by characterizing the locus of points where the importer is indifferent between the two threat points (P and FT), for a given b^{FT}. Clearly, the importer is indifferent between the two threat points when $\gamma = S(b^D) - S(b^{FT})$. This threat-point-indifference curve is depicted in Figure 4. Note that introducing $b^{FT} < 0$ in the contract shifts the threat-point-indifference curve upwards relative to Figure 1.

Next we ask: given b^{FT}, what are the regions of the (b^D, γ) space in which governments renegotiate the contract? Let us first derive the region in which the threat point is P but governments renegotiate toward FT (which we continue to label FT_R). It is immediate to verify that the threat point is P iff $\gamma > S(b^D) - S(b^{FT})$, and that in this case governments will renegotiate to FT iff $\gamma < S(b^D) - S(b^D - \gamma^*) = R(b^D)$. Notice that this latter condition is exactly the same as in the case of b^D-only contracts. Intuitively, conditional on the threat point being P the level of b^{FT} does not affect the outcome. The $R(b^D)$ curve is depicted in Figure 4, and is the same as the $R(b^D)$ curve in Figure 1. Thus, the region FT_R is the region above the $\gamma = S(b^D) - S(b^{FT})$ curve and below the $R(b^D)$ curve. Note for future reference that these curves intersect for $b^D = \gamma^* + b^{FT}$, and note also that if b^{FT} is sufficiently large and negative the FT_R region will be empty. In Figure 4 we depict the case in which the FT_R region overlaps with the positive quadrant, which (as we show below) must be the case at an optimal contract.

We next characterize the region where the threat point is FT but governments renegotiate toward P (which we continue to label P_R). Clearly the threat point is FT iff $\gamma < S(b^D) - S(b^{FT})$, and it is easy to show that in this case governments will renegotiate toward P iff $\gamma < S(b^{FT} + \gamma^*) - S(b^{FT})$. It can be easily verified that $\gamma = S(b^{FT} + \gamma^*) - S(b^{FT})$ is just the horizontal line that goes through the point of intersection between the $\gamma = S(b^D) - S(b^{FT})$ curve and the $R(b^D)$ curve. The P_R region is therefore the region that lies above this horizontal line and below the $\gamma = S(b^D) - S(b^{FT})$ curve, as depicted in Figure 4.

Having characterized the mapping from $(b^D, b^{FT}; \gamma)$ to the four possible outcomes, we can now turn to the characterization of the optimal contract.

We start by extending the result of Proposition 1, which is an intermediate step toward proving Proposition 6. We argue that it can never be strictly optimal to set $b^D > \gamma^* + b^{FT}$, and the optimal contract never induces renegotiation toward P, while it does induce renegotiation

toward FT for an intermediate range of γ.

We will suppose by contradiction that it is strictly optimal to set $b^D > \gamma^* + b^{FT}$ and will show that the initial contract can be (weakly) improved upon. We can write the expected joint surplus as

$$E\Omega(b^D, b^{FT})|_{b^D \geq \gamma^* + b^{FT}} = V(FT) + \int_{S(b^D)-S(b^{FT})}^{\infty} [\gamma - \gamma^* - c(b^D)]dH(\gamma) \tag{7.1}$$

$$+ \int_{S(b^{FT}+\gamma^*)-S(b^{FT})}^{S(b^D)-S(b^{FT})} [\gamma - \gamma^* - c(b^e(b^{FT};\gamma))]dH(\gamma) - \int_{0}^{S(b^{FT}+\gamma^*)-S(b^{FT})} c(b^{FT})dH(\gamma)$$

where (with a slight abuse of notation) $b^e(b^{FT};\gamma)$ denotes the equilibrium transfer in region P_R; note that b^e depends only on b^{FT} and not on b^D. To understand this expression, refer to Figure 4 and notice that if $b^D > \gamma^* + b^{FT}$ there are three relevant intervals of γ: for $\gamma > S(b^D) - S(b^{FT})$, we are in region P and the joint surplus is $V(FT) + \gamma - \gamma^* - c(b^D)$; for $S(b^{FT} + \gamma^*) - S(b^{FT}) < \gamma < S(b^D) - S(b^{FT})$, we are in region P_R and the joint surplus is $V(FT) + \gamma - \gamma^* - c(b^e(b^{FT};\gamma))$; and for $\gamma < S(b^{FT} + \gamma^*) - S(b^{FT})$ we are in region FT and hence the joint surplus is $V(FT) - c(b^{FT})$.

We can now write down the partial derivatives of $E\Omega$:

$$\frac{\partial E\Omega}{\partial b^D}\bigg|_{b^D \geq \gamma^* + b^{FT}} = -c'(b^D)[1 - H(S(b^D) - S(b^{FT}))] + (1 + c'(b^D))[c(b^D) - c(b^e(\cdot))]h(S(b^D) - S(b^{FT}))$$

and

$$\frac{\partial E\Omega}{\partial b^{FT}}\bigg|_{b^D \geq \gamma^* + b^{FT}} = -(1 + c'(b^{FT}))[c(b^D) - c(b^e(\cdot))]h(S(b^D) - S(b^{FT})) \tag{7.2}$$

$$- \int_{S(b^{FT}+\gamma^*)-S(b^{FT})}^{S(b^D)-S(b^{FT})} \frac{dc(b^e(b^{FT};\gamma))}{db^{FT}}dH(\gamma) - c'(b^{FT})H(S(b^{FT} + \gamma^*) - S(b^{FT}))$$

where we have used the fact that Ω is continuous at the border between the FT region and the P_R region (i.e. at $\gamma = S(b^{FT} + \gamma^*) - S(b^{FT})$).

We also note for future reference that, in analogy with the result of Lemma 2, one can show that $b^e(b^{FT};\gamma)$ is increasing in b^{FT}; intuitively, a higher b^{FT} worsens the threat point for the importer and hence the importer gets a worse deal in the renegotiation. Finally, recall that in the P_R region $\gamma^* + b^{FT} < b^e < S^{-1}(\gamma + S(b^{FT}))$.

There are two cases to consider, depending on whether $b^D > 0$ or $b^D = 0$.

Suppose first that $b^D > 0$ at the initial contract. In this case we can improve over the initial contract by lowering b^D to $\max\{0, \gamma^* + b^{FT}\}$. From expression 7.1 it is clear that this will increase $E\Omega$, because it induces no change in the policy and (i) for states γ that lie above $S(b^D) - S(b^{FT})$ before and after the change, the transfer b^D is reduced, and (ii) for states γ that lie below $S(b^D) - S(b^{FT})$ before the change but above $S(b^D) - S(b^{FT})$ after the change, the transfer goes from b^e to b^D, which is an improvement since $b^e > \gamma^* + b^{FT}$.

Next suppose $b^D = 0$ at the initial contract. In this case the initial contract can be dominated by increasing b^{FT} slightly toward zero. To see this, notice that (i) since we have supposed that $b^D \geq \gamma^* + b^{FT}$, we have $b^{FT} < 0$ and hence $c'(b^{FT}) < 0$; and (ii) given $b^D = 0$, in region P_R we have $\gamma < -S(b^{FT})$, and hence $b^e < S^{-1}(\gamma + S(b^{FT})) < 0$, and recalling that b^e is increasing in b^{FT}, this in turn implies $\frac{dc(b^e(b^{FT};\gamma))}{db^{FT}} < 0$. These two observations together imply that $\frac{\partial E\Omega}{\partial b^{FT}}\big|_{b^D \geq \gamma^* + b^{FT}} > 0$ when evaluated at $b^D = 0$.

To summarize, we have just shown that the result of Proposition 1 extends to this more general class of contracts, in the sense that we can focus without loss of generality on contracts with $b^D < \gamma^* + b^{FT}$, and the optimal contract never induces renegotiation toward P, while it does induce renegotiation toward FT for an intermediate range of γ.

We can now turn to proving the claims made in Proposition 6. It is convenient to start with the case of large uncertainty (Proposition 6(ii)).

Large uncertainty.

For our purposes it suffices to focus on the case of full support, i.e. $\gamma \in (0, \infty)$. Recall that we are focusing on the case where $b^D \geq 0$ and $b^{FT} \leq 0$ (we show later that it can never be optimal to set $b^D < 0$ or $b^{FT} > 0$).

Let $\tilde{b}^D$ denote the optimal value of b^D conditional on $b^{FT} = 0$. Given our results above and our focus on $b^D \geq 0$, it follows that $0 \leq \tilde{b}^D \leq \gamma^*$. We now argue that, starting from $(b^{FT} = 0, b^D = \tilde{b}^D)$, we can raise expected joint surplus by making b^{FT} slightly negative. Decreasing b^{FT} slightly has no impact on the policy allocation, but it has two effects on the expected equilibrium transfer. First, for $\gamma < S(b^D)$ there is now a small transfer b^{FT}, which introduces a cost, but this is a second order cost since $c'(0) = 0$. Second, for γ just above $\gamma = S(b^D)$ the threat point switches from P to FT, so for these states, before the change governments renegotiate toward FT and the equilibrium transfer is nonnegligible, and after the change the importer chooses FT without renegotiating and the transfer is negligible (because

b^{FT} is close to zero); this is a first-order beneficial effect. Note that, within the renegotiation region FT_R, decreasing b^{FT} has no impact on the threat point, hence it does not affect the equilibrium transfer.

To see this more formally, let us write down the expected joint surplus as a function of b^D and b^{FT}. As we argued above we can focus on the case $b^D \leq \gamma^* + b^{FT}$. We can then write the expected joint surplus as

$$E\Omega(b^D, b^{FT})|_{b^D \leq \gamma^* + b^{FT}} = V(FT) + \int_{R(b^D)}^{\infty} [\gamma - \gamma^* - c(b^D)]dH(\gamma) \tag{7.3}$$

$$- \int_{S(b^D)-S(b^{FT})}^{R(b^D)} c(b^e(b^D; \gamma))dH(\gamma) - \int_{0}^{S(b^D)-S(b^{FT})} c(b^{FT})dH(\gamma)$$

Differentiating this expression with respect to b^{FT} and evaluating at $(b^{FT} = 0, b^D = \tilde{b}^D)$ we obtain

$$\frac{\partial E\Omega}{\partial b^{FT}}\bigg|_{(b^{FT}=0,b^D=\tilde{b}^D)} = -h(S(\tilde{b}^D))c(b^e(\tilde{b}^D; S(\tilde{b}^D))) < 0$$

where we have used the facts that $S(b^{FT}) = b^{FT} + c(b^{FT})$ and $c(0) = c'(0) = 0$; it can easily be shown that $b^e(\tilde{b}^D; S(\tilde{b}^D)) \neq 0$; and recall that we are assuming a large enough support of γ, hence $h(S(\tilde{b}^D)) > 0$. We can conclude that, when the support of γ is large enough, $b^{FT} = 0$ cannot be optimal, and coupled with the fact that the optimal b^{FT} cannot be positive (as we next argue), this implies that the optimal b^{FT} is strictly negative.

We now rule out the possibility that $b^D = 0$ at an optimal contract. Given our results above, the only case we need to rule out is $b^D = 0 \leq \gamma^* + b^{FT}$. Letting $\tilde{b}^{FT}$ denote the optimal value of b^{FT} conditional on $b^D = 0$, we can write

$$\frac{\partial E\Omega}{\partial b^D}\bigg|_{(b^D=0,b^{FT}=\tilde{b}^{FT})} = -\int_{-S(\tilde{b}^{FT})}^{R(0)} \frac{dc(b^e(b^D; \gamma))}{db^D}dH(\gamma) + h(-S(\tilde{b}^{FT}))[c(b^e(0; -S(\tilde{b}^{FT})) - c(\tilde{b}^{FT})]$$

In this case it is immediate to establish that in the FT_R region $b^e \leq S^{-1}(-\gamma) < 0$. It follows that at the lower border of the FT_R region, where $\gamma = -S(\tilde{b}^{FT})$, it must be $b^e \leq \tilde{b}^{FT} \leq 0$. This implies $c(b^e(0; -S(\tilde{b}^{FT})) \geq c(\tilde{b}^{FT})$, so the second term of the expression above is nonnegative. Also recall that $\frac{\partial b^e(b^D;\gamma)}{\partial b^D} > 0$, hence $\frac{dc(b^e(b^D;\gamma))}{db^D} < 0$. We can conclude that $\frac{\partial E\Omega}{\partial b^D}\big|_{(b^D=0,b^{FT}=\tilde{b}^{FT})} > 0$, and hence $b^D = 0$ cannot be optimal. This, together with the fact that the optimal b^D cannot be negative (as we argue below), implies that the optimal b^D is strictly positive.

41

We now return to our earlier claim that it cannot be optimal to set $b^D < 0$ or $b^{FT} > 0$ (recall that we just ruled out the possibilities $b^D = 0$ and $b^{FT} = 0$, so we can focus on strict inequalities). To establish this claim, we need to rule out several possibilities:

(a) $b^D < 0$ and $b^{FT} < 0$. We need to distinguish two subcases:

(a_i) $b^D \geq \gamma^* + b^{FT}$. In this case it is easy to show that the equilibrium transfer b^e in the renegotiation region (P_R) is negative. Our strategy to improve on the initial contract depends on whether b^D is higher or lower than b^e at the initial contract. If $b^D < b^e$, we can improve on the initial contract by increasing b^D slightly; to see this, refer to expression 7.1 and note that in this case $c(b^D) > c(b^e)$ and $c'(b^D) < 0$, therefore $\frac{\partial E\Omega}{\partial b^D}\big|_{b^D \geq \gamma^* + b^{FT}} > 0$. If $b^D > b^e$, then we can improve on the initial contract by increasing b^{FT} slightly, because $\frac{\partial E\Omega}{\partial b^{FT}}\big|_{b^D \geq \gamma^* + b^{FT}} > 0$; to see this, note that in this case $c(b^D) - c(b^e) < 0$, $c'(b^{FT}) < 0$ and $\frac{dc(b^e(b^{FT};\gamma))}{db^{FT}} < 0$ (and recall the assumption that $1 + c'(\cdot) > 0$ for any transfer level).

(a_{ii}) $b^D < \gamma^* + b^{FT}$. Also in this case the equilibrium transfer b^e in the renegotiation region (FT_R) is negative. Our strategy to improve on the initial contract depends on whether b^{FT} is higher or lower than b^e in absolute level. If $|b^e| > |b^{FT}|$, we can improve on the initial contract by increasing b^D slightly toward zero. This has three first-order beneficial effects: (i) it reduces the transfer (in absolute value) for states $\gamma > R(b^D)$, where the importer chooses P without renegotiating; (ii) it improves the threat point for the importer in the FT_R region and hence it makes b^e less negative; (iii) for states just above $\gamma = S(b^D) - S(b^{FT})$, before the change governments renegotiate toward FT and after the change the importer chooses FT without renegotiating, thus the equilibrium transfer switches from b^e to b^{FT}; since we are focusing on the case $|b^e| > |b^{FT}|$, also this effect is beneficial.

If $|b^e| < |b^{FT}|$, on the other hand, we can improve on the initial contract by increasing b^{FT} slightly toward zero. This has two beneficial first-order effects: (i) it reduces the transfer (in absolute value) for states $\gamma < S(b^D) - S(b^{FT})$, where the importer chooses FT without renegotiating, and (ii) for states just above $\gamma = S(b^D) - S(b^{FT})$, the equilibrium transfer switches from b^{FT} to b^e; since we are focusing on the case $|b^e| < |b^{FT}|$, this effect is beneficial.

(b) $b^D < 0$ and $b^{FT} > 0$.

It can be easily shown that we can lower b^{FT} to zero, and in fact we can make it slightly negative, without affecting the policy allocation or the equilibrium transfer for any γ. This takes us back to the previous case where $b^D < 0$ and $b^{FT} < 0$, which we already ruled out.

(c) $b^D > 0$ and $b^{FT} > 0$. Here our strategy to improve on the initial contract depends on

whether b^D is higher or lower than γ^*. If $b^D < \gamma^*$, we can improve on the initial contract by increasing b^{FT} to zero. And if $b^D < \gamma^*$, the initial contract can be improved upon by increasing both b^D and b^{FT} toward zero in such a way that $S(b^D) - S(b^{FT})$ is kept constant. We leave the proof of these claims to the reader.

Finally, the claim that an optimal contract entails $b^D < \gamma^*$ follows from the fact that an optimal contract entails $b^D < \gamma^* + b^{FT}$ and $b^{FT} < 0$.

Small uncertainty.

We can now turn to the case of small uncertainty (Proposition 6(i)).

The first observation is that a noncontingent allocation (where the same policy is chosen for all γ in the support) can be implemented at zero cost (i.e. with no transfers occurring in equilibrium) with a property rule. Thus, conditional on a noncontingent allocation being optimal, a property rule is optimal. We next show that if the support of γ is sufficiently small, a noncontingent allocation is indeed optimal.

Let the support of γ be given by $(\gamma^* - \varepsilon, \gamma^* + \varepsilon)$; note that we are considering a symmetric support, but the argument is easily extended to the case of an asymmetric support. Consider a contingent allocation with threshold $\hat{\gamma}_\varepsilon \in (\gamma^* - \varepsilon, \gamma^* + \varepsilon)$ (we use the subscript ε because we need to allow this allocation to vary as we drive ε to zero). We have shown above that at an optimum it must be the case that for $\gamma = \hat{\gamma}_\varepsilon$ the importer is indifferent between choosing P without renegotiating and renegotiating toward FT. In other words, it must be $\hat{\gamma}_\varepsilon = R(b^D) = S(b^D) - S(b^D - \gamma^*)$. This implies that for states γ just above $\hat{\gamma}_\varepsilon$ the importer will pay a transfer b^D that is close to $R^{-1}(\hat{\gamma}_\varepsilon)$; clearly, this transfer does not become small as ε gores to zero. For states γ just below $\hat{\gamma}_\varepsilon$ the governments will renegotiate and the equilibrium transfer may be lower, but this transfer is unrelated to ε and hence does not become small as the support shrinks.

Now consider replacing this contingent allocation with a noncontingent allocation where policy FT is chosen in all states (and no transfers are incurred). As $\varepsilon \to 0$ this noncontingent allocation must dominate, because it implies a non-negligible savings in transfer costs for each state γ, while the associated loss in terms of policy efficiency is at most of magnitude ε for each state γ. Note that this argument holds even if the threshold $\hat{\gamma}_\epsilon$ approaches one of the bounds of the support as $\varepsilon \to 0$. We have thus shown that if the support is sufficiently small, a noncontingent allocation must be optimal, and therefore a property rule is optimal. **QED**

References

Bagwell, Kyle and Robert W. Staiger (1990), "A Theory of Managed Trade," *American Economic Review*, September.

Bagwell Kyle and Robert W. Staiger (1999), "An Economic Theory of GATT," *American Economic Review*, March.

Bagwell, Kyle and Robert W. Staiger (2001), "Domestic Policies, National Sovereignty and International Economic Institutions," *Quarterly Journal of Economics* 116(2), pp 519-562.

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 (1987), "Politically Realistic Objective Functions and Trade Policy," *Economics Letters* 24, pp. 287-90.

Baldwin, Richard and Frederic Robert-Nicoud (2007), "Entry and Asymmetric Lobbying: Why Governments Pick Losers," *Journal of the European Economics Association*, vol. 5(5), pp 1064-1093.

Beshkar, Mostafa (2010a), "Optimal Remedies in International Trade Agreements," *European Economic Review* 54(3), pp. 455-466.

Beshkar, Mostafa (2010b), "Third-Party-Assisted Renegotiation of Trade Agreements," unpublished manuscript, Yale University.

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.

Copeland, Brian R. (1990), "Strategic Interaction Among Nations: Negotiable and Non-negotiable Trade Barriers," *Canadian Journal of Economics* 23, pp. 84-108.

Costinot, Arnaud (2008), "A Comparative Institutional Analysis of Agreements on Product Standards," *Journal of International Economics* 75(1),pp. 197-213.

Ederington, Josh (2001), "International Coordination of Trade and Domestic Policies," *The American Economic Review* 91(5), pp. 1580-93.

Fernandez, Raquel and Dani Rodrik (1991). "Resistance to reform: Status quo bias in the presence of individual-specific uncertainty." *American Economic Review*, 81, pp. 1146-1156.

GATT, (1953), "Pakistan Export Fee and Duty on Raw Jute: Note by the Executive Secretary," GATT L/82/Add.1, March 26.

Goldstein, Judith and Lisa L. Martin (2002), "Legalization, Trade Liberalization and Domestic Politics: A Cautionary Note." *International Organization* 54(3): 603-632.

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

Grossman, Gene M. and Elhanan Helpman (1996): "Rent Dissipation, Free Riding, and Trade Policy," *European Economic Review*, 40, 795-803.

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

Hoda, Anwarul (2001), *Tariff Negotiations and Renegotiations under the GATT and the WTO: Procedures and Practices*, Cambridge University Press, Cambridge England.

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

Horn, Henrik (2006), "National Treatment in the GATT," *The American Economic Review* 96(1), pp. 394-404.

Horn, Henrik, Giovanni Maggi and Robert W. Staiger (2010), "Trade Agreements as Endogenously Incomplete Contracts," *American Economic Review*, March.

Howse, Robert and Robert W. Staiger (2005), "United States–Anti-Dumping Act of 1916 – 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.

Hufbauer, Gary C. and Kimberly A. Elliott (1994), *Measuring the Costs of Protection in the United States.* Institute for International Economics: Washington DC.

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.

Klimenko, Mikhail, Garey Ramey and Joel Watson (2008), "Recurrent Trade Agreements and the Value of External Enforcement," *Journal of International Economics* 74(2), pp. 475-499, March.

Krueger, Anne (1990). "Asymmetries in policy between exportables and import-competing goods." in *The political economy of international trade*, edited by Ronald Jones and Anne Krueger. Cambridge MA: Basil Blackwell.

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

Ludema, Rodney (2001), "Optimal international trade agreements and dispute settlement procedures," *European Journal of Political Economy* 17, 355–376.

Maggi, Giovanni (1999), "The Role of Multilateral Institutions in International Trade Cooperation," *American Economic Review* 89, 190-214.

Maggi, Giovanni and Robert W. Staiger (2011), "The Role of Dispute Settlement Procedures in International Trade Agreements," *Quarterly Journal of Economics*, February.

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

Maskin, Eric and John Moore (1999), "Implementation and Renegotiation," *The Review of Economic Studies* 66, 39-56.

Park, Jee-Hyeong (forthcoming), "Enforcing International Trade Agreements with Imperfect Private Monitoring," *Review of Economic Studies.*

Pauwelyn, Joost (2008), *Optimal Protection of International Law: Navigating between European Absolutism and American Volunteerism*, Cambridge University Press, Cambridge UK.

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

Rosendorff, B. Peter and Helen V. Milner (2001), "The Optimal Design of International Trade Institutions: Uncertainty and Escape," *International Organization* 55(4), Winter.

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.

Segal, Ilya R. and Michael D. Whinston (2002), "The Mirrlees Approach to Mechanism Design with Renegotiation (with Applications to Hold-up and Risk Sharing)," *Econometrica* 70, pp. 1-45.

Shavell, Steven (1984), "The Design of Contracts and Remedies for Breach," *Quarterly Journal of Economics*, February.

Sykes, Alan O. (1991), "Protectionism as a 'Safeguard': A Positive Analysis of the GATT 'Escape Clause' with Normative Speculations," *The University of Chicago Law Review* Vol. 58, No. 1 (Winter, 1991), pp. 255-305

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.

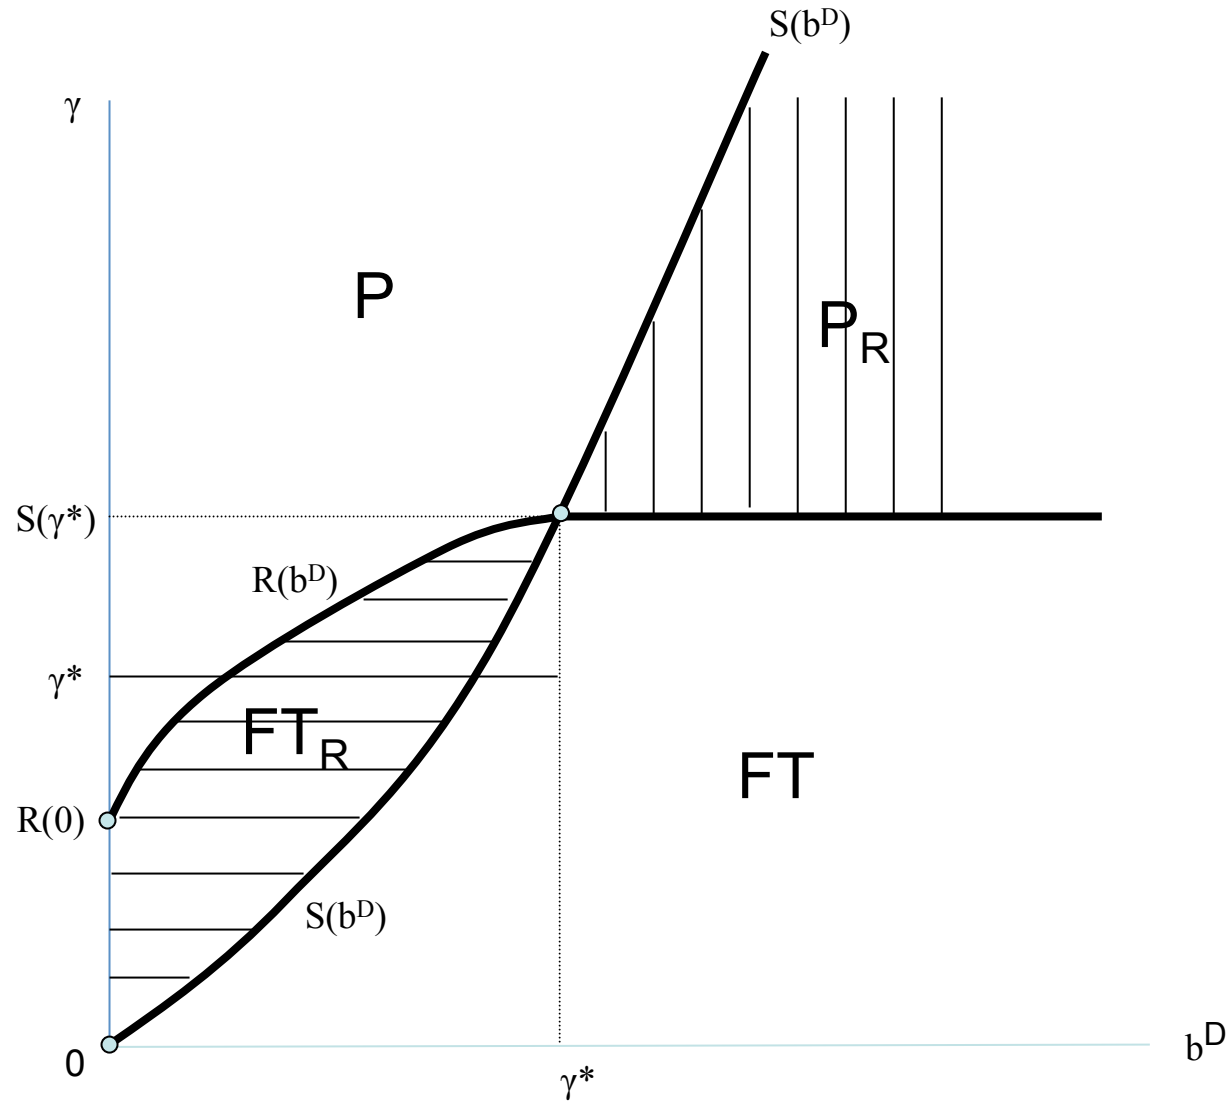**Watson,** Joel (2007), "Contract, Mechanism Design, and Technological Detail," *Econometrica* 75(1), 55-81.
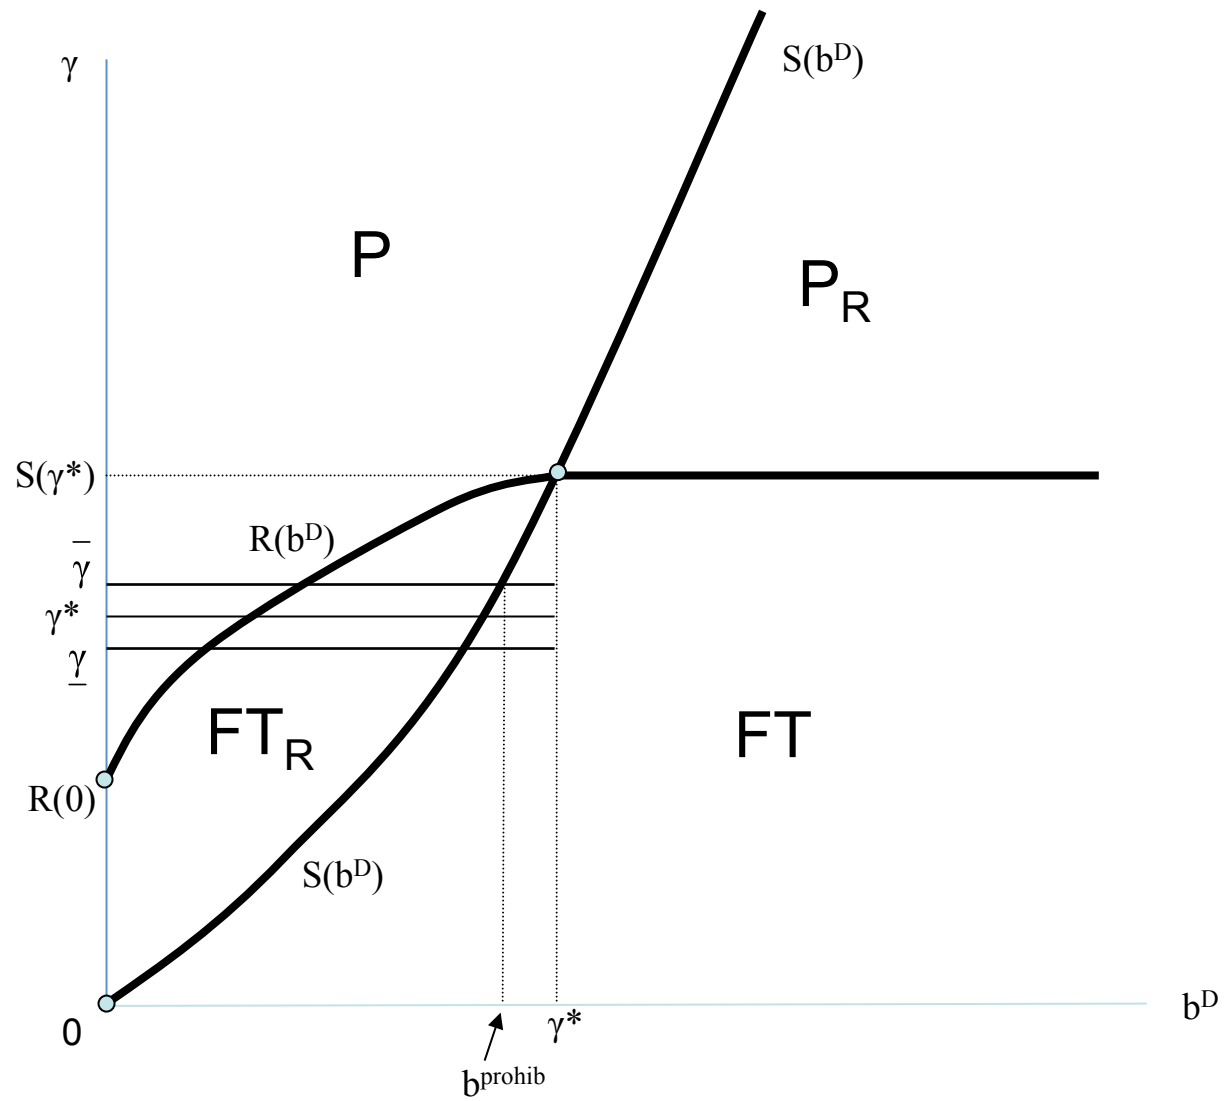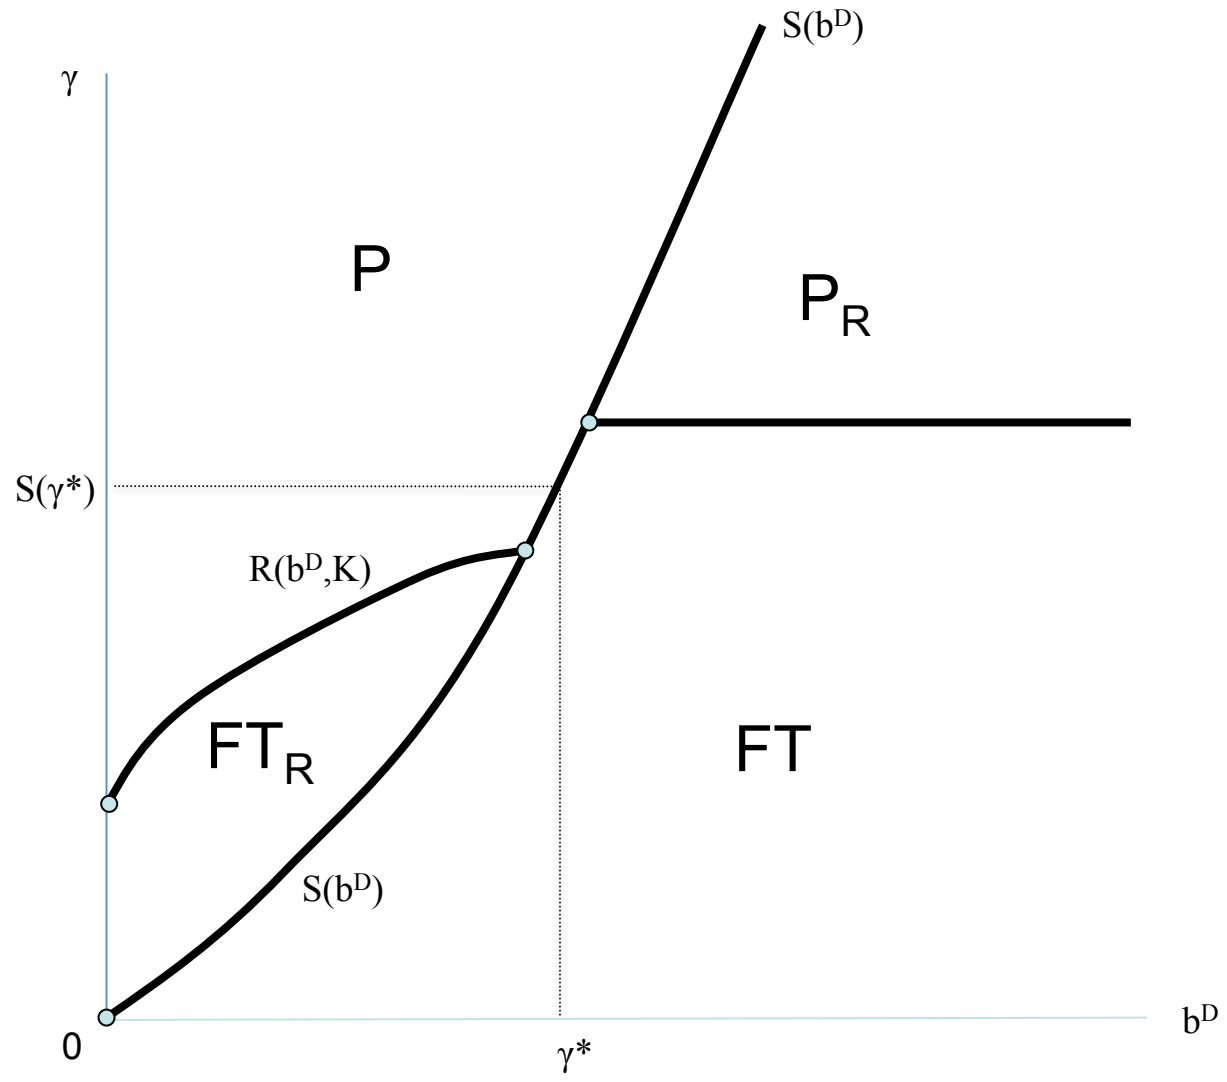
Figure 1

Figure 2

Figure 3

Figure 4