The GATT/WTO as an Incomplete Contract

Henrik Horn (IIES, Stockholm University)
Giovanni Maggi (Princeton University and NBER)
Robert W. Staiger (Stanford University and NBER)

May 2006

Abstract

We propose a model of trade agreements in which contracting is costly, and as a consequence the optimal agreement may be incomplete. In spite of its simplicity, the model yields rich predictions on the structure of the optimal trade agreement and how this depends on the fundamentals of the contracting environment. We argue that taking contracting costs explicitly into account can help explain a number of key features of the observed trade agreements.

1. Introduction

The World Trade Organization (WTO) regulation of trade in goods – the General Agreement on Tariffs and Trade (GATT) – is a highly incomplete contract.[1] It directly binds only trade policies, leaving significant discretion over domestic policy instruments with trade impact to national governments. The policies that are bound, are for the most part bound rigidly, and are thus not highly adaptable to stochastic economic or political shocks. And most of the provisions are vaguely worded: the actual ambit of the agreement is therefore left largely to be determined by adjudicating bodies.

A sizeable economic literature seeks to shed light on various aspects of this incompleteness. The typical approach is to impose exogenous restrictions on the set of policy instruments that can be included in a trade agreement, and examine what the agreement can accomplish given these limitations.[2]

This literature has helped shed light on several important aspects of the agreement. But it also exhibits an important limitation: the incompleteness of the agreement is assumed rather than endogenously derived. It is easy to accept as a general statement that a trade agreement

[1] For simplicity, we will use the term "GATT" in a somewhat imprecise sense. More formally, the WTO regulates trade in goods through "GATT 1994", consisting of the original 1947 GATT agreement plus a number of agreed modifications and interpretations, plus a number of special agreements, such as those governing safeguards, anti-dumping, technical barriers to trade and health measures.

[2] An incomplete list of papers that fall into this category is Copeland (1990), Bagwell and Staiger (2001), Battigalli and Maggi (2003) and Horn (2003).

has to be incomplete because of the immense costs that would be involved in reaching a fully efficient agreement (if this is a practical possibility at all). However, there are many different ways in which an agreement can save on contracting costs, and so there are many different forms that the incompleteness could take.

The broad purpose of this paper is to take the analysis of the GATT as an incomplete contract one step further, by endogenously determining the choice of contract form. In contrast to the existing literature, we will thus not take the structure of the agreement for given, and analyze possible consequences thereof. Instead we seek to characterize circumstances under which various forms of trade agreements may arise.

The GATT consists of a large number of provisions. But we believe that an incomplete contracting perspective can help to shed light on certain structural features of the agreement:

1. The agreement binds the level of trade instruments, whereas domestic instruments are largely left to the discretion of national governments. But internal policies have to respect the National Treatment (NT) clause, and the WTO has introduced a regulation of subsidies.

2. The bindings are rigid. But there are "escape mechanisms" that allow countries to unilaterally impose temporary protection (GATT Art. XIX) or to renegotiate bindings (GATT Art. XXVIII).

3. Tariff bindings are "weak," in the sense that they only impose an upper bound on the tariff level.

4. "Outcome-based" bindings of trade volumes or prices are not featured, but rather bindings of policies are the norm. Nevertheless, some legal instruments may approximate an outcome-based role. The legal instrument that comes closest to such a binding is the "Non-Violation" (NV) provision of GATT Art. XXIII, but this seems to have played a relatively limited role, especially since the advent of the WTO.[3]

5. The agreement requires that bindings fulfill the Most-Favored Nation (MFN) provision.

The ultimate aim of this line of research is to understand why the regulation of goods trade has taken this particular form. We believe that the analysis to follow sheds light on at least some of these core features.

The analytical starting point of the paper is the notion that legislators face two fundamental sources of difficulty when designing a trade agreement. The first is that there is significant uncertainty concerning the circumstances that will prevail during the life-time of the agreement. This uncertainty suggests that the agreement should be highly adaptable to the contingencies

[3]The exact meaning of the NV provision is hard to determine from the text itself. But it is often seen as protecting market access expectations of governments against changes in policies by their trading partners – even when these policies are not contracted over – which would have the effect of upsetting the market access that a government could have reasonably expected based on a prior GATT/WTO negotiation. It provided the central discipline on the use of domestic subsidies during the GATT era, but its role has been substantially diminished within the WTO.

that unfold. The second source of difficulty in designing a trade agreement is the wide array of policy instruments – border measures and especially internal measures – that need to be constrained in order to keep in check the governments' incentives to act opportunistically. This feature in turn suggests that the agreement should be very comprehensive in its coverage of trade-relevant policies.

Of course the complexity of the contracting environment would not constitute a problem if contracting were costless. But in reality there are important costs associated with forming a trade agreement: there are costs in terms of effort and time for working out bargaining proposals and for evaluating proposals made by others, and there are costs involved in verifying that trading partners abide by the rules of the agreement. We will refer to these costs broadly as *contracting costs*. While contracting costs can take many different forms, it is probably safe to say that they tend to be higher when the agreement is more detailed, in terms of the contingencies that it specifies and the number of policies that it seeks to constrain. This basic idea will be reflected in our formalization of contracting costs.

It is hard to dispute that in reality contracting costs constitute a severe constraining factor for the design of a trade agreement. The choice of the structure of the agreement therefore will reflect a trade off between the performance of the contract and the associated contracting costs. A key objective of this paper is hence to highlight the role that contracting costs may play for explaining core features of the GATT.

We will work with a partial equilibrium, competitive, two-country setting, where countries may experience two types of externalities: a consumption externality, and/or a production externality. These externalities give rise to a rationale for policy intervention. For simplicity, we focus on intervention in the Home country, and assume that the Home country government (alone) has access to a full set of taxation instruments in its import-competing sectors, and namely: import tariffs, distinct consumption taxes on domestically-produced products and on imported products, and production subsidies. Uncertainty will play a central role in the analysis. For tractability reasons, we will assume that uncertainty is one-dimensional, and will consider several scenarios: the source of the uncertainty may be in one of the externalities, or in the level of import demand.

In the absence of an agreement, the importing country would manipulate the terms of trade in standard fashion. This of course would lead to a globally inefficient outcome, and hence there is scope for an agreement to restrain governments from behaving opportunistically. Were it not for the production and consumption externalities, the optimal trade agreement (at least in the absence of contracting costs) would be very simple: it would just stipulate free trade and no other policy intervention. But due to the externalities, the contracting problem is substantially more complex: the first best agreement will now involve domestic policy instruments – not just border measures – and will require these policies to be state contingent.

We formalize the notion of contracting costs in a very simple way. Following an approach similar to that of Battigalli and Maggi (2002), we will assume that these costs depend both on the degree to which a contract is state contingent, and on the scope of its coverage of policy instruments. As a result of these costs, the parties may find it worthwhile to use a simpler contract form than the one required to implement the first best outcome. As pointed out by Battigalli and Maggi, there are two essential ways in which the agreement can save on contracting costs. One is that the agreement is (partially or fully) *rigid* – i.e. it is insensitive to changes in the underlying economy. The other is that it leaves *discretion* in the governments' choices of policies. In our framework, one would naturally expect that the cost of making the contract more contingent pushes toward rigidity, and the cost of contracting over policies pushes toward discretion, and this is indeed the case, but we will show that the presence of contracting costs has also much more subtle implications.

The paper proceeds as follows. Section 2 describes our model of the economy and our formalization of contracting costs. The section also presents two benchmark scenarios. One is the no-agreement outcome – that is the Nash equilibrium – and the other is the first-best outcome. The approach of the paper is to view an optimal agreement as one that maximizes global welfare minus contracting costs. The no-agreement and first-best outcomes can thus be seen as the outcomes of two extreme forms of contracting costs, the first-best outcome resulting when contracting costs are zero, and the no-agreement outcome resulting when they are sufficiently high to dominate any gains that could be had from an agreement.

Section 3 characterizes the optimal trade agreement within a simple class of contracts, and analyzes how the optimal agreement depends on contracting costs, on the degree and type of uncertainty and on the features of the underlying economy.

Our first result concerns the way contractual incompleteness varies across policy instruments. We find that the optimal agreement tends to leave more discretion on domestic policy instruments than on border measures. More specifically, while for a range of contracting costs it is optimal to bind import taxes while leaving domestic instruments discretionary, it is never optimal to leave import taxes to discretion and contract only over domestic instruments.

As contracting costs increase from zero, the optimal agreement is initially fully state-contingent, then it becomes increasingly rigid and/or discretionary, and eventually it becomes optimal to have no agreement at all. Whether the optimal agreement tends to feature rigidity or discretion, or a combination of the two, depends crucially on the nature of the uncertainty and on the demand and supply parameters. Intuitively, rigidity is relatively more attractive when uncertainty is small. On the other hand, discretion (over domestic instruments) is relatively more attractive when (i) domestic instruments are less effective at manipulating terms of trade, or in other words, *the degree of substitutability* between these instruments and import taxes is lower, since in this case the *ability* to manipulate terms of trade through domestic instruments is lower; and (ii) the importing country has less *monopoly power* in trade (which is the case when the import demand level is lower), since in this case the *incentive* to distort terms of trade through domestic instruments is lower.

The role of uncertainty depends in subtle ways on its source. We find that, if uncertainty concerns the level of externalities – or more generally state variables that are directly relevant for the first-best policy levels – then rigidity and discretion are *complementary* ways of saving on contracting costs: more specifically, the cost of discretion is lower in the presence of rigidity than in the absence of it. But if uncertainty concerns the level of import demand – or more generally state variables that are not directly relevant for the first-best policy levels – rigidity and discretion are *substitutable*, meaning that the cost of discretion is higher when the agreement is rigid.

When there is uncertainty about the level of import demand, we find that it may be optimal for the agreement to specify an escape-clause type rule, whereby governments are allowed to raise tariffs when the level of import demand is higher. However, our rationale for an escape clause is distinct from those that have been highlighted in the existing theoretical literature.[4] In particular, an escape clause can be appealing in our model when the agreement leaves discretion over domestic instruments. Since, as we have observed above, the incentive to distort these instruments for manipulating the terms of trade is stronger when the import demand level is higher, allowing higher tariffs in the high-import-demand states can be attractive as a way to mitigate this incentive.

The remaining part of the paper extends the analysis to shed light on several other core aspects of the GATT that we believe are best understood from an incomplete contracts perspective.

Section 4 investigates the role of the NT clause as a possible means of saving on contracting costs. We identify one type of NT-based agreement that can be strictly optimal in our setting: this is an agreement that – in addition to imposing the NT rule – ties down (in a rigid or contingent way) the import tariff and the production subsidy, but leaves the (common) consumption tax to the governments' discretion. This type of agreement has the virtue of allowing for some ex post flexibility, since the importing country can set the general consumption tax level in response to stochastic disturbances, but it is not sufficient to implement the first best outcome, since it allows the common tax level to be set opportunistically. We identify a simple set of conditions under which this type of agreement is indeed optimal. The key condition concerns the degree of substitutability between the consumption tax and the tariff for the purposes of manipulating terms of trade: if this degree is sufficiently low (which is the case when demand is more rigid and supply is more elastic), and if the level of contracting costs lies in an intermediate range, then an NT-based contract is strictly optimal.

Section 5 examines the usefulness of an NV provision as a means to economize on contracting costs. The NV clause allows countries to avoid contracting directly over domestic policy instruments, thereby saving on the costs of specifying and verifying the values of these instruments. On the other hand, the clause requires verification of the state of demand and supply, which is costly. Hence, we find that an NV-based contract tends to be optimal when the cost of contracting over a wide set of policy instruments is large relative to the cost of verifying the relevant state of the world.

Finally, in section 6 we argue that the presence of contracting costs may explain why GATT stipulates weak bindings – e.g. maximum tariff levels – rather than strict bindings. More specifically, we show that the optimal agreement may include *rigid* weak bindings. This type of binding combines rigidity and discretion, since the ceiling does not depend on the state of the world, and the government has discretion to set the policy below the ceiling. This finding strengthens the insight – already highlighted above – that rigidity and discretion may be complementary ways to economize on contracting costs.

An appendix provides proofs that are not contained in the body of the paper.

[4] An escape clause could be motivated for distributional reasons if the government lacked better instruments with which to redistribute income. Bagwell and Staiger (1990) show that an escape clause can be motivated for enforcement purposes when trade agreements lack external enforcement mechanisms.

2. The Model

We adopt a partial equilibrium perspective, according to which there are potentially many goods produced, consumed and traded between a home country and a foreign country. For simplicity, we concentrate on a single good, for which the home country is the natural importer. We then characterize the impact of internationally negotiated contracts on the production, consumption and trade of this good.

We are interested in exploring contracting possibilities over a rich set of instruments. To this end, we assume that the home government can use an import tariff (τ), an internal tax on consumption of the domestically produced good (t_h), an internal tax on consumption of the imported product (t_f), and a production subsidy to domestic firms (s). All instruments are expressed in specific terms. For simplicity, we assume that the foreign (exporting) government has no policies available in the sector under consideration, though our results generalize naturally to a setting in which the foreign government also makes policy choices.

The goods markets in the two countries are integrated, and prices differ only to the extent of government intervention. Throughout we focus on non-prohibitive levels of government intervention that do not choke off all trade. Let p and p^* denote the prices paid by consumers in the home and foreign country, respectively, with asterisks denoting variables in the foreign country here and throughout. Due to the possibility of consumer arbitrage, and to the absence of taxation in the foreign country, we have the following relationship between home and foreign consumer prices: $p = p^* + \tau + t_f$. For a foreign firm to sell in both countries, it must receive the same price for sales in the foreign-country market as it receives after taxes for sales in the home-country market: $q^* = p - \tau - t_f$, where q^* is the price received by a foreign firm for sales in the foreign-country market. Due to the absence of taxation or other trade costs in the foreign country, we also have that producer and consumer prices in the foreign country are equalized, or $q^* = p^*$. Finally, let q denote the home-country producer price, i.e., the price received by a home firm for sales in the home-country market. The relationship between the home-country producer price and the home-country consumer price is given by $q = p - t_h + s$.

We can express the above pricing relationships in more compact form as

$$\begin{aligned} p &= p^* + T, \text{ and} \\ q &= p^* + T + S, \end{aligned} \tag{2.1}$$

where $T \equiv \tau + t_f$ and $S \equiv s - t_h$. The arbitrage relationships in 2.1 describe the two key price wedges that play a central role in the analysis to follow; the first one is the wedge between the home-country consumer price and the foreign-country price (equal to T), and the second one is the wedge between the home-country producer price and the foreign-country price (equal to $T + S$). Note that τ and t_f are perfectly substitutable policy instruments (only their sum matters), and the same is true for s and t_h (only their difference matters). Thus, while it is appropriate to refer to τ as a "border measure" and to t_f, t_h and s as "internal measures," we will also sometimes refer to T as the total tax on imports, or simply as the "import tax," and to S as the "effective production subsidy."

We turn next to specifying the demand, supply and market-clearing conditions. Home and foreign demand functions take a simple linear form: $D(p) = \alpha - \beta p$, and $D^*(p^*) = \alpha^* - \beta^* p^*$, where $\alpha > 0$, $\beta > 0$, $\alpha^* > 0$, and $\beta^* > 0$. The home and foreign supply functions are also linear: $X(q) = \lambda q$, and $X^*(q^*) = \lambda^* q^*$, where $\lambda > 0$ and $\lambda^* > 0$. Market clearing requires that world demand equal world supply, or

$$D(p) + D^*(p^*) = X(q) + X^*(q^*). \tag{2.2}$$

The market clearing condition 2.2, together with the two arbitrage relationships in 2.1, yields expressions for the three market clearing prices as functions of T and S:

$$p(T,S) = [\alpha + \alpha^* + (\beta^* + \lambda^*)T - \lambda S]/\Upsilon,$$

$$q(T,S) = [\alpha + \alpha^* + (\beta^* + \lambda^*)T + (\beta + \beta^* + \lambda^*)S]/\Upsilon, \text{ and}$$
$$p^*(T,S) = q^*(T,S) = [\alpha + \alpha^* - (\beta + \lambda)T - \lambda S]/\Upsilon,$$

where $\Upsilon \equiv \lambda + \lambda^* + \beta + \beta^*$. At the market clearing prices, home import volume, M, is equal to foreign export volume, E^*, and is given by

$$M(T,S) = E^*(T,S) = [\alpha(\beta^* + \lambda^*) - \alpha^*(\beta + \lambda) - (\beta + \lambda)(\beta^* + \lambda^*)T - \lambda(\beta^* + \lambda^*)S]/\Upsilon.$$

Note that $M(T = 0, S = 0) > 0$, and hence the home country is a natural importer of the good under consideration, provided that

$$\frac{\alpha}{\lambda + \beta} > \frac{\alpha^*}{\lambda^* + \beta^*}. \tag{2.3}$$

We will henceforth assume that 2.3 holds. For future use we may also define implicitly the locus of policies that prohibit trade according to $M(T^a(S), S) \equiv 0$. Explicit calculations yield

$$T^a(S) = [\frac{\alpha}{\lambda + \beta} - \frac{\alpha^*}{\lambda^* + \beta^*}] - \frac{\lambda}{\lambda + \beta} \cdot S.$$

We assume that each government's objective corresponds to the welfare of its representative citizen. For the foreign-country government, who we recall has no policy instruments of its own in the sector under consideration, this objective is simply the sum of foreign consumer and producer surplus, which we denote by CS^* and PS^*, respectively. Hence, the objective of the foreign-country government, $W^*(T,S)$, is given by

$$W^*(T,S) = CS^*(T,S) + PS^*(T,S),$$

where

$$CS^*(T,S) \equiv \int_{p^*(T,S)}^{\alpha^*/\beta^*} D^*(p^*)dp^*; \text{ and } PS^*(T,S) \equiv \int_0^{q^*(T,S)} X^*(q^*)dq^*.$$

In the home country, in addition to home consumer and producer surplus (CS and PS, respectively), a further surplus consideration is the net revenue generated by the home-government's policy intervention. The home government's net revenue is composed of its revenue from the import tax ($T \cdot M$) minus its expenditure on the effective production subsidy ($S \cdot X$). Moreover, we allow the possibility that there may exist several kinds of externalities in the home country that introduce a divergence between national income and national welfare. Specifically, we assume that there is a positive production externality equal to σX with $\sigma > 0$, and a negative consumption externality equal to $-\gamma D$ with $\gamma > 0$ (nothing substantial in the analysis would change if the signs of the externalities were different). These externalities enter directly and separably into the representative home-country citizen's utility and do not cross borders. Hence, the home-country government's objective, $W(T, S)$, is given by the sum of consumer surplus, producer surplus, tax revenue and the valuation of the externalities associated with home-country production and consumption, or

$$W(T,S) = CS(T,S) + PS(T,S) + T \cdot M(T,S) - S \cdot X(T,S) + \sigma X(T,S) - \gamma D(T,S),$$

where

$$\begin{aligned} D(T,S) &\equiv D(p(T,S));\ X(T,S) \equiv X(q(T,S)); \\ CS(T,S) &\equiv \int_{p(T,S)}^{\alpha/\beta} D(p)dp;\ \text{ and } PS(T,S) \equiv \int_0^{q(T,S)} X(q)dq. \end{aligned}$$

2.1. The Nash equilibrium and efficient policies

We first derive the Nash equilibrium policies, which we take to represent the policy choices made in the absence of any agreement between the home and foreign governments. With the foreign government passive, the Nash equilibrium policies are defined by the two first-order conditions characterizing the home government's best-response policy choices, which simplify to

$$\begin{aligned} \frac{dW(T,S)}{dT} &= 0 \Longrightarrow \frac{E^*(S,T)}{\beta^* + \lambda^*} - T + \frac{\lambda}{\beta + \lambda}(\sigma - S) + \frac{\beta}{\beta + \lambda}\gamma = 0, \text{ and} \\ \frac{dW(T,S)}{dS} &= 0 \Longrightarrow \frac{E^*(S,T)}{\beta^* + \lambda^*} - T + \frac{\beta + \beta^* + \lambda^*}{\beta^* + \lambda^*}(\sigma - S) - \frac{\beta}{\beta^* + \lambda^*}\gamma = 0. \end{aligned}$$

These two first-order conditions define, respectively, the best-response level of T given S, which we denote $T^R(S)$, and the best-response level of S given T, which we label $S^R(T)$. These best-response functions will play an important role in what follows.

Solving the system, we may derive the following expressions for the Nash equilibrium import tax and effective production subsidy choices of the home government, which we denote by T^{NE} and S^{NE}, respectively:[5]

$$\begin{aligned} T^{NE} &= \gamma + \frac{E^*(S^{NE}, T^{NE})}{\beta^* + \lambda^*} = \gamma + \frac{p^*}{\eta^*}, \text{ and} \\ S^{NE} &= \sigma - \gamma, \end{aligned}$$

where η^* is the elasticity of the foreign export supply (itself evaluated at S^{NE} and T^{NE}).

[5] It is not hard to verify that W is jointly concave in (T, S), which ensures that the first-order conditions are sufficient.

Recalling that $T \equiv \tau + t_f$ and $S \equiv s - t_h$, we note that there are many equivalent policy combinations that correspond to the Nash policy choices T^{NE} and S^{NE}. One of these combinations is $\{\tau = \frac{p^*}{\eta^*}, t_h = t_f = \gamma, s = \sigma\}$. This particular policy combination makes it transparent that in the Nash equilibrium the home-country government sets its traditional (Johnson, 1953-54) "optimal tariff" – the inverse of the Nash equilibrium foreign export supply elasticity – to exploit its monopoly power over the terms of trade (p^*), applies a Pigouvian production subsidy at the level of the production externality, and applies a uniform Pigouvian consumption tax at the level of the consumption externality.

Finally, solving for the explicit expression for T^{NE} in terms of the underlying parameters yields

$$T^{NE} = \frac{(\beta^* + \lambda^*)[\alpha + \Upsilon\gamma - \lambda(\sigma - \gamma)] - (\beta + \lambda)\alpha^*}{(\beta^* + \lambda^*)[\Upsilon + (\beta + \lambda)]}.$$

We focus throughout on parameter combinations for which strictly positive trade occurs in the Nash equilibrium. These parameter combinations are defined by the restriction that $T^{NE} < T^a(\sigma - \gamma)$. In light of 2.3, which assures that the home country is a natural importer of the good under consideration, it is direct to verify that $T^{NE} < T^a(\sigma - \gamma)$ is assured in the absence of externalities (i.e., when $\sigma = 0$ and $\gamma = 0$), and that this restriction in effect places upper limits on the magnitude of the externality parameters σ and γ.

Having characterized the Nash equilibrium policy choices, we turn next to the globally efficient policies. The globally efficient policies are those policies that maximize "global welfare," that is, the sum of home and foreign welfare:[6]

$$W^G(T, S) \equiv W(T, S) + W^*(T, S).$$

It is direct to verify that the efficient import tax and effective production subsidy choices of the home government, which we denote by T^{eff} and S^{eff}, respectively, are given by

$$\begin{aligned} T^{eff} &= \gamma, \text{ and} \\ S^{eff} &= \sigma - \gamma. \end{aligned}$$

Hence, efficient policy combinations ensure that the relevant price wedges only reflect externalities, not terms of trade considerations. In particular, the wedge between the domestic consumer price and the foreign price (T) should be equal to the consumption externality γ (Pigouvian consumption tax), and the wedge between the domestic producer price and the foreign price ($S + T$) should be equal to the production externality σ (Pigouvian production subsidy).

At this point it is convenient to emphasize a feature of the Nash policy choices and their relation to the efficient policy choices that will turn out to be important for interpreting our results in the following sections. In particular, notice that the Nash choice of effective production subsidy is efficient ($S^{NE} = S^{eff}$), and the nature of the policy inefficiency associated with the Nash equilibrium is then entirely reflected in a Nash level of import taxes that is too high – and

[6]By defining globally efficient policies in this way, we are implicitly assuming that the two governments can transfer surplus between them in a lump sum fashion.

Nash trade volumes that are therefore too low – relative to their efficient levels ($T^{NE} > T^{eff}$). The inefficiently high level of T reflects in turn the unilateral incentive to manipulate the terms of trade with the choice of import taxes.

Therefore, it is accurate to say that the potential gains from contracting in this setting arise entirely from the ability to control the incentive to utilize import taxes to manipulate the terms of trade. As a consequence of this feature, while the agreements we consider below may impose constraints beyond the choice of import taxes, we will nevertheless refer to these agreements as "trade agreements," because they represent attempts to solve what is evidently at its core a trade – and trade policy – problem.[7]

2.2. Uncertainty

The economy described so far is deterministic. But in actuality, trade agreements are formed with very inadequate knowledge concerning a large number of different aspects of the economy, and the political trade-offs facing future governments. Indeed, a crucial problem when formulating a trade agreement is the fact that it must be appropriate under a wide range of different circumstances. It should therefore ideally be highly adaptable to changes in the underlying environment. But agreements that are adaptable in this sense, while still retaining control over the various policies, are also likely to be very costly to design and implement.

For reasons that will become clear below, it complicates the analysis greatly to include several sources of uncertainty in the model at one time. For this reason, we consider simple stochastic environments where uncertainty is one-dimensional. In particular, we will consider three cases: uncertainty in the consumption externality (γ); uncertainty in the level of domestic demand (α); and uncertainty in the production externality (σ). Exploring these three simple stochastic environments separately allows us to identify a number of central and novel features that arise when trade agreements are configured in a costly contracting environment. As will become clear below, the basic insights generated by the specification with one-dimensional uncertainty are likely to extend to more general stochastic environments. For the sake of expositional convenience, we furthermore assume that each state variable can take one of two values; that is, $\theta = \bar{\theta} \pm \Delta_\theta$, with θ denoting α, γ, or σ. We will sometimes refer to these as the state-of-the-world variables, or simply the "state" variables.

Finally, we denote expected global welfare gross of contracting costs (henceforth simply "expected gross global welfare") by $\Omega(\cdot) \equiv EW^G(\cdot)$.

2.3. The costs of contracting

We formalize contracting costs in a very stylized way. There are two kinds of contracting costs: the costs of including *state* variables in the agreement – that is, α, γ and σ, depending on which one is uncertain – and the costs of including *policy* variables (τ, t_f, s, t_h) in the agreement. We think of the cost of including a given variable in the agreement as capturing both the cost of describing this variable (i.e. defining the variable, how it should be measured etc., along the lines of the "writing costs" emphasized by Battigalli and Maggi, 2002) as well as the cost of

[7] The feature we emphasize above is quite general, as argued in Bagwell and Staiger (2001).

verifying its value ex-post. We assume that, if a variable is included in the agreement, the court automatically verifies its value ex-post, incurring the associated verification cost.[8] A broader interpretation of these contracting costs might also include negotiation costs: it is reasonable to think that negotiation costs are higher when there are more policy instruments on the table, and when there are more relevant contingencies to be discussed for a given policy instrument.

The cost of contracting over a state variable is c_s and the cost of contracting over a policy variable is c_p. We assume that, if a variable is included in the agreement, the associated cost is incurred only once, regardless of how many times that variable is mentioned in the agreement; in other words, there is no cost in "recalling" a given variable after the first time it appears in the agreement.

Summarizing, the cost of writing an agreement is given by

$$C = c_s \cdot n_s + c_p \cdot n_p,$$

where n_s and n_p denote, respectively, the number of state variables and policy variables included in the agreement.

A few examples can be useful to illustrate our assumptions on contracting costs:

Example 1: The agreement $\{\tau = 2\}$ specifies a rigid binding on the level of the home tariff, and costs c_p.

Example 2: The agreement $\{\tau + t_f = 3\}$, or $\{T = 3\}$, specifies a rigid commitment for the level of the total import tax, and costs $2c_p$.

Example 3: The agreement $\{S = \sigma\}$ specifies a state-contingent commitment for the level of the effective production subsidy, and costs $2c_p + c_s$.

It might be reasonable to assume that it is more costly to contract over internal measures (t_f, s, t_h) than over tariffs (τ), because in reality it is easier to verify border measures than internal measures. But as will become clear below, in this case our qualitative results would only be strengthened. So in the interests of parsimony, we do not introduce this distinction, and rather maintain the assumption of a common contracting cost for border and internal measures.

3. Optimal Agreements

We begin by focusing on *instrument-based* agreements, i.e. agreements that impose (possibly contingent) constraints on policy instruments. We defer until a later section the discussion of *outcome-based* agreements, i.e. agreements that impose constraints on equilibrium outcomes such as prices or trade volumes.

Observe first that since the two policy instruments τ and t_f are perfect substitutes and matter only through their sum T, when we view them as contractual variables they are perfect *complements*: constraining one of the two instruments but not the other would have no effect. The same is true for the domestic instruments s and t_h, which matter only through their

[8] Of course this is a strong assumption. In reality, the WTO verifies compliance with the agreement only if there is a complaint by one of the contracting parties. Broadly, we expect that similar qualitative insights would emerge in a richer model with verification "on demand" to the extent that verification occurs in equilibrium at least with some probability, though we discuss later in the paper some potentially important differences that might arise in a richer model of this kind.

difference S. Hence, as a starting point we can think of T and S as the relevant policy variables, and the associated contracting costs are $2c_p$ for each of these variables.

Moreover, for now we restrict our search to agreements that impose separate equality constraints on T and S. To be concrete, we allow for clauses of the type $(T = \gamma)$ or $(S = 10)$, but not for clauses of the type $(T + S = \sigma)$ or for inequality constraints of the type $(T \leq 1)$.[9] We label this class of agreements $\mathcal{A}_0$. In the following sections we will consider broader classes of agreements.

In order to characterize the optimal choice of agreement, we need to introduce some definitions and notation. First, we say that two agreements are *equivalent* if they implement the same outcome and have the same cost. Introducing a notion of agreement equivalence in this context is necessary because, as will become clear shortly, for any given agreement there may exist other agreements that implement the same outcome and have the same cost. Second, we refer to the *efficiently-written first-best* agreement as the least costly among the agreements that implement the first best outcome. We will often label this the "first-best" agreement, or simply the $\{FB\}$ agreement. In a similar vein, we will refer to the case of no agreement as the "empty agreement," which formally is denoted $\{\emptyset\}$. Finally, an *optimal agreement* is an agreement that maximizes expected global welfare net of contracting costs (henceforth simply "expected net global welfare"), that is $\Omega - C$.

Before we impose more structure on the nature of uncertainty, we present some results that hold quite generally. We start with an important observation: an agreement that restricts the level of S (even in a state-contingent way) while leaving T to the government's discretion can *never* be optimal.

To establish this, we now argue that an agreement that constrains only S cannot achieve higher gross global welfare than at the Nash equilibrium, for any state of the world. Recalling that $T^R(S)$ is the best-response level of T given S, the maximal gross global welfare that can be achieved by this type of agreement is given by the value of $W^G(T^R(S), S)$ evaluated at the optimal level of S,[10] with the optimal level of S defined in turn by the associated first-order condition $dW^G/dS = 0 \Longrightarrow T_S^R(S) = -W_S^G/W_T^G$, where subscripts denote derivatives. This first-order condition requires that the slope of $T^R(S)$ be equated with the slope of an iso-W^G curve in (T, S) space. It is direct to verify that the slope of $T^R(S)$ is

$$T_S^R(S) = -\frac{\lambda}{\beta + \lambda}.$$

[9] When there is significant uncertainty, a noncontingent contract of the type $g(\tau, t_f, t_h, s) = 0$ may do better than a noncontingent contract that pins down T and/or S separately, for two reasons. First, if g constrains the relationship between the components of T (for example as in $\tau = 1 - t_f^2$), it can mimic a weak binding ($T \leq \#$), or more generally a constraint that T must lie in a certain subset of the real line. In a later section we will consider weak bindings and show that they can improve over strong bindings, because allowing for downward discretion on T can be a good thing. Second, if g constrains the relationship between T and S, it may achieve a higher gross global surplus than a contract that pins down the exact levels of T and S, again because it introduces some discretion. In a later section we will consider an outcome-based type of contract (NV) that has a similar flavor as an instrument-based contract that constrains the relationship between T and S.

[10] Since we are focusing on a given state of the world, we do not have to make the state of the world explicit in the notation.

The slope of an iso-W^G curve in general is given by $-\frac{W_S^G}{W_T^G} = \frac{W_S + W_S^*}{W_T + W_T^*}$. However, since at the Nash equilibrium $W_S = W_T = 0$, the slope of the iso-W^G curve at the Nash equilibrium point is

$$-\frac{W_S^G}{W_T^G} = -\frac{W_S^*}{W_T^*} = -\frac{\frac{dW^*}{dp^*} \cdot \frac{dp^*}{dS}}{\frac{dW^*}{dp^*} \cdot \frac{dp^*}{dT}} = -\frac{\frac{dp^*}{dS}}{\frac{dp^*}{dT}} = -\frac{\lambda}{\beta + \lambda} \tag{3.1}$$

Therefore, the slope of the iso-W^G curve at the Nash point is equal to the slope of the $T^R(S)$ curve, and as a consequence, the level of S that maximizes $W^G(T^R(S), S)$ is the Nash equilibrium level S^{NE}. We may conclude, then, that an agreement that constrains only S cannot achieve greater surplus than $W^G(T^R(S^{NE}), S^{NE})$, which is just the Nash equilibrium surplus $W^G(T^{NE}, S^{NE})$. The next lemma records this result.[11]

Lemma 1. *An agreement that constrains the effective subsidy S while leaving the import tax T to discretion cannot improve over the Nash equilibrium, and therefore cannot be an optimal agreement.*

At a broad level, the intuition for this result is very simple. Contracting over S alone is useless because, as we have emphasized above, the inefficiency that arises in the noncooperative equilibrium concerns T, not S. We can also be a little more precise about the logic behind Lemma 1. The key steps of the argument are two. First, as equation 3.1 makes transparent, the slope of the iso-W^G curve at the Nash point is equal to the slope of the iso-p^* line. This is an immediate consequence of the fact that Home policies affect Foreign welfare only through terms of trade, and that at the Nash point small changes in home policies have no first-order effect on Home welfare. Second, the best-response import tax line $T^R(S)$ coincides with the iso-p^* line through the Nash point. The reason is that the Nash tariff is the one that implements the optimal (from the Home point of view) terms of trade, and constraining S away from its reaction curve triggers a change in T that brings the terms of trade back to its optimal level.[12]

We emphasize that, in a world of costless contracting, the result we have highlighted in Lemma 1 would be irrelevant, because if agreements are costless they would always be written in a way that placed constraints on all policy instruments that would be set inefficiently absent these constraints. But in a world of costly contracting, one has to consider agreements that place constraints on only a subset of these instruments, and the result of Lemma 1 then gains relevance. In particular, in such a world Lemma 1 implies that it can never be optimal to constrain internal measures but not import taxes.

We next turn to the characterization of the optimal agreement. A natural and convenient way to characterize the optimal agreement is to track how the optimal agreement changes as

[11] Notice that this result is distinct from and not contradictory to the result emphasized by Copeland (1990), that negotiating over tariffs can always generate surplus even if other instruments are non-negotiable. Copeland's result implies that the inclusion of tariffs in the set of instruments over which negotiations occur is *sufficient* for the possibility of gains from negotiations. The result we report in Lemma 1 implies that the inclusion of tariffs in the set of instruments over which negotiations occur is also *necessary* for the possibility of gains from negotiations.

[12] We note as well that the result reported in Lemma 1 is quite general, and in particular that it does not depend on the linearity assumptions of our model (a proof of this result for general non-linear demand and supply functions is available on request).

the level of elementary contracting costs increase. We consider a proportional increase in the two elementary contracting costs (c_p, c_s). To express our results in a simple comparative-statics fashion, we let $c_p \equiv c$, and $c_s \equiv k \cdot c$, where $k \geq 0$ captures the cost of contracting over a state variable relative to that of contracting over a policy variable, while c captures the general level of elementary contracting costs, which we henceforth refer to simply as "contracting costs." In much of the analysis to follow, we keep k fixed and consider changes in c. But our qualitative results would be the same if we allowed c_p and c_s to vary in a non-proportional way, as long as they co-vary. Note that, with this new notation, the total contracting cost can then be expressed as $C = (n_p + k \cdot n_s) \cdot c \equiv m \cdot c$. The parameter m can be seen as a rough measure of the "complexity" of the agreement, in the sense that it captures the number of policy variables and states involved (with the latter weighed by the parameter k).

Intuition suggests that there should be a monotonic relationship between the general level of contracting costs and the complexity of the optimal agreement (as measured by m). The next lemma confirms this intuition, showing that *more complex agreements will be chosen when the general level of contracting costs is lower.* We let $m(A)$ denote the complexity of agreement A, and $\Omega(A)$ the expected gross global welfare associated with agreement A.

Lemma 2. *Consider two non-equivalent agreements A',A'' in class $\mathcal{A}_0$. If A' is optimal for $c = c'$ and A'' is optimal for $c = c'' > c'$, then $m(A'') < m(A')$.*

Proof: Suppose by contradiction that $m(A'') > m(A')$. Since A'' is optimal for $c = c''$, then $\Omega(A'') - m(A'')c'' > \Omega(A') - m(A')c''$. But if $m(A'') > m(A')$ then this inequality holds *a fortiori* if we replace c'' with $c' < c''$, which implies that A'' is optimal also for $c = c'$, a contradiction.■

Next we present a result that provides conditions for an agreement A to be optimal for a range of contracting costs.

Lemma 3. *Consider the class of agreements $\mathcal{A}_0$. An agreement $\hat{A}$ (other than $\{FB\}$ and $\{\emptyset\}$) is optimal for some c if and only if it satisfies the following two conditions:*
(i) $\hat{A}$ is optimal in its complexity class, i.e. there is no agreement A' such that $m(A') = m(\hat{A})$ and $\Omega(A') > \Omega(\hat{A})$;
(ii) for any pair of agreements A', A'' such that $m(A') < m(\hat{A}) < m(A'')$,

$$\Omega(\hat{A}) \geq \frac{m'' - \hat{m}}{m'' - m'}\Omega(A') + \frac{\hat{m} - m'}{m'' - m'}\Omega(A''), \tag{3.2}$$

where $m' \equiv m(A')$, $m'' \equiv m(A'')$ and $\hat{m} \equiv m(\hat{A})$.

Proof: We first argue that conditions (i) and (ii) are both necessary, then we argue that the two conditions together are sufficient. The necessity of condition (i) is obvious. Consider condition (ii). For $\hat{A}$ to be an optimal agreement for some c there must be a value of c such that, for any agreement A' such that $m' < \hat{m}$,

$$\frac{\Omega(\hat{A}) - \Omega(A')}{\hat{m} - m'} \geq c \tag{3.3}$$

and for any agreement A'' such that $m'' > \hat{m}$,

$$\frac{\Omega(A'') - \Omega(\hat{A})}{m'' - \hat{m}} \leq c \tag{3.4}$$

and consequently that

$$\frac{\Omega(A'') - \Omega(\hat{A})}{m'' - \hat{m}} \leq \frac{\Omega(\hat{A}) - \Omega(A')}{\hat{m} - m'}$$

which can be rewritten as (3.2). Differently put, if (3.2) were violated, there would not be a value of c such that $\hat{A}$ simultaneously dominates A' and A'', contradicting the optimality of $\hat{A}$.

To see the sufficiency part, suppose that $\hat{A}$ is optimal in its complexity class and condition (3.2) holds for all agreements $A', A'' \in \mathcal{A}_0$ such that $m' < \hat{m} < m''$. Then there must be a value of c such that (3.3) holds for all agreements such that $m' < \hat{m}$, and (3.4) holds for all agreements such that $\hat{m} < m''$.■

Lemma 3 expresses necessary and sufficient conditions for an agreement to be optimal for a range of c, in terms of the properties of the gross global welfare function Ω. Condition (i) is obvious, as it requires that the candidate agreement not be dominated by some other agreement in its complexity class. Condition (ii) is a kind of concavity requirement on the Ω function. To understand in what sense this is a "concavity" condition, define $\hat{\Omega}(m)$ as the maximum level of Ω that can be attained with an agreement of complexity m – this is the level of gross global welfare as a reduced-form function of m. For an agreement $\hat{A}$ (with associated complexity level $\hat{m}$) to be optimal for some c, it must pass the following test: pick an arbitrary complexity level lower than $\hat{m}$ (call it m'), and one higher than $\hat{m}$ (call it m''); the function $\hat{\Omega}(m)$ must be concave with respect to the three points m', $\hat{m}$ and m''.

The economic interpretation of this "concavity" condition is, broadly speaking, that there must be *declining gains in gross welfare from adding complexity to the agreement.* If this were not the case, then if it paid to move from A' to the more complex $\hat{A}$ it would also pay to take the further step to the even more complex A'', in which case $\hat{A}$ would not be optimal for any c.

From an economic point of view, it may perhaps seem natural that there are diminishing gross returns from including additional variables in the agreement. However, as will be seen, this is typically *not* true in the contracting environment that we consider in this paper.

At this point we need to impose more structure on the stochastic environment, in order to derive sharper predictions on the nature of the optimal agreement.

3.1. Uncertainty about the consumption externality

In this section we assume that only γ is uncertain, and that it can take two possible values with equal probability: a high realization $\bar{\gamma} + \Delta_\gamma$ and a low realization $\bar{\gamma} - \Delta_\gamma$, with $\Delta_\gamma > 0$. To reflect the assumption that γ is the only uncertain parameter, we use a zero subscript for all the other parameters, to indicate that they are deterministic.

The first step is to derive the $\{FB\}$ agreement. It is clear that the agreement $\{T = \gamma; S = \sigma_0 - \gamma\}$ implements the first best outcome. This agreement costs $(4 + k) \cdot c$. But it might be conjectured that the first best outcome could also be implemented without constraining S, and therefore be accomplished more cheaply, since as we have noted previously only T is set inefficiently in the Nash equilibrium. This conjecture is incorrect, but it is instructive to see why. The reason is simply that an agreement that dictates an import tax level T but leaves discretion over S would permit the home government to choose S according to the best-response function $S^R(T)$. As we have observed previously, this best response is equal to S^{eff} when the home government is on its import–tax reaction curve, but more generally is it direct to verify that

$$S^R(T) = S^{eff} - \frac{\Upsilon^2 - (\beta+\lambda)^2}{\Upsilon^2 - \lambda(\beta+\lambda)}(T - T^{NE}). \tag{3.5}$$

From this equation it follows that the difference between $S^R(T)$ and S^{eff} is directly proportional to the distance between T and the home government's best-response import tax T^{NE}.[13] As a consequence, an agreement that attempts to move T towards its efficient level without also constraining S will cause S to become distorted for terms-of-trade purposes.

In fact, it is easy to verify that one cannot implement the first best outcome with an agreement in $\mathcal{A}_0$ that costs less than $(4+k)\cdot c$. We can conclude that $\{T=\gamma; S=\sigma_0-\gamma\}$ is indeed the efficiently-written first best agreement. Note that both S and T are state-contingent in the $\{FB\}$ agreement in this environment.

Next we look for the optimal agreement, that is the agreement that maximizes expected net global welfare in $\mathcal{A}_0$.

The $\{FB\}$ agreement yields expected net global welfare equal to $\Omega(T=\gamma, S=\sigma_0-\gamma) - (4+k)\cdot c$. When contracting costs are zero, the $\{FB\}$ agreement is of course optimal. But for sufficiently high contracting costs the $\{FB\}$ agreement cannot be optimal, because as c rises eventually even the empty agreement (which costs nothing and yields expected global welfare $\Omega(T=T^{NE}, S=\sigma_0-\gamma)$) will yield higher expected net global welfare. The basic question is then whether other (incomplete) agreements will become optimal as contracting costs rise from zero, reflecting an attractive trade-off between the surrender of expected gross global welfare for a reduction in the overall costs of the agreement, and if so how the optimal agreement is shaped by this trade-off.

Even though uncertainty is one-dimensional, the optimal restructuring of an agreement to economize on overall contracting costs can in principle be quite onerous. This difficulty is exacerbated by the inherent non-separability of the contracting problem that we study, a feature that seems an unavoidable part of the problem faced by trade negotiators but which precludes characterizing the optimal contract "task by task" as in Battigalli and Maggi (2002). However, knowledge of the $\{FB\}$ agreement, together with the perfect policy substitutability between τ and t_f and between s and t_h, simplifies the problem considerably. Specifically, between the $\{FB\}$ agreement which costs $(4+k)\cdot c$ and the empty agreement which costs nothing, there are three cost classes of agreements that warrant consideration: agreements costing $(2+k)\cdot c$; agreements costing $4\cdot c$; and agreements costing $2\cdot c$.

We can now use the result of Lemma 1 to further simplify the problem of deriving the optimal agreement. Since we can ignore agreements that constrain S but not T, we only have three kinds of agreements to consider, in addition to $\{FB\}$ and $\{\emptyset\}$: agreements that

[13] Since the foriegn government is assumed to have no (export-sector) policy options, the home government's best-response import tax and the Nash equilibrium import tax T^{NE} are in fact one and the same.

constrain T as a function of γ, which we denote $\{T(\gamma)\}$; agreements that constrain T and S in a non-state-contingent fashion, which we denote $\{T, S\}$; and agreements that constrain T in a non-state-contingent fashion, which we denote $\{T\}$.

The three types of agreements $\{T, S\}$, $\{T(\gamma)\}$ and $\{T\}$ are all incomplete, but they are each incomplete in a different way. To describe these differences, it is useful at this point to recall the distinction, introduced by Battigalli and Maggi (2002), between two forms of contractual incompleteness: *rigidity*, which occurs when contractual obligations do not include state contingencies that are included in the first-best agreement; and *discretion*, which occurs when contractual obligations that are included in the first-best agreement are simply missing. We can thus say that the agreement $\{T, S\}$ is rigid, the agreement $\{T(\gamma)\}$ features discretion (over S), and the agreement $\{T\}$ is both rigid and discretionary.

From this vantage point, the key question we wish to answer may therefore be phrased as follows: As contracting costs rise from zero, what is the optimal way of restructuring the agreement to economize on contracting costs?; Making the agreement rigid, discretionary, or both?; And in what sequence?

In Battigalli and Maggi's (2002) model, as contracting costs rise, first the optimal contract becomes rigid, and then discretion is introduced. One might conjecture that this is true also in the present setting, but this conjecture turns out to be incorrect, as we will argue below. In fact, due to the non-separable nature of the contracting problem we are considering here, it may be optimal to economize on contracting costs by introducing discretion before rigidity, and it may even be the case that rigidity is not optimal for any level of contracting costs.

In order to examine the optimal structure of the agreement as a function of contracting costs, it proves useful to introduce some new concepts. It is natural to define the *cost of rigidity* as the loss of gross global welfare when the state variable γ is excluded from the agreement, and the *cost of discretion* as the loss of gross global welfare when the policy variable S is excluded from the agreement. For the moment we focus on the non-empty agreements, so we can ignore the case in which both T and S are discretionary.

It is important to observe that rigidity and discretion *interact* in non-trivial ways: the cost of rigidity depends on whether or not discretion is present in the agreement, and the cost of discretion depends on whether or not the agreement is rigid. To describe this interaction, we define a zero-one "rigidity" variable R that takes a value of one if the agreement is rigid and zero otherwise, and a zero-one "discretion" variable D that takes a value of one if the effective subsidy S is discretionary and zero otherwise. Next we define $\hat{\Omega}(R, D)$ as the level of gross welfare expressed as a reduced-form function of the degree of rigidity and discretion. We can now introduce the following concepts:

a. The cost of discretion absent rigidity (CD): $\hat{\Omega}(0,0) - \hat{\Omega}(0,1) = \Omega_{\{FB\}} - \Omega_{\{T(\gamma)\}}$;

b. The cost of discretion in the presence of rigidity ($\overline{CD}$): $\hat{\Omega}(1,0) - \hat{\Omega}(1,1) = \Omega_{\{T,S\}} - \Omega_{\{T\}}$;

c. The cost of rigidity absent discretion ($CRTS$): $\hat{\Omega}(0,0) - \hat{\Omega}(1,0) = \Omega_{\{FB\}} - \Omega_{\{T,S\}}$; and

d. The cost of rigidity in the presence of discretion (CRT): $\hat{\Omega}(0,1) - \hat{\Omega}(1,1) = \Omega_{\{T(\gamma)\}} - \Omega_{\{T\}}$.

Notice that these four quantities are linearly dependent, since one can always be expressed as a linear combination of the other three. Finally, it will also prove useful to introduce a last concept that captures the cost of discretion over both T and S, namely, the potential gain from contracting (PGC): $\Omega_{\{FB\}} - \Omega_{\{\emptyset\}}$.

An important manifestation of the interaction between rigidity and discretion is the following: while the cost of discretion absent rigidity, CD, is always positive, the cost of discretion in the presence of rigidity, $\overline{CD}$, *may be negative*. In other words, it is possible that $\Omega_{\{T,S\}}$ is lower than $\Omega_{\{T\}}$: conditional on the agreement being rigid, introducing discretion may increase gross global welfare. Intuitively, introducing discretion (in S) is a way of introducing state-contingency in the agreement, and this beneficial effect may outweigh the negative effect of allowing a government to use S to manipulate the terms of trade. This suggests that *rigidity and discretion are complementary*, in the sense that the presence of rigidity decreases the cost of discretion, possibly making it negative. Indeed, one can show that $\overline{CD} < CD$ for all parameter values. Note also that $\overline{CD} < CD$ immediately implies $CRTS > CRT$. This is the flip-side of the complementarity we just highlighted: the presence of discretion decreases the cost of rigidity.[14]

The complementarity between rigidity and discretion plays an important role in the remaining analysis of this section, but it should be mentioned here that this complementarity depends crucially on the fact that uncertainty concerns a state variable that is directly relevant for the first-best policy levels (in the specific case, γ). As we will explain in the next section, when uncertainty concerns a state variable that is *not* directly relevant for the first-best policy levels (e.g. the import demand level α), rigidity and discretion are substitutable, not complementary. Thus it should be kept in mind that whether rigidity and discretion are complementary or substitutable depends critically on the nature of uncertainty.

We are now ready to return to the question we posed above: What is the optimal sequence of agreements as c increases from zero? This question can be answered by using Lemmas 2 and 3, together with the complementary property that we just highlighted.

Lemma 2 tells us that, as c increases, we must move from more complex to less complex agreements. This immediately implies that the optimal sequence of agreements is a subsequence of $(\{FB\}, \{T,S\}, \{T(\gamma)\}, \{T\}, \{\emptyset\})$. We say "subsequence" because each of these agreements (except $\{FB\}$ and $\{\emptyset\}$) may be "skipped" over as c increases.[15]

Next we use Lemma 3 together with the complementarity property to establish a key result: the agreements $\{T,S\}$ and $\{T(\gamma)\}$ cannot *both* be part of the optimal sequence of agreements.

Let us focus on the case $k = 1$; the extension of the argument to $k \neq 1$ is straightforward. As Lemma 3 establishes, a necessary condition for $\{T,S\}$ to be optimal for some c is that there are diminishing returns in Ω as we move from $\{T\}$ to $\{T,S\}$ to $\{FB\}$. Using 3.2, this requires:

$$2(\Omega_{\{FB\}} - \Omega_{\{T,S\}}) \leq \Omega_{\{T,S\}} - \Omega_{\{T\}}$$

or, using the definitions above,

$$2CRTS \leq \overline{CD} \tag{3.6}$$

[14]In section 6 we highlight a second sense in which rigidity and discretion may be complementary: there we show that, conditional on the agreement being rigid, it may be a good idea to impose an upper bound on the import tax T (weak binding), rather than an equality constraint. In other words, if the agreement is to be rigid, it may be valuable to allow not only for discretion over S, but also for some *downward* discretion on T.

[15]In the above statement we implicitly assumed that $k \leq 2$, so that $\{T(\gamma)\}$ is not more costly than $\{T,S\}$. But the statement is true even if $k > 2$, because as we establish below, $\{T(\gamma)\}$ and $\{T,S\}$ cannot both be part of the optimal sequence.

Similarly, a necessary condition for $\{T(\gamma)\}$ to be optimal for some c is that there are diminishing returns in Ω as we move from $\{T\}$ to $\{T(\gamma)\}$ to $\{FB\}$. Using 3.2, this requires:

$$\Omega_{\{FB\}} - \Omega_{\{T(\gamma)\}} \leq 2(\Omega_{\{T(\gamma)\}} - \Omega_{\{T\}})$$

or

$$CD \leq 2CRT \tag{3.7}$$

Now recall that $\overline{CD} < CD$ and $CRTS > CRT$ (complementarity between rigidity and discretion). Clearly, this implies that 3.6 and 3.7 cannot both be satisfied: if there are diminishing returns to complexity around $\{T, S\}$ there cannot be diminishing returns to complexity around $\{T(\gamma)\}$, and vice-versa. In other words, because of the complementarity between rigidity and discretion, returns to complexity must be *increasing* around (at least) one of these two agreements. As a consequence, $\{T, S\}$ and $\{T(\gamma)\}$ cannot both be part of the optimal sequence of agreements.

We summarize these results in the following proposition:

Proposition 1. *Consider the agreement class $\mathcal{A}_0$, and assume that only γ is uncertain. There exist scalars c_1, c_2, c_3 and c_4 with $0 < c_1 \leq c_2 \leq c_3 \leq c_4 < \infty$ such that the optimal agreement is:*
(a) the $\{FB\}$ agreement for $c \in (0, c_1)$;
(b) of the form $\{T, S\}$ for $c \in (c_1, c_2)$;
(c) of the form $\{T(\gamma)\}$ for $c \in (c_2, c_3)$;
(d) of the form $\{T\}$ for $c \in (c_3, c_4)$; and
(e) the empty agreement for $c > c_4$.
The critical levels of c satisfy: $(c_2 - c_1)(c_3 - c_2) = 0$.

A first important aspect of Proposition 1 concerns the way contractual incompleteness varies across policy instruments. According to Proposition 1 the effective subsidy S tends to be more discretionary than T; more specifically, for a range of sufficiently high contracting costs it may be optimal to contract over import taxes T while leaving the effective production subsidy S to discretion, but as Lemma 1 indicates and Proposition 1 confirms it is *never* optimal to leave T to discretion and contract only over S.

In this way, Proposition 1 provides a general prediction that trade agreements should always include commitments over import taxes, and should only introduce commitments over other internal measures as the agreement becomes more complete. This prediction resonates to some degree with the broad approach taken by the GATT/WTO, which has been to first establish a base of commitments over import tax levels, and only later to broaden the agreement to explicitly take on various internal measures. Notice, too, that our prediction does not rely on an assumption that embodies the commonly-held view that border measures are more transparent than internal measures and are therefore less costly to contract over, an assumption that would only reinforce this prediction.

A second important aspect of Proposition 1 is the "complementary slackness" condition $(c_2 - c_1)(c_3 - c_2) = 0$: if the discretionary contract $\{T(\gamma)\}$ is optimal for a range of contracting costs (i.e., if $c_2 < c_3$), then the rigid contract $\{T, S\}$ is not optimal for any cost level (i.e.,

$c_1 = c_2$), and vice-versa. In light of this condition, the next question that needs to be addressed is the following: As contracting costs rise from zero, when is it optimal to first economize on contracting costs by introducing discretion ($c_2 < c_3$), and when by first introducing rigidity ($c_1 < c_2$)?

To answer this question, we start by developing some intuition. We ask: How do the costs of rigidity and discretion ($CRT, CRTS, CD, \overline{CD}$) depend on the fundamentals of the contracting problem? Our attention will be focused primarily on two parameters: Δ_γ, which captures the degree of uncertainty in the environment, and α, which can be interpreted as the level of import demand. We will at times highlight also the effect of changes in the Home demand and supply slope parameters, β and λ, as a way to illustrate the role – explained below – of the degree of substitutability between policy instruments.

Let us start with the cost of rigidity. Intuitively, the cost of rigidity should be increasing in the degree of uncertainty, Δ_γ. A rigid agreement "gets it right" only on average, and therefore when the environment is more uncertain the cost of rigidity should be higher. Indeed, it is easy to verify analytically that both CRT and $CRTS$ are increasing in Δ_γ. It can also be verified that CRT and $CRTS$ are independent of the import demand level α.

The cost of discretion (over S) depends on the economic environment in more subtle ways. Let us focus first on the cost of discretion absent rigidity, CD. Intuitively, there are two features of the contracting environment that affect CD.

First, CD is higher when the effective subsidy S is a closer substitute for the import tax T. Recall that the gains from contracting in this setting derive entirely from the ability to control the incentive to manipulate the terms of trade. It follows, then, that when S and T are very similar instruments with regard to their effect on the terms of trade, the cost of discretion over S (permitting S to be used for terms-of-trade manipulation) is high. Thus a key determinant of the cost of discretion is *the degree of substitutability between policy instruments*.

Intuition also suggests that a second feature of the contracting environment that affects CD is the wedge between the Nash and efficient tariff, which in turn reflects the degree of home-country *monopoly power* as measured by the (inverse of the) magnitude of the foreign export supply elasticity η^*. For a given level of instrument substitutability between T and S, if the degree of home-country monopoly power is high (η^* is low), then the wedge between the Nash and efficient tariff is large; therefore, the incentive to alter policies to manipulate the terms of trade is high, and hence, controlling this incentive with an agreement over T but leaving S to discretion will result in costly distortions in S for purposes of manipulating the terms of trade (as indicated in 3.5).

How do these two features of the contracting environment depend on the underlying parameters of the model? Let us focus first on the degree of substitutability between T and S. Recalling that the terms of trade is given by $p^*(T, S) = [\alpha + \alpha^* - (\beta + \lambda)T - \lambda S]/\Upsilon$, a rough measure of this substitutability is given by the marginal rate of substitution between T and S with respect to the terms of trade: $\frac{\partial p^*/\partial T}{\partial p^*/\partial S} = \frac{\beta}{\lambda} + 1$. This suggests that S is a closer substitute for T when demand is less elastic (β is low) and when supply is more elastic (λ is high), and that the other parameters of the model – including the import demand level α – have no impact on the degree of instrument substitutability.[16,17]

As far as the home-country monopoly power is concerned, this is directly linked to the volume of trade, and therefore a primary determinant of its magnitude is the level of import demand, α. When α is higher, home-country monopoly power is higher (η^* is lower), and therefore CD should be higher.

The intuition we developed above about the determinants of CD can be confirmed analytically: as expected, CD is higher when the import demand level α is higher, because α affects CD only through the home-country monopoly power channel. The impacts of changes in the demand slope β and the supply slope λ on CD are more complicated, since β and λ affect CD through both the instrument substitutability and the monopoly power channels. Nevertheless, it can be confirmed that the instrument substitutability channel dominates for β and λ sufficiently low. Specifically, CD approaches its minimum value of zero as λ approaches zero. And CD approaches its maximum value of PGC (so that it becomes worthless to contract over T while leaving S to discretion) as β approaches zero.[18]

What can we say about the cost of discretion in the presence of rigidity, $\overline{CD}$? Intuitively, the determinants of $\overline{CD}$ should be similar to those of CD, but in addition to those, uncertainty should now play an important role. Recall from our discussion of the complementarity between rigidity and discretion that $\overline{CD}$ is lower than CD because, when the level of T is rigid, giving discretion over S is a way to introduce state contingency in the agreement. When uncertainty (Δ_γ) is larger this effect is more important.

Having discussed at a broad level the determinants of the costs of rigidity and discretion, we can now come back to the question of which of the two agreements $\{T, S\}$ or $\{T(\gamma)\}$ (if any) is optimal for a range of contracting costs. Recall that a necessary condition for $\{T, S\}$ to be

---

[16] The ratio $\frac{\partial p^*/\partial T}{\partial p^*/\partial S}$ is not a precise measure of instrument substitutability, because it does not take into account the welfare effects of changes in terms of trade. A more accurate measure of instrument substitutability can be constructed as follows. Recall that, if the agreement dictates an import tax level T but leaves discretion over S, the home government will choose S according to the best-response function $S^R(T)$, and therefore the associated level of gross global welfare will be $\Omega(T, S^R(T))$. In this case, then, a change in the agreed-upon level of T has two effects on Ω: a direct effect, captured by the partial derivative $\partial\Omega/\partial T$, and an indirect effect through the induced change in S, captured by $(\partial\Omega/\partial S)(dS^R/dT)$. The total effect of a change in T is given by the total derivative $d\Omega/dT = \partial\Omega/\partial T + (\partial\Omega/\partial S)(dS^R/dT)$. When T and S are close substitutes, the impact on Ω of a change in T will be largely offset by the induced change in S, and hence $d\Omega/dT$ will be close to zero. On the other hand, if T and S are poor substitutes, the indirect effect through S will be small and therefore $d\Omega/dT$ will be close to $\partial\Omega/\partial S$. This suggests using the ratio $\frac{d\Omega/dT}{\partial\Omega/\partial T}$ as a measure of instrument substitutability: when T and S are close substitutes this ratio is close to zero, and when T and S are poor substitutes it is close to one. It is natural to evaluate this ratio at the Nash equilibrium policies, (T^{NE}, S^{NE}), yielding $\left(\frac{d\Omega(T,S^R(T))/dT}{\partial\Omega(T,S^R(T))/\partial T}\right)_{T=T^{NE}} = \frac{\beta\Upsilon^2}{(\beta+\lambda)(\Upsilon^2-\lambda)}$. It can be shown that this ratio is increasing in β and decreasing in λ. This confirms the insight we developed in the text using the simpler measure of instrument substitutability.

[17] The result we reported in Lemma 1 can be interpreted as well through the lens of instrument substitutability. In particular, Lemma 1 indicates that, beginning from the Nash equilibrium, T is a (locally) perfect substitute for S, in the sense that the impact on Ω of a small change in S will be completely offset by the best-response change in T, and as Lemma 1 implies, this is true for all parameters of the model. Notice, too, that instrument substitutability as we define it is not symmetric with regard to instruments (nor should it be).

[18] These properties of CD with respect to β and λ can be confirmed with reference to the measure of instrument substitutability developed in footnote 16.

optimal for some c is $2CRTS < \overline{CD}$ (condition 3.6), and a necessary condition for $\{T(\gamma)\}$ to be optimal for some c is $CD < 2CRT$ (condition 3.7). Thus, according to the discussion above, we expect $\{T, S\}$ to be favored over $\{T(\gamma)\}$ when the level of demand α is high (so that CD and $\overline{CD}$ are high), when uncertainty Δ_γ is low (so that $CRTS$ and CRT are low), and when the demand slope β is low and the supply slope λ is high (so that CD and $\overline{CD}$ are high).

Of course, the arguments we just made are not the whole story. We know from Lemma 3 that conditions 3.6 and 3.7 are necessary but not sufficient for the optimality of (respectively) $\{T, S\}$ and $\{T(\gamma)\}$. But as the following remark shows, the intuition we developed based on those necessary conditions is in good part borne out:

Remark 1. *(i) If the import demand level α is sufficiently high and/or the degree of uncertainty Δ_γ is sufficiently low, then $c_1 < c_2 = c_3$: as contracting costs rise from zero, it is optimal to first introduce rigidity into the agreement ($\{T, S\}$), and possibly later to introduce rigidity and discretion ($\{T\}$), but it is never optimal to introduce discretion alone ($\{T(\gamma)\}$).*
(ii) If the supply slope λ is sufficiently low, then $c_1 = c_2 < c_3$: as contracting costs rise from zero, it is optimal to first introduce discretion into the agreement ($\{T(\gamma)\}$), and possibly later to introduce discretion and rigidity ($\{T\}$), but it is never optimal to introduce rigidity alone ($\{T, S\}$).
(iii) If the demand slope β is sufficiently low, then $c_2 = c_3 = c_4$: as contracting costs rise from zero, it may be optimal to introduce rigidity into the agreement ($\{T, S\}$), but it is never optimal to introduce discretion, either alone ($\{T(\gamma)\}$) or with rigidity ($\{T\}$).

The results summarized in Remark 1 stand in marked contrast to those reported in Battigalli and Maggi (2002). As we indicated above, the contracting problem that they study is separable, and so the contract can be optimized task by task: there, as contracting costs rise, first the optimal contract becomes rigid, and then discretion is introduced. By contrast, the contracting problem faced by trade negotiators is inherently non-separable, because there exist internal policy measures (S) that can at least imperfectly substitute for import taxes (T): trade negotiators must then consider the possible ramifications for the optimal treatment of *all* policy measures when they consider introducing rigidity or discretion into the treatment of any one policy measure, and this consideration leads to the richer possibilities described in Remark 1.

We argued above that rigidity and discretion may be complementary, but we have not yet established whether it can indeed be optimal to combine rigidity and discretion, or in other words, whether it may be optimal to write an agreement of the form $\{T\}$. It is not hard to find sufficient conditions such that $\{T\}$ is optimal for a range of c, i.e. such that $c_3 < c_4$. For example, one simple sufficient condition is that uncertainty (Δ_γ) is sufficiently low *and* policy instruments are not very substitutable (λ is low or β is high). To see this, notice that when Δ_γ gets close to zero, any contingent agreement becomes dominated, leaving only $\{T\}$, $\{T, S\}$ and the empty agreement as candidates for an optimum; and if policy instruments are very dissimilar $\{T\}$ dominates $\{T, S\}$.

Remark 1 highlights how various parameters affect the optimal sequence of agreements as c varies, but it does *not* examine the comparative-static impact of changes in those parameters. For example, Remark 1 does not tell us how the optimal agreement changes as we increase α while keeping all other parameters (including c) constant, which is an interesting question in its own right. In what follows we examine the comparative-static effects of the two parameters we are mostly interested in, that is α and Δ_γ.

We start with changes in the import demand level, α. Since an increase in α increases the volume of trade, it also increases the potential gains from contracting, PGC. Thus one might conjecture that increasing α has similar effects as decreasing the contracting cost c, making the agreement more complete, both in the sense of reducing discretion and reducing rigidity.[19] This conjecture however is only partially correct: an increase in α does reduce discretion, but it does not reduce the degree of rigidity; more specifically, as α increases (it can be shown that) the optimal agreement never switches from $\{T\}$ to $\{T(\gamma)\}$ or from $\{T, S\}$ to $\{FB\}$. This is because α does not affect the cost of rigidity (CRT or $CRTS$), and hence it does not affect the trade-off between $\{T\}$ and $\{T(\gamma)\}$ or between $\{T, S\}$ and $\{FB\}$. Thus all we can say about the effects of an increase in α is the following:

Proposition 2. *As the import demand level α increases (holding all other parameters fixed), the optimal degree of discretion decreases, in the sense that the number of policy instruments specified in the optimal agreement increases (weakly).*

The comparative-statics result of the above proposition can be interpreted as reflecting the impact that the monopoly power effect has on the desirability of discretion, combined with the implication of Lemma 1. In particular, the higher is the degree of home-country monopoly power (the higher is α), the less desirable it is to leave policy instruments to discretion, with the order in which instruments are tied down (first T and then also S) dictated by Lemma 1.

More broadly, it is noteworthy to observe at this point that both of the factors which we have identified as increasing the cost of discretion and hence favoring the agreement $\{T, S\}$ over $\{T(\gamma)\}$ and $\{T\}$ – the instrument-substitutability effect and the monopoly power effect – could arguably be seen to apply most aptly to relatively large developed countries. The essence of the instrument-substitutability effect is that a government has access to a rich array of internal measures which it can use to carry on terms-of-trade manipulation should its import taxes be constrained through an international trade agreement. And the essence of the monopoly power effect is that a country is large in world markets, so that it faces foreign export supply that is far from perfectly elastic. Arguably, both of these features are more likely to be true of large developed countries. When viewed in light of our discussion above, this observation is potentially interesting, because it suggests that the attractiveness of contracting over internal measures (such as those embodied in S in our formal model) may be different for large developed countries than for small developing countries, and in particular that optimal contractual design might entail small developing countries negotiating commitments of the form $\{T(\gamma)\}$ or $\{T\}$ while large developed countries negotiate commitments of the form $\{T, S\}$. While our two-country model cannot really address this issue formally, our results are at least suggestive of the possible benefits of a kind of "special and differential treatment" rule for small/developing countries when it comes to contracting over internal measures (such as subsidies).[20]

[19] In Battigalli and Maggi (2002), for example, the elementary cost of contracting and the potential gains from contracting matter only through their ratio, and hence reducing the former has exactly the same effects as increasing the latter.

[20] We thank Robert Lawrence for first bringing this implication to our attention.

Next we consider the comparative-statics effects of changes in the degree of uncertainty, Δ_γ. It is straightforward to establish the following result:

Proposition 3. *As the degree of uncertainty Δ_γ increases (holding all other parameters fixed), the optimal agreement may switch from a rigid agreement to a contingent agreement, but not vice-versa.*

By itself, the result reported in the above proposition is not particularly surprising: it seems inevitable that increasing uncertainty should reduce the attractiveness of rigid agreements, just as the proposition reports. But there is also a more subtle feature of this result, which is that it concerns uncertainty over a state variable that is directly relevant for the setting of T in the $\{FB\}$ agreement. As we demonstrate in the next sections, the effects of increasing uncertainty over state variables (such as α and σ) that are not directly relevant for the setting of T in the $\{FB\}$ can be very different.

3.2. Uncertainty about the level of import demand

In the previous section we examined a simple stochastic environment where uncertainty concerns only a state variable that affects directly the first-best levels of T and S. Here we explore the implications of a different type of uncertainty, which in some sense is at the opposite extreme as the case considered previously: we now suppose that uncertainty concerns only a state variable – namely the import demand level α – that has no impact on the first-best levels of T or S. Focusing separately on these two extreme cases is instructive, because it highlights how the rigidity-discretion trade-off and the structure of the optimal agreement depend crucially on the nature of the uncertainty in the contracting environment.

We assume that α can take two possible values with equal probability: $\bar{\alpha} + \Delta_\alpha$ and $\bar{\alpha} - \Delta_\alpha$, with $\Delta_\alpha > 0$. In this environment, the $\{FB\}$ agreement takes the form $\{T = \gamma_0; S = \sigma_0 - \gamma_0\}$. Notice that the $\{FB\}$ agreement is no longer state contingent, because it does not depend on the uncertain parameter α. This property reflects our assumption that the distributional impacts of local price movements are unimportant to governments, implying that globally optimal policy intervention exists only to correct externalities (σ and γ).[21]

Also notice that, as an immediate implication of Lemma 1, there are now four types of agreements that can potentially be optimal: (i) the $\{FB\}$ agreement, which is of the type $\{T, S\}$; (ii) agreements of the form $\{T(\alpha)\}$; (iii) agreements of the form $\{T\}$; and (iv) the empty agreement.

Two important new insights emerge in this environment. The first one is the possibility of the agreement $\{T(\alpha)\}$, where $T(\alpha)$ is an increasing function. This has the flavor of an escape-clause type of agreement: when α is high, the underlying import volume is high, and so with $T(\alpha)$ an increasing function the agreement $\{T(\alpha)\}$ allows for the import tariff to rise in states

[21] In particular, these government objectives rule out the possibility of political economy motives. In the presence of such motives, the $\{FB\}$ contract would in general depend on all (eight) state variables in the model, and the characterization of the optimal agreement would as a consequence be considerably more complex. We leave the introduction of political economy motives into our incomplete contracting setting for future work.

of the world in which the underlying import volume is high, broadly analogous to the escape clause provided in GATT Article XIX.

The potential appeal of an escape clause is not obvious in the current setting, since our formal model rules out concerns about income distribution or enforcement.[22] But its appeal can be understood with reference to the monopoly power effect. In particular, recall that if S is left to discretion (as in the agreement $\{T(\alpha)\}$) then it will be used to manipulate the terms of trade, and the incentive to do so will be stronger when α is higher according to the home-country monopoly power effect. A higher T mitigates the incentive to distort S for terms-of-trade purposes, and so allowing for a higher T when α is higher can therefore help to mitigate the use of S for purposes of terms-of-trade manipulation when the incentive is highest to do so. In this way, our model identifies a novel rationale for the desirability of escape clauses in trade agreements: an escape-clause type agreement of the form $\{T(\alpha)\}$ can be attractive relative to a rigid agreement of the form $\{T\}$ because it provides a means of managing the distortions associated with leaving S to discretion.

The second new insight is that rigidity and discretion are no longer complementary, but are instead *substitutable*. Formally, it is easy to see that the cost of discretion in the presence of rigidity, $\overline{CD} = \Omega_{\{T,S\}} - \Omega_{\{T\}}$, is higher than the cost of discretion absent rigidity, $CD = \Omega_{\{FB\}} - \Omega_{\{T(\alpha)\}}$. This is because the $\{FB\}$ agreement is non-contingent, so $\Omega_{\{FB\}} = \Omega_{\{T,S\}}$, and $\Omega_{\{T(\alpha)\}} > \Omega_{\{T\}}$. The substitutability between rigidity and discretion can also be seen from the perspective of the costs of rigidity, $CRTS$ and CRT. Clearly in this setting $CRTS = 0$, because the first-best agreement is non-contingent, and hence $CRTS < CRT$. Recall that this condition is equivalent to the condition $\overline{CD} > CD$.

Intuitively, here the presence of rigidity does not confer any extra value to discretion: introducing discretion in a rigid contract does not generate any benefits in terms of improved state-contingency. To the contrary, now the cost of discretion is lower when the agreement is *not* rigid, and the reason is that the adverse affects of discretion over S can be mitigated – as we explained above – by making the value of T contingent on the state; this mitigation of the cost of discretion is not possible within a rigid contract.

These observations, together with those made in the previous section, suggest a very important insight: the interaction between rigidity and discretion depends crucially on the exact nature of the uncertainty in the contracting environment. When uncertainty concerns variables (such as γ) that are directly relevant for the first-best levels of T and S, rigidity and discretion tend to be complementary, but when uncertainty concerns variables (such as α) that are *not* directly relevant for the first-best levels of T and S, then rigidity and discretion tend to be substitutable.

One consequence of the substitutability between rigidity and discretion in the present stochastic environment is that the "concavity" condition of Lemma 3 now may be satisfied for all relevant complexity levels, and therefore we do not have a "complementary slackness" condition as in the previous section: all four candidate agreements may be part of the optimal sequence as c increases. The following proposition confirms this point:

[22] See footnote 4.

Proposition 4. *Consider the agreement class $\mathcal{A}_0$, and assume that only α is uncertain. There exist scalars c_1, c_2, and c_3 with $0 < c_1 \leq c_2 \leq c_3 < \infty$ such that the optimal agreement is:*
(a) the $\{FB\}$ agreement $\{T = \gamma_0; S = \sigma_0 - \gamma_0\}$ for $c \in (0, c_1)$;
(b) of the form $\{T(\alpha)\}$ for $c \in (c_1, c_2)$;
(c) of the form $\{T\}$ for $c \in (c_2, c_3)$; and
(d) the empty agreement for $c > c_3$.

Since a key new insight in this environment is the possibility of the $\{T(\alpha)\}$ agreement, the next question we want to address is, Under what conditions (if any) is the agreement $\{T(\alpha)\}$ optimal? The following remark identifies conditions on model parameters under which $\{T(\alpha)\}$ is optimal for a range of c:

Remark 2. *(i) If k is sufficiently low* **and** *λ is sufficiently low, then $c_1 < c_2$: an escape-clause-type agreement of the form $\{T(\alpha)\}$ is optimal for some c.*
(ii) If $\bar{\alpha}$ is sufficiently high, or if Δ_α is sufficiently low, or if β is sufficiently low, then $c_1 = c_2$: an agreement of the form $\{T(\alpha)\}$ is not optimal for any c.

The results reported in Remark 2 are intuitive in light of our discussion just above and the interplay between the instrument-substitutability and the monopoly power effect. In particular, if the degree of substitutability between T and S is sufficiently low (as when λ is low) so that leaving S to discretion is an attractive option, then an escape-clause type agreement of the form $\{T(\alpha)\}$ will be optimal for a range of c as long as the relative cost of contracting over state variables (k) is sufficiently low. As described above, the appeal of this type of agreement is that it provides a means of managing the distortions associated with leaving S to discretion. On the other hand, if the degree of substitutability between T and S is sufficiently high (as when β is low) and/or the degree of home country monopoly power is expected to be high (as when α is expected to be high) so that leaving S to discretion is not an attractive option, or if α is not very uncertain, then an escape-clause type agreement will not be optimal for any range of c.

Finally, we note that the comparative-static effects of changes in the expected demand level $\bar{\alpha}$ can be shown to be similar to those described in Proposition 2 of the previous section for the case of changes in α in the presence of γ uncertainty. On the other hand, the effects of changes in the degree of uncertainty over α (Δ_α) differ in an interesting way from the effects of changes in the degree of uncertainty over γ (Δ_γ) as the latter were reported in Proposition 3 of the previous section. To illustrate this difference, note that as uncertainty over α increases, the optimal agreement may switch from a contingent type $\{T(\alpha)\}$ to a rigid type $\{T, S\}$, which as Proposition 3 indicates can never happen with an increase in uncertainty over γ (Δ_γ). Intuitively, this reflects the workings of the monopoly power effect, and the fact that the cost of discretion (CD) is not only rising in α but also convex: hence, as uncertainty over α (Δ_α) rises, CD rises, and it may be optimal to move from a contingent agreement with discretion to a rigid agreement without discretion. This confirms our earlier observation that the effects of uncertainty are quite different depending on the exact nature of the uncertainty, and in particular on whether uncertainty concerns state variables that are directly relevant to the setting of T in the $\{FB\}$ agreement (as in Δ_γ) or rather state variables that are not directly relevant to the setting of T in the $\{FB\}$ agreement (as in Δ_α).

3.3. Uncertainty about the production externality

In the previous two sections we examined two rather extreme stochastic environments: one where uncertainty concerns only a state variable (γ) that is directly relevant to the first-best levels of both T and S, and one where uncertainty concerns only a state variable (α) that is not relevant to the first-best level of either T or S. In this section we consider an alternative stochastic environment that in some sense falls between those two extremes: we suppose that uncertainty concerns the production externality σ, which affects the first-best level of S but not that of T.

In the interests of brevity, we only describe briefly the important differences that arise relative to the previous two sections when only σ is uncertain. We assume that σ can take two possible values with equal probability: a high realization $\bar{\sigma} + \Delta_\sigma$ and a low realization $\bar{\sigma} - \Delta_\sigma$, with $\Delta_\sigma > 0$.

In this environment, the $\{FB\}$ agreement takes the form $\{T = \gamma_0; S = \sigma - \gamma_0\}$, so that only S is now state-contingent. Characterization of the optimal agreement in the environment where only σ is uncertain is exactly analogous to that where only γ is uncertain, in the sense that each of the statements of Proposition 1 applies with σ taking the role of γ. However, as a result of the differences in the $\{FB\}$ agreement across the two environments, the potential appeal of the $\{T(\sigma)\}$ agreement in an environment where σ is uncertain is distinct from the potential appeal of the $\{T(\gamma)\}$ agreement in an environment where γ is uncertain, and arises for reasons analogous to the potential appeal of the escape-clause-type agreement $\{T(\alpha)\}$ in the environment where α is uncertain. In particular, like the $\{T(\alpha)\}$ agreement, the potential appeal of the $\{T(\sigma)\}$ agreement in an environment where σ is uncertain arises because it provides a means of managing the distortions associated with leaving S to discretion.

To see this, recall that if S is left to discretion (as in the agreement $\{T(\sigma)\}$) then it will be distorted from its Pigouvian level and used to manipulate the terms of trade. However, the higher is σ, the lower is the Nash trade volume and hence the higher will be the Nash equilibrium foreign export supply elasticity, and therefore the lower will be the terms-of-trade incentive to distort S away from its Pigouvian level. The implication, then, is that the incentive to distort S for terms-of-trade reasons will be stronger when σ is lower according to the monopoly power effect. A higher T mitigates the incentive to distort S for terms-of-trade purposes, and so allowing for a higher T when σ is lower can therefore help to mitigate the use of S for purposes of terms-of-trade manipulation when the incentive is highest to do so.

Finally, arguing along the lines highlighted in our discussion just above, it can be confirmed that increasing uncertainty over σ (Δ_σ) has impacts similar to those associated with increasing uncertainty over α (Δ_α) as described in the previous section. The reason is that, like the case of increasing Δ_α and unlike the case of increasing Δ_γ, uncertainty over σ is not directly relevant to the setting of T in the $\{FB\}$ agreement.

4. The Role of the National Treatment Clause

A distinguishing feature of the GATT/WTO (albeit increasingly less of the WTO) is that it combines rigid bindings on tariffs with a significant amount of discretion over internal measures.

This discretion is not complete, however. There are two legal instruments that can generally be invoked by Members in order to attack internal measures: the National Treatment (NT) provision in GATT Article III and the Non-Violation (NV) nullification-or-impairment provisions in GATT Article XXIII.1(b).[23] Each of these instruments can naturally be seen as an attempt to remedy problems that arise under the contractual incompleteness of the GATT/WTO. In this section we therefore evaluate the usefulness of the NT clause as a means to economize on contracting costs. In the next section we perform an analogous evaluation of the NV clause.

The ideal theoretical approach to evaluate these provisions would be to extend the set of feasible agreements $\mathcal{A}_0$ by allowing for more general instrument-based agreements as well as outcome-based agreements, and then examine whether and when NT and NV clauses of the particular type observed in the GATT/WTO are part of an optimal agreement within this wider class. Such an evaluation would be extremely complex however, partly because there would be a very large number of agreements to consider. We therefore attempt something more modest, relying on institutional motivation to restrict our attention to studying the potential desirability of just these particular clauses. In this section we expand the class of feasible agreements $\mathcal{A}_0$ to allow for agreements that include the NT clause, and examine conditions under which the optimal agreement in this wider class includes the NT clause. In the next section we expand the class of feasible agreements $\mathcal{A}_0$ to include the NV clause, and undertake an analogous exercise.

For our purposes, the relevant part of the NT clause can be found in GATT Article III.2, which addresses internal taxation. The exact interpretation of this provision is subject of debate in the legal and economic literature,[24] but we take the view that the core of the NT rule can be captured in the context of our model by the simple constraint $t_h = t_f$.[25] It is important to note that, while the NT provision restricts internal taxes to be the same, it *allows the importing country discretion* over the level at which these are set. Also note that, while the common tax level can be set in response to stochastic disturbances, the provision *as such* is not state-contingent.

[23] The WTO introduced a number of legal instruments that provide member governments with additional avenues for addressing specific internal measures (e.g., subsidies) and that were not available under GATT. These new legal instruments can be interpreted from the perspective of the analysis in the previous sections along the lines of eliminating discretion (and possibly adding rigidity) relative to the original agreement. We focus our analysis in this and the next section on NT and NV, respectively, since these legal instruments introduce qualitiative features that cannot be captured in the analysis of the previous sections.

[24] See Horn and Mavroidis (2004) for legal and economic analyses of Article III text and case law as it pertains to taxation.

[25] There are two interpretation issues that can be raised here. First, Article III.2 speaks of "treatment no less favorable," which suggests that a more accurate formalization of the NT provision is given by the inequality constraint $t_h \geq t_f$. However, in our model it can be shown that this constraint would always be binding, so that there would be no gain in allowing for this inequality constraint. Also, Article III.1, which sets the stage for the NT clause, restricts attention to measures that are applied "...so as to afford protection...". This poorly drafted sentence can be read in various ways, including that the provision should only have a bite against "protectionism," and not in cases where higher taxation of a foreign product is motivated by the pursuit of more legitimate objectives. This is not an issue in the context of the current model, since there is no ground for treating the imported product less favorably than the locally produced good in any of the setups we consider. However, if the model were slightly modified, so as to let the the imported product potentially be associated with a more severe regulatory problem than the domestically produced good, this would become an important issue; see Horn (2006).

An important issue is how to formalize the contracting costs associated with the NT provision. A literal application of our assumptions on contracting costs would imply that, since the NT clause refers to two policy instruments (t_h and t_f), the associated cost is $2c_p$, which is the same as specifying numbers for t_h and t_f, for example as in $\{t_h = 3, t_f = 5\}$. We believe however that it is more realistic to assume that the NT clause costs less than specifying numbers for two policy instruments. This would be especially compelling if we had many sectors, in which case specifying numbers for two policy instruments in all sectors would be vastly more complicated than specifying a blanket NT rule. We capture this type of consideration in our single-sector model by assuming that the NT clause costs less than $2c_p$. Formally, we assume the following: if the only constraint on consumption taxes is the NT rule, the associated contracting cost is $c_{NT} = k_{NT} \cdot c$, with $k_{NT} < 2$; otherwise our assumptions on contracting costs are unchanged.[26]

We refer to an agreement that includes the NT clause as an "NT-based" agreement. As indicated above, in this section we focus on an extended set of agreements that includes the class considered in the previous section ($\mathcal{A}_0$) plus the class of NT-based agreements. Letting $\mathcal{A}_{NT}$ denote the class of NT-based agreements, we are thus focusing on the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NT}$.

We begin by observing that the relationships between price wedges and policies are different for non-NT agreements and NT-based agreements. For non-NT agreements, these relationships are the same as in the previous section, as recorded in 2.1. Within this class, as we argued previously, we can focus on agreements that tie down S and/or T. However, for NT-based agreements, the relationships between price wedges and policies become

$$\begin{aligned} p &= p^* + \tau + t, \text{ and} \\ q &= p^* + \tau + s. \end{aligned} \tag{4.1}$$

Within this class, we can focus on agreements that tie down some or all of the instruments (τ, t, s).

In order to evaluate NT-based agreements, it is useful to derive expressions for the equilibrium prices and trade volume under NT. Equilibrium prices are easily derived to be

$$p(\tau, t, s) = [\alpha + \alpha^* + (\beta^* + \lambda^*)\tau - \lambda s + (\beta^* + \lambda + \lambda^*)t]/\Upsilon,$$

$$q(\tau, t, s) = [\alpha + \alpha^* + (\beta^* + \lambda^*)\tau + (\beta + \beta^* + \lambda^*)s - \beta t]/\Upsilon, \text{ and}$$

$$p^*(\tau, t, s) = q^* = [\alpha + \alpha^* - (\beta + \lambda)\tau - \lambda s - \beta t]/\Upsilon,$$

where as before, $\Upsilon \equiv \lambda + \lambda^* + \beta + \beta^*$, and the market-clearing trade volume is given by

$$M(\tau, t, s) = E^*(\tau, t, s) = [\alpha(\beta^* + \lambda^*) - \alpha^*(\beta + \lambda) - (\beta + \lambda)(\beta^* + \lambda^*)\tau - \beta(\beta^* + \lambda^*)t - \lambda(\beta^* + \lambda^*)s]/\Upsilon.$$

Notice that the NT clause by itself has no real effect of any kind, because as we have observed previously τ and t_f are perfectly substitutable policy instruments, and so any constraints placed

[26] If the contract specifies not only the NT rule but also the level of the common consumption tax, as in $\{t_h = t_f = 3\}$, the associated cost is assumed to be $2c_p$, so in this case there is no amendment to our initial assumptions on contracting costs. The only amendment we have introduced is a reduction in contracting cost for the case in which consumption taxes are handled only by the NT rule.

on t_f – such as the NT requirement that t_f remain equal to t_h – if applied in the absence of further constraints can always be undone by appropriate changes in τ. This point carries with it the important implication that both the efficient outcome and the Nash equilibrium derived in the previous section in the absence of NT, can also be implemented with policies that conform to the NT clause. In particular, it is direct to establish that the efficient policies under NT are given by

$$\tau^{eff} = 0, t^{eff} = \gamma, \text{ and } s^{eff} = \sigma.$$

Similarly, the Nash equilibrium outcome can be achieved with policies that conform to NT by setting

$$\tau^{NE} = \frac{p^*}{\eta^*}, \; t_h^{NE} = t_f^{NE} \equiv t^{NE} = \gamma, \text{ and } s^{NE} = \sigma.$$

There is thus no inherent violation of NT in the Nash equilibrium of our model. Therefore, to the extent the NT clause will have a real bite, it must be because *other* contractual obligations yield incentives for the importing country to use internal taxation in a discriminatory way.

It is useful as well to record the expression for the best-response level of t for any τ and s:

$$t^R(\tau, s) = t^{eff} + \frac{\Upsilon^2 - (\beta+\lambda)^2}{\Upsilon^2 - \beta(\beta+\lambda)}(\tau^{NE} - \tau) - \frac{\lambda(\beta+\lambda)}{\Upsilon^2 - \beta(\beta+\lambda)}(s^{eff} - s). \tag{4.2}$$

These expressions will be helpful in interpreting our results below.

There are potentially many kinds of NT-based agreements, but we can reduce the number that must be considered by focusing only on NT-based agreements that can be *strictly* optimal in the class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$. In this regard, we first observe that any NT-based agreement that ties down only one policy instrument is empty; this follows immediately from the fact that there is a degree of redundancy in the NT price relationships (4.1). Moreover, any NT-based agreement that ties down the level of the common consumption tax t is equivalent to or strictly dominated by a non-NT agreement in $\mathcal{A}_0$. This is because, given our assumptions on contracting costs, imposing the NT constraint and specifying the level of t is equivalent to specifying separately the levels of t_h and t_f (e.g. as in $t_h = 3$, $t_f = 3$), and hence we can apply the argument made in section 2 that such an agreement cannot do better than an agreement in class $\mathcal{A}_0$. We can conclude that the only NT-based agreements that can strictly improve over non-NT agreements in class $\mathcal{A}_0$ are NT-based agreements that tie down τ and s, leaving the government free to choose the common level of the consumption tax t.

Our observations thus far are valid regardless of the nature of the uncertainty. But to proceed from here we have to be more specific on the exact source of the uncertainty. As in the previous section, we will focus mostly on the case in which uncertainty concerns only the consumption externality γ; we will discuss briefly at the end of this section how the conclusions may change when the source of the uncertainty is different.

Given that only γ is stochastic, and in light of the arguments made just above, an immediate implication is that there are only *two* NT-based agreements that can be strictly optimal in our model: (i) the agreement that ties down τ and s in a rigid way, denoted $\{NT, \tau, s\}$, which costs $(2+k_{NT})\cdot c$; and (ii) the agreement that ties down τ in a contingent way, denoted $\{NT, \tau(\gamma), s\}$, which costs $(2 + k + k_{NT}) \cdot c$.[27] The next remark records the point.

Remark 3. *Consider the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$, and assume that only γ is uncertain. The only NT-based agreements that can be strictly optimal are $\{NT, \tau, s\}$ and $\{NT, \tau(\gamma), s\}$.*

Our next objective is to determine whether and under what conditions an NT-based agreement is strictly optimal for a range of contracting cost c.

We start with an intuitive discussion of the pros and cons of the NT-based agreements relative to non-NT agreements. Recall that, in addition to the $\{FB\}$ agreement and the empty agreement, there are three non-NT agreements that can be optimal: $\{T, S\}$, $\{T(\gamma)\}$ and $\{T\}$.

We begin with the contingent NT agreement, $\{NT, \tau(\gamma), s\}$. How attractive is this agreement as a way to save on contracting costs? Relative to $\{FB\}$, this agreement implies lower contracting costs (the difference being given by $2 - k_{NT}$). On the other hand, $\{NT, \tau(\gamma), s\}$ does not achieve the first best level of gross global welfare, because it leaves discretion over the common level of the consumption tax, and this discretion will be used by governments to manipulate terms of trade. This suggests that the attractiveness of the $\{NT, \tau(\gamma), s\}$ agreement is determined by similar factors as that of the $\{T(\gamma)\}$ agreement, except that the former entails a cost of *discretion over the consumption tax* t, $\Omega_{\{FB\}} - \Omega_{\{NT,\tau(\gamma),s\}}$, which we denote by CDt, whereas the latter entails the cost of *discretion over* S, which we previously defined as $\Omega_{\{FB\}} - \Omega_{\{T(\gamma)\}}$ and denoted by CD. Intuitively, then, understanding the key conditions under which $\{NT, \tau(\gamma), s\}$ is optimal for a range of c involves a comparison between CDt and CD.

Just as we have seen before that CD depends critically on the degree of instrument substitutability and the magnitude of the home-country monopoly power, it may be seen that CDt depends critically on these two factors as well, and for the same reasons.

To see this, consider first the relevance of the degree of substitutability between t and τ for CDt. Clearly, if t and τ are highly substitutable, then any constraints placed on τ (and s) through an NT-based agreement can be largely undone if t is left to discretion, just as with S and T and non-NT agreements. However, importantly, the underlying parameter conditions that cause t and τ to be highly substitutable are essentially *opposite* to the parameter conditions that cause S and T to be highly substitutable (low β and/or high λ). In particular, as inspection of the expression for $p^*(\tau, t, s)$ confirms, t is a close substitute for τ in manipulating the terms of trade when β is high and/or λ is low (while as with S and T, α has no bearing on the degree of instrument substitutability between t and τ).[28] When λ is close to zero, t and τ are close to perfect substitutes, and hence $\{NT, \tau(\gamma), s\}$ offers essentially no improvement over the empty

[27] As can be seen from the NT pricing relationships in 4.1, an agreement of the type $\{NT, \tau(\gamma), s(\gamma)\}$ is equivalent to an agreement where only τ is contingent, $\{NT, \tau(\gamma), s\}$, and also to an agreement where only s is contingent, $\{NT, \tau, s(\gamma)\}$. To fix ideas we focus on the agreement type $\{NT, \tau(\gamma), s\}$.

[28] We can define an index of substitutability between t and τ in a similar way as we defined the index of substitutability between T and S in the previous section. Namely, we consider the ratio $\frac{d\Omega(\tau, t^R(\tau), s)/d\tau}{\partial\Omega(\tau, t^R(\tau), s)/\partial\tau}$ evaluated at the Nash equilibrium. This yields $\left(\frac{d\Omega(\tau, t^R(\tau), s)/d\tau}{\partial\Omega(\tau, t^R(\tau), s)/\partial\tau}\right)_{\tau=\tau^{NE}, s=s^{NE}} = \frac{\lambda\Upsilon^2}{(\beta+\lambda)[\Upsilon^2 - \beta(\beta+\lambda)]}$. It can be shown that this index takes on a value of one when $\beta = 0$, a value of zero when $\lambda = 0$, and is monotonically decreasing in β and increasing in λ and independent of α. Accordingly, we may conclude as we indicate in the text that the degree of substitutability between τ and t is increasing in β, decreasing in λ and independent of α.

agreement. In this case, CDt approaches its maximum value of PGC. On the other hand, if β is close to zero, t is nearly useless as a surrogate means to distort terms of trade, and hence $\{NT, \tau(\gamma), s\}$ gets close to implementing the first best outcome. In this case, CDt approaches its minimum value of zero.

Consider next how CDt varies with the degree of home-country monopoly power, as measured by the (inverse of the) magnitude of the foreign export supply elasticity η^*. Using the expression in 4.2 for the best-response level of t, $t^R(\tau, s)$, and noting that $\tau^{NE} = \frac{p^*}{\eta^*}$, it can be seen that the distance between $t^R(\tau^{eff}, s^{eff})$ and t^{eff} is proportional to $\frac{p^*}{\eta^*}$. Hence, when $\frac{p^*}{\eta^*}$ is low and the magnitude of home-country monopoly power is therefore high, it follows that $t^R(\tau^{eff}, s^{eff})$ is far from t^{eff} and CDt is high. Therefore the same observation that we made in the previous section regarding the relation between CD and α applies for CDt as well: since a key determinant of the degree of home-country monopoly power is the import demand level, a higher level of α implies a higher CDt, thus making the $\{NT, \tau(\gamma), s\}$ agreement less attractive.

We now turn to consider the other NT-based agreement that can potentially be optimal, namely, the rigid NT agreement $\{NT, \tau, s\}$. It is useful to examine intuitively how this agreement compares with its contingent counterpart, $\{NT, \tau(\gamma), s\}$. The former saves an amount k in contracting costs relative to the latter, but it also achieves a lower level of gross global welfare, because it foregoes the benefit of making τ contingent on γ. In line with our terminology in the previous section, we can think of this difference in gross global welfare as the *cost of rigidity over* τ, $\Omega_{\{NT,\tau(\gamma),s\}} - \Omega_{\{NT,\tau,s\}}$, which we denote by $CR\tau$. How does $CR\tau$ vary with underlying parameters? It is key to recall that the efficient levels of τ and s are given by $\tau^{eff} = 0$ and $s^{eff} = \sigma_0$, and in particular that they do not depend on the consumption externality γ, and so making τ contingent on γ is beneficial only to the extent that this mitigates the incentive to distort t for terms-of-trade purposes. In particular, when γ is lower, the Nash equilibrium trade volume (as well as the efficient trade volume) is higher, and hence the incentive to manipulate the terms of trade is stronger, and this incentive can be mitigated by raising the level of τ. This type of benefit from state-contingency should by now be familiar, because it is similar to the benefit – highlighted in our analysis of non-NT agreements in sections 3.2 and 3.3 – from making T contingent on α or σ in environments where these parameters are uncertain.

An immediate implication of the observations made just above is that, when the degree of substitutability between t and τ is low, $CR\tau$ must be small, since in this case there is little to gain from making τ contingent. This in turn implies that, when t is a poor substitute for τ, the rigid NT agreement $\{NT, \tau, s\}$ gets close to implementing the first best outcome. On the other hand, when t is a close substitute for τ, the rigid NT agreement offers little improvement over the Nash equilibrium, just like the contingent NT agreement. These points can be expressed in terms of underlying parameters as follows: when β is small and therefore t is a poor substitute for τ, the rigid NT agreement gets close to implementing the first best outcome; and when λ is small and therefore t and τ are close to perfect substitutes, the rigid NT agreement provides a negligible improvement over the Nash equilibrium.

We can now put together the considerations developed above and ask under what conditions each of the two relevant NT-based agreements is optimal for a range of contracting costs. Applying the "concavity" condition of Lemma 3 (which is readily extended to the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$), it is straightforward to show that, if β is sufficiently small, so that t is a poor substitute for τ and S is a close substitute for T, the rigid NT agreement $\{NT, \tau, s\}$ is optimal for a range of c. To see this notice that, if β is small, moving from $\{NT, \tau, s\}$ to a more complex agreement can offer at best a negligible gain, while moving from $\{NT, \tau, s\}$ to a less complex agreement necessarily implies a non-negligible loss. This immediately implies that the condition of Lemma 3 is satisfied. Notice that we cannot draw the same conclusion for the contingent NT agreement $\{NT, \tau(\gamma), s\}$: if β is small, moving from this agreement to a more complex one implies a negligible gain, but also, moving from this agreement to the less complex agreement $\{NT, \tau, s\}$ implies a negligible loss, so the condition of Lemma 3 may or may not be satisfied.

It is also easy to see that if λ is sufficiently small, so that t is a close substitute for τ and S is a poor substitute for T, the "concavity" condition of Lemma 3 will be violated for both NT-based agreements, and therefore neither one can be optimal for any c. This is obvious, because if λ is small then both of these agreements are dominated by other agreements (namely, those that tie down only T) both in terms of higher performance and in terms of lower cost.

The following proposition summarizes these insights:

Proposition 5. *Consider the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$, and assume that only γ is uncertain.*
(i) If β is sufficiently small, so that the consumption tax t is a poor substitute for the tariff τ, the agreement $\{NT, \tau, s\}$ is strictly optimal for a range of c.
(ii) If λ is sufficiently small, so that the consumption tax t is a close substitute for the tariff τ, no NT-based agreement can be strictly optimal for any c.

Part (i) of the above proposition identifies a simple condition under which our model can rationalize the use of an NT-based agreement: this kind of agreement is strictly optimal if the degree of substitutability between t and τ is sufficiently small, and the level of contracting costs (c) lies in some intermediate range. On the other hand, as part (ii) of the proposition highlights, if the consumption tax t is a close enough substitute for the tariff τ, then none of the NT-based agreements can be optimal for any level of contracting costs.

We conclude this section by discussing the robustness of Remark 3 and Proposition 5 to the source of the uncertainty. If the source of uncertainty is the level of import demand (α) rather than the consumption externality, these results are valid exactly as stated (except of course that the contingent NT agreement is $\{NT, \tau(\alpha), s\}$ instead of $\{NT, \tau(\gamma), s\}$). There is a small change in results only if the source of uncertainty is the production externality σ, because this parameter directly determines the efficient level of the production subsidy s (recall that $s^{eff} = \sigma$). This has the following implication: when β is small, so that t is a poor substitute for τ, the rigid NT agreement $\{NT, \tau, s\}$ does *not* get close to implementing the first best outcome, because the first best level of s is contingent on σ, whereas the contingent NT agreement does. (Note that the contingent NT agreement in this case can be written as $\{NT, \tau, s(\sigma)\}$ – see footnote 27.) As a consequence, if β is sufficiently small, the contingent NT agreement satisfies the "concavity" condition of Lemma 3 and will be optimal for a range of c, whereas the rigid NT agreement may or may not satisfy that condition. Therefore, Proposition 5 is still valid if $\{NT, \tau, s\}$ is replaced by $\{NT, \tau, s(\sigma)\}$ in part (i).

5. The Role of Non-Violation Complaints

We now turn to an evaluation of the usefulness of the Non-Violation (NV) nullification-or-impairment provisions in GATT Article XXIII.1(b) as a means to economize on contracting costs. To develop this understanding, we begin with a general observation: tariff commitments in the GATT/WTO are interpreted as implying something beyond simple tariff obligations. In particular, when a government makes a tariff commitment in the GATT/WTO, it is understood to be committing to a level of "market access." Market access, in turn, is interpreted as reflecting the "conditions of competition" between domestic and foreign producers, something which is related to but not synonymous with import volume. Evidently, the market access guarantee that accompanies a tariff binding in the GATT/WTO is a subtle concept: it is a promise by the government not to alter its policies in the future in a way that would upset the conditions of competition established at the time of the original negotiations; but it is not an assurance against changes in market conditions, and hence it is not a promise of trade volumes. It is the GATT's NV clause that is often seen as providing this guarantee, by serving to protect the market access expectations of governments against changes in policies by their trading partners – even policies which are not contracted over – when those policy changes would have the effect of upsetting the market access that a government could have reasonably expected based on a prior GATT or WTO negotiation.[29]

In formalizing the NV clause, we seek to capture some rudimentary aspects of the above discussion. In particular, we suppose that the NV clause imposes a constraint on the Home government of the form: "Subsequent to negotiations in which the Home government binds its tariff, if the Home government changes its domestic policies then it must change its tariff to maintain the implied level of market access." Arguing along the lines of Bagwell and Staiger (2001), it can be shown that this constraint is equivalent to a commitment by the Home government that, subsequent to agreeing to a tariff binding, it will not take policy actions which would by themselves alter the foreign price p^*.

To state this constraint more precisely, we need to introduce some notation. Let t_f^0, s^0, and t_h^0 denote the values of the domestic instruments at the time of negotiation, and let θ denote the stochastic state variable – that is, α, γ or σ, depending on which one is uncertain. Finally, let $\tau(\theta)$ denote a (possibly state contingent) contracted reference tariff, and let $\hat{p}^*(\cdot;\theta)$ be the realized foreign price, given the realization of θ and the policies actually chosen by the Home government. We then formally capture the NV constraint on the Home country by the following

[29] For example, Hudec (1990, p. 24) describes the original logic for the inclusion of the NV clause in GATT in these terms: "...The standard trade policy rules could deal with the common types of trade policy measure governments usually employ to control trade. But trade can also be affected by other "domestic" measures, such as product safety standards, having nothing to do with trade policy. It would have been next to impossible to catalogue all such possibilities in advance. Moreover, governments would never have agreed to circumscribe their freedom in all these other areas for the sake of a mere tariff agreement. ... The shortcomings of the standard legal commitments were recognized in a report by a group of trade experts at the London Monetary and Economic Conference of 1933. The group concluded that trade agreements should have another more general provision which would address itself to any other government action that produced an adverse effect on the balance of commercial opportunity..."

restriction:

$$\hat{p}^*(\cdot;\theta) = p^*(\tau(\theta) + t_f^0, s^0 - t_h^0;\theta).$$

For any realized external circumstances θ and for any contracted tariff level $\tau(\theta)$, this constraint holds the Home government to policy choices which ensure that the realized price ($\hat{p}^*(\cdot;\theta)$) is the same as it would be under these same realized external circumstances if the Home government had set its tariff at the negotiated level $\tau(\theta)$ and maintained the internal policies t_f^0, s^0, and t_h^0 it had in place at the time of the original negotiation ($p^*(\tau(\theta) + t_f^0, s^0 - t_h^0;\theta)$). We will say that an agreement including such an obligation is "NV-based," and we denote the NV clause by NV. In this section we focus on an extended set of agreements that includes the class considered in section 3 ($\mathcal{A}_0$) plus the class of NV-based agreements. Letting $\mathcal{A}_{NV}$ denote the class of NV-based agreements, we are thus focusing on the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NV}$.

We begin by developing an important observation. When coupled with a tariff binding, the constraint imposed by the NV clause serves to grant the Home government a degree of unilateral discretion over its internal policies, but only up to a point. Specifically, the exercise of this discretion cannot undermine the implied level of market access that the tariff binding has granted to the foreign government. It is direct to establish within our model that this limited degree of discretion has a very attractive feature: it guarantees that, when coupled with a tariff binding (state-contingent or otherwise), the NV clause ensures that the Home government will always set its internal policies equal to the Pigouvian levels, and then adjust τ to satisfy the NV requirement.[30] Intuitively, as we have argued previously, the only reason that the Home government might distort its internal measures away from their Pigouvian levels is for the purpose of manipulating the terms of trade (p^*), when it is constrained from using its tariff for this purpose. But as the NV clause prevents the Home government's unilateral choice of policies from having any effect on p^*, this distorting motive is eliminated. From this discussion, we may therefore record the following remark:

Remark 4. *Any NV-based agreement that places a constraint on the tariff will deliver Pigouvian levels of the Home government's internal measures, that is, $t_f = \gamma, s = \sigma$, and $t_h = \gamma$.*

To explore further the potential appeal of NV-based agreements, we focus now on specific forms of uncertainty. We begin by considering uncertainty over the level of import demand α.

When α alone is uncertain, recall that the efficiently written first best instrument-based agreement is $\{T = \gamma_0; S = \sigma_0 - \gamma_0\}$ and costs $4 \cdot c$. In light of Remark 4 it is easy to see that, beginning from the Nash policies, the NV-based agreement $\{NV; \tau = 0\}$ also achieves the first best outcome in this environment: under the NV clause, the Home government will chose to maintain its internal measures at the levels $t_f = \gamma, s = \sigma$, and $t_h = \gamma$, which as confirmed in section 2.1 are also its Nash choices, and the NV requirement then requires that its tariff be maintained at $\tau = 0$ as well, which ensures efficiency. Binding the tariff at $\tau = 0$ costs c. But implementation of the NV clause requires verification of the state of import demand α,

[30] More accurately, what is guaranteed under the NV clause is that $S = \sigma - \gamma$, with T then adjusted to meet the NV requirement. However, recalling that $T \equiv \tau + t_f$ and $S \equiv s - t_h$, it can be seen than this is implied if the home government sets its internal policies equal to the Pigouvian levels, so that $t_f = \gamma, s = \sigma$, and $t_h = \gamma$, and then adjusts τ to satisfy the NV requirement.

which is utilized to run the appropriate "but for" counter-factual, and this costs an additional $k \cdot c$. So the cost of achieving the first-best with an NV-based agreement in this environment is $(1+k) \cdot c$. Notice that this implies immediately that the agreement $\{T(\alpha)\}$ is dominated by $\{NV; \tau = 0\}$, since the former costs $(2+k) \cdot c$ and does not achieve the first best. Note also that no other NV-based agreement can be of interest in this environment, since no (non-empty) NV-based agreement can cost less than $(1+k) \cdot c$.

A key observation may now be made. If $k \geq 3$, so that the costs of including states in the agreement is high as compared to the costs of including policies, then there can be no role for NV-based agreements in this environment: the optimal sequence of agreements in the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NV}$ is a subsequence of $(\{T = \gamma_0; S = \sigma_0 - \gamma_0\}, \{T\}, \{\emptyset\})$. Alternatively, if $k < 3$, then there is a role for NV-based agreements in this environment: the optimal sequence of agreements in the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NV}$ is a subsequence of $(\{NV; \tau = 0\}, \{T\}, \{\emptyset\})$.

Intuitively, the key role played by k in determining the appeal of NV-based agreements reflects a basic trade-off between NV-based and instrument-based agreements. On the one hand, by utilizing the NV clause and not specifying any internal policy instruments directly, an agreement of the form $\{NV; \tau = 0\}$ economizes on contracting costs associated with *policies* relative to the agreement $\{T = \gamma_0; S = \sigma_0 - \gamma_0\}$.[31] On the other hand, as we observed above, implementation of the NV clause requires verification of the state of import demand α, and so the agreement $\{T = \gamma_0; S = \sigma_0 - \gamma_0\}$ economizes on contracting costs associated with *states* relative to an agreement of the form $\{NV; \tau = 0\}$. Hence, at a general level a contract of the form $\{NV; \tau = 0\}$ will be a relatively low-cost way of delivering the first best when k, the cost of including states relative to the cost of including policies in the agreement, is low.

The following proposition summarizes these insights:

Proposition 6. *Consider the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NV}$, and assume that only α is uncertain. Then the agreement $\{NV; \tau = 0\}$ is strictly optimal for a range of c if and only if $k < 3$.*

We consider next uncertainty over the level of the consumption externality γ. In this environment, recall that the efficiently written first best instrument-based agreement is $\{T = \gamma; S = \sigma_0 - \gamma\}$ and costs $(4+k) \cdot c$. There are now two NV-based agreements that warrant consideration: $\{NV; \tau(\gamma)\}$, which achieves the first best and costs $(1+k) \cdot c$; and $\{NV; \tau\}$, which costs c. That the agreement $\{NV; \tau(\gamma)\}$ achieves the first best can be seen by noting that, under the first best policies $T = \gamma$ and $S = \sigma_0 - \gamma$, the level of p^* is state-dependent and given by $p^* = [\alpha + \alpha^* - (\beta + \lambda) \cdot \gamma - \lambda \cdot (\sigma_0 - \gamma)]/\Upsilon$. Therefore, $\tau(\gamma)$ can be set to achieve the first-best state-dependent p^*, with the NV clause then ensuring that the Home government makes the efficient policy choices to deliver the first-best p^* in each state. In this light, the

[31] A similar idea is expressed by Sykes (2003, p. 18) in the context of comparing the NV clause as a method to discipline domestic subsidies in relation to the more direct approach taken under the WTO Agreement on Subsidies and Countervailing Measures: "A nice feature of the nonviolation doctrine is the fact that it does not require subsidies to be carefully defined or measured. A complaining member need simply demonstrate that an unanticipated government program has improved the competitive position of domestic firms at the expense of their foreign competition. The administration of the doctrine is thus reasonably straightforward, and the fighting issue is likely to be whether the government policy in question was foreseen by trade negotiatiors."

NV-based agreement $\{NV;\tau\}$ saves $k \cdot c$ on contracting costs relative to $\{NV;\tau(\gamma)\}$, but fails to achieve the first best as a result of the rigid p^* that this agreement implies.

In this environment, it is direct to confirm that the NV-based agreement $\{NV;\tau(\gamma)\}$ dominates both the efficiently written first best instrument-based agreement $\{T = \gamma; S = \sigma_0 - \gamma\}$ and the agreement $\{T(\gamma)\}$. Moreover, it can also be confirmed that the NV-based agreement $\{NV;\tau\}$ dominates $\{T, S\}$: neither agreement permits a state-contingent p^*, but the agreement $\{NV;\tau\}$ costs c and always delivers (state-contingent) Pigouvian levels of the Home government's internal measures, while the agreement $\{T, S\}$ costs $4 \cdot c$ and cannot achieve Pigouvian intervention it each state. Hence, in the environment where γ is uncertain, the optimal sequence of agreements in the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NV}$ is a subsequence of $(\{NV;\tau(\gamma)\}, \{T\}, \{NV;\tau\}, \{\emptyset\})$.

Clearly, then, there is a role for NV-based agreements in the environment where γ is uncertain, because for c sufficiently low the optimal agreement is the first best NV-based agreement $\{NV;\tau(\gamma)\}$. The remaining question is whether the NV clause can ever be part of an optimal incomplete agreement, and if so, under what conditions. This amounts to the question of whether the NV-based agreement $\{NV;\tau\}$ is ever strictly optimal for a range of c.

Consider first the case where $k \leq 1$, so that the cost of including states relative to the cost of including policies in the agreement is low. In this case $\{NV;\tau(\gamma)\}$ dominates $\{T\}$, and so the optimal sequence of agreements in the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NV}$ is a subsequence of $(\{NV;\tau(\gamma)\}, \{NV;\tau\}, \{\emptyset\})$. To evaluate when the agreement $\{NV;\tau\}$ will be strictly optimal for a range of c, it is helpful to define *the cost of rigidity over* p^*, $\Omega_{\{FB\}} - \Omega_{\{NV,\tau\}}$,which we denote by CRp^*. Like its counterparts CRT, $CRTS$, and $CR\tau$ in the instrument-based agreements considered in previous sections, the concept embodied in CRp^* is the natural measure of the cost of rigidity in an outcome-based contract such as an NV-based agreement. Using Lemma 3 and recalling that the potential gain from contracting, PGC, is defined by $\Omega_{\{FB\}} - \Omega_{\{\emptyset\}}$, it is direct to confirm that the agreement $\{NV;\tau\}$ will be strictly optimal for a range of c if and only if $PGC \geq [(1+k)/k]CRp^*$. But observing that the expression for the state-dependent value of p^* under the first-best policies given just above can be simplified to $p^* = [\alpha + \alpha^* - \beta \cdot \gamma - \lambda \cdot \sigma_0)]/\Upsilon$, it is immediate that the value of p^* under the first-best policies approaches a fixed (non-state-dependent) value as β approaches zero, and hence CRp^* must approach zero as well (while PGC remains strictly positive). We may thus conclude for the case of $k \leq 1$ that $\{NV;\tau\}$ will be optimal for a range of c provided that β (and hence CRp^*) is sufficiently low.

Consider next the case where $k > 1$, so that the cost of including states relative to the cost of including policies in the agreement is high. In this case, the optimal sequence of agreements in the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NV}$ is a subsequence of $(\{NV;\tau(\gamma)\}, \{T\}, \{NV;\tau\}, \{\emptyset\})$. Again using Lemma 3, and recalling that the cost of discretion (CD) and the cost of rigidity in the presence of discretion (CRT) are defined as $\Omega_{\{FB\}} - \Omega_{\{T(\gamma)\}}$ and $\Omega_{\{T(\gamma)\}} - \Omega_{\{T\}}$, respectively, it is straightforward to confirm that the agreement $\{NV;\tau\}$ will be strictly optimal for a range of c if and only if (i) $PGC \geq [(1+k)/k]CRp^*$ and (ii) $PGC \geq [CRp^* - CD] + [\ CRp^* - CRT]$. But with CD and CRT strictly positive, it can be confirmed that both (i) and (ii) are satisfied provided that β (and hence CRp^*) is sufficiently low.

The following proposition summarizes these insights:

Proposition 7. *Consider the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NV}$, and assume that only γ is uncertain.*

(i) Then the agreement $\{NV; \tau(\gamma)\}$ is the efficiently written first best agreement, and it is strictly optimal for a range of low c.
(ii) If β is sufficiently small, so that the degree of state-contingency of p^* under the first-best policies is low and therefore CRp^* is also low, then the agreement $\{NV; \tau\}$ is strictly optimal for a range of c.

We conclude this section by observing that the analysis of NV-based agreements in the presence of uncertainty over the production externality σ yields conclusions which are identical to those we have reported when uncertainty concerns the consumption externality γ. The reason is simply that, for NV-based agreements, what is important about the nature of uncertainty is whether it effects p^* directly (as with uncertainty over α) or only indirectly through efficient policy choices (as with uncertainty over γ or σ). If uncertainty takes (only) the first form, then whether or not the efficiently written first best agreement will be instrument-based or NV-based will depend on k, the relative cost of including states in the agreement relative to policies. On the other hand, if uncertainty takes (only) the second form, then the efficiently written first best agreement will always be an NV-based agreement, and whether the NV clause will also be part of an optimal *incomplete* agreement for some contracting costs will depend on the relative cost of including states in the agreement relative to policies.

6. The Role of Weak Bindings

In the previous sections we focused on agreements that impose equality constraints, as in $\{T = 2\}$ or $\{NT, \tau = 0, s = \sigma\}$. This, of course, would be without loss of generality in a world of costless contracting, for in this case it would be optimal to implement the first best outcome, and this can be achieved with an agreement that imposes equality constraints (as in $\{T = \gamma, S = \sigma - \gamma\}$), so there would be nothing to gain from using inequality constraints.[32] In the presence of contracting costs, however, it may not be optimal to implement the first-best outcome, and in a second-best environment it may be preferable to impose policy *ceilings* (as in $\{T \leq 2\}$) rather than equality constraints. Below we will make this claim more formal, but as a first step we develop some intuition through a simple example.

Suppose the only uncertain parameter is the production externality σ, which can take two values, σ_0 and σ_1, and set $\gamma \equiv 0$. Moreover, assume that c_s is very high, so that writing a state contingent agreement is too costly to be worthwhile, and that $c_p = 0$, so that it is costless to write rigid contracts. Now suppose that $\Pr(\sigma_0)$ is small, so that, if one ignores the possibility of inequality constraints, the optimal contract is of the form $\{\bar{T}, \bar{S}\}$, where $\bar{T}$ is close to zero and $\bar{S}$ is close to σ_1. Consider now an alternative agreement that imposes the same restriction on T, but only a ceiling on S: $\{\bar{T}, S \leq \bar{S}\}$. Clearly, this contract is at least as good as the previous one, because given the constraint on T, the government's incentive to distort S is upwards: the government is tempted to "oversubsidize" production. Is this contract *strictly* preferable to the previous one? This depends crucially on the subsidy reaction function $S^R(T; \sigma)$ and on the support of σ. In particular, if $S^R(0; \sigma_0)$ is below σ_1, then in state σ_0 the agreement $\{\bar{T}, S \leq \bar{S}\}$ induces a more efficient outcome than the agreement $\{\bar{T}, \bar{S}\}$. Intuitively, if the government is

[32] To distinguish bound levels, we denote these by a bar in this section.

free to go below the ceiling level S, it will choose to do so if the externality σ is low enough, and this is beneficial for global welfare, even though the government will not decrease S all the way to the Pigouvian level.

The above reasoning suggests that inequality constraints should dominate equality constraints when the support of the state vector is large, so that imposing an equality constraint creates opposite incentives to deviate in different states: an incentive to go above the binding $\bar{S}$ in some states and an incentive to go below $\bar{S}$ in other states. The next proposition, which is proved in the Appendix, confirms and extends this intuition. We will refer to a constraint of the kind $S \leq \bar{S}$ as a "rigid weak binding" and to a constraint of the kind $S = \bar{S}$ as a "rigid strong binding."

Proposition 8. *Rigid weak bindings are weakly preferable to rigid strong bindings. The preference is strict if the support of the state vector is sufficiently large.*

It is interesting to note that a rigid weak binding combines rigidity and discretion, since the ceiling does not depend on the state of the world and a government has discretion to set the policy below the ceiling. Thus, our result highlights another sense in which rigidity and discretion may be complementary ways to economize on contracting costs: in section 3.1 we showed that, if uncertainty concerns γ, the cost of discretion over domestic instruments is lower in the presence of rigidity; here we have shown that, conditional on the agreement being rigid, it may be valuable to concede downward discretion in the setting of the relevant policies.

In reality, the support of the state vector that is relevant for a trade agreement is likely to be very large. The Proposition thus suggests that the constraints imposed by trade agreements should predominantly take the form of weak bindings. This suggestion is consistent with the observed nature of GATT/WTO policy commitments, which essentially all take the form of weak bindings.[33]

7. Conclusion

To be written

[33] We note here that this is not the only possible explanation for the use of weak bindings. Maggi and Rodriguez-Clare (2005) propose an alternative explanation based on political-economy considerations: their basic idea is that weak bindings allow governments to extract rents from lobbies after the agreement is signed, since they allow a government to credibly threaten its domestic lobbies to lower the level of trade protection below the ceiling. We also note that the explanation proposed here is somewhat related to the explanation proposed in Bagwell and Staiger (2005), where weak bindings may be preferred to strong bindings in the presence of political-economy shocks that are privately observed by governments.

8. Appendix

Proof of Proposition 8

We will consider those agreements that include rigid strong bindings (RSBs) and are optimal for some c, and show that, if we replace RSBs with rigid weak bindings (RWBs), efficiency is weakly increased. We can focus on the following contracts: (a) the best $\{T = \bar{T}\}$ agreement; (b) the best $\{T = \bar{T}; S = \bar{S}\}$ agreement; and (c) the best $\{NT; \tau = \bar{\tau}; s = \bar{s}\}$ agreement. Note that for the NV-based agreement, a rigid binding on the tariff $\bar{\tau}$ serves only as a reference point, therefore nothing changes if it is replaced with $\tau \leq \bar{\tau}$.

Let us start with case (a). Consider replacing $\{T = \bar{T}\}$ with $\{T \leq \bar{T}\}$. This can decrease efficiency only if in some state the government chooses $T < \bar{T}$ and this implies lower global welfare than $T = \bar{T}$. But T will only be set below $\bar{T}$ if the Nash import tax $T^{NE} < \bar{T}$, in which case the importing country will set $T = T^{NE}$. Let us show that Ω decreases in T in the range $(T^{NE}, \bar{T})$. Recalling that the subsidy is set as $S = S^R(T)$, we need to evaluate the derivative

$$\frac{d}{dT}\Omega(T, S^R(T)) = W_T(T, S^R(T)) + \frac{d}{dT}W^*(T, S^R(T))$$

where we have used the envelope theorem to set $\frac{d}{dT}W(T, S^R(T)) = W_T(T, S^R(T))$. Clearly $W_T < 0$ for $T > T^{NE}$. Also, the sign of $\frac{d}{dT}W^*(T, S^R(T))$ is the same as the sign of $\frac{d}{dT}p^*(T, S^R(T))$. It is direct to verify that this derivative is negative, which in turn implies $\frac{d}{dT}\Omega(T, S^R(T)) < 0$ for $T > T^{NE}$. We can conclude that switching to a weak binding cannot decrease Ω.

Next consider case (b), and consider replacing $\{T = \bar{T}; S = \bar{S}\}$ with $\{T \leq \bar{T}; S \leq \bar{S}\}$. For a given state, there are four relevant possibilities for how the importing country sets (T, S) under an agreement $\{T \leq \bar{T}; S \leq \bar{S}\}$:

(i) the importing country chooses $(T = \bar{T}, S = \bar{S})$: In this case there is of course no change in Ω relative to $\{\bar{T}; \bar{S}\}$.

(ii) the importing country chooses $(T = \bar{T}, S = S^R(T))$: Here it must be that $S > S^R(T)$. Let us evaluate $\Omega_S = W_S + W^*_S$.Clearly, $W_S < 0$ for $S > S^R(T)$, and $W^*_S < 0$, hence $\Omega_S < 0$ in this region, which in turn implies that switching to weak bindings increases Ω.

(iii) the importing country chooses $(T = T^R(S), S = \bar{S})$: Here it holds that $T > T^R(S)$. Let us evaluate $\Omega_T = W_T + W^*_T$. Since $W_T < 0$ for $T > T^R(S)$, and $W^*_T < 0$, it follows that $\Omega_T < 0$ in this region, which ensures that switching to weak bindings increases Ω.

(iv) the importing country chooses $(T = T^{NE}, S = S^{NE})$: The same result can be shown by combining the arguments we just made for cases (ii) and (iii).

Consequently, a switch from $\{T = \bar{T}; S = \bar{S}\}$ to $\{T \leq \bar{T}; S \leq \bar{S}\}$ cannot decrease Ω.

Finally, consider case (c). Since the NT agreement fixes the wedge $q - p^*$ and leaves the wedge $p - p^*$ discretionary, it is convenient to re-define variables as follows:

$$\begin{aligned} p - p^* &\equiv z \\ q - p^* &\equiv v \end{aligned}$$

We can think of z and v as the policy instruments and of the NT agreement as imposing a constraint $v = \bar{v}$. Also, it is useful to rewrite the world price as a function of v and z as

$$p^* = \frac{1}{\Upsilon}(\alpha + \alpha^* - \beta z - \lambda v)$$

and the reaction function for z as

$$z^R(v) = \gamma + \frac{M + \lambda(v - \sigma)}{\beta^* + \lambda + \lambda^*}$$

Let us now replace the agreement $\{NT; \tau = \bar{\tau}; s = \bar{s}\}$ with $\{NT, \tau \leq \bar{\tau}, s \leq \bar{s}\}$. In the new notation, this means replacing the constraint $v = \bar{v}$ with the constraint $v \leq \bar{v}$. We can apply a similar argument as for case (a): it suffices to show that, for any given state, $\Omega(v, z^R(v))$ is decreasing in v for $v > v^{NE}$. Thus we need to study the derivative

$$\frac{d}{dv}\Omega(v, z^R(v)) = W_v(v, z^R(v)) + \frac{d}{dv}W^*(v, z^R(v))$$

Clearly, $W_v < 0$ for $v > v^{NE}$. Also, $\frac{d}{dv}W^*(v, z^R(v))$ has the same sign as $\frac{d}{dv}p^*(v, z^R(v))$. It can be verified that

$$\frac{d}{dv}p^*(v, z^R(v)) = -\frac{\lambda}{\beta + \beta^* + \lambda^*} < 0$$

This implies that switching to weak bindings cannot decrease Ω.

We have shown that replacing RSBs with RWBs cannot decrease the expected surplus. It remains to show that there is some state $\boldsymbol{\theta}$ for which replacing RSBs with RWBs increases Ω strictly. Applying the arguments developed above, we know that a sufficient condition for this to happen is that the noncooperative level for a policy is below the ceiling of that policy. It is easy to show that there exists a state of the world for which this is the case, and hence if the support is large enough it will encompass such a state. **QED**