

Trade Agreements as Endogenously Incomplete Contracts

Henrik Horn (Research Institute of Industrial Economics, Stockholm, and CEPR)
Giovanni Maggi (Yale University, NBER and CEPR)
Robert W. Staiger (Stanford University and NBER)*

Revised: July 2008

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 real trade agreements.

*This paper has benefited from the detailed and helpful comments of three referees and the Co-Editor, as well as from the helpful comments of Kyle Bagwell, Pierpaolo Battigalli, Gene Grossman, Elhanan Helpman, Robert Lawrence, Andres Rodriguez-Clare, Johan Stennek and seminar participants at Calgary, CEMFI, Ente Luigi Einaudi, Harvard, Minneapolis Fed, Penn State, Princeton, UBC, UC Davis, UCSD and Yale, and from participants in an ERWIT conference and in a conference at CREI (Universitat Pompeu Fabra). Horn gratefully acknowledges financial support from the Marianne and Marcus Wallenberg Foundation. Maggi gratefully acknowledges financial support from the NSF (SES-0351586) and thanks the Ente Luigi Einaudi for hospitality during part of this project. Staiger gratefully acknowledges financial support from the NSF (SES-0518802).

1. Introduction

The World Trade Organization (WTO) regulation of trade in goods – the General Agreement on Tariffs and Trade (GATT) – is obviously a highly incomplete contract. And while the GATT/WTO is the most central trade agreement in the world trading system, this characterization applies as well to every other entry in the vast catalogue of existing trade agreements. In the economics literature there exist formal models that examine various aspects of this incompleteness, but 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.¹ This literature illuminates the *consequences* of the incompleteness of trade agreements, but it cannot *explain* the particular forms that the incompleteness has taken, because the incompleteness is assumed rather than endogenously derived.

The broad purpose of this paper is to take the analysis of trade agreements as incomplete contracts one step further, by *endogenously* determining the choice of contract form. A more specific purpose is to demonstrate that an incomplete-contracting perspective can help explain some core features of the GATT/WTO, including the following: (i) The agreement binds the levels of trade instruments. In contrast, domestic instruments are largely left to the discretion of governments, with two important exceptions: first, internal policies have to respect the National Treatment clause; and second, the WTO has introduced a regulation of domestic subsidies; (ii) The bindings are largely rigid (i.e. not state-contingent). But there are “escape clauses” that allow countries to unilaterally impose temporary protection or to renegotiate bindings; and (iii) The bindings only stipulate upper bounds on the tariffs that can be applied, thus leaving governments with discretion to go below the bounds.

An important aspect of the incompleteness of the GATT/WTO, which is embodied in the above features but also reflected to varying degrees in other trade agreements, is that the agreement displays an interesting combination of *rigidity*, in the sense that contractual obligations are largely insensitive to changes in economic (and political) conditions, and *discretion*, in the sense that governments have substantial leeway in the setting of many policies.

In this paper we propose a simple theoretical framework where the incompleteness of the agreement, and in particular the manner and degree in which discretion and rigidity are present in the agreement, is optimally determined. The analytical starting point of the paper is the

¹An incomplete list of papers that fall into this category is Copeland (1990), Bagwell and Staiger (2001), Battigalli and Maggi (2003), Horn (2006) and Costinot (2008).

notion that governments face two fundamental sources of difficulty when designing a trade agreement. The first is that there is a *wide array of policy instruments* – border measures and especially internal “domestic” measures – that must be constrained to keep in check each government’s incentives to act opportunistically. This feature suggests that the agreement should be comprehensive in its policy coverage. The second is that there is *significant uncertainty* concerning the circumstances that will prevail during the life-time of the agreement. This feature suggests that the agreement should be highly adaptable to the contingencies that unfold.

Of course these features would not pose a problem if contracting were costless. But in reality there are important costs associated with forming a trade agreement. While these costs can take a variety of forms, it is likely that they are higher when the agreement is more detailed, both in terms of the number of policies that it seeks to constrain and the number of contingencies that it specifies. We explicitly incorporate the costs of contracting over policies and contingencies into an analysis of the optimal structure of a trade agreement.

An objection might be raised that, when it comes to trade agreements, the costs of contracting are likely to be small relative to the potential gains from an agreement, and so the costs of contracting are unlikely to have important effects on the structure of trade agreements. But one should keep in mind the vast number of products, countries, policy instruments and contingencies that are involved when designing and implementing such an agreement. Indeed, the WTO Agreement, which by all accounts is considered to be an extremely incomplete agreement, still fills some 24,000 pages, and it took approximately 8 years of negotiations to complete. Hence, we believe that it is reasonable to view the contracting costs associated with trade agreements as significant even relative to the potential benefits of the agreement, and that these costs are then likely to shape the nature of the agreement.²

We work within a competitive two-country setting, where both consumption and production may create (localized) externalities, thereby giving rise to a variety of possible efficiency

²The difficulties associated with writing an agreement that is comprehensive in policy coverage and is highly contingent have been emphasized in the trade-law literature. For example, Hudec (1990) writes: “...The standard trade policy rules could deal with the common type 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” (p. 24). Also, Schwartz and Sykes (2001) write: “...Many contracts are negotiated under conditions of considerable complexity and uncertainty, and it is not economical for the parties to specify in advance how they ought to behave under every conceivable contingency ... The parties to trade agreements, like the parties to private contracts, enter the bargain under conditions of uncertainty. Economic conditions may change, the strength of interest group organization may change, and so on” (pp. 181-4).

rationales for policy intervention. For simplicity we focus on intervention in import sectors, and assume that governments possess a complete set of taxation instruments: we focus first on import tariffs and production subsidies, but later also consider consumption taxes in order to evaluate the National Treatment clause. Uncertainty plays a central role. To bring out the main points, we consider three sources of uncertainty: the magnitude of the consumption externality, the magnitude of the production externality, and the magnitude of the underlying trade volume between countries.

We formalize the notion of contracting costs in a simple way. Following an approach similar to that of Battigalli and Maggi (2002), we assume that these costs are increasing in the number of state variables and policies included in the agreement, and we characterize the agreement that maximizes expected global welfare minus contracting costs (the “optimal” agreement).

We begin by examining the first-best outcome and the no-agreement outcome – that is, the noncooperative equilibrium. In the absence of an agreement, the importing country would use its policy instruments to correct the externalities in a Pigouvian fashion, but it would also utilize its tariff to manipulate the terms of trade. 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 externalities, the first-best agreement would be very simple: it would just stipulate laissez-faire across all policy instruments and under all circumstances. But due to the externalities, the contracting problem is substantially more complex: the first-best agreement involves the use of policy instruments, and it requires these policies to be state-contingent if the externalities are uncertain.

As a result of contracting costs, the governments may find it worthwhile to write an agreement that is simpler than the first-best agreement. As our discussion above suggests, there are two ways to save on contracting costs: the agreement can be made (partially or fully) rigid; and/or some of the policies can be left to the discretion of governments. The focus of our analysis hence consists of examining the optimal degrees of rigidity and discretion in the trade agreement, and how these depend on contracting costs and features of the underlying economy.

Our first main result is that it cannot be optimal to contract over domestic subsidies while leaving tariffs to discretion. Intuitively, this finding reflects a kind of “targeting principle” logic (Bhagwati and Ramaswami, 1963 and Johnson, 1965): contracting over domestic policies alone is suboptimal, because as we have described above it is the tariff that is the source of the (terms-of-trade driven) inefficiency in the noncooperative equilibrium. This finding accords well

with the emphasis on trade over domestic instruments that characterizes the GATT/WTO; and while this feature is often explained informally as deriving from distinct levels of contracting costs that reflect differences in transparency across these instruments, our model imposes no such distinction, and so it identifies in this respect a more fundamental explanation.

In light of our first result, two broad questions remain for the design of the optimal agreement: whether subsidies should (also) be constrained by the agreement; and whether the agreement should be state-contingent and, if so, what state variables should be included. Regarding the first of these questions, we show that governments face the following tradeoff. On one hand, there are benefits from leaving the subsidy to discretion: a first benefit is the direct savings in contracting costs; but we also identify a second benefit which applies whenever the agreement would be rigid with respect to the externalities, and this is the *indirect state contingency* for the subsidy that is accomplished when the subsidy is left out of the (rigid) agreement. On the other hand, the cost of leaving the subsidy to discretion takes the form of distortions in the level of the subsidy for terms-of-trade manipulation: we identify *monopoly power*, *trade volume*, and *instrument substitutability* effects as key features of the trade-agreement contracting environment that determine the severity of these distortions and hence the costs of discretion.

Using these effects, our second main result is that the optimal agreement will leave domestic subsidies to discretion if: (i) countries have little monopoly power in trade, in which case they have little ability to manipulate terms of trade; or (ii) they trade little, in which case they gain little from exploiting their power over terms of trade; or (iii) subsidies are a poor substitute for tariffs as a tool for manipulating the terms of trade.

The result that a larger trade volume makes discretion over domestic subsidies less attractive suggests a possible explanation for the fact that the WTO has introduced a regulation of domestic subsidies that was not present in GATT: broadly speaking, the explanation is that a general increase in trade volumes over time has increased the cost of discretion, thereby heightening the need to constrain domestic policies. And in combination with the instrument substitutability effect, the model suggests as well a reason why developing countries may have been largely exempted from the WTO's (rigid) regulation of domestic subsidies through "special and differential treatment" clauses: the typical developing country may lack both the size in world markets to wield substantial market power and the array of domestic policy instruments necessary to find easy substitutes for tariffs.

We next consider the question whether or not the agreement should be state-contingent

and, if so, what state variables should be included. Intuitively, the cost of specifying a given state variable in the agreement must be weighed against the benefit of making the agreement contingent on that state variable, and as might be expected this benefit is in turn determined in large part by the degree of uncertainty over the magnitudes of the externalities. But whenever the agreement constrains the tariff while leaving the domestic subsidy to discretion, our model also suggests a more subtle insight concerning *why* state-contingent tariff commitments may be beneficial, and therefore *which* contingencies to introduce in the agreement. In particular, as we have observed, the incentive to distort the domestic subsidy for terms-of-trade purposes grows with the underlying trade volume, and this suggests that there may be a benefit from introducing contingent tariff commitments into the agreement as a way to mitigate this incentive against especially high trade volumes. We label this the *indirect incentive management* effect.

The identification of the indirect incentive management effect leads to our third main result: conditional on leaving the domestic subsidy to discretion, and provided that the cost of including state-contingencies is not too high, the optimal agreement will specify tariff commitments that are contingent on state variables that have implications for trade volume but are *irrelevant* to the first-best tariff level. This result implies that, when there is substantial uncertainty about the level of import demand, 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 high. Our rationale for an escape clause is, however, quite different from those that have been highlighted in the existing theoretical literature.³ An escape clause can be appealing in our model as a way to manage the higher incentives to distort domestic instruments for terms of trade purposes in periods of high underlying import volume.

We derive these first three results in a fairly general economic environment, but to fully characterize the optimal agreement and evaluate how it changes as fundamental parameters change, we focus on a parameterized (linear) version of the model with one-dimensional uncertainty. In this setting we show that the indirect state-contingency effect tends to make rigidity and discretion *complementary* methods of saving on contracting costs, while the indirect incentive-management effect tends to make rigidity and discretion *substitutes*. We also show how these effects arise either in isolation or together depending on the exact source of the uncertainty. Finally, we demonstrate that, depending on its source, increased uncertainty can

³An 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.

lead to either less or more rigidity in the optimal agreement. In this way, we establish that the role of uncertainty in shaping the optimal agreement depends in subtle ways on its source.

For the remainder of the paper we return to our more general model and extend the analysis to shed light on two other core aspects of the GATT/WTO: the presence of a National Treatment (NT) clause for internal taxes, and the fact that tariffs are constrained by “weak” bindings (i.e. upper bounds) rather than by “strong” bindings (i.e. exact levels).

We evaluate the NT clause as a means of saving on contracting costs. To this end, we introduce distinct consumption taxes on domestically-produced and imported goods, and we interpret the NT clause as requiring that consumption taxes be equalized across these goods.

We first show that an agreement which includes the NT clause but does not bind the consumption tax offers a novel form of discretion (over the consumer price wedge) that *cannot* be achieved without the NT clause (where a decision not to bind the production subsidy implies discretion over the producer price wedge). We then investigate the circumstances under which this type of discretion is desirable, and use the results of this investigation to derive a simple condition under which the optimal agreement includes the NT clause. This condition describes circumstances in which the NT-based agreement gets close to the first best without utilizing costly state-contingencies, by utilizing instead the indirect state-contingency associated with discretion over internal taxes constrained only by the NT clause.

Finally, we argue that the presence of contracting costs may explain why GATT stipulates weak bindings rather than strong bindings. More specifically, we show that the optimal agreement may include *rigid* weak bindings. The appeal of this type of binding is that it combines rigidity and discretion in a novel fashion, since the ceiling does not depend on the state of the world, and the government has (downward) discretion to set the policy below the ceiling.⁴

Section 2 lays out the basic model and characterizes the role of rigidity and discretion in the optimal agreement. Sections 3 and 4 evaluate the NT clause and weak bindings, respectively. Section 5 concludes. The Appendix provides proofs not contained in the body of the paper.

⁴Our rationale for weak bindings is again quite distinct from others that have been recently proposed. Maggi and Rodriguez-Clare (2007) propose an alternative explanation based on political-economy considerations, whereby weak bindings may be desirable because they induce lobbies to pay contributions even after the agreement is signed. More closely related to our rationale is the explanation proposed by Bagwell and Staiger (2005), where weak bindings may be preferred in the presence of political-economy shocks that are privately observed by governments. However, there are important differences: for one, there the appeal of weak bindings is due to the combination of private information and the absence of international transfers, whereas here it is due to the presence of contracting costs; for another, there only import tariffs are considered, while here the appeal of weak bindings is shown to extend to domestic subsidies as well.

2. The Basic Model

We consider two countries, Home and Foreign. There are three goods, a numeraire good and two non-numeraire goods (labeled 1 and 2). Home is a natural importer of good 1 and Foreign a natural importer of good 2. Markets are perfectly competitive, but we allow for the presence of a production externality and a consumption externality. In this way, we introduce multiple economic rationales for policy intervention; as we explain later, this allows us to highlight how the exact source of uncertainty is relevant for the nature of the optimal agreement.

We start by describing the supply structure in the Home country. The numeraire good is produced one-for-one from labor, with the supply of labor large enough to ensure strictly positive production at all times; therefore the equilibrium wage is equal to one. Each non-numeraire good $j \in \{1, 2\}$ is produced from labor according to the concave production function $X_j = f_j(L_j)$ with $f_j' > 0$ and $f_j'' < 0$, where X_j is the production of good j and L_j is the labor employed in the production of good j . With q_j denoting the producer price for good j and with the wage fixed at one, the supply and profit functions for good j can then be expressed as increasing functions of q_j , and we denote these functions by $X_j(q_j)$ and $\Pi_j(q_j)$ respectively. We assume a similar supply structure for the Foreign country, and let asterisks denote Foreign variables: $X_j^* = f_j^*(L_j^*)$ with $f_j^{*'} > 0$ and $f_j^{*''} < 0$, with associated supply and profit functions given by $X_j^*(q_j^*)$ and $\Pi_j^*(q_j^*)$ respectively.

As noted above, we allow for the possibility of a (positive) production externality. We assume that the externality is linear in aggregate domestic production, enters directly and separably into the representative citizen's utility, and does not cross borders. Producers ignore the effects of their production on the level of aggregate production, and so the externality does not affect supply functions.⁵ Also, we assume that the production of good 1 generates an externality only in Home, and good 2 only in Foreign: hence, the value of the production externality in Home is $\sigma_1 X_1$, while in Foreign it is $\sigma_2^* X_2^*$, with the parameters σ_1 and σ_2^* (defined positively) capturing the strength of the production externality in each country.⁶

In each country, the representative citizen's utility function is linear in the numeraire good

⁵See Markusen (1975) and Ederington (2001) for analogous representations of production externalities.

⁶The assumption of no cross-border externalities is substantive, as it allows us to focus on the traditional (terms-of-trade) motive for negotiating trade agreements. In the concluding section we briefly discuss other (non terms-of-trade) motives for international agreements. On the other hand, our assumption that externalities are experienced only by the importing country does not play a critical role in our results, but seems natural in light of the focus on import-sector intervention that we introduce below.

and separable in the non-numeraire goods. We also allow for the possibility of a (negative) consumption externality. In analogy with the production externality described above, we assume that the consumption externality is linear in aggregate domestic consumption and does not cross borders. Also, as with the production externality, we assume that consumption of good 1 generates an externality only in Home, and good 2 only in Foreign.

Formally, the representative citizens of the two countries enjoy the following utility:

$$U = c_0 + \sum_{j=1}^2 u_j(c_j) - \gamma_1 C_1, \quad U^* = c_0^* + \sum_{j=1}^2 u_j^*(c_j^*) - \gamma_2^* C_2^*,$$

where c_j and C_j denote respectively individual and aggregate consumption of good j . The parameters γ_1 and γ_2^* (defined positively) capture the strength of the consumption externality in each country. Consumers ignore the effects of their individual consumption on aggregate consumption, so the externality does not affect demand functions. We assume that the subutility functions are concave, so that the implied Home and Foreign demands are decreasing functions of the Home and Foreign consumer prices p_j and p_j^* , respectively. We let $D_j(p_j)$ and $D_j^*(p_j^*)$ denote the Home and Foreign demands. Assuming that the population in each country is a continuum of measure one, the consumer surplus associated with good j in Home and Foreign respectively is $\Gamma_j(p_j) = u_j(D_j(p_j)) - p_j D_j(p_j)$ and $\Gamma_j^*(p_j^*) = u_j^*(D_j^*(p_j^*)) - p_j^* D_j^*(p_j^*)$.

We assume that each government can intervene only in its import sector, but within this sector we allow each government to use a pair of instruments, namely, an import tariff (τ) and a domestic production subsidy (s). Both instruments are expressed in specific terms. While we could also introduce a consumption tax, we choose to postpone the treatment of consumption taxes until section 3, when we consider the rationale for a National Treatment clause. We note, though, that τ and s together already comprise a complete set of taxes for the import sector.

At this point we impose a strong symmetric structure on the model: we assume that the two non-numeraire sectors are mirror-images of each other. This allows us to focus on a single sector and drop subscripts from now on. We focus on sector 1, where Home is the natural importer, but it should be kept in mind that in the background there is a mirror-image sector with identical equilibrium conditions, except that the two countries' roles are reversed. The symmetry of the model is inessential, and could be relaxed at the cost of extra notation.

Throughout the paper we focus on non-prohibitive levels of government intervention. In the sector under consideration, due to the absence of taxation by the Foreign government, Foreign producer and consumer prices are equalized, or $q^* = p^*$. In addition, for a firm in Foreign to

sell in both countries, it must receive the same price for sales in Foreign that it receives after taxes for sales in Home, or $p^* = p - \tau$. And finally, the relationship between the Home producer price and the Home consumer price is given by $q = p + s$.

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

$$\begin{aligned} p &= p^* + \tau, \text{ and} \\ q &= p^* + \tau + s. \end{aligned} \tag{2.1}$$

The arbitrage relationships in (2.1) describe the two central price wedges in the model; the first is the wedge between the Home consumer price and the Foreign price (equal to τ), and the second is the wedge between the Home producer price and the Foreign price (equal to $\tau + s$).⁷

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), determines the three market clearing prices as functions of τ and s : $p(\tau, s)$, $q(\tau, s)$ and $p^*(\tau, s)$. At the market clearing prices, Home import volume, $M \equiv D - X$, is equal to Foreign export volume, $E^* \equiv X^* - D^*$. Finally, using $p(\tau, s)$, $q(\tau, s)$ and $p^*(\tau, s)$, we may define economic magnitudes directly as functions of policies. With a slight abuse of notation, we define:

$$\begin{aligned} D(\tau, s) &\equiv D(p(\tau, s)), \quad X(\tau, s) \equiv X(q(\tau, s)), \quad M(\tau, s) \equiv D(\tau, s) - X(\tau, s), \\ \Gamma(\tau, s) &\equiv \Gamma(p(\tau, s)), \quad \Pi(\tau, s) \equiv \Pi(q(\tau, s)); \end{aligned}$$

and similarly for the Foreign country:

$$\begin{aligned} D^*(\tau, s) &\equiv D^*(p^*(\tau, s)), \quad X^*(\tau, s) \equiv X^*(p^*(\tau, s)), \quad E^*(\tau, s) \equiv X^*(\tau, s) - D^*(\tau, s), \\ \Gamma^*(\tau, s) &\equiv \Gamma^*(p^*(\tau, s)), \quad \text{and } \Pi^*(\tau, s) \equiv \Pi^*(p^*(\tau, s)). \end{aligned}$$

Note that $M(\tau = 0, s = 0) > 0$ under our assumption that the Home country is a natural importer of the good under consideration.

We assume that each government maximizes the welfare of its representative citizen. Since the welfare function is separable across sectors, we can focus again on sector 1. In this sector,

⁷The relationships in (2.1) also confirm that τ and s together comprise a complete set of taxes for the import sector: as is well known, an import tariff acts as both a tax on consumption and a subsidy to producers of the import-competing good, and together with a production subsidy the consumer and producer margins can be independently targeted with the two instruments.

Home welfare can be written as the sum of consumer surplus, profits, net revenue (i.e. revenue from the import tariff τ minus expenditures on the production subsidy s), and the valuation of the externalities. Therefore we can write the Home government's objective as:

$$W(\tau, s) \equiv \Gamma(\tau, s) + \Pi(\tau, s) + \tau \cdot M(\tau, s) - s \cdot X(\tau, s) + \sigma \cdot X(\tau, s) - \gamma \cdot D(\tau, s).$$

Recalling that in the sector under consideration the Foreign country has no externalities and no policy instruments of its own, Foreign welfare is the sum of consumer surplus and profits:

$$W^*(\tau, s) \equiv \Gamma^*(\tau, s) + \Pi^*(\tau, s).$$

Notice that, as can be confirmed from the definitions of $\Gamma^*(\tau, s)$ and $\Pi^*(\tau, s)$, Home's policies affect Foreign welfare only through the *terms of trade* p^* .

2.1. The efficient policies and the noncooperative equilibrium

We first derive the globally efficient policies, which we define as the policies that maximize the sum of Home and Foreign payoffs:⁸

$$W^G(\tau, s) \equiv W(\tau, s) + W^*(\tau, s).$$

We assume that both $W(\tau, s)$ and $W^G(\tau, s)$ are concave in τ and s . It is direct to verify that the efficient levels of τ and s , which we denote by τ^{eff} and s^{eff} , are respectively given by

$$\begin{aligned} \tau^{eff} &= \gamma, \text{ and} \\ s^{eff} &= \sigma - \gamma. \end{aligned} \tag{2.3}$$

Hence, efficient policy combinations ensure that the relevant price wedges reflect the externalities. In particular, as a comparison of (2.1) and (2.3) confirm, the wedge between the Home consumer price and the Foreign price (τ) is equal to the consumption externality γ (Pigouvian consumption tax), and the wedge between the Home producer price and the Foreign price ($s + \tau$) is equal to the production externality σ (Pigouvian production subsidy).

⁸In our symmetric setting, it is natural to define efficiency in this way. Recall that there is another sector that mirrors exactly the one under consideration, and in which Foreign is the importer. Therefore, a combination of policies that is Pareto-efficient and gives the same welfare to the two countries must maximize the sum of Home and Foreign payoffs in each sector (with the uncertainty that we introduce in the next section, our assumption of ex-ante symmetry across sectors need not imply ex-post symmetry, but in the presence of uncertainty the relevant notion of efficiency is defined according to ex-ante welfare, and the same statement applies). More generally, this notion of efficiency would also be appropriate in asymmetric settings, provided that international lump sum transfers were available.

Next we turn to the noncooperative equilibrium policies, which we take to represent the choices made in the absence of an agreement. With the Foreign government passive (in the sector under consideration), the Home government's optimal unilateral⁹ policies are defined by

$$\begin{aligned}\frac{dW(\tau, s)}{d\tau} &= 0 \implies \gamma + \frac{E^*}{E^{*'}} + \frac{X'}{D' - X'} \cdot [s + \gamma - \sigma] - \tau = 0, \text{ and} \\ \frac{dW(\tau, s)}{ds} &= 0 \implies \sigma - \gamma + \frac{E^{*'}}{E^{*'} - D'} \cdot \left[\gamma + \frac{E^*}{E^{*'}} - \tau\right] - s = 0,\end{aligned}\tag{2.4}$$

where here and throughout, a prime denotes the derivative of a function with respect to the relevant price. The first condition in (2.4) defines the optimal unilateral choice of τ given s , which we denote $\tau^R(s)$, and the second condition in (2.4) defines the optimal unilateral choice of s given τ , denoted $s^R(\tau)$.

>From the above system we may derive the Home government's noncooperative equilibrium policies, which we denote by τ^N and s^N :

$$\begin{aligned}\tau^N &= \gamma + \frac{p^*}{\eta^*}, \text{ and} \\ s^N &= \sigma - \gamma,\end{aligned}\tag{2.5}$$

where $\eta^* \equiv \frac{p^* E^{*'}}{E^*}$ is the elasticity of Foreign export supply (evaluated at the optimal unilateral policies). Recalling the relationships in (2.1), it is apparent from (2.5) that in the noncooperative equilibrium the Home country employs τ and s to efficiently address the externalities, and then applies its traditional (Johnson, 1953-54) “optimal tariff” – the inverse of the Foreign export supply elasticity – and thereby exploits its monopoly power over the terms of trade (p^*).¹⁰

Notice from (2.3) and (2.5) that the expressions for the efficient and noncooperative levels of s are the same, and that it is only the optimal tariff motivation (as contained in the term p^*/η^*) that drives a wedge between τ^N and τ^{eff} . Therefore, the potential gains from contracting in this setting arise entirely from the ability to control the incentive to utilize import taxes as a means of reducing import volume to manipulate the terms of trade. As a consequence of this feature – which is quite general, as argued in Bagwell and Staiger (2001) – we will refer to international agreements as “trade agreements,” even though they may impose constraints beyond the choice of import taxes, because they represent attempts to solve what is evidently at its core a trade – and trade policy – problem.

⁹We use interchangeably the words “noncooperative” and “unilateral.”

¹⁰It is direct to verify that our focus on non-prohibitive levels of government intervention in effect places an upper limit on the magnitude of the externality parameters σ and γ .

2.2. Uncertainty

We consider three possible sources of uncertainty: the production externality (σ), the consumption externality (γ) and the level of domestic demand. To capture domestic demand shocks, we parametrize the Home demand function (with a slight abuse of notation) by $D(p; \alpha)$, where $D_\alpha > 0$, so that a higher α corresponds to a higher-demand state.

Uncertainty about σ and γ can be interpreted as uncertainty about the efficiency rationale for policy intervention, while shocks to α can be interpreted as shocks to the underlying trade volume.¹¹ Focusing on uncertainty in σ , γ and α while abstracting from other sources of uncertainty helps to illustrate some general principles for understanding how the optimal agreement depends on the source of uncertainty. We sometimes refer to σ , γ and α as the state-of-the-world variables, or simply the “state” variables. For the moment we do not need to impose any structure on the distribution of these variables; we will do so at a later stage.

We consider the following simple timing: (1) the agreement is drafted: (2) uncertainty is resolved; and (3) policies are chosen subject to the constraints set by the agreement. Implicit in this timing is the assumption that agreements are perfectly enforceable: in this paper we abstract from issues of self-enforcement of the agreements.

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

2.3. The costs of contracting

Before we formalize the costs of contracting, we need to specify what type of contracts we will consider. Throughout the paper we focus on *instrument-based* agreements, i.e. agreements that impose (possibly contingent) constraints on policy instruments. In the concluding section we briefly discuss the possibility of *outcome-based* agreements, i.e. agreements that impose constraints on equilibrium outcomes such as trade volumes.¹²

As a first step we consider a relatively narrow class of agreements: we consider agreements that impose separate equality constraints on τ and s . To be concrete, we allow for clauses of the type $(\tau = \gamma)$ or $(s = 10)$, but not for clauses of the type $g(\tau, s) = 0$ or for inequality

¹¹We could alternatively consider a shock that shifts the domestic supply function, but the qualitative results would not change.

¹²We also abstract from agreements that are based on both instruments and outcomes, in the sense that they constrain directly the relationship between policy instruments and equilibrium outcomes, such as for example an agreement that constrains τ to be a direct function of the import volume M .

constraints of the type ($\tau \leq 1$).¹³ We label this class of agreements \mathcal{A}_0 . In later sections we consider broader classes of agreements.

We formalize the contracting costs associated with a trade agreement in a very stylized way. Our central assumption is that these costs are higher, the more policy instruments the agreement involves, and the more contingencies it includes.

More specifically, we assume that there are two kinds of contracting costs: the costs of including *state* variables in the agreement – that is, the random variables σ , γ and α – and the costs of including *policy* variables – that is, τ and s . 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 verifying its value ex-post.¹⁴ 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.

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 of “recall.”¹⁵ Summarizing, the cost of writing an agreement is $C = c_s \cdot n_s + c_p \cdot n_p$, where n_s and n_p are, respectively, the number of state and policy variables in the agreement. We could allow C to be a more general increasing function of n_s and n_p , but we choose the linear specification to simplify the analysis and the exposition of our results.¹⁶

¹³We consider agreements that impose inequality constraints of the type ($\tau \leq 1$) in section 4. When there is significant uncertainty, a noncontingent contract of the type $g(\tau, s) = 0$ may do better than a noncontingent contract that pins down τ and/or s separately, because the former contract type introduces some discretion. This has the flavor of an outcome-based contract, which we discuss in the concluding section.

¹⁴The interpretation of contracting costs as verification costs is “tight” only if the court automatically verifies (ex post) the values of the variables included in the contract. In the WTO, the Trade Policy Review Mechanism provides periodic reviews of the member countries’ trade policies. But a more thorough verification process in the WTO occurs 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.

¹⁵Our main results (Propositions 1-5) do not depend on the assumption of no recalling costs. If we allowed for recalling costs the only change would be that, in the parametrized specification of section 2.5, we would need to consider agreements where one instrument is contingent but the other is not (which we can ignore under the no-recalling-cost assumption), but again the main insights of the analysis would not change.

¹⁶Also, it might be reasonable to suppose that it is more costly to contract over internal measures (s) 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.

A couple of examples may be useful to illustrate our assumptions on contracting costs:

Ex. 1: The agreement $\{\tau = 3\}$ specifies a rigid commitment for the tariff, and costs c_p .

Ex. 2: The agreement $\{\tau = \gamma, s = 5\}$ specifies a state-contingent commitment for the tariff and a rigid commitment for the subsidy, and costs $2c_p + c_s$.

Overall, our approach to modeling the costs of contracting has advantages and also limitations. On the plus side, our approach preserves tractability while adding some generality relative to other approaches in the literature.¹⁷ On the minus side, our approach abstracts from some potentially important considerations: for example, we assume that the number of state variables n_s summarizes the costs of state-contingency, but in reality this cost might depend as well on the “coarseness” of the contingencies (e.g. it might be easier to verify a clause like $(\tau = 0 \text{ if } \gamma \leq 1)$ then a clause $\tau = \gamma$). On balance, however, we believe that the basic feature that contracting costs are increasing in the number of state variables and policies included in the agreement is likely to be preserved in most reasonable models of these costs, and for this reason we believe that our approach provides a good starting point for the analysis of trade agreements as endogenously incomplete contracts.

2.4. Optimal Agreements

To characterize the optimal agreement, we need to introduce some definitions and notation. First, we refer to the *efficiently-written first-best* agreement as the least costly among the agreements that implement the first best outcome. We label this simply the $\{FB\}$ agreement. In a similar vein, we 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 “net global welfare”), that is $\omega \equiv \Omega - C$.

The first step is to derive the $\{FB\}$ agreement. Recall that the first-best policies are defined by (2.3). We can conclude that an agreement of the form $\{\tau = \gamma, s = \sigma - \gamma\}$ achieves the first-best outcome. This agreement has $n_p = 2$ and $n_s = 2$ and therefore costs $2c_p + 2c_s$. Moreover, it is clear that the first-best outcome cannot be implemented with an agreement that costs less than $2c_p + 2c_s$, and so $\{\tau = \gamma, s = \sigma - \gamma\}$ is indeed the $\{FB\}$ agreement in the class \mathcal{A}_0 .

The $\{FB\}$ agreement yields net global welfare equal to $\Omega^{FB} - (2c_p + 2c_s)$, where Ω^{FB} denotes

¹⁷For example, Battigalli and Maggi (2002) associate a cost c with each “primitive sentence” included in the contract, and the analogue in our setting would be to associate a cost c with each state variable or policy included in the contract. Under this analogy, the form of contracting costs adopted by Battigalli and Maggi is a special case of our approach in which $c_s = c_p$.

the gross global welfare implied by the first-best policies. Clearly, when contracting costs are sufficiently small the $\{FB\}$ agreement is optimal; and if they are sufficiently high, the empty agreement (which costs nothing and yields the noncooperative equilibrium outcome) is optimal. The interesting question is what happens between these two extremes: What is the optimal way to save on contracting costs?

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 state variables are missing from the agreement; and *discretion*, which occurs when policy variables are missing from the agreement. Thus, for example, the agreement $\{\tau = 0, s = 5\}$ is fully rigid; the agreement $\{s = g(\sigma, \gamma, \alpha)\}$ features discretion over τ , and the agreement $\{\tau = 3\}$ is both rigid and discretionary (over s).¹⁸ With these notions of rigidity and discretion, the question we posed above can be rephrased as: What is the optimal combination of rigidity and discretion?

Given that we have two policy variables (τ and s) and three state variables (σ , γ and α), and that specifying each of these variables in the contract is costly, in principle there are many types of contract that we should consider. Indeed, a complete characterization of the optimal contract will have to wait until we impose more structure on the stochastic environment and on the demand and supply functions (which we do in section 2.5). Nonetheless, we are able in this general setting to derive a number of insights about the structure of the optimal agreement.

Our first result (proved in the Appendix) records an important feature of the trade-agreement contracting environment under study: if an agreement is to achieve any improvement over the noncooperative equilibrium, it must constrain import taxes. More formally:

Proposition 1. *An agreement that constrains the subsidy s (even in a state-contingent way) while leaving the import tariff τ to discretion cannot improve over the noncooperative equilibrium, and therefore cannot be an optimal agreement.*

At a broad level, the intuition behind Proposition 1 is very simple, and reflects a kind of “targeting principle” logic (Bhagwati and Ramaswami, 1963 and Johnson, 1965): contracting over s alone is suboptimal because, as we have emphasized in section 2.1, the inefficiency in the noncooperative equilibrium concerns τ , not s .

To develop a more precise understanding of this result, consider an agreement that imposes a small exogenous change in s starting from the noncooperative equilibrium. This triggers a

¹⁸Notice that rigidity and discretion do not necessarily imply a loss of gross surplus relative to the first best. For example, the $\{FB\}$ contract is not contingent on the demand parameter α , so it is rigid with respect to α .

change in the Home government's choice of τ , and in particular we show in the Appendix that τ adjusts to the exogenous change in s so as to maintain p^* at the noncooperative level. Recalling that Home's policies affect Foreign welfare only through the terms of trade p^* , this implies that Foreign welfare is unchanged; and since the imposition of a constraint on s can only reduce Home welfare, global welfare goes down as a consequence. Thus a small exogenous change in s cannot improve over the noncooperative equilibrium.¹⁹

In a world of costless contracting, the result highlighted in Proposition 1 would be irrelevant, because if agreements were costless they would always be written in a way that placed the needed constraints on all policy instruments. But with costly contracting this result gains relevance. In particular, as Proposition 1 indicates, any (nonempty) agreement must include commitments over import taxes, and should only introduce commitments over domestic policies if it is optimal to make the agreement more complete. Notice, too, that this prediction does not rely on an assumption that embodies the commonly-held view that border measures are more transparent than domestic policies and are therefore less costly to contract over, an assumption that would only reinforce this prediction. Instead, the prediction arises as a consequence of the nature of the inefficiency that governments attempt to address with their agreement.²⁰

Given the result of Proposition 1, there are two remaining questions that must be answered in designing the optimal agreement: (i) whether or not s should also be constrained by the agreement; and (ii) whether or not the agreement should be state-contingent and, if so, what state variables should be included. We consider each of these remaining questions in turn.

To answer the first of these questions, it is helpful to begin by recording the expression for $s^R(\tau)$, the optimal unilateral choice of s given τ , which is the choice of s that would be made if τ is constrained by the agreement but s is left unconstrained. This choice solves $dW(\tau, s)/ds = 0$, and using (2.4), it is direct to derive the following expression:

$$s^R(\tau) = (\sigma - \gamma) - \left(\tau - \gamma - \frac{p^*}{\eta^*}\right) \cdot \frac{E^{*'}}{E^{*'} - D'}, \quad (2.6)$$

¹⁹We note that the result stated in Proposition 1 is distinct from and not contradictory to Copeland's (1990) result that negotiating over tariffs can generate surplus even if other instruments are non-negotiable. Copeland's result implies that contracting over tariffs is *sufficient* to generate some surplus, whereas Proposition 1 implies that it is also *necessary*. As we discuss in the concluding section, this result must be qualified when political economy forces are introduced, but it is still the case that constraining s alone is suboptimal, at least provided that political economy forces are not too strong.

²⁰In particular, this prediction reflects the structure of the terms-of-trade driven Prisoners' Dilemma that governments attempt to solve in this setting; it is not clear that it would arise as naturally under alternative theories of trade agreements such as the commitment theory (see Bagwell and Staiger, 2002, Ch.2, for a review of these theories).

where all right-hand-side magnitudes are evaluated at τ and $s^R(\tau)$.²¹ On the other hand, the efficient level of s conditional on τ , which we denote by $s^{eff}(\tau)$, solves $dW^G(\tau, s)/ds = 0$, and it is direct to verify that

$$s^{eff}(\tau) = (\sigma - \gamma) - (\tau - \gamma) \cdot \frac{E^{*'}}{E^{*'} - D'}, \quad (2.7)$$

where all right-hand-side magnitudes are evaluated at τ and $s^{eff}(\tau)$. Notice that the only difference between $s^R(\tau)$ and $s^{eff}(\tau)$ is that the noncooperative tariff level, $\gamma + \frac{p^*}{\eta^*}$, is replaced by the efficient tariff level, γ . It is straightforward to show that $s^{eff}(\tau) < s^R(\tau)$ for all τ .

Clearly, in a costly contracting environment a key ingredient in answering the question whether or not s should also (in addition to τ) be constrained by the agreement is the extent to which $s^R(\tau)$ implies a loss in global surplus relative to $s^{eff}(\tau)$. For a given state of the world, this loss (defined positively) can be written as:²²

$$W^G(s^{eff}(\tau), \tau) - W^G(s^R(\tau), \tau) = - \int_{s^{eff}(\tau)}^{s^R(\tau)} W_s^G(s, \tau) ds. \quad (2.8)$$

If this loss is sufficiently small for all relevant values of the tariff τ and of the state variables, then it is optimal to omit s from the agreement, since in this case the savings in contracting costs (which are at least c_p , and which may be higher if s is specified in a state-contingent way) will exceed the cost of leaving discretion over s .²³

Under what conditions, then, will the expression in (2.8) be small? Recalling that $s^{eff}(\tau) < s^R(\tau)$ and noting that $W_s^G(s^{eff}(\tau), \tau) = 0$ and that W^G is concave in s , a sufficient condition for the expression in (2.8) to be small is that $|W_s^G(s^R(\tau), \tau)|$ is small. It is direct to verify that $W_s^G(s^R(\tau), \tau) = W_s(s^R(\tau), \tau) + \frac{\partial p^*}{\partial s} \cdot M$. Noting that $W_s(s^R(\tau), \tau) = 0$ and after some manipulation, we therefore have:

$$|W_s^G(s^R(\tau), \tau)| = -\frac{\partial p^*}{\partial s} \cdot M = \frac{1}{\frac{1}{X'}(\frac{\eta^*}{p^*} + \frac{|D'|}{M}) + \frac{1}{M}} \equiv \mathcal{B}, \quad (2.9)$$

where all magnitudes in \mathcal{B} are evaluated at τ and $s^R(\tau)$. Intuitively, as (2.9) indicates, $|W_s^G(s^R(\tau), \tau)|$ is just the income gain enjoyed by Home as a result of the terms-of-trade movement triggered by a small rise in s beginning from $s^R(\tau)$, and hence the cost of leaving s to

²¹Note that this expression is valid also if τ is constrained in a contingent way, in which case τ will be a function of (some or all of) the state variables; the same applies to the expression for $s^{eff}(\tau)$ below.

²²For simplicity we do not highlight the state variables in the notation.

²³In the presence of uncertainty, it is the expected value of the loss in (2.8) that is relevant for determining whether or not s should be omitted from the agreement. However, to keep the exposition simple we present a sufficient condition that ensures that this loss is small for each state of the world. Such a condition is stronger than we need, but it is the most transparent.

discretion will be small when the magnitude of this terms-of-trade effect, which can be re-expressed as \mathcal{B} , is small. Hence, we may conclude that the expression in (2.8) is small, and therefore that it is optimal to omit s from the agreement, if \mathcal{B} as defined in (2.9) is small.²⁴

In combination with Proposition 1, the conditions that make \mathcal{B} small provide immediate insight into a number of the key forces that shape the nature of the optimal agreement. Specifically, (2.9) points to three circumstances under which the cost of discretion over s will be small, and hence to circumstances where leaving commitments on s out of the trade agreement is an attractive way to save on contracting costs.

First, \mathcal{B} will be small if p^*/η^* (Johnson’s optimal tariff) is sufficiently small. This describes the “small country” case in which Home has little international monopoly power, and hence little ability to manipulate terms of trade. If countries are sufficiently small in world markets, the cost of leaving s to discretion is small. We refer to this as the *monopoly power* effect.

Second, \mathcal{B} will be small if M is sufficiently low. This describes the case in which Home has little trade volume over which to apply its international monopoly power, and hence gains little from exploiting its ability to manipulate the terms of trade. If the volume of trade is sufficiently low, the cost of leaving s to discretion is small. We refer to this as the *trade volume* effect.

Third, \mathcal{B} will be small if X' is sufficiently low or $|D'|$ is sufficiently high. Recalling that s distorts only the producer margin, while τ distorts both the producer and the consumer margin, this describes the case in which Home’s ability to utilize s rather than τ as an instrument for terms-of-trade manipulation is limited. If the subsidy s is a sufficiently poor substitute for τ as an instrument for manipulating the terms of trade, the cost of leaving s to discretion is small. We refer to this as the *instrument substitutability* effect.

Together with Proposition 1, the monopoly power, trade volume and instrument substitutability effects describe key features of the trade-agreement contracting environment that help to determine whether commitments on subsidies should be included in an optimal agreement. These effects highlight a tradeoff between a direct benefit of leaving s to discretion

²⁴Our discussion in the text abstracts from a technical issue: as \mathcal{B} becomes small, it must be assured that the range of integration in (2.8), $s^R(\tau) - s^{eff}(\tau)$, does not blow up “too fast.” For this reason, some care is required when considering changes in demand/supply functions that drive \mathcal{B} to zero but might also drive $s^R(\tau) - s^{eff}(\tau)$ to infinity. Using (2.6) and (2.7), it can be shown that this is not an issue for changes in M , p^*/η^* or D' , but when considering changes in X' this issue becomes relevant, because if $X'/X \rightarrow 0$ both $s^R(\tau)$ and $s^{eff}(\tau)$ go to infinity. In this case, it suffices to consider the limit of a sequence of supply functions such that $s^R(\tau) - s^{eff}(\tau)$ does not go to infinity. It is easy to show that this is always possible: for example, with linear ($X = \lambda q$) or exponential ($X = \kappa e^{\lambda q}$) supply functions, this problem does not arise when taking the limit as $\lambda \rightarrow 0$.

– namely the associated savings in contracting costs c_p – and the distortions caused by the unilateral choice of s when τ is constrained. Broadly speaking, these effects suggest that leaving a country’s domestic subsidies out of the trade agreement is an attractive way to save on contracting costs if the country has little monopoly power in trade, or if it trades little, or if domestic subsidies are a poor substitute for import tariffs as tools to manipulate terms of trade.

Finally, in addition to the direct savings on the cost of contracting over s , there is also a second, indirect, benefit of leaving s out of the agreement, which applies whenever c_s is sufficiently high that the optimal agreement itself is rigid with respect to γ and/or σ : in this case, leaving s to discretion has the benefit of indirectly introducing state-contingency in the agreement. We refer to this benefit of discretion over s as the *indirect state-contingency* effect.

To illustrate this effect, we shut down the direct benefit of discretion highlighted above by setting $c_p = 0$, and consider the case in which c_s is prohibitively high, so that the optimal agreement is not state-contingent. In this case there are only two types of (nonempty) agreement that could be optimal: an agreement that constrains rigidly τ and s , and one that constrains rigidly only τ . As a comparison of (2.6) and (2.7) confirms, under these circumstances, if s is left to discretion the Home country will distort s to manipulate terms of trade, but the advantage of this discretion is that, like s^{eff} , s will be responsive to any changes in γ and/or σ . And if \mathcal{B} is small, so that the cost of the terms-of-trade manipulation is small, then it is optimal to leave s to discretion (provided there is at least some uncertainty over γ and/or σ).²⁵

Our discussion of whether s should also be constrained by the agreement leads to:

Proposition 2. (a) If $c_p > 0$ and \mathcal{B} is sufficiently small, it is optimal to leave discretion over the subsidy s . (b) If c_s is sufficiently high (so that contingent contracts are suboptimal) and there is some uncertainty over γ and/or σ , then if \mathcal{B} is sufficiently small, it is optimal to leave discretion over s , even if $c_p = 0$.

Proposition 2, together with the analysis that leads up to it, sheds light on the main forces that determine whether to leave domestic subsidies out of the trade agreement. The potential benefits of omitting s from the agreement accrue in the form of direct savings on the costs of contracting over s and the attainment of indirect state contingency in s . The potential costs

²⁵More precisely, the condition is that \mathcal{B} is small for all τ and all states of the world. With a slight abuse of language we omit this qualifier in the statements that follow. Also notice that, if \mathcal{B} is small, the empty agreement could be optimal, but this is consistent with the statement that it is optimal to leave s to discretion.

take the form of terms-of-trade manipulation, and the magnitude of these costs depend on the strength of the market power, trade volume and instrument substitutability effects.

Notice that, at a broad level, these forces suggest a possible explanation for an important aspect of the evolution from GATT to the WTO, namely, the fact that the WTO has introduced a substantial effort to regulate the use of domestic subsidies that was not present in GATT, and is moving toward further constraints on domestic policies more generally.²⁶ Proposition 2 suggests that this evolution could be explained by an increase in trade volumes over time (and the implied rise in \mathcal{B}) which, by raising the cost of discretion, has given rise to the need to constrain subsidies and other domestic policies in the agreement. Similarly, Proposition 2 suggests an interesting cross-country prediction. The essence of low monopoly-power/trade-volume is that a country imports small volume from a relatively elastic source of Foreign export supply, while the essence of low instrument substitutability is that the government has limited domestic policy options at its disposal. Arguably, these conditions (and the implied low level of \mathcal{B}) are most likely to apply to small developing countries, and accordingly Proposition 2 suggests that contracting over domestic policies (such as s) is likely to be more attractive for large developed countries than for small/developing countries: this points to the possible benefits of a kind of “special and differential treatment” rule for small/developing countries when it comes to contracting over domestic policies, especially if the value of indirect state contingency over domestic policies is high in these countries.²⁷

We now turn to the second question posed above, and consider whether or not the agreement should be state-contingent and, if so, what state variables should be included. Intuitively, the cost of specifying a given state variable in the agreement (c_s) must be compared with the benefit of making the agreement contingent on that state variable. The benefit of introducing state variables in the agreement is in turn determined in large part by the degree of uncertainty in the contracting environment. For example, the more uncertain are the state variables γ and/or σ , the more uncertain will be the first-best levels of the policy instruments as given by (2.3), and in general the more beneficial it is to write a state-contingent contract. This is not a surprising statement. But whenever the agreement constrains τ while leaving s to discretion, the model

²⁶During the GATT era, subsidies were subject primarily to the disciplines of countervailing duties and non-violation nullification-or-impairment claims, and the WTO’s Subsidies and Countervailing Measures (SCM) Agreement is a significant strengthening of these disciplines (see Sykes, 2005 and Bagwell and Staiger, 2006).

²⁷In fact, Part VIII of the WTO’s SCM Agreement introduces just such an exemption from subsidy commitments for developing country members. We thank Robert Lawrence for pointing this out.

also suggests a more subtle insight concerning *why* state-contingent tariff commitments may be beneficial, and therefore *which* contingencies to introduce in the agreement.

To develop this last point, we consider an environment in which it is optimal to constrain only τ , not s . Above we presented sufficient conditions for this to be the case. In this setting, as we have observed, the unilateral choice of s will be distorted above s^{eff} as a way to manipulate the terms of trade, but recall as well that this distortion will tend to be more severe if the trade volume is higher, owing to the trade-volume effect highlighted above. But then intuitively, it might be desirable to allow τ to change with α – the trade volume shift parameter – as a way to dampen the trade volume in high-volume states of the world and thereby mitigate the incentive to distort s for terms-of-trade purposes; moreover, a similar observation applies to σ as well, to the extent that changes in σ imply changes in trade volume (through changes in $s^R(\tau)$). We refer to this as the *indirect incentive-management* effect. The interesting point is that in general it can be optimal to make the tariff τ contingent on α and/or σ even though α and σ are *per se* irrelevant for the first-best level of the tariff τ^{eff} (as (2.3) confirms).

Following this logic, it is straightforward to establish the next result:

Proposition 3. *Conditional on the agreement constraining τ but leaving discretion over s , if c_s is sufficiently low then it is optimal to make τ contingent on α and/or σ , even though the first-best level of τ does not depend on α or σ .*

The indirect incentive-management effect that underlies Proposition 3 can give rise to an escape-clause type of agreement: under some conditions the tariff level will be increasing in α , so the agreement will allow the import tariff to rise in states of the world in which the underlying import volume is high, broadly analogous to the escape clause provided in GATT Article XIX.²⁸ This suggests a novel rationale for the desirability of escape clauses in trade agreements: an escape-clause type agreement that makes τ contingent on the import demand level α can be attractive because it provides an indirect means of managing the distortions associated with leaving s to discretion.²⁹

²⁸Proposition 3 establishes conditions under which it is optimal to make τ contingent on α (and/or σ). Less clear is whether τ is increasing in α . The reason is that an increase in α has a direct effect on the cost of discretion through the trade volume, but may also have indirect effects through the slopes of demand and supply functions evaluated at the equilibrium point. In the linear specification of section 2.5, we show that τ is indeed increasing in α , but with general nonlinear demand and supply functions it cannot be guaranteed that the direct effect of a change in α will dominate the indirect effects. The point we emphasize here is that it *may* be optimal to have τ increasing in α , so the model can explain an escape-clause type of agreement.

²⁹We say that an agreement where τ is increasing in α is an “escape-clause type” agreement because there

2.5. A linear specification of the model

Propositions 1 through 3 identify the broad forces that shape the nature of the optimal agreement and illustrate these forces in action. But to fully characterize the optimal agreement and analyze the comparative-statics effect of changes in the economic and contracting environment, we now impose more structure and work with a parametric specification of the model.

Maintaining our focus on sector 1 where Home is the importer (and suppressing subscripts), we suppose that Home's production function takes the form $X = \sqrt{2\lambda L}$, implying the linear supply function $X(q) = \lambda q$ and associated profit function $\Pi(q) = \frac{1}{2}\lambda q^2$, with a similar supply structure in the Foreign country yielding $X^*(q^*) = \lambda^* q^*$ and $\Pi^*(q^*) = \frac{1}{2}\lambda^* q^{*2}$. Analogously on the demand side, we suppose that sub-utility functions are quadratic, so that implied demand functions are linear: $D(p) = \alpha - \beta p$ and $D^*(p^*) = \alpha^* - \beta^* p^*$. All parameters are positive.

We also impose more structure on the stochastic environment. We consider separately two cases: uncertainty in the consumption externality γ ; and uncertainty in α , which notice under our linear specification is the intercept of the domestic demand and import demand functions. There are two reasons for focusing separately on these two sources of uncertainty: the first reason is tractability; the second reason is that it allows us to highlight that the implications of uncertainty depend in crucial ways on the *source* of the uncertainty; and as we will explain shortly, uncertainty in γ and uncertainty in α represent two polar cases. At the end of the section we then briefly consider the case of uncertainty in the production externality (σ) which, as we will argue, shares some features with each of the two previous cases. Viewed together, our separate analyses of uncertainty over these three state variables illuminate the manner in which the source of the uncertainty faced by governments is relevant for the optimal agreement.

Finally, a convenient way to characterize the optimal agreement is to track how it changes as the general level of contracting costs increases. We consider a proportional increase in the contracting costs (c_p, c_s) . To express our results in a simple comparative-statics fashion, we let

are some important features of GATT Article XIX that are not captured by this kind of contract. For instance, Article XIX links the possibility of tariff increases directly to increases in import volume (rather than indirectly through changes in underlying market conditions such as α), a contracting possibility we abstract from in this paper (see note 12). Moreover, Article XIX includes an "injury" test, which has no counterpart in our model (but we note that an explanation for the injury test is also lacking in other theoretical interpretations of the escape clause, such as Bagwell and Staiger, 1990). And finally, under Article XIX a country is *allowed* to raise its tariff in case of an import surge, whereas the contract considered here technically leaves no discretion on τ . This feature, however, can be captured by our model in a straightforward manner: as we argue in section 4, imposing the equality constraint $\tau = \tau(\cdot)$ is equivalent to imposing the inequality constraint $\tau \leq \tau(\cdot)$. Under the latter, when α is higher the government is allowed to raise τ up to a higher level, but is not forced to do so.

$c_p \equiv c$, 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 contracting costs. In much of the analysis to follow, we keep k fixed and consider changes in c .

Uncertainty about γ .

We first focus on the case where only the consumption externality γ is uncertain. For the sake of expositional simplicity, we set the production externality σ to zero, and we assume that γ can take two possible values with equal probability: a high realization $\bar{\gamma} + \Delta_\gamma$ and a low realization $\bar{\gamma} - \Delta_\gamma$, with $\Delta_\gamma > 0$. Let $\bar{\alpha}$ denote the deterministic value of α .

The $\{FB\}$ agreement in this setting is $\{\tau = \gamma; s = -\gamma\}$. This agreement has $n_s = 1$ and $n_p = 2$ and therefore costs $(2 + k)c$. Clearly, if c is small enough $\{FB\}$ is optimal. From this starting point, we seek to characterize the optimal agreement as a function of c : What is the optimal way of restructuring the agreement as c rises from zero?

By Proposition 1, we can ignore agreements that constrain s but not τ . Also, our assumption of no recalling costs implies that we can ignore agreements where one policy instrument is contingent but the other one is not (e.g. $\{\tau(\gamma), s\}$). Therefore we only have three kinds of agreements to consider, in addition to $\{FB\}$ and $\{\emptyset\}$: (i) agreements that constrain τ and s in a rigid fashion, which we denote $\{\tau, s\}$; (ii) agreements that constrain τ as a function of γ , which we denote $\{\tau(\gamma)\}$; and (iii) agreements that constrain τ in a rigid fashion, which we denote $\{\tau\}$. Notice that the agreement $\{\tau, s\}$ is rigid, the agreement $\{\tau(\gamma)\}$ features discretion (over s), and the agreement $\{\tau\}$ is both rigid and discretionary. We have the following result:

Remark 1. *When only γ is uncertain, there exist 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 $\{\tau, s\}$ for $c \in (c_1, c_2)$;*
- (c) *of the form $\{\tau(\gamma)\}$ for $c \in (c_2, c_3)$;*
- (d) *of the form $\{\tau\}$ for $c \in (c_3, c_4)$; and*
- (e) *the empty agreement for $c > c_4$.*

Moreover, either $c_2 = c_1$ or $c_3 = c_2$ (or both).

The proof of Remark 1 is in the Appendix. There are two features of Remark 1 that deserve emphasis. First, Remark 1 confirms and expands an insight that was anticipated by Proposition

1: the subsidy s tends to be more discretionary than the import tariff τ . More specifically, for a range of contracting costs it may be optimal to contract over τ while leaving s to discretion, but it is never optimal to contract over s and leave τ to discretion.

The second feature that deserves emphasis is the “complementary slackness” property reported at the end of Remark 1: $\{\tau, s\}$ and $\{\tau(\gamma)\}$ cannot *both* be part of the optimal sequence of agreements as c increases. This feature reflects the non-separability of our contracting problem across instruments. Thus, Remark 1 indicates that the optimal way to save on contracting costs as c rises can be either to first introduce rigidity and then consider adding discretion, or to first introduce discretion and then consider adding rigidity; we will shortly examine how the choice between these two possibilities is determined by the underlying economic environment.³⁰

The source of the complementary slackness property in Remark 1 lies ultimately in the indirect state-contingency effect highlighted in Proposition 2(b). This can be understood in two steps. First, the indirect state-contingency effect implies that *rigidity and discretion are complementary*, in the sense that the cost (in terms of lost surplus relative to $\{FB\}$) of discretion and rigidity when they occur together (as in $\{\tau\}$) is *lower* than the sum of the costs when they occur separately (as in, respectively, $\{\tau(\gamma)\}$ and $\{\tau, s\}$); this derives from the fact that when the uncertainty concerns γ , omitting s from a rigid agreement introduces valuable indirect state-contingency, because γ is relevant for s^{eff} . To see the second step, suppose, to fix ideas, that $\{\tau(\gamma)\}$ is less costly than $\{\tau, s\}$ (i.e. $k < 1$); then it is intuitive (and easily shown) that a necessary condition for both $\{\tau(\gamma)\}$ and $\{\tau, s\}$ to be optimal for some c is that Ω must increase in a concave manner as we move from $\{\tau\}$ to $\{\tau(\gamma)\}$ to $\{\tau, s\}$ to $\{FB\}$. But the complementarity between rigidity and discretion implies that this concavity cannot be satisfied, and hence $\{\tau, s\}$ and $\{\tau(\gamma)\}$ cannot both be optimal for some c .

In light of the complementary slackness condition, we may now ask: As c rises from zero, when is it optimal to first economize on contracting costs by introducing rigidity ($c_1 < c_2$), and when by first introducing discretion ($c_2 < c_3$)? The answer to this question depends on a comparison between the costs of rigidity and the costs and benefits of discretion (over s).

Consider first the costs and benefits of discretion. These can be understood through the market power, trade volume, instrument substitutability and indirect state-contingency effects identified in Proposition 2. In our linear model, both M and p^*/η^* are increasing in the level

³⁰This feature stands in marked contrast to the findings of Battigalli and Maggi (2002): there, separability ensures that, as contracting costs rise, first the optimal contract becomes rigid, and then discretion is introduced.

of import demand $\bar{\alpha}$, so the cost of discretion is increasing in $\bar{\alpha}$ through both the trade volume and monopoly power effects. On the other hand, according to the instrument substitutability effect, a low level of β (the slope of domestic demand) should imply a high cost of discretion, while the cost of discretion should be low when λ (the slope of domestic supply) is low. And finally, a low level of Δ_γ (uncertainty in γ) should diminish the appeal of discretion, because the added benefit of discretion as embodied in the indirect state-contingency effect is then small.

Consider next the cost of rigidity. This is determined largely by the degree of uncertainty: a rigid agreement “gets it right” only on average, and therefore when the environment is more uncertain (i.e. when Δ_γ is higher) the cost of rigidity (with or without discretion) is higher.

Consistent with these intuitions, we find that as c increases: it is optimal to first introduce rigidity ($c_1 < c_2 = c_3$) when $\bar{\alpha}$ is sufficiently high and/or Δ_γ is sufficiently low; it is optimal to first introduce discretion ($c_1 = c_2 < c_3$) when the supply slope λ is sufficiently low; and it may be optimal to introduce rigidity, but discretion is never optimal ($c_2 = c_3 = c_4$) if the demand slope β is sufficiently low.

Thus far we have discussed how various parameters affect the optimal sequence of agreements as c varies, but we have not described the effects of changing a parameter while holding constant all other parameters (including c). We turn now to two of the more illuminating comparative-static results, namely, those for $\bar{\alpha}$ and Δ_γ .

We start with changes in $\bar{\alpha}$. As noted above, an increase in $\bar{\alpha}$ increases the cost of discretion through both the trade volume and market power effects, so it should lead to less discretion in the optimal agreement. Also, $\bar{\alpha}$ does not affect the cost of rigidity, and so $\bar{\alpha}$ does not affect *directly* the degree of rigidity in the optimal agreement.³¹ Nevertheless, the complementarity between rigidity and discretion suggests that, as $\bar{\alpha}$ increases and discretion falls, rigidity may *also* fall, as in the movement from $\{\tau\}$ to $\{FB\}$. The following remark confirms this intuition:

Remark 2. *As the import demand level $\bar{\alpha}$ increases: (i) The optimal degree of discretion decreases (weakly), in the sense that the number of policy instruments specified in the optimal agreement increases (weakly); (ii) The optimal degree of rigidity decreases (weakly).*

Remark 2(i) reflects the impact of trade volume and monopoly power on the cost of discretion. As we noted earlier in our discussion after Proposition 2, this prediction resonates with the

³¹Indeed it can be shown that, as $\bar{\alpha}$ increases, the optimal agreement never switches from $\{\tau\}$ to $\{\tau(\gamma)\}$ or from $\{\tau, s\}$ to $\{FB\}$, and so it is never the case that a change in $\bar{\alpha}$ changes the degree of rigidity without affecting the degree of discretion.

GATT/WTO's evolution away from an initial emphasis on border measures towards an increasing focus on internal/domestic policies.³² Remark 2(ii) is a consequence of the complementarity between rigidity and discretion. This implies that rising trade volumes and monopoly power (higher $\bar{\alpha}$) may make it worthwhile to add contingencies to the agreement, but only because it is now worthwhile to contract over domestic policies, and the value of adding state-contingencies is enhanced as a result.

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

Remark 3. *As the degree of uncertainty Δ_γ increases, the optimal agreement may switch from a rigid agreement to a contingent agreement, but not vice-versa.*

By itself, this result is not particularly surprising: it seems inevitable that increasing uncertainty should reduce the attractiveness of rigid agreements. But there is also a more subtle feature of this result, which is that it concerns uncertainty over a state variable that is relevant for the level of the efficient tariff τ^{eff} . As we demonstrate below, the effects of increasing uncertainty over state variables (such as α) that are not relevant for the level of τ^{eff} can be very different.³³

Uncertainty about α .

Above we examined a stochastic environment where uncertainty concerns only a state variable that affects the first-best levels of s and τ . We now suppose that uncertainty concerns only a state variable (α) that has no impact on the first-best level of either instrument. We assume that α can take two possible values with equal probability: $\bar{\alpha} + \Delta_\alpha$ and $\bar{\alpha} - \Delta_\alpha$, with $\Delta_\alpha > 0$. We let $\bar{\gamma}$ denote the deterministic value of γ , and for simplicity we continue to set $\sigma \equiv 0$.

In this environment, the $\{FB\}$ agreement takes the form $\{\tau = \bar{\gamma}; s = -\bar{\gamma}\}$. Notice that the $\{FB\}$ agreement is no longer state contingent, because it does not depend on the uncertain parameter α , and $\bar{\gamma}$ is a deterministic value, and so the cost of the $\{FB\}$ agreement is now $2c$.

³²In our discussion after Proposition 2 we emphasized the impacts of rising trade volumes alone on the cost of discretion, with market power held fixed. In our linear model an increase in $\bar{\alpha}$ increases both trade volume and market power, but one way to generate rising trade volumes while holding market power fixed within the linear model is to increase $\bar{\alpha}$, λ^* , and β^* in an appropriate fashion. It can be confirmed with these parameter changes that the statements of Remark 2 continue to apply.

³³We note here that an increase in Δ_γ has an ambiguous impact on the optimal degree of discretion, because: (1) conditional on the agreement being rigid, it makes discretion more attractive through the indirect state-contingency effect; but (2) it tends to make the agreement less rigid (Remark 3), and since rigidity and discretion are complementary, this can lead to less discretion in the agreement.

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

Two new features emerge in this environment. The first is the possibility of the agreement $\{\tau(\alpha)\}$. It can be shown that, if the optimal agreement takes this form, τ is strictly increasing in α . Intuitively, in the linear model the expression for \mathcal{B} in (2.9) is increasing in α , and so the $\{\tau(\alpha)\}$ agreement has τ increasing in α as a result of the indirect incentive-management effect identified in Proposition 3. This confirms the insight developed in the general model above, that the presence of contracting costs can explain an escape-clause type of agreement.

The second new feature is that rigidity and discretion are no longer complementary, but are instead *substitutable*, in the sense that the cost (in terms of lost surplus relative to $\{FB\}$) of discretion and rigidity when they occur together (as in $\{\tau\}$) is *higher* than the sum of the costs when they occur separately (as in, respectively, $\{\tau(\alpha)\}$ and $\{\tau, s\}$). Intuitively, this reversal reflects two differences across the γ -uncertainty and α -uncertainty environments. First, when α (alone) is uncertain the presence of rigidity does not confer any extra value to discretion, because s^{eff} does not depend on α , and hence the indirect state-contingency effect – which underpins the complementarity between rigidity and discretion in the γ -uncertainty case – is inoperative. And second, in the α -uncertainty case the indirect incentive-management effect *is* operative, as we have observed just above, and this makes discretion more costly when the agreement is also rigid, hence implying that discretion and rigidity are substitutes in this environment.

The following remark characterizes the optimal agreement as a function of c :

Remark 4. *When only α is uncertain, there exist 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 for $c \in (0, c_1)$;*
- (b) *of the form $\{\tau(\alpha)\}$ (with τ increasing in α) for $c \in (c_1, c_2)$;*
- (c) *of the form $\{\tau\}$ 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 escape-clause-type agreement $\{\tau(\alpha)\}$, we next ask, Under what conditions (if any) is the agreement $\{\tau(\alpha)\}$ optimal? It is direct to show that if k is sufficiently small and λ is sufficiently low (but strictly positive), then $c_1 < c_2$: an escape-clause-type agreement of the form $\{\tau(\alpha)\}$ is optimal for some c . This is

consistent with the result of Proposition 3: if λ is low then s is a poor substitute for τ , and hence leaving s to discretion is an attractive option, which in turn implies that an agreement of the form $\{\tau(\alpha)\}$ is optimal for a range of c as long as the added cost of contracting over state variables (k) is sufficiently small.

Finally, we note that 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 reported in Remark 3. Specifically, as Δ_α increases, the optimal agreement may switch from a contingent type $\{\tau(\alpha)\}$ to a rigid type $\{\tau, s\}$, which as Remark 3 indicates can never happen with an increase in Δ_γ . We report this in:

Remark 5. *As the degree of uncertainty Δ_α increases, the optimal agreement can switch from a contingent agreement to a rigid agreement.*

Intuitively, Remark 5 reflects the combined workings of the monopoly-power/trade-volume effects and the indirect incentive-management effect, and the fact that the cost of discretion in our linear model is not only rising in α but also convex. The key point is that as uncertainty over α rises, the cost of discretion rises, and as a consequence it may be optimal to move from a contingent agreement with discretion – where the contingencies provide indirect incentive management – to an agreement without discretion where the contingencies are no longer beneficial and where rigidity therefore becomes preferred.

Uncertainty about σ .

Here we consider briefly the case of uncertainty over the production externality σ . We assume that σ can take two possible values with equal probability: $\bar{\sigma} + \Delta_\sigma$ and $\bar{\sigma} - \Delta_\sigma$, with $\Delta_\sigma > 0$; we again let $\bar{\alpha}$ denote the deterministic value of α , and for simplicity we set $\gamma \equiv 0$. In this environment, the $\{FB\}$ agreement is $\{\tau = 0; s = \sigma\}$. Evidently, σ is relevant for s^{eff} but not for τ^{eff} , and this distinguishes the σ -uncertainty case from both the γ -uncertainty case (where γ is relevant for both s^{eff} and τ^{eff}) and the α -uncertainty case (where α is not relevant for either s^{eff} or τ^{eff}). This difference has subtle implications for the optimal agreement.

A first difference is the following. While the set of candidate optimal agreements (aside from $\{FB\}$) is analogous to those for the case of γ -uncertainty considered above, the potential appeal of making τ contingent on σ is distinct from the potential appeal of making τ contingent on γ , and arises for reasons analogous to the potential appeal of the escape-clause-type agreement $\{\tau(\alpha)\}$, that is, because of the indirect incentive-management effect as Proposition 3 indicates.

A second difference concerns the complementarity/substitutability between rigidity and discretion. As observed just above, the indirect incentive-management effect is present in the σ -uncertainty case, and as discussed in the case of α -uncertainty this effect tends to make rigidity and discretion substitutes. But the indirect state-contingency effect is also present in the σ -uncertainty case, as Proposition 2(b) indicates, and as discussed in the case of γ uncertainty this effect tends to make rigidity and discretion complements. As a result, rigidity and discretion can either be complements or substitutes in the σ -uncertainty case.

These observations, together with those made for the case of γ - and α -uncertainty, suggest an important insight. When uncertainty concerns variables (such as γ) that are relevant for τ^{eff} and s^{eff} , the indirect state-contingency effect is operative while the indirect incentive-management effect is not, and rigidity and discretion tend to be complements. When uncertainty concerns variables (such as α) that are not relevant for either τ^{eff} or s^{eff} , the indirect incentive-management effect is operative while the indirect state-contingency effect is not, and rigidity and discretion tend to be substitutes. Finally, as noted above, σ is relevant for s^{eff} but not τ^{eff} , and so both forces are at work; as a consequence, uncertainty about σ has ambiguous implications for the complementarity/substitutability between rigidity and discretion.

To sum up, *the nature of the optimal agreement, and in particular the interaction between rigidity and discretion, depends crucially and subtly on the source of uncertainty, and specifically on whether and how the uncertain variable is relevant for first-best intervention.*

3. The Role of the National Treatment Clause

Thus far we have focused on production subsidies as the central internal measure that governments must address along with tariffs when designing a trade agreement. But consumption taxes are of course an important policy instrument as well, and constraining the relationship between consumption taxes on domestically-produced and imported goods is the purpose of one of the GATT/WTO's central provisions, the National Treatment clause. In this section we evaluate the National Treatment (NT) clause as a means to economize on contracting costs.

To this end, we now suppose that, in addition to its tariff (τ) and production subsidy (s), the Home government has at its disposal an internal tax on consumption of the domestically produced good (t_h) and an internal tax on consumption of the imported good (t_f). As noted just above, the NT clause constrains the relationship between the separate consumption taxes t_h

and t_f , but evaluating the merits of the NT clause requires that we first explore the contracting possibilities in the absence of such a constraint. In fact, as we next show, an examination of the pricing relationships that must hold in this richer policy environment permits a simple reinterpretation of all of our earlier results to the present (non-NT) policy setting.

We return to the general model of section 2, augmented to this richer policy environment. As was the case in our earlier analysis, in the sector under consideration Foreign producer and consumer prices are equalized, or $q^* = p^*$, due to the absence of taxation by the Foreign government. And as before, for a Foreign firm to sell in both countries, it must receive the same price for sales in the Foreign country that it receives after taxes for sales in the Home country: now, however, with the richer set of Home policies, this condition implies $p^* = p - \tau - t_f$. And finally, the relationship between the Home producer and consumer price is now given by $q = p - t_h + s$. Nevertheless, despite the apparent differences that arise in this richer policy environment, we can express these new (non-NT) pricing relationships in the familiar form

$$\begin{aligned} p &= p^* + T, \text{ and} \\ q &= p^* + T + S, \end{aligned} \tag{3.1}$$

where $T \equiv \tau + t_f$ and $S \equiv s - t_h$.

Evidently, as a comparison between (3.1) and (2.1) reveals, the two central price wedges of the model are unchanged in this richer policy environment when the NT clause is absent, except that the role of the import tariff τ is now played by T , the “total tax on imports,” and the role of the production subsidy s is now played by S , the “effective production subsidy.” Hence, in the absence of an NT clause, each of the results of the previous sections can be reinterpreted as applying to T and S , with the the cost of including $T \equiv \tau + t_f$ in an agreement given by $2c_p$, and similarly the cost of including $S \equiv s - t_h$ in an agreement given by $2c_p$.³⁴ In analogy with our earlier analysis, we only consider agreements that impose separate equality constraints on T and S ; with a slight abuse of notation, we let \mathcal{A}_0 denote this class of agreements.

We next turn to the NT clause. For our purposes, the relevant part of the NT clause can be found in GATT Article III.2, which addresses internal taxation. Within the context of our model, we represent the core of the NT rule by the simple constraint $t_h = t_f$.³⁵ It is important

³⁴Note that τ and t_f are perfect substitutes and the same is true for s and t_h , and so T and S define the relevant policies for contracting in this richer policy setting absent an NT clause.

³⁵The NT clause can be interpreted as permitting a foreign product to be taxed more heavily in some cases, but only if and to the extent that this is motivated by legitimate policy objectives. This is not an issue in

to note that, while the NT provision restricts internal taxes to be the same, it does not constrain the common level at which these taxes are set (which we will denote t).

If the NT clause is included in an agreement, therefore, it transforms the set of policy instruments from (τ, s, t_h, t_f) to (τ, s, t) . We assume that including the NT clause costs $2c_p$ (because it is a constraint of the form $t_h = t_f$, hence it involves two policy instruments); and in the presence of the NT clause, we continue to assume that the inclusion of a policy instrument (τ , s , or t) in the agreement costs c_p .³⁶

For simplicity, we rely on institutional motivation to restrict our attention to just this particular clause: that is, we expand the class of feasible agreements \mathcal{A}_0 to allow for agreements that include the NT clause, and search for conditions under which the optimal agreement in this wider class includes the NT clause. We refer to an agreement that includes the NT clause as an “NT-based” agreement. Letting \mathcal{A}_{NT} denote the class of NT-based agreements, we thus focus on the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NT}$.

We begin with a key observation: the relationships between price wedges and policies for NT-based agreements are different from those that apply for non-NT agreements. For non-NT agreements, these relationships are given just above by (3.1). However, for NT-based agreements, these relationships become

$$\begin{aligned} p &= p^* + \tau + t, \text{ and} \\ q &= p^* + \tau + s. \end{aligned} \tag{3.2}$$

Notice a crucial difference between (3.1) and (3.2): as (3.2) indicates, with an NT-based agreement that ties down τ and s and leaves t to discretion, it is possible to tie down the producer price wedge $q - p^*$ while leaving discretion over the consumer price wedge $p - p^*$; but as (3.1) indicates, this is *not* possible with a non-NT agreement. For this reason, an NT-based agreement can offer something that cannot be achieved in the absence of the NT clause; put differently, leaving discretion *just* over the consumer price wedge *requires* that the NT clause be included in

the context of our model, since there is no efficiency rationale for treating the imported product less favorably than the locally produced good. For a model where this is a possibility, see Horn (2006). See also Horn and Mavroidis (2004) for legal and economic analyses of Article III text and case law.

³⁶It could be argued that including an NT clause in the agreement should cost less than specifying exact levels for t_h and t_f . By abstracting from this consideration we are stacking the deck against NT: if including the NT clause costs less than $2c_p$, the parameter region under which NT is optimal will be wider. Similarly, in the presence of the NT clause it might be argued that the cost of including t should be lower than c_p . As will become clear below, the main result of this section does not depend on the cost of including t in an NT-based agreement, and so we adopt the simplest assumption concerning this cost.

the agreement. The remaining question is then under what conditions this feature is desirable.³⁷

To answer this question, we begin by observing that the $\{FB\}$ agreement is given by the non-NT agreement $\{T = \gamma; S = \sigma - \gamma\}$, which costs $4c_p + 2c_s$. The first-best outcome can also be implemented by the NT agreement $\{NT; \tau = 0; t = \gamma; s = \sigma\}$, but this costs $5c_p + 2c_s$, so it is not efficiently written. From this starting point, we now ask: Is there a region of parameters for which it is (strictly) optimal to include the NT clause in an agreement? Consider the NT-based agreement $\{NT; \tau = 0; s = \sigma\}$. This agreement ties down the producer price wedge $q - p^*$ while leaving discretion over the consumer price wedge $p - p^*$, and costs $4c_p + c_s$, marking a savings of c_s over the $\{FB\}$ agreement as a result of the exclusion of γ from the agreement. Clearly, if the discretion over t associated with the agreement $\{NT; \tau = 0; s = \sigma\}$ does not lead the Home government to significantly distort t away from its efficient level for terms-of-trade purposes, so that when left to discretion Home sets t sufficiently close to γ , then $\{NT; \tau = 0; s = \sigma\}$ can achieve close to the first-best outcome and dominates the $\{FB\}$ agreement by exploiting the indirect state-contingency effect to save on contracting costs.

A key question is then what conditions will ensure that t is not significantly distorted for terms-of-trade purposes when left to discretion. Denoting the efficient level of t for given levels of τ and s by $t^{eff}(\tau, s)$, and denoting the optimal unilateral (unconstrained) choice of t for given levels of τ and s by $t^R(\tau, s)$, the loss in global surplus implied by $t^R(\tau, s)$ relative to $t^{eff}(\tau, s)$, defined positively, is given by

$$W^G(t^{eff}(\tau, s), \tau, s) - W^G(t^R(\tau, s), \tau, s) = - \int_{t^{eff}(\tau, s)}^{t^R(\tau, s)} W_t^G(t, \tau, s) dt. \quad (3.3)$$

Following steps similar to those leading up to (2.9), we may observe that a sufficient condition for this loss in global surplus to be small is that $|W_t^G(t^R(\tau, s), \tau, s)|$ is small. And $|W_t^G(t^R(\tau, s), \tau, s)|$ can in turn be written as:

$$|W_t^G(t^R(\tau, s), \tau, s)| = -\frac{\partial p^*}{\partial t} \cdot M = \frac{1}{\frac{1}{|D'|} \left[\frac{\eta^*}{p^*} + \frac{X'}{M} \right] + \frac{1}{M}} \equiv \mathcal{N}, \quad (3.4)$$

³⁷Notice that no other constraint on the relationship between t_h and t_f can achieve this feature (e.g., it is easy to confirm that a constraint of the form $t_h = 2t_f$ cannot accomplish this). For this reason, our analysis provides a rationale for a constraint along the lines of NT , not simply for “linkage” between t_h and t_f . It might also be wondered whether an NT-based agreement could be an efficient way to tie down $p - p^*$ while leaving $q - q^*$ to discretion. The answer is no. To see why, note that this can be achieved with a non-NT agreement by tying down T , which costs $2c_p$; but it costs $4c_p$ to achieve this with an NT-based agreement, because this would require the inclusion of the NT clause (which costs $2c_p$), and then tying down τ and t (which costs an additional $2c_p$).

where all magnitudes in \mathcal{N} are evaluated at τ , s and $t^R(\tau, s)$. Thus, we may conclude that the expression in (3.3) is small, and therefore that the loss in global surplus in omitting t from an NT-based agreement is small, if \mathcal{N} as defined in (3.4) is small.³⁸

Evidently, as (3.4) indicates, the conditions that lead the cost of discretion over t to be small can be understood in terms of the monopoly power, trade volume and instrument substitutability effects familiar from (2.9), with one important difference that reflects the difference in instruments across these two expressions: as (3.4) reflects, the substitutability between t and τ is low when X' is high or when $|D'|$ is low, because in each of these cases t (which distorts only the consumer margin) is a poor substitute for τ (which distorts both the producer and the consumer margin) as an instrument for manipulating the terms of trade.

Note an interesting point: while trade volume and monopoly power have similar impacts on the desirability of discretion over t in the NT-based agreement and over S in the non-NT agreement, the conditions that make t a poor substitute for the tariff (low price-sensitivity of demand, high price-sensitivity of supply) are essentially *opposite* to those that make S a poor substitute for the import tax (low price-sensitivity of supply, high price-sensitivity of demand).

We can now turn to conditions under which the NT clause is part of the optimal agreement. Note that, if we drive $|D'|$ to zero, while keeping the other magnitudes in (3.4) strictly positive and finite, the agreement $\{NT; \tau = 0; s = \sigma\}$ (which recall costs $4c_p + c_s$) approximates the first-best outcome, and no other agreement with the same (or lower) cost can accomplish this. In particular, as we have observed, a non-NT agreement with S left to discretion cannot approach the first-best outcome in these circumstances; and recall that the $\{FB\}$ agreement costs $4c_p + 2c_s$, because it is contingent on γ as well as σ . The key point is that the agreement $\{NT; \tau = 0; s = \sigma\}$ gets close to the first best outcome *despite not being contingent on γ* , as a consequence of the indirect state-contingency effect. From here, it is a small step to conclude that there is a range of contracting costs such that $\{NT; \tau = 0; s = \sigma\}$ is strictly optimal.³⁹

Clearly, the same argument as above holds also if we drive X' to infinity, while keeping the other magnitudes in (3.4) strictly positive and finite. The following proposition summarizes:

³⁸As before, the condition is that \mathcal{N} is small for all τ and all states of the world.

³⁹The statement of this result can be made more precise, at the cost of being more cumbersome. For example, let us focus on the condition that $|D'|$ is “sufficiently small.” Consider a parametric specification of the demand and supply functions, and let θ be the vector of demand/supply parameters. Assume that there exists θ^0 such that, as $\theta \rightarrow \theta^0$, $D' \rightarrow 0$ for all p and α , while E^* , X' and M stay strictly positive and finite. (This is not a strong assumption if the parametric specification is rich enough. It is satisfied for example in the linear model of section 2.5; there, it suffices to take β to zero keeping the other parameters constant.) Then if θ is close enough to θ^0 , there is a range of contracting costs such that an NT-based agreement is strictly optimal.

Proposition 4. *If $|D'|$ is sufficiently small or X' is sufficiently large, there is a range of contracting costs for which it is optimal to include the NT clause in the agreement.*

Proposition 4 identifies a simple condition under which our model can rationalize the use of an NT-based agreement. This condition describes circumstances in which the NT-based agreement gets close to the first best while avoiding the need to utilize costly state-contingencies, by utilizing instead the indirect state-contingency associated with discretion over internal taxes constrained only by the NT clause.

Finally, notice that the NT-based agreement on which we have focused, $\{NT; \tau = 0; s = \sigma\}$, includes constraints on s . This feature fits comfortably with the WTO, because the SCM Agreement places significant constraints on subsidies, but the NT clause was also a central feature of the (pre-WTO) GATT, and there subsidies were largely unconstrained. This raises the question: Can our model account for an agreement of the form $\{NT, \tau\}$ that includes the NT clause but leaves s (and t) to discretion? We first observe that the pair of instruments s and t represent a complete set of taxes in the presence of NT, and so an agreement that left both s and t to discretion would be empty for any sector where both s and t were readily available to the government. This observation implies that the introduction of frictions in the use of sector-specific domestic policies (e.g., administrative costs) is a necessary ingredient – for any model – in accommodating an agreement of the form $\{NT, \tau\}$; but it is also easy to see how such an agreement could be understood within a multi-sector generalization of our model that allowed for such frictions. In particular, suppose that in some sectors there are frictions in the use of consumption taxes, while in other sectors it is problematic to use production subsidies, so that in each sector the government may have a complete set of tax instruments at its disposal but just not a redundant set. Then the analysis of section 2 would apply to the first type of sector, while the analysis of the present section would apply to the second type of sector;⁴⁰ and an agreement of the form $\{NT, \tau\}$ (applied across sectors) could potentially be optimal, if the conditions identified in Proposition 2 (Proposition 4) held for sectors where production subsidies (consumption taxes) were the available domestic policy.⁴¹

⁴⁰It is direct to confirm that, if production subsidies were not available, our analysis of NT-based agreements would be unchanged, except that the optimal NT-based agreement would take the form $\{NT; \tau = \sigma\}$.

⁴¹In the setting just described, if trade volumes increase over time, it may be optimal to switch from an agreement of the type $\{NT, \tau\}$ to one of the type $\{NT, \tau, s\}$, as the governments' incentives to distort subsidies (in sectors where subsidies are available) increase. In this sense, the explanation of the evolution from the GATT to the WTO described in section 2 would still be valid in this setting. More broadly, the general prediction that our model offers in this regard is that, if trade volumes increase, the agreement should tend to introduce more

4. The Role of Weak Bindings

In the previous sections we focused on agreements that impose equality constraints (“strong bindings”), as in $\{T = 2\}$ or $\{NT; \tau = 0, s = \sigma\}$. In a world of costless contracting, the first-best outcome would be implemented, and hence there would be nothing to gain from using inequality constraints. In the presence of contracting costs, however, it may not be optimal to implement the first-best outcome, and as we argue in this section, in a second-best environment it may be preferable to impose policy ceilings (“weak bindings”) rather than strong bindings. Below we formalize this claim, but we first develop some intuition through a simple example.

Consider the model of section 3. Suppose for the moment that only γ is uncertain, and let us focus on (non-NT) agreements that constrain the import tax T . As a first observation, we note that weak bindings can achieve at least the same level of net global welfare as strong bindings. Intuitively, this is because the purpose of the agreement is to prevent governments from *raising* import taxes above their efficient level. The next question is: Can weak bindings offer a *strict* improvement over strong bindings? To answer this question, we need to distinguish between contingent and rigid bindings.

It is clear that a contingent weak binding (e.g. $\{T \leq \gamma\}$) cannot offer a strict improvement over a contingent strong binding (e.g. $\{T = \gamma\}$). The reason is that a contingent strong binding can position the policy variable exactly where it is optimal to place it for all realizations of the state variable, and so the added ex-post flexibility that a weak binding offers cannot be of value.

When it comes to rigid bindings, however, the situation is different. Compare a *rigid strong binding* of the form $\{T = \bar{T}\}$ with the corresponding *rigid weak binding* $\{T \leq \bar{T}\}$. We will argue that for some model configurations, the latter can offer a strict improvement over the former.⁴² Let $T^N(\gamma)$ denote the noncooperative equilibrium level of T as a function of γ , and let T_{\max}^N and T_{\min}^N denote, respectively, the highest and lowest values of $T^N(\gamma)$ over all possible realizations of γ . Intuitively, the optimal level of the strong binding \bar{T} must be below T_{\max}^N , but it may be above T_{\min}^N . If \bar{T} is above T_{\min}^N , then a weak binding is strictly preferable, because in the lowest states the government sets T below the binding, and this improves global welfare. Clearly there exist model configurations for which this is the case.

constraints on domestic instruments. We note here that, among other changes introduced relatively recently in the WTO, there is a strengthening of the constraints imposed on the use of consumption-side policies such as product standards; this too seems broadly consistent with the results of our model.

⁴²By “model configuration” we mean a configuration of demand functions (D, D^*) , supply functions (X, X^*) , contracting costs (c_s, c_p) and distribution of the state vector.

Following up on this intuition, weak bindings should be appealing not only for the total import tax T or the tariff τ , but also for the other policy instruments that the agreement may need to bind, namely the subsidies S and s . The reason is that governments are tempted to distort production subsidies in import-competing industries *upwards*, so the relevant constraint is an upper bound on the subsidy; and allowing a government to go below the ceiling can only be good for global welfare.⁴³ Finally, the appeal of weak bindings should extend to the case of multidimensional uncertainty: intuitively, to the extent that a binding is at least partially rigid (i.e. not contingent on all state variables) it should be desirable to make it weak.

We show in the Appendix that the intuition developed above is valid provided that two simple conditions are satisfied. The first one is a kind of no-Lerner-paradox requirement; specifically, we require the import tax T to have the standard (favorable) impact on terms of trade even when S is discretionary: formally, letting $p^*(T, S)$ denote the equilibrium world price as a function of T and S , we require $\frac{d}{dT}p^*(T, S^R(T)) < 0$. The second condition is similar: the production subsidy s must also have the intuitive (favorable) impact on terms of trade even when the consumption tax t is discretionary: formally, letting $p^*(s, \tau, t)$ denote the world price as a function of s , τ and t under the NT constraint, we require $\frac{d}{ds}p^*(s, \tau, t^R(s, \tau)) < 0$.⁴⁴ We will henceforth assume both of these conditions.

To state the result in a concise way, we let: $\mathcal{A}^S \equiv \mathcal{A}_0 \cup \mathcal{A}_{NT}$ denote the class of agreements we have considered thus far and \mathcal{A}^W denote the same class of agreements except that strong bindings are replaced by weak bindings. We may now state:

Proposition 5. (i) *Weak bindings cannot do worse than strong bindings: $\max_{A \in \mathcal{A}^W} \omega(A) \geq \max_{A \in \mathcal{A}^S} \omega(A)$.* (ii) *For some model configurations, weak bindings perform strictly better than strong bindings: $\max_{A \in \mathcal{A}^W} \omega(A) > \max_{A \in \mathcal{A}^S} \omega(A)$. A weak binding can improve strictly over the corresponding strong binding only if the binding is at least partially rigid.*

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, Proposition 5 highlights another sense in which rigidity and discretion may be

⁴³As applied to trade taxes, this argument would also remain valid in an export sector. However, it would have to be qualified with respect to the domestic instruments, because in export sectors the terms-of-trade motives lead to domestic interventions of reverse signs (i.e., taxes on domestic production of the export good, and subsidies on domestic consumption of the export good).

⁴⁴It is easy to verify that these conditions are satisfied, for example, if demand and supply functions are linear.

complementary ways to economize on contracting costs: if the agreement is rigid, it may be valuable to give governments *downward* discretion in the setting of the relevant policies.

In light of the above result, our model suggests that the constraints imposed by trade agreements should predominantly take the form of weak bindings. This prediction is broadly consistent with the observed nature of the GATT/WTO contract, where policy commitments are essentially all in the form of weak bindings.

5. Conclusion

Our model abstracts from some important elements that should be incorporated into a more complete theory. We conclude the paper with a brief discussion of a number of these elements.

We have worked within a two-country setting. This precludes the study of one of the foundational provisions of the GATT/WTO, its MFN rule, and by implication precludes as well the study of its most important exception to the MFN rule under which free trade areas and customs unions may form. Extending our analysis to a multi-country environment would permit an exploration of these and related topics.

We have focused on instrument-based contracts, excluding outcome-based contracts from our analysis. Outcome-based bindings of trade volumes are not emphasized in real-world trade agreements, and so this is a natural starting point. But there are provisions of the GATT/WTO (most notably the non-violation provision in GATT Article XXIII) that do have this flavor, and such provisions warrant investigation within an incomplete-contracts setting.

We have focused on import-sector policies, abstracting from export policies. But when it comes to export policies, the GATT/WTO exhibits a curious mix of rigidity (export subsidies are banned) and discretion (export taxes are generally left unconstrained), and an important question is whether these features can be understood from an incomplete contracts perspective.

In a similar vein, we have emphasized tax instruments, but we have not considered quantity instruments such as quotas, which are essentially banned by GATT Article XI. In the competitive setting we consider, there is an equivalence between tariffs and quotas, and so an agreement might naturally seek to ban one of these instruments. But a first-best agreement that banned quotas and specified tariffs would require fewer contingencies than one that banned tariffs and specified quotas (γ and σ in the former; γ , σ and α in the latter), suggesting more generally that banning quotas and contracting over tariffs might be a way to save on contracting costs

by reducing the number of required contingencies. This too seems like an idea worth exploring.

In this paper, we have adopted the view that trade agreements serve to provide an escape for governments from a terms-of-trade driven Prisoner's Dilemma. An alternative view is that trade agreements help governments make commitments to their private sectors (e.g., political lobbies or unions). Exploring the implications of this alternative view for the optimal design of trade agreements in the presence of contracting costs would constitute a fascinating project in its own right. Also, we have ruled out the existence of (non-pecuniary) cross-border externalities associated with production and consumption: such externalities could alter the nature of the optimal trade agreement (if those externalities were not handled in another international forum), and their inclusion would be a valuable extension to explore.

We have abstracted from political economy motives, which are clearly important considerations for real-world trade policy determination. If political economy motives can be represented as an extra weight on producer surplus (see Baldwin, 1987), then there is a close similarity between the presence of such motives and the case of production externalities that we have considered, since the domestic producer surplus is closely related to the domestic output. Indeed, we have explored the implications of a simple political-economy extension of the linear model of section 2.5, and find that our main results are unaffected. The one qualification is that, when political economy motives are present, an agreement that constrains only the domestic subsidy can now improve over the noncooperative equilibrium, contrary to the case where governments maximize social welfare. Nevertheless, it is still the case that contracting over subsidies alone is suboptimal, as Proposition 1 indicates, at least as long as the political benefits of import protection are not too convex. In the linear model, this is ensured provided that the extra weight placed by the government on domestic producer surplus is below a critical level.⁴⁵

Our formal analysis does not identify an explicit role for a dispute settlement body. But it is often observed informally that the Dispute Settlement Body of the WTO plays an important role in helping to “complete” the incomplete WTO contract. Our contracting costs are modeled as a “black box,” but introducing an explicit role for a dispute settlement body into our analysis would require disentangling contract writing costs from costs of interpreting and enforcing the

⁴⁵We emphasize that this critical level need not be “small.” For example, if we denote this critical level by $\hat{\xi}$, then in the linear model it is easy to demonstrate that $\hat{\xi}$ is the prohibitive level – and hence contracting over subsidies alone is suboptimal for all non-prohibitive parameter values – provided that the slope of the Foreign export supply curve is not too high relative to the magnitude of the slope of the Home import demand curve; moreover, $\hat{\xi}$ can never be less than 1/2 for any parameter values.

contract. This is a difficult task, but it could add an important new dimension to our analysis (see Maggi and Staiger, 2008, for a beginning in this direction).

Finally, our paper explains contract incompleteness on the basis of contracting costs, but other approaches are possible. In the contract-theoretic literature, it is standard to assume that there is asymmetric information between the contracting parties and the court, so that some variables observed by the parties are not “verifiable,” and to then characterize the optimal contract by means of mechanism-design techniques. We can relate this “standard” approach to our approach with a simple example. Consider our model of section 2 and suppose there is a single uncertain variable, say γ . The standard approach would assume that γ is not verifiable, so that the contract cannot be made contingent on γ . A contract is then a *menu* of policy combinations (τ, s) , from which the (importing) government can choose. With one-dimensional uncertainty, this contract is typically a nonlinear function, say $g(\tau, s) = 0$. Under some conditions, the optimal contract induces self-selection (separation) of the different government “types,” that is, the government chooses a different point in the menu depending on the value of γ .

At this point it is easy to see the relationship between the standard approach and the approach taken in our paper. The key links are two: (1) In the standard approach, the only impediment to contracting is the nonverifiability of γ . In terms of our model, this is analogous to assuming a prohibitive cost of contracting over γ (e.g. $c_s = \infty$) and zero cost of contracting over policies ($c_p = 0$). In this sense, our approach can be seen as more general, since it allows for a non-prohibitive cost of contracting over state variables, and perhaps even more importantly, for a positive cost of contracting over policies. (2) The standard approach allows for contracts that impose general constraints of the form $g(\tau, s) = 0$, whereas in the present paper we have focused on a simpler class of contracts for tractability reasons.

Notice that, as a consequence of these differences in assumptions, the predictions also differ. The standard approach predicts that the optimal contract always takes the form $g(\tau, s) = 0$; thus, the optimal contract is never directly contingent on state variables such as γ , and it always includes *all* policy instruments, because contracting over policies is assumed costless. On balance, then, our modeling of contracting costs is arguably a richer formalization of the impediments to contracting relative to the standard approach; but this comes at the price of focusing on a narrower class of contracts. Ideally, one would retain our framework of contracting costs while allowing for a more general class of contracts of the form $g(\tau, s) = 0$, thus achieving the best of both approaches. We see this as an ambitious avenue for future research.

6. Appendix

Proof of Proposition 1:

Consider the Home government's best response when s is exogenously moved away from its Nash equilibrium level. It is convenient to think of the Home government's choice variable as being the import volume M , rather than the tariff. We will write the objective as a function of M , with s an exogenous parameter. Let $p(M; s)$ be the inverse import demand function, defined implicitly by $D(p) - X(p + s) = M$. Let $p^*(M)$ be the inverse export supply function; note that s does not affect this function directly. For future reference, note that $p_s(M; s) = \frac{X'}{D' - X'}$ and $p_M(M; s) = \frac{1}{D' - X'}$. The maximization problem can be written as

$$\max_M W = \Gamma(p(\cdot)) + \Pi(p(\cdot) + s) + [p(\cdot) - p^*(\cdot)]M - (s - \sigma)X(p(\cdot) + s) - \gamma D(p(\cdot)).$$

Letting $\hat{M}(s)$ be the optimal import level as a function of s , the claim is that $\hat{M}'(s) = 0$ when evaluated at the Nash equilibrium subsidy ($s = \sigma - \gamma$). This amounts to showing that $W_{Ms} = 0$ when evaluated at the Nash equilibrium policies.

After some algebra, we have

$$W_M = -(s - \sigma + \gamma)X'(\cdot)p_M(\cdot) + [p(\cdot) - p^*(\cdot)] - Mp_M^*(\cdot) - \gamma.$$

The cross derivative is

$$W_{Ms} = -(s - \sigma + \gamma)\frac{d[X'p_M]}{ds} - X'p_M + p_s.$$

Noting that the first term is zero at the Nash equilibrium subsidy, and plugging in the expressions for p_M and p_s , we find that $W_{Ms}|_{s=\sigma-\gamma} = -\frac{X'}{D'-X'} + \frac{X'}{D'-X'} = 0$. This proves the claim. QED

Proof of Remark 1

We start by stating necessary and sufficient conditions for an agreement to be optimal for a range of c . To state these conditions, note first that the total contracting cost can be expressed as $C = c \cdot (n_p + k \cdot n_s) \equiv c \cdot m$. The variable m can be seen as a measure of the ‘‘complexity’’ of the agreement. Let $m(A)$ and $\Omega(A)$ denote respectively the level of complexity and the level of Ω associated with agreement A . Then it is not hard to show that an agreement \hat{A} is optimal for some c if and only if:

- (A) There is no agreement \hat{A} such that $m(\hat{A}) = m(\hat{A})$ and $\Omega(\hat{A}) > \Omega(\hat{A})$; and
(B) for any pair of agreements \hat{A}, \check{A} such that $m(\hat{A}) \leq m(\hat{A}) \leq m(\check{A})$, the following holds:

$$\Omega(\hat{A}) \geq \frac{m(\check{A}) - m(\hat{A})}{m(\check{A}) - m(\hat{A})} \Omega(\hat{A}) + \frac{m(\hat{A}) - m(\hat{A})}{m(\check{A}) - m(\hat{A})} \Omega(\check{A}).$$

Condition (A) states that the candidate agreement must be optimal in its complexity class. Condition (B) states that, for an agreement \hat{A} to be optimal for some c , it must pass the following test: if we pick an agreement \hat{A} with a lower complexity level than \hat{A} and an agreement \check{A} with a higher complexity level than \hat{A} , the gross surplus Ω must increase in a concave way when moving from \hat{A} to \hat{A} to \check{A} .

The next observation is that the optimal degree of complexity is decreasing in c . It can easily be shown that, for any two non-equivalent agreements, \hat{A} and \check{A} , if \hat{A} is optimal for $c = \acute{c}$ and \check{A} is optimal for $c = \check{c} > \acute{c}$, then $m(\check{A}) < m(\hat{A})$.

We are now ready to prove Remark 1. We consider separately two cases: $k \leq 1$ and $k > 1$. Suppose first $k \leq 1$. We observed above that, as c increases, we must move from more complex to less complex agreements. This implies that the optimal sequence of agreements is a subsequence of $(\{FB\}, \{\tau, s\}, \{\tau(\gamma)\}, \{\tau\}, \{\emptyset\})$.

To proceed further, it proves useful to define:

- The cost of discretion absent rigidity (CD): $\Omega_{\{FB\}} - \Omega_{\{\tau(\gamma)\}}$;
- The cost of discretion in the presence of rigidity (\overline{CD}): $\Omega_{\{\tau, s\}} - \Omega_{\{\tau\}}$;
- The cost of rigidity absent discretion (CR): $\Omega_{\{FB\}} - \Omega_{\{\tau, s\}}$; and
- The cost of rigidity in the presence of discretion (\widetilde{CR}): $\Omega_{\{\tau(\gamma)\}} - \Omega_{\{\tau\}}$.

The next observation is that rigidity and discretion are complementary: indeed, it can be confirmed that $\overline{CD} < CD$ for all parameter values (and \overline{CD} may even be negative). Note also that $\overline{CD} < CD$ is equivalent to $CR > \widetilde{CR}$.

We now argue that $\{\tau, s\}$ and $\{\tau(\gamma)\}$ cannot both be part of the optimal sequence of agreements. Suppose that $\{\tau, s\}$ is part of the optimal sequence of agreements. Condition (B) then implies $(\Omega_{\{FB\}} - \Omega_{\{\tau, s\}}) \leq k(\Omega_{\{\tau, s\}} - \Omega_{\{\tau\}})$, or, using the definitions above,

$$CR \leq k \cdot \overline{CD}. \quad (6.1)$$

Similarly, suppose that $\{\tau(\gamma)\}$ is part of the optimal sequence of agreements. Then condition (B) implies $k(\Omega_{\{FB\}} - \Omega_{\{\tau(\gamma)\}}) \leq (\Omega_{\{\tau(\gamma)\}} - \Omega_{\{\tau\}})$, or

$$k \cdot CD \leq \widetilde{CR}. \quad (6.2)$$

Now recall that $\overline{CD} < CD$ and $CR > \widetilde{CR}$. This implies that (6.1) and (6.2) cannot both be satisfied. The claim follows.

The argument for the case $k > 1$ is similar to the one developed above, except for the first step: condition A implies that the optimal sequence of agreements is $(\{FB\}, \{\tau(\gamma)\}, \{\tau, s\}, \{\tau\}, \{\emptyset\})$. But then one can establish, with analogous steps as above, that $\{\tau, s\}$ and $\{\tau(\gamma)\}$ cannot both be part of the optimal sequence of agreements, and hence the statement in Remark 1 follows. QED

Proof of Proposition 5

(i) Let \tilde{A}^S be the optimal agreement in class \mathcal{A}^S . To prove the claim it suffices to show that, if we replace strong bindings with weak bindings in \tilde{A}^S , global welfare Ω cannot decrease. In what follows we will omit the uncertain parameters from the arguments of the relevant functions, as this should not cause confusion.

Agreement \tilde{A}^S can be one of the following types: (a) $\{T = \tilde{T}\}$; (b) $\{T = \tilde{T}; S = \tilde{S}\}$; or (c) $\{NT; \tau = \tilde{\tau}; s = \tilde{s}\}$. The bindings \tilde{T} , \tilde{S} , $\tilde{\tau}$, \tilde{s} may be contingent, but again we omit the state variables from the notation.

Let us start with case (a). Consider replacing $\{T = \tilde{T}\}$ with $\{T \leq \tilde{T}\}$. This can decrease Ω only if in some state the government chooses $T < \tilde{T}$ and this implies a lower level of Ω than $T = \tilde{T}$. But T will only be set below the ceiling if the noncooperative import tax T^N is lower than the ceiling, in which case the importing country will set $T = T^N$. Let us show that Ω decreases in T for $T > T^N$. Recall that the subsidy is set as $S = S^R(T)$ and note that

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

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 $T > T^N$. Also, the sign of $\frac{d}{dT}W^*(T, S^R(T))$ is the same as the sign of $\frac{d}{dT}p^*(T, S^R(T))$, which is negative by assumption. This in turn implies $\frac{d}{dT}\Omega(T, S^R(T)) < 0$ for $T > T^N$. We can conclude that switching to a weak binding cannot decrease Ω .

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

(1) it chooses $(T = \tilde{T}, S = \tilde{S})$: In this case there is of course no change in Ω relative to the strong-binding agreement.

(2) it chooses $(T = \tilde{T}, S = S^R(T))$: Here it must be that $S^R(T)$ is lower than the ceiling. Let us evaluate $\Omega_S = W_S + W_S^*$. Clearly, $W_S < 0$ for $S > S^R(T)$, and $W_S^* < 0$, hence $\Omega_S < 0$ for $S > S^R(T)$, which in turn implies that switching to weak bindings increases Ω .

(3) it chooses $(T = T^R(S), S = \tilde{S})$: Here it must be that $T^R(S)$ is below the ceiling. 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 .

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

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

Finally, consider case (c). Since the NT-based agreement constrains the wedge $q - p^*$ and leaves the wedge $p - p^*$ discretionary, it is convenient to re-define variables as follows: $p - p^* \equiv z$ and $q - p^* \equiv v$. We can think of z and v as the policy instruments and of the NT-based agreement as imposing a constraint $v = \tilde{v}$.

Let us now replace the agreement $\{NT; \tau = \tilde{\tau}; s = \tilde{s}\}$ with $\{NT, \tau \leq \tilde{\tau}, s \leq \tilde{s}\}$. Using the new notation, this is equivalent to replacing the constraint $v = \tilde{v}$ with the constraint $v \leq \tilde{v}$. In other words, the NT-based agreement with weak bindings effectively places an upper bound on the producer price wedge. 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 $v > v^N$ (where v^N denotes the unilateral optimum for v and $z^R(v)$ the unilateral optimum for z given v). Note that

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

Clearly, $W_v < 0$ for $v > v^N$. Next note that v and z affect W^* only through p^* , and $\frac{d}{dv}W^*(v, z^R(v))$ has the same sign as $\frac{d}{dv}p^*(v, z^R(v))$. It is direct to verify that our assumptions imply $\frac{d}{dv}p^*(v, z^R(v)) < 0$. This in turn implies that switching to weak bindings cannot decrease Ω .

(ii) To prove this claim it suffices to show that there exists a model configuration for which (a) agreement \tilde{A}^S contains some rigid strong bindings, and (b) replacing these with rigid weak bindings increases Ω strictly. By making c_s very high and c_p very low we can ensure condition (a). Next, from the arguments developed above, we know that a sufficient condition for (b) to be satisfied is that for some state of the world the noncooperative level of a policy is below the (rigid) binding for that policy. It is easy to show that there exists a model configuration for which this is the case. QED

References

- [1] Bagwell, Kyle and Robert W. Staiger, "A Theory of Managed Trade," *The American Economic Review*, vol. 80, pp. 779-795, 1990.
- [2] Bagwell, Kyle and Robert W. Staiger, "Enforcement, Private Political Pressure and the GATT/WTO Escape Clause," *Journal of Legal Studies*, vol. 34(2), pp. 471-514, 2005.
- [3] Bagwell, Kyle and Robert W. Staiger, "Domestic Policies, National Sovereignty and International Economic Institutions," *Quarterly Journal of Economics*, 116(2), pp 519-562, 2001.
- [4] Bagwell, Kyle and Robert W. Staiger, *The Economics of the World Trading System*, MIT Press, 2002.
- [5] Bagwell, Kyle and Robert W. Staiger, "Will International Rules on Subsidies Disrupt the World Trading System?," *The American Economic Review*, vol. 96(3), pp. 877-895, 2006.
- [6] Baldwin, Richard, "Politically Realistic Objective Functions and Trade Policy," *Economics Letters*, vol. 24, pp. 287-90, 1987.
- [7] Battigalli, Pierpaolo and Giovanni Maggi, "Rigidity, discretion and the costs of writing contracts," *The American Economic Review*, vol. 92(4), pp. 798-817, 2002.
- [8] Battigalli, Pierpaolo and Giovanni Maggi, "International Agreements on Product Standards: an Incomplete-Contracting Theory," NBER Working Paper No. 9533, 2003.
- [9] Bhagwati, Jagdish and V.K. Ramaswami, "Domestic Distortions, Tariffs and the Theory of Optimum Subsidy," *Journal of Political Economy* 71 (1), February 1963, pp. 44-50.
- [10] Copeland, Brian, "Strategic Interaction Among Nations: Negotiable and Non-negotiable Trade Barriers," *Canadian Journal of Economics*, vol. 23, pp. 84-108, 1990.
- [11] Costinot, Arnaud, "A Comparative Institutional Analysis of Agreements on Product Standards," *Journal of International Economics* vol. 75(1),pp. 197-213, 2008.
- [12] Ederington, Josh, "International Coordination of Trade and Domestic Policies," *The American Economic Review*, vol. 91(5), pp. 1580-93, 2001.

- [13] Horn, Henrik, "National Treatment in the GATT," *The American Economic Review*, vol. 96(1), pp. 394-404, 2006.
- [14] Horn, Henrik and Petros Mavroidis, "Still Hazy After All These Years: The Interpretation Of National Treatment In The GATT/WTO Case-Law On Tax Discrimination," *European Journal of International Law*, vol. 15, pp. 39-69, 2004.
- [15] Hudec, Robert, *The GATT Legal System and World Trade Diplomacy* (2nd Ed.), Salem NH: Butterworth, 1990.
- [16] Johnson, Harry G., "Optimal Trade Intervention in the Presence of Domestic Distortions," in Richard E. Caves, Harry G. Johnson and Peter B. Kenen (eds.), *Trade, Growth and the Balance of Payments*, Amsterdam: North-Holland Publishing Company, 1965.
- [17] Maggi, Giovanni and Andres Rodriguez-Clare, "A Political-Economy Theory of Trade Agreements," *The American Economic Review*, vol. 97(4), pp. 1374-1406, 2007.
- [18] Maggi, Giovanni and Robert W. Staiger, "On the Role and Design of Dispute Settlement Procedures in International Trade Agreements," NBER Working Paper No. 14067, June 2008.
- [19] Markusen, James R., "International Externalities and Optimal Tax Structures," *Journal of International Economics*, vol.5, pp. 15-29, 1975.
- [20] Schwartz, Warren F. and Alan O. Sykes, "The Economics Structure of Renegotiation and Dispute Resolution in the WTO/GATT System," *The Journal of Legal Studies*, Vol. 31, pp. 179-204, 2001.
- [21] Sykes, Alan O., "The Economics of WTO Rules on Subsidies and Countervailing Measures," in Patrick F.J. McCrory, Arthur E. Appleton, and Michael G. Plummer, eds., *The World Trade Organization: Legal, economic and political analysis*. Vol. 2. Springer: New York, 2005.