

Horn, Henrik; Maggi, Giovanni; Staiger, Rikard W.

Working Paper

Trade Agreements as Endogenously Incomplete Contracts

IFN Working Paper, No. 689

Provided in Cooperation with:

Research Institute of Industrial Economics (IFN), Stockholm

Suggested Citation: Horn, Henrik; Maggi, Giovanni; Staiger, Rikard W. (2006) : Trade Agreements as Endogenously Incomplete Contracts, IFN Working Paper, No. 689, Research Institute of Industrial Economics (IFN), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/81185>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Trade Agreements as Endogenously Incomplete Contracts

Henrik Horn, Giovanni Maggi and Robert W. Staiger

Trade Agreements as Endogenously Incomplete Contracts

Henrik Horn (Research Institute of Industrial Economics, Stockholm, and CEPR)

Giovanni Maggi (Princeton University, NBER and CEPR)

Robert W. Staiger (Stanford University, University of Wisconsin and NBER)*

November 2006

Abstract

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

*This paper has benefited 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, UCSD and Yale, as well as 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. A sizeable economics literature examines various aspects of this incompleteness. The typical approach is to impose exogenous restrictions on the set of policy instruments that can be included in a trade agreement, and examine what the agreement can accomplish given these limitations.¹ This literature has advanced our understanding of 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 to shed light on core features of the GATT/WTO, including:

1. 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.

2. The bindings are largely rigid (i.e. not state-contingent). But there are “escape clauses” that allow countries to unilaterally impose temporary protection (GATT Art. XIX) or to renegotiate bindings (GATT Art. XXVIII).

3. 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.

¹An incomplete list of papers that fall into this category is Copeland (1990), Bagwell and Staiger (2001), Battigalli and Maggi (2003), Costinot (2004) and Horn (2006).

In this paper we propose a simple theoretical framework where the incompleteness of the agreement, and in particular the ways and degrees in which discretion and rigidity are present in the agreement, is determined endogenously. This approach, we believe, can help explain why the regulation of goods trade in the GATT/WTO has been structured along the lines described above. And while in this paper we take the GATT/WTO as our institutional focus, we believe that the framework and results we develop here can contribute as well to an understanding of the forms taken by contractual incompleteness in trade agreements more generally.

The analytical starting point of the paper is the 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 the governments’ incentives to act opportunistically. This feature suggests that the agreement should be comprehensive in its coverage of trade-relevant policies. The second source of difficulty 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 while the gains from an agreement are likely to be quite large, and so the costs of contracting are unlikely to have important effects on the structure of trade agreements. But it should be kept in mind that these costs include in principle the cost of negotiation delays, the cost of lawyers, the cost of dispute panels, and the like, and that in reality these costs must be multiplied by a vast number of products, countries, policy instruments, and contingencies. 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. To take just one illustrative example, what might seem on its surface to be a relatively straightforward task of finding a workable agreement to limit the use of subsidies has preoccupied member governments of the GATT/WTO for over

50 years, and a concise definition of exactly what is meant by a “subsidy” continues to elude negotiators. 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 that is negotiated.²

We work within a competitive two-country setting, where countries may experience a consumption externality. The role of this externality is to provide an efficiency rationale for policy intervention. As we explain in more detail in a later section, the model would generate similar insights if, instead of a consumption externality, we considered a production externality or introduced political-economy motives in the governments’ objectives. What is essential is the presence of an (economic or political) efficiency rationale for policy intervention.

For simplicity we focus on intervention in import sectors, and assume that governments have access to a rich set of taxation instruments, namely: import tariffs, distinct consumption taxes on domestically-produced goods and on imported goods, and production subsidies. Uncertainty plays a central role. To bring out the main points, we focus for much of our analysis on one-dimensional uncertainty, and we contrast two cases: one in which the source of the uncertainty is the level of the externality, and one in which it is the underlying level of import demand. We later consider how the insights derived in the settings with one-dimensional uncertainty generalize to a setting of multidimensional uncertainty.

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).

The first step of our analysis is to examine two benchmark scenarios: one is the no-agreement outcome – that is the noncooperative equilibrium – and the other is the first-best outcome. In the absence of an agreement, the importing country would use its policy instruments to

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, Robert 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, Warren Schwartz and Alan 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).

manipulate the terms of trade in standard fashion. This of course would lead to a globally inefficient outcome, and hence there is scope for an agreement to restrain governments from behaving opportunistically. Were it not for the consumption externality, 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 externality, the contracting problem is substantially more complex: the first best agreement will now involve the use of policy instruments, and it will require these policies to be state-contingent if the externality is 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 we hinted above, there are two essential ways to save on contracting costs: one is to make the agreement (partially or fully) rigid, and the other is to leave some of the policies to the discretion of governments. The key part 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 result is that the optimal agreement tends to leave more discretion on domestic policy instruments than on import taxes. This result accords well with the traditional emphasis on border measures over domestic instruments that characterizes the GATT/WTO; moreover, 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.

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

As a consequence of the monopoly-power effect, the model predicts that the optimal degree

of discretion is lower if the import demand level, and hence the volume of trade, is higher, because the incentive to distort domestic instruments for manipulating the terms of trade is then stronger. This in turn 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 also suggests a reason why developing countries may have been largely exempted from the WTO 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 rich array of domestic policy instruments necessary to find easy substitutes for tariffs.

The role of uncertainty in shaping the optimal agreement depends in subtle ways on its source. This can be traced to a key observation: in our contracting environment, rigidity and discretion *interact*, and they do so in different ways depending on the source of the uncertainty.

In particular, when uncertainty involves the level of the consumption externality – or more generally state variables that are directly relevant for the first-best levels of domestic policy instruments – rigidity and discretion are *complementary* ways of saving on contracting costs, in the sense that the cost of rigidity is lower when the agreement features discretion. The intuition for this result is that introducing discretion (over domestic instruments) into a rigid agreement is a way to achieve *indirect state-contingency*: with discretion, the unilateral setting of domestic instruments varies in the “right” way with the level of the externality; and for this reason, introducing discretion mitigates the cost of rigidity.

On the other hand, when uncertainty involves the level of import demand – or more generally state variables that are not directly relevant for the first-best levels of import taxes – rigidity and discretion are *substitutes*, meaning that the cost of rigidity is higher when the agreement features discretion (over domestic instruments). The intuition is that in this case rigidity diminishes the ability of the agreement to provide *indirect incentive management*: as we have observed above, the incentive to distort domestic instruments for manipulating the terms of trade is stronger when the import demand level is higher, and so allowing higher import taxes in the high-import-demand states can be valuable as a way to mitigate this incentive, something that rigidity precludes.

Thus, the indirect state-contingency effect tends to make rigidity and discretion comple-

mentary, while the indirect incentive-management effect tends to make rigidity and discretion substitutes. With the aid of these effects, we are able to describe how the source of uncertainty – consumption externality, import demand level, and in a later section production externality/political economy motive – determines whether rigidity and discretion are complements or rather substitutes in a given environment, and how this shapes the optimal agreement.

Of special interest is our finding 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 higher. Our rationale for an escape clause is distinct from those that have been highlighted in the existing theoretical literature.³ In particular, an escape clause can be appealing in our model due precisely to the indirect incentive management effect, that is, as a way to manage the higher incentives to distort domestic instruments for terms of trade purposes in periods of high underlying import volume.

In the final part of the paper we extend the analysis to shed light on two other core aspects of the GATT/WTO: the presence of a National Treatment (NT) rule 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 possible means of saving on contracting costs. We interpret the NT rule as requiring equal internal taxation of the imported and the domestically produced good. We show that there is only one type of NT-based agreement that can be strictly optimal in our setting: this is an agreement that imposes the NT rule and ties down the import tariff and the production subsidy, but leaves the (common) consumption tax to the governments’ discretion. We identify a simple set of conditions under which this type of agreement is indeed optimal. The key condition concerns the degree of substitutability between the consumption tax and the tariff for the purposes of manipulating terms of trade: if this degree is sufficiently low (which is the case when demand is more rigid and supply is more elastic), and if the level of contracting costs lies in an intermediate range, then an NT-based agreement is strictly optimal.

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. This type of binding combines rigidity and discretion,

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.

since the ceiling does not depend on the state of the world, and the government has discretion to set the policy below the ceiling. This finding strengthens the insight – highlighted above – that rigidity and discretion may be complementary ways to economize on contracting costs.

The paper is structured in the following way: in section 2 we lay out the basic model and derive the two benchmarks of no-agreement outcome and first-best outcome; in section 3 we characterize the optimal agreement; in section 4 we examine the role of the NT clause; in section 5 we examine the role of weak bindings; in section 6 we discuss the case of production externalities and that of political-economy motives in the governments' objectives; in the Conclusion we discuss a number of simplifying assumptions made in the model and suggest directions for further research; the Appendix provides proofs that are not contained in the body of the paper.

2. The Model

We consider a perfectly competitive world with two countries, Home and Foreign. There are three goods, a numeraire good and two non-numeraire goods (which we label 1 and 2). Home is a natural importer of good 1 and Foreign a natural importer of good 2.

We start by describing the supply structure in the Home country. The numeraire good is produced one-for-one from labor. The supply of labor is large enough that this good is always produced in positive amount; therefore the equilibrium wage is equal to one. Each non-numeraire good is produced from labor with diminishing returns. In particular, we assume the following production function for each good j :

$$X_j = \sqrt{2\lambda_j L_j}, \quad j = 1, 2,$$

where X_j is the production of good j and L_j is the labor employed in the production of good j . This supply structure is convenient because it implies linear supply functions. In particular, if q_j denotes the producer price for good j , the supply function for good j is

$$X_j(q_j) = \lambda_j q_j, \quad j = 1, 2.$$

The associated profit function for good j is

$$\Pi_j(q_j) = \frac{1}{2} \lambda_j q_j^2, \quad j = 1, 2.$$

We assume a similar supply structure also for the Foreign country, and let asterisks denote foreign variables:

$$X_j^*(q_j^*) = \lambda_j^* q_j^*, \quad \Pi_j^*(q_j^*) = \frac{1}{2} \lambda_j^* q_j^{*2}, \quad j = 1, 2.$$

The representative citizen's utility function is linear in the numeraire good and separable in the non-numeraire goods. Also, to create an efficiency rationale for policy intervention, we allow for the possibility of a (negative) consumption externality. As we argue in section 6, the main qualitative results of the model would be the same if we had a production externality, or if we had political-economy motives in the governments' objectives. What is important is the existence of an (economic or political) efficiency rationale for policy intervention.

We assume that the externality is linear in aggregate domestic consumption and does not cross borders. Hence, we model the consumption externality as a purely domestic "eyesore" pollutant, along the lines of Markusen (1975) and especially Ederington (2001). Also, we assume that good 1 generates an externality only in Home, and good 2 only in Foreign.⁴

Formally, we are assuming the following utility functions for the two countries:

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

where c_j and C_j denote respectively individual and aggregate consumption of good j . The parameters γ_1 and γ_2^* capture the strength of the externalities in the two countries. Each consumer ignores the effect of her individual consumption on aggregate consumption, so the externalities do not affect demand functions. We assume that the sub-utility functions are quadratic, so that the implied demand functions are linear:

$$D_j(p_j) = \alpha_j - \beta_j p_j, \quad D_j^*(p_j^*) = \alpha_j^* - \beta_j^* p_j^*,$$

where p_j and p_j^* denote consumer prices. All parameters introduced thus far are positive.

Assuming that the population in each country is a continuum of measure one, we can write the consumer surplus associated with good j in Home and Foreign respectively as:

$$CS_j(p_j) = u_j(D_j(p_j)) - p_j D_j(p_j), \quad CS_j^*(p_j^*) = u_j^*(D_j^*(p_j^*)) - p_j^* D_j^*(p_j^*).$$

We assume that each government can intervene only in its import sector, but within this sector we allow each government to use a rich set of taxation instruments, namely: an import tariff (τ), an internal tax on consumption of the domestically produced good (t_h), an internal

⁴We could relax this assumption and allow each good to have externalities in both countries, but this would only complicate the analysis without adding to the insights of the model. Notice also that, if one considers political-economy motives (as in section 6) instead of consumption externalities, such an asymmetric structure would capture situations where import-competing interests are organized but export interests are not.

tax on consumption of the imported good (t_f), and a production subsidy to domestic firms (s). All instruments are expressed in specific terms. The reason we allow for this rich set of policies is twofold. First, we want to represent a situation in which governments can use a variety of instruments to act opportunistically and manipulate terms of trade, so that there is an interesting trade-off between the benefits of contracting over many policies and the costs of doing so. And second, we consider in a later section an NT rule that constrains the relationship between taxes on the consumption of domestically-produced and imported goods. Evaluating the merits of such a rule requires that we start with separate consumption taxes t_h and t_f .

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 the reader should keep 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 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 $p^* = p - \tau - t_f$. And finally, the relationship between the Home producer price and the Home consumer price is given by $q = p - t_h + s$.

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

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

where $T \equiv \tau + t_f$ and $S \equiv s - t_h$. The arbitrage relationships in (2.1) describe the two central price wedges in the model; the first is the wedge between the Home consumer price and the Foreign price (equal to T), and the second is the wedge between the Home producer price and the Foreign price (equal to $T + S$). Note that τ and t_f are perfectly substitutable policies (only their sum matters), and the same is true for s and t_h (only their difference matters). Thus, while it is appropriate to refer to τ as a “border measure” and to t_f , t_h and s as “internal measures,” we will also sometimes refer to T as the total tax on imports, or simply as the “import tax,” and to S as the “effective production subsidy” or the set of “domestic policies.”

Market clearing requires that world demand equal world supply, or

$$D(p) + D^*(p^*) = X(q) + X^*(q^*). \quad (2.2)$$

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

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

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

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

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

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

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

$$\frac{\alpha}{\lambda + \beta} > \frac{\alpha^*}{\lambda^* + \beta^*}. \quad (2.3)$$

We will henceforth assume that (2.3) holds.

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, Home welfare can be written as the sum of consumer surplus, profits, net revenue (i.e. revenue from the import tax T minus expenditures on the effective production subsidy S), and the valuation of the externality:

$$W(T, S) \equiv CS(T, S) + \Pi(T, S) + T \cdot M(T, S) - S \cdot X(T, S) - \gamma D(T, S),$$

where the notation emphasizes the dependence of each component of welfare on policies. Recalling that in the sector under consideration the Foreign country has no externality and no policy instruments of its own, Foreign welfare is simply the sum of consumer surplus and profits:

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

2.1. The noncooperative equilibrium and the efficient policies

We first derive the noncooperative equilibrium policies, which we take to represent the choices made in the absence of international agreements. With the foreign government passive (in the sector under consideration), the home government's optimal unilateral⁵ policies are defined by

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

The first condition defines the optimal unilateral choice of T given S , which we denote $T^R(S)$, and the second condition defines the optimal unilateral choice of S given T , denoted $S^R(T)$.

From the above system we may derive the home government's noncooperative equilibrium policies, which we denote by T^N and S^N :

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

where η^* is the elasticity of the foreign export supply (itself evaluated at S^N and T^N).⁶

Recalling that $T \equiv \tau + t_f$ and $S \equiv s - t_h$, there are many equivalent policy combinations that correspond to the unilateral policy choices T^N and S^N . One of these combinations is $\{\tau = \frac{p^*}{\eta^*}, t_h = t_f = \gamma, s = 0\}$, making it transparent that in the noncooperative equilibrium the home country sets its traditional (Johnson, 1953-54) "optimal tariff" – the inverse of the foreign export supply elasticity – to exploit its monopoly power over the terms of trade (p^*) and applies a uniform Pigouvian consumption tax equal to the consumption externality. It 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 parameter γ .

Next we turn to the globally efficient policies, which we define as the policies that maximize "global welfare," i.e. the sum of home and foreign welfare:⁷

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

⁵We use interchangeably the words "noncooperative" and "unilateral."

⁶It is not hard to verify that W is jointly concave in (T, S) , which ensures that the first-order conditions are sufficient. For completeness we report here the explicit expression for T^N : $T^N = \frac{(\beta^* + \lambda^*)[\alpha + \Upsilon\gamma + \lambda\gamma] - (\beta + \lambda)\alpha^*}{(\beta^* + \lambda^*)[\Upsilon + (\beta + \lambda)]}$.

⁷In 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 utility to the two countries must maximize the sum of Home and Foreign welfare in each sector. More generally, this notion of efficiency would also be appropriate in asymmetric settings, provided that international lump sum transfers were available.

The efficient levels of T and S , which we denote by T^{eff} and S^{eff} , are respectively given by

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

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

Notice that the noncooperative level of S is equal to the efficient level ($S^N = S^{eff}$), and the noncooperative policies differ from the efficient policies only in that import taxes are too high – and noncooperative trade volumes are therefore too low – relative to their efficient levels ($T^N > T^{eff}$). The inefficiently high level of T reflects in turn the unilateral incentive to manipulate the terms of trade. Therefore, the potential gains from contracting in this setting arise entirely from the ability to control the incentive to utilize import taxes to manipulate the terms of trade. As a consequence of this feature – 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.

2.2. Uncertainty

We consider two possible sources of uncertainty: the consumption externality (γ) and the level of domestic demand (α). Uncertainty about γ can be interpreted as uncertainty about the true efficiency rationale for policy intervention, while shocks to α can be interpreted as shocks to the trade volume that are irrelevant for efficient policy levels.⁸ Focusing on uncertainty in 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. In later sections we discuss the extension of our results to more general stochastic environments.

We sometimes refer to γ and α as the state-of-the-world variables, or simply the “state” variables. For now, we impose no further structure on the nature of uncertainty.

⁸We could alternatively consider a shock that shifts the domestic supply function, but the qualitative results would not change.

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.⁹

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 prices or trade volumes.¹⁰

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 α and γ – and the costs of including *policy* variables – that is, τ, t_f, s and t_h . 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.¹¹

⁹Also implicit here is the assumption that governments can costlessly observe the realized values of the uncertain parameters. One might object that this is unrealistic for externality parameters, which are presumably difficult to ascertain even for a government, let alone the WTO courts. But our results remain valid even in the extreme opposite case in which governments do not learn anything about γ ex-post: in this case the agreement cannot be made contingent on γ , and the situation would be essentially equivalent to one where there is no uncertainty. Also notice that this objection would have less force if we had political-economy motives rather than externalities, because governments are likely to know well the political-economy pressures they face.

¹⁰We also abstract from agreements that are based on both instruments and outcomes, in the sense that they constrain the relationship between policy instruments and equilibrium outcomes, such as for example an agreement that constrains T to be a function of the import volume M .

¹¹The 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.

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.”¹² 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.¹³

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

Example 1: The agreement $\{\tau + t_f = 3\}$, or equivalently $\{T = 3\}$, specifies a rigid commitment for the total import tax, and costs $2c_p$.

Example 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.¹⁴ 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 $(T = 0 \text{ if } \gamma \leq 1)$ then a clause $T = \gamma$). On balance, however, we believe that the basic feature

¹²Relaxing the assumption of no recalling costs would introduce additional agreements for consideration (i.e. agreements that differ by the number of times a given variable appears), but the main qualitative insights of our analysis are unlikely to change.

¹³The only results that could change under a more general cost function of the form $C(n_s, n_p)$ are those concerning the complementarity/substitutability between rigidity and discretion; for example if the cross-derivative of C with respect to n_s and n_p were positive and sufficiently high, rigidity and discretion would be always substitutes. Also, it might be reasonable to assume that it is more costly to contract over internal measures (t_f, s, t_h) than over tariffs (τ) , because in reality it is easier to verify border measures than internal measures. But as will become clear below, in this case our qualitative results would only be strengthened.

¹⁴For 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 $c_s = c_p$.

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.

3. Optimal Agreements

In order to characterize the optimal choice of agreement, we need to introduce some definitions and notation. First, we say that two agreements are *equivalent* if they implement the same outcome and have the same cost. Second, we refer to the *efficiently-written first-best* agreement as the least costly among the agreements that implement the first best outcome. We will 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 $\omega \equiv \Omega - C$.

A natural and 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 $c_p \equiv c$, 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 (we henceforth refer to c simply as “contracting costs”). In much of the analysis to follow, we keep k fixed and consider changes in c . Note that, with this new notation, the total contracting cost can then be expressed as $C = c \cdot (n_p + k \cdot n_s) \equiv c \cdot m$. The variable m provides a rough measure of the “complexity” of the agreement, in that it captures the number of policy variables and states involved (with the latter weighed by the parameter k).

Before we impose more structure on the nature of uncertainty and the set of agreements under consideration, we present three results that hold quite generally. Our first result provides necessary and sufficient conditions for an agreement to be optimal for a range of contracting costs. To develop these conditions, we begin by observing that any agreement A is characterized by a level of complexity $m(A)$ and a level of gross global welfare $\Omega(A)$, and therefore it is associated with a point in (Ω, m) space. If we plot this point for each feasible agreement, we obtain a set that describes all the feasible combinations of Ω and m . Denote this set by \mathcal{F} .

Lemma 1. Consider a feasible set of agreements \mathcal{F} . An agreement \hat{A} in \mathcal{F} is optimal for some c if and only if it satisfies the following two conditions:

- (i) \hat{A} is optimal in its complexity class, i.e. there is no agreement A' such that $m(A') = m(\hat{A})$ and $\Omega(A') > \Omega(\hat{A})$;
- (ii) for any pair of agreements A', A'' such that $m(A') \leq m(\hat{A}) \leq m(A'')$,

$$\Omega(\hat{A}) \geq \frac{m'' - \hat{m}}{m'' - m'}\Omega(A') + \frac{\hat{m} - m'}{m'' - m'}\Omega(A'') \quad (3.1)$$

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

Lemma 1 is illustrated by Figure 1. Each point drawn in the (Ω, m) space is a feasible agreement. Consider the candidate agreement \hat{A} . Condition (i) is that the candidate agreement must be optimal in its complexity class. This is reflected in Figure 1 by the fact that \hat{A} attains the highest level of Ω on the vertical line that contains it. Condition (ii) is a kind of concavity requirement on the Ω function. To explain in what sense this is a concavity condition, we first define $\bar{\Omega}(m)$ as the maximum level of Ω that can be attained with an agreement of complexity m – this is the level of gross global welfare as a reduced-form function of m . Condition (ii) states that, for an agreement \hat{A} (with associated complexity level \hat{m}) to be optimal for some c , it must pass the following test: pick an arbitrary complexity level lower than \hat{m} (call it m'), and one higher than \hat{m} (call it m''); the function $\bar{\Omega}(m)$ must be concave with respect to the three points m' , \hat{m} and m'' . The agreement \hat{A} satisfies condition (ii) as well and, as illustrated, is the optimal agreement for $c = \hat{c}$ (as well as for a range of c around \hat{c}) achieving the maximum net global welfare $\omega(\hat{A})$. By contrast the agreement \tilde{A} , while satisfying condition (i), does not satisfy condition (ii), and so is not an optimal agreement for any c .¹⁵

The economic interpretation of the “concavity” condition (ii) is, broadly speaking, that there must be declining gains in gross welfare from adding complexity to the agreement. Otherwise, if it paid to move from A' to the more complex \hat{A} it would also pay to take the further step to the even more complex A'' , in which case \hat{A} would not be optimal for any c . From an economic point of view, it may seem natural that there should be diminishing gross returns from including additional variables in the agreement. However, as will be seen, this is often *not* true in the contracting environment considered here.

¹⁵Note that conditions (i) and (ii) together are equivalent to the condition that the agreement A lie on the upper boundary of the convex hull of \mathcal{F} , depicted by the dashed line in Figure 1. This condition is also reflected in Figure 1: agreement \hat{A} is on the dashed line, while agreement \tilde{A} is below it.

Our second result concerns the relationship between the level of contracting costs and the complexity of the optimal agreement (as measured by m). Intuition suggests that this relationship should be monotonic. The following result (whose proof is straightforward and omitted) confirms this intuition, showing that more complex agreements will be chosen when the level of contracting costs is lower.

Lemma 2. *Consider two non-equivalent agreements, A' and A'' . If A' is optimal for $c = c'$ and A'' is optimal for $c = c'' > c'$, then $m(A'') < m(A')$.*

To state our third result, we first observe that, since the two policy instruments τ and t_f are perfect substitutes and matter only through their sum T , constraining one of the two instruments but not the other would have no effect. The same is true for the domestic instruments s and t_h , which matter only through their difference S . Hence, we can think of T and S as the relevant policy variables, with the inclusion of each variable in the contract costing $2c_p$.

We may now state our third result: 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 effective subsidy S (even in a state-contingent way) while leaving the import tax T to discretion cannot improve over the noncooperative equilibrium, and therefore cannot be an optimal agreement.*

At a broad level, the intuition for this result is very simple. Contracting over S alone is useless because, as we emphasized above, the inefficiency in the noncooperative equilibrium concerns T , not S . To develop a more precise understanding of Proposition 1, let us begin at the noncooperative equilibrium and consider an agreement that imposes a small exogenous change in S . This triggers a change in the Home government's choice of T . In particular, as we explain below, T will adjust 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.

Why is T adjusted so that p^* remains unchanged? To see this intuitively, consider for simplicity the case in which there is no externality ($\gamma = 0$). It is convenient to think of the choice variable as being the import volume M , rather than the tariff. Consider the first-order

condition for the choice of M if the subsidy is undistorted, i.e. $S = 0$. This condition can easily be derived as $Mp^{*'}(M) = p(M) - p^*(M)$, where $p(M)$ is the inverse import demand function and $p^*(M)$ the inverse export supply function. The interpretation of this condition is standard: $Mp^{*'}(M)$ is the terms-of-trade gain from a marginal decrease in M , and $p(M) - p^*(M)$ (which is equal to the tariff) is the deadweight loss from a marginal decrease in M .

Now consider the optimal choice of M in the presence of an exogenous subsidy S . The first order condition in this case becomes $Mp^{*'}(M) = [p(M) - p^*(M)] - SX'(p)p'(M)$. The additional term $-SX'(p)p'(M)$ captures the increase in subsidy expenditures generated by a marginal decrease in M . We can now examine how an exogenous increase in S – starting from $S = 0$ – affects these marginal effects. First note that increasing S does not affect the marginal terms-of-trade gain ($Mp^{*'}(M)$). Second, increasing S decreases the marginal deadweight loss because it reduces $p(M)$, the price at which the Home country is willing to import M ; the amount of this reduction can be found by differentiating the condition $D(p) - X(p+S) = M$ for fixed M , yielding $(dp/dS)_{M=const} = \frac{X'}{D'-X'}$. Third, introducing a small S increases the marginal subsidy-expenditure term by $X'(p)p'(M) = \frac{X'}{D'-X'}$; but this exactly offsets the reduction in the marginal deadweight loss. Thus, a small exogenous S does not affect the net marginal benefit of changing M , and hence does not change the optimal import volume, which in turn implies that the optimal p^* is unaffected.¹⁶

We emphasize that, 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 constraints on all policy instruments. But with costly contracting the result of Proposition 1 gains relevance, as we show below.

At this point we impose more structure on the stochastic environment. It is convenient to consider separately two cases: uncertainty in the externality (γ) and uncertainty in the level of domestic demand (α). In a later section we consider multidimensional uncertainty.

¹⁶Notice that this intuitive argument does not rely on the linearity of the model, and indeed, it can be shown that Proposition 1 extends to a setting of general non-linear demand and supply functions. Also, this result carries through in a setting where governments use tariff and subsidy instruments to pursue additional policy goals such as revenue needs or distributional concerns. What is crucial for the result is that a government has sufficient instruments to target the domestic ($q - p$) and foreign ($p - p^*$) price wedge. Finally we note that our result 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*.

3.1. Uncertainty about the consumption externality

In this subsection we focus on the case where only γ is uncertain. For the sake of expositional simplicity, 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 $\Delta_\gamma > 0$.

We also impose more structure on the set of agreements under consideration: for the remainder of this section we consider only agreements that impose separate equality constraints on T and S . To be concrete, we allow for clauses of the type $(T = \gamma)$ or $(S = 10)$, but not for clauses of the type $(T + S = 3)$ or for inequality constraints of the type $(T \leq 1)$.¹⁷ We label this class of agreements \mathcal{A}_0 . In later sections we consider broader classes of agreements.

The first step is to derive the efficiently-written first-best agreement ($\{FB\}$). Clearly the agreement $\{T = \gamma; S = -\gamma\}$ implements the first best outcome. This agreement has $n_s = 1$ and $n_p = 4$ and therefore costs $(4+k)c$. But it might be conjectured that the first best outcome could also be implemented without constraining S , and therefore be accomplished more cheaply, since as we have noted previously in the noncooperative equilibrium only T differs from its efficient level. This conjecture is incorrect, but it is instructive to see why. The reason is that an agreement that constrains T but leaves discretion over S would permit the home government to choose S according to its unilateral optimum given T , which we have previously labeled $S^R(T)$. As we have observed, this unilateral optimum is equal to S^{eff} if the choice of T is unconstrained, but more generally it is direct to verify that

$$S^R(T) = S^{eff} + \frac{\Upsilon^2 - (\beta + \lambda)^2}{\Upsilon^2 - \lambda(\beta + \lambda)}(T^N - T). \quad (3.2)$$

By (3.2), the difference between $S^R(T)$ and S^{eff} is proportional to the difference between T^N and T . As a consequence, an agreement that attempts to move T towards its efficient level without also constraining S will cause S to become distorted for terms-of-trade purposes.

In fact, one cannot implement the first best outcome with an agreement that costs less than $(4+k)c$, and so $\{T = \gamma; S = -\gamma\}$ is indeed the $\{FB\}$ agreement. Note that both S and T are

¹⁷When there is significant uncertainty, a noncontingent contract of the type $g(\tau, t_f, t_h, s) = 0$ may do better than a noncontingent contract that pins down T and/or S separately, for two reasons. First, if g constrains the relationship between the components of T (for example as in $\tau = 1 - t_f^2$), it can mimic a weak binding ($T \leq \#$), or more generally a constraint that T must lie in a certain subset of the real line. In a later section we consider weak bindings and show that they can improve over strong bindings, because allowing for downward discretion on T can be good. Second, if g constrains the relationship between T and S , it may achieve higher gross global welfare than a contract that pins down the exact levels of T and S , again because it introduces some discretion. This has the flavor of an outcome-based contract, which we discuss in the concluding section.

state-contingent in the $\{FB\}$ agreement.

We next seek to characterize the optimal agreement as a function of the level of contracting costs: What is the optimal way of restructuring the agreement as c rises from zero?

The $\{FB\}$ agreement yields net global welfare equal to $\Omega(T = \gamma, S = -\gamma) - (4 + k)c$. Clearly, when c is sufficiently small the $\{FB\}$ agreement is optimal; and if c is sufficiently high, the empty agreement (which costs nothing and yields global welfare $\Omega(T = T^N, S = -\gamma)$) is optimal. The interesting question is then: What happens between these two extremes?

For intermediate levels of c , there are three cost classes of agreements that warrant consideration: agreements costing $(2 + k)c$; agreements costing $4c$; and agreements costing $2c$. By Proposition 1, we can ignore agreements that constrain S but not T . 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. $\{T(\gamma), S\}$). Therefore we only have three kinds of agreements to consider, in addition to $\{FB\}$ and $\{\emptyset\}$, as candidate optimal agreements: (i) agreements that constrain T as a function of γ , which we denote $\{T(\gamma)\}$; (ii) agreements that constrain T and S in a non-state-contingent fashion, which we denote $\{T, S\}$; and (iii) agreements that constrain T in a non-state-contingent fashion, which we denote $\{T\}$.

The three types of agreements $\{T, S\}$, $\{T(\gamma)\}$ and $\{T\}$ are all incomplete, but they are each incomplete in a different way. To describe these differences, it is useful at this point to recall the distinction, introduced by Battigalli and Maggi (2002), between two forms of contractual incompleteness: *rigidity*, which occurs when some of the policy variables are constrained in a non-contingent way; and *discretion*, which occurs when some of the policy variables are missing from the agreement.¹⁸ We can thus say that the agreement $\{T, S\}$ is rigid, the agreement $\{T(\gamma)\}$ features discretion (over S), and the agreement $\{T\}$ is both rigid and discretionary.

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

We now make an important observation: in our contracting environment, rigidity and discretion *interact* in non-trivial ways. The cost of rigidity depends on whether or not discretion is

¹⁸Notice that, in general, rigidity and discretion as defined in the text do not necessarily imply a loss of gross surplus relative to the first best. For example, if the demand parameter α is uncertain (as in section 3.2), then the first-best contract is rigid, because it does not include the state variable α .

present in the agreement, and the cost of discretion depends on whether or not the agreement is rigid. To describe this interaction, we introduce the following definitions:

- a. The cost of discretion absent rigidity (CD): $\Omega_{\{FB\}} - \Omega_{\{T(\gamma)\}}$;
- b. The cost of discretion in the presence of rigidity (\overline{CD}): $\Omega_{\{T,S\}} - \Omega_{\{T\}}$;
- c. The cost of rigidity absent discretion ($CRTS$): $\Omega_{\{FB\}} - \Omega_{\{T,S\}}$; and
- d. The cost of rigidity in the presence of discretion (CRT): $\Omega_{\{T(\gamma)\}} - \Omega_{\{T\}}$.

Notice that these four quantities are linearly dependent, since one can always be expressed as a linear combination of the other three.

An interesting manifestation of the interaction between rigidity and discretion is the following: while the cost of discretion absent rigidity, CD , is always positive, the cost of discretion in the presence of rigidity, \overline{CD} , *may be negative*. In other words, it is possible that $\Omega_{\{T,S\}}$ is lower than $\Omega_{\{T\}}$: conditional on the agreement being rigid, introducing discretion may increase gross global welfare. Intuitively, introducing discretion (in S) into a rigid agreement is an indirect way of introducing state-contingency in the agreement, and this is beneficial, because the unilateral choice of S – even if distorted – varies with γ in the “right” way (i.e. both $S^R(T)$ and S^{eff} are decreasing in γ). This beneficial effect of discretion may outweigh the negative effect of allowing a government to use S to manipulate the terms of trade. This suggests that *rigidity and discretion are complementary*, in the sense that the presence of rigidity mitigates the cost of discretion, possibly making it negative. Indeed, it can be confirmed that $\overline{CD} < CD$ for all parameter values. Note also that $\overline{CD} < CD$ is equivalent to $CRTS > CRT$. This is the flip-side of the complementarity we just highlighted: the presence of discretion decreases the cost of rigidity.

We note that the *indirect state-contingency* effect identified above, which is responsible for the complementarity between rigidity and discretion, depends crucially on the fact that the uncertain state variable (γ) affects the first-best setting of domestic policies (S). As we explain in the next subsection, if uncertainty concerns a state variable that does *not* affect the first-best level of S (e.g. the demand level α), then the indirect state-contingency effect is inoperative. This suggests – and the next subsection confirms – that the complementarity between rigidity and discretion hinges on the source of uncertainty, and need not arise in all uncertain environments.

We are now ready to return to the question we posed above: What is the optimal sequence of agreements as c increases from zero? This question can be answered by using Lemmas 1 and

2, together with the complementarity property that we just highlighted.

Lemma 2 tells us that, as c increases, we must move from more complex to less complex agreements. This immediately implies that the optimal sequence of agreements is a subsequence of $(\{FB\}, \{T, S\}, \{T(\gamma)\}, \{T\}, \{\emptyset\})$. We say “subsequence” because each of these agreements (except $\{FB\}$ and $\{\emptyset\}$) may be “skipped” over as c increases.¹⁹

Next we use Lemma 1 together with the complementarity property identified above to establish the following result: the agreements $\{T, S\}$ and $\{T(\gamma)\}$ cannot *both* be part of the optimal sequence of agreements. To see this, suppose that the agreement $\{T, S\}$ is part of the optimal sequence. Then Figure 2a illustrates a consequence of the “concavity” condition of Lemma 1: the point associated with $\{T, S\}$ must lie (weakly) above the line connecting the points associated with $\{T\}$ and $\{FB\}$, which in turn implies $2(\Omega_{\{FB\}} - \Omega_{\{T, S\}}) \leq k(\Omega_{\{T, S\}} - \Omega_{\{T\}})$, or, using the definitions above,

$$2CRTS \leq k\overline{CD} \tag{3.3}$$

Similarly, suppose that $\{T(\gamma)\}$ is part of the optimal sequence of agreements. Then, with reference to Figure 2b, Lemma 1 implies that the point associated with $\{T(\gamma)\}$ must lie (weakly) above the line connecting the points associated with $\{T\}$ and $\{FB\}$. Clearly this implies: $k(\Omega_{\{FB\}} - \Omega_{\{T(\gamma)\}}) \leq 2(\Omega_{\{T(\gamma)\}} - \Omega_{\{T\}})$, or

$$kCD \leq 2CRT \tag{3.4}$$

Now recall that $\overline{CD} < CD$ and $CRTS > CRT$ (complementarity between rigidity and discretion). This implies that (3.3) and (3.4) cannot both be satisfied: if there are diminishing returns to complexity around $\{T, S\}$ there cannot be diminishing returns to complexity around $\{T(\gamma)\}$, and vice-versa. Hence, returns to complexity must be *increasing* around (at least) one of these two agreements.

Figures 2a and 2b reflect this conclusion: if, as in Figure 2a, $\{T, S\}$ is positioned above the dotted line, $\{T(\gamma)\}$ must lie below it; and if, as in Figure 2b, $\{T(\gamma)\}$ is positioned above the dotted line, $\{T, S\}$ must lie below it.²⁰ Intuitively, the agreement $\{T\}$ features both rigidity *and* discretion. But the complementarity between rigidity and discretion means that these two

¹⁹In the above statement we implicitly assume that $k \leq 2$, so that $\{T(\gamma)\}$ is not more costly than $\{T, S\}$. But the statement is true also with $k > 2$, because as we establish below, $\{T(\gamma)\}$ and $\{T, S\}$ cannot both be part of the optimal sequence.

²⁰Each figure is drawn under the assumption that $k < 2$, but as can be seen from the preceding arguments this is immaterial for the result.

features are less costly when they occur together than when they occur separately, as in the agreements $\{T, S\}$ (which features only rigidity) and $\{T(\gamma)\}$ (which features only discretion). Therefore, $\{T, S\}$ and $\{T(\gamma)\}$ cannot both lie above the line connecting $\{T\}$ and $\{FB\}$.

We summarize these results in the following proposition:

Proposition 2. *Assume that only γ is uncertain. There exist scalars c_1, c_2, c_3 and c_4 with $0 < c_1 \leq c_2 \leq c_3 \leq c_4 < \infty$ such that the optimal agreement in \mathcal{A}_0 is:*

- (a) *the $\{FB\}$ agreement for $c \in (0, c_1)$;*
- (b) *of the form $\{T, S\}$ for $c \in (c_1, c_2)$;*
- (c) *of the form $\{T(\gamma)\}$ for $c \in (c_2, c_3)$;*
- (d) *of the form $\{T\}$ for $c \in (c_3, c_4)$; and*
- (e) *the empty agreement for $c > c_4$.*

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

An important feature of Proposition 2 concerns the way contractual incompleteness varies across policy instruments: the effective subsidy S tends to be more discretionary than the import tax T . More specifically, for a range of contracting costs it may be optimal to contract over T while leaving S to discretion, but it is never optimal to contract over S and leave T to discretion. In this way, Proposition 2 predicts that trade agreements should always include commitments over import taxes, and should only introduce commitments over domestic policies as the agreement becomes more complete. This prediction resonates broadly with the approach taken by the GATT/WTO, which has been to first establish a base of commitments over import tax levels, and only later to take on various domestic policies (most notably subsidies). Notice, too, that our 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.

Consider next the significance of the “complementary slackness” condition of Proposition 2: if $\{T(\gamma)\}$ is optimal for a range of c (i.e., if $c_2 < c_3$), then $\{T, S\}$ is not optimal for any c (i.e., $c_1 = c_2$), and vice-versa. In light of this condition, we may now pose the following question: As contracting costs rise from zero, when is it optimal to first economize on contracting costs by introducing discretion ($c_2 < c_3$), and when by first introducing rigidity ($c_1 < c_2$)?

To provide answers to this question in terms of underlying model parameters, we first relate at a broad intuitive level the way in which the cost of rigidity and the cost of discretion depend

on the fundamentals of the contracting environment.

Let us start with the cost of rigidity. Intuitively, 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.

The determinants of the cost of discretion (over S) are more subtle.²¹ First, intuitively the cost of discretion (with or without rigidity) is higher when the effective subsidy S and the import tax T are closer substitutes with regard to their effect on the terms of trade. Thus a key determinant of the cost of discretion is *the degree of substitutability between policy instruments*. Recalling that the terms of trade are given by $p^*(T, S) = [\alpha + \alpha^* - (\beta + \lambda)T - \lambda S]/\Upsilon$, a rough measure of this substitutability is given by the marginal rate of substitution between T and S with respect to the terms of trade: $\frac{\partial p^*/\partial T}{\partial p^*/\partial S} = \frac{\beta}{\lambda} + 1$. This suggests that S is a closer substitute for T – and hence the cost of discretion is higher – when demand is less elastic (β is low) and when supply is more elastic (λ is high).

A second determinant of the cost of discretion is the wedge between the noncooperative and efficient tariff, which in turn reflects the degree of home-country *monopoly power* as measured by the (inverse of the non-cooperative) foreign export supply elasticity η^* . If the home country has large monopoly power, then this wedge is large; therefore, the incentive to alter domestic policies to manipulate the terms of trade is high, and hence the cost of discretion is high. Notice that the home country’s monopoly power in our model is directly linked to the volume of trade, and thus a primary determinant of its magnitude is the level of import demand, α .²² When α is higher, η^* is lower, and so the cost of discretion (with or without rigidity) is higher.

Finally, our earlier discussion of the indirect state-contingency effect suggests that there is an additional determinant of the cost of discretion, albeit one that applies only in the presence of rigidity (\overline{CD}): the extent to which the state-contingency introduced into a rigid agreement by discretion is beneficial. Intuitively, this effect is stronger when uncertainty in γ is higher.

According to the intuition we developed above on the costs of rigidity and discretion, we therefore expect $\{T, S\}$ to be favored over $\{T(\gamma)\}$ when the level of import demand α is high (so that the cost of discretion is high) and when uncertainty Δ_γ is low (so that the cost of

²¹Note that leaving discretion over S effectively leaves discretion over the producer price wedge $q - p^*$ (see the pricing relationships 2.1). As will be seen in section 4, there exist agreements outside class \mathcal{A}_0 that leave discretion (only) on the consumer price wedge $p - p^*$, but this cannot be accomplished by agreements in \mathcal{A}_0 .

²²Note that α is the intercept of the Home country’s import demand function. For this reason we refer to α as the import demand “level.”

discretion is high and the cost of rigidity is low). Also, we highlighted above that the cost of discretion tends to be high when the effective subsidy S is a close substitute for the import tax T , which in turn is the case when the demand slope β is low or the supply slope λ is high. The following remark confirms this intuition:

Remark 1. (i) *If the import demand level α is sufficiently high and/or the degree of uncertainty Δ_γ is sufficiently low, then $c_1 < c_2 = c_3$: the optimal sequence of agreements always includes $\{T, S\}$ and never includes $\{T(\gamma)\}$.*

(ii) *If the demand slope β is sufficiently low (so that S is a close substitute for T), then $c_2 = c_3 = c_4$: the optimal sequence of agreements may include $\{T, S\}$, but it never includes $\{T(\gamma)\}$ or $\{T\}$.*

(iii) *If the supply slope λ is sufficiently low (so that S is a poor substitute for T), then $c_1 = c_2 < c_3$: the optimal sequence of agreements always includes $\{T(\gamma)\}$ and never includes $\{T, S\}$.*

The results in Remark 1 stand in marked contrast to those of Battigalli and Maggi (2002): there, as contracting costs rise, first the optimal contract becomes rigid, and then discretion is introduced. Here, due to the non-separability of the contracting problem across instruments, it may be optimal to economize on contracting costs by introducing discretion before rigidity, and it may be that rigidity is not optimal for any level of contracting costs.

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

Remark 1 highlights how various parameters affect the optimal sequence of agreements as c varies, but it does not describe the effects of changing a parameter while holding constant all other parameters (including c). We conclude this subsection by presenting two of the more illuminating comparative-static results, namely, those for α and Δ_γ .

We start with changes in the import demand level, α . As we argued above, an increase in α tends to increase the cost of discretion. Intuitively, then, increasing α while keeping all

other parameters constant should lead to a lower degree of discretion. Also, it can be verified that α does not affect the costs of rigidity (CRT or $CRTS$); thus, intuitively, α does not affect *directly* the degree of rigidity. More specifically, as α increases, it can be shown that the optimal agreement never switches from $\{T\}$ to $\{T(\gamma)\}$ or from $\{T, S\}$ to $\{FB\}$, and so it is never the case that a change in α changes the degree of rigidity without affecting the degree of discretion. Nevertheless, the complementarity between rigidity and discretion implies that, as α increases and discretion falls, rigidity may *also* fall, as in the movement from $\{T\}$ to $\{FB\}$. The following proposition confirms this intuition:

Proposition 3. *As the import demand level α increases (holding all other parameters fixed):*
(i) The optimal degree of discretion decreases, in the sense that the number of policy instruments specified in the optimal agreement increases (weakly); (ii) The optimal degree of rigidity decreases (weakly).

The result of Proposition 3(i) reflects the monopoly power effect. In particular, the higher is the degree of monopoly power (the higher is α), the less desirable it is to leave policy instruments to discretion, with the order in which instruments are tied down (first T and then also S) dictated by Proposition 2.

Proposition 3(i) suggests 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. The possible explanation for this highlighted by Proposition 3(i) is that the increase in trade volumes over time (which in our model can be captured by an increase of α) has increased the cost of discretion, which in turn has augmented the need to constrain subsidies and other domestic policies in the agreement.²³

Our analysis developed thus far also suggests a broader insight. The essence of high instrument substitutability is that a government has access to a rich array of domestic policies which it can use to manipulate terms of trade if import taxes are constrained by a trade agreement. And the essence of high monopoly-power is that it faces a relatively inelastic foreign export

²³This interpretation might be more convincing if rising trade volumes increase the cost of discretion when market power is held fixed, but in our linear model an increase in α increases both trade volume and market power. One way to generate rising trade volumes while holding market power fixed within our model is to increase α , λ^* , and β^* in an appropriate fashion. Using (3.2), it can be confirmed with these parameter changes that rising trade volume with fixed market power does indeed increase the cost of discretion over domestic policies (S).

supply. Arguably, both of these conditions are most likely to apply to large developed countries. Therefore, it is more likely that contracting over domestic policies (such as S) is attractive for large developed countries than for small/developing countries. While our two-country model cannot address this issue formally, these results are at least suggestive of the possible benefits of a kind of “special and differential treatment” rule for small/developing countries when it comes to contracting over domestic policies (such as subsidies).²⁴

Finally, as we highlighted above, Proposition 3(ii) is a consequence of the complementarity between rigidity and discretion. This implies that rising trade volumes (higher α) 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 to the agreement is enhanced as a result.²⁵

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

Proposition 4. *As the degree of uncertainty Δ_γ increases (holding all other parameters fixed), the optimal agreement may switch from a rigid agreement to a contingent agreement, but not vice-versa.*

By itself, 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 directly relevant for the setting of T in the $\{FB\}$ agreement. As we demonstrate in the next subsection, the effects of increasing uncertainty over state variables (such as α) that are not directly relevant for the setting of T in the $\{FB\}$ can be very different.

3.2. Uncertainty about the level of import demand

In the previous subsection we examined a stochastic environment where uncertainty concerns only a state variable that affects directly the first-best levels of T and S . Here we explore the

²⁴In fact, when it comes to subsidies, Part VIII of the WTO Subsidies and Countervailing Measures Agreement introduces just such an exemption from commitments for developing country members. We thank Robert Lawrence for first bringing this implication to our attention.

²⁵We also note that it would *not* be accurate to say that an increase in α reduces both the degrees of discretion and rigidity because it increases the surplus from contracting, and hence for α sufficiently high the first-best agreement becomes optimal. It is not hard to verify that, starting from a parameter configuration for which $\{T, S\}$ is optimal, increasing α will *not* cause a switch from this contract to $\{FB\}$.

implications of a different source of uncertainty, which in some sense is at the opposite extreme: we now suppose that uncertainty concerns only a state variable (α) that has no impact on the first-best levels of T or S . We assume that α can take two possible values with equal probability: $\bar{\alpha} + \Delta_\alpha$ and $\bar{\alpha} - \Delta_\alpha$, with $\Delta_\alpha > 0$.

In this environment, the $\{FB\}$ agreement takes the form $\{T = \bar{\gamma}; S = -\bar{\gamma}\}$. 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. Also notice that, as an immediate 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 $\{T, S\}$; (ii) agreements of the form $\{T(\alpha)\}$; (iii) agreements of the form $\{T\}$; and (iv) the empty agreement.

Two important new insights emerge in this environment. The first is the possibility of the agreement $\{T(\alpha)\}$, where $T(\alpha)$ is an increasing function. This has the flavor of an escape-clause type of agreement: when α is high, the underlying import volume is high, and so with $T(\alpha)$ an increasing function the agreement $\{T(\alpha)\}$ allows for the import tariff to rise in states of the world in which the underlying import volume is high, broadly analogous to the escape clause provided in GATT Article XIX.²⁶

To understand the potential appeal of an escape clause in the current setting, recall that if S is left to discretion (as in the agreement $\{T(\alpha)\}$) then it will be used to manipulate the terms of trade, and the incentive to do so will be stronger when α is higher. A higher T mitigates the incentive to distort S for terms-of-trade purposes, and so allowing for a higher T when α is higher can help to mitigate the use of S for purposes of terms-of-trade manipulation when the incentive is highest to do so. Hence, the agreement $\{T(\alpha)\}$ provides a degree of *indirect incentive management*, and in this way our model identifies a novel rationale for the desirability of escape clauses in trade agreements: an escape-clause type agreement of the form $\{T(\alpha)\}$ can be attractive relative to a rigid agreement of the form $\{T\}$ because it provides an indirect means of managing the distortions associated with leaving S to discretion.

²⁶We say that $\{T(\alpha)\}$ has the “flavor” of an escape-clause type of agreement, because there are some important features of GATT Article XIX that are not captured by $\{T(\alpha)\}$. For instance, Article XIX includes an “injury” test, which has no counterpart in $\{T(\alpha)\}$ (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). Also, under Article XIX a country is *allowed* to raise its tariff in case of an import surge, whereas $\{T(\alpha)\}$ technically leaves no discretion on T ; but this feature can be captured by our model in a straightforward manner: as we argue in section 5, imposing the equality constraint $T = T(\alpha)$ is equivalent to imposing the inequality constraint $T \leq T(\alpha)$. Under the latter, when α is higher the government is allowed to raise T up to a higher level, but is not forced to do so.

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

Intuitively, 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 the first-best level of S 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, because the first-best level of T does not depend on α , and this effect (which is not present in the γ -uncertainty case because the first-best level of T *does* depend on γ) makes the cost of discretion *higher* when the agreement is rigid: the reason is that the adverse effects of discretion over S can be mitigated by making the value of T contingent on α , and this mitigation of the cost of discretion is not possible within a rigid contract.

These observations, together with those made in the previous subsection, suggest an important insight: the interaction between rigidity and discretion depends crucially on the *source* of uncertainty, and this dependence can be understood from the perspective of the indirect state-contingency effect (which pushes toward complementarity) and the indirect incentive-management effect (which pushes toward substitutability). When uncertainty concerns variables (such as γ) that are directly relevant for the first-best levels of both domestic instruments (S) and import taxes (T), rigidity and discretion tend to be complementary, because the indirect state-contingency effect is operative while the indirect incentive management effect is not. And when uncertainty concerns variables (such as α) that are *not* directly relevant for the first-best levels of either S or T , rigidity and discretion tend to be substitutable, because the indirect incentive-management effect is operative while the indirect state-contingency effect is not. (The remaining cases of uncertainty over variables that are directly relevant for first-best levels of either S or T but not both can be understood from this perspective as well, as we

describe in a later section).

One consequence of the substitutability between rigidity and discretion in the present stochastic environment is that the “concavity” condition implied by Lemma 1 now may be satisfied for all relevant complexity levels, and therefore we do not have a “complementary slackness” condition as in the previous section: all four candidate agreements may be part of the optimal sequence as c increases. The following proposition confirms this point:

Proposition 5. *Consider the agreement class \mathcal{A}_0 , and assume that only α is uncertain. There exist scalars c_1, c_2 , and c_3 with $0 < c_1 \leq c_2 \leq c_3 < \infty$ such that the optimal agreement is:*

- (a) *the $\{FB\}$ agreement $\{T = \bar{\gamma}; S = -\bar{\gamma}\}$ for $c \in (0, c_1)$;*
- (b) *of the form $\{T(\alpha)\}$ for $c \in (c_1, c_2)$;*
- (c) *of the form $\{T\}$ for $c \in (c_2, c_3)$; and*
- (d) *the empty agreement for $c > c_3$.*

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

Remark 2. *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 $\{T(\alpha)\}$ is optimal for some c .*

The result reported in Remark 2 is intuitive. If the degree of substitutability between T and S is sufficiently low (as when λ is low) so that leaving S to discretion is an attractive option, then an escape-clause type agreement of the form $\{T(\alpha)\}$ is optimal for a range of c as long as the added complexity 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 Proposition 4. Specifically, as Δ_α increases, the optimal agreement may switch from a contingent type $\{T(\alpha)\}$ to a rigid type $\{T, S\}$, which as Proposition 4 indicates can never happen with an increase in Δ_γ . Intuitively, this reflects the workings of the monopoly power effect and the indirect incentive management effect, and the fact that the cost of discretion (CD) is not only rising in α but also convex. The key point is that as uncertainty over α rises, CD rises, and it may therefore 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. This confirms our earlier observation that the effects of uncertainty on the optimal agreement depend on whether this uncertainty concerns state variables that are directly relevant to first best policy levels (as in Δ_γ) or rather state variables that are irrelevant to the first best policy levels (as in Δ_α).

3.3. Multidimensional uncertainty

In sections 3.1 and 3.2 we focused on the case of one-dimensional uncertainty. Here we consider how the main results are modified when both γ and α are uncertain.

A first observation is that the $\{FB\}$ agreement continues to be $\{T = \gamma; S = -\gamma\}$ and therefore costs $(4 + k)c$, just as in the case where only γ is uncertain. A second observation is that, besides the $\{FB\}$ agreement and the empty agreement, there are now five agreements that can potentially be optimal, four that we have considered already and a new one. The four that we have seen already are $\{T, S\}$, $\{T(\gamma)\}$, $\{T(\alpha)\}$ and $\{T\}$; the new one is $\{T(\alpha, \gamma)\}$. Hence, the main difference compared to the case of one-dimensional uncertainty is that there are three potentially optimal contingent- T agreements ($\{T(\gamma)\}$, $\{T(\alpha)\}$ and $\{T(\alpha, \gamma)\}$), instead of just one. Intuitively, it may be appealing to make T contingent on α , on γ or on both, depending on the exact distribution of uncertainty.

It is simple to establish sufficient conditions under which the agreement $\{T(\alpha, \gamma)\}$ is optimal for a range of contracting costs: for example, this is ensured if (i) k is sufficiently small, (ii) the cost of discretion is not too high (for example because S is a poor substitute for T), and (iii) α and γ are not perfectly correlated.²⁷

It is worth emphasizing the possibility that, even though the $\{FB\}$ agreement has T contingent on γ , the optimal agreement may be $\{T(\alpha)\}$: in the presence of contracting costs it may be optimal to make T contingent on the “wrong” state variable, i.e. the one that is not relevant for the first best policy levels.²⁸

Aside from the changes in results highlighted above, the other qualitative insights that we derived in the context of one-dimensional uncertainty generalize in a natural way to a setting where both α and γ are uncertain. For example, a higher α leads to a lower degree of discretion

²⁷Consider the extreme case $k = 0$. Then $\{T(\alpha, \gamma)\}$ dominates $\{T(\alpha)\}$, $\{T(\gamma)\}$ and $\{T\}$. Moreover, $\{T, S\}$ is dominated by $\{FB\}$. As a consequence, $\{T(\alpha, \gamma)\}$ is the only agreement that can be optimal, besides $\{FB\}$ and the empty agreement. And if the cost of discretion is not too high, $\{T(\alpha, \gamma)\}$ is optimal for some c .

²⁸In reality it is probably the case that α is easier to describe and verify, so the cost of contracting over α is lower than the cost of contracting over γ . But this would of course only strengthen our point, which holds even in the absence of such an asymmetry in the cost of contracting over state variables.

in the optimal contract, similarly to the one-dimensional uncertainty case. Also our results about the complementarity or substitutability between rigidity and discretion generalize in an intuitive way: rigidity and discretion tend to be complementary if uncertainty in γ is important relative to uncertainty in α , and vice-versa. Overall, then, the essential insights of our one-dimensional uncertainty analysis are preserved in an environment where both γ and α are uncertain.

4. The Role of the National Treatment Clause

In this section we evaluate the National Treatment (NT) clause as a means to economize on contracting costs. 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 $t_h = t_f$.²⁹ It is important to note that, while the NT provision restricts internal taxes to be the same, it allows the importing country discretion over the level at which these taxes are set. In line with our assumptions on contracting costs, we assume that including the NT clause in the agreement costs $2c_p$.³⁰

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 examine 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. As indicated above, we focus on an extended set of agreements that includes the class considered in the previous section (\mathcal{A}_0) plus the class of NT-based agreements. Letting \mathcal{A}_{NT} denote the class of NT-based agreements, we thus focus on the set of agreements $\mathcal{A}_0 \cup \mathcal{A}_{NT}$.

²⁹There are two interpretation issues that can be raised here. First, Article III.2 speaks of “treatment no less favorable,” which suggests that a more accurate formalization of the NT provision is given by the inequality constraint $t_h \geq t_f$. However, in our model this constraint would always be binding, so there would be no gain in allowing for this inequality constraint. Also, Article III.1 restricts attention to measures that are applied “...so as to afford protection...”. This sentence can be read in various ways, including that a foreign product can be taxed more heavily as long as this is motivated by legitimate policy objectives. This is not an issue 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.

³⁰It could be argued that including an NT clause in the agreement should cost less than specifying exact levels for t_h and t_f , so it might be more realistic to assume that the NT clause costs less than $2c_p$. 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.

The points we make in this section do not depend on the exact nature of the uncertainty, so we will allow for multidimensional uncertainty as in section 3.3. Also, we continue to focus on the case of a consumption externality, but the main insights would not change if we had a production externality or political-economy motives along the lines of section 6.

We begin by observing that the relationships between price wedges and policies are different for non-NT agreements and NT-based agreements. For non-NT agreements, these relationships are given by (2.1) as recorded in section 2, and within this class we can focus on agreements that tie down S and/or T . However, for NT-based agreements, the arbitrage conditions become

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

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

Notice that both the efficient outcome and the noncooperative equilibrium, derived in section 2.1 in the absence of NT, can also be implemented with policies that conform to the NT clause. In particular, the efficient policies under NT are given by $\tau^{eff} = 0$, $t^{eff} = \gamma$, and $s^{eff} = 0$, and the noncooperative equilibrium outcome can be achieved with policies that conform to NT according to $\tau^N = p^*/\eta^*$, $t_h^N = t_f^N \equiv t^N = \gamma$, and $s^N = 0$. Hence, there is no inherent violation of NT in the noncooperative equilibrium of our model, and if the NT clause has any real bite, it must be because *other* contractual obligations create incentives for the importing country to use internal taxation in a discriminatory way.

There are many kinds of NT-based agreements, but we can reduce the number that must be considered by focusing only on NT-based agreements that can be *strictly* optimal in the class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$. It turns out that the only type of NT-based agreement that can be strictly optimal is one that ties down τ and s , leaving the common consumption tax t to discretion. We denote this type of agreement by $\{NT, \tau, s\}$. The next remark states the point.

Remark 3. *Consider the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$. The only NT-based agreements that can be strictly optimal are of the form $\{NT, \tau, s\}$.*

The intuition for this result can be understood as follows. An agreement of the type $\{NT, \tau, s\}$ leaves discretion over the common consumption tax t . But expression (4.1) reveals a more fundamental feature: this agreement leaves discretion over the consumer price wedge $p - p^*$, while tying down the producer price wedge $q - p^*$. And as we remarked earlier

(see footnote 21), this kind of discretion *cannot* be generated by a non-NT agreement. This is a subtle point that bears emphasis: an agreement that imposes separate constraints on (some or all of) the policy instruments (s, τ, t_h, t_f) , or equivalently on T and/or S , cannot leave discretion over the consumer price wedge while tying down the producer price wedge. This can be done *only* by constraining the relationship between the internal consumption taxes t_h and t_f ; and imposing NT is a simple way of doing this. This is why an NT-based agreement can potentially achieve a strict improvement over non-NT agreements. It is also not hard to see that the *only* NT-based agreement that can strictly improve over non-NT agreements is $\{NT, \tau, s\}$.³¹

We next seek conditions under which an NT-based agreement of the form $\{NT, \tau, s\}$ is strictly optimal for a range of contracting cost c . We start with an intuitive discussion.

How attractive is $\{NT, \tau, s\}$ as a way to save on contracting costs? Relative to $\{FB\}$, the agreement $\{NT, \tau, s\}$ implies lower contracting costs (as long as $\Delta_\gamma > 0$). On the other hand, $\{NT, \tau, s\}$ cannot achieve the first best outcome, because it leaves discretion over consumption taxes, and this discretion will be used by governments to manipulate terms of trade. In what follows we refer to the “cost of discretion over t ” as the difference in gross global surplus between $\{FB\}$ and $\{NT, \tau, s\}$, that is $\Omega_{\{FB\}} - \Omega_{\{NT, \tau, s\}}$. Just as we saw in section 3.1 that the cost of discretion over S depends critically on the degree of instrument substitutability and the magnitude of the home-country monopoly power, it may be seen that the cost of discretion over t depends critically on these factors as well, and for the same reasons.

To see this, consider first the degree of substitutability between t and τ . Clearly, if t and τ are highly substitutable, then any constraints placed on τ (and s) through an NT-based agreement can be largely undone if t is left to discretion, just as with S and T for non-NT agreements. Importantly, though, the underlying parameter conditions that cause t and τ to be highly substitutable are a high β and/or low λ , and these are essentially *opposite* to the conditions that cause S and T to be highly substitutable (low β and/or high λ). This can be confirmed from the equilibrium terms of trade, expressed now as a function of τ , t and s :

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

When λ is close to zero, t and τ are close to perfect substitutes, and hence $\{NT, \tau, s\}$ offers little improvement over the empty agreement. If β is close to zero, t is nearly useless as a means to distort terms of trade, and hence the cost of discretion over t approaches zero.

³¹Note in particular that, since the NT clause costs $2c_p$, a contingent NT-based agreement of the form $\{NT, \tau(\cdot), s(\cdot)\}$ cannot be strictly optimal, because it costs as much as the $\{FB\}$ agreement.

Now consider the degree of home-country monopoly power. This can be shown to affect the cost of discretion over t in a similar way as the cost of discretion over S : when the home country enjoys higher monopoly power over trade, the incentive to distort t in order to manipulate terms of trade is stronger. And since a key determinant of the degree of home-country monopoly power is the import demand level, when α is higher the cost of discretion over t is higher, thus making $\{NT, \tau, s\}$ less attractive.

Having discussed at a broad intuitive level the pros and cons of NT-based agreements of the form $\{NT, \tau, s\}$, the next question is whether there exists a parameter region in which the optimal agreement is indeed of the type $\{NT, \tau, s\}$. The answer is yes. To see why, notice first that, if β is small, $\{NT, \tau, s\}$ can implement an outcome close to the first best ($t^{eff} = \gamma$, $\tau^{eff} = 0$, $s^{eff} = 0$), because the importing government will set t close to γ . Recalling the “concavity” condition of Lemma 1, it is easy to see that, if β is sufficiently small, moving from $\{NT, \tau, s\}$ to a more complex agreement can offer at best a negligible gain, while moving from $\{NT, \tau, s\}$ to a less complex agreement necessarily implies a non-negligible loss. This immediately implies that the concavity condition implied by Lemma 1 is satisfied. And since $\{NT, \tau, s\}$ is undominated in its complexity class, both conditions of Lemma 1 are met.

The following proposition records this result:

Proposition 6. *Consider the agreement class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$. If $\Delta_\gamma > 0$ and β is sufficiently small, then there is an intermediate range of c for which the optimal agreement includes the NT clause.*

Proposition 6 identifies a simple condition under which our model can rationalize the use of an NT-based agreement: an agreement of this kind is strictly optimal if the degree of substitutability between t and τ is sufficiently small, and the level of elementary contracting costs c lies in some intermediate range.³² Intuitively, this condition describes a world 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.

³²The reader might wonder whether our model can rationalize a contingent NT agreement of the form $\{NT, \tau(\cdot), s\}$ as a strictly optimal agreement. As remark 3 states, under our contracting cost assumptions the answer is no. But if we modified slightly our contracting cost assumptions by assuming that the NT clause costs less than $2c_p$ – see footnote 30 – then there would exist a parameter region where a contingent NT agreement is optimal. Incidentally, note from the NT pricing relationships in 4.1 that an agreement where both τ and s are contingent is equivalent to one where only τ or s is contingent.

5. The Role of Weak Bindings

In the previous sections we focused on agreements that impose equality constraints ("strong bindings"), as in $\{T = 2\}$ or $\{NT, \tau = 0, s = 0\}$. In a world of costless contracting where the first best outcome would be implemented, 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 make this claim more formal, but as a first step we develop some intuition through a simple example.

Suppose that only γ is uncertain, and focus on 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 configurations of parameters the latter can offer a strict improvement over the former. Let $T^N(\gamma)$ denote the noncooperative equilibrium level of T as a function of γ . Intuitively, and noting that T^N is increasing in γ , the optimal level of the strong binding \bar{T} must be below $T^N(\bar{\gamma} + \Delta_\gamma)$, but it may be above $T^N(\bar{\gamma} - \Delta_\gamma)$. If \bar{T} is above $T^N(\bar{\gamma} - \Delta_\gamma)$, then a weak binding is strictly preferable, because in the low state the government sets T below the binding, and this improves global welfare. Clearly there exist configurations of parameters for which this is the case.

The next proposition confirms and extends this intuition. In particular, we show that the above results are valid not only for the import tax T or the tariff τ , but also for the other policy instruments that the agreement may need to bind, namely the production subsidies.

The intuition for the extended result is similar: governments are tempted to distort production subsidies in import-competing industries *upwards*, and hence the relevant constraint is an upper bound on the subsidy.³³ The proposition is valid regardless of which state (γ or α) is uncertain, so we do not specify the source of uncertainty.

To express 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; \mathcal{A}^W denote the same class of agreements except that strong bindings are replaced by weak bindings; and \mathcal{A}^{RW} denote the same class except that rigid strong bindings are replaced by rigid weak bindings.

Proposition 7. (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) *Rigid weak bindings can offer a strict improvement over rigid strong bindings: $\max_{A \in \mathcal{A}^{RW}} \omega(A) > \max_{A \in \mathcal{A}^S} \omega(A)$ for some configurations of parameters.*

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 7 highlights another sense in which rigidity and discretion may be 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.³⁴

³³As applied to trade taxes, the proposition 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).

³⁴We note here that this is not the only possible explanation for the use of weak bindings. Maggi and Rodriguez-Clare (2005) propose an alternative explanation based on political-economy considerations: their basic idea is that weak bindings may be desirable because they induce lobbies to pay contributions even after the agreement is signed, since a government can credibly threaten to lower tariffs below the ceiling levels, and this in turn reduces the net return to capital in the protected sectors, thereby mitigating the allocative distortions caused by lobbying. We also note that the explanation proposed here is closely related to the one proposed in Bagwell and Staiger (2005), where weak bindings may be preferred to strong bindings in the presence of political-economy shocks that are privately observed by governments. One key difference is that private information and the absence of international transfers prevents governments from implementing the first best in the Bagwell-Staiger paper, thereby opening the possibility that weak bindings may be attractive, whereas here the presence of contracting costs makes it too costly to implement the first best. Another difference is that, while Bagwell and Staiger's model considers only import tariffs, here we establish that weak bindings are preferable to strong bindings also when applied to domestic subsidies.

6. Production Externalities and Political Economy

Thus far we have assumed that a consumption externality provides the efficiency rationale for policy intervention. In this section we consider two alternative possibilities: first, that a production externality provides the efficiency rationale for policy intervention, and second, that political motives serve this role. To keep the discussion simple, we return to the class of agreements \mathcal{A}_0 , but the arguments we present here are easily extended to allow for the possibility of NT-based agreements and weak bindings.

Consider first the case of a production externality. We set $\gamma \equiv 0$ and suppose that there is a positive production externality in the home country equal to σX with $\sigma > 0$: as before, this externality enters directly and separably into the representative home-country citizen's utility and does not cross borders. We allow that both σ and α may be uncertain.

In this environment, it is straightforward to establish that the $\{FB\}$ agreement takes the form $\{T = 0; S = \sigma\}$, so that only S is now state-contingent. Notice the difference between this case and that of an uncertain γ , where the first-best levels of *both* S and T are state-contingent. As anticipated by our discussion in previous sections, this difference has subtle implications for the nature of the optimal agreement, and we detail these below.

A first difference is the following. While the set of candidate optimal agreements (aside from the $\{FB\}$ agreement) are similar to those for the case of a consumption externality identified in subsection 3.3, except of course that $\{T(\gamma)\}$ and $\{T(\gamma, \alpha)\}$ are replaced respectively by $\{T(\sigma)\}$ and $\{T(\sigma, \alpha)\}$, the potential appeal of making T contingent on σ is distinct from the potential appeal of making T contingent on γ , and arises for reasons analogous to the potential appeal of the escape-clause-type agreement $\{T(\alpha)\}$. In particular, making T contingent on σ is potentially attractive because it provides an indirect means of managing the distortions associated with leaving S to discretion (i.e., because of the indirect incentive management effect).³⁵ Notice also that the three agreements $\{T(\alpha)\}$, $\{T(\sigma)\}$ and $\{T(\sigma, \alpha)\}$ are potentially attractive for a similar reason; intuitively, then, the performance of these agreements relative to each other will depend on the exact nature of the uncertainty, including which of the two

³⁵To see this, recall that if S is left to discretion (as for example in the agreement $\{T(\sigma)\}$) then it will be distorted from its Pigouvian level and used to manipulate the terms of trade. However, the higher is σ , the lower is the noncooperative trade volume and hence the lower will be the terms-of-trade incentive to distort S away from its Pigouvian level. The implication, then, is that the incentive to distort S for terms-of-trade reasons will be stronger when σ is lower according to the monopoly power effect. A higher T mitigates the incentive to distort S for terms-of-trade purposes, and so allowing for a higher T when σ is lower can help to mitigate the use of S for purposes of terms-of-trade manipulation when the incentive is highest to do so.

sources of uncertainty is more important.

A second difference concerns the complementarity/substitutability between rigidity and discretion. As noted previously, uncertainty about state variables that affect the first-best level of S (e.g. γ) tends to generate complementarity (through the indirect state-contingency effect), whereas uncertainty about state variables that do not affect the first-best level of T (e.g. α) tends to generate substitutability (through the indirect incentive management effect). As noted above, σ affects the first-best level of S but not T , and so both forces are at work: thus, uncertainty about σ has ambiguous implications for the complementarity/substitutability between rigidity and discretion. Since uncertainty about γ unambiguously pushes toward complementarity, broadly speaking rigidity and discretion are less likely to be complementary in the case of production externalities than in the case of consumption externalities.

Next we consider a simple political-economy extension of the model, in which each government maximizes a modified welfare function that attaches an extra weight to domestic profits. More specifically, we consider the possibility that each government maximizes an objective function of the form $\tilde{W} \equiv W + \xi \cdot \Pi$. We allow for the possibility that both the import demand level α and the political-economy parameter ξ may be uncertain.

There is a close similarity between this case and the case of production externality considered above, since the domestic producer surplus is closely related to the domestic output X . This similarity is reflected in the feature that, in both cases, the first-best outcome (from the point of view of the governments' objectives) entails a price wedge only on the producer side and not on the consumer side, and therefore the first-best level of T is zero and the first-best level of S is positive. But there is also an important difference: the first-best level of S does not depend on demand and supply parameters in the production externality case (it is equal to σ), but in the case of political economy considerations the first-best level of S is proportional to domestic output X , and this *does* depend on demand and supply parameters. As a consequence, the first-best agreement is $\{T = 0; S = S(\xi, \alpha)\}$, where $S(\xi, \alpha)$ is increasing in both arguments. Note that this agreement costs $2(2 + k)c$.

The fact that the $\{FB\}$ agreement is contingent on both state variables implies that, between $\{FB\}$ and the empty agreement, there is now one more cost class to consider, namely agreements that cost $(4 + k)c$, where both policies are contingent on a single state variable (α or γ). Also, since in this environment α is now relevant for the first-best level of S but not T , both the indirect state-contingency and incentive management effects are present with regard

to α uncertainty, and so uncertainty over α now has an ambiguous effect on the complementarity/substitutability between rigidity and discretion. Aside from these differences, the main qualitative results are the same as in the case of a production externality.

Finally, we emphasize again the possibility of an escape-clause type agreement $\{T(\alpha)\}$, because in this political-economy setting it has a particularly interesting interpretation. The $\{FB\}$ agreement dictates that governments practice free trade ($T = 0$) and offer domestic subsidies contingent on political (ξ) and import demand (α) shocks. But in the presence of contracting costs, it may be optimal to leave domestic subsidies to discretion, allow governments to interfere with trade, and permit higher trade barriers in case of a surge in import demand.

7. Conclusion

This paper takes a first step in the analysis of trade agreements as endogenously incomplete contracts. We have shown that an incomplete contracting perspective provides a novel explanation for the emphasis on border measures that has traditionally characterized real world trade agreements, and provides as well a novel explanation for the appeal of escape clauses in the presence of uncertain import demand. We have established that the nature of the optimal agreement in an incomplete contracting setting depends on features of the underlying economic environment that have simple economic interpretations: the degree of substitutability across instruments; the extent of monopoly power on world markets; the extent to which discretion facilitates indirect state-contingency; the extent to which rigidity interferes with indirect incentive management; and the source of uncertainty. Employing these features, we have identified conditions under which the two essential methods for saving on contracting costs – introducing rigidity and/or introducing discretion into the agreement – can be either complements or substitutes. We have found that the optimal agreement tends to leave less discretion over domestic instruments when trade volumes are higher. And finally, we have shown that the appeal of some of the more subtle clauses and features of the GATT/WTO, such as its NT provision and its emphasis on weak bindings, can be understood from the incomplete contracting perspective.

Our model abstracts from some important elements that should be incorporated into a more complete theory. We conclude with a brief discussion of a number of these elements, and suggest directions for further research.

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 restricted our focus to instrument-based contracts, excluding outcome-based contracts from our analysis. Outcome-based bindings of trade volumes or prices 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, as they are suggestive of attempts to economize on contracting costs.

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., unions or political lobbies). Under this alternative view, the nature of the first-best contract would be quite different, and so naturally the nature of the optimal agreement in an incomplete contracting environment is likely to be quite different as well.

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 contract. This is a difficult task, but it could add an important new dimension to our analysis.

Finally, our paper explains contract incompleteness on the basis of contracting costs, but other approaches are possible. In the contract-theoretic literature, a common approach is 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 the approach taken here with a simple example. Consider our basic model of section 2 and suppose there is a single uncertain variable, say γ . The standard approach would be to 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 (T, S) , from which the (importing) government can choose. With one-dimensional uncertainty, this contract is typically a nonlinear function,

which we can represent as $g(T, S) = 0$. 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 $g(T, S) = 0$, whereas in the present paper we have focused on a simpler class of contracts for tractability reasons.

Notice that, as a consequence of the above differences in assumptions, the predictions are also very different across the two approaches. In particular, the standard approach predicts that the optimal contract always takes the form $g(T, S) = 0$; thus, the optimal contract is never directly contingent on state variables such as γ , and it always includes *all* policy instruments, because contracting over policy instruments 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 $g(T, S) = 0$, thus achieving the best of both approaches. We see this as an ambitious avenue for future research.

8. Appendix

Proof of Lemma 1: We first argue that conditions (i) and (ii) are both necessary, then we argue that the two conditions together are sufficient. The necessity of condition (i) is obvious. Consider condition (ii). For \hat{A} to be an optimal agreement for some c there must be a value of c such that, for any agreement A' such that $m' < \hat{m}$,

$$\frac{\Omega(\hat{A}) - \Omega(A')}{\hat{m} - m'} \geq c \quad (8.1)$$

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

$$\frac{\Omega(A'') - \Omega(\hat{A})}{m'' - \hat{m}} \leq c \quad (8.2)$$

and consequently that

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

which can be rewritten as (3.1).

To see the sufficiency part, suppose that \hat{A} is optimal in its complexity class and condition (3.1) holds for all agreements $A', A'' \in \mathcal{A}_0$ such that $m' < \hat{m} < m''$. Then there must be a value of c such that (8.1) holds for all agreements such that $m' < \hat{m}$, and (8.2) holds for all agreements such that $\hat{m} < m''$. ■

Proof of Proposition 1

We will argue that an agreement that constrains only S cannot achieve higher gross global welfare than the noncooperative equilibrium (T^N, S^N) , for any state of the world. The maximal gross global welfare that can be achieved by this type of agreement is given by $\max_S W^G(T^R(S), S)$.³⁶ The first-order condition for this problem yields $T^{R'}(S) = -W_S^G/W_T^G$ (the second-order condition can be shown to hold). This condition requires that the slope of $T^R(S)$ be equal to the slope of an iso- W^G curve in (T, S) space. It is direct to verify that $T^{R'}(S) = -\frac{\lambda}{\beta + \lambda}$. The slope of an iso- W^G curve is given by $-\frac{W_S^G}{W_T^G} = \frac{W_S + W_S^*}{W_T + W_T^*}$. However, at the noncooperative equilibrium we have $W_S = W_T = 0$ and hence

$$-\frac{W_S^G}{W_T^G} = -\frac{W_S^*}{W_T^*} = -\frac{\frac{dW^*}{dp^*} \cdot \frac{dp^*}{dS}}{\frac{dW^*}{dp^*} \cdot \frac{dp^*}{dT}} = -\frac{\frac{dp^*}{dS}}{\frac{dp^*}{dT}} = -\frac{\lambda}{\beta + \lambda} \quad (8.3)$$

³⁶Since we are focusing on a given state of the world, we do not have to make the state of the world explicit in the notation.

Therefore, the slope of the iso- W^G curve at the noncooperative equilibrium equals the slope of $T^R(S)$, hence the level of S that maximizes $W^G(T^R(S), S)$ is given by S^N . It follows that an agreement that constrains only S cannot achieve greater surplus than $W^G(T^R(S^N), S^N)$, which is just the noncooperative equilibrium surplus $W^G(T^N, S^N)$. ■

Proof of Proposition 7

(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. The arguments made in this proof are valid whether uncertainty concerns γ or α or both, so we will generically refer to the “state” to indicate the realization of the uncertain vector. Also, we will omit the uncertain parameters from the arguments of the relevant functions, as this should not cause confusion.

Agreement \tilde{A}^S must be one of the following types: (a) $\{T = T(\cdot)\}$; (b) $\{T = T(\cdot); S = S(\cdot)\}$; or (c) $\{NT; \tau = \bar{\tau}; s = \bar{s}\}$. A dot in parenthesis means that the constraint may be contingent.

Let us start with case (a). Consider replacing $\{T = T(\cdot)\}$ with $\{T \leq T(\cdot)\}$. This can decrease Ω only if in some state the government chooses $T < T(\cdot)$ and this implies a lower level of Ω than $T = T(\cdot)$. 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 $\frac{d}{dT}p^*(T, S^R(T))$. It is direct to verify that this derivative is negative, which in turn implies $\frac{d}{dT}\Omega(T, S^R(T)) < 0$ for $T > T^N$. We can conclude that switching to a weak binding cannot decrease Ω .

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

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

(ii) it chooses $(T = T(\cdot), S = S^R(T))$: 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 .

(iii) it chooses $(T = T^R(S), S = S(\cdot))$: 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 .

(iv) 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 $\{T = T(\cdot); S = S(\cdot)\}$ to $\{T \leq T(\cdot); S \leq S(\cdot)\}$ cannot decrease Ω .

Finally, consider case (c). Since the NT agreement fixes the wedge $q - p^*$ and leaves the wedge $p - p^*$ discretionary, it is convenient to re-define variables as follows: $p - p^* \equiv z$ and $q - p^* \equiv v$. We can think of z and v as the policy instruments and of the NT agreement as imposing a constraint $v = \bar{v}$. Also, it is useful to rewrite the world price as a function of v and z as $p^* = \frac{1}{\tau}(\alpha + \alpha^* - \beta z - \lambda v)$.

Let us now replace the agreement $\{NT; \tau = \bar{\tau}; s = \bar{s}\}$ with $\{NT, \tau \leq \bar{\tau}, s \leq \bar{s}\}$. Using the new notation, this is equivalent to replacing the constraint $v = \bar{v}$ with the constraint $v \leq \bar{v}$. We can apply a similar argument as for case (a): it suffices to show that, for any given state, $\Omega(v, z^R(v))$ is decreasing in v for $v > v^N$ (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 $v > v^N$. Also, $\frac{d}{dv}W^*(v, z^R(v))$ has the same sign as $\frac{d}{dv}p^*(v, z^R(v))$. It can be verified that $\frac{d}{dv}p^*(v, z^R(v)) = -\frac{\lambda}{\beta + \beta^* + \lambda^*} < 0$. This implies that switching to weak bindings cannot decrease Ω .

(ii) To prove this claim it suffices to show that there is some configuration of parameters 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 the noncooperative level of a policy is below the (rigid) binding for that policy. It is easy to show that there exists a configuration of parameters for which this is the case. ■

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] Battigalli, Pierpaolo and Giovanni Maggi, "Rigidity, discretion and the costs of writing contracts," *The American Economic Review*, vol. 92(4), pp. 798-817, 2002.
- [5] Battigalli, Pierpaolo and Giovanni Maggi, "International Agreements on Product Standards: an Incomplete-Contracting Theory," NBER Working Paper No. 9533, 2003.
- [6] Copeland, Brian, "Strategic Interaction Among Nations: Negotiable and Non-negotiable Trade Barriers," *Canadian Journal of Economics*, vol. 23, pp. 84-108, 1990.
- [7] Costinot, Arnaud, "A Comparative Institutional Analysis of Agreements on Product Standards," mimeo, 2004.
- [8] Ederington, Josh, "International Coordination of Trade and Domestic Policies," *The American Economic Review*, vol. 91(5), pp. 1580-93, 2001.
- [9] Horn, Henrik, "National Treatment in the GATT," *The American Economic Review*, vol. 96(1), pp. 394-404, 2006.
- [10] 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.
- [11] Hudec, Robert, *The GATT Legal System and World Trade Diplomacy* (2nd Ed.), Salem NH: Butterworth, 1990.
- [12] Maggi, Giovanni and Andres Rodriguez-Clare, "A Political-Economy Theory of Trade Agreements," NBER Working Paper No. 11716, 2005.

- [13] Markusen, James R., "International Externalities and Optimal Tax Structures," *Journal of International Economics*, vol.5, pp. 15-29, 1975.
- [14] 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.
- [15] 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.




Figure 1




Figure 2a




Figure 2b