The GATT/WTO as an Endogenously Incomplete Contract

Henrik Horn (Stockholm University and CEPR)
Giovanni Maggi (Princeton University and NBER)
Robert W. Staiger (Stanford, University of Wisconsin and NBER)*

August 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 seminar participants at Calgary, CEMFI, Ente Luigi Einaudi, Penn State, Princeton, UBC, UCSD and Yale, as well as participants in an ERWIT conference and in a conference at CREI (Universitat Pompeu Fabra). 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)[1] – is obviously a highly incomplete contract. A sizeable economic 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.[2] But while this literature has advanced our understanding of the *consequences* of the incompleteness of trade agreements, it cannot *explain* the particular form that the incompleteness has taken in the GATT/WTO, because the incompleteness is assumed rather than endogenously derived.

The broad purpose of this paper is to take the analysis of the GATT/WTO as an incomplete contract one step further, by *endogenously* determining the choice of contract form. We believe that an incomplete contracting perspective can help to shed light on core features of the agreement, including:

1. The agreement binds the level of trade instruments, whereas domestic instruments are left largely to the discretion of national governments. But there are 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 mechanisms" 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 transpires from the above features, is that the agreement displays an interesting combination of *rigidity*, in the sense that contractual obligations are largely insensitive to changes in economic (and political) conditions, and *discretion*, in the sense that governments have substantial leeway in the setting of many policies.

In this paper we propose a simple theoretical framework where the incompleteness of the agreement, and in particular the 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 has been structured along the lines described above.

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 measures – that need to

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

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

be constrained in order 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 in designing a trade agreement 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 constitute a problem if contracting were costless. But in reality there are important costs associated with forming a trade agreement. While these costs can take many different forms, it is probably safe to say that they tend to be higher when the agreement is more detailed, both in terms of the number of policies that it seeks to constrain and in terms of the number of contingencies that it specifies. The modeling approach we take in this paper is to explicitly incorporate these two forms of contracting costs into an analysis of the optimal structure of a trade agreement. The main message of the paper is that the presence of contracting costs can contribute to explain the core features of the GATT/WTO highlighted above.

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.[3] 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.

In what follows we outline the structure of our model and the main results.

We work within a competitive two-country setting, where countries may experience a consumption externality. The role of this externality in the model is to provide an efficiency rationale for policy intervention.[4] For simplicity, we focus on intervention in import sectors, and assume that governments have access to a full set of taxation instruments, namely: import tariffs, distinct consumption taxes on domestically-produced products and on imported products, 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

[3] 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.

[4] As we point out in a later section, the model would generate similar insights if, instead of a consumption externality, we had a production externality or political-economy motives in the governments' objectives (in the form of an extra weight attached to the domestic producers' surplus). What is essential to the model is the presence of an (economic or political) efficiency rationale for policy intervention.

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 very simple way. Following an approach similar to that of Battigalli and Maggi (2002), we assume that these costs are an increasing function of the number of state variables and policies included in the agreement, and we seek to characterize the agreement that maximizes 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 Nash equilibrium – and the other is the first-best outcome. In the absence of an agreement, the importing country would use its policy instruments to 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 externalities, the first best agreement would be very simple: it would just stipulate laissez-faire across all policy instruments and under all circumstances. But due to the externality, the contracting problem is substantially more complex: the first best agreement will now involve the use of domestic 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 in examining the optimal degrees of rigidity and discretion in the trade agreement, and how these depend on contracting costs, on the degree and type of uncertainty and on the features of the underlying economy.

Our first result is that the optimal agreement tends to leave more discretion on domestic policy instruments than on border measures. More specifically, while for a range of contracting costs it is optimal to bind import taxes while leaving domestic instruments to discretion, we show that it is *never* optimal to leave import taxes to discretion and contract only over domestic instruments. This result accords well with a basic feature of the GATT/WTO, namely, the traditional emphasis on border measures over domestic instruments; 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 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 (no rigidity or discretion), then it becomes increasingly rigid and/or discretionary, and eventually it becomes optimal to have no agreement at all. Whether the optimal agreement tends to feature rigidity or discretion, or a combination of the two, depends crucially on the nature of the uncertainty and on the demand and supply conditions. Intuitively, rigidity is relatively more attractive when uncertainty is small. On the other hand, discretion (over domestic instruments) is relatively more attractive when (i) domestic instruments are less effective at manipulating terms of trade, or in other words, *the degree of substitutability* between these instruments and import taxes is lower, since in this case the *ability* to manipulate terms of trade through domestic instruments is lower; and (ii) the importing country has less *monopoly power* in trade (which is the case when the import demand level is lower), since in this case the *incentive* to distort terms of trade through domestic instruments is lower.

As a consequence of the monopoly-power effect described just above, the model predicts that the optimal degree of discretion is lower if the import demand level, and hence the volume of trade, is higher. 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 the 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 described just above, 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 depends in subtle ways on its source. We find that, if uncertainty concerns primarily the level of the externality – or more generally state variables that are directly relevant for the first-best policy levels – then rigidity and discretion are *complementary* ways of saving on contracting costs: more specifically, the cost of discretion is lower in the presence of rigidity than in the absence of it. In this case, rigidity and discretion tend to co-vary as underlying exogenous parameters vary. But if uncertainty concerns primarily the level of import demand – or more generally state variables that are not directly relevant for the first-best policy levels – rigidity and discretion are *substitutable*, meaning that the cost of discretion is higher when the agreement is rigid.

Also, to the extent that uncertainty concerns the level of the externality, it has an important effect on the cost of discretion: leaving domestic instruments to discretion is an indirect way to introduce state-contingency in the agreement, since the unilateral setting of domestic instruments varies in the "right" way with the level of the externality; in this sense, the presence of uncertainty mitigates the cost of discretion.

When there is substantial uncertainty about the level of import demand, we find that it may be optimal for the agreement to specify an escape-clause type rule, whereby governments are allowed to raise tariffs when the level of import demand is higher. However, our rationale for an escape clause is distinct from those that have been highlighted in the existing theoretical literature.[5] In particular, an escape clause can be appealing in our model when the agreement leaves discretion over domestic instruments. Since, as we have observed above, the incentive to distort these instruments for manipulating the terms of trade is stronger when the import demand level is higher, allowing higher tariffs in the high-import-demand states can be attractive as a way to mitigate this incentive.

In the final part of the paper we extend the analysis to shed light on two other core aspects of the GATT/WTO that we believe are best understood from an incomplete contracting

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

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

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 externality 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). The Home country will be a natural importer of good 1 and the Foreign country 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:

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

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

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

The associated profit function for good j is

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

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

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

On the demand side, we assume that the representative citizen's utility function is linear in the numeraire good and separable in the non-numeraire goods. Also, we allow for the possibility of a (negative) consumption externality. The role of this externality in the model is to create an efficiency rationale for policy intervention. As we will 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 (in the form of an extra weight attached to the domestic producers' surplus). What is important for the points we want to make is the existence of a (political or economic) efficiency rationale for policy intervention.

We assume that the externality is linear and does not cross borders.[6] Also, each good generates an externality only in the country that imports it; thus, good 1 generates an externality in Home, and good 2 in Foreign.[7]

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

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

where lower-case cs denote individual consumption and upper-case Cs denote aggregate consumption. 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. We write the demand functions as:

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

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

---

[6] Hence, we model the consumption externality as a purely domestic "eyesore" pollutant along the lines of the modeling of production externalities discussed in Markusen (1973) and especially Ederington (2001).

[7] This assumption is unrealistic but is not essential for our main points. Notice also that, if we had political-economy motives instead of a consumption externality, such an asymmetry would be relatively natural, as it would capture situations where import-competing interests are organized but export interests are not.

The consumer surplus associated with good j in Home and Foreign can respectively be written as:

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

We assume that each government can intervene only in its import sector. We allow each government to use a full set of taxation instruments in the import sector, namely: an import tariff (τ), an internal tax on consumption of the domestically produced good (t_h), an internal tax on consumption of the imported product (t_f), and a production subsidy to domestic firms (s). All instruments are expressed in specific terms. The reason we allow for a relatively rich set of policies is that we want to represent in a stylized way a situation in which governments can use a variety of instruments to act opportunistically and manipulate terms of trade. Such a situation calls for an agreement that constrains a whole host of policies, thus creating a tradeoff between the benefits of contracting over many policies and the costs of doing so.

At this point we impose a strong symmetric structure on the model: in particular, we assume that the two non-numeraire sectors are mirror-images of each other. The symmetric and separable structure allows us to focus on a single sector, say sector 1, and drop subscripts from now on. While the analysis will focus only on sector 1, where Home is the natural importer, the reader should keep in mind that in the background there is a mirror-image sector where the equilibrium conditions are exactly the same, except that the two countries' roles are reversed. The symmetry of the model is not crucial for our results, and could be relaxed at the cost of extra notation.

The goods markets in the two countries are integrated and there are no trade costs, hence prices differ only to the extent of government intervention. Throughout we focus on non-prohibitive levels of government intervention that do not choke off all trade. In the sector under consideration, due to the absence of taxation by the foreign government, producer and consumer prices in the foreign country 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 market as it receives after taxes for sales in the home-country market: $p^* = p - \tau - t_f$. And finally, the relationship between the home-country producer price and the home-country consumer price is given by $q = p - t_h + s$.

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

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

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

Market clearing requires that world demand equal world supply, or

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

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

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

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

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

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

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

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

We will henceforth assume that 2.3 holds.

We assume that each government's objective corresponds to 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 the welfare function on the policy instruments.

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 Nash equilibrium and efficient policies

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

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

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

From the above system, we may derive the following expressions for the Nash equilibrium import tax and effective production subsidy choices of the home government, which we denote by T^{NE} and S^{NE}, respectively:[8]

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

where η^* is the elasticity of the foreign export supply (itself evaluated at S^{NE} and T^{NE}).[9]

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

We focus throughout on parameter combinations for which strictly positive trade occurs in the Nash equilibrium. It is direct to verify that the Nash equilibrium trade volume is strictly positive in the absence of the externality (i.e., when $\gamma = 0$), and that the restriction to positive volumes in effect places upper limits on the magnitude of the externality parameter γ.

Having characterized the Nash equilibrium policy choices, we turn next to the globally efficient policies. The globally efficient policies are those policies that maximize "global welfare," that is, the sum of home and foreign welfare:[10]

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

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

[9]For completeness we report here the explicit expression for T^{NE} in terms of the underlying parameters: $T^{NE} = \frac{(\beta^*+\lambda^*)[\alpha+\Upsilon\gamma+\lambda\gamma]-(\beta+\lambda)\alpha^*}{(\beta^*+\lambda^*)[\Upsilon+(\beta+\lambda)]}$.

[10]In 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.

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

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

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 Nash choice of effective production subsidy is efficient ($S^{NE} = S^{eff}$), and the Nash equilibrium is inefficient only in that import taxes are too high – and Nash trade volumes are therefore too low – relative to their efficient levels ($T^{NE} > T^{eff}$). The inefficiently high level of T reflects in turn the unilateral incentive to manipulate the terms of trade with the choice of import taxes. Therefore, it is accurate to say that the potential gains from contracting in this setting arise entirely from the ability to control the incentive to utilize import taxes to manipulate the terms of trade.[11]

2.2. Uncertainty

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

We will consider two possible sources of uncertainty: the consumption externality (γ) and the level of domestic demand (α). These are two natural types of uncertainty to consider: uncertainty about γ can be interpreted as uncertainty about the true efficiency rationale for policy intervention, while shocks in α can be interpreted as shocks in the volume of trade, that are *per se* not relevant for the efficient policy levels. As will become clear, these two types of uncertainty have different implications for the optimal design of a trade agreement. Focusing on uncertainty only in γ and α while abstracting from other sources of uncertainty is helpful to illustrate some general principles on how the optimal agreement depends on the nature of uncertainty. In section 3.3 we will discuss how results can be generalized 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 additional structure on the nature of uncertainty, though we will later do so to derive sharper results.

[11] As a consequence of this feature – which is quite general, as argued in Bagwell and Staiger (2001) – we will refer to international agreements in this setting 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.

We will consider the following simple timing: (1) the agreement is drafted: (2) uncertainty is realized; 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.[12]

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

2.3. The costs of contracting

Before we formalize the costs of contracting, we need to specify what type of contracts we will consider. Throughout the paper we will focus on *instrument-based* agreements, i.e. agreements that impose (possibly contingent) constraints on policy instruments. In the concluding section we will briefly discuss the possibility of *outcome-based* agreements, i.e. agreements that impose constraints on equilibrium outcomes such as prices or trade volumes.[13]

We propose a very stylized way of formalizing the contracting costs associated with a trade agreement. 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.[14] A broader interpretation of these contracting costs might also include negotiation costs: it is reasonable to think that negotiation costs are higher when there are more policy instruments on the table, and when there are more relevant contingencies to be discussed.

The cost of contracting over a state variable is c_s and the cost of contracting over a policy variable is c_p. We assume that, if a variable is included in the agreement, the associated cost is incurred only once, regardless of how many times that variable is mentioned in the agreement;

[12] Also implicit here is the assumption that governments can costlessly observe the realized values of the uncertain parameters. One might object that this is not very realistic for externality parameters, which are presumably difficult to ascertain even for a government, let alone the WTO courts. But our results will be 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 ojection would not have much force if we had political-economy motives rather than externalities, because governments are likely to know well the political-economy pessures they face.

[13] We 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.

[14] The 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.

in other words, there is no cost in "recalling" a given variable after the first time it appears in the agreement.

Summarizing, the cost of writing an agreement is given by

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

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

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

Example 1: The agreement $\{\tau + t_f = 3\}$, or $\{T = 3\}$, 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 $2c_p + c_s$.

The linearity of the contracting cost function is not essential for the main results. We could, at the cost of heavier notation, allow C to be a more general function of n_s and n_p, insisting only that it be increasing in each argument. We chose the linear specification to simplify the analysis and the exposition of our results.[15]

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

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.[16] 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 ($T = 0$ if $\gamma \leq 1$) then a clause $T = \gamma$). While our approach captures some of the most salient features of contracting costs, it abstracts from these and other more subtle but also potentially important considerations. On balance, however, we believe that our basic assumption that contracting costs are increasing in the number of state variables and policies included in the agreement is likely to be preserved in any reasonable model 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.

[15] The only results that could change under a more general cost function are the ones concerning the complementarity/substitutability between rigidity and discretion. As we will argue, rigidity and discretion tend to be complementary if the source of uncertainty is γ and substitutable if the source of uncertainty is α. These results may change if the cross-derivative of C with respect to n_s and n_p is high in absolute value. For example, we will show that, if the source of uncertainty is γ, discretion and rigidity are complementary; but if the cross-derivative of C with respect to n_s and n_p is positive and high, this can overturn the sign of the interaction between rigidity and discretion and make them substitutable.

[16] For 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 Maggi and Battigalli (2002) is a special case of our approach in which $C \equiv c(n_s + n_p)$.

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 often label this simply the "first-best" agreement, or the $\{FB\}$ agreement. In a similar vein, we will refer to the case of no agreement as the "empty agreement," which formally is denoted $\{\emptyset\}$. Finally, an *optimal agreement* is an agreement that maximizes expected global welfare net of contracting costs (henceforth simply "expected net global welfare"), that is $\omega \equiv \Omega - C$.

A natural and convenient way to characterize the optimal agreement is to track how the optimal agreement 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, which we henceforth refer to simply as "contracting costs." In much of the analysis to follow, we keep k fixed and consider changes in c. 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 parameter m can be seen as a rough measure of the "complexity" of the agreement, in the sense that it captures the number of policy variables and states involved (with the latter weighed by the parameter k).[17]

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.[18] 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 $m(A') \leq m(\hat{A}) \leq m(A'')$,

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

[17] The reason we focus on a proportional change in contracting costs – rather than varying contracting costs in some other way – is that we want a comparative-statics exercise that takes us from the benchmark of zero contracting costs (so that the first-best agreement is optimal) to that of prohibitive costs (so that the empty agreement is optimal). Varying c_s and c_p in a proportional way is a natural way to do so.

[18] We will later define a particular feasible set of agreements for consideration, but for now we remain noncommital on what the feasible set is since our first results hold for any feasible set.

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

Figure 1 illustrates the content of Lemma 1. With Ω on the vertical axis and m on the horizontal axis, each point in the figure is associated with a feasible agreement A, and the collection of points defines the set $\mathcal{F}$. The two conditions of Lemma 1 are illustrated for the candidate agreement $\hat{A}$. The first condition is obvious: the candidate agreement must not be dominated by some other agreement in its complexity class in terms of gross global welfare Ω. 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. The second condition is a kind of concavity requirement on the Ω function. To see 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 $\hat{A}$ in Figure 1 satisfies condition (ii) as well. By contrast, the agreement labeled $\tilde{A}$ in Figure 1, while satisfying condition (i), does not satisfy condition (ii), and so is not an optimal agreement for any c.[19]

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*. If this were not the case, then if it paid to move from A' to the more complex $\hat{A}$ it would also pay to take the further step to the even more complex A'', in which case $\hat{A}$ would not be optimal for any c.

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

Our second result concerns the relationship between the general 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 general 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, when we view them as contractual variables they are perfect *complements*: constraining one of the two instruments but not the other would have no effect. The same is true for the domestic instruments s and t_h, which matter only through their 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 $2c_p$.

[19] Note 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 $\mathcal{F}$. This is reflected in figure 1: the dashed line is the upper boundary of the convex hull of $\mathcal{F}$; agreement $\hat{A}$ is on the dashed line, while agreement $\tilde{A}$ is below it.

We may now state our third result: if an agreement is to achieve any improvement over the Nash 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 Nash equilibrium, and therefore cannot be an optimal agreement.*

At a broad level, the intuition for this result is very simple. Contracting over S alone is useless because, as we have emphasized above, the inefficiency that arises in the noncooperative equilibrium concerns T, not S. We can also be a little more precise about the logic behind Proposition 1. Starting from the Nash equilibrium, consider a small exogenous change in S. Since the Nash tariff is the one that implements the optimal terms of trade p^* (from Home's point of view), constraining S away from Home's reaction curve triggers a change in T that brings p^* back to its optimal level. Recalling that Home's policies affect Foreign welfare only through terms of trade, this implies that Foreign welfare is unchanged. And since the imposition of a constraint on S obviously reduces Home welfare, we can conclude as well that global welfare goes down as a consequence. Thus a small exogenous change in S cannot improve over the Nash equilibrium.[20,21]

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, one has to consider agreements that place constraints on only a subset of these instruments, and the result of Proposition 1 then gains relevance. In particular, Proposition 1 implies that if contracting is costly it can never be optimal to constrain internal measures but not import taxes.

At this point we will impose more structure on the stochastic environment in order to derive sharper predictions on the nature of the optimal agreement. To make our points in the most transparent way, 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 will consider the more general case of multidimensional uncertainty.

3.1. Uncertainty about the consumption externality

In this section we focus on the case where only γ is uncertain. For the sake of expositional simplicity, we assume that it can take two possible values with equal probability: a high realization $\bar{\gamma} + \Delta_\gamma$ and a low realization $\bar{\gamma} - \Delta_\gamma$, with $\Delta_\gamma > 0$.

We also impose at this point more structure on the set of agreements under consideration. In particular, for the remainder of this section we restrict our search to agreements that impose

[20] Notice 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.

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

separate equality constraints on T and S. To be concrete, we allow for clauses of the type $(T = \gamma)$ or $(S = 10)$, but not for clauses of the type $(T + S = \sigma)$ or for inequality constraints of the type $(T \leq 1)$.[22] We label this class of agreements $\mathcal{A}_0$. In later sections we will consider broader classes of agreements.

The first step is to derive the $\{FB\}$ agreement. It is clear that the agreement $\{T = \gamma; S = -\gamma\}$ implements the first best outcome. This agreement has $n_s = 1$ and $n_p = 4$ and therefore costs $c \cdot (4 + k)$. But it might be conjectured that the first best outcome could also be implemented without constraining S, and therefore be accomplished more cheaply, since as we have noted previously only T is set inefficiently in the Nash equilibrium. This conjecture is incorrect, but it is instructive to see why. The reason is simply that an agreement that dictates an import tax level T but leaves discretion over S would permit the home government to choose S according to the best-response function $S^R(T)$. As we have observed previously, this best-response is equal to S^{eff} when the home government's choice of T is unconstrained, but more generally is it direct to verify that

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

From this equation it follows that the difference between $S^R(T)$ and S^{eff} is directly proportional to the distance between T and the optimal unilateral import tax T^{NE}. As a consequence, an agreement that attempts to move T towards its efficient level without also constraining S will cause S to become distorted for terms-of-trade purposes.

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

Next we look for the optimal agreement, that is the agreement that maximizes expected net global welfare in $\mathcal{A}_0$. The $\{FB\}$ agreement yields expected net global welfare equal to $\Omega(T = \gamma, S = -\gamma) - (4 + k)c$. Clearly, when c is sufficiently small the $\{FB\}$ agreement is optimal. And it is also clear that, if c is sufficiently high, the empty agreement (which costs nothing and yields expected global welfare $\Omega(T = T^{NE}, S = -\gamma))$ is optimal. The basic question is then what happens between these two extremes: What is the optimal way of restructuring the agreement as contracting costs rise from zero?

Even though uncertainty is one-dimensional, the optimal restructuring of an agreement to economize on overall contracting costs can in principle be quite onerous. This difficulty is exacerbated by the inherent non-separability of the contracting problem that we study. This

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

feature, which is a consequence of the fact that internal policy measures (S) can at least imperfectly substitute for import taxes (T), precludes characterizing the optimal contract "task by task" (in our case, instrument by instrument) as in Battigalli and Maggi (2002).

However, knowledge of the $\{FB\}$ agreement, together with the result of Proposition 1, simplifies the problem considerably. Specifically, between the $\{FB\}$ agreement which costs $(4+k)c$ and the empty agreement which costs nothing, there are three cost classes of agreements that warrant consideration: agreements costing $(2+k)c$; agreements costing $4c$; and agreements costing $2c$. Moreover, Proposition 1 implies that we can ignore agreements that constrain S but not T, therefore we only have three kinds of agreements to consider, in addition to $\{FB\}$ and $\{\emptyset\}$: (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 contractual obligations do not include state contingencies that are included in the first-best agreement; and *discretion*, which occurs when policies that are included in the first-best agreement are missing. We can thus say that the agreement $\{T, S\}$ is rigid, the agreement $\{T(\gamma)\}$ features discretion (over S), and the agreement $\{T\}$ is both rigid and discretionary.

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

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

Next we examine at a broad intuitive level what determines the cost of rigidity and the cost of discretion, and how they depend on the fundamentals of the contracting problem.

Let us start with the cost of rigidity. Intuitively, the cost of rigidity should be higher when the environment is more uncertain, i.e. when Δ_γ is higher. A rigid agreement "gets it right" only on average, and therefore when the environment is more uncertain the cost of rigidity should be higher.

The determinants of the cost of discretion (over S) are more subtle.[23] Intuitively, the cost of this type of discretion is affected in an important way by two features of the contracting environment.

First, the cost of discretion is higher when the effective subsidy S is a closer substitute

[23] Note 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 $\mathcal{A}_0$.

for the import tax T. Recall that the gains from contracting in this setting derive entirely from the ability to control the incentive to manipulate the terms of trade. It follows, then, that when S and T are very similar instruments with regard to their effect on the terms of trade, the cost of discretion over S is high. Thus a key determinant of the cost of discretion is *the degree of substitutability between policy instruments*. Recalling that the terms of trade is given by $p^*(T,S) = [\alpha + \alpha^* - (\beta+\lambda)T - \lambda S]/\Upsilon$, a rough measure of this substitutability is given by the marginal rate of substitution between T and S with respect to the terms of trade: $\frac{\partial p^*/\partial T}{\partial p^*/\partial S} = \frac{\beta}{\lambda} + 1$. This suggests that S is a closer substitute for T when demand is less elastic (β is low) and when supply is more elastic (λ is high).[24]

A second feature of the contracting environment that affects the cost of discretion is the wedge between the Nash and efficient tariff, which in turn reflects the degree of home-country *monopoly power* as measured by the (inverse of the) foreign export supply elasticity η^*. If the home country has large monopoly power over trade, then the wedge between the Nash and efficient tariff is large; therefore, the incentive to alter policies to manipulate the terms of trade is high, and hence the cost of discretion is high. Notice that the home country's monopoly power is directly linked to the volume of trade, and thus a primary determinant of its magnitude is the level of import demand, α. When α is higher, η^* is lower, and therefore the cost of discretion should be higher.

The next important observation is that rigidity and discretion *interact* in non-trivial ways: the cost of rigidity depends on whether or not discretion is present in the agreement, and the cost of discretion depends on whether or not the agreement is rigid. To describe this interaction, we introduce the following concepts:

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 important manifestation of the interaction between rigidity and discretion is the follow-

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

ing: while the cost of discretion absent rigidity, CD, is always positive, the cost of discretion in the presence of rigidity, $\overline{CD}$, *may be negative*. In other words, it is possible that $\Omega_{\{T,S\}}$ is lower than $\Omega_{\{T\}}$: conditional on the agreement being rigid, introducing discretion may increase gross global welfare. Intuitively, introducing discretion (in S) is a way of introducing state-contingency in the agreement, and this 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, from the perspective of gross global welfare, *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.[25]

The discussion above suggests also that there is a third determinant of the cost of discretion, in addition to the instrument-substitutability effect and the monopoly-power effect, but one that applies only to the cost of discretion in the presence of rigidity ($\overline{CD}$): the extent to which discretion introduces state-contingency into a rigid agreement. This effect is stronger when there is more uncertainty in γ and – it is worth emphasizing again – is responsible for the complementarity between rigidity and discretion.

It should be mentioned here that the implicit-state-contingency effect highlighted above, and the complementarity between rigidity and discretion that is implied by it, depends crucially on the fact that uncertainty concerns a state variable that is directly relevant for the first-best policy levels (in the specific case, γ). As we will explain in the next subsection, if uncertainty concerns a state variable that is *not* directly relevant for the first-best policy levels (e.g. the domestic demand level α), then rigidity and discretion are substitutable. Thus it should be kept in mind that whether rigidity and discretion are complementary or substitutable depends critically on the nature of uncertainty.

We are now ready to return to the question we posed above: What is the optimal sequence of agreements as c increases from zero? This question can be answered by using Lemmas 1 and 2, together with the complementary property that we just highlighted.

Lemma 2 tells us that, as c increases, we must move from more complex to less complex agreements. This immediately implies that the optimal sequence of agreements is a subsequence of $(\{FB\}, \{T,S\}, \{T(\gamma)\}, \{T\}, \{\emptyset\})$. We say "subsequence" because each of these agreements (except $\{FB\}$ and $\{\emptyset\}$) may be "skipped" over as c increases.[26]

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

[25] In Section 5 we highlight a second sense in which rigidity and discretion may be complementary: there we show that, conditional on the agreement being rigid, it may be a good idea to impose an upper bound on the import tax T (weak binding), rather than an equality constraint. In other words, if the agreement is to be rigid, it may be valuable to allow not only for discretion over S, but also for some *downward* discretion on T.

[26] In 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.

optimal sequence. Then Figure 2a, which depicts the relevant (m, Ω) combinations, 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 $\{FB\}$. It is straightforward to show that this 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 the agreement $\{T(\gamma)\}$ 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). Clearly, 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. In other words, because of the complementarity between rigidity and discretion, returns to complexity must be *increasing* around (at least) one of these two agreements.

Figures 2a and 2b reflect this point: if, as in Figure 2a, $\{T, S\}$ is positioned above the dotted line, $\{T(\gamma)\}$ must be positioned below it; and if, as in Figure 2b, $\{T(\gamma)\}$ is positioned above the dotted line, $\{T, S\}$ must be positioned below it.[27] Intuitively, the agreement $\{T\}$ features both rigidity *and* discretion. But the complementarity between rigidity and discretion means that these two features are less costly when they are featured together than when they are featured 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 be positioned above the line connecting $\{T\}$ and $\{FB\}$.

We summarize these results in the following proposition:

Proposition 2. *Consider the agreement class $\mathcal{A}_0$, and assume that only γ is uncertain. There exist scalars c_1, c_2, c_3 and c_4 with $0 < c_1 \leq c_2 \leq c_3 \leq c_4 < \infty$ such that the optimal agreement is:*
(a) the $\{FB\}$ agreement for $c \in (0, c_1)$;
(b) of the form $\{T, S\}$ for $c \in (c_1, c_2)$;
(c) of the form $\{T(\gamma)\}$ for $c \in (c_2, c_3)$;
(d) of the form $\{T\}$ for $c \in (c_3, c_4)$; and
(e) the empty agreement for $c > c_4$.
The critical levels of c satisfy: $(c_2 - c_1)(c_3 - c_2) = 0$.

A first important aspect of Proposition 2 concerns the way contractual incompleteness varies across policy instruments. According to Proposition 2 the effective subsidy S tends to be more

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

discretionary than T; more specifically, for a range of sufficiently high contracting costs it may be optimal to contract over import taxes T while leaving the effective production subsidy S to discretion, but as Proposition 1 indicates and Proposition 2 confirms it is *never* optimal to leave T to discretion and contract only over S.

In this way, Proposition 2 provides a general prediction that trade agreements should always include commitments over import taxes, and should only introduce commitments over other internal measures as the agreement becomes more complete. This prediction resonates to some degree with the broad approach taken by the GATT/WTO, which has been to first establish a base of commitments over import tax levels, and only later to broaden the agreement to explicitly take on various internal measures (most notably, perhaps, 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 internal measures and are therefore less costly to contract over, an assumption that would only reinforce this prediction.

A second interesting aspect of Proposition 2 is the "complementary slackness" condition $(c_2 - c_1)(c_3 - c_2) = 0$: if the discretionary contract $\{T(\gamma)\}$ is optimal for a range of contracting costs (i.e., if $c_2 < c_3$), then the rigid contract $\{T, S\}$ is not optimal for any cost level (i.e., $c_1 = c_2$), and vice-versa. In light of this condition, the next question that needs to be addressed is the following: As contracting costs rise from zero, when is it optimal to first economize on contracting costs by introducing discretion ($c_2 < c_3$), and when by first introducing rigidity ($c_1 < c_2$)?

We focus mostly on the impact of the degree of uncertainty (Δ_γ) and of the import demand level (α). According to the intuition developed above on the determinants of the costs of rigidity and discretion, we 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 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 summarized in Remark 1 stand in marked contrast to those reported in Battigalli and Maggi (2002): in that model, as contracting costs rise, first the optimal contract becomes rigid, and then discretion is introduced. In the present setting, due to the non-separable nature of the contracting problem, it may be optimal to economize on contracting costs by introducing discretion before rigidity, and it may even be the case that rigidity is not optimal for any level of contracting costs.

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

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

We start with changes in the domestic 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 $\{FB\}$. In other words, 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) can be interpreted as reflecting the impact that the monopoly power effect has on the desirability of discretion, combined with the implication of Proposition 1. In particular, the higher is the degree of home-country monopoly power (the higher is α), the less desirable it is to leave policy instruments to discretion, with the order in which instruments are tied down (first T and then also S) dictated by Proposition 1.

More broadly, it is noteworthy to observe at this point that both of the factors which we have identified as increasing the cost of discretion and hence favoring the agreement $\{T, S\}$ over $\{T(\gamma)\}$ and $\{T\}$ – the instrument-substitutability effect and the monopoly power effect – could arguably be seen to apply most aptly to relatively large developed countries. The essence of the instrument-substitutability effect is that a government has access to a rich array of internal measures which it can use to carry on terms-of-trade manipulation should its import taxes be constrained through an international trade agreement. And the essence of the monopoly power effect is that a country is large in world markets, so that it faces foreign export supply that is far from perfectly elastic. Arguably, both of these features are more likely to be true of large developed countries. When viewed in light of our discussion above, this observation is potentially interesting, because it suggests that the attractiveness of contracting over internal measures (such as those embodied in S in our formal model) may be different for large developed countries than for small developing countries, and in particular that optimal contractual design might entail small developing countries negotiating commitments of the form $\{T(\gamma)\}$ or $\{T\}$ while large developed countries negotiate commitments of the form $\{T, S\}$. While our two-country model cannot really address this issue formally, our results are at least suggestive of the possible benefits of a kind of "Special and Differential Treatment" rule for small/developing countries when it comes to contracting over internal measures (such as subsidies).[28]

Proposition 3(i) also 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 internal measures 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 domestic subsidies and other internal measures in the agreement.

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 internal measures, and the value of adding state-contingencies to the agreement is enhanced as a result.[29]

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, the result reported in the above proposition is not particularly surprising: it seems inevitable that increasing uncertainty should reduce the attractiveness of rigid agreements, just as the proposition reports. But there is also a more subtle feature of this result, which is that it concerns uncertainty over a state variable that is directly relevant for the setting of T in the $\{FB\}$ agreement. As we demonstrate in the next 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.

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

[29] We 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 $\{FB\}$.

3.2. Uncertainty about the level of import demand

In the previous subsection we examined a simple stochastic environment where uncertainty concerns only a state variable that affects directly the first-best levels of T and S. Here we explore the implications of a different type of uncertainty, which in some sense is at the opposite extreme: we now suppose that uncertainty concerns only a state variable – namely the domestic demand level α – that has no impact on the first-best levels of T or S. Focusing separately on these two extreme cases is instructive, because it highlights how the rigidity-discretion trade-off and the structure of the optimal agreement depend crucially on the nature of the uncertainty in the contracting environment. In section 3.3 we show how the insights gained from these two extreme cases allows an anticipation of the results that arise under multidimensional uncertainty.

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 one is the possibility of the agreement $\{T(\alpha)\}$, where $T(\alpha)$ is an increasing function. This has the flavor of an escape-clause type of agreement: when α is high, the underlying import volume is high, and so with $T(\alpha)$ an increasing function the agreement $\{T(\alpha)\}$ allows for the import tariff to rise in states of the world in which the underlying import volume is high, broadly analogous to the escape clause provided in GATT Article XIX.[30]

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

The second new insight is that rigidity and discretion are no longer complementary, but are instead *substitutable*. Formally, it is easy to see that the cost of discretion in the presence

[30] We 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, under Article XIX a country is *allowed* to raise its tariff in case of import surge, whereas $\{T(\alpha)\}$ leaves no discretion on T (though our later results on weak bindings suggest that this feature could be captured in a straightforward manner). More importantly, Article XIX includes an "injury" test, which has no counterpart in $\{T(\alpha)\}$ (however, an explanation for the injury test is also lacking in other theoretical interpretations of the escape clause, such as Bagwell and Staiger, 1990).

of rigidity, $\overline{CD} = \Omega_{\{T,S\}} - \Omega_{\{T\}}$, is higher than the cost of discretion absent rigidity, $CD = \Omega_{\{FB\}} - \Omega_{\{T(\alpha)\}}$. This is because the $\{FB\}$ agreement is non-contingent, so $\Omega_{\{FB\}} = \Omega_{\{T,S\}}$, and $\Omega_{\{T(\alpha)\}} > \Omega_{\{T\}}$. The substitutability between rigidity and discretion can also be seen from the perspective of the costs of rigidity, $CRTS$ and CRT. Clearly in this setting $CRTS = 0$, because the first-best agreement is non-contingent, and hence $CRTS < CRT$. Recall that this condition is equivalent to the condition $\overline{CD} > CD$.

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

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

One consequence of the substitutability between rigidity and discretion in the present stochastic environment is that the "concavity" condition 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, the next question we want to address is, Under what conditions (if any) is the agreement $\{T(\alpha)\}$ optimal? The following remark identifies conditions on model parameters under which $\{T(\alpha)\}$ is optimal for a range of c:

Remark 2. *If k is sufficiently small* ***and*** *λ is sufficiently low, then $c_1 < c_2$: an escape-clause-type agreement of the form $\{T(\alpha)\}$ is optimal for some c.*

The result reported in Remark 2 is intuitive in light of our discussion just above. In particular, if the degree of substitutability between T and S is sufficiently low (as when λ is low) so that leaving S to discretion is an attractive option, then an escape-clause type agreement of the form $\{T(\alpha)\}$ will be optimal for a range of c as long as the added complexity of contracting over state variables (k) is sufficiently small.

Finally, we note that the comparative-static effects of changes in the expected demand level $\bar{\alpha}$ can be shown to be similar to those described in Proposition 3 of the previous subsection for the case of changes in α in the presence of γ uncertainty. On the other hand, the effects of changes in the degree of uncertainty over α (Δ_α) differ in an interesting way from the effects of changes in the degree of uncertainty over γ (Δ_γ) as reported in Proposition 4. To illustrate this difference, note that 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 fact that the cost of discretion (CD) is not only rising in α but also convex: hence, as uncertainty over α rises, CD rises, and it may be optimal to move from a contingent agreement with discretion to a rigid agreement without discretion. This confirms our earlier observation that the effects of uncertainty depend on whether it concerns state variables that are directly relevant to the first best policy levels (as in Δ_γ) or rather state variables that are not directly relevant to the first best policy levels (as in Δ_α).

3.3. Multidimensional uncertainty

In sections 3.1 and 3.2, for pedagogical reasons we focused on two extreme stochastic environments: one where uncertainty concerns only the externality (γ), and one where uncertainty concerns only the level of import demand (α). Here we consider how the main results derived in those sections are modified when both γ and α are uncertain.

Let us examine what types of agreement can potentially be optimal in this setting. First observe 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. The 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 $\{T(\alpha, \gamma)\}$.

So the main difference in results with respect 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 worth emphasizing the possibility that an agreement $\{T(\alpha, \gamma)\}$ may be optimal: this is a case in which the presence of contracting costs may lead the agreement to be *contingent on more variables* than the first best. The reason why it may be desirable to have such an over-contingent agreement is that, if it is worthwhile leaving discretion over domestic instruments to save on contracting costs, it may also be worthwhile making the import tax T contingent on the domestic demand level α – in addition to the externality γ – in order to mitigate the distortions associated with discretion.

Let us consider briefly the conditions under which an agreement of the type $\{T(\alpha, \gamma)\}$ is optimal for a range of contracting costs c. Note that this agreement costs $2(1 + k)c$. Thus, if $k \geq 2$, so that the cost of adding a state-contingency is large relative to the cost of adding a policy, the agreement $\{T(\alpha, \gamma)\}$ costs at least as much as the $\{FB\}$ agreement, and hence can never be optimal. Thus $k < 2$ is a necessary condition for $\{T(\alpha, \gamma)\}$ to be optimal for some c.

Also, it is not hard to establish sufficient conditions under which this agreement is optimal for some c: for example, this is the case if (i) k is sufficiently small, (ii) the cost of discretion is not too high (for example because S is not a good substitute for T), and (iii) α and γ are not perfectly correlated.[31]

It is worth emphasizing also the possibility that, even though the first best outcome calls for T to be 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.[32]

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, it can be shown that a higher expected level of domestic demand leads to a lower degree of discretion in the optimal contract, similarly as in the one-dimensional uncertainty case. Also our results about the complementarity or substitutability between rigidity and discretion generalize in an intuitive way. In the previous sections we showed that, if only γ is uncertain, rigidity and discretion are complementary, while they are substitutable if only α is uncertain. Intuitively, then, rigidity and discretion tend to be complementary if uncertainty in γ is important relative to uncertainty in α, and vice-versa.

Overall, we may state that the essential insights of our one-dimensional uncertainty analysis are preserved in an environment where both γ and α are uncertain. The same insights are likely to be preserved also with more general multidimensional uncertainty, except of course that the optimal agreement may be contingent on state variables other than α or γ.

4. The Role of the National Treatment Clause

A distinguishing feature of the GATT/WTO (albeit increasingly less of the WTO) is that it combines rigid bindings on tariffs with a significant amount of discretion over internal measures. This discretion is not complete, however. An important legal instruments that can be invoked by Members in order to address internal measures is the National Treatment (NT) provision in GATT Article III. This instrument can naturally be seen as an attempt to remedy problems that arise under the contractual incompleteness of the GATT/WTO. In this section we evaluate the usefulness of the NT clause as a means to economize on contracting costs.

The ideal theoretical approach to evaluate the NT provision would be to extend the set of feasible agreements $\mathcal{A}_0$ to allow for a more general class of agreements, and then to examine whether and when an NT clause of the type observed in the GATT/WTO is part of an optimal agreement within this wider class. Such an evaluation would be extremely complex however, partly because there would be a very large number of agreements to consider. We therefore

[31] Consider 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 clearly, if the cost of discretion is not too high, $\{T(\alpha, \gamma)\}$ will be optimal for some c.

[32] In 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.

attempt something more modest, relying on institutional motivation to restrict our attention to studying the potential desirability of just this particular clause: 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.

For our purposes, the relevant part of the NT clause can be found in GATT Article III.2, which addresses internal taxation. The exact interpretation of this provision is subject to debate in the legal and economic literature, but we take the view that the core of the NT rule can be captured in the context of our model by the simple constraint $t_h = t_f$.[33] It is important to note that, while the NT provision restricts internal taxes to be the same, it *allows the importing country discretion* over the level at which these are set. In line with our assumptions on contracting costs, we assume that including the NT clause in the agreement costs $2c_p$.[34]

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

The points we will 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. We focus on the case of a consumption externality, as in our basic model, 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 the same as in the previous section, as recorded in 2.1. Within this class, as we argued previously, we can focus on agreements that tie down S and/or T. However, for NT-based agreements, the arbitrage conditions become

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

[33] There are two interpretation issues that can be raised here. First, Article III.2 speaks of "treatment no less favorable," which suggests that a more accurate formalization of the NT provision is given by the inequality constraint $t_h \geq t_f$. However, in our model it can be shown that this constraint would always be binding, so that there would be no gain in allowing for this inequality constraint. Also, Article III.1, which sets the stage for the NT clause, restricts attention to measures that are applied "...so as to afford protection...". This poorly drafted sentence can be read in various ways, including that the provision should only have a bite against "protectionism," and not in cases where higher taxation of a foreign product is motivated by the pursuit of more legitimate objectives. This is not an issue in the context of the current model, since there is no ground for treating the imported product less favorably than the locally produced good in any of the setups we consider. However, if the model were slightly modified, so as to let the imported product potentially be associated with a more severe regulatory problem than the domestically produced good, this would become an important issue; see Horn (2006). See also Horn and Mavroidis (2004) for legal and economic analyses of Article III text and case law as it pertains to taxation.

[34] It 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.

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

Notice that the NT clause, if applied in the absence of further constraints, has no real effect because it does not prevent the importing country from manipulating both price wedges at will. As a consequence, both the efficient outcome and the Nash equilibrium derived in the previous section in the absence of NT, can also be implemented with policies that conform to the NT clause. In particular, the efficient policies under NT are given by

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

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

$$\tau^{NE} = \frac{p^*}{\eta^*}, \ t_h^{NE} = t_f^{NE} \equiv t^{NE} = \gamma, \text{ and } s^{NE} = 0.$$

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

There are potentially many kinds of NT-based agreements, but we can reduce the number that must be considered by focusing only on NT-based agreements that can be *strictly* optimal in the class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$.

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 is simple. An agreement of the type $\{NT, \tau, s\}$ leaves discretion over the common consumption tax t. At a more fundamental level, as can be seen from the expressions (4.1), this agreement leaves discretion over the consumer price wedge $p - p^*$, while tying down the producer price wedge $q - p^*$. As we remarked earlier (see footnote 23), this kind of discretion can*not* be generated by a non-NT agreement. This is a simple but 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* type of NT-based agreement that can strictly improve over non-NT agreements is $\{NT, \tau, s\}$.[35]

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

[35] Note 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.

We start with an intuitive discussion of the pros and cons of the NT-based agreement $\{NT, \tau, s\}$ relative to non-NT agreements. How attractive is the NT-based agreement as a way to save on contracting costs? Relative to $\{FB\}$, this agreement implies lower contracting costs (as long as there is some uncertainty, i.e. $\Delta_\gamma > 0$), but on the other hand, it 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 have seen in previous sections 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 role of 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. However, importantly, the underlying parameter conditions that cause t and τ to be highly substitutable are essentially *opposite* to the parameter conditions that cause S and T to be highly substitutable (low β and/or high λ). To see this, consider the equilibrium terms of trade, which is easily derived to be

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

Inspection of this expression confirms that t is a close substitute for τ when β is high and/or λ is low. When λ is close to zero, t and τ are close to perfect substitutes, and hence an agreement of the type $\{NT, \tau, s\}$ offers essentially no improvement over the empty agreement. On the other hand, if β is close to zero, t is nearly useless as a surrogate means to distort terms of trade, and hence the cost of discretion over t approaches zero.

The degree of home-country monopoly power can be shown to affect the cost of discretion over t in a very 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 the agreement $\{NT, \tau, s\}$ less attractive.

Having discussed at a broad intuitive level what are the pros and cons of NT-based agreements, 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 frst 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 implied by 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 (so that the consumption tax t is a poor substitute for the tariff τ), then there is an intermediate range of c for which the optimal agreement includes the NT clause.*

The above proposition 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.[36] 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 discretion over internal taxes afforded by the NT clause.

5. The Role of Weak Bindings

In the previous sections we focused on agreements that impose equality constraints, as in $\{T = 2\}$ or $\{NT, \tau = 0, s = 0\}$. This restriction, of course, would be without loss of generality in a world of costless contracting, for in this case it would be optimal to implement the first best outcome, and this can be achieved with an agreement that imposes equality constraints (as in $\{T = \gamma, S = -\gamma\}$), so 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 will argue in this section, in a second-best environment it may be preferable to impose policy *ceilings* rather than equality constraints. Below we will make this claim more formal, but as a first step we develop some intuition through a simple example.

As a first observation, we note that a "weak" binding cannot offer a strict improvement over a "strong" binding that is state-contingent (e.g., $\{T \leq \gamma\}$ cannot improve over $\{T = \gamma\}$). The reason is simply that a state-contingent strong binding can be structured so as to position the policy variable exactly where it is optimal to place it for every possible realization of the state variable, and so the added ex-post flexibility that a weak binding offers cannot be of value.

Suppose that only γ is uncertain, and that the optimal agreement in class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$ is of the form $\{T = \bar{T}\}$. Now consider replacing this strong binding with the weak binding $\{T \leq \bar{T}\}$. This agreement clearly performs at least as well as the previous one, because the purpose of the agreement is to prevent governments from *raising* import taxes above their efficient level. Can a weak binding be *strictly* preferable to a strong binding? This depends on the exact configuration of parameters. Let $T^{NE}(\gamma)$ denote the Nash equilibrium level of T as a function of γ. Intuitively, and noting that T^{NE} is increasing in γ, the optimal level of the strong binding $\bar{T}$ must be below $T^{NE}(\bar{\gamma}+\Delta_\gamma)$, but it may not be below $T^{NE}(\bar{\gamma}-\Delta_\gamma)$. If $\bar{T}$ is above $T^{NE}(\bar{\gamma}-\Delta_\gamma)$, then a weak binding is strictly preferable, because in the low state the government will choose a

[36] The 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 we note that, if we modified slightly our contracting cost assumptions by assuming that the NT clause costs less than $2c_p$ – see footnote 34 – 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 is contingent, and also to one where only s is contingent.

level of 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 are able to show that weak bindings are preferable to strong bindings 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 is similar as the one we provided above for the import tax: governments are tempted to distort production subsidies in import-competing industries *upwards*, and hence the relevant constraint that needs to be imposed on governments is an upper bound on the subsidy.

The proposition is valid regardless of which state (γ or α) is uncertain, therefore we do not need to be specific about the source of uncertainty.[37] We will refer to a constraint of the kind $x \leq \bar{x}$ (where x is some policy instrument and $\bar{x}$ a noncontingent level) as a "rigid weak binding" and to a constraint of the kind $x = \bar{x}$ as a "rigid strong binding."

Proposition 7. *Suppose the optimal agreement in the class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$ includes some rigid strong binding. Then this agreement can be weakly improved upon by replacing the rigid strong bindings with rigid weak bindings, and the improvement is strict 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, our result highlights another sense in which rigidity and discretion may be complementary ways to economize on contracting costs: in Section 3.1 we argued that, if uncertainty concerns state variables that are directly relevant to the first-best policy levels, the cost of discretion over domestic instruments is lower in the presence of rigidity than in its absence; here we have shown that, conditional on the agreement being 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.[38]

[37] As 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).

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

6. Production Externalities and Political-Economy Motives

In our basic model, the efficiency rationale for policy intervention is given by a consumption externality. In this section we consider two alternative possibilities: first, that the efficiency rationale for policy intervention is given by a production externality, and second, that the efficiency rationale for policy intervention is of a political nature.

To keep the discussion simple we focus on the class of agreements $\mathcal{A}_0$, that is, we abstract from the possibility of NT-based agreements and weak bindings, but the arguments we present here are easily extended to allow for these possibilities.

Consider first the case of a production externality. Let us set $\gamma \equiv 0$ and suppose that there is a positive production externality in the home country equal to σX with $\sigma > 0$, and that as before this externality enters directly and separably into the representative home-country citizen's utility and does not cross borders. We allow the possibility that both the externality σ and the domestic demand level 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 with respect to the case of an (uncertain) consumption externality, where the first-best levels of *both* S and T were state-contingent. This difference has some subtle implications for the nature of the optimal agreement, as we discuss below.

To a large extent, the qualitative results for this case are similar to those for the case of a consumption externality. In particular, the candidates for an optimal agreement are the same as in section 3.3, except of course that $T(\gamma)$ and $T(\gamma, \alpha)$ are replaced respectively by $T(\sigma)$ and $T(\sigma, \alpha)$ (and the $\{FB\}$ agreement is different, as we already highlighted). But there is an interesting difference between the two cases: the potential appeal of making T contingent on σ (as in the agreements $T(\sigma)$ and $T(\sigma, \alpha)$) is distinct from the potential appeal of making T contingent on γ in the case of a consumption externality, 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 a means of managing the distortions associated with leaving S to discretion.[39] Notice that the three agreements $T(\alpha)$, $T(\sigma)$ and $T(\sigma, \alpha)$ are potentially attractive for a similar reason; intutively, then, the performance of these agreements relative to each other will depend on the exact nature of the uncertainty, including which of the two sources of uncertainty is more important.

Another point of difference between the production and consumption externality cases concerns the complementarity/substitutability between rigidity and discretion. As we argued in the previous sections, uncertainty about state variables that affect the first-best level of S (e.g. γ) tends to generate complementarity, whereas uncertainty about state variables that do not affect the first-best level of T (e.g. α) tends to generate substitutability. As we highlighted

[39] To 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 Nash 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.

above, σ affects the first-best level of S but not that of T, and this implies that both forces are at work, thus in general uncertainty about σ has ambiguous implications for the complementarity/substitutability between rigidity and discretion. Since uncertainty about γ unambiguously pushes toward complementarity, broadly speaking we can say that in the case of production externalities rigidity and discretion are less likely to be complementary, as compared with the case of consumption externality.

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 producer surplus. More specifically, we consider the possibility that each government maximizes an objective function of the form $\tilde{W} \equiv W + \zeta \cdot \Pi$. We allow for the possibility that both the import demand level α and the political-economy parameter ζ may be uncertain.

Clearly, 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: while in the case of production externality the first-best level of S is equal to σ and does not depend on the demand and supply conditions, 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 will take the form $\{T = 0; S = S(\xi, \alpha)\}$, 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 one more cost class to consider relative to the cases examined previously, namely the agreements that cost $(4+k)c$. These are the agreements where both policies are contingent on a single state variable (α or γ). Aside from this, however, the main qualitative results are the same as in the case of production externality.

Before concluding, we want to 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 efficient way to take care of the governments' distributional concerns is to practice free trade ($T = 0$) and offer domestic subsidies contingent on political shocks (ζ) and on import demand shocks (α). 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 surging import demand. We have also established that the nature of the optimal agreement in an incomplete contracting setting depends on features of the underlying economic environment that can be given simple economic interpretations: the degree of substitutability across instruments; the extent of monopoly power on world markets; the extent to which discretion facilitates state-contingency; and the source (and not just the level) 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.

Of course, for both analytical feasibility and pedagogical reasons, we have adopted many strong simplifying assumptions, and so our model abstracts from many 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.

Our analysis has been performed within a two-country setting. This precludes the study of one of the foundational provisions of the GATT/WTO, its MFN rule. Extending our analysis to a multi-country environment would also permit a formal exploration of a point suggested by our two-country analysis, namely, that a "special and differential treatment" rule might be warranted for small/developing countries when it comes to contracting over internal measures (such as subsidies).

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 the GATT/WTO, and so this is a natural starting point. But there are certain provisions of the GATT/WTO (most notably the Non-Violation provision of GATT Article XXIII) that do have this flavor, and such provisions warrant further investigation within an incomplete-contracts setting, since they seem highly 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.

Finally, 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. We leave these and other important tasks for future work.

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 $m' < \hat{m}$,

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

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

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

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). Differently put, if (3.1) were violated, there would not be a value of c such that $\hat{A}$ simultaneously dominates A' and A'', contradicting the optimality of $\hat{A}$.

To see the sufficiency part, suppose that $\hat{A}$ is optimal in its complexity class and condition (3.1) 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

To establish this, we argue that an agreement that constrains only S cannot achieve higher gross global welfare than at the Nash equilibrium, for any state of the world. Recalling that $T^R(S)$ is the best-response level of T given S, the maximal gross global welfare that can be achieved by this type of agreement is given by the value of $W^G(T^R(S), S)$ evaluated at the optimal level of S,[40] with the optimal level of S defined in turn by the associated first-order condition $dW^G/dS = 0 \Longrightarrow T^R_S(S) = -W^G_S/W^G_T$, where subscripts denote derivatives. This first-order condition requires that the slope of $T^R(S)$ be equated with the slope of an iso-W^G curve in (T, S) space. It is direct to verify that the slope of $T^R(S)$ is

$$T^R_S(S) = -\frac{\lambda}{\beta + \lambda}.$$

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

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

is

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

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

Proof of Proposition 7

An agreement that is optimal in the class under consideration (i.e., in the class $\mathcal{A}_0 \cup \mathcal{A}_{NT}$) and includes rigid strong bindings (RSBs) must be one of the following: (a) the best $\{T = \bar{T}\}$ agreement; (b) the best $\{T = \bar{T}; S = \bar{S}\}$ agreement; and (c) the best $\{NT; \tau = \bar{\tau}; s = \bar{s}\}$ agreement. As a first step we show that, for each of these agreements, if we replace RSBs with rigid weak bindings (RWBs), efficiency is weakly increased.

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.

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

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

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

Next consider case (b), and consider replacing $\{T = \bar{T}; S = \bar{S}\}$ with $\{T \leq \bar{T}; S \leq \bar{S}\}$. For a given state, there are four relevant possibilities for how the importing country sets (T, S) under an agreement $\{T \leq \bar{T}; S \leq \bar{S}\}$:
(i) the importing country chooses $(T = \bar{T}, S = \bar{S})$: In this case there is of course no change in Ω relative to $\{\bar{T}; \bar{S}\}$.
(ii) the importing country chooses $(T = \bar{T}, S = S^R(T))$: Here it must be that $S > S^R(T)$. Let us evaluate $\Omega_S = W_S + W_S^*$.Clearly, $W_S < 0$ for $S > S^R(T)$, and $W_S^* < 0$, hence $\Omega_S < 0$ in this region, which in turn implies that switching to weak bindings increases Ω.
(iii) the importing country chooses $(T = T^R(S), S = \bar{S})$: Here it holds that $T > T^R(S)$. Let us evaluate $\Omega_T = W_T + W_T^*$. Since $W_T < 0$ for $T > T^R(S)$, and $W_T^* < 0$, it follows that $\Omega_T < 0$ in this region, which ensures that switching to weak bindings increases Ω.
(iv) the importing country chooses $(T = T^{NE}, S = S^{NE})$: The same result can be shown by combining the arguments we just made for cases (ii) and (iii).

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

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

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

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

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

and the reaction function for z as

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

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

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

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

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

This implies that switching to weak bindings cannot decrease Ω.

We have shown that replacing RSBs with RWBs cannot decrease the expected surplus. It remains to show that there is some configuration of parameters for which replacing RSBs with RWBs increases Ω strictly. Applying the arguments developed above, we know that a sufficient condition for this to happen is that the noncooperative level of a policy is below the ceiling for that policy. It is easy to show that there exists a configuration of parameters for which this is the case. ■

9. References

Bagwell, Kyle and Robert W. Staiger, “A Theory of Managed Trade,” *The American Economic Review*, vol. 80, pp. 779-795, 1990.

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.

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

Battigalli, Pierpaolo and Giovanni Maggi, “Rigidity, discretion and the costs of writing contracts,” *The American Economic Review*, vol. 92(4), pp. 798-817, 2002.

Battigalli, Pierpaolo and Giovanni Maggi, "International Agreements on Product Standards: an Incomplete-Contracting Theory," NBER Working Paper No. 9533, 2003.

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

Costinot, Arnaud, "A Comparative Institutional Analysis of Agreements on Product Standards," mimeo, 2004.

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

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.

Hudec, Robert, *The GATT Legal System and World Trade Diplomacy* (2nd Ed.), Salem NH: Butterworth, 1990.

Maggi, Giovanni and Andres Rodriguez-Clare, “A Political-Economy Theory of Trade Agreements,” NBER Working Paper No. 11716, 2005.

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.

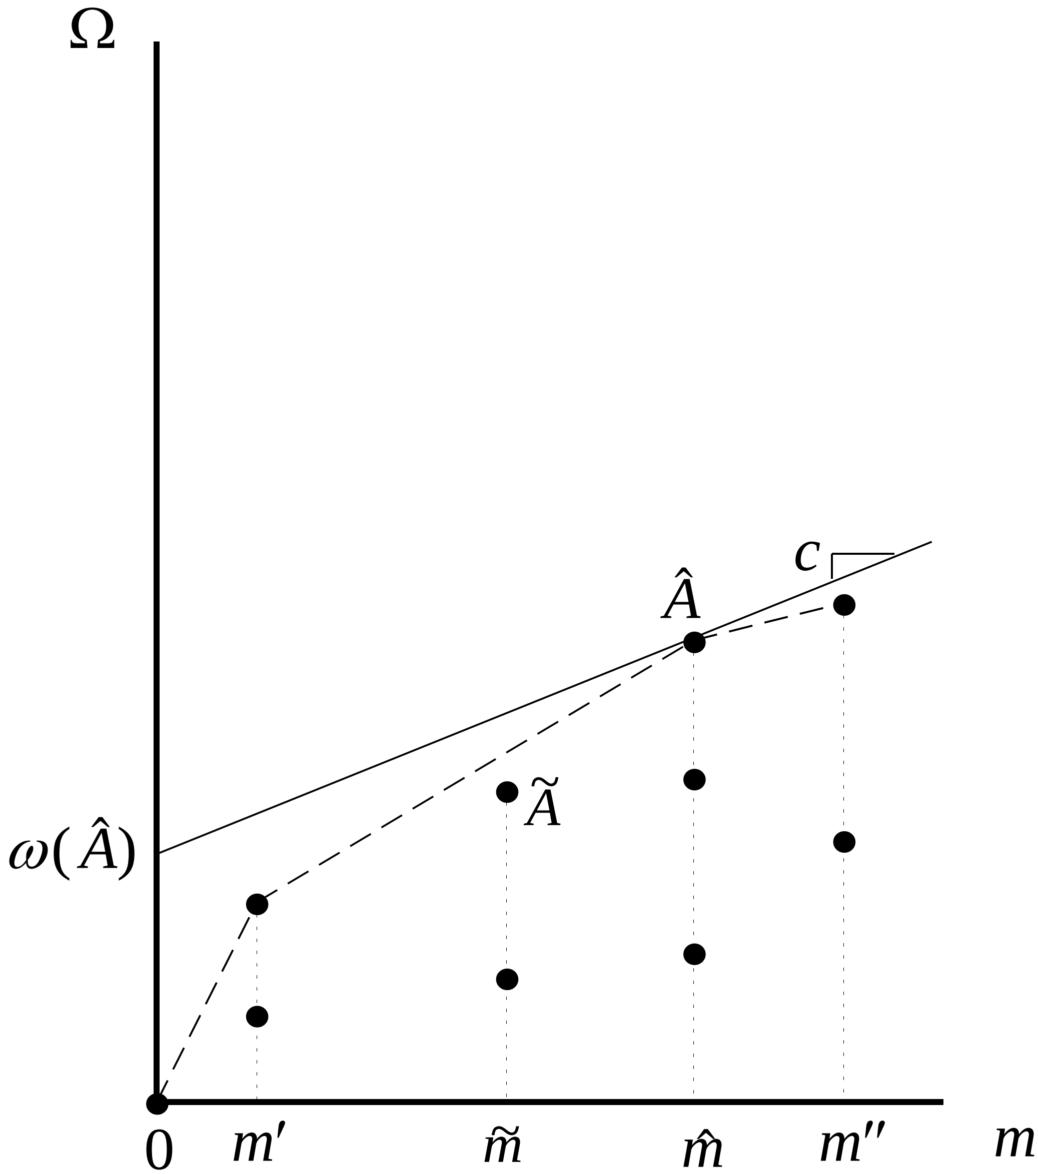

Figure 1

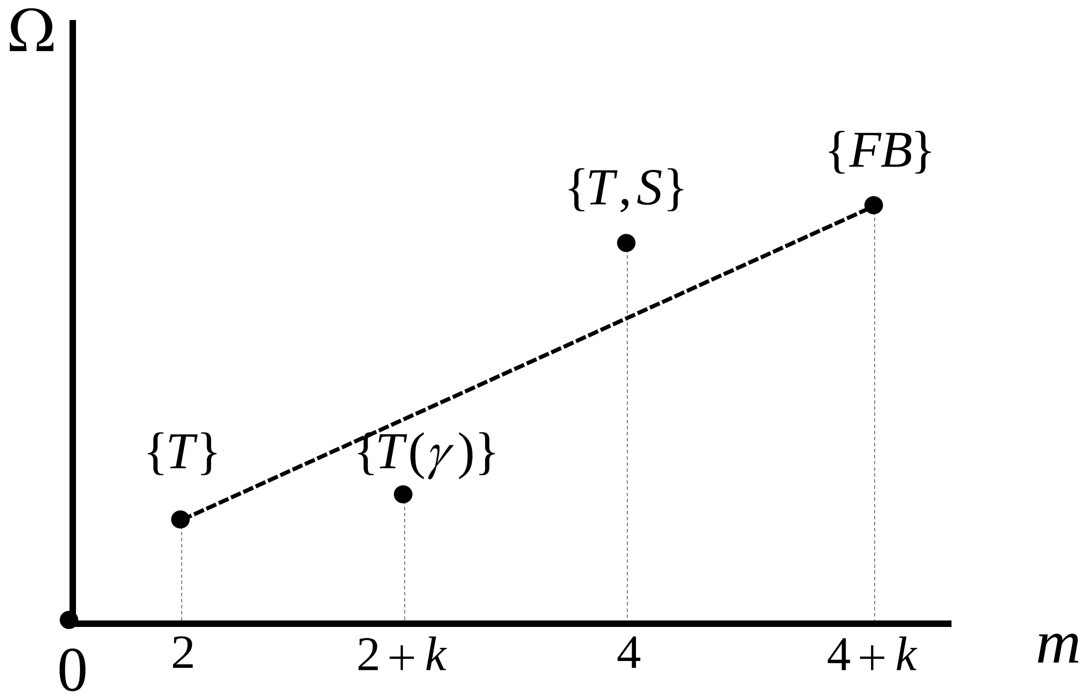

Figure 2a

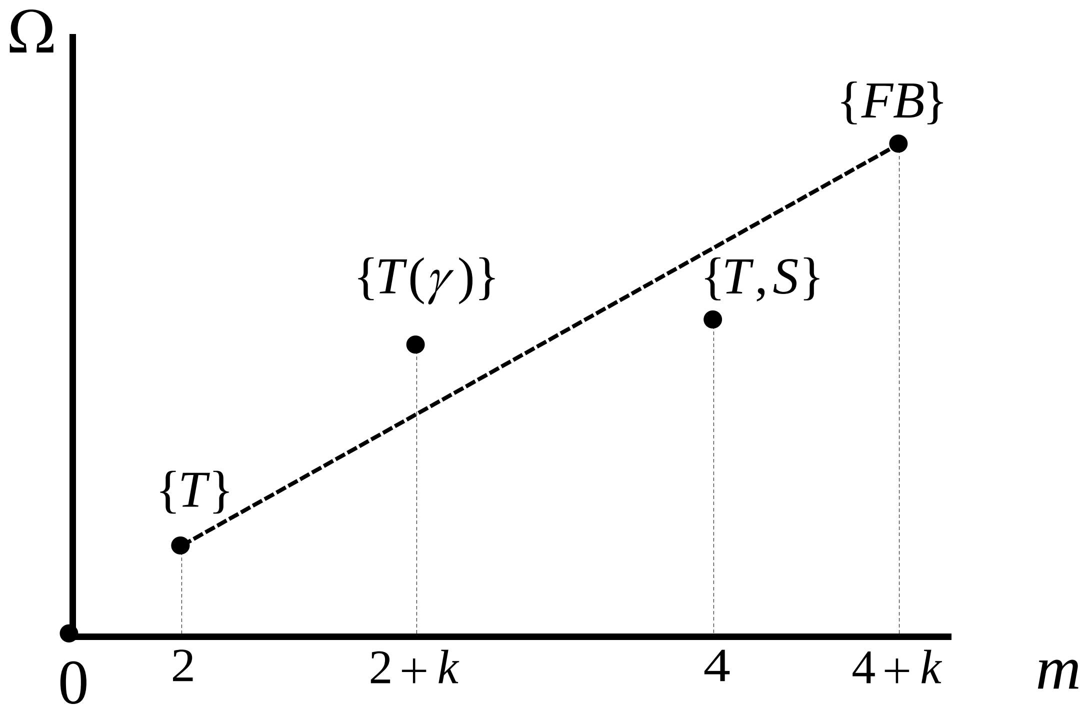

Figure 2b