

International agreements on product standards: an incomplete-contracting theory*

Pierpaolo Battigalli, Bocconi University
Giovanni Maggi, Princeton University and NBER

August 2002

Abstract

In this paper we take the view that trade agreements are highly incomplete contracts. Contractual incompleteness is particularly severe for agreements that regulate complex policies, such as product standards. Focusing on the case of product standards, we examine two potential remedies to the problem of contractual incompleteness. First, we consider a dispute settlement procedure (DSP) that provides arbitration in states of the world that are not covered by the ex ante agreement. We identify conditions under which a DSP can provide ex-ante efficiency gains, and examine how these gains depend on the fundamentals of the problem. A second potential remedy to contractual incompleteness is given by *rigid rules*, i.e. rules that are not product-specific. We show that the national treatment rule is the only rule of this kind that increases ex ante efficiency for any probability distribution over potential products. Finally we show that, under some plausible conditions, the optimal ex-ante agreement is structured in three parts: (i) a set of clauses that specify standards for existing products; (ii) a rigid national treatment rule, and (iii) a dispute settlement procedure. Although the model focuses on the case of product standards, the analysis suggests a more general theory of international trade institutions.

*We thank Kyle Bagwell, Gene Grossman, John McLaren and Robert Staiger for very helpful comments and discussions. We also thank the participants in seminars at Duke University, the University of Maryland and the University of Wisconsin-Madison for their feedback. Giovanni Maggi acknowledges financial support from the National Science Foundation.

1. Introduction

Trade agreements are highly incomplete contracts. This is particularly true for agreements that attempt to regulate complex policies, such as product standards. One manifestation of this incompleteness is that the current agreements say very little about standards for products that may arise in the future, whereas a complete agreement would specify standards for all potential future products. For example Robert Hudec, a prominent scholar of the GATT-WTO, describes the difficulties faced by the original GATT drafters in the following terms (1990, p.24): “...*The standard trade policy rules could deal with the common types of trade policy measure governments usually employ to control trade. But trade can also be affected by other ‘domestic’ measures, such as product safety standards, having nothing to do with trade policy. It would have been next to impossible to catalogue all such possibilities in advance. ...*”

In this paper we examine two potential remedies to the incompleteness of agreements: first, a *dispute settlement mechanism* that provides arbitration in states of the world that are not covered by the ex ante agreement, and second, the use of *rigid rules*, i.e. rules that are not product specific.

The idea that an ex post dispute settlement mechanism can help remedy the incompleteness of agreements has been expressed informally by several scholars. It has also been pointed out that such mechanisms are highly imperfect because the state of the world may not be verifiable ex post, even if it is observed by the contracting parties. For example, Dixit (1996) writes: “...*All feasible contracts are necessarily incomplete. Therefore, ex post institutions (dispute settlement mechanisms) are very important. But the very complexity of nature makes these mechanisms less than perfect – the state of the world may not be observable ex post or even if observed by the parties to the contract, may not be verifiable to an outsider whose job is to enforce it.*”¹ A primary objective of this paper is to examine this idea more rigorously. Whether efficiency can be enhanced by a dispute settlement procedure (DSP) is not obvious, even if the state of the world is verifiable: after all, if an unforeseen contingency arises, governments can negotiate *ex post*, thus filling the gaps of the original agreement without the help of third parties. We analyze the conditions under which a DSP can provide a Pareto-improvement for

¹For statements along similar lines, see for example Kovenock and Thursby (1993) and Ethier (2001).

the participating countries, and examine how the gains from such a system depend on the fundamentals of the problem.

We address these questions in the context of a simple two-country model where each government chooses product standards to maximize national welfare. The resulting noncooperative equilibrium is inefficient, and thus an international agreement is called for. Agreements are enforceable and can cover existing products, but not future potential products. When a new product arises, governments can bargain efficiently over the standards for the new product, thus the outcome is ex-post efficient. However, due to contractual incompleteness, the outcome is inefficient from an ex-ante perspective. This inefficiency can be understood as arising from *missing issue linkage*. Suppose governments could bargain ex ante over future potential products, and think about each potential product as a distinct “issue.” Then the ex ante bargain would link these issues together. This linkage is missing if a complete contingent contract is not available.

We then ask whether a DSP can remedy the inefficiency caused by the incompleteness of the agreement. We think of the DSP as an arbitrator. Each government can invoke the DSP during the negotiation process, in which case the DSP’s decision is binding on both parties. Thus the DSP influences the outcome of the ex post bargaining game by affecting the players’ outside options. In the ex-ante agreement governments specify the principle that guides the DSP in its decisions. In our setting, it is natural to assume that this principle is the maximization of the two countries’ joint welfare, which is what governments would do if they could write a complete contingent contract; we will often refer to this as the “umbrella” principle. The application of this principle however is imperfect, because the DSP has limited information about the state of the world, even ex post.

In spite of the fact that governments bargain efficiently and have more information than the DSP, we show that the presence of the DSP can improve ex-ante efficiency, provided the imperfection in its information is not too serious. The key point is that the presence of the DSP can indirectly recover some of the gains from issue linkages that are missing from the incomplete agreement. This can be achieved by shifting the ex-post bargaining outcome in favor of the “weaker” country, which in turn can be done by offering a favorable outside option

to this country. In sum, the DSP can increase ex-ante utilities by smoothing out ex-post power imbalances. The DSP can then be interpreted as pursuing a “rules-based” approach, as opposed to a “power-based” system, in the sense that it makes decisions based on an anonymous rule (maximization of the countries’ joint welfare) that is defined without reference to the government’s ex-post bargaining positions.²

The above argument is subject to an important qualification: there is a potential role for an ex-post arbitration system only if the governments’ utilities are nontransferable. If pure international transfers are available, the incompleteness of the ex ante agreement does not cause any inefficiency. Also, the efficiency gains from an arbitration system are smaller when the ex post bargaining is more symmetric. If the ex post bargaining is fully symmetric, a DSP can offer no value added.

We emphasize that the DSP in our model plays a very different role than in most existing models of the DSP (e.g. Hungerford, 1991, Kovenoch and Thursby, 1993, Maggi, 1999). In these models, the DSP is desirable only if it can increase the severity of punishments for violations of the agreements. In our model the DSP is not concerned with punishments for past actions, but tries to implement more efficient policies for the future. Also, in the previous models any “out of court” settlement by governments is undesirable, as this undermines the punishment mechanism. In our model, out-of-court settlement is efficient. Thus, our theory rationalizes a trade institution that facilitates genuine dispute *settlement*, rather than retaliation.

A second potential remedy to contractual incompleteness is given by *rigid rules*, that is, rules that are not product-specific. A prominent example of rigid rule is the nondiscrimination (ND) rule. In our two-country model, this takes the form of a “national-treatment” rule, which requires a country to apply the same standard to imported and domestically produced units

²This is reminiscent of John Jackson’s characterization of the WTO as a rules-based system as opposed to a power-based system. Jackson (1990, pp 85-86) writes: “One can roughly divide the various techniques for the peaceful settlement of international disputes into two types: settlement by negotiation and agreement with reference (explicitly or implicitly) to relative power status of the parties; or settlement by negotiation or decision with reference to norms or rules to which both parties have previously agreed. ... In both techniques negotiation and private settlement of disputes is the dominant mechanism for resolving differences; but the key is the perception of the participants as to what are the bargaining chips.” Jackson goes on to explain that under a “rules-oriented” system, an arbitration body takes over for the parties if these reach an impasse in their negotiations, while under a “power-oriented” system the parties are left basically to rely on their respective “power” positions. Our model provides an efficiency rationale for this type of rules-oriented approach.

of a given good. As a first step, we consider the effects of the ND rule abstracting from the possibility of a DSP. Imposing the ND rule is not sufficient to implement the first best outcome, however we show that this rule has the desirable property of increasing ex-ante efficiency for any probability distribution over potential products. Furthermore, we find that the ND rule is the *only* feasible rigid rule that satisfies this property.

Based on the above considerations, we single out the ND rule as the only candidate rigid rule, and examine under what conditions it is desirable to combine it with a dispute settlement procedure. We show that, under some plausible conditions about the DSP's informational capabilities, a combination of DSP and ND rule is indeed the best option. In sum, the optimal ex-ante agreement is structured in three parts: (i) a set of clauses that specify standards for existing products; (ii) a rigid national treatment rule, and (iii) a clause that institutes a dispute settlement procedure.

One possible interpretation of our theory is normative, since we focus on the agreement that maximizes ex-ante efficiency, but we believe that the theory also has some positive explanatory power. Next we briefly discuss whether the predictions of the model are consistent with the WTO's approach to the regulation of product standards.

Within the WTO, the main agreements on product standards are the Agreement on Technical Barriers to Trade (TBT) and the Agreement on Sanitary and Phytosanitary Measures (SPM). First, the SPM agreement endorses a few sectoral agreements that set standards for a variety of existing products, such as the FAO-WHO Codex Alimentarius, the International Office of Epizootics and the International Plant Convention. Second, the TBT and SPM agreements include a blanket nondiscrimination rule (national treatment and MFN). This is arguably the only rigid rule in these agreements. Third, these agreements contain an umbrella principle stating that standards should not create unnecessary obstacles to trade. For example, Article 2.2 of the TBT agreement states: “...*Members shall ensure that technical regulations are not prepared, adopted or applied with a view to or with the effect of creating unnecessary obstacles to international trade.*” This “no-disguised-protectionism” principle is further articulated in articles 2 and 5 of the SPM agreement. Several scholars have interpreted this principle as saying that standards should maximize global real income, or alternatively, that standards should be

those that a welfare-maximizing government would choose if all producers were domestic (see Baldwin, 1970, Engel, 1999, and Fischer and Serra, 2001). If one accepts this interpretation, this is analogous to the umbrella principle that guides the DSP in our model.³ Overall, we think that these observations are broadly consistent with our model. In the concluding section we will expand our discussion of the relationship between the model's predictions and the WTO's approach to product standards.⁴

Although in this paper we focus on the issue of product standards, the model suggests a more general incomplete-contracting theory of trade agreements. Our analysis of the DSP is fairly general, and should apply to a wide set of trade-relevant policies. The other general insight is that an incomplete-contracting approach can explain the presence of rigid rules in trade agreements. The basic idea is that rigid rules can enhance ex-ante efficiency by affecting the status quo in ex-post negotiations. Of course, characterizing the optimal rigid rules is inevitably a policy-specific exercise. But we think the theoretical framework presented here suggests a useful method for studying agreements in other policy areas.

Next we discuss the related literature. Ethier (2001) considers a model that is similar to ours in that trade agreements are incomplete and a dispute settlement body intervenes when an unforeseen contingency arises. However the focus of the paper and the structure of the model are very different from ours. For example, in Ethier's model the DSP is equivalent to a mechanical rule which specifies that, if a country takes a trade-related action that reduces

³A subtle issue arises when comparing the predictions of the model with the reality of WTO disputes. In reality the DSP rules only on whether the agreement has been violated, whereas in (the literal interpretation of) our model the DSP provides arbitration for contingencies that are not covered by the agreement. However we think that, if appropriately interpreted, our model is fully consistent with this aspect of reality. In the model, the "no-disguised-protectionism" principle can be construed as a "rule", and the DSP can be viewed as determining whether the (noncooperative) standards in place at the time of the dispute are consistent with this rule. The important feature of the "no-disguised-protectionism" rule is that welfare functions are not verifiable, hence the rule is not directly enforceable, and its application requires subjective judgement by the DSP.

⁴Other trade agreements where product standards play an important role are NAFTA and the EU. The approach followed by NAFTA in the area of product standards is very similar to the one followed by WTO (see for example Sykes, 1995, pp.108-109). The approach followed by the EU is somewhat different. For example, the EU adopts a principle of mutual recognition rather than one of national treatment (mutual recognition is essentially a transfer of authority over import standards to the exporting country), and actively pursues a full harmonization of standards. These differences can probably be explained by the fact that the EU has reached a much deeper level of economic integration than the WTO or NAFTA. For this reason, a theoretical understanding of the EU approach would require a different model than the one we present here. Among other things, such a model would have to allow for more sophisticated authority structures, including the possibility of delegating the choice of harmonized standards to supra-governmental bodies.

the effective degree of trade liberalization, the injured country can retaliate with a proportional tariff increase after a fixed amount of time. In contrast, the DSP in our model is a genuine arbitrator that, if invoked, gathers information and maximizes an objective function. Another difference is that we examine the role of rigid rules and their interaction with the DSP, which Ethier does not.

Bagwell and Staiger (2001a) examine a model where governments choose standards and import tariffs. If chosen noncooperatively, both instruments are set at inefficient levels, because they exert international externalities through terms of trade. Bagwell and Staiger show that globally efficient outcomes can be achieved through a tariff agreement alone, provided this agreement is complemented by a market-access rule of the following kind: after the tariff agreement is signed, each country can unilaterally change its standards and tariffs subject to the only constraint that this does not undermine the market access levels implied by the tariff agreement. The main difference between our model and Bagwell and Staiger's is that we focus on the potential remedies to contractual incompleteness. We note that market-access rules *à la* Bagwell and Staiger cannot solve the problems related to contractual incompleteness. This is because the efficient level of market access is product specific, thus it cannot be specified *ex ante* for future potential products. Another important difference is that in our model the international externality associated with product standards does not travel simply through terms of trade, thus it is not clear that a market-access rule can achieve efficient outcomes even for existing products. We will come back to this issue in section 4, where we focus on the implications of tariffs in our model.

Somewhat related to our work is also Bagwell and Staiger (1999), which proposes an efficiency explanation for the most-favored-nation (MFN) rule in tariff negotiations. Bagwell and Staiger show that, if the MFN rule is used in concert with the rule of reciprocity, the outcome of tariff negotiations is a point on the Pareto frontier (the “politically-optimal” outcome) that is defined without reference to the governments’ negotiating powers. They then argue that, if production requires irreversible investment (sunk costs), a system based on these rules may be preferred *ex-ante* not only by weaker countries but also by powerful countries, because it encourages weaker countries to participate in the agreement.⁵

⁵Bagwell and Staiger (2001b) points out another efficiency rationale for the MFN rule. This paper considers

There is a growing literature on the use of product standards as a protectionist tool. One example is Fischer and Serra (2001). Their basic setting is similar to ours; the main difference is that they have oligopolistic, rather than competitive, markets. They show that if a national-welfare-maximizing government chooses standards in a noncooperative way, the resulting standard levels will be inefficient from a global point of view, even if standards are constrained to satisfy national treatment (a result that is fully consistent with our analysis). Fischer and Serra do not consider international agreements on product standards, which is the main focus of our paper. Another example of this literature is Sturm (2000). This paper considers a setting where the choice of product standards is inefficient because of political-economy forces influencing the government. Sturm also considers the effect of a mutual-recognition rule, whereby the authority over standards on imported goods is effectively transferred to the exporting country, but does not have a systematic analysis of international agreements on product standards.

Finally, our theoretical approach bears an interesting relation with a school of thought in law-and-economics that views courts as “filling the gaps” of incomplete contracts (e.g. Egglestone, Posner and Zeckhauser, 2000). According to these authors, courts fill contractual gaps based on the principle of what the parties would have agreed to if they had contracted *ex ante* (what we have called the “umbrella” principle). These authors further argue that the courts should take this active role only if they possess enough information on the parties’ objectives. This argument finds a parallel in our results on the efficiency role of the DSP. But note the difference with respect to our setting: as we have argued, the umbrella principle for product standards is explicitly stated in the agreement, so in our case the DSP need not be construed as an “activist” court that imposes fresh commitments on the contracting parties, but can be seen as one that applies – though with subjective criterions – the agreed-upon rules.⁶

a multi-country world where bilateral tariff negotiations exert negative externalities on third countries and show that imposing MFN and reciprocity on bilateral negotiations helps preserve the welfare of third countries, and hence increases ex-ante efficiency.

⁶We should also mention the literature on labor arbitration (e.g. Crawford, 1979, Farber, 1980, Gibbons, 1988). The similarity between these models and ours is that the presence of an arbitrator influences the outcome of the negotiation between two parties, but the structure of the game is quite different. More importantly, these models do not *explain* the role of arbitration. They simply postulate the presence of an arbitrator with an (exogenously specified) objective function, and study its implications on the outcome of negotiations. In particular, contractual incompleteness plays no role in explaining the role of arbitration in this literature.

The remainder of the paper is organized as follows: in section 2 we present the basic model, focusing on the complete-contracting benchmark; in section 3 we extend the model to allow for contract incompleteness; in section 4 we discuss the robustness of the model to some extensions.

2. The basic model

Consider two countries, Home and Foreign. Asterisks denote the foreign country. In each country there is a population of identical citizens which is a continuum of measure L . Each citizen owns one unit of labor. We assume that L is large enough that every good is produced in equilibrium. There are $N + N^* + 1$ sectors. Good $j = 0$ is the numeraire, which is produced one-for-one from labor. Each nonnumeraire good is produced with labor (which is perfectly mobile within a country but cannot cross borders) and a specific factor, K_j . Technology is identical in the two countries. The supply of each specific factor is fixed and equal to one in both countries: $\bar{K}_j = \bar{K}_j^* = 1$ for all $j = 1, \dots, N + N^*$.

For goods $j \in \{1, \dots, N\} \equiv \mathcal{G}$ there is demand only in Home, so these goods are imported by Home in equilibrium. For goods $j \in \{N + 1, \dots, N + N^*\} \equiv \mathcal{G}^*$ there is demand only in Foreign, so these goods are imported by Foreign in equilibrium. Under this assumption, for each given product the standard set in the exporting country is irrelevant, so we can focus on standards in the importing country. This is useful for expositional purposes: our points are more transparent in a setting where standards in the exporting country are not an issue. In section 4 we will examine how results change when there is demand also in the exporting country.⁷

Each good j comes in a continuum of versions. Each version is characterized by the amount of a (good-specific) toxic, which can be a poisonous ingredient or a pollutant;⁸ we let e_j denote the amount of toxic per unit of good j . For example, e_j could be the amount of lead contained in a unit of wall paint, or the level of carbon monoxide emissions from a car, or the amount of pesticide in an apple, or the amount of nicotine in a cigarette. Producing a safer product is

⁷We could have simplified further, by assuming that there is no import-competing production. But this is a complication that we like to have, because we want to be able to talk about national-treatment rules.

⁸In the case of a pollutant, we assume that this is released at the act of consuming the good (as in the case of automobile emissions). This is slightly different from the more standard models of pollution, e.g. Copeland and Taylor (1994), where pollution occurs at the act of production.

more expensive; we capture this by postulating the following production function for good j :

$$y_j = A_j(e_j)F^j(k_j, l_j)$$

where F^j is a smooth function with constant returns to scale, $F_k^j, F_l^j > 0$, $F_{kk}^j, F_{ll}^j < 0$, $e_j \in [0, 1]$ (normalization) and A_j is a strictly increasing, smooth function. Further, to help ensure that the equilibrium standards are interior, we assume the following boundary conditions: $F_l^j(1, 0) < \infty$, $F_{ll}^j(1, 0) > -\infty$, $A_j(0) = 0$, $0 < A'_j(0) < \infty$ and $A'_j(1) = 0$.⁹

The utility function for an individual i in the home country is:

$$U^i = x_0^i + \sum_{j=1}^N [u_j(x_j^i) - H_j(E_j)]$$

where x_j^i is the individual consumption of good j , E_j is the aggregate amount of toxic associated with consumption of good j in the home country and H_j is a damage function that captures the adverse health effects of toxics. Note that this damage is a pure externality: an individual's health depends only on the aggregate amount of toxic E_j .¹⁰ H_j is assumed to be smooth and strictly increasing, with $H'_j(0) < \infty$. In general the standard on imports may differ from the standard on domestic production, therefore we write

$$E_j = e_j y_j + e_j^m m_j$$

where y_j is domestic production of good j , m_j denotes imports of good j , e_j is the standard for domestic producers of good j and e_j^m is the standard for imports. The sub-utility functions u_j are smooth and satisfy $u'_j > 0$, $u''_j < 0$ and $u'_j(0) = \infty$.

Since the toxic in a good generates a pure externality, an individual's demand function will depend only on the good's price, not on e_j . For this reason, in the absence of regulation firms will choose $e_j = 1$. Because of the externality, a government may want to impose a maximum limit on e_j . We will address the endogenous determination of standards later on; for the moment we think of standards as exogenously fixed.¹¹ We abstract from the possibility of trade taxes; we will discuss in a later section the robustness of the results to the introduction of trade taxes.

⁹All derivative functions are extended by continuity at the frontier.

¹⁰A more realistic assumption would be that an individual's utility depends separately on the aggregate amount of toxics and on the toxics generated by the individual's consumption. But this would complicate the analysis without adding much insight.

¹¹We ignore alternative policy instruments that could correct this externality, such as Pigouvian taxes. The main conclusions of the analysis would not change if these were allowed.

Next we write the utility function for a representative individual in the foreign country. Now that the nature of the externality is clear, we can omit the i superscript:

$$U^* = x_0^* + \sum_{j=N+1}^{N+N^*} [u_j(x_j^*) - H_j(E_j^*)]$$

where $E_j^* = e_j^* y_j^* + e_j^{*m} m_j^*$.

We can now turn to the analysis of the model. We first characterize the competitive equilibrium for given standards, and then endogenize the choice of standards.

First note that wages are given by $w = w^* = 1$. We can derive the market equilibrium separately for each good j . Let us focus on goods imported by Home. The amount of labor employed by Home in sector j , denoted by L_j , is determined implicitly by:

$$1 = p_j A_j(e_j) F_l^j(1, L_j)$$

where p_j is the price of good j . We can think of the equilibrium level of L_j as a function of $p_j A_j(e_j)$, which we denote $L_j(p_j A_j(e_j))$. Clearly, this function is continuous and increasing. The amount of labor employed by Foreign in sector j can be derived similarly as a function of $p_j A_j(e_j^m)$. Given our symmetry assumptions, this is given by $L_j(p_j A_j(e_j^m))$.

The demand function for good j is implicitly defined by $u_j'(x_j) = p_j$. We denote this demand function by $d_j(p_j)$, and let $\varepsilon_j^d \equiv -d_j'(p_j)p_j/d_j(p_j)$ denote its elasticity. To simplify the analysis of the equilibrium choice of standards, we assume that demand is elastic, i.e. $\varepsilon_j^d > 1$.

Domestic and foreign supply of good j are respectively given by $A_j(e_j)F^j(1, L_j(p_j A_j(e_j)))$ and $A_j(e_j^m)F^j(1, L_j(p_j A_j(e_j^m)))$. In equilibrium therefore we must have:

$$A_j(e_j)F^j(1, L_j(p_j A_j(e_j))) + A_j(e_j^m)F^j(1, L_j(p_j A_j(e_j^m))) = d_j(p_j) \quad (2.1)$$

This condition determines p_j as a function of standards, $p_j(e_j, e_j^m)$, $j \in \mathcal{G}$. Similarly one can derive the equilibrium prices for goods imported by Foreign, which we denote $p_j^*(e_j^*, e_j^{*m})$, $j \in \mathcal{G}^*$. It is direct to verify that each price function is decreasing in both arguments. This is intuitive. Consider for example $p_j(e_j, e_j^m)$. An increase in either e_j or e_j^m increases supply for two reasons: a direct increase of productivity (i.e. an increase in A_j), and an increase in labor demand. As a consequence, the price must fall.

Next we can find the equilibrium rate of return for each specific factor. For goods imported by Home, this is given by $\pi_j(p_j A_j(e_j)) = p_j A_j(e_j) F_k^j(1, L_j(p_j A_j(e_j)))$. Given the symmetry in technology and endowments, the corresponding rental rate for the foreign capital is given by $\pi_j(p_j A_j(e_j^m))$. For goods imported by Foreign, the rental rates are respectively $\pi_j(p_j^* A_j(e_j^{*m}))$ and $\pi_j(p_j^* A_j(e_j^*))$.

2.1. The noncooperative standards game

We can now turn to the endogenous determination of standards. Each government maximizes national welfare. Letting $s_j(p_j) = u_j(d_j(p_j)) - p_j d_j(p_j)$ be the consumer surplus for good j , we can write the aggregate indirect utilities for Home and Foreign as follows

$$\begin{aligned}\Omega(\mathbf{e}, \mathbf{e}^*) &= L + \sum_{j=1}^N [\pi_j(p_j A_j(e_j)) + s_j(p_j) - H_j(e_j y_j + e_j^m m_j)] + \sum_{j=N+1}^{N+N^*} \pi_j(p_j^* A_j(e_j^{*m})), \\ \Omega^*(\mathbf{e}, \mathbf{e}^*) &= L + \sum_{j=1}^N \pi_j(p_j A_j(e_j^m)) + \sum_{j=N+1}^{N+N^*} [\pi_j(p_j^* A_j(e_j^*)) + s_j(p_j^*) - H_j(e_j^* y_j^* + e_j^{*m} m_j^*)],\end{aligned}$$

where \mathbf{e} and \mathbf{e}^* denote the vectors of home and foreign standards and p_j, y_j, m_j (p_j^*, y_j^*, m_j^*) are the market-equilibrium values given e_j and e_j^m (e_j^* and e_j^{*m}).

We seek to characterize the Nash equilibrium of the game in which governments choose standards simultaneously. Note first that, given the separability of payoffs, a government's optimal choice of standards is independent of the standards in the other country. Hence we have a dominant strategy equilibrium. In particular, the equilibrium standards for a good imported by Home solve:

$$\max_{0 \leq e_j, e_j^m \leq 1} [\pi_j(p_j A_j(e_j)) + s_j(p_j) - H_j(e_j y_j + e_j^m m_j)]$$

and the equilibrium standards for a good imported by Foreign solve:

$$\max_{0 \leq e_j^*, e_j^{*m} \leq 1} [\pi_j(p_j^* A_j(e_j^*)) + s_j(p_j^*) - H_j(e_j^* y_j^* + e_j^{*m} m_j^*)].$$

Generically, these problems have unique solutions, hence the Nash equilibrium is generically unique.¹² Note also that, given our assumptions, the equilibrium (domestic and import) stan-

¹²In what follows we suppose that the equilibrium is unique, but we note that, if there were multiple equilibria, our results would hold for any equilibrium.

dards are interior. Our next result states that for each good the importing country imposes a stricter standard on imported units than on domestically produced units. In other words, governments discriminate against foreign producers in the noncooperative equilibrium. The proof of this and all other results can be found in Appendix.

Proposition 1. *At the Nash equilibrium, $0 < e_j^m < e_j < 1$ for all $j \in \mathcal{G}$ and $0 < e_j^{*m} < e_j^* < 1$ for all $j \in \mathcal{G}^*$.*

To gain intuition for this result, consider starting from a position where $e_j = e_j^m$. Then the importing country can increase its welfare by reducing e_j^m and increasing e_j . Consider a small variation $\Delta e_j = -\Delta e_j^m > 0$: the increase in domestic production is approximately compensated by the decrease in imports and the first order effect on price, consumer surplus and negative externality is null, but the first order effect on Home producer surplus is positive.

2.2. Complete-agreement benchmark

Our next objective is to characterize the distortions involved in the Nash equilibrium and examine how an international agreement can address these distortions.

The signs of the distortions in the noncooperative equilibrium depend on the signs of the externalities that the importing country's policies exert on the exporting country. To determine these signs, let us focus on a good imported by Home. The first policy externality to consider is the effect of the domestic standard on the exporting country, $\frac{\partial \Omega^*}{\partial e_j}$. Since Ω^* depends on e_j only through p_j , and p_j is decreasing in e_j , then $\frac{\partial \Omega^*}{\partial e_j} < 0$. This implies that the Nash equilibrium entails excessively lax domestic standards, in the sense that both countries can be made better off by tightening domestic standards on all goods.

Next we consider the effect of the import standard on the exporting country, $\frac{\partial \Omega^*}{\partial e_j^m}$. Since e_j^m affects foreign welfare only through foreign producer surplus, and this is an increasing function of $p_j A_j(e_j^m)$, then $\frac{\partial \Omega^*}{\partial e_j^m} > 0$ if and only if $p_j A_j(e_j^m)$ increases with e_j^m . This condition is ensured given our assumption that the demand elasticity is higher than one, $\varepsilon_j^d > 1$ (see Lemma 2 in

the Appendix).¹³ Therefore the Nash equilibrium entails excessively strict import standards, in the sense that both countries can be made better off by relaxing import standards on all goods. The model can then explain the occurrence of “green protectionism”, which is a frequent complaint in reality: noncooperative import standards are inefficiently strict.

Next we model explicitly the bargaining process that leads to an agreement. We assume that agreements are perfectly enforceable.

We consider a bargaining game *à la* Rubinstein (1982) with alternating offers, assuming that governments play the one-shot Nash equilibrium standards until an agreement is reached.¹⁴ We will think of the bargaining process as compressed in a small period of time. It is well known that, as the interval between offers goes to zero, the equilibrium outcome of this game converges to the Nash bargaining solution with threat point given by the Nash equilibrium and with equal bargaining powers. The reason we model explicitly the noncooperative structure of the game, rather than adopting a reduced-form cooperative approach such as the Nash bargaining solution, is that in the next section we will introduce a further element in the bargaining process, namely the option of invoking arbitration by the DSP, and at that point there will be no ready-made cooperative bargaining notion to invoke.

The equilibrium agreement is the one that solves the following problem:

$$\max_{\mathbf{e}, \mathbf{e}^*} (\Omega - \Omega^{NE})(\Omega^* - \Omega^{*NE})$$

where Ω^{NE} and Ω^{*NE} denote the governments’ one-shot Nash equilibrium payoffs. Clearly, the outcome of the bargaining is a point on the Pareto frontier. The exact point selected by the agreement depends on where the Nash equilibrium point lies relative to the frontier, which in

¹³If demand elasticity is lower than one, foreign producer surplus may be increasing or decreasing in e_j^m . Other things equal, it is more likely to be increasing when demand is more elastic, when supply is more elastic, and when import penetration is lower. We have assumed $\varepsilon_j^d > 1$ to simplify the exposition and to focus on (what we think is) the more interesting and realistic case: after all, in reality exporters often complain that import standards are too strict, they never complain that they are too lax. At any rate, the key insights of the analysis do not depend on this restriction.

¹⁴There may well be equilibria where countries do not play the static Nash equilibrium when an agreement is not reached. In particular, if the relevant discount factor is high, as we assume below, an efficient outcome can be supported as a subgame perfect equilibrium by Nash reversion strategies that never reach a binding agreement. But since we are assuming that binding agreements can be signed, it seems natural to focus on externally enforced agreements and hence on strategies playing the static Nash equilibrium whenever a binding agreement is not reached.

turn depends on the fine details of the model. The only general statement we can make is the following: for each good, the equilibrium agreement relaxes the import standard and/or tightens the domestic standard relative to the Nash equilibrium. This is a straightforward implication of the argument made above that the Nash equilibrium entails excessively tight import standards and excessively lax domestic standards.¹⁵

3. Incomplete agreements

Thus far we have considered a simple complete-contracting benchmark. Now, to capture the role of contract incompleteness, we add a further stage to the model. At time t_0 , the model is the same as in the previous section. At time t_1 a new product appears. This product is drawn randomly from a large set of potential products, all of which satisfy the assumptions we made at the outset on technology, demand and externality (including the assumption that demand is concentrated in one country).¹⁶

Governments are ex-ante symmetric with respect to future potential goods. In particular, for each potential good there is a 50% probability that demand will be concentrated in Home, so that Home will import the good, and a 50% probability that demand will be concentrated in Foreign, so that Foreign will import it. Thus we can think of governments as contracting ex ante under a veil of ignorance. We will later consider the case in which governments are ex-ante asymmetric.

We suppose that the characteristics of the new product at t_1 can be captured by a vector of parameters θ . This is a (potentially very long) vector that includes all the relevant information about the new product, including technology, demand and externalities. Viewed from time t_0 , θ is a random variable distributed according to the (common-knowledge) c.d.f. $\Phi(\theta)$. We will

¹⁵The reason for the “and/or” statement is that we cannot rule out an agreement where, for example, both the import standard and the domestic standard for a given good are looser than in the Nash equilibrium. Intuitively, starting from the Nash equilibrium, both countries can be made strictly better off if all import standards are relaxed and the domestic standards are not changed. By continuity, a Pareto-improvement can still be achieved by slightly relaxing some domestic standards.

¹⁶In a more complete dynamic model with more than two periods, we would have a stochastic process for the appearance of new products; however, to the extent that new products tend to appear one at a time, we would expect our qualitative results to survive.

use the symbol θ as index of the new good, and let Ω_θ^m (resp. Ω_θ^x) be the incremental payoff function associated with good θ for the importing (resp. exporting) country.

The notion of "new" product need not be interpreted literally. Similar results would obtain if we had a new version of an existing product. For example, suppose at t_0 product j comes in a single, "safe" version e_j^0 (e.g. non hormone-treated beef) and at t_1 a new version e_j^θ appears (e.g. hormone-treated beef), which is cheaper and possibly less safe. Also suppose that it is efficient to allow sales of version e_j^0 , while it may or may not be efficient to allow sales of version e_j^θ , depending on θ . Then the model would yield analogous results, with some adaptations due to the zero-one nature of the policy choice. Yet another interpretation of the "new" product is an old product with changed features. For example, suppose at t_1 new scientific evidence is received about the health damage caused by an existing product j . This means that the H function associated with this product has changed. In this case we would effectively have a new product replacing an old one, and the analysis would apply virtually unchanged. In sum, although we chose to focus on the case of a new product line, the analysis can also be interpreted as applying to unforeseen contingencies concerning an existing product.

We assume that the agreement at t_0 does not cover future products. This assumption can be justified in terms of contracting costs. If there is a very large number of potential products and each of them has a very small probability of appearing tomorrow, then it is more efficient to wait until tomorrow and write an agreement after the new product has appeared, rather than writing a complete contingent agreement today.¹⁷ Alternatively, one might take the view that governments are not able to foresee or describe future potential products, in which case contracting at t_0 over future products is not even feasible. We are assuming for simplicity that governments have in mind a probability distribution over future potential products, but our main arguments can be extended to a setting of unforeseen contingencies.

At t_1 , governments get together to bargain over standards for the new product. To simplify the analysis, we do not allow governments to renegotiate policies for the old products at t_1 . This assumption seems reasonable given that (i) in practice there are important costs of renegotiating

¹⁷If we introduced contracting costs explicitly in the model, we would find that for sufficiently high contracting costs it is optimal to leave some existing products out of the agreement at t_0 , but this would not change anything substantial in our conclusions.

agreements and implementing changes in product standards, and (ii) we interpret the "ex-post" stage of the model as capturing a whole stream of future stages, so that renegotiating every time would be extremely costly. At any rate, our qualitative results survive even if renegotiation is costless – more on this later in the section.

The bargaining process at t_1 has the same structure as the one at t_0 : governments make alternating offers in rapid sequence, and keep the standards for the new product at their non-cooperative levels until an agreement is reached (meanwhile the standards for the old goods are kept at their agreed-upon levels, since there is no renegotiation and agreements are enforceable). We suppose for the moment that international transfers are not available; we will examine the role of transfers later in the section. We allow governments to randomize the choice of standards – randomization will play no direct role in the analysis, but it ensures that the set of feasible utilities is convex.

Given that payoffs are separable across goods, and given the no-renegotiation assumption, we can focus separately on the agreements at t_0 and t_1 . At time t_0 , the outcome of the bargaining is the same as in the previous section.

At time t_1 , the outcome of the bargaining can be determined focusing on the new good in isolation. The key step is to identify the location of the threat point, that is, the noncooperative payoff point. Since in a noncooperative situation the standards for the new good are both chosen by the importing country, it follows that the threat point is on the Pareto frontier. This in turn implies that the Nash bargaining outcome coincides with the threat point. In Figure 1, we depict the set of feasible utilities in the $(\Omega_\theta^x, \Omega_\theta^m)$ space. Point N, that is, the point that maximizes Ω_θ^m , is the Nash equilibrium and Nash bargaining outcome. Notice that it is not possible to Pareto-improve over the threat point.

The outcome at t_1 is ex-post Pareto efficient, but inefficient from an ex-ante perspective. This inefficiency is due to the incompleteness of the ex-ante agreement. To better understand the nature of this inefficiency, let us derive the first-best outcome, i.e. the outcome that would obtain if governments could costlessly write a complete contract at t_0 . For the sake of simplicity, we suppose that Home and Foreign are symmetric not only with respect to future products, but also with respect to existing products, so that the bargaining problem at t_0 is fully symmetric.

In this case, a complete contract at t_0 would maximize the governments' common expected (incremental) payoff, which is given by $\int (\frac{1}{2}\Omega_\theta^x + \frac{1}{2}\Omega_\theta^m)d\Phi(\theta)$.¹⁸ This boils down to maximizing $\Omega_\theta^x + \Omega_\theta^m$, the unweighted sum of the countries' ex-post utilities, for each potential product θ . In contrast, under incomplete contracts the outcome maximizes Ω_θ^m . Graphically, the first-best outcome is the point on the bargaining frontier that has slope minus one (point W in Figure 1).

The nature of this inefficiency can be understood in terms of *missing issue linkage*. The ex-post bargaining problem involves a single issue, which is very asymmetric: the exporting country has nothing to offer in exchange for more favorable policies by the importing country, hence no mutual gains are possible relative to the noncooperative equilibrium. With a complete contract, in contrast, governments can exchange “concessions” across states of the world. Each state of the world is associated with an “issue,” with a corresponding bargaining frontier. A complete contract would link all these (potential) issues together, and this beneficial linkage is what is missing in an incomplete-contracting world.¹⁹

3.1. The role of DSP

Next we ask whether an ex post dispute settlement procedure can remedy, at least in part, the inefficiency created by incomplete trade agreements. We think about the DSP in the following terms: when a new product appears, governments bargain over policies for the new product, but each government has the option of invoking arbitration by the DSP, in which case the DSP's decision is binding. More precisely, we modify the rules of the bargaining game as follows. Governments make alternating offers in rapid sequence. At any point during this process each government has the option of invoking the DSP. If the option is exercised by either government, the bargaining ends and the DSP immediately issues a determination, which is binding on both parties. Until an agreement is reached or the DSP is invoked, policies for the new product are

¹⁸If countries were asymmetric with respect to existing products, the ex-ante agreement about potential products could be biased in favour of one country. We will discuss the role of our assumption of ex-ante symmetry in section 4.

¹⁹The reader may be tempted to interpret this inefficiency in terms of incomplete risk sharing, but this would be misleading. First, governments are risk neutral in our setting. Second, the gains from issue linkage arise also with linear utility frontiers, thus they are not due to concavity of the utility frontiers. See for example Horstmann, Markusen and Robles (2000), who analyze in detail the gains from issue linkage in bargaining.

kept at their noncooperative levels. This bargaining game is similar to the bargaining game “with outside options” as presented by Osborne and Rubinstein (1990) (see also the survey by Binmore *et al.* (1992)).

The DSP, if called upon to arbitrate, makes decisions based on a general *principle*, that is an objective function to be maximized. Given that the first best outcome maximizes $\Omega_\theta^x + \Omega_\theta^m$, it is natural to assume that this is the DSP’s objective.²⁰ This principle can be interpreted as applying the “spirit” of the ex-ante agreement, i.e. what the parties *would have agreed to* had they contracted ex ante on the contingency at hand. In the case of product standards in the WTO, we would argue that this principle of global welfare maximization is stated in the actual agreement, in the form of the “no-disguised-protectionism” principle (see our discussion in the introduction).

The DSP has only limited information about the new product θ and hence about the welfare functions Ω_θ^m and Ω_θ^x .²¹ In particular, if the DSP is invoked it will observe a noisy signal of θ , say $s = \theta + \varepsilon$, where ε is a vector of noise terms that, conditional on θ , are *iid* with mean zero and variance $\sigma_\varepsilon^2(\theta)$ (thus the noise in the signal can depend on the nature of the product). We let σ_ε^2 denote the unconditional variance of each noise term (this captures the noise in the DSP’s information as viewed from an ex ante perspective). The signal s can be interpreted as the information that the DSP is able to gather through an independent investigation. For the moment we do not consider the possibility of eliciting information from the governments by setting up a message-based game (hearing); we will discuss this possibility later in the section. Let $\Phi(\theta|s)$ be the posterior c.d.f. for θ given s . The DSP, if called upon to arbitrate, will then solve

$$\max_{e, e^m} \int [\Omega^x(e, e^m, \theta) + \Omega^m(e, e^m, \theta)] d\Phi(\theta|s),$$

where $\Omega^c(e, e^m, \theta)$ ($c = x, m$) is the parametrized incremental payoff function of country c for

²⁰We would argue that this is the natural objective for the DSP not only if governments can foresee all potential products, but also in the case that they cannot, as long as they see themselves as ex-ante symmetric. What is required is that, for each state of the world, there is another state of the world which is equally likely and has mirror-image payoff implications. It is *not* required that governments have in mind a probability distribution over all possible states of the world.

²¹The reader may wonder why we did not simplify the model by assuming verifiable welfare functions. The reason is that, if welfare functions were verifiable, the first best outcome could be easily implemented by a simple clause of the form “standards must maximize $\Omega_\theta^m + \Omega_\theta^x$ ”, as this clause would be perfectly enforceable. The assumption of nonverifiable welfare functions makes the incomplete-contracting problem interesting.

product θ . Governments have more information than the DSP. A simple way to capture this asymmetry of information is to assume that at t_1 (before bargaining takes place) governments observe the true realization of θ . From the governments' point of view, then, the expected payoff pair resulting from the DSP's decision can be thought of as a point in $(\Omega_\theta^x, \Omega_\theta^m)$ space. If the DSP could observe θ perfectly, this point would be the first-best point (W in Figure 1). But because the DSP's information is imperfect, this point is (generically) below the Pareto frontier. Governments bargain efficiently, therefore the equilibrium outcome will still be on the Pareto frontier, but the DSP in general will not be able to implement the first-best point W.²²

Before proceeding with the analysis, we pause to discuss our assumption on the behavior of the DSP. We assumed that, if invoked, the DSP imposes exact levels of standards on the importing government. We could make different assumptions, that would perhaps be more realistic but theoretically less natural. For example, we could assume that the DSP determines only whether the (noncooperative) standards currently in place are acceptable and, if not, in what direction they should be changed; and that the DSP can be invoked repeatedly.²³ We would expect the results to be similar in this case. Note that, if the disputed standard is of the zero-one type (e.g., the issue is whether or not it is legitimate to ban hormone-treated beef, as in the example discussed in section 3), then there is no difference between this mechanism and the one we assumed.

We can now analyze the bargaining game between the governments. In the presence of the DSP, we have a game of bargaining with outside options, where the status quo is the one-shot Nash equilibrium (point N in Figure 1) and the outside option for each player is given by the DSP expected payoff pair. Again, we look at the limit (subgame perfect) equilibrium outcome of sequential bargaining as the interval between offers becomes small. It turns out that, if the DSP point is dominated by point N , then the DSP is irrelevant and the (limit) outcome of

²²Given that the DSP's decision influences the bargaining outcome by creating an outside option, it is possible that committing strategically to an objective that differs from expected welfare might achieve a higher expected welfare. However this commitment may not be credible, in the following sense: suppose the DSP selects a judge to hear the case, and the judge genuinely cares about expected welfare; then, if the dispute subgame is actually reached, the judge will have an incentive to maximize expected welfare.

²³A case of iterative dispute, where the DSP has been invoked more than once by the exporting country because the importing country (allegedly) did not go far enough in changing their standards, is given by the U.S.-Japan dispute on quarantine regulations for fruit. Outside the area of product standards, the U.S.-E.U. banana dispute has been a well-known case of iterative dispute.

bargaining is still N . If, on the other hand, the DSP gives one of the players a higher utility than at N , then the equilibrium is the point on the frontier where this player gets the DSP utility. To understand the effect of DSP on the equilibrium outcome, we divide the set of feasible (incremental) payoff pairs in three regions (see Figure 1):

Region I: Ω_θ^x is lower than at point N .

Region II: Ω_θ^x is higher than at point N and $\Omega_\theta^x + f(\Omega_\theta^x)$ is higher than total welfare at N , where $(\Omega_\theta^x, f(\Omega_\theta^x))$ is the point on the Pareto frontier with abscissa Ω_θ^x .

Region III: Ω_θ^x is higher than at point N and $\Omega_\theta^x + f(\Omega_\theta^x)$ is lower than total welfare at point N .

Let $(\Omega_{DSP}^x, \Omega_{DSP}^m)$ denote the expected DSP point. If $(\Omega_{DSP}^x, \Omega_{DSP}^m)$ is in region *I*, the bargaining outcome is N and DSP is irrelevant; if it is in region *II* or *III*, the bargaining outcome is $(\Omega_{DSP}^x, f(\Omega_{DSP}^x))$; therefore if the DSP point is in region *II*, DSP improves aggregate welfare; if it is in region *III*, the DSP decreases aggregate welfare. Note that the DSP can be harmful only if, due to imperfect information, the DSP decision is expected to be very biased in favour of the exporting country. This may well happen if, given the realization of θ , the DSP noise $\sigma_\varepsilon^2(\theta)$ is high. Intuitively, however, this event is ex ante unlikely.²⁴

We can now examine the desirability of DSP from an ex-ante standpoint, as viewed from time t_0 . At t_0 , governments know only the ex-ante joint distribution of θ and ε . Based on these distribution, they attach some probability to the event that the expected DSP point will fall in region *III* for the new product, hence that the DSP will be harmful. This probability is guaranteed to be small (and possibly zero) if σ_ε^2 is small. We can conclude that the DSP improves ex-ante welfare if σ_ε^2 is sufficiently low:

²⁴This intuition is easily explained with an example. Suppose potential goods are parametrized by the marginal externality of consumption, e.g., $H_\theta(E) = \theta E$, and the prior distribution on θ is concentrated on intermediate values. Now suppose the realization of θ is extremely high, so that the optimal (common) standard is very strict. In terms of payoffs, point W is close to point N , because the exporting country's welfare is small at both points. If $\sigma_\varepsilon^2(\theta)$ is sufficiently high, countries expect the observed value of θ to be in the right tail of the DSP distribution $\Phi(\cdot|s)$, hence expect the DSP to rule for a permissive standard. But then the bargaining outcome will be to the far right of point W , at a point where total welfare is lower than at N . This bad outcome is caused by an unexpectedly large realization of θ , an ex ante unlikely event.

Proposition 2. *If σ_ε^2 is sufficiently low, a dispute settlement procedure increases ex ante efficiency.*

In spite of the fact that governments bargain efficiently ex post and have more information than the DSP, the analysis shows that the presence of the DSP can enhance ex ante efficiency, provided the imperfection in the DSP's information is not too serious. It is worth summarizing the logic of this result. Due to the incompleteness of the ex ante agreement, governments fail to link issues across states of the world. The DSP can improve ex-ante efficiency by affecting the *distribution* of utilities in the ex post bargaining, which it can do by offering an outside option to the bargainers. In particular, the DSP will shift the outcome in favor of the “weaker” country (which in our model is the exporting country), and this indirectly recovers some of the gains from issue linkage that were left unexploited by the incomplete agreement. Notice that the DSP does not intervene in equilibrium: it is the threat of invoking the DSP that changes the outcome of the bargaining.

The role of the DSP can also be interpreted in terms of influencing the balance of power between governments. In the absence of DSP, there is a strong imbalance of power in the ex-post bargaining; in particular, the importing country has much more power, since it attains the maximum feasible utility in the bargain. The DSP can increase ex-ante utilities by smoothing out ex-post power imbalances. The DSP can then be viewed as replacing a “power-based” system with a “rules-based” system, in the sense that it makes decisions based on an anonymous rule (joint welfare maximization) that is independent of the countries' ex-post bargaining positions.

The above reasoning suggests that the value of DSP depends crucially on the degree of asymmetry in the ex post bargaining game. In our basic model, the ex post imbalances of power are extreme. We can capture situations where ex post negotiations are more balanced with a simple extension of the model: suppose that with probability ρ the new good comes in two symmetric versions, each exported by one country. Thus the state of the world can be of two types: asymmetric (with probability $1 - \rho$) or symmetric (with probability ρ). This is a simple way to capture the (expected) degree of ex post asymmetry with a single parameter (ρ). While this structure is exceedingly simple, it allows us to make a point that is quite general.

Consider the ex post bargaining game in a symmetric state. In figure 2, Ω_{t_1} and $\Omega_{t_1}^*$ represent the domestic and foreign incremental payoffs at t_1 . The threat point N is below the bargaining frontier. The bargaining outcome in the absence of DSP (point B) is the point of the frontier that has slope minus one, which maximizes $\Omega_{t_1} + \Omega_{t_1}^*$. Thus, in a symmetric state the possibility to invoke the DSP can only do harm, if it affects the bargaining outcome at all. It is then immediate to prove the following result:

Proposition 3. *Suppose there is a probability ρ that the state is symmetric at t_1 , and σ_ε^2 is sufficiently low. Then the efficiency gains from a dispute settlement procedure are decreasing in ρ .*

This provides a qualification to Dixit’s (1996) argument that a DSP can be a partial remedy to contract incompleteness: the analysis shows that this argument must rely on the presence of asymmetries in the ex post bargaining problem, and that the value of arbitration is higher when these asymmetries are more pronounced. Intuitively, if ex post bargaining is more symmetric, there is less scope for linking issues across states, and hence arbitration can offer fewer efficiency gains.

Our basic model depicts the DSP investigation as a pure fact-finding activity. A complementary way to gather information about the state of the world θ might be to set up a hearing where governments can send cheap-talk messages regarding θ . There exists a small literature on cheap-talk games where two “experts” send messages to a less informed decision maker, the most prominent examples being Krishna and Morgan (2001) and Battaglini (2002). The results of these papers suggest that, in a setting where the experts have conflicting interests, as in our model, the decision maker (DSP) may be able to elicit *some* information from the experts (governments) but is unlikely to learn the state of the world with certainty.²⁵ This suggests

²⁵We can be a little more precise here. The above mentioned papers analyze situations where the state θ and the policy action of the decision maker are vectors with the same dimension. Negative results (non-existence of fully revealing equilibria in the hearing game) are obtained in the one-dimensional case with strong conflict of interests (Krishna-Morgan, 2001, Propositions 4 and 5; Battaglini, 2002, Proposition 2). In our model, θ can be multi-dimensional, but the policy decision is essentially uni-dimensional (it is easy to show that the DSP always chooses a common standard). The conflict of interests in our model is likely to be strong, because the exporting country prefers the most permissive common standard (it can be shown that $\pi(p(e, e)A(e))$ is increasing in e), while the importing country prefers a much stricter standard. Although our model does not perfectly fit the

a reinterpretation of the signal s observed by the DSP in our model: s can be thought of as incorporating the information that the DSP is able to extract from the hearings, and $\sigma_\varepsilon^2(\theta)$ can be interpreted as capturing the DSP's *residual* uncertainty about θ after the hearings. The unconditional variance σ_ε^2 again captures the overall noise in the DSP's information as viewed from an ex ante perspective.

Before proceeding, we want to discuss the notion of “arbitration” in our model. If invoked, the DSP selects policies according to the principle of joint-welfare maximization. One might object that in our model the DSP is just “enforcing” a clause of the ex-ante contract that reads “policies should maximize the joint welfare of the countries.” The reason we do not think of this process as enforcement is that the countries' welfare functions are not verifiable. By enforcement we mean that a “court” imposes stiff penalties if a (verifiable) rule is violated, whereas in our model the DSP only offers an outside option to the governments. In this sense, our model can rationalize the DSP as a genuine arbitrator, not a straight enforcer. This is in contrast with the previous theoretical literature on the DSP, in which the DSP is beneficial only to the extent that it can bring about more severe punishments for violations of the agreements. Also notice that, in the models that focus exclusively on the enforcement role of the DSP, any “out of court” settlement by the governments is not desirable, as it diminishes the punishment for violations. In our model, on the other hand, out-of-court settlement is desirable. Thus, our model can rationalize a trade institution that facilitates genuine dispute *settlement*, rather than retaliation.

International transfers

In this section we consider the case in which international transfers are available. We think of transfers as occurring in terms of the numeraire good. Consider the bargaining game at t_1 in the absence of DSP. When transfers are available, the Pareto frontier is the line with slope minus one that is tangent to the Pareto frontier of the game without transfers. It follows immediately that the equilibrium policies for the new good are the ones that maximize the joint surplus of the two governments. This in turn implies that the DSP is redundant in the

frameworks of Battaglini and Krishna-Morgan, their results suggest that full revelation would not occur in our case.

presence of transfers:

Proposition 4. *If international transfers are available, there are no efficiency gains from a dispute settlement procedure.*

This provides a second qualification to Dixit's (1996) intuition that a dispute settlement procedure can be a remedy to contract incompleteness: in the context of our model, a DSP can enhance efficiency only if international transfers are not available. It is worth pointing out however that this result depends on the assumption of quasi-linear utility. If utility is nonlinear in all goods, or governments are risk averse, then even if lump sum transfers are available the ex post bargaining frontier is not linear, and hence a DSP can provide ex ante efficiency gains.

At any rate, an interesting empirical question is whether and to what extent international transfers are available within trade organizations like the WTO. It is often argued that international monetary transfers are rarely used, but concessions in other policy areas can play the role of payments. This idea can be examined in our model by supposing that governments at time t_1 can renegotiate policies for the existing products, which we have thus far not allowed. The next subsection discusses this extension of the model.

Renegotiation of standards on existing products

In the basic model we assumed that it is too costly to renegotiate standards for the old products at t_1 . Here we briefly consider the implications of renegotiation on existing products. We will argue that the DSP may have a role to play even if this renegotiation is costless.

Suppose for simplicity that the bargaining problem at t_0 is symmetric, in the sense that both the Pareto frontier and the threat point are symmetric. Let Γ^O be the frontier for the old products. Clearly, the outcome of the bargaining at t_0 is the point on Γ^O with slope minus one. Now suppose that at t_1 governments bargain over standards for the new product, and can renegotiate standards for the old products. Let Γ^θ be the frontier for the new product. Recall that, in the absence of renegotiation, the outcome at t_1 is the point on Γ^θ with slope zero. If governments can renegotiate costlessly, they can achieve a Pareto improvement by moving

the point on Γ^O to the left (where the derivative is higher than minus one) and the point on Γ^θ to the right (where the derivative is lower than zero). In fact, it is easy to see that the overall bargaining will yield points on Γ^O and Γ^θ that have equal slopes, and the common slope will be between zero and minus one (if the two points did not have the same slope, a Pareto improvement would be possible). But we know that ex-ante Pareto efficiency requires that both slopes be equal to minus one, therefore the outcome of the overall bargaining is not ex-ante Pareto efficient. It follows that the DSP can provide an ex-ante Pareto improvement, provided σ_ε^2 is sufficiently small.

The intuition is simple: if utilities are nontransferable, first-best efficiency requires that governments include in the bargaining not only existing, but also potential, products. Renegotiation allows governments to include in the bargain all existing products, but not the potential products that did not materialize ex-post.

3.2. Rigid rules

There is another potential remedy to contract incompleteness: *rigid rules*, i.e. blanket rules that are not product-specific and apply to any present and future product. To make the exposition more clear, in this section we focus exclusively on rigid rules, supposing that DSP is not available. In the next section we will consider the possibility of using both rigid rules and DSP.

In the case of more traditional trade policies, such as tariffs, a focal rigid rule is free trade for all products. This is the simple solution adopted in free trade areas and customs unions. In multilateral agreements such as the WTO, rigid tariff bindings or tariff-reduction formulas have often been adopted (e.g., all tariffs in a certain sector are subject to a common tariff binding of $t\%$, or are slashed by $x\%$ relative to pre-agreement levels). One interpretation of these rigid rules for tariffs is that they may serve to cope with contract incompleteness and transaction costs.

In the case of nontariff barriers, and particularly product standards, there is more limited scope for rigid rules. Standards are measured in different units for different goods, and cannot

be specified in ad-valorem terms. Standards are specific in nature, and this makes it hard to impose common bounds on standards for different goods. A rigid rule that is adopted by most existing agreements, and most notably the WTO, is the *nondiscrimination* principle, which has two components: (i) a most-favored-nation rule, which requires that import standards for a given good be the same regardless of the source of imports, and (ii) a national-treatment rule, which requires that the same standard be applied to imported units and domestically produced units of each good. Our two-country model is able to capture the second form of nondiscrimination (ND) rule, which amounts to imposing a constraint $e_j = e_j^m$ for all goods.

The ND rule is not the only feasible rigid rule in our model. A rigid rule in general is a constraint on (e_j, e_j^m) that is independent of the product j . However, if the rule is to specify unambiguous, verifiable obligations, the constraint must be independent of the scale of measurement. This in turn implies that the constraint must be *ordinal* in nature, because only ordinal constraints are invariant to monotonic transformations of the measure.²⁶

To derive the set of possible ordinal constraints in our model, consider the following three types of rule:

(A) Rules that constrain the ranking between e and e^m : (A1) $e = e^m$; (A2) $e \geq e^m$; (A3) $e \leq e^m$.²⁷ Notice right away that, by Proposition 1, constraint A2 is not binding and constraint A3 is equivalent to constraint A1. Therefore, we can neglect A2 and A3 without loss of generality.

(B) Rules that constrain a government to impose *no* regulation on imported products or domestic products: (B1) $e^m = 1$; (B2) $e = 1$.²⁸

²⁶Note that even a constraint on the relative standard e^m/e , for example $e^m/e \geq r$, is not unit-free. (We owe this observation to Gene Grossman.) This point can be illustrated with an example. Suppose e (e^m) is the minimum temperature at which domestic (imported) milk must be processed, and consider a constraint $e^m/e \geq 2$. This constraint is not unit free, because the value of e^m/e changes depending on whether temperature is measured in C⁰ or F⁰ degrees. (This example is inspired by a real dispute occurred in 1991 between the U.S. and Canada within the U.S.-Canada Free Trade Agreement).

²⁷We can ignore strict inequality constraints, because they are either not binding or they cause non-existence of equilibrium.

²⁸Note that these are indeed ordinal constraints, as they refer to the maximum value of e . Also note that, even if the maximum value of e were different across products, it would still be possible to express the rule in a non-product-specific way as “a government should impose no regulation on...”.

(C) Rules that constrain a government to impose the strictest feasible standard on imported products or domestic products: (C1) $e^m = 0$; (C2) $e = 0$. Further rigid rules can be obtained by combining rules of type (B) and (C) (e.g. $e = e^m = 1$). We will show that ND is the only rigid feasible rule that satisfies a desirable efficiency property.

We assume that a rigid rule constrains the choices of governments unless they agree to waive it ex post. Even if the rule is waived ex post, however, it still influences the final outcome by affecting the status quo in the ex post negotiation.²⁹ We could allow governments to commit not to waive the rule – but one can easily show that they can never gain from doing this.

One possible approach at this point would be to look for the rigid rule that maximizes ex-ante welfare. This approach however would be questionable on two counts: first, the solution would depend in a delicate way on the details of the probability distribution $\Phi(\theta)$; second, this approach does not work if one takes the view that governments cannot foresee all the potential products. In light of these considerations, we take a “robustness” approach instead: we ask whether there exists a rigid rule that increases ex-ante welfare for any distribution $\Phi(\theta)$. The next result answers in the affirmative:

Proposition 5. *Suppose DSP is not available. Imposing the ND rule is not sufficient to implement the first best outcome, but it improves ex-ante efficiency for any $\Phi(\theta)$. Furthermore, ND is the only rule that satisfies this property.*

The first part of the result follows from two observations. First, the ND rule constrains the choice of standards by the importing country, thus lowering its welfare in case of disagreement. This shift in the disagreement point increases total welfare, provided the shift does not go too far. Second, the first-best standards satisfy the ND constraint, and this ensures that the shift in the disagreement point never goes too far. The second part of the result follows from the observation that, for any alternative rule R , one can find a potential product θ such that R would push the disagreement point too far.

²⁹An alternative assumption would be that a rule is enforced only if a country invokes it (e.g., it files a violation complaint). In this case, the rule affects the outcome by providing an outside option. Our qualitative results would still be valid.

Figure 3 illustrates the first part of the result in more detail. Figure 3a depicts the situation in (e_θ, e_θ^m) space. Point N is the unconstrained Nash equilibrium. The importing country's iso-welfare functions are centered around point N. The ND-constrained Nash equilibrium is point ND. This is the point where the diagonal is tangent with an importer's iso-welfare function. The standards that maximize $\Omega_\theta^m + \Omega_\theta^x$ are given by point W. This point is on the diagonal because the joint-welfare maximizing standards are nondiscriminatory. Drawn in bold is the efficiency locus (where the two countries' iso-welfare curves are tangent), which runs through points W and N. The outcome of the bargaining when the threat point is ND is a point like B(ND), which lies on the efficiency locus and Pareto-improves over point ND. Figure 3b depicts the situation in $(\Omega_\theta^x, \Omega_\theta^m)$ space. The meaning of the labels is the same as in Figure 3a. Note that the constrained Pareto frontier when the ND rule is imposed is everywhere inside the original frontier, except at point W where the two frontiers are tangent. Points N and ND are respectively the peaks of the two frontiers. Point B(ND) lies on the original frontier between N and W, hence it involves a higher joint welfare than point N.³⁰

One remark is in order about the robustness of this result. We have proved the result exploiting the symmetry of $\Omega_\theta^m + \Omega_\theta^x$ with respect to domestic and import standards, which in turn is due to the assumption of symmetric technology and factor endowments. But intuitively, even if supply conditions are asymmetric, the ND rule is likely to be beneficial for a large class of potential products, because it shifts the ex-post balance of power in the right direction, that is from the importing country to the exporting country. Therefore, we think we have established a presumption that the ND rule is likely to generate an ex-ante Pareto improvement under quite general conditions.

³⁰The reader may question the realism of one of the model's predictions, namely that in equilibrium the negotiated standards do not conform to the ND rule. This might seem at odds with the GATT-WTO's strict national treatment principle. It is not clear however that in reality standards always conform to national treatment. First, there is plenty of anecdotal evidence that the national treatment rule is easy to get around in practice. This is even more so if both trading partners agree not to enforce the rule. Second, the GATT-WTO allows a number of exceptions to the national-treatment rule (e.g. GATT's Article XX); if two trading partners want to get around the rule, there is little that prevents them from appealing to some of these exceptions.

3.3. The optimal ex-ante agreement

In this section we allow for a dispute settlement mechanism as well as a rigid ND rule in the ex ante agreement. In the previous section we argued that the ND rule is the only rigid rule that can increase ex ante welfare for any probability distribution over future products. Based on this consideration, we feel justified in focusing on the ND rule as the only candidate rigid rule.

Since we assumed that at t_0 countries are symmetric with respect to both existing and future products, we can focus separately on two parts of the ex ante agreement: the part that concerns existing products and the part that concerns future products. For the former part, the same analysis as in section 2.2 applies. For the latter part, we have only four options to compare: (i) a combination of DSP and ND rule, (ii) ND rule alone, (iii) DSP alone, and (iv) neither of the two. In what follows, with a slight abuse of terminology we will use the word “optimal” to indicate the best of these four alternatives from the point of view of ex-ante welfare.

The ND rule and the DSP influence the ex post bargaining outcome in two different ways: the ND rule changes the *status quo* for the negotiation, while the DSP influences the outcome by offering an *outside option* to the bargainers. Let us study the interaction between these two effects in greater detail.

As before, we look at the limit (subgame perfect) equilibrium outcome of sequential bargaining as the interval between offers becomes small. We have two possible disagreement points: the unconstrained Nash equilibrium N , where the payoff of the importing country is maximized, and the Nash equilibrium ND that would obtain under the ND rule. We let D denote a generic disagreement point. In the presence of the DSP, we have a game of bargaining with outside options, where the status quo is D (a one-shot Nash equilibrium) and the outside option for each player is given by the expected payoffs from the DSP decision. In what follows we refer to this pair of expected payoffs simply as “DSP point”. We let $B(D, DSP)$ denote the corresponding bargaining outcome, whereas $B(D)$ is the bargaining solution with disagreement point D and no outside option. If the DSP point is dominated by point $B(D)$, then the DSP is irrelevant and the (limit) outcome of bargaining is still $B(D)$. If, on the other hand, one of

the players expects a higher utility from DSP than at $B(D)$, then the equilibrium is the point on the frontier where this player gets the DSP expected utility.

The DSP point can fall in one of the following three regions of the set of feasible (Ω^x, Ω^m) pairs (we omit the subscript θ for simplicity). See Figure 4:

Region I: The payoff Ω^x of the exporting country is lower than in $B(ND)$.

Region II: The payoff Ω^x of the exporting country is higher than in $B(ND)$ and $\Omega^x + f(\Omega^x)$ is higher than total welfare in $B(ND)$, where $(\Omega^x, f(\Omega^x))$ is the point on the Pareto frontier with ascissa Ω^x .

Region III: The payoff Ω^x of the exporting country is higher than in $B(ND)$ and $\Omega^x + f(\Omega^x)$ is lower than total welfare in $B(ND)$.

First notice that, if σ_ε^2 is low, the DSP point falls in region *II* with high probability, and in this case the bargaining outcome is more efficient than in the absence of the DSP (whether or not a ND rule is in place), hence instituting a DSP is desirable from an ex ante perspective. Next notice that, if the DSP point falls in region *I* (i.e., $\Omega_{DSP}^x < \Omega_{B(ND)}^x$), having a ND rule is strictly beneficial. Finally note that, if the DSP point falls in region *III*, the presence of a DSP lowers aggregate welfare. However, the probability of the DSP point falling in region *III* is low if σ_ε^2 is low. The next proposition draws the conclusion:

Proposition 6. *If σ_ε^2 is sufficiently low and there is a positive probability that $\Omega_{DSP}^x < \Omega_{B(ND)}^x$, then the optimal ex-ante agreement has three components: (a) a set of clauses that specify standards for existing products; (b) a rigid nondiscrimination clause, and (c) a clause that institutes a dispute settlement procedure.*

The above result states that, under some conditions on the quality of the DSP's information, the DSP and the ND rule are complementary tools for coping with the incompleteness of agreements. The sufficient condition stated in the proposition is somewhat restrictive, but intuitively the result should hold quite generally. The critical level of σ_ε^2 below which the DSP is ex-ante beneficial is likely to be high, and possibly infinite, as we discussed in section 3.1.

On the other hand, the condition for the ND rule to be beneficial is that there is some state θ for which the expected DSP point falls in region I ; this is likely to be the case if there is some state θ for which the conditional noise $\sigma_{\varepsilon}^2(\theta)$ is sufficiently high, a rather weak condition.

4. Extensions

In this section we discuss the robustness of the results to some extensions of the model. We start with the assumption that there is no demand in the exporting country. Let us focus on a new good. If there is demand in both countries, there are three relevant standards to consider for this good, two standards chosen by the importing country and one by the exporting country. It is easy to see that in this case the noncooperative payoff will generally lie below the bargaining frontier, and the bargaining outcome will be less asymmetric. In other words, the imbalance of power in favor of the importer will be less pronounced. However, ex-post bargaining will still generically lead to ex-ante inefficient outcomes, and a DSP will still have a potential efficiency role to play. As for the national treatment rule, our intuition is that it should still be beneficial under fairly general conditions, since its key effect is that of shifting bargaining power from the importing country to the exporting country, and this is likely to increase ex-ante efficiency.

Next we discuss our assumption that countries are ex-ante symmetric. As a preliminary remark, we believe that the assumption of “veil of ignorance” is quite reasonable for the purposes of studying the optimal design of international institutions. Problems of institutional design are not dissimilar in this respect from problems of constitutional design, and the assumption of veil of ignorance is widely used in constitutional design theory. Setting up an institution such as a DSP is a long term commitment, so if the contracting parties view the future as very important and very uncertain, the assumption of a symmetric future is not unreasonable. Further, even if the parties’ present conditions are asymmetric, if they are sufficiently patient the considerations about the future will dominate those about the present, and their ex-ante bargaining powers will be similar. Having said this, the presence of ex-ante asymmetries need not invalidate the qualitative results of the model. In what follows we discuss this point in more detail.

The first remark is that a DSP continues to have an efficiency role to play for virtually any pattern of ex-ante asymmetries. This is not hard to see formally, with a minor notation change in our basic model. Suppose θ incorporates not only the characteristics of the new product, but also the identity of the country that exports it. Let Ω_θ and Ω_θ^* denote respectively the Home and Foreign (incremental) payoffs. Each state θ is characterized by a bargaining frontier in $(\Omega_\theta, \Omega_\theta^*)$ space, which we denote by Γ^θ , and a bargaining outcome point B^θ . Let γ^θ denote the absolute value of the slope of Γ^θ at point B^θ . Absent DSP, if the frontiers Γ^θ are concave, γ^θ will vary across states of the world (generically). But this is ex-ante inefficient, because (it can be shown that) ex-ante Pareto efficiency requires $\gamma^\theta = \gamma$ for all θ .³¹ As for the role of the DSP, it is clear that if σ_ε^2 is sufficiently small the DSP can improve ex-ante efficiency. Suppose for the sake of argument that σ_ε^2 is zero, and that the DSP, if invoked, maximizes a weighted aggregate welfare function of the form $\gamma\Omega_\theta + \Omega_\theta^*$. This will implement the point on the bargaining frontier with slope $-\gamma$ for each θ . It follows that, if γ is chosen appropriately, the DSP can generate an ex-ante Pareto improvement.

Our second remark is that it may be reasonable to assume that the ex-ante agreement maximizes the countries' joint welfare even if countries are ex-ante asymmetric. This will be the case if standards are negotiated in the context of a comprehensive trade negotiation where countries discuss many issues simultaneously, so that concessions in policy areas other than standards can be used as a compensation tool. In this case utility is effectively transferable, even if monetary transfers are unavailable, and as a consequence the ex ante agreement will maximize aggregate welfare. Further, if there are important costs of renegotiation (as we assumed in our basic model), utility is *not* transferable in the ex-post negotiations, hence our results about DSP and the ND rule will still be valid. (Note that in this case the argument about the DSP in the previous paragraph would apply, and the ex-ante agreement would set $\gamma = 1$.)³²

³¹Intuitively, if there are two states θ' and θ'' such that $\gamma^{\theta'} > \gamma^{\theta''}$, one can make a Pareto improvement by moving the payoff point along the bargaining frontier upwards in state θ' and downwards in state θ'' .

³²Even if the ex-ante agreement does not maximize joint welfare, but some weighted sum of the countries' welfare levels, the qualitative conclusions of the model need not be upset. The DSP would still have an efficiency role to play, as we argued in the text. As for the ND rule, this will still be desirable if the two countries's weights are not too different. The reason is that this is an all-or-nothing rigid rule. Intuitively, if ex ante asymmetries are small it is better to have the ND rule than nothing.

Next we discuss the robustness of our results to political-economy considerations. We distinguish two different ways of formalizing political-economy forces. One is to suppose that governments maximize an adjusted welfare function that assigns extra weight to the well-being of politically powerful groups (owners of certain specific factors), as for example in Grossman and Helpman (1994). We think this model is a good representation of the process that determines short-run policies, but may not apply equally well to the determination of long-term policy regimes, such as trade institutions, because industrial lobbies tend to have a short-run nature (since factors are "specific" only in the short run). An alternative political-economy framework, that might be a better fit for a setting of institutional design as in the present paper, is one where governments maximize welfare at the stage of designing the institution, but they anticipate that *ex post* they will be subject to short-run political pressures.³³ What would happen to our results if we extended the model in this direction is not obvious, but intuitively the main results should come out strengthened. First, the presence of ex-post political pressures should increase the potential gains from a DSP, because these pressures tend to increase the divergence between ex-post bargaining outcomes and ex ante efficient outcomes. Second, the ND rule should be even more valuable in this case, because ex post political pressures (if they take the form of an extra weight assigned to import-competing firms) will make noncooperative standards even more discriminatory, hence the ND rule will have more bite.

Thus far we have ignored the possibility of using trade taxes. We start by considering the implications of tariffs in our model. If governments can choose standards and tariffs, Proposition 1 may change. In particular, it is possible that $e^m \geq e$ in the noncooperative equilibrium. The reason is that the importing country can use the tariff to discriminate against imports, and this changes qualitatively the nature of the tradeoff in choosing e and e^m . However, the main results about trade agreements will remain unchanged, with the only amendment that the ex-ante agreement will specify also a rigid zero tariff for all present and future products.

We come back now to the question of what can be accomplished with a market-access mechanism *à la* Bagwell and Staiger (2001). As we noted in the introduction, this type of mechanism cannot solve the problems related to contractual incompleteness. Here we discuss

³³See for example Maggi and Rodriguez-Clare (1998) for a theory of trade agreements based on a similar approach.

whether it can work in a complete-contracting setting (e.g. with no future products). First note that, because the international externality associated with standards does not travel simply through terms of trade (changing a standard affects the exporting country not only through the product price but also directly), a rule that requires preservation of import volumes will not work. We can think of only one mechanism that can implement efficient outcomes: after tariffs have been negotiated, a government is free to change its standards subject to the constraint that foreign employment is kept constant.³⁴ However, the efficiency of this mechanism is an artifact of the simplifying assumptions of our model. If we had a more general production function (say $y = F(L, K, e)$, or if there were more than one mobile factor), the rule would not work, because there would be no variable that is linked one-to-one with foreign welfare.

Our qualitative results would be affected if tariffs *and* export subsidies were allowed. The reason is that, if these two tools are used jointly, lump-sum transfers can indirectly be effected between countries. As a consequence, utility would be transferable in the ex-post bargaining game, and the outcome would be ex-ante efficient. However, we think this does not invalidate the points of the paper. First, export subsidies are banned by the WTO, so this argument has only theoretical relevance. Our model is not rich enough to explain the WTO ban on export subsidies, but a model cannot explain everything. Second, even if export subsidies are available, if utility is nonlinear in all goods or if governments are risk averse, then utility remains nontransferable and our results will remain valid.³⁵

³⁴The reason this mechanism can implement efficient outcomes is that, in our model, the level of employment in the exporting country is linked one-to-one with the exporting country's welfare. For this reason, if the importing country is constrained to keep foreign employment constant, it is effectively constrained to keep foreign welfare constant, and this is sufficient to implement efficient outcomes. We are grateful to Bob Staiger for making this point to us.

³⁵Our results could also potentially be affected if tariff-rate quotas were allowed. Suppose that, when the new product appears, tariff-rate quotas are in place on old products. Then it may be possible to effect a lump sum transfer by varying the within-quota tariff rate on some old product. This argument however is subject to the same considerations we made for export subsidies (recall that quotas are banned by the WTO), plus an additional one: if renegotiation costs are high then governments will not want to revise the trade policies on existing products everytime a new product arises.

5. Conclusion

In this section we expand the discussion that we started in the introduction on the key articles of the TBT and SPM agreements in the WTO. We already mentioned article 2.2 of TBT, that states the general principle that standards should not result in disguised restrictions to trade. This principle is further articulated in various articles of the SPM agreement. We discuss some of these in turn.

Article 5.4 of the SPM agreement states that "*Members should, when determining the appropriate level of sanitary or phytosanitary protection, take into account the objective of minimizing negative trade effects.*" This article explicitly requires that the standard-setting country take into account the effects of its standards on the volume of trade. Article 5.6 states that "*...Members shall ensure that such measures are not more trade-restrictive than required to achieve their appropriate level of sanitary or phytosanitary protection...*" This principle is more narrow, but probably easier to apply, than article 5.4: it states that a standard is illegitimate if there exists another measure that achieves the same level of protection but is less trade-restrictive. Article 5.6 played a key part in the dispute between the US and Japan on quarantine regulations for fruit. The WTO's appellate body, in its February 1999 report, found against Japan's policies mostly on the basis of this article.³⁶ Article 2.2 states that "*Members shall ensure that any sanitary or phytosanitary measure is applied only to the extent necessary to protect human, animal or plant life, or health, is based on scientific principles and is not maintained without sufficient scientific evidence...*" This article played a key role in the dispute between the US and the EU on beef hormones. Even though international trade is not explicitly mentioned in this article, it is legitimate to argue that the notion of "sufficient" evidence refers implicitly to a tradeoff between domestic welfare objectives and trade restrictions. After all, if trade considerations were not relevant for this article, its presence would be puzzling: why would the WTO be interested in teaching a government how to pursue its own domestic objectives?

³⁶A similar principle is stated in article 2.2 of TBT: "*...technical regulations shall not be more trade-restrictive than necessary to fulfil a legitimate objective, taking account of the risks non-fulfilment would create. Such legitimate objectives are, inter alia: national security requirements; the prevention of deceptive practices; protection of human health and safety, animal or plant life or health, or the environment.*" Note that the "legitimate objectives" do not include domestic producers' interests, so this principle cannot be construed as saying "standards should not be more trade-restrictive than necessary to achieve maximization of domestic welfare." As a consequence, this rule will impose a binding constraint even on a welfare-maximizing government.

The common theme across these articles is that an importing country must give consideration to the trade-restrictive effects of its standards. But since trade restrictions directly hurt foreign exporters, these articles require that foreign welfare be taken into account when choosing standards. Even though it is not clear from the literal reading of these articles how much weight the importing country is supposed to assign to foreign exporters, this reading is consistent with the idea that the standard-setting country should give foreign exporters a similar weight as domestic producers, at least as a general benchmark principle. As we mentioned in the introduction, this is the interpretation given by a number of economists to the general no-disguised-protectionism principle.

Overall, we think of the cluster of rules we discussed above as more specific implications of the general no-disguised-protectionism principle. For example, the least-trade-restrictive rule is a necessary condition for standards to be globally efficient. Also the national treatment rule can be viewed as a necessary condition for global efficiency. A natural question then arises: why would the WTO drafters feel the need to spell out these more specific rules, rather than writing a single general principle that standards should be consistent with global welfare maximization? One reason we have in mind is that general principles are more difficult to verify than specific rules: it is very hard to verify welfare functions (let alone the problem of defining exactly what welfare means), while it is easier to verify that, for example, the least-trade-restrictive rule has been violated. The model we presented in this paper is too simple to fully address this question, but we think an extension of our theory in this direction would be an interesting avenue for future research.

6. Appendix

To simplify the notation, in the proofs we focus on the Home country and we omit the index j whenever this causes no confusion. We start by proving two preliminary results invoked in the main text.

Lemma 1. *Equilibrium standards are interior: $0 < e < 1$, $0 < e^m < 1$.*

Proof. Focusing on the Home country, we first show that at a maximum $e > 0$. Suppose $e = 0$ and $e^m = \bar{e} > 0$. Then Home can improve its payoff by setting $e = \bar{e} > 0$ and $e^m = 0$. In this way aggregate supply, price, consumer surplus and negative externality (H) do not change, but producer surplus increases. Also $e = e^m = 0$ cannot be a maximum because a slight increase in e increases consumer surplus more than the negative externality (recall that $A'(0) > 0$, $u'(0) = \infty$ and $H'(0) < \infty$). We can also exclude the extreme values $e = 1$ and $e^m = 1$: Since $A'(1) = 0$ a small decrease in e (e^m) from 1 would have only a second order negative effect on producer and consumer surplus, but it would have a first order effect in decreasing the negative externality. Finally we need to rule out $e^m = 0$. Writing $y = y(p, e)$, $m = y(p, e^m)$ and applying the envelope theorem, we obtain

$$\frac{\partial \Omega}{\partial e^m} = \left[-y(p, e^m) - \widehat{H} \cdot (ey_p(p, e) + e^m y_p(p, e^m)) \right] \frac{\partial p}{\partial e^m} - \widehat{H} \cdot [y(p, e^m) + e^m y_e(p, e^m)],$$

where $\widehat{H} = H'(ey + e^m m) > 0$. Note that

$$y_e(p, e^m) = -A'(e^m) \frac{(F_l)^2}{F_u}$$

Our assumptions $0 < A'(0) < \infty$, $0 < F_l(1, 0) < \infty$ and $-\infty < F_u(1, 0) < 0$ ensure that $0 < y_e(p, 0) < \infty$. Also, note that $y(p, 0) = 0$. Therefore,

$$\left. \frac{\partial \Omega}{\partial e^m} \right|_{e^m=0} = -\widehat{H} e y_p \frac{\partial p}{\partial e^m}$$

We have just shown that $e > 0$ in equilibrium. Also, it is not hard to show that y_p is strictly positive. Let us examine the sign of $\frac{\partial p}{\partial e^m}$ evaluated at $e^m = 0$:

$$\left. \frac{\partial p}{\partial e^m} \right|_{e^m=0} = -\frac{y_e(p, 0)}{y_p(p, e) + y_p(p, 0) - d'(p)}$$

Clearly, also the denominator of this expression is strictly positive. It follows that $\frac{\partial \Omega}{\partial e^m} \Big|_{e^m=0} > 0$, which implies $e^m > 0$ in equilibrium. ■

Proof of Proposition 1

Let (\bar{e}, \bar{e}^m) be the maximizing choice for Home. We only have to show that $\bar{e} > \bar{e}^m$. As argued in the main text, starting from any point on the 45⁰ line, a small change $\Delta e = \Delta e^m > 0$ has a second order effect on the negative externality and consumer surplus and a positive first order effect on Home producer surplus, thus it increases Home welfare; if $\Delta e = \Delta e^m < 0$ we have the opposite effect. This implies that (\bar{e}, \bar{e}^m) is not on the 45⁰ line.

Next we show that if the domestic standard is more strict than the import standard, then Home can increase its welfare simply by switching standards. Note that Home and Foreign have the same production technology and endowments. This implies that the equilibrium price function $p(e, e^m)$ and the function $E(e, e^m) \equiv ey(e, p(e, e^m)) + e^m m(e^m, p(e, e^m))$ are symmetric in their two arguments [see equation (2.1)]. Suppose that $e = e' < e'' = e^m$. By switching standards, the Home government leaves price, consumer surplus and negative externality unchanged: $s(p(e', e'')) = s(p(e'', e'))$, $H(E(e', e'')) = H(E(e'', e'))$. On the other hand, Home producer surplus increases with e : $\pi(pA(e'')) > \pi(pA(e'))$. Therefore, the maximizing choice (\bar{e}, \bar{e}^m) must be below the diagonal. ■

Lemma 2. *Since $\varepsilon^d > 1$, $p(e, e^m)A(e)$ is increasing in e^m .*

Proof. Differentiating the market equilibrium condition we have

$$\frac{\partial p}{\partial e^m} = \frac{A'(e^m)}{A(e^m)} \frac{m - \frac{F_l^*}{pF_u^*}}{\frac{F_l}{p^2 F_u} + \frac{F_l^*}{p^2 F_u^*} + d'(p)},$$

where F_l, F_u are computed at $(1, L(pA(e)))$ and F_l^*, F_u^* are computed at $(1, L(pA(e^m)))$. From

$$\frac{\partial(p(e, e^m)A(e^m))}{\partial e^m} = \frac{\partial p}{\partial e^m} A(e^m) + pA'(e^m)$$

we obtain

$$\begin{aligned} \frac{\partial(p(e, e^m)A(e^m))}{\partial e^m} &= A'(e^m) \frac{m - \frac{F_l^*}{pF_{ll}^*}}{\frac{F_l}{p^2F_{ll}} + \frac{F_l^*}{p^2F_{ll}^*} + d'(p)} + pA'(e^m) \\ &= A'(e^m) \frac{d(p) - y + pd'(p) + \frac{F_l}{pF_{ll}}}{\frac{F_l}{p^2F_{ll}} + \frac{F_l^*}{p^2F_{ll}^*} + d'(p)} = A'(e^m) \frac{(\varepsilon^d - 1)d(p) + y + \left| \frac{F_l}{pF_{ll}} \right|}{\left| \frac{F_l}{p^2F_{ll}} + \frac{F_l^*}{p^2F_{ll}^*} + d'(p) \right|} > 0. \end{aligned}$$

■

Proof of Proposition 5

Let Ω_N^m, Ω_N^x be the payoffs in the unconstrained Nash equilibrium, $\Omega_{ND}^m, \Omega_{ND}^x$ the payoffs in the Nash equilibrium under the ND rule, and Ω_W^m, Ω_W^x the payoffs at the first-best point W . First note that $\Omega_N^m \geq \Omega_{ND}^m \geq \Omega_W^m$. The first inequality holds because the ND rule imposes a constraint on the importer's maximization problem. The second inequality follows from the fact that, given our assumptions, the first best satisfies ND ($e = e^m$); therefore the importer's utility at W cannot be higher than Ω_{ND}^m , otherwise the optimality of the importer's choice would be contradicted. From $\Omega_N^m \geq \Omega_{ND}^m \geq \Omega_W^m$ it follows that the Nash bargaining outcome when the status quo is $(\Omega_{ND}^m, \Omega_{ND}^x)$ lies between point N and point W , and therefore (by concavity of the Pareto frontier) entails a higher ex-ante welfare than the unconstrained Nash equilibrium. That the ND rule is not sufficient to implement the first-best outcome is obvious.

The second part of the result can be shown by exhibiting, for every alternative rule R , a parameter vector θ such that the induced bargaining equilibrium is either on the left of the $B - ND$ or far on the right of W . ■

Proof of Proposition 6

First note that, if σ_ε^2 is small, having a DSP is better than having nothing. Also, Proposition 5 implies that a simple ND rule is better than nothing. Therefore we only have to compare the simple DSP and ND rules with the composite ND+DSP rule. If there is no ND rule and $DSP = (\Omega_{DSP}^x, \Omega_{DSP}^m)$ falls in region I , the bargaining outcome $B(N, DSP)$ will be either N (if Ω_{DSP}^x is lower than in N) or the point $(\Omega_{DSP}^x, f(\Omega_{DSP}^x))$. In both cases the ND+DSP rules turns

out to yield a higher aggregate welfare: since DSP is in region I then $B(ND, DSP) = B(ND)$. If DSP falls in region II or III the bargaining outcome with the DSP outside option is $(\Omega_{DSP}^x, f(\Omega_{DSP}^x))$ whether or not the ND rule is in place. The ND+DSP rule yields a higher aggregate welfare than the simple ND rule in region II and a lower aggregate welfare in region III . We conclude that the ND+DSP rule is at least as good in terms of aggregate welfare as the simple DSP rule, and strictly better if DSP falls in region I . Whether or not the DSP rule is worth adding to the ND rule from an ex ante perspective depends on the probability of a DSP point falling in the “extreme” region III . If this probability is sufficiently low then the ND+DSP rule is optimal. But this is ensured if σ_{ε}^2 is sufficiently low. ■

7. References

Bagwell, K. and Staiger, R., “An Economic Theory of GATT”, *The American Economic Review*, March 1999, 89(1), pp.215-248.

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

Bagwell, K. and Staiger, R., “Multilateral Trade Negotiations, Bilateral Opportunism and the Rules of GATT/WTO,” mimeo, 2001(b).

Baldwin, R., *Nontariff Distortions of International Trade*, Washington, D.C.: Brookings Institution, 1970.

Battaglini, M., “Multiple Referrals and Multidimensional Cheap Talk,” *Econometrica*, forthcoming.

Binmore, K., Osborne, M and Rubinstein, A., “Noncooperative Models of Bargaining,” in (Aumann, R. and Hart, S., Eds.) *Handbook of Game Theory with Economic Applications*, Vol. 1, 1992, Amsterdam: Elsevier.

Copeland, B. and Taylor, S., “International Trade and the Environment: a Framework for Analysis,” NBER Working Paper No 8540, 2001.

Crawford, V., "On Compulsory-Arbitration Schemes", *Journal of Political Economy*, 1979, vol. 87, pp. 131-159.

Dixit, Avinash, *The Making of Economic Policy: A Transaction-Cost Politics Perspective*, 1996, Cambridge: The MIT Press.

Egglestone, K., E.A. Posner and R. Zeckhauser, "Simplicity and Complexity in Contracts," mimeo, 2000.

Engel, E., "Poisoned Grapes, Mad Cows and Protectionism," NBER Working Paper No 6959, 1999.

Ethier, W., "Punishments and Dispute Settlement in Trade Agreements", mimeo, 2001.

Farber, H., "An Analysis of Final-Offer Arbitration," *Journal of Conflict Resolution*, December 1980, 35, pp. 683-705.

Fischer, R. and P. Serra, "Standards and Protection ", *Journal of International Economics*, forthcoming, 2001.

Gibbons, R., "Learning in Equilibrium Models of Arbitration," *The American Economic Review*, December 1988, 78(5), pp.896-912.

Grossman and Helpman, "Protection for Sale", *The American Economic Review*, March 1994, 84(4), 833-850.

Horstmann, I., J. Markusen and J. Robles, "Multi-issue Bargaining and Linked Agendas: Ricardo Revisited or No Pain No Gain", mimeo, 2000.

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

Hungerford, T., "GATT: A Cooperative Equilibrium in a Noncooperative Trading Regime?", *Journal of International Economics*, November 1991, 31(3-4), pp. 357-69.

Jackson, John, *The World Trading System*, 1989, Cambridge: The MIT Press.

Kovenoch, D. and M. Thursby, "GATT, Dispute Settlement and Cooperation", *Economics and Politics*, March 1993, 4(1), pp. 151-70.

Krishna, V. and J. Morgan, "A Model of Expertise," *Quarterly Journal of Economics*, May 2001, 116(2), pp 747-775.

Maggi, G., "The Role of Multilateral Institutions in International Trade Cooperation", *The American Economic Review*, March 1999, 89(1), pp.190-214.

Maggi, G. and A. Rodriguez-Clare, "The Value of Trade Agreements in the Presence of Political Pressures," *Journal of Political Economy*, June 1998, vol. 106(3), pp. 574-601.

Osborne, Martin and Rubinstein, Ariel, *Bargaining and Markets*, 1990, San Diego: Academic Press.

Rubinstein Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, vol 50, 97-109.

Sturm, D., "Product Standards, Trade Disputes and Protectionism," mimeo, 2001.

Sykes, A., *Product Standards for Internationally Integrated Goods Markets*, Washington D.C.: Brookings Institution, 1995.

Figure 1




Figure 2




Figure 3a




Figure 3b




Figure 4

