

OFFSHORING AND THE VALUE OF TRADE AGREEMENTS*

Pol Antràs

Harvard University and NBER

Robert W. Staiger

Stanford University and NBER

PRELIMINARY VERSION – COMMENTS WELCOME

September 25, 2007

Abstract

We study the trade policy choices of governments in an environment in which some of the trade flows being taxed or subsidized involve the exchange of customized inputs, and the contracts governing these transactions are incomplete. We show that the second-best policies that emerge in this environment entail free trade in final goods but *not* in intermediate inputs, since import or export subsidies targeted to inputs can alleviate the international hold-up problem. We next show that the Nash equilibrium policy choices of governments do not coincide with internationally efficient choices, and that the Nash policies imply an inefficiently low level of intermediate input trade across countries. The reason is that in our environment trade policy choices serve a dual role: they can enhance investment by suppliers but, because of ex-post bargaining over prices, they can also be used to redistribute profits across countries. The inefficiencies inherent in the Nash policy choices of governments not only result in suboptimal input subsidies, but also in positive distortions in final-good prices, even when countries cannot affect world (untaxed) prices in those goods. As a result, an international trade agreement that brings countries to the efficiency frontier will necessarily increase trade in inputs, but it may require a reduction in final-goods trade. When countries ignore the *direct* impact of their policies on ex-post negotiated international input prices, the resulting policy choices are efficient, and hence a modified terms-of-trade interpretation of the purpose of trade agreements can be offered, but only when governments maximize real national income. If governments are motivated by political economy (distributional) concerns, the purpose of a trade agreement becomes more complex, and cannot be reduced to solving a simple terms-of-trade problem.

*We thank Kyle Bagwell and Alan Sykes for helpful discussions. Staiger gratefully acknowledges financial support from the NSF (SES-0518802).

1 Introduction

International trade in intermediate inputs is a dominant feature of the world economy. Using OECD input-output tables, Ramanarayanan (2006) concludes that in the late 1990s imports of intermediate goods comprised about forty to sixty percent of total merchandise imports for a large number of OECD countries. Similarly, a thorough examination of highly disaggregated trade data led Yeats (2001) to conclude that intermediate input trade accounted for roughly thirty percent of world trade in manufacturing goods in 1995. Furthermore, several authors have noted that the share of intermediate inputs in world trade appears to have increased significantly in recent years.¹

Recent developments in international trade theory have attempted to bridge the apparent gap between the characteristics of international trade in the data and the standard representation of these trade flows in terms of final goods in neoclassical trade theory. One branch of this new literature has focused on incorporating input trade in otherwise standard models with perfectly competitive markets and frictionless contracting.² Another branch of the literature has stressed that modelling offshoring as simply an increase in the fragmentation of production across countries misses important aspects of the characteristics of intermediate input trade.³ First, intermediate inputs tend to be much more customized to their intended buyers than final goods, and hence, input trade embodies a disproportionate amount of relationship-specific investments. Second, because contracts involving international transactions are especially hard to enforce, the cross-border exchange of specialized intermediate inputs cannot generally be governed by the same contractual safeguards that usually accompany similar exchanges occurring within borders. A third distinguishing feature of international offshoring is that it is often associated with the costly search for suitable suppliers that can provide the required inputs.

The purpose of this paper is to study the Nash equilibrium and internationally efficient trade policy choices of governments in an environment in which some of the trade flows being taxed or subsidized involve the exchange of customized inputs, contracts governing these transactions are incomplete, and the matching between final-good producers and suppliers may involve search frictions. By identifying circumstances under which the Nash and internationally efficient policies do not correspond to one another in this environment, we thereby seek to understand the value of trade agreements in the presence of offshoring.

The first result of the paper is that the second-best policies that emerge in this environment entail free trade in final goods but *not* in intermediate inputs. Intuitively, the combination of

¹See, for instance, Feenstra and Hanson (1996b), Feenstra (1998), Campa and Goldberg (1997), Hummels, Ishii and Yi (2001) and Borja and Zeile (2004).

²See for instance the work of Feenstra and Hanson (1996a), Jones (2000), Deardorff (2001), Antràs, Garicano and Rossi-Hansberg (2006), or Grossman and Rossi-Hansberg (2006).

³Theoretical developments include the work of McLaren (2000), Grossman and Helpman (2002, 2005), Antràs (2003, 2005), and Antràs and Helpman (2004). The empirical work of Feenstra and Hanson (2005), Levchenko (2007) and Nunn (2007) substantiate the empirical relevance of these non-standard features of offshoring.

relationship-specific investments and incomplete contracts results in a hold-up problem that leads to an inefficiently low volume of input trade across countries, and import or export subsidies targeted to these inputs can help bring countries closer to the efficiency frontier. In fact, we show that by choosing these subsidies appropriately, the international hold-up problem can be fully resolved and the efficiency frontier or first-best can be attained.

We next tackle the question of whether the Nash equilibrium policy choices of governments will coincide with the internationally efficient policies. We answer this question in the negative, and we show that the Nash equilibrium in trade policies always involves an inefficiently low level of intermediate input trade across countries. The broad intuition for the result is simple. Trade policy serves a dual role in a world with incomplete contracting in vertical relationships. On the one hand, as mentioned above, *subsidies* to the exchange of intermediate inputs can serve as a substitute for more standard contractual safeguards available in domestic transactions and can thus increase the volume of input trade toward its efficient level. On the other hand, in a world where input prices are not fixed in an enforceable initial (or ex-ante) contract, we show that trade *taxes* can also be used to redistribute surplus across countries. For instance, although an export tax in the input-producing country may reduce the incentives to invest of suppliers, in their ex-post bargaining with their buyers, these suppliers will be able to pass part of the cost of the tax on to final-good producers abroad. There is hence a basic tension that each government faces in its unilateral trade policy choices between correcting the hold-up problem and capturing surplus from its trading partner, and this tension prevents governments from making internationally efficient policy choices in the absence of a trade agreement (i.e., in the Nash equilibrium).

In order to illustrate these forces in the simplest possible way, we develop a model with two countries and a single taxable homogenous final good produced with a customized input, which can also be taxed. We assume that the two countries are small in the sense that their choices of trade taxes have no effect on the world (untaxed) price of the final good. For simplicity, we assume that all final-good producers are located in one of the two countries (“Home”), while all input producers are located in the other country (“Foreign”).

In the absence of frictions in input markets, welfare-maximizing governments would have no incentive to tax final-good trade, and as mentioned above, second-best trade policies would also involve free trade in the final good. We show, however, that in the Nash equilibrium in trade policies, the input-importing country will generally have an incentive to tax or subsidize not only the intermediate input, but also the final good. This follows from the fact that by affecting the surplus and outside options of both producers, trade taxes on the final good can also serve a surplus-shifting role. In particular, by imposing an export tax or import subsidy on the final good, the Home government is able to reduce the domestic price of the final-good, while not bearing the full cost of this policy.

With these findings we therefore identify two novel purposes for trade agreements. First, to the extent that international input trade requires relationship-specific investments between domestic and foreign firms that cannot be contracted over on an ex-ante basis, the volume of input trade

will be too low from the perspective of international efficiency when governments set trade policy unilaterally: by expanding input trade volume a trade agreement can therefore move countries toward the international efficiency frontier. And second, in the presence of such relationship-specific investments in inputs and the bargaining over price that is implied, governments face incentives to distort final-good trade as well, as a means of extracting bargaining surplus from foreign firms: by eliminating these additional policy distortions as well, a trade agreement can help countries achieve the international efficiency frontier. An interesting feature of this second purpose of trade agreements is that it need not imply that trade volumes in final goods should be increased under the agreement: indeed, we identify circumstances under which a country must agree to policy changes which *restrict* trade volume in the final good in an internationally efficient agreement.

This general reasoning stands in sharp contrast to the existing theories of trade agreements, which cast the problem that a trade agreement exists to correct as either stemming from a terms-of-trade driven Prisoners' Dilemma between the governments of countries that are large in world markets (the terms-of-trade theory), or stemming from an inability of governments to make commitments to the private sector (the commitment theory).⁴ We make assumptions that preclude a commitment role for trade agreements in this paper. On the other hand, the role that we identify here is related to the terms-of-trade theory, in that there and here the purpose of a trade agreement is to induce those governments that have the ability to affect foreign exporter (international) prices with their trade policy choices to behave in an internationally efficient way. But our theory exhibits some key differences.

First, foreign exporter prices are determined by competitive (or sometimes non-competitive) *market* conditions in the terms-of-trade theory, while here (the relevant) foreign exporter prices are determined as a result of ex-post bilateral producer-supplier *bargaining*. This leads to distinct predictions across the theories concerning the determinants of the degree of trade liberalization required to bring countries to the international efficiency frontier. According to the terms-of-trade theory, it is the Nash foreign export supply elasticities faced by each country that predict the inefficiencies in its Nash tariff choices, and therefore that predict the changes in its policies necessary to reach the international efficiency frontier (see, for example, Bagwell and Staiger, 1999, 2006). In the theory we develop here it is the details of the ex-post bilateral producer-supplier bargaining that predict the inefficiencies in each country's Nash tariff choices: for example, we show that the input-producing Foreign country will have an incentive to use an export tax to appropriate surplus from the Home country only when its suppliers have sufficiently weak bargaining power relative to final-good producers at Home. One implication of these differences is that the terms-of-trade theory predicts that the fundamental task of a trade agreement is always to increase trade volumes from their Nash levels, while as we have observed above it is possible under the theory we develop here that a trade agreement should lead to *reductions* in some (i.e., the final good) trade volumes relative to their Nash levels.

Second, according to the terms-of-trade theory there is a single international problem that

⁴See Bagwell and Staiger (2002, Chapter 2) for a description of the major theories of trade agreements.

a trade agreement must solve, and that is terms-of-trade manipulation, while here there is also a second international problem that a trade agreement must address, namely, the international hold-up problem. According to the terms-of-trade theory, if governments ignored the effect of their policies on world prices, the resulting tariffs would coincide with efficient choices. This feature is not shared by our model, because with relationship-specific investments and incomplete contracting, active manipulation of the terms of trade also serves the valuable role of enhancing the incentives of suppliers to invest. In fact, as long as governments seek to maximize aggregate real national income with their trade policies, the role of an international trade agreement in our model is precisely to nullify the “bad” terms-of-trade manipulation, while *maintaining* the “good” terms-of-trade manipulation inherent in the second-best subsidies to input trade.

A final difference between our theory and the terms-of-trade theory arises when political economy motives are introduced. According to the terms-of-trade theory, preventing governments from manipulating the terms of trade with their trade policy choices is *the* task of a trade agreement, regardless of whether the member governments are concerned only with maximizing aggregate real national income or rather have distributional (e.g., political) concerns as well (see Bagwell and Staiger, 1999). In our model, as we have indicated above, the role of a trade agreement can be interpreted as preventing governments from engaging in “bad” terms-of-trade manipulation provided governments seek to maximize aggregate real national income. But contrary to the terms-of-trade theory, we show that when governments have political economy motives the task of a trade agreement according to the theory we develop here becomes more complicated, and cannot be reduced to solving a simple terms-of-trade problem.

To our knowledge, this is the first paper to consider *endogenous* trade policy choices in a model with intermediate input trade in an environment with relationship-specific investments and incomplete contracting. Ornelas and Turner (forthcoming) develop a model in which import tariffs on intermediate inputs are shown to aggravate the hold-up problem in international vertical relationships, with the implication that trade liberalization may lead to a larger increase in trade flows than in standard models. Ornelas and Turner do not however study optimal trade policies or the possibility of trade agreements in their framework.⁵ McLaren (1997) studies the desirability of announcing a future trade liberalization in a model where producers incur sunk costs to service foreign markets, but his framework emphasizes commitment problems from which we completely abstract.⁶

The rest of the paper is organized as follows. In section 2, we develop a Benchmark Model that introduces the international hold-up problem and illustrates the valuable role of a trade agreement in a simple framework that allows for lump-sum transfers between producers. In section 3, we consider a more realistic model without lump-sum transfers, but show that it delivers very similar

⁵Similarly, Antràs and Helpman (2004) study the effect of trade frictions on the choice of organizational form of firms contemplating offshoring, but they also treat trade frictions as exogenous.

⁶Yarbrough and Yarbrough (1992) also emphasize commitment problems associated with trade relationships that involve substantial relationship- (or market-) specific investments, but they focus on how these issues affect the choice between unilateral liberalization, bilateral agreements and multilateral agreements.

conclusions. In section 4, we consider a variety of extensions of the model, and we offer some concluding remarks in section 5.

2 The Benchmark Model

We begin this section by describing a benchmark free-trade two-small-country trade model in which final-good producers in the home country import inputs from suppliers in the foreign country. We refer to this model as the *Benchmark Model*. While simple and special along a number of dimensions, the Benchmark Model is meant to highlight the essential features of the basic international hold-up problem which arises under free trade. After presenting the setup and characterizing the free-trade equilibrium, we consider the trade policies for the home (importing) government and the foreign (exporting) government which are unilaterally optimal. We then characterize the Nash policy equilibrium and explore the possible role for international trade agreements.

2.1 Setup

We consider a world of two small countries, Home (H) and Foreign (F), and a large rest-of-world whose only role in the model is to fix the prices at which varieties of a final good 1 are available to H and F on world markets (the direction of trade in good 1 is not specified and is immaterial). Consumer preferences are identical in H and F and given by

$$U^j = c_0^j + u \left(c_{1T}^j + \delta c_{1G}^j \right), \quad (1)$$

where c_i^j is consumption of good $i \in \{0, 1_T, 1_G\}$ in country $j \in \{H, F\}$, where $u' > 0$ and $u'' < 0$, and where $\delta \in (0, 1)$. Good 0, which we take to be the numeraire, is assumed to be costlessly traded and available in sufficient quantities that it is always consumed in positive amounts in both H and F . Good 1 comes in two types: a customized type T and a generic type G . Note that the preferences in (1) are such that consumers are willing to buy both types of good 1 only if the price of the generic relative to that of the customized type is equal to δ . This is analogous to consumers perceiving the two goods as perfect substitutes up to a quality shifter. By choice of units for measuring the quantity of the customized type of good 1, we set its (fixed) price on world markets equal to 1. For now we assume that trade in good 1 is free, so that the price of the customized type of good 1 is equal to 1 everywhere in the world.

Regardless of the type, good 1 is produced with an intermediate input x according to the production function $y(x)$, with $y'(x) > 0$ and $y''(x) < 0$.⁷ The key difference between the two types of good 1 is that the production of a generic good G uses an intermediate input x that is not customized to the producer's needs. We suppose that the home country H is inhabited by a unit measure of producers (or entrepreneurs, or distributors) of the final good 1, while the foreign country F is inhabited by a unit measure of suppliers (or managers, or wholesalers) of the

⁷In order to ensure that the second-order conditions are met, we will later impose additional assumptions on $y(x)$.

intermediate input x . Hence, to produce the final good 1, producers in H must import inputs from suppliers in F . We assume that the marginal cost of input production (generic or customized) in F is constant and, through choice of the units in which inputs are measured, we normalize it to 1. For now, we also assume that trade in x is free.

We next turn to focus on the nature of the bilateral relationship between a final-good producer in H and an input supplier in F , which comprises the essence of the model. We adopt a setting of incomplete contracts between final-good producers and input suppliers. In our Benchmark Model, contractual incompleteness can be rationalized in the following simple way. Following Grossman and Helpman (2002) and Antràs (2003), we assume that, when investing in the supply of (customized or non-customized) x , the supplier can choose between manufacturing a high-quality or a low-quality input, and the latter can be produced at a negligible cost but is useless to final-good producers. The quantity of x is observable to everyone and therefore verifiable by third-parties, but we assume that the quality of x is only observable to the supplier and producer in the particular bilateral relationship, and so quality-contingent contracts are not available. Although parties could still sign a contract specifying a price and a quantity, if they did so, the supplier would always have an incentive to produce the low quality input (at lower cost) and still receive the same contractually stipulated price.⁸

Hence, in this environment, no contracts are signed between suppliers and producers prior to the initial supplier investment decisions. And without an initial contract, the price at which each supplier in F sells its inputs to a producer in H is then decided ex-post (through bargaining) once quality has been chosen. We follow the bulk of the literature in assuming that the bargained price is determined through symmetric Nash bargaining. Because parties have symmetric information at the bargaining stage, ex-post efficiency ensures that low-quality production will never be chosen by an input supplier in equilibrium, and so this dimension of the model can be kept in the background henceforth.

We now describe the structure of the bilateral producer-supplier relationship in detail. We assume that all agents have an ex-ante zero outside option. The sequence of events is as follows:

stage 1. The unit measure of producers in H and suppliers in F are randomly matched, producing a unit measure of matches. Each producer in H makes a take-it-or-leave-it offer of a lump sum transfer between it and its matched supplier. If the supplier rejects the offer, the producer and supplier exit; if the supplier accepts the offer, the producer retains the supplier and provides it with a list of customized input specifications.

stage 2. Each supplier decides on the amount x of customized input to be produced (at marginal cost of 1) – the marginal cost of production for a generic input (not customized to the matched producer’s needs) is the same, so there is no benefit in not customizing the input

⁸There is a large literature proposing a variety of mechanism design resolutions to the hold-up inefficiencies caused by incomplete contracts. These resolutions however generally rely on the ability of parties to commit not to renegotiate an initial contract. Bolton and Dewatripont (2005, Chapter 12) offer an excellent overview of the insights and limitations of this literature.

for the matched producer at this point.

stage 3. Each producer-supplier pair bargains over the price of the intermediate input (we assume symmetric Nash bargaining).

stage 4. A small number (formally, a measure-zero countable infinity) n of the bilateral pairs are exogenously dissolved and randomly rematched in a secondary market. They bargain again, but each supplier has full bargaining power at this point. No further inputs can be produced; the amount produced in *stage 2* is perceived as generic in the secondary market because it was tailored to another producer's specifications with probability one.

stage 5. Each producer in H imports x from its partner-supplier and produces the final good with the acquired x , and payments agreed in *stages 3* and *4* are settled.

This 5-stage game generates the simple hold-up problem that forms the heart of our analysis. A number of features of this setup are worth noting at this point.

First, we assume that lump-sum transfers can be made between producers and suppliers in *stage 1* and that producers in H have all of the bargaining power at this stage. The possibility of lump-sum transfers turns out to be important, as we demonstrate when it is relaxed in the next section, and it is an assumption that may be particularly hard to defend in the international context that we study, where such transfers and the obligations associated with them might be difficult to enforce. But as we confirm below, it is an assumption that serves a useful pedagogical purpose, which is why we begin with it.

Second, the role of *stage 4* is to pin down the outside options of the producer and the supplier should their *stage-3* bargaining break down. In light of the broader structure of our benchmark free-trade model, it is easy to see that the breakup of a single bargaining pair in *stage 3* would result in each member of the pair being rematched with probability 1 with a random partner in *stage 4*, and therefore that *stage 4* implies outside options of 0 and $\delta y(x)$, respectively, for the producer in H and the supplier in F as they engage in Nash bargaining in *stage 3*.⁹ Beyond this, *stage 4* plays no role, and in particular only the customized type of good 1 will be produced with positive measure in equilibrium.

Finally, we note that production-side efficiency requires that the intermediate input is used to produce the customized final good T , and that the customized input is produced at a level x^E which satisfies

$$y'(x^E) = 1, \tag{2}$$

and thereby equates the marginal revenue generated from an additional unit of the input (recall that the price of the customized final good is fixed by world markets and equal to 1 under free trade) with the marginal cost of producing an additional unit of the input (which is constant and normalized to 1).

⁹In section 4, we relax the assumption that suppliers have full bargaining power in the secondary market.

2.2 Free Trade Equilibrium

We now characterize the subgame perfect equilibrium of the 5-stage game described in the previous sub-section. The characterization follows very simply from a few key observations. We consider a representative producer in H and supplier in F that are matched in *stage 1*.

First, if the producer uses the supplier's input to produce the final good in *stage 5*, its revenue is given by $y(x)$. Second, as observed in the previous subsection, the outside options of the producer and the supplier in their *stage-3* Nash bargain are 0 and $\delta y(x)$, respectively, so the sum of their outside options is then equal to $\delta y(x)$. Hence, the quasi-rents over which the producer and supplier bargain in *stage 3* (recall that the cost of producing x is sunk at this point) are $(1 - \delta)y(x)$. Therefore, in the symmetric Nash bargain of *stage 3*, the final-good producer in H obtains $\frac{1}{2}(1 - \delta)y(x)$ while the input supplier in F receives $\frac{1}{2}(1 + \delta)y(x)$.

Next, rolling back to *stage 2*, observe that the input supplier chooses x to maximize $\frac{1}{2}(1 + \delta)y(x) - x$, so the optimal quantity \hat{x} of input satisfies

$$\frac{1}{2}(1 + \delta)y'(\hat{x}) = 1. \quad (3)$$

Given the concavity of $y(x)$, it is clear from a comparison of (3) with (2) that $\hat{x} < x^E$ provided that $\delta < 1$. This is the under-investment associated with the hold-up problem, and it reflects the fact that the producer and supplier bargain over the price of the input after the supplier has already sunk investment in input supply. It is also clear from (3) that \hat{x} is increasing in δ , and hence the hold-up problem is alleviated by increasing the outside option of the supplier. In the limit in which $\delta \rightarrow 1$, full production efficiency is restored.

Now consider *stage 1*. With ex-ante lump-sum transfers and full bargaining power for the final-good producer at this stage, the producer can extract a lump-sum transfer of $T = \frac{1}{2}(1 + \delta)y(\hat{x}) - \hat{x}$, in exchange for retaining the supplier and providing it with a list of customized input specifications. This leaves the input supplier in F with a total payoff of $\pi^F = 0$, indifferent between participating or not, and it leaves the final-good producer in H with a total payoff of

$$\pi^H = \frac{1}{2}(1 - \delta)y(\hat{x}) + T = y(\hat{x}) - \hat{x}.$$

Notice that

$$\frac{\partial \pi^H(\hat{x})}{\partial \delta} = (y'(\hat{x}) - 1) \frac{\partial \hat{x}}{\partial \delta} > 0 = \frac{\partial \pi^F(\hat{x})}{\partial \delta},$$

and hence with ex-ante lump-sum transfers and full bargaining power for the final-good producer in *stage 1*, producers in H bear the entire cost of the underinvestment caused by the hold-up problem, and as a consequence enjoy the entire benefit to the extent that the hold-up problem is alleviated. We summarize this discussion with:

Proposition 1 *In the Benchmark Model, a hold-up problem exists under free trade, leading to insufficient investment in the production of imported foreign inputs ($\hat{x} < x^E$), and producers in H*

bear the entire cost of the productive inefficiency.

We next turn to consider trade intervention as a possible means of alleviating the hold-up problem.

2.3 Import Policy

Two features of our Benchmark Model help to highlight the possibility of a novel role for unilateral trade intervention. First, the assumption that H and F are small in world markets for the final good 1 ensures that neither country has an incentive to impose trade taxes on good 1 for traditional (Johnson, 1953-54) optimal tariff reasons. And second, as we observed in the previous subsection and report in Proposition 1, our assumption that ex-ante lump-sum transfers are available and that the final-good producer has full bargaining power in *stage* 1 ensures that the importer of intermediate inputs in H bears the entire cost of the underinvestment caused by the hold-up problem. Hence, if a novel role for unilateral trade policy intervention arises in our Benchmark Model, it is likely to arise with respect to the trade policy choices of the home government.

We now explore the possible role of trade policy in the home country, by considering its incentives to impose trade taxes. To this end, we let τ_x^H denote the trade tax imposed by H on imports of the input x (positive if an import tariff, negative if an import subsidy) defined in specific terms. And we let τ_1^H denote the trade tax imposed by H on the home country's trade in the final good 1 (positive if an import tariff or export subsidy, negative if an import subsidy or export tax) also defined in specific terms. Observe that the price of the final good 1 in H is now given by $p_1^H = 1 + \tau_1^H$, whereas the price of the input x continues to be determined by Nash bargaining between producers and suppliers (though trade taxes may affect this negotiated price).

How does the introduction of these trade taxes affect the equilibrium characterized in the previous subsection? To explore this question, we introduce the following *stage* 0 which occurs prior to *stage* 1 of the 5-stage game described in subsection 2.1:

stage 0. The home government H selects a trade tax τ_1^H on the final good 1, and a trade tax τ_x^H on the imported input x .

After the home government has selected its tariffs in *stage* 0, the sequence of events is as outlined in subsection 2.1 (with trade taxes collected at the time of importation and production/sales in *stage* 5).¹⁰

¹⁰Implied by the timing of the home government's tariff choices is the assumption that the government can make tariff commitments to the private sector. If the government did not have this ability, then as is well known a separate commitment role for trade agreements might arise (see Bagwell and Staiger, 2002, Chapter 2, for a review of this literature). The particular commitment problems that governments face when trade requires relationship-specific investments are emphasized by Yarbrough and Yarbrough (1992) as providing a reason for trade agreements to exist, and by McLaren (1997) as creating the possibility of perverse negotiating outcomes. Our assumed timing permits us to abstract from the possible commitment role of trade agreements throughout this paper, so that we may focus on other issues.

Consider now how H 's trade policy choices in *stage 0* affect the equilibrium outcome of the game. In their *stage-3* bargaining, if the producer and supplier reach an agreement they stand to obtain a joint payoff of (recalling again that the cost of producing x is sunk at that point)

$$(1 + \tau_1^H)y(x) - \tau_x^H x.$$

A positive import tariff or export subsidy on the final good ($\tau_1^H > 0$) raises the joint surplus of the producer and supplier because it raises the price at which the final good is sold in H . Conversely, a positive import tariff on inputs ($\tau_x^H > 0$) reduces the joint surplus of the producer and supplier because it transfers part of the surplus to the home government.

If the producer and the supplier do not reach an agreement, the final-good producer is left with nothing from its *stage-4* bargain while the input supplier obtains

$$\delta(1 + \tau_1^H)y(x) - \tau_x^H x.$$

These expressions are valid provided they are non-negative, and here and throughout the body of the paper we characterize results for the case where these non-negativity constraints are non-binding. In the Appendix we consider the cases where a non-negativity constraint is binding (so that the associated payoff is zero), and show that our qualitative results carry through for those cases as well.

Notice that the ex-post gains from trade (or quasi-rents) are now given by $(1 - \delta)(1 + \tau_1^H)y(x)$. Hence, a final-good producer in H obtains $\frac{1}{2}(1 - \delta)(1 + \tau_1^H)y(x)$ in the Nash bargain of *stage 3*, while the payoff to the input supplier in F is now equal to $\frac{1}{2}(1 + \delta)(1 + \tau_1^H)y(x) - \tau_x^H x$. The input supplier's choice of x in *stage 2* must then satisfy

$$\frac{1}{2}(1 + \delta)(1 + \tau_1^H)y'(\hat{x}) = 1 + \tau_x^H. \quad (4)$$

It is clear from (4) that \hat{x} is increasing in δ as before, but now it is also increasing in τ_1^H and decreasing in τ_x^H . Intuitively, incomplete contracting leads to rent-sharing between the producer and supplier, and hence the latter's incentives to invest tend to be higher whenever the surplus is higher. We will see in later sections that the positive dependence of \hat{x} on τ_1^H and negative dependence of \hat{x} on τ_x^H hold for a variety of specifications of the game played between the producer and supplier.

At *stage 1*, the final-good producer in H can retain the input supplier in F by suggesting an initial lump-sum transfer of $T = \frac{1}{2}(1 + \delta)(1 + \tau_1^H)y(\hat{x}) - (1 + \tau_x^H)\hat{x}$, which will leave the supplier in F with a payoff of $\pi^F = 0$ and the final-good producer in H with a payoff of

$$\pi^H = (1 + \tau_1^H)y(\hat{x}) - (1 + \tau_x^H)\hat{x}. \quad (5)$$

Now consider the measure of social welfare in H implied by our Benchmark Model: namely, consumer surplus plus profits plus trade tax revenue. Using (1), we have that home demand for good

1 is given by $D_1(p_1^H) \equiv u'^{-1}(p_1^H)$, with consumer surplus then defined as $CS(p_1^H) \equiv \int_{p_1^H}^{\bar{p}} D_1(p) dp$ where \bar{p} is the “choke” price for home demand of good 1. Home-country welfare may then be represented by

$$W^H = CS(p_1^H) + \pi^H + \tau_1^H [D_1(p_1^H) - y(\hat{x})] + \tau_x^H \hat{x}.$$

Using (5), the optimal choice of τ_1^H and τ_x^H will then maximize

$$W^H = CS(p_1^H) + \tau_1^H D_1(p_1^H) + y(\hat{x}) - \hat{x},$$

which implies that the chosen home-country tariffs satisfy¹¹

$$\begin{aligned} \frac{\partial W^H}{\partial \tau_1^H} &= \tau_1^H \frac{\partial D_1}{\partial p_1^H} + [y'(\hat{x}) - 1] \frac{\partial \hat{x}}{\partial \tau_1^H} = 0, \text{ and} \\ \frac{\partial W^H}{\partial \tau_x^H} &= [y'(\hat{x}) - 1] \frac{\partial \hat{x}}{\partial \tau_x^H} = 0. \end{aligned} \tag{6}$$

The first-order conditions in (6) which define the unilaterally optimal choice of τ_1^H and τ_x^H are instructive. Recalling from (4) that $\partial \hat{x} / \partial \tau_1^H > 0$, it is clear from (6) that the optimal choice of τ_1^H is strictly positive, provided that $[y'(\hat{x}) - 1] > 0$ which by (2) implies that $\hat{x} < x^E$: this suggests that an import tariff or export subsidy on trade in the final good 1 could raise welfare in H , by increasing \hat{x} toward x^E and thereby helping to ameliorate the hold-up problem at the cost of lost consumer surplus. However, recalling from (4) that $\partial \hat{x} / \partial \tau_x^H < 0$, it is clear from (6) that the optimal choice of τ_x^H must ensure that $[y'(\hat{x}) - 1] = 0$, thereby achieving productive efficiency: there is no associated loss in consumer surplus when the tariff on imported inputs τ_x^H is used to increase \hat{x} , and the optimal choice of τ_x^H therefore solves completely the hold-up problem and achieves productive efficiency. This in turn leaves no reason for home-country intervention with regard to trade in the final good 1. Hence, the optimal choice of τ_1^H , which we denote by $\hat{\tau}_1^H$, is $\hat{\tau}_1^H = 0$.

What about the optimal level of the tariff on imports of x , which we denote by $\hat{\tau}_x^H$? With $\hat{\tau}_1^H = 0$, it follows from (4) that $\tau_x^H \geq 0$ implies $y'(\hat{x}) > 1$, and hence a tax on imported inputs is never optimal according to (6). On the other hand for sufficiently low $\tau_x^H < 0$, (4) implies that we must have $y'(\hat{x}) < 1$, and thus according to (6) a sufficiently high subsidy to imported inputs is not optimal either. In sum, there is a unique interior optimal subsidy on the importation of inputs, and it is given by:

$$\hat{\tau}_x^H = \frac{1}{2} (1 + \delta) - 1 = -\frac{1}{2} (1 - \delta) < 0.$$

Notice also that the trade policies $\hat{\tau}_1^H$ and $\hat{\tau}_x^H$ are not only optimal from the point of view of the home country but, as they solve the hold-up problem and achieve productive efficiency without introducing any other distortions into the world economy, they are also internationally efficient. We may thus state:

¹¹It is easily checked that second-order conditions are satisfied (see Appendix).

Proposition 2 *In the Benchmark Model, when only H intervenes with trade policy, its unilaterally optimal policy choices maintain free trade in the final good ($\hat{\tau}_1^H = 0$) and subsidize importation of the input ($\hat{\tau}_x^H = -\frac{1}{2}(1 - \delta)$) so as to solve the hold-up problem and achieve productive efficiency. Furthermore, these policies are internationally efficient.*

The intuition for Proposition 2 is simple. The hold-up problem between producers in H and suppliers in F results in a level of imported inputs which is inefficiently low. Trade intervention which raises the home-country price of the final good can increase the volume of imported inputs, but at a cost of reduced home-country consumer surplus. A subsidy to imported inputs does not reduce consumer surplus, but it nevertheless succeeds in increasing the volume of imported inputs by increasing the surplus over which the parties negotiate in the ex-post (*stage-3*) bargain while at the same time transferring bargaining power to the foreign supplier. This last point can be seen by noting that with $\tau_1^H = 0$, when a subsidy is offered on the home-country importation of inputs ($\tau_x^H < 0$), the home producer receives $[\frac{1}{2}(1 - \delta)y(x)]$ as under free trade, while the foreign supplier now receives $[\frac{1}{2}(1 + \delta)y(x) - \tau_x^H x]$ – as opposed to $\frac{1}{2}(1 + \delta)y(x)$ under free trade – and therefore enjoys more surplus, owing to its enhanced outside option. In effect, within our Benchmark Model the foreign input supplier enjoys the benefit of the home import subsidy whether or not it stays in the bilateral relationship with the home producer for which its input is customized, while the same is not true for the home producer, and this difference provides the input supplier with more power at the bargaining table, and therefore greater incentive to invest in x .¹² As a consequence, a subsidy to imported inputs targets just the distorted margin, and in analogy with the targeting principle (Bhagwati and Ramaswami, 1963, Johnson, 1965) is hence the first-best method of addressing the problem. And given that H enjoys all of the surplus from the producer-supplier relationships and therefore (as Proposition 1 emphasizes) bears all of the cost of the hold-up-induced inefficiency, the government in H has the incentive to utilize an import subsidy to achieve the first best.¹³

We have thus identified a novel role for unilateral trade policy intervention, namely, as a means of addressing the international hold-up problem that arises when international trade requires relationship-specific investments between domestic producers and their foreign suppliers. In the Benchmark Model, this is the sole motivation for the home government’s unilateral use of trade policy, and when only the home government is policy active the Benchmark Model leads to a striking result: the inefficiencies associated with international hold-up are solved by the home government’s unilateral trade policy intervention, and as a consequence there is no reason for a trade agreement. A remaining question is whether this result survives once the foreign government’s incentives to intervene are taken into account. We consider these incentives in the next subsection, and then return to evaluate the potential role of trade agreements in the Benchmark Model once we have characterized the Nash policies of the home and foreign governments.

¹² As shown in the Appendix, even when the import tariff τ_x^H does not affect the outside option of the supplier (as would be the case if the secondary market were domestic in nature), it is still the case that a subsidy ($\tau_x^H < 0$) raises the equilibrium value of x by increasing the surplus over which the parties bargain.

¹³ Notice too that $\hat{\tau}_x^H$ is increasing in δ . Hence, the larger the ex-post bargaining power of suppliers, the smaller the optimal level of the home country subsidy on imported inputs.

2.4 Export Policy

We next consider optimal unilateral trade policy intervention on the part of the foreign country F , returning for now to the assumption that H maintains a policy of free trade. It is intuitively clear (and is easily shown) that F has no incentive to impose trade taxes on trade in the final good 1, since such trade taxes could only alter the domestic price of good 1 in F (owing to F 's small size on world markets) and that price has no impact on the hold-up problem between F 's input suppliers and H 's final good producers. But what about an export tax/subsidy imposed by F on its exports of inputs to H ? Presumably, an export subsidy on F 's exports of x to H could help solve the hold-up problem just as we showed above that H 's import subsidy could serve this role, but it seems unlikely that F would have an incentive to intervene in this way when H 's producers capture all of the surplus from the bilateral producer-supplier relationship in any event. A more promising possibility is that F might benefit by taxing its exports of x , if in doing so it can collect revenue from producers in H .

To sort out these possibilities, we now set $\tau_1^H = \tau_x^H \equiv 0$ and suppose instead that the foreign government taxes trade in inputs. Let τ_x^F denote the foreign export tax (if positive) or subsidy (if negative) on exports of the input x from F to H . As before, we introduce the following *stage 0* which occurs prior to *stage 1* of the 5-stage game described in subsection 2.1:

stage 0. The foreign government F selects a trade tax τ_x^F on the exported input x .

As before, after the foreign government has selected its tariffs in *stage 0*, the sequence of events is as outlined in subsection 2.1 (with trade taxes collected at the time of importation and production/sales in *stage 5*).

Consider now how F 's trade policy choices in *stage 0* affect the equilibrium outcome of the game. In their *stage-3* bargaining, if the producer and supplier reach an agreement they stand to obtain a joint payoff (recalling once more that the cost of producing x is sunk at that point) of $y(x) - \tau_x^F x$. If the parties don't reach an agreement, the home producer is left with nothing, while the foreign input supplier obtains $\delta y(x) - \tau_x^F x$.

As with the case of free trade, we again have that the ex-post gains from trade (or quasi-rents) are $(1 - \delta)y(x)$. Hence, the producer in H again obtains $\frac{1}{2}(1 - \delta)y(x)$. The payoff to the supplier in F is now equal to $\frac{1}{2}(1 + \delta)y(x) - \tau_x^F x$, and hence its choice of x in *stage 2* must satisfy

$$\frac{1}{2}(1 + \delta)y'(\hat{x}) = 1 + \tau_x^F. \quad (7)$$

We now have that \hat{x} is decreasing in τ_x^F , because a foreign tax on exported inputs (like a home tariff on imported inputs) reduces the surplus over which the parties negotiate as well as the outside option of the supplier. Clearly then, as (7) indicates, it is possible for τ_x^F to be set in such a way as to solve the hold-up problem and bring \hat{x} up to x^E , and doing so would require an export subsidy ($\tau_x^F < 0$).

At *stage 1*, the final-good producer in H can retain the input supplier in F by suggesting an initial lump-sum transfer of $T = \frac{1}{2}(1 + \delta)y(\hat{x}) - (1 + \tau_x^F)\hat{x}$, which will leave the input supplier in F with a payoff of $\pi^F = 0$.

We may now define foreign welfare as a function of the foreign trade policy intervention. Given that the foreign government has no incentive to intervene in its final-good market, foreign consumer surplus is tied down by the fixed world price of good 1, and so our measure of foreign social welfare is simply the sum of consumer surplus evaluated at $p^F = 1$, foreign profits (which, as we have just established, is zero) and trade tax revenue:

$$W^F = CS(1) + \pi^F + \tau_x^F \hat{x} = CS(1) + \tau_x^F \hat{x}.$$

Hence, the unilaterally optimal export policy for F simply maximizes the collection of trade tax revenue, and therefore is defined by¹⁴

$$\frac{\partial W^F}{\partial \tau_x^F} = \hat{x} + \tau_x^F \frac{\partial \hat{x}}{\partial \tau_x^F} = 0. \quad (8)$$

Because $\partial \hat{x} / \partial \tau_x^F < 0$, it is clear from (8) that the optimal foreign trade policy on exports of x , which we denote by $\hat{\tau}_x^F$, must involve a positive export tax: $\hat{\tau}_x^F > 0$. Defining $\hat{\tau}_1^F$ as the optimal foreign trade policy on trade in the final good 1, and recalling from our earlier discussion that $\hat{\tau}_1^F = 0$, we may thus state:

Proposition 3 *In the Benchmark Model, when only F intervenes with trade policy, its unilaterally optimal policy choices maintain free trade in the final good ($\hat{\tau}_1^F = 0$) and tax exports of the input ($\hat{\tau}_x^F > 0$) so as to generate maximal export-tax revenue.*

The intuition for Proposition 3 is as follows. Absent government intervention, foreign input suppliers gain nothing from their bilateral relationships with home final good producers in our Benchmark Model, and as a consequence welfare in F is unaffected by the hold-up problem. This gives the foreign government no incentive to intervene to try to fix the problem, even though as (7) indicates, it could do so with an appropriate export subsidy on x . Instead, as the foreign government can commit to an export policy before its suppliers are matched with producers in the home country (i.e., before *stage 1*), its only goal is to collect maximal trade tax revenue from the exportation of x , because the incidence of its tax is borne completely by final good producers in H (through a reduction in the transfer T that can be extracted from the foreign supplier by the domestic producer).^{15,16}

¹⁴For the second-order conditions to be satisfied in this case, we need to impose that $2 + xy'''(x)/y''(x) > 0$ for all x (see Appendix). As an example, if $y(x) = \lambda x^\eta$, this condition is met for all $\eta > 0$.

¹⁵Naturally, the ability of the foreign government to extract surplus from home producers partially relies on the assumption that home producers can seek suppliers only in the foreign country. In section 4, we develop an extension of the model in which home producers can search for suppliers in one of two foreign countries. We show there that both foreign countries will continue to set a positive export tax even when one of these countries becomes infinitesimally small in relation to the size of world markets.

¹⁶It is interesting to note that our Benchmark Model suggests a novel explanation for the observation that developing

Of course, the foreign export tax policy is an internationally inefficient beggar-thy-neighbor policy, because while it transfers surplus from H to F , it also reduces the level of overall surplus by worsening the hold-up problem between foreign suppliers and home producers.^{17,18} This raises the question of whether unilaterally optimal foreign export intervention of this kind would interfere with the unilaterally optimal home-country import intervention that was shown in the previous subsection to solve the hold-up problem. To answer this question and thereby understand the potential role for trade agreements in this environment, we turn in the next subsection to characterizing the Nash policy equilibrium between the home and foreign governments.

2.5 Nash Policies and the Nature of Trade Agreements

We suppose now that both the home and foreign governments are policy active, and we characterize their non-cooperative Nash equilibrium trade policy choices. In light of the analysis of the previous two subsections, this is easily done.

As before, we introduce the following *stage 0* which occurs prior to *stage 1* of the 5-stage game described in subsection 2.1:

stage 0. The home government H selects a trade tax τ_1^H on the final good 1, and a trade tax τ_x^H on the imported input x ; simultaneously, the foreign government F selects a trade tax τ_x^F on the exported input x .

After the home and foreign governments have selected their respective tariffs in *stage 0*, the sequence of events is as outlined in subsection 2.1 (with trade taxes collected at the time of importation and production/sales in *stage 5*).

countries often appear to use trade policy as a means to raise revenue, and therefore set revenue-maximizing trade taxes. The standard explanation is that these countries, while small in world markets, impose taxes on trade because their governments have revenue needs but do not have access to more appropriate measures of the kind typically available to developed-country governments for taxing citizens to raise revenue (e.g., an income tax); and so these governments are forced to use trade taxes as a second-best measure for raising revenue from their own citizens. In our Benchmark Model, governments have no revenue needs of their own: rather, they are assumed only to use trade taxes as a way to increase real national income. Nevertheless, the foreign country in our Benchmark Model could be interpreted as a less-developed country, both because its firms produce the inputs that are finished for sale in the home (developed) country, and because we have assumed that the home country producers have the (*stage 1*) bargaining power to capture all of the surplus from their bilateral relationship with foreign suppliers. Interpreted in this way, our Benchmark Model then indicates that the use of tariffs by developing-country governments for the purpose of raising revenue may in fact not be a second-best measure for raising revenue from their own citizens to fund public expenditures, but rather represents a first-best measure for extracting revenue from developed-country producers in light of the poor bargaining power of the less-developed-country supplier firms in their bilateral relationship with developed-country producers.

¹⁷It may be thought that a more efficient arrangement would involve a lower export tax by the foreign government in “exchange” for a lower ex-ante lump-sum transfer charged by home producers to their foreign suppliers. This agreement would however violate subgame perfection, since foreign producers would have no incentive to carry out their part of the “bargain.” We will show however that a similar agreement between the home and foreign governments will indeed restore efficiency.

¹⁸In the Appendix, we show that \hat{x} is increasing in δ , just as in the model with free trade. Hence, the larger the ex-post bargaining power of suppliers, the smaller the equilibrium hold-up inefficiencies, even taking into account foreign’s optimal export tax.

In their *stage-3* bargaining, if the producer and supplier reach an agreement they now stand to obtain a joint payoff of $(1 + \tau_1^H)y(x) - (\tau_x^H + \tau_x^F)x$. If the negotiations instead break down, the final-good producer is left with nothing, while the input supplier obtains a payoff of $\delta(1 + \tau_1^H)y(x) - (\tau_x^H + \tau_x^F)x$ in the secondary market.

We again have that the ex-post gains from trade (or quasi-rents) are $(1 - \delta)(1 + \tau_1^H)y(x)$, and hence, the final-good producer obtains $\frac{1}{2}(1 - \delta)(1 + \tau_1^H)y(x)$, while the payoff to the supplier is now equal to $\frac{1}{2}(1 + \delta)(1 + \tau_1^H)y(x) - (\tau_x^H + \tau_x^F)x$. Consequently, the input supplier's choice of x at *stage 2* now satisfies

$$\frac{1}{2}(1 + \delta)(1 + \tau_1^H)y'(\hat{x}) = 1 + \tau_x^H + \tau_x^F. \quad (9)$$

At *stage 1*, the final-good producer in H can retain the input supplier in F by suggesting an initial lump-sum transfer of $T = \frac{1}{2}(1 + \delta)(1 + \tau_1^H)y(\hat{x}) - (1 + \tau_x^H + \tau_x^F)\hat{x}$, which will leave the input supplier in F with a payoff of $\pi^F = 0$, while the final-good producer in H receives a payoff of

$$\pi^H = (1 + \tau_1^H)y(\hat{x}) - (1 + \tau_x^H + \tau_x^F)\hat{x}.$$

Home-country welfare may then be represented by

$$\begin{aligned} W^H &= CS(p_1^H) + \pi^H + \tau_1^H[D_1(p_1^H) - y(\hat{x})] + \tau_x^H\hat{x} \\ &= CS(p_1^H) + \tau_1^H D_1(p_1^H) + y(\hat{x}) - (1 + \tau_x^F)\hat{x}, \end{aligned}$$

while foreign-country welfare is given by

$$W^F = CS(1) + \pi^F + \tau_x^F\hat{x} = CS(1) + \tau_x^F\hat{x}.$$

Hence, the Nash trade policy choices of H and F , which we denote by $\hat{\tau}_1^{HN}$, $\hat{\tau}_x^{HN}$, and $\hat{\tau}_x^{FN}$, are determined by the following three first-order conditions¹⁹

$$\begin{aligned} \frac{\partial W^H}{\partial \tau_1^H} &= \tau_1^H \frac{\partial D_1}{\partial p_1^H} + [y'(\hat{x}) - (1 + \tau_x^F)] \frac{\partial \hat{x}}{\partial \tau_1^H} = 0; \\ \frac{\partial W^H}{\partial \tau_x^H} &= [y'(\hat{x}) - (1 + \tau_x^F)] \frac{\partial \hat{x}}{\partial \tau_x^H} = 0; \text{ and} \\ \frac{\partial W^F}{\partial \tau_x^F} &= \hat{x} + \tau_x^F \frac{\partial \hat{x}}{\partial \tau_x^F} = 0. \end{aligned} \quad (10)$$

The Nash first-order conditions in (10) are illuminating. Observe first that they immediately imply $\hat{\tau}_1^{HN} = 0$: there is no role for intervening in final goods trade (for either H or F) in the Nash equilibrium. Observe next that $\hat{\tau}_x^{FN} > 0$: the foreign tax on exported inputs remains positive in the Nash equilibrium. And finally, note (using (9)) that the home import subsidy remains positive in the Nash equilibrium ($\hat{\tau}_x^{HN} < 0$), but that it is no longer high enough to fully correct the international hold-up problem and bring \hat{x} up to the level of x^E . This is due to the positive foreign

¹⁹ Again, for the second-order conditions to be satisfied we need to impose that $2 + xy'''(x)/y''(x) > 0$ for all x .

export tax, as can be confirmed by noting that $\hat{\tau}_x^{HN}$ is set so that $y'(\hat{x}) = (1 + \hat{\tau}_x^{FN})$, whereas efficiency requires $y'(\hat{x}) = 1$ as indicated by (2).²⁰

Hence, in our Benchmark Model, the Nash trade policy equilibrium is inefficient from an international perspective. The inefficiency arises from the hold-up problem between foreign suppliers and home producers, and it results in a volume of input trade which is inefficiently low. The Nash intervention of the home government works to mitigate the hold-up problem, but the Nash intervention of the foreign government exacerbates the hold-up problem, and the net effect of Nash intervention falls short of a full solution (and may on net make the problem worse than under free trade).

In our Benchmark Model, then, a reason for a trade agreement between H and F exists, and its purpose is to better help governments solve the international hold-up problem between producers and suppliers. However, notice that a particularly simple form of the trade agreement is suggested by our Benchmark Model, provided that the governments of H and F have the ability to exchange lump sum transfers as a means of redistributing the gains from trade liberalization between them: according to the Benchmark Model, a ban on export taxes, leaving all other trade policies unconstrained, would be sufficient to solve the hold-up problem and bring H and F to the international efficiency frontier.²¹

We summarize this discussion with:

Proposition 4 *In the Benchmark Model, the international hold-up problem continues to exist in the Nash policy equilibrium between H and F , leading to inefficiently low trade in inputs between H and F . A trade agreement can then be useful as a means of helping governments better address the international hold-up problem. If governments have the ability to exchange lump sum transfers between them, a ban on export taxes would be sufficient to solve the hold-up problem and bring H and F to the international efficiency frontier.*

As Proposition 4 records, our Benchmark Model indicates that international hold-up problems between home producers and foreign suppliers can give rise to inefficiencies that are not fully corrected by government intervention in non-cooperative settings, hence pointing to a novel reason for mutually beneficial trade agreements. In effect, as we have demonstrated, each government has the ability to correct the international hold-up problem with its unilateral trade policy choices, but neither government acting alone can capture fully the benefits from solving the problem, and so

²⁰In fact, recalling that $\hat{\tau}_x^H$ is the tariff chosen by the home government when it alone is policy active, and recalling as well from Proposition 2 that the hold-up problem is completely solved in that case, it can be shown that $\hat{\tau}_x^{HN} = (1 + \hat{\tau}_x^{FN}) \hat{\tau}_x^H$. Hence, the home government responds to the foreign export tax by raising its import subsidy, thereby partially accommodating the negative effect of the foreign export tax on input production, but not by enough to maintain a complete solution to the hold-up problem.

²¹Direct transfers are rarely used as part of international trade agreements. But it is not uncommon for countries to grant market access through tariff concessions in return for concessions from their trading partners that are unrelated to market access (e.g., tighter intellectual property rights protection, or beefed-up drug control efforts), and such non-market-access concessions can be interpreted as serving the role of international transfers that redistribute the gains from trade liberalization across countries.

neither government has sufficient incentive in the Nash equilibrium to fully correct the international hold-up problem and bring the world to the international efficiency frontier.²²

As we have observed, the nature of the unilateral incentives for trade policy intervention are starkly different across H and F in our Benchmark Model. This feature has served a useful pedagogical purpose, by highlighting the potential role for trade policy in solving the international hold-up problem, and by highlighting as well the various incentives for trade policy intervention that arise. However, this also gives rise to a stark implication: in principle, a simple ban on export taxes could bring H and F to the international efficiency frontier. This implication, in turn, derives from our assumption that lump-sum transfers can be made between producers and suppliers in *stage 1* (and that producers in H have all of the bargaining power at this stage). As we have indicated above, the ability of producers and suppliers to exchange ex-ante lump-sum transfers may be questioned, particularly in the international context that we study, where such transfers and the obligations associated with them might be especially difficult to enforce. Therefore, in the next section we depart from the Benchmark Model by assuming that such transfers between producers and suppliers are infeasible, and we reconsider the unilateral policy incentives of H and F in this alternative environment and the nature of trade agreements that is implied.

3 The Basic International Hold-Up Model

We now introduce a single departure from the Benchmark Model analyzed in the previous section: we no longer permit the home producer and foreign supplier to make lump-sum transfers in *stage 1*. Everything else is unchanged from before. We refer to this model as the Basic International Hold-up Model, or the *Basic Model*. As discussed previously, disallowing ex-ante lump sum transfers between producers and suppliers seems particularly reasonable in the international context of the hold-up problem that we study.

For reference, we present the complete structure of the bilateral producer-supplier relationship with the altered *stage 1* of the Basic Model. As before, we assume that all agents have an ex-ante zero outside option. The sequence of events is as follows:

stage 1. The unit measure of producers in H and suppliers in F are randomly matched, producing a unit measure of matches. The producer provides the supplier with a list of customized input specifications.

stage 2. Each supplier decides on the amount x of customized input to be produced (at marginal cost of 1) – the marginal cost of production for a generic input (not customized to the

²²In the Appendix we show that not only is \hat{x} suboptimal, but the equilibrium value of \hat{x} is *decreasing* in δ , which is in sharp contrast with the results under free trade and unilateral export policy. Intuitively, as we have seen, the incentive of the home government to solve the hold-up problem and raise \hat{x} to x^E is diminished in the presence of a foreign export tax, and so the degree of underinvestment in the Nash equilibrium rises monotonically with the Nash level of the foreign export tax $\hat{\tau}_x^{FN}$. From (9), a high δ makes \hat{x} less sensitive to τ_x^F , and therefore makes a high τ_x^F more attractive from the perspective of foreign trade tax collection, which in turn implies a bigger Nash foreign export tax $\hat{\tau}_x^{FN}$ and hence a bigger underinvestment problem in the Nash equilibrium.

matched producer's needs) is the same, so there is no benefit in not customizing the input for the matched producer at this point.

stage 3. Each producer-supplier pair bargains over the price of the intermediate input (we assume symmetric Nash bargaining).

stage 4. A small number (formally, a measure-zero countable infinity) n of the bilateral pairs are exogenously dissolved and randomly rematched in a secondary market. They bargain again, but each supplier has full bargaining power at this point. No further inputs can be produced; the amount produced in *stage 2* is perceived as generic in the secondary market because it was tailored to another producer's specifications with probability one.

stage 5. Each producer in H imports x from its partner-supplier and produces the final good with the acquired x , and payments agreed in *stages 3* and *4* are settled.

The lack of lump-sum transfers in *stage 1* of this alternative 5-stage game implies that in the Basic Model the home producer will now get the payoff that results from its *stage-3* Nash bargain with the foreign supplier, while the foreign supplier will now get the payoff that results from its *stage-3* Nash bargain with the home producer less its cost of (*stage-2*) input supply. The question is then how these differing payoffs translate into differing incentives for the home and foreign government to intervene with trade policy.

Unilateral Home Policy

Let us begin by considering the unilateral incentives of the home government to intervene with trade policy (as before, in a prior *stage 0*) in the Basic Model, assuming for the moment that the foreign government remains passive. The final-good producer in H now has a payoff of

$$\pi^H = \frac{1}{2} (1 - \delta) (1 + \tau_1^H) y(\hat{x}), \quad (11)$$

with \hat{x} defined by (4), which implies that home welfare is now given by

$$W^H = CS(p_1^H) + \frac{1}{2} (1 - \delta) (1 + \tau_1^H) y(\hat{x}) + \tau_1^H [D_1(p_1^H) - y(\hat{x})] + \tau_x^H \hat{x}.$$

The optimal choice of τ_1^H and τ_x^H , which we again denote by $\hat{\tau}_1^H$ and $\hat{\tau}_x^H$, hence must satisfy the first-order conditions

$$\begin{aligned} \frac{\partial W^H}{\partial \tau_1^H} &= 0 = \tau_1^H \frac{\partial D_1}{\partial p_1^H} - \frac{1}{2} (1 + \delta) y(\hat{x}) - \\ &\quad \left[\frac{1}{2} (1 + \delta) y'(\hat{x}) \tau_1^H - \frac{1}{2} (1 - \delta) y'(\hat{x}) - \tau_x^H \right] \frac{\partial \hat{x}}{\partial \tau_1^H}, \text{ and} \\ \frac{\partial W^H}{\partial \tau_x^H} &= 0 = \hat{x} - \left[\frac{1}{2} (1 + \delta) y'(\hat{x}) \tau_1^H - \frac{1}{2} (1 - \delta) y'(\hat{x}) - \tau_x^H \right] \frac{\partial \hat{x}}{\partial \tau_x^H}, \end{aligned}$$

where recall that \hat{x} is given by equation (4).²³ Applying the implicit function theorem (twice) to (4) delivers

$$\frac{\partial \hat{x} / \partial \tau_1^H}{\partial \hat{x} / \partial \tau_x^H} = -\frac{1}{2} (1 + \delta) y'(\hat{x}),$$

which can be used to manipulate the above first-order conditions to obtain:

$$\begin{aligned} \hat{\tau}_1^H &= \frac{\frac{1}{2} (1 + \delta) \hat{x} [y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}]}{-\partial D_1 / \partial p_1^H}, \text{ and} \\ \hat{\tau}_x^H &= y'(\hat{x}) \hat{\tau}_1^H - \frac{1}{2} (1 - \delta) (1 + \hat{\tau}_1^H) y'(\hat{x}) - \frac{\hat{x}}{\partial \hat{x} / \partial \tau_x^H}. \end{aligned} \quad (12)$$

The expressions in (12) reflect an interesting logic. As was the case in the Benchmark Model where ex-ante transfers are available to the producer and supplier, in the Basic Model the home country's goal in intervening with τ_1^H and/or τ_x^H is to raise \hat{x} towards its efficient level x^E , but now there is an offsetting leakage of surplus to the foreign supplier that must be taken into account in the optimal unilateral policy choices for the home government. This leads to two observations: first, it is no longer optimal to deliver the chosen \hat{x} using only τ_x^H , and the use of τ_1^H reflects a new and independent source of international inefficiency associated with the unilateral policy choices of the home country; and second, it is no longer optimal for the home country to raise \hat{x} all the way to its efficient level x^E .

The first observation can be understood as follows. In the Basic Model the home government must now concern itself with two tasks as it considers its policy choices: first, as in the Benchmark Model, it must face foreign suppliers with the appropriate *marginal* incentives for investment in the supply of x so as to achieve the desired \hat{x} ; but in addition, as its final-good producers are unable in the Basic Model to extract surplus from foreign suppliers with lump-sum transfers, the home government must also concern itself with extracting *infra-marginal* surplus from foreign suppliers. With its two tariff instruments τ_1^H and τ_x^H , the home government can orchestrate infra-marginal foreign surplus extraction with adjustments in τ_1^H and τ_x^H that hold \hat{x} fixed, so that $d\tau_x^H(\tau_1^H)/d\tau_1^H = -\frac{\partial \hat{x} / \partial \tau_1^H}{\partial \hat{x} / \partial \tau_x^H}$, and can extract foreign surplus in this fashion at the rate

$$\frac{\partial W^F(\tau_1^H, \tau_x^H(\tau_1^H))}{\partial \tau_1^H} \Big|_{d\hat{x}=0} = -\frac{1}{2} (1 + \delta) \hat{x} [y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}]. \quad (13)$$

Evidently, with the concavity of $y(x)$ implying $[y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}] < 0$, it follows from (13) that for any given level of \hat{x} , additional surplus can be extracted from the foreign country by reducing τ_1^H and accompanying this with a reduction in τ_x^H which preserves the level of \hat{x} . What, then, prevents the home country from lowering τ_1^H in this fashion indefinitely, until all of the surplus has been extracted from foreign suppliers? The impact on home-country welfare of these tariff changes is

²³The second-order conditions for each of the maximization problems in the Basic Model do not reduce to simple expressions as was the case with the Benchmark Model. In the Appendix, we present the second-order conditions for the Basic Model and show that they are satisfied for a simple parameterized example.

given by

$$\frac{\partial W^H(\tau_1^H, \tau_x^H(\tau_1^H))}{\partial \tau_1^H} \Big|_{d\hat{x}=0} = \tau_1^H \frac{\partial D_1^H}{\partial p_1^H} + \frac{1}{2} (1 + \delta) \hat{x} [y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}]. \quad (14)$$

As (14) makes clear, what eventually stops this process of foreign surplus extraction is the growing home-country final-good demand distortions that are associated with $\tau_1^H < 0$. It is for these reasons that (12) implies $\hat{\tau}_1^H < 0$: in the Basic Model, it is now unilaterally optimal for the home government to utilize trade policy to distort downward the price of the final good 1 in the home market (through either an import subsidy or an export tax on the final good) as a means of extracting bargaining surplus from foreign suppliers. Notice as well that this incentive is present regardless of the level of \hat{x} , and in particular is present even when \hat{x} is set at its efficient level x^E . Hence, in the Basic Model there is a new and independent source of international inefficiency when the home country sets its tariffs unilaterally: the attempt to extract bargaining surplus from foreign suppliers by distorting the home market price of the final good 1.

The second observation above, that it is no longer optimal for the home country to raise \hat{x} all the way to its efficient level x^E , can be confirmed by considering the expression for $\hat{\tau}_x^H$ in (12). This expression is of indeterminate sign, indicating that $\hat{\tau}_x^H$ can now be either negative (an import subsidy on inputs of x) or positive (an import tariff on inputs of x): this reflects the tension that arises for the home-country government in the Basic Model between correcting the hold-up problem and capturing surplus from the foreign input supplier, a tension that is absent in the Benchmark Model where ex-ante transfers between home producers and foreign suppliers are available. Nevertheless, substituting (4) into the expression for $\hat{\tau}_x^H$ in (12) and simplifying yields

$$y'(\hat{x}) = 1 - \frac{\hat{x}}{\partial \hat{x} / \partial \tau_x^H} > 1.$$

Hence, in the Basic Model, and in contrast to the Benchmark Model, when the exporting government remains passive it is unilaterally optimal for the home government to utilize its trade policies in a way that does *not* fully correct the international hold-up problem.

Contrary to the Benchmark Model, then, we may conclude in the Basic Model that when only H intervenes international efficiency is not achieved. Instead, there are now *two* sources of international inefficiency that arise: an inefficiently low input trade volume that results from the continued existence of the international hold-up problem; and distortions in the final good market that arise as a result of the home-country's attempts to extract bargaining surplus from foreign suppliers. We may thus state:

Proposition 5 *In the Basic Model, when only H intervenes with trade policy, its unilaterally optimal policy choices reduce the local price of the final good – with an import subsidy or an export tax on trade in good 1 – and may either tax or subsidize importation of the input, but in any case result in (i) insufficient investment in the production of imported foreign inputs ($\hat{x} < x^E$) and the continued existence of the international hold-up problem, and (ii) distortions in the final good market that arise as a result of the home-country's attempts to extract bargaining surplus from*

foreign suppliers.

Proposition 5 stands in marked contrast to Proposition 2, and reflects a simple point. Unless home-country producers have the ability to extract all ex-ante surplus from their bilateral relationship with foreign suppliers (as was the case in our Benchmark Model), the unilateral incentives of the home-country government to intervene with trade policy to mitigate the international hold-up problem will be muted by the fact that foreign suppliers enjoy some of the benefits of this intervention. In this environment, the home-country's unilateral intervention must be concerned as well with capturing foreign surplus, and therefore the home country cannot be counted on to solve the international hold-up problem on its own. Moreover, the home-country's attempts to extract bargaining surplus from foreign suppliers will spill-over into the final good market as well, and introduce additional distortions there.

Foreign Retaliation and Nash Equilibrium

We turn next to consider the unilateral incentives of the foreign government to intervene with a trade tax τ_x^F (as before, in a prior *stage* 0) in the Basic Model. Anticipating the analysis of the Nash equilibrium below, we now skip the analysis of F 's tariff choices when H remains passive and instead consider F 's incentive to intervene under the assumption that H 's policy choices are given by a pair (τ_1^H, τ_x^H) . In this case, the input supplier in F now has a payoff of

$$\pi^F = \frac{1}{2} (1 + \delta) (1 + \tau_1^H) y(\hat{x}) - (1 + \tau_x^H + \tau_x^F) \hat{x},$$

with \hat{x} defined by (9), which implies that foreign welfare is now given by

$$W^F = CS(1) + \frac{1}{2} (1 + \delta) (1 + \tau_1^H) y(\hat{x}) - (1 + \tau_x^H) \hat{x}.$$

The optimal choice of τ_x^F , which we denote by $\hat{\tau}_x^F$, hence must satisfy the first-order condition

$$\frac{\partial W^F}{\partial \tau_x^F} = 0 = \left[\frac{1}{2} (1 + \delta) (1 + \tau_1^H) y'(\hat{x}) - (1 + \tau_x^H) \right] \frac{\partial \hat{x}}{\partial \tau_x^F}. \quad (15)$$

Recalling that $\partial \hat{x} / \partial \tau_x^F < 0$, the first-order condition in (15) is satisfied only when $\frac{1}{2} (1 + \delta) (1 + \tau_1^H) y'(\hat{x}) = (1 + \tau_x^H)$, which by (9) then implies that $\hat{\tau}_x^F = 0$. Hence, $\hat{\tau}_x^F = 0$ represents the best-response of the foreign government to *any* selection of policies τ_1^H and τ_x^H by the home government. A direct implication of this result is that the optimal home policies $\hat{\tau}_1^H$ and $\hat{\tau}_x^H$ derived above (which represent the home government's best-response tariff choices to $\tau_x^F = 0$) also correspond to the home government's Nash policy choices $\hat{\tau}_1^{HN}$ and $\hat{\tau}_x^{HN}$, while the foreign government's Nash policy choice is simply $\hat{\tau}_x^{FN} = 0$. Furthermore, from our results above, this Nash equilibrium involves suboptimal trade in intermediate inputs, $\hat{x} < x^E$. Hence we may state:

Proposition 6 *In the Nash equilibrium of the Basic Model, F adopts an export policy of laissez faire, while H adopts trade intervention in both the final good and the intermediate input. There is insufficient investment in the production of imported foreign inputs ($\hat{x} < x^E$) that results from the continued existence of the international hold-up problem, and there are distortions in the home-country final good market that arise as a result of the home-country's attempts to extract bargaining surplus from foreign suppliers.*

The Nash equilibrium of the Basic Model described in Proposition 6 exhibits some interesting features that deserve further explanation. To begin with, Why doesn't the foreign government intervene in its export market (that is, Why is $\hat{\tau}_x^{FN} = 0$)? The essential reason is as follows: for a fixed level of x , the Nash bargaining payoff of the home producer (see (11)) is independent of the foreign input tariff $\hat{\tau}_x^{FN}$, and hence the foreign government cannot use an export tax in the Basic Model to collect revenue from the home producers; and while the foreign export tax could be used to alter the level of x chosen by the foreign supplier, this level is already chosen by the foreign supplier to maximize foreign profits, and so there is no gain to the foreign country from manipulating this choice with export-sector intervention. As we show in section 4, this laissez-faire result depends crucially on our assumption that the foreign supplier has full bargaining power in the secondary market. When home producers have some bargaining power in that market, τ_x^F can be (and will be) used to transfer surplus ex-post. Abstracting from these issues in the Basic Model serves to highlight that, even when the foreign government maintains a policy of laissez faire in its export sector, the home government does not intervene sufficiently on its own to solve the international hold-up problem. This is in stark contrast with our results in the Benchmark model.

Given that the foreign government maintains a policy of laissez faire in its export sector, Why is unilateral policy-setting by the home government not sufficient in the Basic Model to solve the international hold-up problem? The answer, to which we have already alluded above, is that solving the international hold-up problem in the Basic Model inevitably involves shifting some of the surplus in the bilateral producer-supplier relationship towards the foreign supplier, and so in the Basic Model the unilateral incentives of the home-country government to intervene to fix the hold-up problem are muted by the fact that foreign suppliers enjoy some of the benefits of this intervention.

A final interesting feature of Proposition 6 is that the home government chooses to distort the price of the final good in its local market, reflecting an underlying incentive that: (i) operates even though the home country is small in world markets for the final good; and (ii) operates even if the home government makes other policy adjustments to hold its input trade volume fixed. What accounts for this? After all, (i) implies directly that there can be no terms-of-trade gains for the home country in the final good market, and under market clearing (ii) would imply that movements in the foreign exporter price of traded inputs are ruled out and therefore could not be a source of terms-of-trade gains for the home country either. The answer is that the foreign exporter price of the input is determined by *bargaining*, not market clearing, and this severs the link between input price and input trade volume that market clearing would have implied, allowing the foreign

exporter price of the input to be altered by changes in home-country policies (and thus allowing bargaining surplus to be extracted from the foreign supplier) even when those changes leave input trade volume fixed.

The Role of a Trade Agreement

In the Basic Model, then, as in the Benchmark Model where ex-ante transfers between home producers and foreign suppliers are allowed, a reason for a trade agreement between H and F exists, but its purpose is now twofold: to better help governments solve the international hold-up problem between producers and suppliers; and to prevent governments from introducing distortions in an effort to extract bargaining surplus from foreign firms. Moreover, when ex-ante transfers are not feasible as in the Basic Model, the simple form of trade agreements suggested by the Benchmark Model – a ban on export taxes – is no longer effective in getting H and F to the international efficiency frontier. Indeed, as we have shown, in the Basic Model, the foreign country adopts a trade policy of laissez faire. And so, in contrast to the Benchmark Model, simply banning export taxes on the foreign input has no affect here. Interestingly, in the Basic Model it is now feasible for governments to reach the international efficiency frontier with an agreement that emphasizes *import* policies regarding the traded input: in particular, an agreement that leaves the foreign export policy τ_x^F completely unconstrained can achieve the international efficiency frontier. We summarize with the following:

Proposition 7 *In the Basic Model, a trade agreement that imposes free trade in the final good ($\tau_1^H = 0$) and a subsidy on imports of inputs equal to $\tau_x^H = -\frac{1}{2}(1 - \delta)$ increases world welfare and brings H and F to the international efficiency frontier. This is achieved by (i) solving the international hold-up problem, and (ii) preventing governments from introducing distortions in an effort to extract bargaining surplus from foreign firms. Moreover, in such a trade agreement, F 's export tax τ_x^F can be left completely unconstrained.*

Proposition 7 identifies two novel purposes for trade agreements. First, to the extent that international trade requires relationship-specific investments between domestic and foreign firms that cannot be contracted over on an ex-ante basis, the volume of this trade will be too low from the perspective of international efficiency when governments set trade policy unilaterally, and by expanding this trade volume a trade agreement can move countries toward the international efficiency frontier. And second, in the presence of such relationship-specific investments and the bargaining over price that is implied, governments face incentives to use distortionary means to extract bargaining surplus from foreign firms, and by eliminating these additional policy distortions a trade agreement can help countries achieve the international efficiency frontier. An interesting feature of this second purpose of trade agreements is that it need not imply that trade volumes should be increased under the agreement: indeed, according to Proposition 7 the home country must agree to policy changes which increase the price of the final good 1 in its local markets, and

this would amount to an agreement to *restrict* trade volume in the final good 1 if the home country is an importer of this good.²⁴

This general reasoning stands in sharp contrast to the existing theories of trade agreements, which cast the problem that a trade agreement exists to correct as either stemming from a terms-of-trade driven Prisoners' Dilemma between the governments of countries that are large in world markets (the terms-of-trade theory), or stemming from an inability of governments to make commitments to the private sector (the commitment theory).²⁵ As previously mentioned (see note 10), we have made assumptions that preclude a commitment role for trade agreements in this paper.

The role that we identify here is related to the terms-of-trade theory, in that there and here the purpose of a trade agreement is to induce those governments that have the ability to affect foreign exporter (international) prices with their trade policy choices to behave in an internationally efficient way, but there are also some key differences: for one, foreign exporter prices are determined by competitive (or sometimes non-competitive) *market* conditions in the terms-of-trade theory, while here, as we have emphasized, (the relevant) foreign exporter prices are determined as a result of ex-post bilateral producer-supplier *bargaining*; for another, according to the terms-of-trade theory there is a single international problem that a trade agreement must solve, and that is terms-of-trade manipulation, while here there is also a second international problem that a trade agreement must address, namely, the international hold-up problem.

These differences are substantive, as they lead to distinct predictions across the theories concerning the determinants of the degree of trade liberalization required to bring countries to the international efficiency frontier: according to the terms-of-trade theory, it is the Nash foreign export supply elasticities faced by each country that predict the inefficiencies in its Nash tariff choices, and therefore that predict the changes in its policies necessary to reach the international efficiency frontier (see, for example, Bagwell and Staiger, 1999, 2006); in the theory we develop here it is the details of the ex-post bilateral producer-supplier bargaining that predict the inefficiencies in each country's Nash tariff choices. As one stark example of this difference, the terms-of-trade theory predicts that the fundamental task of a trade agreement is always to increase trade volumes from their Nash levels, while as we have observed above it is possible under the theory we develop here that a trade agreement should lead to *reductions* in some (i.e., final good 1) trade volumes relative to their Nash levels.

Further understanding of the difference between the terms-of-trade theory and the theory we present here can be gained by defining here the analogue to what Bagwell and Staiger (1999) call the *political optimum*. Politically optimal tariffs are the tariffs that governments would choose were they to act "as if" they were small in international markets and therefore could not affect international prices (their terms of trade), and Bagwell and Staiger show in the context of the

²⁴In the Appendix, we show that when the location of the secondary market is instead in the foreign country, the foreign government now also has an incentive to manipulate the local price of final good 1 in its market as a means of shifting surplus toward its input supplier firms, but seeks to raise rather than lower this price with an import tariff or export subsidy. In this case, if the foreign country is an exporter of the final good 1, a trade agreement that achieves the efficiency frontier would restrict the foreign country's trade volume in that good by prohibiting export subsidies.

²⁵See Bagwell and Staiger (2002, Chapter 2) for a description of the major theories of trade agreements.

terms-of-trade theory that these tariffs attain the international efficiency frontier defined according to the underlying preferences of the governments. It is for this reason that, according to the terms-of-trade theory, preventing governments from manipulating the terms of trade with their trade policy choices is *the* task of a trade agreement, regardless of whether the member governments are concerned only with maximizing aggregate real national income or rather have distributional (e.g., political) concerns as well. It is immediate that, according to the theory we develop here, inducing governments to act “as if” they were small in international markets and could not affect international prices will *not* deliver governments to the international efficiency frontier, because such governments would then not actively be engaged in intervening to help solve the international hold-up problem (an intervention which *exploits* the ability of governments to alter international prices with their trade policy choices). Nevertheless, there is a sense in which terms-of-trade manipulation is the problem here too: it is just that here the “bad” manipulation must be separated from the “good” manipulation.

In fact, as we next show, it is possible to define a variant of politically optimal tariffs under which governments do achieve the international efficiency frontier in the Basic Model. This exercise is instructive for two reasons. First, it formalizes in the Basic Model the fundamental difference between the terms-of-trade theory and the theory we develop here in terms of the purpose of a trade agreement. And second, as we establish below in section 4.4, it permits a simple interpretation of why a further difference between the two theories arises once political economy forces are introduced.

To formalize our variant of politically optimal tariffs, we begin by defining the *international price* of the input x , which we denote by p_x^* . In words, p_x^* is the price negotiated in *stage 3* for the exchange of inputs between the foreign supplier and the home producer. It is easy to see that in the Basic Model this price is given by

$$p_x^* = \frac{\pi^F}{\hat{x}} + (1 + \tau_x^F) = \frac{1}{2} (1 + \delta) (1 + \tau_1^H) \frac{y(\hat{x})}{\hat{x}} - \tau_x^H \equiv p_x^*(\hat{x}, \tau_1^H, \tau_x^H),$$

with the price of x in the home and foreign country, respectively, then given by $p_x^H = p_x^*(\hat{x}, \tau_1^H, \tau_x^H) + \tau_x^H \equiv p_x^H(\tau_x^H, p_x^*)$ and $p_x^F = p_x^*(\hat{x}, \tau_1^H, \tau_x^H) - \tau_x^F \equiv p_x^F(\tau_x^F, p_x^*)$. With these price definitions, and with $p_1^H = 1 + \tau_1^H \equiv p_1^H(\tau_1^H)$, the home and foreign welfare expressions for the Basic Model can be rewritten in the following form:

$$\begin{aligned} W^H &= CS(p_1^H) + [p_1^H y(\hat{x}) - p_x^H \hat{x}] + (p_1^H - 1)[D(p_1^H) - y(\hat{x})] + (p_x^H - p_x^*)\hat{x} \\ &= CS(p_1^H) + p_1^H y(\hat{x}) + (p_1^H - 1)[D(p_1^H) - y(\hat{x})] - p_x^* \hat{x} \equiv W^H(p_1^H(\tau_1^H), p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x}); \text{ and} \\ W^F &= CS(1) + \hat{x}[p_x^F - 1] + (p_x^* - p_x^F)\hat{x} = CS(1) + (p_x^* - 1)\hat{x} \equiv W^F(p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x}). \end{aligned}$$

Notice that there is no reason for the home government to care about the home-country price p_x^H independently of the quantity \hat{x} , and similarly there is no reason for the foreign government to care about the foreign-country price p_x^F independently of the quantity \hat{x} , and as a consequence p_x^H does not enter as an independent argument in $W^H(p_1^H(\tau_1^H), p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x})$, and likewise p_x^F does not

enter as an independent argument in $W^F(p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x})$.

Using these welfare functions, we may now define *politically optimal tariffs* in the Basic Model as those tariffs that would be chosen unilaterally by each government if it were to “ignore” (i.e., not value) any *direct* impact of its tariff choices on the international price p_x^* , but not the indirect impact on p_x^* running through \hat{x} . Formally, politically optimal tariffs are defined as the solution to the following first-order conditions (where subscripts on W denote partial derivatives):²⁶

$$\begin{aligned} W_{p_1^H}^H + [W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^H] \frac{\partial \hat{x}}{\partial \tau_1^H} &= 0; \\ [W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^H] \frac{\partial \hat{x}}{\partial \tau_x^H} &= 0; \text{ and} \\ [W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^F] \frac{\partial \hat{x}}{\partial \tau_x^F} &= 0. \end{aligned}$$

It is straightforward to verify from these first-order conditions that the politically optimal tariffs so defined, which we denote with a superscript “*PO*,” are given by $\tau_1^{HPO} = 0$, $\tau_x^{HPO} = -\frac{1}{2}(1 - \delta)$, and $\tau_x^{FPO} = 0$, and hence achieve the international efficiency frontier.

Intuitively, in the Basic Model the foreign government is already selecting its politically optimal tariffs in the Nash equilibrium, because its tariff choices have no direct impact on p_x^* anyway, and these choices lead it to a policy of *laissez faire*; and when the home government ignores the direct impact of its tariff choices on p_x^* as demanded by the political optimum, it is induced to internalize fully the benefits of solving the international hold-up problem with its tariff choices, and these choices then lead it to solve the hold-up problem completely.

With our definition of politically optimal tariffs, we have hence formalized in the Basic Model the sense in which terms-of-trade manipulation is the problem that must be solved by a trade agreement that delivers governments to the international efficiency frontier. Contrary to the terms of trade theory, the task of a trade agreement in this setting is not to make all countries “act as if they were small” and thereby eradicate terms-of-trade manipulation altogether; rather, the task of a trade agreement in this setting is to prevent only the “bad” terms-of-trade manipulation – defined as sensitivity to the impact of tariff choices on p_x^* when \hat{x} is held fixed – from occurring, while permitting the “good” terms-of-trade manipulation – defined as sensitivity to the impact of tariff choices on p_x^* that runs through \hat{x} – to continue. Finally, not only is the nature of the (efficient) political optimum – and hence the specific purpose of a trade agreement – distinct across the two theories, but as we have also emphasized above, in the theory we develop here it is the details of the ex-post producer-supplier bargaining rather than market conditions that determines the extent to which each government deviates from its politically optimal policy choices in the Nash

²⁶The assumption here is not that governments fail to understand the full effects that their tariff choices have on the international price p_x^* ; rather, in analogy with Bagwell and Staiger (1999) the hypothetical situation we consider is that governments are not motivated by the direct impact of their tariff choices on p_x^* , in the specific sense that $W_{p_x^*}^H \frac{\partial p_x^*}{\partial \tau_1^H} = W_{p_x^*}^H \frac{\partial p_x^*}{\partial \tau_x^H} \equiv 0$ and similarly for W^F . We then identify the tariffs that would be chosen by governments with these hypothetical preferences and ask whether these tariffs are efficient with respect to actual government preferences.

equilibrium.

Of course, both our Benchmark and Basic Models impose many strong assumptions, and it is important to know whether the insights developed thus far reflect special features of these strong assumptions or rather are more general. In the next section, we explore this issue.

4 Sensitivity

In this section we consider the generality of our central findings to various alternative modeling assumptions. For the most part, we focus on studying extensions of the Basic Model.

4.1 Alternative Specifications of the Secondary Market

Thus far we have assumed that the outside options for the home-country final good producer and the foreign input supplier are determined in a secondary market in the home country in which the supplier has all the bargaining power. We now explore alternative assumptions. Below we consider the case of general ex-post bargaining power. In the Appendix we also consider the possibility that the secondary market for the foreign supplier is located in the foreign country. In each case, we show that helping governments solve the international hold-up problem between producers and suppliers and avoid the policy distortions that would be introduced by attempts to extract bargaining surplus from foreign firms remain the fundamental purposes of a trade agreement in the model.

We assume that bargaining in the secondary market of *stage 4* is characterized by the generalized Nash bargaining solution with weights α and $(1 - \alpha)$ for the home producer and foreign supplier, respectively, where $\alpha \in (0, 1)$. As before, we assume that all agents have an ex-ante zero outside option. The sequence of events is identical to that of the Basic Model, but we replace *stage 4* with:

stage 4. A small number (formally, a measure-zero countable infinity) n of the bilateral pairs are exogenously dissolved and randomly rematched in a secondary market, but now with Nash bargaining weights α and $(1 - \alpha)$ for the home producer and foreign supplier, respectively. No further inputs can be produced; the amount produced in *stage 2* is perceived as generic in the secondary market because it was tailored to another producer's specifications with probability one.

We focus directly on deriving Nash policy choices, assuming as before that the home and foreign governments select their respective tariffs simultaneously in a prior *stage 0*. Following analogous steps as in previous sections, it is easy to see that the final-good producer in H now has a *stage-3* payoff of

$$\left[\frac{1}{2} + \left(\alpha - \frac{1}{2} \right) \delta \right] (1 + \tau_1^H) y(x) - \alpha (\tau_x^H + \tau_x^F) x,$$

with the supplier in F now receiving a *stage-3* payoff of

$$\left[\frac{1}{2} + \left(\frac{1}{2} - \alpha \right) \delta \right] (1 + \tau_1^H) y(x) - (1 - \alpha) (\tau_x^H + \tau_x^F) x.$$

As a result, we have that the *stage-2* choice of \hat{x} is now defined by

$$\left[\frac{1}{2} + \left(\frac{1}{2} - \alpha \right) \delta \right] (1 + \tau_1^H) y'(\hat{x}) = 1 + (1 - \alpha) (\tau_x^H + \tau_x^F), \quad (16)$$

and hence the *stage-1* payoffs of the home and foreign firm are given by

$$\begin{aligned} \pi^H &= \left[\frac{1}{2} + \left(\alpha - \frac{1}{2} \right) \delta \right] (1 + \tau_1^H) y(\hat{x}) - \alpha (\tau_x^H + \tau_x^F) \hat{x}, \text{ and} \\ \pi^F &= \left[\frac{1}{2} + \left(\frac{1}{2} - \alpha \right) \delta \right] (1 + \tau_1^H) y(\hat{x}) - (1 - \alpha) (\tau_x^H + \tau_x^F) \hat{x} - \hat{x}. \end{aligned}$$

Home and foreign welfare may now be written as

$$\begin{aligned} W^H &= CS(p_1^H) + \pi^H + \tau_1^H [D_1(p_1^H) - y(\hat{x})] + \tau_x^H \hat{x}, \text{ and} \\ W^F &= CS(1) + \pi^F + \tau_x^F \hat{x}. \end{aligned}$$

The first-order conditions that define the Nash policies $\hat{\tau}_1^{HN}$, $\hat{\tau}_x^{HN}$, and $\hat{\tau}_x^{FN}$ can be manipulated using (16) to yield²⁷

$$\begin{aligned} \hat{\tau}_1^{HN} &= \frac{\left[\frac{1}{2} + \left(\frac{1}{2} - \alpha \right) \delta \right] \hat{x} \left[y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}} \right]}{-\partial D_1 / \partial p_1^H}, \\ \hat{\tau}_x^{HN} &= \frac{\left[\frac{1}{2} + \left(\frac{1}{2} - \alpha \right) \delta \right] (1 + \hat{\tau}_1^{HN}) y'(\hat{x})}{(1 - \alpha)} - \frac{y'(\hat{x})}{(1 - \alpha)} + \frac{\alpha \hat{\tau}_x^{FN}}{(1 - \alpha)} - \frac{\hat{x}}{\partial \hat{x} / \partial \tau_x^H}, \text{ and} \\ \hat{\tau}_x^{FN} &= -\alpha \frac{\hat{x}}{\partial \hat{x} / \partial \tau_x^F}. \end{aligned}$$

The expression for $\hat{\tau}_1^{HN}$ is negative: it is similar to the expression derived in the Basic Model and contained in (12), and it carries with it the same underlying intuition. The expressions for $\hat{\tau}_x^{HN}$ and $\hat{\tau}_x^{FN}$ differ from those derived in the Basic Model mainly in that now $\hat{\tau}_x^{FN} > 0$: that is, the foreign government will again resort to an export tax, just as it did in the Benchmark Model. The intuition for this result is that whenever foreign suppliers have less than full bargaining power in the secondary market, they bear less than the full cost of any increase in their marginal cost of production. Hence, the foreign government can pass part of the cost of an export tax on to the home country, and it will choose to do so in a non-cooperative tariff equilibrium.

It bears re-emphasis here that although the goal of an export tax in our model is similar to that in the traditional terms-of-trade theory, the determinants of the size of this tax are quite distinct. In particular, the traditional literature emphasizes the importance of size (Johnson, 1953, Kennan and Riezman, 1988, Bond and Syropoulos, 1996): larger countries will tend to set higher export taxes. The reason is that in the standard terms-of-trade theory, exporter prices are determined through market clearing conditions, and so a country can alter these prices only if it is big enough to make a difference to the market clearing price when it alters its excess supplies to the market. In

²⁷Throughout the extension section, we simply assume that second-order conditions are met.

our theory, by contrast, exporter prices are determined as a result of bilateral bargaining, and so a country can alter exporter prices if it can alter the conditions of bilateral bargaining that determine those prices. Such a country need not be “big” in the traditional sense, and in fact in our theory a prerequisite for export taxes to be used is for the foreign country suppliers to have weak (or at least less-than-complete) bargaining power in the secondary market for inputs.

Leaving aside the differences that the specifics of the bargaining process make to the equilibrium choices of $\hat{\tau}_1^{HN}$, $\hat{\tau}_x^{HN}$, and $\hat{\tau}_x^{FN}$, it is important to emphasize two general results that continue to hold regardless of the value of α . First, the international hold-up problem persists and the volume of international input trade is inefficiently low as a consequence. To see this, one can manipulate the first-order conditions and use the expression for $\partial\hat{x}/\partial\tau_x^H$ and $\partial\hat{x}/\partial\tau_x^F$ implied by (16) to derive

$$y'(\hat{x}) = 1 - \frac{\hat{x}}{\partial\hat{x}/\partial\tau_x^H} > 1,$$

which implies that $\hat{x} < x^E$. Second, our model continues to predict that there are distortions in the final good market that arise as a result of the home-country’s attempts to extract bargaining surplus from foreign suppliers. The purpose of a trade agreement therefore remains to help governments better solve these two problems.²⁸

It is also possible to use the extended model developed in this section to make a broader point. Up to now we have not taken a stance as to whether the home producer and foreign supplier are vertically related or not. According to the transaction-cost approach to the boundaries of the firm (c.f., Coase, 1936, Williamson, 1985), vertical integration would arise precisely when the hold-up inefficiencies that we have modelled above become large relative to the larger “governance” costs of running an integrated organization. According to that view, our novel rationale for trade agreements would disappear because production of the final good could then be characterized by neoclassical production theory.²⁹ Nevertheless, the property-rights approach to the theory of the firm (c.f., Grossman and Hart, 1986, Hart and Moore, 1990) has persuasively argued that firm boundaries are better understood as determining the relative bargaining power of producers (via the allocation of residual rights of control inherent in the ownership of productive physical assets) rather than as affecting the space of contracts available to economic agents. Under this interpretation, the rationale for trade agreements that we propose in this paper would very much apply to vertical integrated cross-border production relationships. A crude way to capture the essence of the property-rights theory of the firm in terms of the extended model developed in this section would be to associate international outsourcing relationships with a low value of α (the

²⁸We have also explored an extension of the Benchmark Model in which instead of allowing home firms to make *stage-1* take-it-or-leave-it offers to their foreign suppliers, we let the ex-ante distribution of joint profits be determined by Nash bargaining, with home producers capturing a fraction β of joint profits. This version of the model also delivers a Nash equilibrium featuring hold-up inefficiencies and distortions in final-good trade. The main new feature is the possibility (when β is low) that the foreign government uses export subsidies in equilibrium.

²⁹This assumes that home producers could hire foreign suppliers in a competitive market at a given price, which is consistent with the transaction-cost assumption of a frictionless integrated structure (see Grossman and Helpman, 2002, for a general equilibrium treatment).

bargaining power of home producers) as compared to the value of α applying to international insourcing relationships. With this interpretation, our finding that the fundamental purpose of a trade agreement does not depend on the value of α then suggests as well that the presence or absence of vertical integration would not alter the fundamental purpose of a trade agreement.³⁰

4.2 Adding an Alternative Source for Inputs

Our analysis has been restricted to situations in which home producers can only search for suppliers in F . It is straightforward to show that at least some of our results could be overturned when this restriction is relaxed. To illustrate this in the starkest possible way, consider again our Benchmark Model with H following laissez-faire policies. As shown in section 2, in this case the foreign government has an incentive to set a positive export tax. Now consider the case in which there is a second “foreign” country, denoted by S for “South,” with an additional unit measure of potential suppliers identical to those in F . As long as $\tau_x^F > \tau_x^S$, all home producers will prefer to match with southern suppliers over suppliers in F . It is straightforward to show then that the optimal foreign and southern export taxes that emerge from this variant of the model are 0.³¹ As a result, in this extended variant of the Benchmark Model, home producers internalize the whole gains from the offshoring relationship, and the home government chooses the first-best policies described in Proposition 2. The rationale for a trade agreement in this extended variant of the Benchmark Model has vanished.

This example, however, is special in a number of ways. To begin with, the assumption that F and S are symmetric is not innocuous: if one of the two foreign countries has a comparative advantage in supplying inputs, it can (and will) maintain a positive export tax (analogous to “limit pricing” in the case of Bertrand competition among firms), and the result of our Benchmark Model is then preserved. More importantly, the structure of the example above imposes that home producers find a match with probability one, no matter where they search for suppliers. As emphasized by Grossman and Helpman (2005), an important feature of offshoring relationships is the costly search for suitable partners. The same characteristics that make offshoring relationships contractually difficult (i.e., customization, international enforceability of contracts, etc.) preclude the existence of a frictionless competitive market for inputs or for suppliers.

To illustrate these issues in a simple way, we now assume that F contains a measure ρ of potential suppliers, while S contains a measure $1 - \rho$. If k home producers search for matches in F , the total number of successful matches there is given by the matching function $m(k, \rho) \leq \min\{k, \rho\}$, where $m(\cdot)$ is increasing in both arguments and features constant returns to scale. For simplicity, we assume that both H and S adopt laissez-faire policies. Will this force F to give up the use of an

³⁰This is not to say that the presence or absence of vertically integrated home producers and foreign suppliers would be irrelevant for the nature of trade agreements. On the contrary, to the extent that international factor ownership associated with vertically integrated multinational firms alters the objective functions of each government, the nature of trade agreements could be very much affected (see Blanchard, 2006). Rather, our point is simply that vertical integration does not by itself obviate the need for a trade agreement to address the international hold-up problem.

³¹The logic is analogous to that behind the fact that Bertrand competition implies marginal-cost pricing.

export tax? As we now demonstrate, the answer is “No.”

To show this, we begin by noting that, for home producers to be indifferent between searching in F and in S , we need:

$$\frac{m(k, \rho)}{k} (y(\hat{x}^F) - (1 + \tau_x^F) \hat{x}^F) = \frac{m(1 - k, 1 - \rho)}{1 - k} (y(\hat{x}^S) - \hat{x}^S), \quad (17)$$

where \hat{x}^F is such that $\frac{1}{2}(1 + \delta)y'(\hat{x}^F) = 1 + \tau_x^F$, while \hat{x}^S is such that $\frac{1}{2}(1 + \delta)y'(\hat{x}^S) = 1$. Equation (17) defines a negative relationship between k and τ_x^F : intuitively, an increase in the foreign export tax should be matched by an increase in the probability of finding a match in that country, which in turn requires a decrease in the measure of home producers searching for partners in that country. To see this formally, note that using the assumption of constant-return-to scale in the matching function, we can write:

$$\frac{dk}{d\tau_x^F} = \frac{(y'(\hat{x}^F) - (1 + \tau_x^F))(-\partial\hat{x}^F/\partial\tau_x^F) + \hat{x}^F}{-\frac{\rho\mu'(\rho/k)}{k\mu(\rho/k)}\frac{1}{k}(y(\hat{x}^F) - (1 + \tau_x^F)\hat{x}^F) - \frac{1}{(1-k)^2}\frac{(1-\rho)\mu'((1-\rho)/(1-k))}{\mu(\rho/k)}(y(\hat{x}^S) - \hat{x}^S)} < 0, \quad (18)$$

where $\mu(\rho/k) \equiv m(1, \rho/k)$.

In order to explore the implications of this framework for the optimal choice of an export tax in F , we first define welfare in F as the sum of consumer surplus and tariff revenue collected from all the matched bilateral pairs:

$$W^F = CS(1) + m(k, \rho)\tau_x^F\hat{x}^F,$$

It thus follows that optimal choice of τ_x^F (denoted $\hat{\tau}_x^F$) will now satisfy:

$$\frac{\partial W^F}{\partial \tau_x^F} = \frac{\partial m(k, \rho)}{\partial k} \frac{dk}{d\tau_x^F} \tau_x^F \hat{x}^F + m(k, \rho) \hat{x}^F + m(k, \rho) \tau_x^F \frac{\partial \hat{x}^F}{\partial \tau_x^F} = 0,$$

which in turn implies:

$$\hat{\tau}_x^F = \frac{\hat{x}^F}{-\frac{\partial \hat{x}^F}{\partial \tau_x^F} - \frac{\partial m(k, \rho)}{\partial k} \frac{1}{m(k, \rho)} \frac{dk}{d\tau_x^F} \hat{x}^F} > 0.$$

In sum, provided that $\frac{dk}{d\tau_x^F}$ remains bounded, the optimal export tax will be positive. It is straightforward to show that for well-behaved matching functions, the export tax will remain positive even for infinitesimally small countries. In particular, notice from equation (17) that whenever the elasticity of $m(\cdot)$ with respect to both of its arguments is positive, we will have that when $\rho \rightarrow 0$, and hence as F becomes infinitesimally small, $\frac{dk}{d\tau_x^F}$ goes to 0 as well, and thus

$$\hat{\tau}_x^F \rightarrow \frac{\hat{x}^F}{-\partial\hat{x}^F/\partial\tau_x^F},$$

which is the expression we derived in our Benchmark Model with only one country being the source

of inputs.³²

4.3 Ad Valorem Tariffs

We next return to the Basic Model but depart from the specific tariff analyzed in the previous sections, and consider instead an ad-valorem import tariff on intermediate inputs. We show here that ad valorem tariffs introduce a novel channel through which bargaining between the home producer and foreign supplier can be affected. Despite this novel channel, however, we confirm that the role played by an international trade agreement remains the same.

To this end, with the “international” (foreign exporter) price p_x^* still denoting the price negotiated in *stage 3* for the exchange of intermediate inputs between the foreign supplier and the home producer, we now let t_x^H and t_x^F denote, respectively, the home-country and foreign-country taxes on trade in the intermediate good x expressed in ad valorem terms. With this notation we highlight explicitly that the *stage-3* negotiation between producer and supplier divides surplus between them by agreeing on the *price* at which the foreign supplier sells the x units of intermediate input to the home producer. To focus on the novel aspects of ad valorem tariffs, we now ignore tariffs on the final good and set $\tau_1^H = \tau_1^F \equiv 0$. With this assumption, according to the Basic Model there would be only one problem for a trade agreement to solve in the presence of specific tariffs on trade in the intermediate input x , namely, the elimination of the international hold-up problem, and we now confirm that this remains the case when the tariffs take an ad valorem form.

Specifically, if the producer and supplier reach an agreement in their *stage-3* bargaining that specifies a price level \tilde{p}_x^* , then the home-country producer receives a *stage-3* payoff of $\omega^H = y(x) - (1 + t_x^H)\tilde{p}_x^*x$ while the foreign supplier receives a *stage-3* payoff of $\omega^F = \frac{\tilde{p}_x^*}{(1+t_x^F)}x$. Notice that this implies a bargaining frontier defined by $\omega^H = y(x) - (1 + t_x^H)(1 + t_x^F)\omega^F$: because the level of the exporter price p_x^* is used by the home producer and foreign supplier to shift surplus between them, a positive ad valorem import tariff or export tax makes the *slope* of the bargaining frontier between the home producer and the foreign supplier steeper, while a negative ad valorem tariff (an import or export subsidy) makes the slope of the bargaining frontier flatter.³³ In effect, then, ad valorem trade taxes penalize the producer and supplier for shifting surplus toward the foreign supplier (with a high p_x^*), while ad valorem trade subsidies encourage surplus-shifting in this direction, suggesting a novel channel through which ad valorem trade taxes can affect the severity of the international hold-up problem. This channel is not present when a specific tariff is instead utilized, because the slope of the bargaining frontier between producer and supplier is -1 independent of the level of

³²In particular, under the maintained assumptions, $\frac{\rho\mu'(\rho/k)}{k\mu(\rho/k)}$ is positive and bounded below 1, while $\frac{(1-\rho)\mu'((1-\rho)/(1-k))}{\mu(\rho/k)}$ goes to infinity when $\rho \rightarrow 0$.

³³We abstract here from the possibility that firms might engage in transfer-pricing-type behavior in order to avoid trade taxes or collect trade subsidies. In our setting, this amounts to assuming that firms do not have other (non-price) means to transfer surplus between them in their bilateral bargain. If they did have such means, then the price they negotiate would be determined completely by the sign of the trade taxes subject only to the ability of governments to regulate such behavior. Even without such means, the firms in our model do respond to government trade taxes by negotiating different prices, but at least when these firms are taken to be engaged in arms-length transactions this would not be interpreted as transfer pricing in the traditional sense.

the specific tariffs τ_x^H and τ_x^F .

On the other hand, if the producer and supplier fail to reach an agreement in their *stage-3* bargaining, then according to the Basic Model the home-country producer is left with nothing. The disagreement payoff of the foreign supplier is determined by the zero-profit condition for the secondary-market producer, $\delta y(x) = (1 + t_x^H)(1 + t_x^F)\hat{p}_x^*x$, which implies a disagreement payoff for the foreign supplier of $\delta y(x)/[(1 + t_x^H)(1 + t_x^F)]$.

The *stage-3* Nash bargaining problem between the home producer and foreign supplier can now be characterized as follows:

$$\begin{aligned} \text{Max}_{\omega^H, \omega^F} & (\omega^H - 0) (\omega^F - \delta y(x)/[(1 + t_x^H)(1 + t_x^F)]) \\ \text{s.t.} & \quad \omega^H = y(x) - (1 + t_x^H)(1 + t_x^F)\omega^F. \end{aligned}$$

The solution to this bargaining problem yields $\omega^H = \frac{1}{2}(1 - \delta)y(x)$ and $\omega^F = \frac{\frac{1}{2}(1 + \delta)y(x)}{(1 + t_x^H)(1 + t_x^F)}$, and an implied foreign exporter price of $\hat{p}_x^* = \frac{\frac{1}{2}(1 + \delta)y(x)}{(1 + t_x^H)x}$. The choice of x at *stage 2* is then governed by

$$\frac{1}{2}(1 + \delta)y'(\hat{x}) = (1 + t_x^H)(1 + t_x^F), \quad (19)$$

and hence \hat{x} continues to be decreasing in the (ad valorem) tariffs t_x^H and t_x^F , despite the novel channel through which the ad valorem tariffs affect the bargaining between home producer and foreign supplier. With this, we can now write the *stage-1* payoffs of the home and foreign firm as

$$\begin{aligned} \pi^H &= \frac{1}{2}(1 - \delta)y(\hat{x}), \text{ and} \\ \pi^F &= \frac{\frac{1}{2}(1 + \delta)y(\hat{x})}{(1 + t_x^H)(1 + t_x^F)} - \hat{x}. \end{aligned}$$

We consider next the Nash tariff choices. With $\tau_1^H = \tau_1^F \equiv 0$, home and foreign welfare are now given by

$$\begin{aligned} W^H &= CS(1) + \pi^H + \tau_x^H \hat{p}_x^* \hat{x} = \frac{1}{2}(1 - \delta)y(\hat{x}) + \frac{\frac{1}{2}\tau_x^H(1 + \delta)y(\hat{x})}{(1 + t_x^H)}, \\ W^F &= CS(1) + \pi^F + \frac{\tau_x^F}{1 + \tau_x^F} \hat{p}_x^* \hat{x} = \frac{\frac{1}{2}(1 + \delta)y(\hat{x})}{(1 + t_x^H)} - \hat{x}. \end{aligned}$$

Using the expression for \hat{x} in (19), it is direct to show that the first order condition for t_x^F implies $\hat{t}_x^{FN} = 0$: as was the case with specific tariffs, when tariffs take an ad valorem form the foreign country has no unilateral incentive to intervene in the Basic Model.

To check whether \hat{t}_x^{HN} might achieve international efficiency in light of $\hat{t}_x^{FN} = 0$, we make use of (19) to observe that, in combination with $t_x^F = 0$, international efficiency would require $t_x^H = -$

$\frac{1}{2}(1 - \delta)$. But differentiating W^H with respect to τ_x^H and evaluating at $t_x^F = 0$ yields

$$\frac{\partial W^H}{\partial \tau_x^H} = \frac{1}{2} \left\{ (1 - \delta) y'(\hat{x}) \frac{\partial \hat{x}}{\partial t_x^H} + \frac{(1 + \delta) y(\hat{x}) + \tau_x^H (1 + \delta) y'(\hat{x}) \frac{\partial \hat{x}}{\partial t_x^H}}{(1 + t_x^H)} - \frac{\tau_x^H (1 + \delta) y(\hat{x})}{(1 + t_x^H)^2} \right\},$$

which is strictly positive when evaluated at the internationally efficient level of $t_x^H = -\frac{1}{2}(1 - \delta)$: by implication, then, \hat{t}_x^{HN} is higher than the internationally efficient level.

Hence, while the mechanisms through which specific and ad valorem tariffs on traded inputs influence the international hold-up problem are distinct, the broad conclusions are the same. Under Nash policy choices, the international hold-up problem persists, the volume of international input trade is inefficiently low as a consequence, and the purpose of a trade agreement (when $\tau_1^H = \tau_1^F \equiv 0$) remains to help governments better solve the problem.³⁴

4.4 Political Economy

We have thus far assumed that each country's government is benevolent and seeks to maximize the aggregate welfare of its residents. Both casual and formal evidence suggest, however, that it is more realistic to formulate a social welfare function that weights asymmetrically the welfare of different groups in society. The political economy literature has stressed the role of special interest groups in generating these biases in policy (Baron, 1994, Grossman and Helpman, 1996). In this section, we explore the effects of stipulating a government welfare function that places a higher weight on producer welfare than on consumer welfare. Implicitly, we assume that producers are in a better position to solve the "collective action" problem and hence can better coordinate their demands on the government. We also assume that the ownership of productive assets is highly concentrated, so that we can ignore the role of producers as consumers and as receivers of lump-sum tax rebates. Finally, for simplicity we take the weight placed on producer welfare to be the same across countries. In particular, we let

$$W^j = CS^j + \gamma \pi^j + \text{Trade Tax Revenue}^j, \quad \text{with } \gamma > 1, \text{ for } j \in \{H, F\}. \quad (20)$$

In terms of the Basic Model, this implies that we can write the welfare of the home and foreign governments as

$$W^H = CS(p_1^H) + \frac{\gamma}{2} (1 - \delta) (1 + \tau_1^H) y(\hat{x}) + \tau_1^H [D_1(p_1^H) - y(\hat{x})] + \tau_x^H \hat{x}$$

³⁴It is interesting to observe that the novel channel through which ad valorem tariffs alter the bargaining outcome between home producer and foreign supplier – namely, the slope of the bargaining frontier – also suggests that these policy instruments may have a broader class of applicability with regard to their ability to mitigate international hold-up problems than is the case for specific tariffs. For example, if x were reinterpreted as the unverifiable quality of a fixed unit to be traded, so that tariff policy could not then be conditioned on x , a specific tariff on trade in x would lose its ability to affect the hold-up problem, but an ad valorem tariff would continue to be useful in this regard.

and

$$W^F = CS(1) + \frac{\gamma}{2} (1 + \delta) (1 + \tau_1^H) y(\hat{x}) - \gamma(1 + \tau_x^H) \hat{x} - (\gamma - 1) \tau_x^F \hat{x},$$

respectively.³⁵ Straightforward manipulation of the first-order conditions that define the Nash policies $\hat{\tau}_1^{HN}$, $\hat{\tau}_x^{HN}$, and $\hat{\tau}_x^{FN}$ delivers:

$$\hat{\tau}_1^{HN} = \frac{\frac{\gamma-1}{2} (1 - \delta) y(\hat{x}) + \frac{1}{2} (1 + \delta) \hat{x} [y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}]}{-\partial D_1 / \partial p_1^H}. \quad (21)$$

Notice that as $\gamma \rightarrow 1$, we obtain the same expression for $\hat{\tau}_1^{HN}$ as in the Basic Model, and thus for low enough $\gamma > 1$, the home government finds it optimal in the Nash equilibrium to set a positive export tax (or import subsidy) on the final good. Nevertheless, when the weight that the government welfare function places on producer surplus becomes sufficiently high, $\hat{\tau}_1^{HN}$ flips sign according to (21) and becomes positive. In such a case, the home government puts in place a Nash trade policy that leads to an *increase* in the domestic price of the final good (i.e., an import tariff or export subsidy). As we have shown above, these policies tend to transfer surplus from the home country to the foreign country, but a sufficiently politically influenced home government is willing to allow this because consumers bear a disproportionate part of the cost of this rent-dissipation.

Manipulation of the first-order conditions also delivers

$$\hat{\tau}_x^{FN} = \frac{(\gamma - 1) \hat{x}}{\partial \hat{x} / \partial \tau_x^F} < 0,$$

which indicates that for $\gamma > 1$, the foreign government no longer adopts free trade policies in the Nash equilibrium but rather chooses to subsidize exports of intermediate inputs. Intuitively, although a subsidy reduces foreign tariff revenue by an amount which is strictly larger than the amount by which foreign profits increase, a politically influenced foreign government weights the latter effect disproportionately more, and thus sets a positive export subsidy in the Nash equilibrium.

The fact that the magnitude and even the sign of Nash policies are sensitive to political economy considerations is not particularly surprising: analogous findings are reported for example in Grossman and Helpman (1994) and Bagwell and Staiger (2002, Chapter 10). Nevertheless, as we next demonstrate, the introduction of political economy concerns into the Basic Model creates a further distinction between the terms-of-trade theory and the theory we develop here concerning the fundamental purpose of a trade agreement. In particular, when political economy considerations are present, the political optimum still delivers governments to the international efficiency frontier according to the terms-of-trade theory, while in the theory we develop here this is no longer the case.

This can be seen by rewriting the home and foreign welfare expressions for the Basic Model,

³⁵It is straightforward to show that our introduction of political economy into the Basic Model does not create a reason for F to utilize τ_1^F , and so we continue to set $p_1^F \equiv 1$ and focus only on τ_1^H , τ_x^H and τ_x^F .

now augmented with political economy concerns, in the following form:

$$\begin{aligned}
W^H &= CS(p_1^H) + \gamma[p_1^H y(\hat{x}) - p_x^H \hat{x}] + (p_1^H - 1)[D(p_1^H) - y(\hat{x})] + (p_x^H - p_x^*)\hat{x} \\
&\equiv W^H(p_1^H(\tau_1^H), p_x^H(\tau_x^H, p_x^*), p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x}); \\
W^F &= CS(1) + \gamma\hat{x}[p_x^F - 1] + (p_x^* - p_x^F)\hat{x} \equiv W^F(p_x^F(\tau_x^F, p_x^*), p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x}).
\end{aligned}$$

Notice that, in contrast to the analogous representation of welfare functions presented in section 3 for the case where political economy concerns are absent, in the presence of political economy concerns there *is* a reason for the home government to care about the home-country price p_x^H independently of the quantity \hat{x} , because this price has implications for the income of politically represented producers, and similarly there *is* a reason for the foreign government to care about the foreign-country price p_x^F independently of the quantity \hat{x} . As a consequence, p_x^H now enters as an independent argument in $W^H(p_1^H(\tau_1^H), p_x^H(\tau_x^H, p_x^*), p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x})$, and likewise p_x^F now enters as an independent argument in $W^F(p_x^F(\tau_x^F, p_x^*), p_x^*(\hat{x}, \tau_1^H, \tau_x^H), \hat{x})$. Using these welfare functions, it is straightforward to verify that the efficient policy choices (i.e., those policies that maximize the sum of home and foreign welfare when evaluated in light of the objectives of the governments) must satisfy the following two conditions:

$$\begin{aligned}
W_{p_x^H}^H \frac{dp_x^*}{d\tau_x^F} + [W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^H] \frac{d\hat{x}}{d\tau_x^F} + W_{p_x^F}^F \frac{dp_x^*}{d\tau_x^H} + [W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^F] \frac{d\hat{x}}{d\tau_x^H} &= 0; \text{ and} \\
W_{p_1^H}^H + W_{p_x^H}^H \frac{dp_x^*}{d\tau_1^H} + [W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^H] \frac{d\hat{x}}{d\tau_1^H} + W_{p_x^F}^F \frac{dp_x^*}{d\tau_1^H} + [W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^F] \frac{d\hat{x}}{d\tau_1^H} &= 0,
\end{aligned}$$

where subscripts on W denote partial derivatives.³⁶ The first condition defines the efficient choice of the sum $\tau_x^H + \tau_x^F$ (τ_x^H and τ_x^F are perfect substitutes from the perspective of world welfare); the second condition defines the efficient choice of τ_1^H .

We next turn to the politically optimal tariff choices. We again define *politically optimal tariffs* in this variant of the Basic Model as those tariffs that would be chosen unilaterally by each government if it were to ignore any *direct* impact of its tariff choices on the international price p_x^* , but not the

³⁶Defining the efficiency frontier in this way fits well with the “member-driven” nature of the WTO, and it is the approach to evaluating the performance of trade agreements taken by most of the literature, but it is not the only approach. An alternative (pursued for example by Aghion, Antràs, and Helpman, 2007 and by Ornelas, forthcoming) is to evaluate the performance of trade agreements on the basis of whether or not the agreement guides governments to a point on an efficiency frontier that is defined with regard to a set of preferences that are unrelated to government preferences (e.g., the maximization of real world income). Such an approach would have some potentially interesting implications in the current setting, because politically motivated governments tend to adopt trade policies that promote producer surplus, and this tends to lead to levels of input production that are increasing in the political weight γ . In a world in which input production is too low as a result of hold up, the policy bias of politically motivated governments could therefore have some attractive consequences from the perspective of real world income. In principle, governments with a particular set of political preferences could adopt policies in the Nash equilibrium which deliver efficient levels of input production in the sense of achieving $\hat{x} = x^E$. Nevertheless, in general there is no reason to expect that political pressures would align themselves in this particular way, and so even from this perspective there would as a general matter be inefficiencies in the Nash equilibrium which a trade agreement could correct.

indirect impact on p_x^* running through \hat{x} .³⁷ Exploiting the fact that $\frac{dp_x^H}{d\tau_x^H} = \frac{dp_x^*}{d\tau_x^*}$, $\frac{dp_x^F}{d\tau_x^F} = \frac{dp_x^*}{d\tau_x^*}$, and $\frac{d\hat{x}}{d\tau_x^H} = \frac{d\hat{x}}{d\tau_x^F}$, the politically optimal tariffs τ_x^{HPO} , τ_x^{FPO} and τ_1^{HPO} can be characterized respectively by the three conditions:³⁸

$$\begin{aligned} W_{p_x^*}^H \frac{dp_x^*}{d\tau_x^F} + [W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^H] \frac{d\hat{x}}{d\tau_x^F} &= 0; \\ W_{p_x^*}^F \frac{dp_x^*}{d\tau_x^H} + [W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^F] \frac{d\hat{x}}{d\tau_x^H} &= 0; \text{ and} \\ W_{p_1^H}^H + W_{p_x^*}^H \frac{dp_x^*}{d\tau_1^H} + [W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^H] \frac{d\hat{x}}{d\tau_1^H} &= 0. \end{aligned}$$

It can be confirmed that the politically optimal tariffs τ_x^{HPO} and τ_x^{FPO} satisfy the first efficiency condition above. However, the politically optimal tariff τ_1^{HPO} fails to satisfy the second efficiency condition above unless $\gamma = 1$ (in which $W_{p_x^*}^F = 0$).

Intuitively, the efficient choice of each policy must take account of the effects that small adjustments in that policy have on the sum of home and foreign welfare. By definition, when a government sets a policy at its politically optimal level, it ignores the effects of small adjustments in that policy on the welfare of its trading partner, but it also ignores any impacts on its *own* welfare associated with movements in p_x^* that are generated directly from changes in the policy (as opposed to indirectly through the impact that the policy has on \hat{x}). From this backdrop, Why are τ_x^{HPO} and τ_x^{FPO} efficient? Consider τ_x^{HPO} . When the home government makes a small adjustment to τ_x^H , there are three impacts on foreign welfare: the effect on foreign welfare of the induced change in p_x^* that runs directly through τ_x^H ; the effect on foreign welfare of the change in p_x^F associated with the induced change in p_x^* ; and the (total) effect on foreign welfare of the induced change in \hat{x} . The home government ignores the first effect, but according to the political optimum it ignores this effect on its own welfare as well, and as the two effects cancel out in terms of overall world welfare, no inefficiency arises from this source. The home government also ignores the second and third effects of its adjustment in τ_x^H on foreign welfare, but these effects are identical to the effects on foreign welfare that the foreign government considers when it sets its politically optimal level of τ_x^F , and so the first-order condition that defines τ_x^{FPO} ensures that no inefficiency arises from the second and third effects either. A similar logic explains why τ_x^{FPO} is also efficient. Finally,

³⁷ A subtle issue arises with regard to the most appropriate definition of the political optimum in the Basic Model when political economy forces are present, relating to how the derivatives $\frac{dp_x^H}{d\tau_x^H}$ and $\frac{dp_x^F}{d\tau_x^F}$ should be calculated at the political optimum. In the characterization of the political optimum that follows we take $\frac{dp_x^H}{d\tau_x^H}$ and $\frac{dp_x^F}{d\tau_x^F}$ to include both the indirect *and* direct effects of the tariff on p_x^* , and we suppose that governments ignore the direct effects of tariffs on p_x^* only when valuing the movements in p_x^* according to $W_{p_x^*}^H \frac{\partial p_x^*}{\partial \hat{x}}$ and $W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}}$. As will become clear below, this characterization permits the cleanest interpretation of the source of the inefficiency of the political optimum when political economy forces are introduced into the Basic Model. But the inefficiency persists as well under an alternative characterization of politically optimal tariffs that has governments also ignore the direct effect of tariffs on p_x^* in their calculations of $\frac{dp_x^H}{d\tau_x^H}$ and $\frac{dp_x^F}{d\tau_x^F}$.

³⁸ We note that the three relationships $\frac{dp_x^H}{d\tau_x^H} = \frac{dp_x^*}{d\tau_x^*}$, $\frac{dp_x^F}{d\tau_x^F} = \frac{dp_x^*}{d\tau_x^*}$, and $\frac{d\hat{x}}{d\tau_x^H} = \frac{d\hat{x}}{d\tau_x^F}$, which underlie our inefficiency results for the political optimum in the presence of political economy forces, hold as well in the generalized ex-post bargaining model of section 4.1.

Why does τ_1^{HPO} fail to be efficient? The reason is that, in contrast to the case of τ_x^H , the foreign government does not possess a policy that can replicate the effects on foreign welfare of τ_1^H at the political optimum, namely, the combined effect of the change in p_x^F associated with the induced change in p_x^* and the (total) effect of the induced change in \hat{x} that comes from a small adjustment in τ_1^H ; and hence, there is no first-order condition for the foreign country at the political optimum that would prevent inefficiencies from being associated with these effects of a small adjustment in τ_1^H .

As noted above, in the absence of political economy forces (i.e., when $\gamma = 1$), it follows that $W_{p_x^F}^F = 0$, and as was argued in section 3 and can be confirmed from the first-order conditions above, this implies that τ_1^{HPO} now satisfies the conditions for efficiency as well. The reason is that, with no political economy concerns, the foreign government no longer values the redistribution of its economy's surplus toward its input suppliers that an increase in p_x^F implies, and therefore a small adjustment in τ_1^H affects foreign welfare at the political optimum only through the (total) effect of the impact of τ_1^H on \hat{x} , i.e., only through $[W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^F] \frac{d\hat{x}}{d\tau_1^H}$. But with $\frac{d\hat{x}}{d\tau_1^H} < 0$, the foreign government's choice of τ_x^{FPO} now guarantees that $[W_{p_x^*}^F \frac{\partial p_x^*}{\partial \hat{x}} + W_{\hat{x}}^F] = 0$, and therefore in the absence of political economy concerns τ_1^{HPO} is efficient as well.³⁹

Hence, in our model the purpose of a trade agreement becomes more complicated when governments possess political economy motives. In this sense, it could be said that we have formally identified a separate “political externality” for a trade agreement to address, somewhat along the lines described in Ethier (2004).⁴⁰ This is in stark contrast to the terms-of-trade theory, where the presence or absence of political economy has no impact on the purpose of a trade agreement. The essential difference between the terms-of-trade theory and the theory we develop here that generates this distinction can be traced to the difference across the two theories in the way that international prices are determined. According to the terms of trade theory, international prices are determined through market clearing conditions, with local prices in each economy then determined by arbitrage conditions that link international prices with local prices through a country's own policies. When countries choose politically optimal policies according to the terms of trade theory, they therefore ensure that the local-price effects of small adjustments in the policies of their trading partners can have no first-order effect on their welfare, because these local price effects could already have been generated by their own policy adjustments, and their politically optimal first-order-condition ensures that the welfare impacts are zero. By contrast, in the theory we develop here, international prices are determined by bilateral bargaining between foreign suppliers and home producers, and given the more complex channels through which a trading partner's policies can induce local effects

³⁹In fact, as can be inferred from our discussion, it is the *foreign* political economy forces that prevent the politically optimal choice of τ_1^H from being efficient. More generally, however, in the presence of symmetric home-supplier/foreign-producer relationships, political economy forces in either country will interfere with the efficiency properties of the political optimum.

⁴⁰As Ethier (2004, p. 305) puts it, “Political externalities,’ by my definition, arise when policymakers in one country believe that their political status (whatever that might be specified to mean) is directly sensitive, to some degree, to actions by policymakers in another country.” See also Bagwell and Staiger (2002, Chapter 2) for a further discussion of these ideas.

in a country's economy, the country's politically optimal first-order conditions cannot – except for the special case in which governments maximize real national income – ensure that the welfare effects will be zero, because there is no longer any guarantee that the local effects of a trading partner's policy adjustments could have been generated by a country's own policy adjustments.

Put differently, when international prices are determined through market clearing conditions as in the terms-of-trade theory, any redistribution toward favored producers occurs through changes in local prices, and a government does not need its trading partners to help it exercise control over its own local prices. By contrast, when international prices are determined by bilateral bargaining between foreign suppliers and home producers as in the theory we develop here, any redistribution toward favored producers occurs through changes in the outcome of bilateral bargaining, and a government *may* need its trading partners to help it exercise control over the outcome of that bargaining process.

This last point can be illustrated starkly by considering the limiting case in which $\delta \rightarrow 1$, so that there is no hold-up problem. In this case, as (11) indicates, $\pi^H = 0$ in the Basic Model and so trade policies are simply incapable of transferring surplus to producers in H : as a consequence, when $\delta \rightarrow 1$ the home-country government has *no political use of its own* for trade policy in the Basic Model regardless of its underlying political motivations.⁴¹ Nevertheless, it is straightforward to establish that the conditions for international efficiency reduce in this case to:

$$\begin{aligned} (\gamma - 1)\hat{x} + [1 - y'(\hat{x})]\frac{d\hat{x}}{d\tau_x^F} &= 0; \text{ and} \\ (\gamma - 1)\hat{x}\left[\frac{y(\hat{x})}{\hat{x}} - y'(\hat{x})\right] + \tau_1^H \frac{\partial D_1}{\partial p_1^H} &= 0. \end{aligned}$$

The first condition for international efficiency dictates that the sum $\tau_x^H + \tau_x^F$ must maintain productive efficiency only if $\gamma = 1$: when $\gamma > 1$ so that political economy motives are present, it is optimal to accept some over-production of inputs ($y'(\hat{x}) < 1$) as a necessary cost of distributing surplus to foreign producers. Hence, in order to achieve the international efficiency frontier it is necessary that the sum $\tau_x^H + \tau_x^F$ is sensitive to foreign political considerations, and it is plausible that this efficiency condition could be satisfied under the political optimum (as we show it is more generally above) with the appropriate sensitivity in τ_x^F . But the second condition for international efficiency dictates that τ_1^H should be set to zero only if $\gamma = 1$ (recall that concavity implies $[\frac{y(\hat{x})}{\hat{x}} - y'(\hat{x})] > 0$): when $\gamma > 1$ so that political economy motives are present, it is optimal to accept some final-good demand distortions in the *home* country as a necessary cost of distributing surplus to *foreign* producers. Hence, in order to achieve the international efficiency frontier it is necessary that τ_1^H is sensitive to *foreign* political considerations. Clearly, there is no possible definition of the political optimum that would yield this sort of sensitivity, and thereby allow the purpose of a trade agreement to be reduced to solving a simple terms-of-trade problem: in our model, political motives complicate the task of a trade agreement.

⁴¹This feature is reflected in (21), where it is apparent that the direct dependence of $\hat{\tau}_1^{HN}$ on γ disappears when $\delta \rightarrow 1$.

5 Conclusion

In this paper, we have initiated the study of trade agreements in the presence of offshoring. Our findings suggest that a number of the salient features of trade in inputs – including the prominent role played by relationship-specific investments and the associated contracting difficulties that are likely to give rise to international hold-up problems – have important implications for the nature and purpose of trade agreements. If governments seek to maximize real national income, then the purpose of a trade agreement in this environment is to induce governments to solve the international hold-up problem by exploiting their ability to affect international prices with their trade policy choices, while at the same time preventing governments from seeking to exploit their power over international prices for traditional terms-of-trade reasons. If governments are instead motivated by political economy/distributional concerns, the purpose of a trade agreement becomes more complex, and cannot be reduced to solving a simple terms-of-trade problem. Finally, regardless of the objectives of governments, the degree of inefficiency in the unilateral trade policy choices of governments depends on the details of the ex-post producer-supplier bargaining environment, and while input trade must always be increased in an efficient international agreement, efficiency may require that final goods trade is diminished. As we have observed, these findings are at odds with existing theories of trade agreements, and they suggest a novel interpretation of the value of trade agreements when offshoring is prevalent.

Our paper raises many new questions, both theoretical and empirical. Are international prices best thought of as determined through countless bilateral bargains between buyers and sellers, or rather through anonymous market clearing mechanisms? To the extent that it is the former, Do the trade policy stances of governments in practice have systematic impacts on bargaining outcomes and, through this channel, on trade volumes? To what extent can the architecture of the WTO, including its emphasis on reciprocity and non-discrimination, be understood from the perspective of the theory we develop here? Does the changing nature of international trade indicate the need for fundamental changes in the international institutions that govern the world trading system? These and related questions strike us as especially fertile areas for further research.

Appendix

5.1 Non-Negativity Constraints

In the main text, we have ignored situations in which equilibrium trade policies might violate the non-negativity constraints on the outside options and the surplus available to agents in the negotiation. The purpose of this Appendix is to explore those situations and show how they do not invalidate the main results of the paper. To save space, we focus on an analysis of the Benchmark and Basic International Hold-Up models developed in sections 2 and 3. In those models, the surplus over which the producer and the supplier bargain is given by

$$(1 + \tau_1^H) y(\hat{x}) - (\tau_x^H + \tau_x^F) \hat{x}, \quad (\text{A1})$$

where remember that the equilibrium \hat{x} satisfies

$$\frac{1}{2} (1 + \delta) (1 + \tau_1^H) y'(\hat{x}) = 1 + \tau_x^H + \tau_x^F. \quad (\text{A2})$$

Our first result is that *regardless* of the equilibrium values of τ_1^H , τ_x^H , and τ_x^F , the surplus in equation (A1) will always be non-negative. To see this, note that using (A2) we can write

$$(1 + \tau_1^H) y(\hat{x}) \geq \left(\frac{1}{2} (1 + \delta) (1 + \tau_1^H) \frac{y'(\hat{x}) \hat{x}}{y(\hat{x})} \right) y(\hat{x}) = \hat{x} + \hat{x} \tau_x^H + \hat{x} \tau_x^F \geq \hat{x} \tau_x^H + \hat{x} \tau_x^F,$$

where we have used that the concavity of $y(\cdot)$ implies $y'(\hat{x}) \hat{x} < y(\hat{x})$. Hence, the non-negativity constraint on the surplus can be ignored hereafter. Intuitively, no matter how distortionary trade taxes are, the level of investment x will adjust to ensure a positive joint surplus of the relationship.

Matters are not as simple with regards to the outside option of the supplier (remember that the producer's outside option is assumed to be 0). In particular, we are now careful to define this outside options as follows:

$$\max \{ \delta(1 + \tau_1^H) y(x) - (\tau_x^H + \tau_x^F) x, 0 \}.$$

It is straightforward to see that whenever $\delta \rightarrow 0$, the supplier may find it optimal to ignore the secondary market and simply throw away the amount x of input produced. In such a case, both outside options are zero and the (ex-post) payoff to both the producer and the supplier is given by

$$\pi^F = \pi^S = \frac{1}{2} ((1 + \tau_1^H) y(x) - (\tau_x^H + \tau_x^F) x).$$

As a result, the optimal choice of x at *stage 2* now satisfies

$$\frac{1}{2} (1 + \tau_1^H) y'(\tilde{x}) = 1 + \frac{1}{2} (\tau_x^H + \tau_x^F). \quad (\text{22})$$

Notice that \tilde{x} differs from \hat{x} in (A2), but it continues to be the case that \tilde{x} increases with τ_1^H and decreases with τ_x^H and τ_x^F .

We can now proceed to study equilibrium trade taxes whenever the non-negativity constraint on the supplier's outside option binds. In the Benchmark model, the ex-ante lump-sum transfer ensures that the objective function of both governments remains unaltered, and hence, the first-order conditions continue to be given by (6), with \tilde{x} (in eq. (22)) replacing \hat{x} . Consequently, all the qualitative results derived in section 2 continue to hold even when the trade policy choices drive the supplier's outside option to 0.

In the ‘‘Basic’’ model with no ex-ante transfers, the non-negativity constraint directly affects the objective function of each government. In particular, we now have that

$$W^H = CS(p_1^H) + \frac{1}{2} [(1 + \tau_1^H)y(\hat{x}) - (\tau_x^H + \tau_x^F)\hat{x}] + \tau_1^H[D_1^H(p_1^H) - y(\hat{x})] + \tau_x^H\hat{x}.$$

and

$$W^F = CS(1) + \frac{1}{2} [(1 + \tau_1^H)y(\hat{x}) - (\tau_x^H + \tau_x^F)\hat{x}] - \hat{x} + \tau_x^F\hat{x}.$$

These objective functions as well the optimality condition (22) are a particular case of the model with general ex-post bargaining power that we develop in section 4.1, where one needs to set $\alpha = 1/2$. As we discuss in that section, the main results of the paper are robust to a specification with general ex-post bargaining power.

5.2 Second-Order Conditions and Comparative Statics

In this Appendix we provide a discussion of the second-order conditions and the comparative statics associated with the main tariff setting games developed in the main text.

Benchmark Model with Unilateral Import Policy

It is easily verified that the second order conditions associated with the first-order conditions in (6) are satisfied. Simply note that evaluated at the equilibrium, we have

$$\begin{aligned} \frac{\partial^2 W^H}{\partial (\tau_1^H)^2} &= \frac{\partial D_1}{\partial p_1^H} + y''(\hat{x}) \left(\frac{\partial \hat{x}}{\partial \tau_1^H} \right)^2 < 0 \\ \frac{\partial^2 W^H}{\partial (\tau_x^H)^2} &= y''(\hat{x}) \left(\frac{\partial \hat{x}}{\partial \tau_x^H} \right)^2 < 0 \\ \frac{\partial^2 W^H}{\partial \tau_x^H \partial \tau_1^H} &= y''(\hat{x}) \frac{\partial \hat{x}}{\partial \tau_1^H} \frac{\partial \hat{x}}{\partial \tau_x^H} < 0 \end{aligned}$$

and thus $(\partial^2 W^H / \partial (\tau_1^H)^2) (\partial^2 W^H / \partial (\tau_x^H)^2) - (\partial^2 W^H / \partial \tau_x^H \partial \tau_1^H)^2 = (\partial D_1 / \partial p_1^H) y''(\hat{x}) (\partial \hat{x} / \partial \tau_x^H)^2 > 0$. As for the comparative statics, it was shown in the main text that $\hat{x} = x^E$ is independent of δ , while $\hat{\tau}_x^H$ is increasing in δ .

5.2.1 Benchmark Model with Unilateral Export Policy

Now consider the program characterized by the first-order condition in (8). Differentiating further with respect to τ_x^F , we have

$$\frac{\partial^2 W^F}{\partial (\tau_x^F)^2} = 2 \frac{\partial \hat{x}}{\partial \tau_x^F} + \tau_x^F \frac{\partial^2 \hat{x}}{\partial (\tau_x^F)^2}. \quad (23)$$

But using the implicit function theorem, we have

$$\frac{\partial \hat{x}}{\partial \tau_x^F} = \frac{1}{\frac{1}{2}(1 + \delta) y''(\hat{x})}, \quad (24)$$

which implies

$$\frac{\partial^2 \hat{x}}{\partial (\tau_x^F)^2} = \frac{1}{\frac{1}{2}(1 + \delta) (y''(\hat{x}))^2} y'''(\hat{x}) \frac{\partial \hat{x}}{\partial \tau_x^F}.$$

Using these expressions as well as $\tau_x^F = -\hat{x} / (\partial\hat{x}/\partial\tau_x^F)$ – see equation (8) –, we can write (23) as

$$\frac{\partial^2 W^F}{\partial (\tau_x^F)^2} = \frac{\partial\hat{x}}{\partial\tau_x^F} \left(2 + \frac{\hat{x}y'''(\hat{x})}{y''(\hat{x})} \right),$$

which is negative only if $2 + \hat{x}y'''(\hat{x})/y''(\hat{x}) > 0$, as claimed in footnote 14. As an example, assume that $y(x) = x^\eta/\eta$, with $\eta \in (0, 1)$. In such case, we have $y'(x) = (\eta - 1)x^{\eta-2}$ and $y''(x) = (\eta - 2)(\eta - 1)x^{\eta-3}$, and hence $2 + \hat{x}y'''(\hat{x})/y''(\hat{x}) = \eta$, which is indeed positive.

As for the comparative statics, we simply note that we can use (7), to write the first-order condition in (8) as follows:

$$\hat{x} + \frac{\frac{1}{2}(1 + \delta)y'(\hat{x}) - 1}{\frac{1}{2}(1 + \delta)y''(\hat{x})} = 0.$$

This defines an implicit function in \hat{x} and δ . It is clear that this function is decreasing in δ . But it is easily verified that the same condition we need for the second-order condition to hold, imposes that this function is increasing in \hat{x} .⁴² In sum, we can conclude that \hat{x} is an increasing function of δ , as claimed in footnote 18. Can we say anything about the effect of δ on the equilibrium Foreign tariff $\hat{\tau}_x^F$? In general, we cannot. For the particular case in which $y(x) = x^\eta/\eta$, in fact it can be verified that $\hat{\tau}_x^F = (1 - \eta)/\eta$, and is thus independent of δ .

5.2.2 Benchmark Model with Retaliation

The first-order conditions that characterize this equilibrium are found in (10). It is straightforward to verify that the second-order conditions for the choice of τ_1^H and τ_x^H are always verified, even when $\tau_x^F > 0$ (the expressions are actually identical to those above with $\tau_x^F = 0$). On the other hand, the first-order condition associated with the choice of τ_x^F is identical to that in the case without retaliation, and hence the same condition $2 + \hat{x}y'''(\hat{x})/y''(\hat{x}) > 0$ is necessary in this case as well.

We can now focus on the effect of δ on the equilibrium variables. Using equation (24) and the fact that $y'(\hat{x}) = 1 + \hat{\tau}_x^F$, we can write the first-order condition as follows:

$$\hat{x} + \frac{y'(\hat{x}) - 1}{\frac{1}{2}(1 + \delta)y''(\hat{x})} = 0.$$

This again defines \hat{x} implicitly in terms of δ . Partially differentiating this function with respect to \hat{x} and using the first-order condition to simplify, we find that the implicit function is increasing in \hat{x} provided that

$$1 + \frac{1}{\frac{1}{2}(1 + \delta)} + \frac{xy'''(\hat{x})}{y''(\hat{x})} > 0,$$

which is a condition implied by the second-order condition (since $\frac{1}{2}(1 + \delta) < 1$). On the other hand, it is clear that the implicit function is increasing in δ . Overall, we thus have that \hat{x} is now *decreasing* in δ , in sharp contrast to the case with unilateral policies. Furthermore, because $y'(\hat{x}) = 1 + \hat{\tau}_x^F$, we can also conclude that $\hat{\tau}_x^F$ is increasing in δ .

⁴²To see this, simply differentiate the expression with respect to \hat{x} , and use $\frac{1}{2}(1 + \delta)y'(\hat{x}) - 1 = \tau_x^F = \frac{1}{2}(1 + \delta)\hat{x}$ to simplify.

5.2.3 Basic Model

We next move to a discussion of the second-order conditions in the model without ex-ante lump-sum transfers. Using $\frac{1}{2}(1+\delta)(1+\hat{\tau}_1^H)y'(\hat{x}) = 1 + \hat{\tau}_x^H$, we can simplify the first-conditions to conditions to obtain:

$$\begin{aligned}\frac{\partial W^H}{\partial \tau_1^H} &= 0 = \tau_1^H \frac{\partial D_1}{\partial p_1^H} - \frac{1}{2}(1+\delta)y(\hat{x}) - [1 - y'(\hat{x})] \frac{\partial \hat{x}}{\partial \tau_1^H} \\ \frac{\partial W^H}{\partial \tau_x^H} &= 0 = \hat{x} - [1 - y'(\hat{x})] \frac{\partial \hat{x}}{\partial \tau_x^H}, \\ \frac{\partial W^F}{\partial \tau_x^F} &= 0 = \left[\frac{1}{2}(1+\delta)(1+\tau_1^H)y'(\hat{x}) - (1+\tau_x^H) \right] \frac{\partial \hat{x}}{\partial \tau_x^F}.\end{aligned}$$

It is straightforward to check that the second-order condition for the choice of τ_x^F is satisfied, i.e., $\partial^2 W^H / \partial (\tau_x^F)^2 < 0$. Conversely, the fact that we now have $\hat{\tau}_1^H \neq 0$ implies that the second-order conditions for the choice of τ_1^H and τ_x^H are quite cumbersome to characterize, as they will now also involve properties of the demand function. Throughout the paper, we simply assume that they are satisfied without providing the exact conditions needed.

We next, however, develop a particular case of our model where the second order conditions are easy to characterize and simple comparative statics can be obtained. In particular, we make the simplifying assumption that demand for the final-good is perfectly elastic, which implies that $\hat{\tau}_1^H = 0$. Under this assumption note that it is sufficient to check that $\partial^2 W^H / \partial (\tau_x^H)^2 < 0$, which requires that

$$\frac{\partial^2 W^H}{\partial (\tau_x^H)^2} = \frac{\partial \hat{x}}{\partial \tau_x^H} + y''(\hat{x}) \left(\frac{\partial \hat{x}}{\partial \tau_x^H} \right)^2 - [1 - y'(\hat{x})] \frac{\partial^2 \hat{x}}{\partial (\tau_x^H)^2} < 0.$$

We can apply the implicit function theorem to $\frac{1}{2}(1+\delta)y'(\hat{x}) = 1 + \hat{\tau}_x^H$ and simplify the above expression to:

$$\frac{\partial^2 W^H}{\partial (\tau_x^H)^2} = \frac{\partial \hat{x}}{\partial \tau_x^H} \left(1 + \frac{1}{\frac{1}{2}(1+\delta)} + \frac{\hat{x}y'''(\hat{x})}{y''(\hat{x})} \right) < 0,$$

which requires the term in parenthesis to be positive. Because $\frac{1}{2}(1+\delta) < 1$, the required assumption on $y(x)$ is weaker than before and is satisfied for $y(x) = x^\eta/\eta$ with $\eta \in (0, 1)$.

Notice next that the first-order condition $\partial W^H / \partial \tau_x^H = 0$ allows us to write

$$\hat{x} + \frac{y'(\hat{x}) - 1}{\frac{1}{2}(1+\delta)y''(\hat{x})} = 0,$$

which implicitly defines \hat{x} in terms of δ . Following the same analysis as in the benchmark model with retaliation, we can confirm again that the equilibrium value of x is a decreasing function of δ . Furthermore, since

$$\frac{1}{2}(1+\delta)y'(\hat{x}) = 1 + \hat{\tau}_x^H,$$

we can conclude that when $\delta, \hat{\tau}_x^H$ increases as well, just as in the first-best case.

5.3 Location of the Secondary Market

We consider here the possibility that the secondary market for the foreign supplier is located in the foreign country. This implies that, in the event of disagreement with the final-good producer in H , the input supplier

in F sells the inputs locally in the foreign country rather than exporting to an alternative buyer in H . There are a number of reasons to think that this possibility could be reflected in a richer model (e.g., as a result of search frictions associated with finding international partners on short notice that can be avoided with domestic matches), but rather than attempting to model these reasons explicitly we simply assume outright that there exists a secondary market in the foreign country (only) where a match with a local producer results in the production of an amount $y(x)$ of the generic good, and we return to the assumption that the supplier has all of the bargaining power at this point.

The key difference now is in the outside option of the supplier, which is no longer $\delta(1 + \tau_1^H)y(x) - (\tau^H + \tau^F)x$ but is instead $\delta(1 + \tau_1^F)y(x)$. Arguing as before, it is easy to see that the final-good producer in H now has a *stage-3* payoff of

$$\frac{1}{2} [(1 + \tau_1^H) - \delta(1 + \tau_1^F)] y(x) - \frac{1}{2}(\tau_x^H + \tau_x^F)x,$$

with the supplier in F now receiving a *stage-3* payoff of

$$\frac{1}{2} [(1 + \tau_1^H) + \delta(1 + \tau_1^F)] y(x) - \frac{1}{2}(\tau_x^H + \tau_x^F)x,$$

so that the *stage-2* choice of \hat{x} is now defined by

$$\frac{1}{2} [(1 + \tau_1^H) + \delta(1 + \tau_1^F)] y'(\hat{x}) = 1 + \frac{1}{2}(\tau_x^H + \tau_x^F), \quad (25)$$

and hence the *stage-1* payoffs of the home and foreign firm are given by

$$\begin{aligned} \pi^H &= \frac{1}{2} [(1 + \tau_1^H) - \delta(1 + \tau_1^F)] y(\hat{x}) - \frac{1}{2}(\tau_x^H + \tau_x^F)\hat{x}, \text{ and} \\ \pi^F &= \frac{1}{2} [(1 + \tau_1^H) + \delta(1 + \tau_1^F)] y(\hat{x}) - \frac{1}{2}(\tau_x^H + \tau_x^F)\hat{x} - \hat{x}. \end{aligned}$$

Anticipating that F may now have reason to alter p_1^F with its choice of τ_1^F (for reasons analogous to H 's incentive to alter p_1^H with its choice of τ_1^H) and hence affect foreign consumer surplus $CS(p_1^F)$, and noting that none (or to be precise, a measure 0) of good 1 is actually produced in F in equilibrium, home and foreign welfare are then given by

$$\begin{aligned} W^H &= CS(p_1^H) + \pi^H + \tau_1^H [D_1(p_1^H) - y(\hat{x})] + \tau_x^H \hat{x}, \text{ and} \\ W^F &= CS(p_1^F) + \pi^F + \tau_1^F D(p_1^F) + \tau_x^F \hat{x}. \end{aligned}$$

The first-order conditions that define the Nash policies $\hat{\tau}_1^{HN}$, $\hat{\tau}_x^{HN}$, $\hat{\tau}_1^{FN}$ and $\hat{\tau}_x^{FN}$ can be manipulated to yield

$$\begin{aligned} \hat{\tau}_1^{HN} &= \frac{\frac{1}{2}\hat{x}[y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}]}{-\partial D_1/\partial p_1^H}, \\ \hat{\tau}_x^{HN} &= \delta(1 + \hat{\tau}_1^{FN})y'(\hat{x}) + \hat{\tau}_1^{HN}y'(\hat{x}) - 1 - \frac{1}{2}\frac{\hat{x}}{\partial \hat{x}/\partial \tau_x^H}, \\ \hat{\tau}_1^{FN} &= \frac{\frac{1}{2}\delta\hat{x}[y'(\hat{x}) - \frac{y(\hat{x})}{\hat{x}}]}{-\partial D_1/\partial p_1^F}, \text{ and} \\ \hat{\tau}_x^{FN} &= -\frac{1}{2}\frac{\hat{x}}{\partial \hat{x}/\partial \tau_x^F}. \end{aligned}$$

Again, the expression for $\hat{\tau}_1^{HN}$ is negative, and is similar to the expression derived in the Basic Model and contained in (12). The intuition is also analogous to that in the Basic Model: the home government finds it optimal to set a negative $\hat{\tau}_1^{HN}$ as a means of shifting surplus from foreign suppliers to the home country. The dual role that $\hat{\tau}_x^{HN}$ plays in alleviating the hold-up problem and at same time transferring surplus implies again that its sign is in general ambiguous.

This extension of the model delivers more interesting implications for the Nash policies adopted by the foreign government, which no longer adopts a laissez-faire approach. First, as in the model with general ex-post bargaining power developed above, the foreign government now has an incentive to set a positive export tax on the intermediate input ($\hat{\tau}_x^{FN} > 0$), because the foreign input supplier can now pass part of this cost on to home producers by threatening not to deliver the intermediate input. The key for this is that the outside option for the supplier is not reduced one to one with τ_x^F . In the model with general ex-post bargaining power, this was ensured by $\alpha < 1$. In the present variant of the model, this is simply due to the fact that the secondary market does not involve trade flows. Second, and contrary to all of the models explored above, foreign taxes on the final good 1 can now affect the distribution of surplus between home and foreign producers. As a result, the foreign government now chooses to optimally balance the relative roles of $\hat{\tau}_x^{FN}$ and $\hat{\tau}_1^{FN}$ in extracting surplus from home firms in the same way that the home government balances $\hat{\tau}_x^{HN}$ and $\hat{\tau}_1^{HN}$ in extracting surplus from foreign firms. For the foreign government this implies the use of a foreign import tariff or export subsidy ($\hat{\tau}_1^{FN} > 0$) on the final good in order to raise p_1^F and thus improve the outside option (and bargaining position) of foreign suppliers.

Although we have shown that the location of the secondary market has implications for the Nash equilibrium values of home and foreign trade policies, it is important again to emphasize the two general features of our model that continue to hold in this extension as well. First, manipulating the above first-order conditions and applying the implicit function theorem to (25), we find

$$y'(\hat{x}) = 1 - \frac{\hat{x}}{\partial \hat{x} / \partial \tau_x^H} > 1,$$

which indicates that again, under Nash policy choices, the international hold-up problem persists and the volume of international input trade is inefficiently low as a consequence. Second, as we have indicated our model predicts the equilibrium use of taxes in the final good market and these distortions arise as a result of each country's attempts to extract bargaining surplus from firms abroad. Once again therefore, the purpose of a trade agreement remains to help governments better solve these two problems.

References

- Antràs, P. (2003), "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics*, 118:4, pp. 1375-1418.
- Antràs, Pol (2005), "Incomplete Contracts and the Product Cycle," *American Economic Review*, 95:4, pp. 1054-1073.
- Antràs, Pol, Luis Garicano and Esteban Rossi-Hansberg (2006), "Offshoring in a Knowledge Economy," *Quarterly Journal of Economics*, 121:1, pp. 31-77.
- Antràs, Pol, and Elhanan Helpman (2004), "Global Sourcing," *Journal of Political Economy* 112:3, pp. 552-580.
- Aghion, Philippe, Pol Antràs, and Elhanan Helpman (2007), "Negotiating Free Trade," *Journal of International Economics*, Vol. 73, No. 1, pp. 1-30.
- Bagwell, K. and R.W. Staiger (1999), "An Economic Theory of GATT," *American Economic Review*, Vol. 89, No. 1, pp. 215-248.
- Bagwell, K. and R.W. Staiger (2002), *The Economics of the World Trading System*, MIT Press.
- Bagwell, K. and R.W. Staiger (2006), "What Do Trade Negotiators Negotiate About? Empirical Evidence from the World Trade Organization," NBER Working Paper No. 12727.
- Baron, D. (1994), "Electoral Competition with Informed and Uninformed Voters," *American Political Science Review*, 88(1), pp. 33-47.
- Blanchard, E. (2007), "Foreign Direct Investment, Endogenous Tariffs, and Preferential Trade Agreements," working paper University of Virginia.
- Bolton, P. and M. Dewatripont (2005), *Contract Theory*, MIT Press.
- Bond, E. W. and C. Syropoulos (1996), "The Size of Trading Blocs: Market Power and World Welfare Effects," *Journal of International Economics*, 40:3, pp. 411-438.
- Bhagwati, J. and Ramaswami, V. K. (1963), "Domestic Distortions, Tariffs, and the Theory of Optimum Subsidy," *Journal of Political Economy*, 71:1, pp., 44-50.
- Campa, Jose M., and Goldberg, Linda S. "The Evolving External Orientation of Manufacturing Industries: Evidence from Four Countries." *Federal Reserve Bank of New York Economic Policy Review* 3 (July 1997): 53-81.
- Coase, Ronald H. (1937), "The Nature of the Firm," *Economica*, 4:16, pp. 386-405.
- Deardorff, Alan V. (2001), "Fragmentation in Simple Trade Models," *North American Journal of International Economics and Finance*, 12(2), p. 121-137.
- Economist, The*. "The Ins and Outs of Outing." (August 31, 1991): 54-56.
- Ethier, W.J. (2004), "Political Externalities, Nondiscrimination, and a Multilateral World," *Review of International Economics*, Vol. 12, No. 3, pp. 303-320.
- Feenstra, Robert C. "Integration of Trade and Disintegration of Production in the Global Economy." *Journal of Economic Perspectives* 12 (Autumn 1998), : 31-50.

- Feenstra, Robert C., and Gordon H. Hanson, "Foreign Investment, Outsourcing and Relative Wages," in Robert C. Feenstra, Gene M. Grossman, and Douglas A. Irwin, eds., *Political Economy of Trade Policy: Essays in Honor of Jagdish Bhagwati* (Cambridge, MA: MIT Press, 1996a), pp. 89-127.
- Feenstra, Robert C., and Hanson, Gordon H. "Globalization, Outsourcing, and Wage Inequality." *American Economic Review* 86 (May 1996b): 240-245.
- Feenstra, Robert C., and Hanson, Gordon H. (2005), "Ownership and Control in Outsourcing to China: Estimating the Property Rights Theory of the Firm," *Quarterly Journal of Economics*, 120(2): 729-762.
- Grossman, S. J., and O. D. Hart (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94:4, pp. 691-719.
- Grossman, G.M. and Helpman, E. (1994), "Protection for Sale," *American Economic Review*, Vol. 84, No. 4, pp. 833-850.
- Grossman, G.M. and E. Helpman (1996), "Electoral Competition and Special Interest Politics," *The Review of Economic Studies*, vol:63 pg:265.
- Grossman, G.M. and Helpman, E. (2002), "Integration vs. Outsourcing in Industry Equilibrium," *Quarterly Journal of Economics* 117 (1), 85-120.
- Grossman, G.M. and Helpman, E. (2005), "Outsourcing in a Global Economy," *Review of Economic Studies*, 72:1, pp. 135-159.
- Grossman, Gene M. and Esteban Rossi-Hansberg (2006), "Trading Tasks: A Simple Theory of Offshoring," mimeo Princeton University.
- Hart, Oliver and John Moore (1990), "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98:6, pp. 1119-1158.
- Hummels, David; Ishii, Jun; and Yi, Kei-Mu. "The Nature and Growth of Vertical Specialization in World Trade." *Journal of International Economics*, 54 (June 2001): 75-96.
- Johnson, H. G. (1953-54), "Optimum Tariffs and Retaliation," *Review of Economic Studies*, 21:2, pp. 142-153.
- Johnson, H. G. (1965), "Optimal Intervention in the Presence of Domestic Distortions," in R. E. Caves, P. B. Kenen, and H.G. Johnson, eds., *Trade, Growth and the Balance of Payments: Essays in Honor of Gottfried Haberler*, Amsterdam: North-Holland, pp. 3-34.
- Jones, Ronald W. Jones (2000), *Globalization and the Theory of Input Trade*, Cambridge: MIT Press.
- Kennan, J. and Riezman, R. (1988), "Do Big Countries Win Tariff Wars?," *International Economic Review*, Vol. 29, No. 1, pp. 81-85.
- Levchenko, Andrei (2007), "Institutional Quality and International Trade," *Review of Economic Studies*, 74:3, 791-819.
- McLaren, J. (1997), "Size, Sunk Costs, and Judge Bowker's Objection to Free Trade," *American Economic Review*, Vol. 87, No. 3, pp. 400-420.
- Nunn, Nathan (2007), "Relationship-Specificity, Incomplete Contracts and the Pattern of Trade," *Quarterly Journal of Economics*, Vol. 122, No. 2, pp. 569-600.

- Ornelas, E. (forthcoming), "Feasible Multilateralism and the Effects of Regionalism," *Journal of International Economics*.
- Ornelas, E. and J. Turner (forthcoming), "Trade Liberalization, Outsourcing, and the Hold Up Problem," *Journal of International Economics*.
- Ramanarayanan, Ananth (2006), "International Trade Dynamics with Intermediate Inputs," mimeo University of Minnesota.
- Williamson, Oliver E. (1985), *The Economic Institutions of Capitalism*, Free Press: New York.
- Yarbrough, B.V. and R.M. Yarbrough (1992), *Cooperation and Governance in International Trade: The Strategic Organizational Approach*, Princeton University Press.
- Yeats, Alexander J. "Just How Big Is Global Production Sharing?" In Arndt, Sven W., and Kierzkowski, Henryk eds., *Fragmentation: New Production Patterns in the World Economy*, Oxford University Press, 2001.