

Quantitative Analysis of Multi-Party Tariff Negotiations

Kyle Bagwell*

Robert W. Staiger[†]

Ali Yurukoglu^{‡§}

February 19, 2021

Abstract

We develop a model of international tariff negotiations to study the design of the institutional rules of the GATT/WTO. A key principle of the GATT/WTO is its most-favored-nation (MFN) requirement of nondiscrimination, a principle that has long been criticized for inviting free-riding behavior. We embed a multi-sector model of international trade into a model of interconnected bilateral negotiations over tariffs and assess the value of the MFN principle. Using 1990 trade flows and tariff outcomes from the Uruguay Round of GATT/WTO negotiations, we estimate the model and use it to simulate what would happen if the MFN requirement were abandoned and countries negotiated over discriminatory tariffs. We find that if tariff bargaining in the Uruguay Round had proceeded without the MFN requirement, it would have wiped out the world real income gains that MFN tariff bargaining in the Uruguay Round produced and would have instead led to a small reduction in world real income relative to the 1990 status quo.

Keywords: multilateral bargaining, tariff determination, quantitative trade

*Department of Economics, Stanford University and NBER.

[†]Department of Economics, Dartmouth College and NBER.

[‡]Graduate School of Business, Stanford University and NBER.

[§]This research was funded under NSF Grant SES-1326940. We have benefited from the detailed comments of the Editor, Aviv Nevo, and five referees, and from helpful input from Dave Donaldson, Ralph Ossa, Ulf Pehrsson and participants at many seminars and conferences. Ohyun Kwon provided outstanding research assistance.

1 Introduction

Multilateral tariff bargaining is complicated. A grand multilateral bargain, whereby all the countries of the world bargain freely over all the tariffs that affect them, would be fraught with difficulty. Decentralizing the bargaining into a collection of bilateral negotiations creates its own challenges, because bilateral agreements to reduce tariffs still affect countries that are not party to those agreements, due to spillovers that travel through the induced changes in international prices. These spillovers create an environment of bilateral bargaining with externalities, with unknown implications for efficiency.

In this paper, we evaluate a key principle of the World Trade Organization (WTO) and its predecessor GATT, a principle that was adopted in part to facilitate a decentralized approach to tariff bargaining through simultaneous bilateral bargains. That principle is the most-favored-nation (MFN) clause of GATT.¹ Article I of the GATT agreement lays out the MFN clause, a rule stipulating that the tariffs that GATT/WTO members impose on each other must be non-discriminatory. We conduct a quantitative evaluation of the welfare effects of the MFN principle embodied in GATT/WTO tariff negotiations.

To evaluate the MFN principle, we combine a quantitative trade model based on Caliendo and Parro (2015) with a model of interconnected bilateral tariff negotiations based on Horn and Wolinsky (1988). We estimate our model using 1990 as a benchmark year together with the outcomes of MFN negotiations from the GATT Uruguay Round (1986-1994), the largest and last-completed GATT/WTO multilateral negotiating round. We then counterfactually simulate discriminatory bilateral negotiations to compare welfare outcomes across discriminatory and non-discriminatory MFN bargaining.

For the Uruguay Round setting, our main finding is that MFN negotiations lead to better aggregate outcomes than do counterfactual discriminatory negotiations. According to our estimates, all countries gained under the MFN tariff negotiations in the Uruguay Round relative to the 1990 status quo. We find that such gains would not have occurred, however, if the Uruguay Round had proceeded without the MFN requirement. Indeed, our estimates indicate that the Uruguay Round without the MFN requirement would have

¹See Bagwell et al. (2020a) for a description of the decentralized approach to tariff bargaining engineered by GATT and the role that the MFN clause played in facilitating it. This decentralized approach was featured in the first five GATT rounds of multilateral tariff negotiations. It was used as a complement to multilateral bargaining methods in the last three GATT rounds, as well as in the now-suspended WTO Doha Round. Among the last three GATT rounds, the Uruguay Round, for example, employed multilateral bargaining methods that included “zero-for-zero” tariff commitments in specific sectors.

led to a small *reduction* in world real income relative to the 1990 status quo.

To explain the economics behind this finding, we first show that there are positive externalities associated with MFN tariff bargaining. These positive externalities are driven by the increased market access and more favorable international prices (“terms of trade”) that third-party countries enjoy when the tariffs in a bilateral bargain are liberalized on a non-discriminatory basis. We show that, as a consequence of these positive externalities and the “free-riding” that they invite, countries do not liberalize enough to reach the efficiency frontier. We find that the Uruguay Round achieved roughly one fifth of the 0.34% gains in world real income relative to the 1990 status quo that would have been enjoyed had the negotiations been able to achieve their world-income maximizing potential.

We then show that abandoning MFN and engaging in discriminatory tariff bargaining instead generates negative externalities, due to the reduction in market access and less favorable terms of trade for third-party countries that these discriminatory tariff cuts imply. The resulting bargaining equilibrium exhibits two sources of inefficiency – one that is associated with the over-liberalization of tariffs induced by these negative externalities, and one that is associated with the resulting tariff discrimination itself.

The quantitative analysis that we develop here provides answers that neither theory alone nor existing quantitative models can provide. Theory can offer a guide to the implications of different sets of rules for the outcomes of tariff bargaining, but theory alone cannot quantify the welfare implications of different rules.² Ossa (2014) and Ossa (2016) initiate the study of non-cooperative and cooperative tariffs in a multi-country quantitative trade model. Ossa’s papers compute Nash equilibrium tariffs and fully cooperative tariffs. Other papers in this literature include Caliendo and Parro (2015), Spearot (2016) and Caliendo et al. (2017). None of these papers attempt to model the GATT/WTO tariff bargaining system as a nexus of bilateral negotiations to compare outcomes under MFN and counterfactual discriminatory negotiations.

The Caliendo and Parro (2015) quantitative trade model on which we build incorporates intermediate goods into the Ricardian models of Eaton and Kortum (2002) and Costinot et al. (2011). To model bilateral tariff bargaining in this environment, we adopt

²Bagwell et al. (2020b) develop a theoretical analysis of equilibrium bilateral tariff bargaining in a three-country trade model and show that, due to the nature of the externalities associated with discriminatory tariffs, tariff bargaining without MFN always results in inefficient over-liberalization. In their working paper version, Bagwell et al. (2018) show that tariff bargaining under an MFN rule typically leads to inefficient outcomes that can exhibit either over- or under-liberalization. See also Bagwell and Staiger (2005, 1999a, 2018) on the implications of different sets of rules for the outcomes of tariff bargaining.

the solution concept of Horn and Wolinsky (1988). This solution concept is commonly employed by the Industrial Organization literature to characterize the division of surplus in bilateral oligopoly settings with externalities across firms and agreements. The concept is referred to as a “Nash-in-Nash” solution, because it can be thought of as a Nash equilibrium between separate bilateral Nash bargains.³ As Bagwell et al. (2020b) discuss, the Nash-in-Nash approach is not without limitations when applied to tariff bargaining, but it does offer a simple characterization of simultaneous bilateral bargaining outcomes in settings with interdependent payoffs. The approach thereby makes the analysis of bilateral tariff bargaining in the GATT/WTO context tractable in a quantitative trade model.

The data underlying estimation of the trade model consist of trade flows, production, and tariffs at the country-sector level – aggregated into 49 traded and 18 non-traded sectors for six key countries, with the rest of the world aggregated into five additional regions – together with data on input-output flows between sectors and a set of gravity variables. We estimate the taste, productivity, iceberg cost, and trade deficit parameters to match 1990 trade shares by country-sector and observed trade deficit levels subject to satisfying equilibrium conditions of the model.

Our estimated trade model determines welfare payoffs for each country given a set of tariffs – observed or counterfactual – where we take the welfare of each country to be defined by its real national income. We therefore abstract from the political economy and distributional concerns that permeate the trade policy choices of real world governments. We adopt these abstractions because we are interested in determining how well the Uruguay Round tariff outcomes can be understood as resulting from the choices of real-income-maximizing governments, and also because introducing political economy or distributional concerns into the Caliendo and Parro (2015) model is not straightforward.⁴ In focusing on tariffs, we are also abstracting from the broader set of issues that formed

³The Nash-in-Nash solution concept has been used by Crawford and Yurukoglu (2012) and by Crawford et al. (2018) to explore negotiations between cable television distributors and content creators, and by Grennan (2013), Gowrisankaran et al. (2015), and Ho and Lee (2017) to consider negotiations in health care. It is broadly related to the pairwise-proof requirements that are indirectly implied under the requirement of passive beliefs in vertical contracting models (Hart and Tirole (1990)) and directly imposed in contracting equilibria (Cremer and Riordan, 1987). See McAfee and Schwartz (1994) for further discussion. Micro-foundations for the Nash-in-Nash approach are developed by Collard-Wexler et al. (2019) in the context of negotiations that concern bilateral surplus division.

⁴Ossa (2014) introduces political economy into a quantitative trade model, but unlike the model of Caliendo and Parro (2015), Ossa’s model does not feature free entry and allows political economy interests to focus on the rents earned by fixed factors.

the entirety of the Uruguay Round negotiating agenda. These limitations on the scope of our analysis should be kept in mind when interpreting our results.

To provide context, we first use the estimated trade model to generate a series of benchmark counterfactual outcomes, including welfare under autarky and in the absence of trade frictions, as well as under Nash tariffs and the tariffs that would maximize world real income. We find that, according to our calculations and beginning from Nash tariffs, the GATT rounds up to but not including the Uruguay Round had achieved by 1990 roughly 60 percent of the aggregate world-wide real income gains that were possible with changes to the tariffs under negotiation in the Uruguay Round.

We then turn to our main task and use the estimated trade model to analyze bilateral tariff negotiations between countries. We first calculate the Horn-Wolinsky bargaining solution beginning from the 1990 tariffs under three institutional constraints reflected in the tariff-bargaining environment of the Uruguay Round, namely, that countries (i) are restricted to bargain over MFN tariffs, (ii) must respect existing GATT tariff commitments and not raise their tariffs, and (iii) abide by the “principal supplier rule,” which guides each importing country to limit its negotiations on a given product to the exporting country that is the largest supplier of that product to its market. We estimate the bargaining parameters such that the bargaining solution best fits the MFN tariff changes between 1990 and 2000, our measure of the tariff bargaining outcomes of the Uruguay Round.

We find that all countries gained from the Uruguay Round relative to the 1990 status quo according to our model predictions. This result is not guaranteed. First, according to our model, some countries do not bargain in the Uruguay Round. Second, under the Nash-in-Nash bargaining solution, countries that do bargain are guaranteed only to do better than their disagreement point, but this point is impacted by the externalities across bargaining pairs and can thus move relative to the 1990 status quo. In terms of bargaining power, we find evidence of asymmetric bargaining positions across bilaterals. We describe how our bargaining power estimates reflect, for each bilateral, an interplay between the relative depth of the tariff cuts that each country agrees to in the bilateral and the slope of the bilateral bargaining frontier, where this slope reflects the details of each country’s position in our model world economy.

Armed with our trade-model and bargaining-power parameters, we then simulate the counterfactual discriminatory bargaining protocol. We find that on a trade-weighted basis average tariffs drop further under discriminatory negotiations than under MFN negotiations, as expected given that discriminatory negotiations do not suffer from the

free-riding behavior that MFN invites. But we find that the liberalizing force unleashed by the abandonment of MFN is *overly* strong and that, together with the cost of the trade diversion implied by discriminatory tariffs, it would have wiped out the world real income gains produced by MFN tariff bargaining in the Uruguay Round and would have instead led to a 0.02% reduction in world real income relative to the 1990 status quo itself.

Our findings also illustrate the dampening impact that MFN has on the expression of bargaining power. We find that the strongest countries would have the most to gain and the weakest the most to lose from the abandonment of MFN. And we demonstrate that it is only when the full equilibrium consequences of the abandonment of MFN are taken into account that the true losses become apparent: if each of the bargaining countries were to consider its discriminatory bilaterals in isolation, that is, as a collection of single-pair discriminatory bargains, abandoning MFN would look like a winning proposition.

More broadly, our quantitative results lead to an important conclusion: while the free-rider issue and associated drag on tariff liberalization created by the positive third-party externality from MFN is widely emphasized as a shortcoming of the GATT/WTO approach, we find for the Uruguay Round that abandoning MFN in tariff bargaining would create negative third-party externalities that are even more powerful, and that would have ultimately led to tariff bargaining outcomes that are worse for world welfare.

2 Model

Our model world economy consists of the multi-sector version of Eaton and Kortum (2002) with intermediate goods following Caliendo and Parro (2015).⁵ This model world economy is then embedded into an equilibrium model of tariff bargaining.

2.1 Model World Economy

We consider a world economy with $i = 1, \dots, N$ countries and $k = 1, \dots, K$ sectors, K_{traded} of which are tradeable. We allow each country to impose an import tariff (possibly discriminatory across trading partners) in each traded sector k . Because our model world economy follows Caliendo and Parro (2015), we provide only a minimal description here, and refer readers to that paper for additional model details.

⁵See also Costinot et al. (2011) for a multi-sector version of Eaton and Kortum (2002) without intermediate goods

Each country i has an immobile-across-countries labor endowment L_i . Within each sector k , there is a continuum of intermediate varieties indexed by ω^k . Intermediate varieties are produced using labor and composite goods from all sectors with a constant-returns-to-scale technology. Composite goods are aggregates of intermediate varieties, and are used for production in all sectors and as final products for consumption. Furthermore, an infinite number of firms, all with the same productivity, exist to produce each variety in each sector, ensuring perfect competition.

The technology for each variety is drawn from a Fréchet distribution with CDF

$$F_i^k(z) = \exp\left(-\left(\frac{z}{z_i^k}\right)^{-\theta^k}\right) \quad (1)$$

where z_i^k is country i 's sector- k level productivity parameter and θ^k is a sector-specific productivity shape parameter. We reference specific draws from these distributions as $z_i^k(\omega^k)$, that is, country i 's productivity for variety ω^k in sector k . While the first and second moments of the distribution of productivity depend on both the z and the θ parameters, the ratio of expected variety productivity for the same sector between two countries is equal to the ratio of their z^k parameters in sector k . Higher values of θ^k imply lower heterogeneity in within-sector productivity, and hence more responsiveness of trade flows with respect to changes in fundamentals (and hence higher trade elasticities).

The production function for each intermediate variety ω^k is:

$$q_i^k(\omega^k) = z_i^k(\omega^k)(l_i^k(\omega^k))^{\gamma_i^k} \prod_{n=1}^K m_i^{k,n}(\omega^k)^{\gamma_i^{k,n}} \quad (2)$$

where $l_i^k(\omega^k)$ and $m_i^{k,n}(\omega^k)$ are respectively the labor and the composite intermediate goods from sector n used for the production of intermediate variety ω^k in country i . These varieties are aggregated within sector into composite goods for sector k . The parameter $\gamma_i^{k,n}$ is the share of the composite good from sector n required to produce in sector k , with labor's share given by $\gamma_i^k = 1 - \sum_n \gamma_i^{k,n}$. Composite good producers in country i 's sector k aggregate intermediate varieties as follows:

$$Q_i^k = \left(\int r_i^k(\omega^k)^{1-\frac{1}{\sigma^k}} d\omega^k \right)^{\frac{\sigma^k}{\sigma^k-1}} \quad (3)$$

where $\sigma^k > 0$ is the elasticity of substitution between intermediate goods within sector k with $1 + \theta^k > \sigma^k$ and where $r_i^k(\omega^k)$ is the demand for intermediate variety ω^k .

Producers face iceberg trading costs and potentially tariffs when serving other coun-

tries, with $d_{ji}^k \geq 1$ denoting the iceberg trade cost for country j 's sector- k exports to country i and with $d_{ii}^k = 1, \forall k$ and $d_{ji}^k = \infty$ for non-traded goods. With perfect competition in each country-sector-variety, country i 's price of variety ω^k in sector k is found by solving country i 's sourcing decision, and is equal to:

$$p_i^k(\omega^k) = \min_{j \in \{1, \dots, N\}} \frac{c_j^k}{z_j^k(\omega^k)} \kappa_{ji}^k \quad (4)$$

where c_j^k is country j 's cost of a sector- k input bundle and $\kappa_{ji}^k \equiv d_{ji}^k(1 + t_{ji}^k)$ with t_{ji}^k equal to the ad valorem tariff levied by country i on sector- k imports from country j .⁶ When $t_{ji}^k \neq t_{hi}^k$, the tariffs that country i applies to imports of k from countries j and h are discriminatory; an MFN tariff is nondiscriminatory and satisfies $t_{ji}^k = t_{hi}^k \equiv t_{MFN,i}^k$.

Turning to the demand side of the model, a representative consumer in each country chooses consumption levels of each composite good in each sector to maximize the following utility function that is Cobb-Douglas across sectors:

$$u_i = \prod_{k=1}^K C_i^k \alpha_i^k \quad (5)$$

where C_i^k is country i 's consumption of the composite good produced in sector k and α_i^k is country i 's taste parameter for sector k . Consumers take prices for each composite good as given. They choose consumption to maximize (5) subject to their budget constraint that total expenditure must be weakly less than their country's labor income plus tariff revenue plus an endowment which represents observed trade deficits.⁷

An equilibrium consists of a vector of wages \mathbf{w} and a vector of sector- and country-level price indexes \mathbf{P} such that all markets clear and consumers and firms behave optimally.

Finally, we discuss the role in the model of the triangle inequality, which states that $\kappa_{ji}^k \kappa_{ih}^k \geq \kappa_{jh}^k$ for all i, j, h and k or, using the definition of κ_{ji}^k , that $d_{ji}^k(1 + t_{ji}^k) \times d_{ih}^k(1 + t_{ih}^k) \geq d_{jh}^k(1 + t_{jh}^k)$. For $i \neq j \neq h$ the triangle inequality involves trade impediments among three countries and describes a relationship that is "baked in" to the model via the

⁶More specifically, $c_j^k = \Upsilon_j^k w_j^{\gamma_j^k} \prod_{n=1}^K P_j^n \gamma_j^{k,n}$ where w_j is the wage of labor in country j , $P_j^n \equiv A^n [\sum_{i=1}^N z_i^n (c_i^n \kappa_{ij}^n)^{-\theta^n}]^{-1/\theta^n}$ is the price in country j of the composite intermediate good from sector n , and where A^n and Υ_j^k are constants. Also, with this specification we are assuming that the ad valorem tariff is applied to the delivered price of the import good at the importing country's border.

⁷We do not model the forces underlying the existence of trade deficits. As discussed by Ossa (2016), this raises the question of how to handle trade deficits in counterfactual analysis. For reasons described by Ossa, we set all trade deficits to zero before performing any of our counterfactuals.

specification of the sourcing decision in (4), which assumes that only direct trade routes need be considered when searching for the lowest cost supplier to a market. In principle discriminatory tariffs (with $t_{ih}^k < t_{jh}^k$) could be an important source of violations of the triangle inequality. But the presence of such tariffs does not invalidate the model provided that it is assumed that rules of origin are put in place to prevent third parties from exploiting the implied roundabout arbitrage opportunities.

For $i \neq j = h$ the triangle inequality becomes $\kappa_{ji}^k \kappa_{ij}^k \geq 1$ and describes a “round-trip” inequality between any two countries i and j , $d_{ji}^k(1 + t_{ji}^k) \times d_{ij}^k(1 + t_{ij}^k) \geq 1$, or equivalently

$$(1 + t_{ji}^k) \times (1 + t_{ij}^k) \geq \frac{1}{d_{ji}^k d_{ij}^k} \quad (6)$$

for all i, j and k , a condition that is assured as long as $t_{ji}^k \geq 0$ and $t_{ij}^k \geq 0$ (no import subsidies) which holds in our 1990 data. However, in our counterfactuals we will consider bilateral bargaining over discriminatory tariffs and in that setting import subsidies may arise that potentially could violate the round-trip inequality in (6). Hence, we must consider the model’s implications when the round-trip inequality is violated and determine whether special handling of tariffs that would violate this inequality is warranted.

When tariffs violate (6), if one of these countries exports variety ω^k to the other, then the price of ω^k in the exporting country must be higher than the cost of purchasing ω^k in the importing country and delivering it to the exporting country, creating a profit opportunity for traders in the importing country to re-export ω^k back to the exporting country.⁸ Since ω^k is not produced in a third country by assumption, rules of origin cannot prevent this arbitrage activity. Moreover, if the round-trip inequality is violated even slightly and round-trip arbitrage were allowed, the export and re-export would not stop until the price of ω^k in the two countries was driven to zero. Finally, the model’s specification of the sourcing decision in (4) does not bake in the exclusion of this round-trip sourcing activity. With traded intermediate goods, the recursive nature of production allows imported varieties to in effect make multiple round trips across the border, generating a manifestation of these arbitrage opportunities in model outcomes when (6) is violated.

For these reasons, in our discriminatory tariff bargaining counterfactuals we will limit

⁸Suppose, for example, that country 1 exports good 2 to country 2 and that (6) is violated so that $\kappa_{21}^2 \kappa_{12}^2 < 1$. If the price of good 2 in country 1 is p_1^2 , then the price of good 2 in country 2 must be $p_2^2 = \kappa_{12}^2 p_1^2$, and the cost of purchasing good 2 in country 2 and delivering it back to country 1 is then $\kappa_{21}^2 p_2^2 = \kappa_{21}^2 \kappa_{12}^2 p_1^2 < p_1^2$.

the tariff levels that can be considered by a bargain between countries i and j to those that satisfy the round-trip inequality in (6), under the assumption that the cost of preventing the round-trip arbitrage that would be possible if (6) were violated is prohibitive.

2.2 Tariff Bargaining

We assume that in a round of tariff negotiations, countries negotiate bilaterally and simultaneously over tariff vectors. This bargaining structure was featured in the first five GATT rounds of multilateral tariff negotiations, and it was used as a complement to multilateral bargaining methods in the last three GATT rounds, including the Uruguay Round, as well as in the now-suspended WTO Doha Round.

As all tariffs affect all countries through the model's trade equilibrium, the payoffs from each bilateral negotiation depend on the outcomes of the other bilateral negotiations. We follow Bagwell et al. (2020b) and apply the solution concept of Horn and Wolinsky (1988) to this tariff bargaining problem. According to this solution, each pair of negotiating countries maximizes its Nash product given the actions of the other pairs.

To formalize the Horn and Wolinsky (1988) bargaining solution in our setting, we let $\pi_i(\mathbf{t})$ be the welfare of country i when the world vector of tariffs is given by \mathbf{t} . We measure a country's welfare by the maximized utility of its representative citizen as defined in (5), and $\pi_i(\mathbf{t})$ simply reflects the fact that consumption is pinned down once \mathbf{t} is given. When country i negotiates with county j , the two countries select levels for the vector of tariffs \mathbf{t}_{ij} which are the subject of their negotiations so as to maximize their Nash product:

$$\begin{aligned} \max_{\mathbf{t}_{ij}} \quad & (\pi_i(\mathbf{t}_{ij}, \mathbf{t}_{-ij}) - \pi_i(\mathbf{t}_{ij}^0, \mathbf{t}_{-ij}))^{\zeta_{ij}} (\pi_j(\mathbf{t}_{ij}, \mathbf{t}_{-ij}) - \pi_j(\mathbf{t}_{ij}^0, \mathbf{t}_{-ij}))^{1-\zeta_{ij}} \\ \text{s.t.} \quad & \pi_i(\mathbf{t}_{ij}, \mathbf{t}_{-ij}) - \pi_i(\mathbf{t}_{ij}^0, \mathbf{t}_{-ij}) \geq 0 \\ & \pi_j(\mathbf{t}_{ij}, \mathbf{t}_{-ij}) - \pi_j(\mathbf{t}_{ij}^0, \mathbf{t}_{-ij}) \geq 0 \end{aligned} \tag{7}$$

where ζ_{ij} is the bargaining power parameter of country i in its bilateral bargain with country j and where we have partitioned the vector of tariffs \mathbf{t} into those that are the subject of negotiation between i and j (the vector \mathbf{t}_{ij}) and all other tariffs (the vector \mathbf{t}_{-ij}). The vector \mathbf{t}_{ij}^0 represents the level for the tariffs under negotiation between i and j that will prevail if i and j fail to reach an agreement. We set \mathbf{t}_{ij}^0 to be the levels of these tariffs in place when the negotiating parties entered the round.

We now define the Horn and Wolinsky (1988) tariff bargaining equilibrium:

Definition (Tariff Bargaining Equilibrium). *An equilibrium in tariffs consists of a vector*

of tariffs \mathbf{t}^{HW} such that for each pair ij the tariffs \mathbf{t}_{ij}^{HW} negotiated by this pair solves the program in (7) when $\mathbf{t}_{-ij} = \mathbf{t}_{-ij}^{HW}$.

The key assumption in the Horn and Wolinsky (1988) bargaining equilibrium is that, when evaluating a candidate \mathbf{t} , the pair ij holds the vector \mathbf{t}_{-ij} fixed. In other words, if ij were to not reach agreement, or were to deviate from a tariff vector specified by the equilibrium, then the other tariffs do not adjust. This equilibrium notion is sometimes referred to as “Nash-in-Nash,” because it is the Nash equilibrium to the synthetic game where each pair constitutes a player, the payoff function is the pair’s Nash bargaining product, and the strategies of each player are the tariffs being negotiated by the pair associated with that player.⁹

To reflect the tariff bargaining environment of the Uruguay Round, we introduce three institutional constraints to our tariff bargaining solution. First, we assume that countries are restricted to bargain over MFN tariffs and cannot engage in bilateral bargains over discriminatory tariffs.¹⁰ Second, we assume that countries honor their pre-existing GATT commitments and are not allowed to raise tariffs above their initial levels. And third, we assume that only the largest supplier of good k to country i prior to the round – the “principal supplier of good k to country i ” in GATT parlance – can negotiate with country i over $t_{MFN,i}^k$, country i ’s MFN tariff in sector k .¹¹ We also assume that international transfers are unavailable to countries in the context of their tariff bargaining.¹² In the absence of transfers, the principal supplier rule implies that a “double coincidence of wants” must exist between any viable pair of bargaining partners, in the sense that each country in the bargaining pair must be a principal supplier of at least one good to the other

⁹The described synthetic tariff-bargaining game is a “generalized game,” since, due to participation constraints, the feasible strategy set for one player is affected by the choices of other players. For further discussion, see Bagwell et al. (2020b).

¹⁰GATT members engage in bilateral bargains over discriminatory tariffs when they negotiate preferential trade agreements, which under GATT/WTO rules are permissible under certain conditions. But in the Uruguay Round, negotiations were restricted to MFN tariffs.

¹¹In their examination of the bargaining data from the GATT Torquay Round, Bagwell et al. (2020a) find that the average number of exporting countries bargaining with an importing country over a given tariff was 1.25, suggesting that our assumption is a reasonable approximation. A potential caveat is that the findings of Bagwell et al. (2020a) apply at the 6-digit HS level of trade, whereas here we are operating at a more aggregate sectoral level; we return to this point in the Conclusion.

¹²We thus abstract from various economic issues at play during the Uruguay Round that went beyond negotiated tariff cuts (e.g., negotiated agreements on intellectual property, agriculture and services) and that may have played a role of imperfect international transfer mechanisms.

country in the pair. Using the trade patterns predicted by our model, this requirement defines the set of Uruguay Round bilateral tariff bargains in our model economy.¹³

3 Data

To operationalize our model, we use sector-level data for each country on trade flows, production, input-output shares, value added, and tariffs, together with country level population and pairwise distances between country. Details of the data cleaning and aggregation are contained in Appendix A. Table 1 provides summary statistics.

Table 1: Summary Statistics

Country	Pop(M)	Import ratio	Relative Value Added per Capita	1990 Average Tariffs	1990 Trade Weighted Tariffs	2000 Average Tariffs	2000 Trade Weighted Tariffs	Largest Trading Partner
US	249.623	0.191	1.000	0.045	0.048	0.032	0.043	EU
EU	364.544	0.163	0.743	0.065	0.058	0.041	0.033	Eur. NES
Japan	123.537	0.122	1.363	0.053	0.027	0.035	0.019	US
South Korea	42.869	0.196	0.440	0.110	0.090	0.083	0.049	US
Australia	17.065	0.156	0.598	0.101	0.102	0.043	0.048	Japan
Canada	27.791	0.339	0.737	0.080	0.081	0.041	0.030	US
Africa NES	485.009	0.047	0.011	0.152	0.137	0.114	0.105	EU
America NES	428.338	0.060	0.109	0.129	0.115	0.128	0.107	US
Asia NES	2669.638	0.101	0.018	0.150	0.136	0.083	0.070	EU
Eur. NES	364.884	0.171	0.095	0.089	0.074	0.071	0.056	EU
MENA NES	463.298	0.141	0.031	0.156	0.142	0.150	0.114	EU

Notes: Average tariffs are computed for non-agricultural sectors.

We partition the world economy of 1990 into six countries and five regional entities, and beginning with SITC2 two-digit codes we aggregate trade flows into 49 traded sectors listed in Appendix Table A1, with an additional 18 non-traded sectors which produce inputs into the production of traded goods listed in Appendix Table A2, resulting in 67 sectors. The six countries are the United States (US), the European Union (EU) which for purposes of tariff negotiations is treated in the GATT/WTO as a single country, Japan,

¹³Omitted from the institutional constraints that we impose on tariff bargaining is the GATT/WTO norm of reciprocity. Reciprocity refers to a negotiation norm under which each country experiences an increase in export volume that is the same value as the increase in its import volume. In Section 7 we use our results to explore the extent to which the Horn-Wolinsky bargaining outcomes deviate from a form of reciprocity, and in the Conclusion we also discuss the possibility of augmenting our representation of the tariff bargaining protocol to include a reciprocity norm.

South Korea, Australia and Canada.¹⁴ The rest of the world is aggregated into five “not-elsewhere-specified” (NES) regional entities. Our six countries were the six largest by GDP in 1990 and all GATT members, and are the focus of our tariff bargaining analysis.

3.1 Trade, Production, Input-Output, and Value Added Data

The starting point for our data is the NBER world trade flows data from Feenstra et al. (2005) for the year 1990. We compute the gross value in 1990 dollars of each country’s bilateral imports at the sector level according to our country and sector definitions. The NBER data do not provide information on production or consumption. We impute each country’s sector-level production by extracting the ratio of exports to total production at the country-sector level from the Global Trade Analysis Project (GTAP) database (Dimaranan and McDougall (2001)).

Turning to input-output data, for the share of intermediate goods in each sector we face a tradeoff between maintaining a high level of sectoral detail and allowing country-level variation in the input-output coefficients. Country-level input-output tables are available for many countries at a high level of sectoral detail, but sectoral classification systems differ widely by country. For our 1990 base year, the Eora26 database (Lenzen et al. (2012, 2013)) contains input-output tables aggregated to a common 26-sector classification for a large number of countries. To our knowledge, this is the data source for 1990 input-output tables that provides the greatest level of sectoral detail harmonized across a large number of countries. However, for purposes of capturing the externalities associated with bilateral tariff bargaining, we place a high value on achieving granularity in sectoral trade flows. For this reason, we opt for a high level of sectoral detail in the input-output coefficients at the expense of country-level variation in these coefficients, and use the Eora purchaser-price measured input-output table for the United States in 1990, aggregated at the sectoral level to our 49 traded and 18 non-traded goods, under the assumption that the share of intermediate goods in each sector implied by the 1990 US input-output table applies to all countries and regional entities in our model world economy.

¹⁴We define the EU as consisting of the 12 members of the European Economic Community in 1990 – Belgium, Denmark, France, (West) Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain and the United Kingdom – plus Austria, Finland and Sweden, who were simultaneously negotiating in the Uruguay Round and for accession in the EU and became EU members on January 1 1995. Our decision to include these three countries in the EU is made for simplicity.

3.2 Tariff Data

We obtain country-sector tariff equivalent applied MFN *ad valorem* tariffs from the UNCTAD TRAINS database on tariffs for 1990 and 2000. We use the 1990 applied tariffs as the pre-Uruguay Round MFN tariffs, and the 2000 applied tariffs as the negotiated outcomes from the Uruguay Round. We assume that any preferential trade agreements (PTAs) that were in place in 1990 are not impacted by GATT tariff negotiations, and we include those tariff preferences and hold them fixed throughout our analysis.

There is an important distinction between the tariffs that countries actually apply to imports into their markets, and the tariff bindings that they negotiate in the GATT/WTO.¹⁵ While introducing this distinction into a quantitative trade model would be a worthwhile project in its own right, it is well beyond the scope and focus of our paper. In addition, the results of GATT/WTO tariff negotiating rounds are typically phased in over an implementation period that can last years. In the Uruguay Round, phase-in periods ranged across countries and sectors up to a maximum of roughly a decade. With the implementation period of the Uruguay Round commencing on January 1 1995, our decision to use the difference between the applied MFN tariffs in place in 1990 and the applied MFN tariffs in place in 2000 as a measure of the negotiating outcomes of the round represents an attempt to capture these complex features in a way that maintains the tractability of our quantitative model and its use for studying tariff bargaining.

Finally, while we estimate the parameters of our trade model utilizing data on trade flows, production, input-output shares, value added, and tariffs for the full coverage of products, for our bargaining analysis we focus attention on bargaining over tariffs for non-agricultural products (product categories 10-11 and 13-49 as defined in Table A1).¹⁶

¹⁵A tariff binding represents a legal cap on the tariff that a country agrees not to exceed when it applies its tariff. The tariff it applies may be at or below the cap. For most industrialized countries, the vast majority of applied tariffs are at the cap, but for many emerging and especially developing countries, applied tariffs are often well below the cap.

¹⁶The reason for not analyzing bargaining over agricultural tariff changes is twofold. First, much of the focus of negotiations on agriculture in the Uruguay Round was on domestic supports such as domestic subsidy and price support programs, and our model is not well-suited to analyze negotiations over such policies. And second, many of the agricultural tariffs that were under negotiation in the Uruguay Round were defined in specific rather than ad-valorem terms. To operationalize the model, we require ad valorem tariffs. However, ad valorem equivalents of specific agricultural tariffs display large fluctuations in levels between 1990 and 2000 due to world price movements rather than changes in specific tariffs, further complicating the analysis of negotiated agricultural tariff changes.

3.3 Gravity Data

We use data on distances between countries and preferential trade agreements from the Geography module of the CEPII Gravity Dataset (Head and Mayer, 2014). This data set constructs distances between countries based on distances between pairs of large cities and the population shares of those cities. For the regional entities, we construct the distance with a partner as the population-weighted average distance between the countries forming the regional entity and the partner in question. For two regional entities, we use the population-weighted average distance across all pairs between the two regional entities.

4 Estimation

We estimate the model in two steps. First, we estimate the taste, production, iceberg cost, and trade deficit parameters of our trade model. Given these estimates, we then estimate the bargaining parameters. We split the estimation process into two steps because the bargaining model is computationally much more intensive than the trade model, as solving the bargaining model once involves potentially thousands of computations of a trade equilibrium at differing tariff levels. Because the trade model has roughly seven hundred parameters, joint estimation with the bargaining model is prohibitively expensive. For feasibility, we thus sacrifice some efficiency by not jointly estimating the trade and bargaining parameters.

4.1 Estimation of Trade Model Parameters

We estimate the trade model to minimize the distance between the data and model predictions, focusing on two categories of predictions: (1) trade shares – the ratio of each country’s bilateral imports in each sector to the country’s total consumption in that sector and (2) trade deficits – the ratio of trade deficits to tariff revenue. Our focus on the ratio of trade deficits to tariff revenue eliminates units from consideration and thereby avoids the mismatch in units between the model (which uses labor in the US as the numeraire) and the data (where trade deficits are measured in dollars).

We parameterize iceberg costs to depend on an origin effect, a destination effect, a sector-specific border effect, and a sector-specific distance effect. It is often noted that the so-called “Quad” countries of the US, the EU, Canada and Japan had an outsized impact on the shape of the Uruguay Round due to their status as major traders and

special trading relationships with each other. We attempt to capture this with inclusion of effects, each common across sectors, for shipments between each of the Quad-country pairs. For $i \neq j$, our parameterization of iceberg trade costs is given by:

$$d_{ji}^k = 1 + \exp \left(\beta_{0,j} + \beta_{0,i} + \beta_0^k + \beta_d^k \text{dist}_{ji} + \sum_{n \in Q} \beta_{q,n} \text{Quad}_{n,ji} \right).$$

The parameters $\beta_{0,j}$, $\beta_{0,i}$, and β_0^k represent fixed effects for origin, destination, and products, respectively. The variable dist_{ji} is the distance between countries j and i , and Q is the set of pairs of the members of the Quad, with $\text{Quad}_{n,ji}$ equal to one whenever countries j and i make up the pair n .

More specifically, the model parameters to estimate consist of a vector of taste parameters ($\boldsymbol{\alpha}$), a vector of input-output shares ($\boldsymbol{\gamma}$), a vector of productivity parameters (\mathbf{z}), a vector of dispersion of productivity parameters ($\boldsymbol{\theta}$), a vector of iceberg cost parameters ($\boldsymbol{\beta}$), and a vector of trade deficits \mathbf{D} . We estimate $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ “off-line” as these are implied by the observed input-output shares, total production by country-sector, and income for each country. We set wages equal to the observed value added per capita of each country relative to the value added per capita of the US, and require that the model parameters satisfy the market clearing conditions for equilibrium at these wages.

We now define

$$G(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D}) = \begin{bmatrix} \frac{\hat{x}_{ij}^k(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D})}{\sum_i \hat{x}_{ij}^k(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D})} - \frac{x_{ij}^k}{\sum_i x_{ij}^k} \\ \hat{b}_i(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D}) - b_i \end{bmatrix}$$

where x_{ij}^k is country i 's exports of sector- k goods to country j , b_i is the ratio of country i 's trade deficit to its tariff revenue, and where a hat over a variable denotes the model's prediction for that variable. Given the $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ estimates, we then choose the remaining parameters to solve the following optimization problem:

$$\begin{aligned} \min_{\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D}} \quad & G(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D})' W G(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D}) + \lambda \sum_{\delta_p \in P} \delta_p^2 \\ \text{s.t.} \quad & F(\mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{D}, \mathbf{w}) = 0 \end{aligned} \tag{8}$$

where W is a weighting matrix and the set P is the set of terms δ_p subject to regularization. We weight all trade shares equally to each other, and weight each deficit ratio

difference by 0.05.¹⁷ The function F consists of the market clearing conditions given the model parameters and wages. We augment the loss function $G(\cdot)'WG(\cdot)$ with a regularization term that penalizes certain parameters deviating from zero. We use the L^2 norm for penalizing, thus employing a non-linear analog of ridge regression (Dagenais, 1983; Caner and Zhang, 2014). The purpose of the ridge regularization is to avoid co-linearity issues from the large number of parameters.¹⁸ We set the penalty weight λ at the lowest level in increments of 0.0001 that ensures smooth convergence of the estimation problem. Specifically, we penalize heterogeneity in the iceberg cost parameters by regularizing $\beta_{0,j}$ and $\beta_{0,i}$ towards their means across countries, β_0^k and β_d^k towards their means across sectors, the $\beta_{q,n}$ parameters towards zero, and the θ^k parameters towards the mean θ . We estimate standard errors by a residual bootstrap procedure that treats deviations from the model's predicted country-sector level import shares as measurement error.

Embedded inside each evaluation of the estimation objective function is the solution to the trade model, which evaluates trade flows and trade deficit to tariff revenue ratios at the wage vector with these evaluations then used to match the model to the data. This procedure differs from the “exact-hat algebra” approach (Dekle et al. (2008)) typically employed by the quantitative trade modeling literature. That approach uses linear regression to estimate the θ^k parameters, and employs exact-hat algebra to perform counterfactuals using the estimates of the θ^k parameters along with a parsimonious subset of other model parameters that can be inferred directly from data. The nested procedure we adopt is somewhat costlier from a computational stand point. We now discuss the advantages and disadvantages of the different approaches in more detail.

The sector level θ^k parameters are the key parameters to be estimated under the exact-hat algebra approach. Typically, the approach taken to estimating θ^k is to derive a linear estimating equation where the left-hand-side variable is a non-linear transformation of bilateral trade flows at the country pair-sector-direction level and the right-hand-side variable is a non-linear transformation of either productivities (Costinot et al. (2011))

¹⁷This weight was chosen so the deficit ratios do not dominate the objective function. Our estimates are not sensitive to changing this weight by a factor of ten in either direction. Focusing on the θ parameters, the correlations of the main estimates with the adjusted weights are 0.997 and 0.998. The mean θ 's are 5.25 in both cases, compared to 5.24 in the main estimates that we present below.

¹⁸In a linear regression, software can determine which fixed effects are not separately identified from each other and drop them automatically. However, this is not possible in a non-linear setting. Ridge regularization resolves these co-linearity issues in our non-linear setting. In the same spirit, we set the z_i^k 's and θ^k 's for the non-traded sectors equal to the means of their estimated values in the traded sectors.

or tariffs (Caliendo and Parro (2015)). The parameter θ^k is the coefficient on the right-hand-side variable in these formulations. Caliendo and Parro (2015) use a rich set of fixed effects to isolate variation in tariffs that is within country-sector, and thus require some countries to have discriminatory tariffs.

We depart from the exact-hat algebra approach for several reasons. First, while our approach yields estimates of all model parameters, the iceberg costs and country productivity levels are not separately estimated under the exact-hat algebra approach. For many counterfactuals this feature can be attractive, as one can then avoid separately estimating the large numbers of additional model parameters that characterize the iceberg costs and country productivities. Indeed, the regularization we use in estimation is necessary because many of these parameters are not well identified from each other. But as discussed above, in our counterfactuals we want to constrain tariffs to satisfy the round-trip inequality in (6), and to impose this constraint we must have estimates of the iceberg costs. Beyond this, having separate estimates of the iceberg costs and country productivity levels also allows us to evaluate whether the estimates make sense and conform with available estimates of the same objects in other papers.

A second benefit of our approach is that we do not have to drop observations that feature zero trade between two countries in a sector. Such censoring, which can produce biased and inconsistent estimates as discussed in Silva and Tenreyro (2006), is an unattractive feature of the exact-hat algebra approach, which relies on a log transformation of the left-hand-side variable. Even at our level of aggregation of two digit sectors and 11 countries/regional entities, 7.5% of the observed trade flows are equal to zero.¹⁹

An additional consideration is that the exact-hat algebra approach attributes all idiosyncratic differences in a country pair’s trade flows to model primitives, leaving no role for measurement error in explaining the trade flows. The estimates of θ^k are not biased by multiplicative measurement error in that approach, but the fact that the counterfactuals are performed relative to the observed trade flows implies that all trade flows are driven by model primitives. This attribution is attractive if trade flows are measured precisely and measurement error is in fact insignificant, because this approach then has the advantage of accounting for unmodeled primitives that determine trade flows; but if measurement

¹⁹Compared to the trade settings investigated in Silva and Tenreyro (2006), the percent of zeros in our trade data is relatively low. Nonetheless, in the context of estimating differentiated product demand systems, Gandhi et al. (2019) report Monte Carlo simulations where substantial bias results from dropping zeros when the percent of zeros in the data is less than ten percent.

error is substantial, the attribution will be misplaced.²⁰

Our estimation relies on patterns in the data familiar from the previous literature, such as, conditional on observable components specified in the iceberg costs, the covariance of trade shares with tariffs. Identification thus relies on tariff data and the assumption that unobserved factors affecting import shares are not correlated with tariff variation, as would be the case, for example, from classical measurement error. While tariffs are endogenously chosen according to our bargaining model (and hence the tariffs that we use to estimate the trade model would also have been endogenously chosen in previous rounds according to our bargaining model), they are chosen optimally for the productivity and iceberg cost parameters which are accounted for in the trade parameter estimation.

The deficit conditions are used to estimate the vector of trade deficit parameters \mathbf{D} . We follow the recommendation in Ossa (2016) to allow for trade deficits in estimating the trade model but to then zero them out and compute counterfactuals relative to a zero trade deficit benchmark. Our approach here is somewhat novel in that treating the deficits as parameters to be estimated resolves the issue raised in Ossa (2016) about matching units between the model’s numeraire and observed trade deficits. Finally, setting the wages equal to the observed relative value added per capita can be interpreted as a set of over-identifying conditions which help pin down the relative productivity of countries in finite samples. We find that including the value added targets in estimation provides more realistic estimates of country level productivity. Furthermore, setting wages equal to observed relative value added per capita significantly speeds up estimation. Instead of solving for a fixed point of the trade equilibrium for each evaluation of the estimation objective function, we only need evaluate the market clearing conditions given the observed wages and parameter vector. We constrain the estimation so that, at the estimated parameter vector, the observed wages are equilibrium wages.

4.2 Estimation of Bargaining Model Parameters

With estimates of the trade model in hand, we estimate the bargaining power parameters between pairs of countries in a second step.²¹ For this step we employ a method of

²⁰As Silva and Tenreyro (2006) note, trade data is plagued by many possible forms of measurement error; see Feenstra et al. (1997) for an inventory of common possibilities.

²¹Gowrisankaran et al. (2015) derive estimating equations for joint estimation of the bargaining model with the demand model. Their approach cannot be applied in our setting because of the sequential timing of the trade equilibrium after the bargaining equilibrium.

moments estimation. Using the estimated trade model parameters, we can solve the bargaining model for predicted tariffs given any vector of bargaining power parameters. We numerically search over these parameters to minimize the distance between the mean tariffs negotiated by each partner in each bilateral of the Uruguay Round and the same mean tariff bargaining outcomes predicted by our model. We weight each distance between means by the number of tariffs under negotiation by that partner in that bilateral, as the variance of the mean is proportional to this number. Thus, letting R be the number of bilaterals in the Uruguay Round and letting the set B_h contain the pair of countries negotiating in bilateral h , we estimate the bargaining power parameters by solving

$$\min_{\zeta} \sum_{h=1}^R \sum_{i \in B_h} n_i^h (\bar{t}_{MFN,i}^k(\zeta) - \bar{t}_{MFN,i}^k)^2$$

where ζ is the vector of bargaining power parameters, n_i^h is the number of country i 's tariffs under negotiation in bilateral h , $\bar{t}_{MFN,i}^k(\zeta)$ is the model's prediction for country i 's MFN tariff in sector k for a candidate vector of bargaining power parameters ζ , $t_{MFN,i}^k$ is the observed MFN tariff of country i in sector k in the year 2000, and where a bar over a variable denotes its mean. The estimated bargaining parameters are thus driven by the observed degree of tariff cutting by each country in a pair together with the estimated slope of the bargaining frontier for that pair. We discuss the determinants of the slope of each bargaining frontier in the context of our parameter estimates in section 6.2.

We interpret the errors between our predicted tariffs and the observed tariffs as measurement error in the tariffs.²² We compute standard errors again by a residual bootstrap procedure. We first draw a bootstrapped trade parameter estimate to account for uncertainty in the trade parameters, and then re-estimate the bargaining parameters on a residual bootstrap data set of tariff changes to generate a draw from the distribution of bargaining parameters. We report the standard deviation of these draws as the standard error of the bargaining parameters.

5 Trade Model Estimates and Benchmarks

In this section we present our estimates for the trade model parameters, and we report a number of benchmark quantifications implied by these estimates.

²²While the product line tariffs are likely measured with little error, aggregation to the sector level induces measurement error due to measurement error in the trade flows.

5.1 Trade Model Estimates

The estimated average iceberg cost across all traded sectors and country-pairs is 494%. The average-across-sectors incurred iceberg cost is 202%, as lower iceberg cost country pairs trade with each other more. Among the Quad countries of the US, EU, Japan and Canada, the average and average incurred iceberg costs are 102% and 99% respectively. These latter numbers are broadly consistent with available estimates in the literature.²³

We also estimate the cross-country fundamental productivity levels. The rankings of the productivity distributions follow the observed value added per capita of the countries. Productivity levels are positively correlated across sectors, so the higher productivity countries in agriculture also tend to be the higher productivity countries in manufacturing.

Table 2: Productivity Shape Estimates by Industry.

Sector	$\hat{\theta}$	SE	Sector	$\hat{\theta}$	SE
Feeding stuff	10.12	1.21	Non-metallic mineral manufactures	4.95	0.71
Plumbing, heating and lighting	9.56	1.23	Seafood	4.78	1.64
Travel goods and bags	9.01	2.74	Scientific instruments	4.76	0.65
Live animals	8.25	3.04	Power generating machinery	4.75	0.44
Other transport equipment	7.49	1.15	Footwear	4.72	0.93
Meat	7.38	1.16	Office machines	4.71	0.70
Electrical machinery	6.90	0.88	Misc manufactures	4.39	0.67
Non-ferrous metals	6.83	1.10	Pharmaceutical	4.23	0.67
Pulp and waste paper	6.43	0.84	Iron and steel	4.20	0.76
Sugar	6.36	3.84	Specialized Machinery	4.14	0.57
Misc. Edible	6.31	0.71	Organic chemicals	4.13	0.53
Furniture and parts thereof	6.20	0.77	Hides and skins	4.11	0.81
Dairy	6.08	0.85	All others	4.06	0.49
Cereals	6.04	0.88	Inorganic chemicals	4.01	0.59
Coal	5.99	1.28	Vegetables and fruit	3.94	0.55
Road vehicles	5.75	0.90	Textile fibres	3.93	0.66
Tobacco	5.64	1.03	Chemical	3.58	0.45
Petroleum	5.63	0.57	Dyeing and tanning	3.41	0.38
Paper manufactures	5.62	0.83	Rubber manufactures	3.38	0.43
Cork and wood	5.59	1.05	Fertilizers	3.19	0.46
Resins	5.58	0.87	Coffee, Tea, Spices	2.97	0.35
Beverages	5.52	1.53	Crude rubber	2.87	0.40
Wood manufactures	5.37	0.77	Fabrics	2.03	0.21
Crude materials,n.e.s.	5.19	0.66	Metal Ores	1.73	0.20
Animal oils and fats	5.17	0.89			
Overall Mean	5.24	0.44			

Notes: Estimates of productivity shape parameter θ by sector in descending order of estimate.

Table 2 presents the within-country dispersion of productivity parameter estimates by

²³For example, Novy (2013) derives a micro-founded measure of bilateral trade costs that can be used to indirectly infer trade costs from observable data. Adopting a θ value of 8 for his calculations, he reports an average iceberg cost of 108% in 1990 for a broadly similar set of industrialized countries.

sector, ordered by descending θ^k (descending trade elasticity). Across sectors, the mean θ is 5.24, with a standard deviation of 0.44 and range from 10.12 to 1.73.

The range of θ estimates in the literature is arguably quite wide and comparison from paper to paper is difficult due to different degrees of product or geographical aggregation. That said, the Eaton and Kortum (2002) estimate of θ across sectors is 8.28. Costinot et al. (2011) put the estimate at 6.53. Caliendo and Parro (2015) estimate an aggregate θ of 4.55 with a range from 50.01 to 0.37. Ossa (2014) estimates a mean of 3.42 with a range from 10.07 to 1.19. Overall, the θ values we estimate tend to be in the mid-range of those found in the literature.

5.2 Trade Model Benchmarks

We compute various benchmarks implied by the 1990-based estimated trade model. The first and second columns of Table 3 report respectively the changes in welfare that would result if, with regard to all products, the world reverted to autarky, or if all iceberg trade costs and tariffs were removed and trade were frictionless. These are standard benchmarks in the quantitative trade modeling literature, and are useful for positioning the broad predictions from our quantitative trade model within that literature.

Our estimates indicate that, relative to the magnitudes implied under the status-quo 1990 tariffs, moving to autarky would reduce world real income by 5.09% with the associated reduction in country welfare averaging 5.42% on a population-weighted basis, while moving to a world of frictionless trade would raise world real income by 57.77% with the average increase in country welfare amounting to 69.66%.²⁴ For the US, moving to autarky reduces country welfare by 3.68%, which is somewhat larger than the 0.7% to 1.4% range computed by Arkolakis et al. (2012) under the assumption of only final goods and of a single trade elasticity in the range of 5 to 10 applied to all sectors. This number is lower, however, than the 13.5% estimated in Ossa (2015), whose analysis, like ours, has intermediate goods and heterogeneity in θ across sectors but is carried out with 2007 data and has more sectors and a substantially lower estimated trade elasticity of 3.63.

We are also interested in a benchmark tariff bargain that reaches the efficiency frontier. There are no market imperfections or political economy forces or distributional concerns in our model, and so achieving free trade in all tariffs would place the world on the efficiency

²⁴We calculate world real income as the sum across countries of the country-level real income, where country-level real income is defined as the country's nominal income deflated by its model-based exact price index according to the utility function in (5) and is equal to country-level welfare.

Table 3: Model Benchmarks

	Autarky	Frictionless Trade	Free Trade	World Income Maximizing	Nash
Welfare by Country					
US	-3.68%	36.54%	-0.09%	0.70%	-0.35%
EU	-3.38%	43.61%	-0.09%	-3.06%	-0.20%
Japan	-4.95%	41.32%	0.30%	1.36%	-0.41%
South Korea	-13.74%	87.86%	0.53%	2.85%	-0.75%
Australia	-8.47%	147.43%	0.00%	-1.15%	-0.71%
Canada	-11.66%	103.54%	0.12%	0.50%	-0.71%
Africa NES	-9.23%	144.31%	0.28%	1.67%	-1.13%
America NES	-3.83%	79.02%	0.16%	0.47%	-0.40%
Asia NES	-6.11%	77.74%	0.62%	2.01%	-0.93%
Europe NES	-8.06%	114.27%	0.61%	2.27%	-1.54%
MENA NES	-10.51%	173.53%	0.76%	1.64%	-1.92%
Mean	-5.42%	69.66%	0.39%	1.00%	-0.73%
World Real Income	-5.09%	57.77%	0.17%	0.34%	-0.50%

Notes: All changes are relative to model-predicted 1990 status quo. In column 1 we set iceberg costs for all countries in all sectors to 5000%. In column 2 we set iceberg costs and tariffs to zero for all countries in all sectors. In column 3 we set to zero all non-agricultural tariffs that were under negotiation in the Uruguay Round according to our model, while in column 4 we solve for the world-real-income maximizing levels of these tariffs and in column 5 we compute a Nash equilibrium in these tariffs. Tariffs in columns 4 and 5 are non-discriminatory. The mean across countries is weighted by population.

frontier. But as a benchmark with which to compare our model outcomes, global free trade is not particularly meaningful, both because our bargaining analysis is limited to tariffs on non-agricultural products and because tariff bargaining on those products was far from universal in the Uruguay Round given that the bargaining was limited to GATT members and proceeded under the principal supplier rule. We consider instead a bargain that sets to zero only the non-agricultural tariffs under negotiation in the Uruguay Round according to our model, and an alternative benchmark that sets these same tariffs equal to the levels that would maximize world real income in light of transport costs and the existing distortions implied by the fixed levels of all other tariffs in the world. Finally, we can solve for the non-cooperative Nash tariff equilibrium over this same set of tariffs, holding all other tariffs fixed at their 1990 levels. To keep these benchmarks comparable with each other and with what might be achieved in GATT tariff negotiations, we continue to assume that the tariffs that we allow to change under the world-welfare maximizing and Nash benchmarks are constrained by the MFN requirement.²⁵

²⁵With certain exceptions, GATT members are obligated to apply all of their tariffs on an MFN basis even if they have not bound those tariffs as a result of GATT negotiations.

The third, fourth and fifth columns of Table 3 report benchmark welfare effects under these free-trade, world-real-income maximizing and Nash benchmarks, respectively. As mentioned, we limit the tariff changes to those tariffs on non-agricultural products that were “under negotiation in the Uruguay Round,” where we consider the non-agricultural tariff $t_{MFN,i}^k$ to be under negotiation in the Uruguay Round if, according to our estimated trade model and focusing on the industrialized countries, country i is the principal supplier of at least one non-agricultural good to the principal supplying country of product k into country i so that there is a double coincidence of wants between the two countries (and this pair of countries is thus included in the set of Uruguay Round bilateral tariff bargains in our model economy).

The third column of Table 3 reports the welfare results from reducing all the tariffs that were under negotiation in the Uruguay Round from their 1990 levels to zero, with all other tariffs held fixed at their 1990 levels. The move to free trade increases world real income by 0.17% and results in an average increase in country welfare of 0.39%. The distribution of the gains across countries in column 3 is uneven. This pattern serves to highlight the two-pronged effects of MFN tariff liberalization, namely, the reduction in domestic price distortions and the effects on the terms of trade. The former effects are enjoyed by countries that liberalize, and these effects contribute to an increase in national welfare; the latter effects are experienced by all countries, and these effects can either diminish or add to national welfare and tend to be negative for countries that liberalize and positive for countries that do not liberalize.

The fourth column of Table 3 reports results when we solve for the levels of the tariffs negotiated in the Uruguay Round that would maximize world real income. Relative to 1990 levels, world real income rises by 0.34% and country welfare levels increase by an average of 1.00%, more than under our free trade benchmark for two reasons. First, there are pre-existing distortions associated with the tariffs not under negotiation at Uruguay which remain fixed at their 1990 levels under both benchmarks. We find that the world-real-income maximizing benchmark on average entails further tariff liberalization than the free trade benchmark (i.e., on average, import subsidies are needed to offset the trade restricting effects of the tariffs not under negotiation at Uruguay such as agricultural tariffs and the tariffs of the developing world). Second, the terms-of-trade effects of the world-real-income maximizing tariffs redistribute income towards those countries with lower national price indices. Since countries do not have sufficient policy instruments with which to effect international lump-sum transfers, maximizing world real income balances

the benefits of such redistribution against the associated distortion costs of deviations from free trade. The large EU losses in the fourth column reflect this calculus and that according to our model the EU has a relatively high price index.²⁶

The fifth column of Table 3 reports the welfare results from increasing all the tariffs that were under negotiation in the Uruguay Round from their 1990 levels to their Nash equilibrium levels, with all other tariffs held fixed at their 1990 levels. According to our estimates, world real income would be reduced by 0.50% and all countries would be worse off, with the average fall in country welfare amounting to 0.73%. This relatively modest loss reflects the fact that the move to Nash tariffs is only allowed for products that were under negotiation in the Uruguay Round, and that the Nash tariffs are sizable but not prohibitive. For example, US tariffs rise on average from 4.5% to 17.5%. EU tariffs rise on average from 6.5% to 21.9%.

Together our estimates in the third and fifth columns of Table 3 suggest that, beginning from Nash tariffs, the GATT rounds up to but not including the Uruguay Round had achieved by 1990 roughly three quarters of the potential world-wide gains in real income from the complete elimination of the tariffs that were under negotiation in the Uruguay Round. Compared to a benchmark that sets these same tariffs equal to the levels that would maximize world real income, the fourth and fifth columns of Table 3 suggest that, beginning from Nash tariffs, the GATT rounds leading up to the Uruguay Round achieved roughly 60 percent of the aggregate world-wide real income gains that were possible with changes to the tariffs under negotiation in the Uruguay Round.

6 Bargaining Model Estimates

We now turn to our second step and estimate the parameters of our bargaining model.

6.1 The Structure of Bilateral Bargaining Pairs

According to the principal supplier rule, a double coincidence of wants must exist between any viable pair of bargaining partners, in the sense that each country in the bargaining pair must be a principal supplier of at least one good to the other country in the pair. Focusing on the industrialized countries in our model world economy, we let the model

²⁶According to our model, a country's sectoral price indices reflect its wage and its productivities and the tariff and iceberg trade costs that it faces. As a result of this combination of factors, we find that the EU has the highest national price index among our six countries.

predictions regarding principal supplier status dictate the set of viable pairs of bargaining partners that will comprise the structure of Uruguay Round tariff bargains within the context of our model world economy.²⁷

Table 4 displays the observed and predicted pattern of principal supplier status at the level of product aggregation in our data. In defining the principal suppliers relevant for Uruguay Round negotiations, we have netted out trade with fellow PTA members (e.g., US exports to Canada are excluded when calculating the identity of principal suppliers into Canada).²⁸

Table 4: Principal Supplier Relationships

	US	EU	Japan	South Korea	Australia
US					
EU	[27,27]				
Japan	[19,5]	[13,3]			
South Korea	[13,2]	[5,2]	[16,4]		
Australia	[10,1]	[22,1]	[3,3]	[1,2]	
Canada	0,0	[30,3]	3,0	1,0	[1,1]
Africa NES	2,0	34,0	0,0	0,0	0,0
America NES	34,0	4,0	1,0	0,0	0,0
Asia NES	7,0	17,0	12,0	0,0	2,0
Eur. NES	0,0	39,0	0,0	0,0	0,0
MENA NES	3,0	34,0	0,0	0,0	0,0

	US	EU	Japan	South Korea	Australia
US					
EU	[26,22]				
Japan	[14,12]	[19,10]			
South Korea	[12,2]	3,0	[20,3]		
Australia	7,0	21,0	[7,1]	[2,1]	
Canada	0,0	30,0	4,0	0,0	0,0
Africa NES	2,0	34,0	0,0	0,0	0,0
America NES	19,0	17,0	2,0	0,0	0,0
Asia NES	5,0	19,0	7,0	0,0	2,0
Eur. NES	0,0	39,0	0,0	0,0	0,0
MENA NES	2,0	36,0	0,0	0,0	0,0

Notes: The top panel presents principal supplier relationships according to the data. The bottom panel represents principal supplier relationships according to the trade model at the estimated parameter vector. For each cell in the table, the first entry gives the number of products for which the column country is the principal supplier into the row country, and the second entry gives the number of products for which the row country is the principal supplier into the column country. For the numbers in this table, trade with fellow PTA members and NES entities has been netted out. Square brackets indicate the bilateral relationships where both entries are positive.

²⁷Specifically, we exclude the regional entities when determining principal supplier status for the purpose of defining the bargaining pairs in the Uruguay Round, and define the principal supplier of product k into country i as the largest supplier of product k into country i among the industrialized countries other than i .

²⁸What is relevant for GATT tariff negotiations is the identity of the principal supplying country of product k into country i among those GATT member countries that face country i 's MFN tariff on product k . This would exclude a PTA partner.

The top panel of Table 4 records 12 country-pairs where both entries are non-zero. These 12 pairings involve each of our six countries – the US, the EU, Canada, Japan, Australia and South Korea. At the level of aggregation in our data, these 12 bargaining pairs negotiated over 214 tariffs covering 61% of world trade in industrialized goods.²⁹

The bottom panel of Table 4 displays the predicted pattern of principal supplier status according to our trade model estimates. The principal supplier relations predicted by our model capture seven of the 12 country pairings in the data: US-EU, US-Japan, US-South Korea, EU-Japan, Japan-Australia, Japan-South Korea and Australia-South Korea. And together these seven predicted pairings involve five of the six countries that are represented in pairings according to the data: three of the four Quad members – the US, the EU and Japan – and Australia and South Korea. More broadly, our set of bargaining countries includes the major industrialized countries that were arguably the key actors in the tariff negotiations of the Uruguay Round, the exclusion of Canada from this set being potentially the most important omission. The 151 tariffs under negotiation in these bilaterals also covered 55% of 1990 world trade in industrialized goods. Overall the seven bargaining pairs predicted by our model therefore seem to comprise a fair representation of the structure of Uruguay Round tariff bargains within the context of our model world economy.³⁰

6.2 Bargaining Power

With the structure of Uruguay Round tariff bargaining represented by these eight bargaining pairs, Table 5 displays our bargaining parameter estimates for each of the negotiating countries. According to the point estimates, three bilaterals are characterized by interior bargaining power and the remaining four amount to take-or-leave bargaining, though these point estimates should be taken with a grain of salt because a number of the standard errors are large. In terms of model performance, the average tariff reduction on the 151 tariffs in the data is 2.5 percentage points. Despite not featuring a constant, the

²⁹This level of coverage is broadly consistent with available assessments of the Uruguay Round. Fieleke (1995) reports that the industrialized countries “agreed to reduce tariffs on industrial imports amounting to 64 percent of the total value of their imports of such products.”

³⁰As mentioned, we compute our bargaining equilibria within a zero-trade-deficit benchmark of our model world economy. For the purpose of identifying principal supplier status, we chose to make use of our model’s predicted trade relations prior to zeroing out trade deficits. This allows us to make clean comparisons of our model’s predicted principal supplier relations with those found in the data, but the principal supplier relations predicted by the model in the zero-trade-deficit benchmark are very similar.

estimated bargaining model produces an average reduction of 2.9 percentage points. The correlation of the model’s predicted tariff changes with observed tariff changes is 0.38.

Table 5: Bargaining Model Parameter Estimates

Country Pair	Bargaining Parameter	SE
US - EU	0.01	0.47
US - Japan	0.05	0.33
US - South Korea	0.01	0.25
EU - Japan	0.01	0.14
Japan - Australia	0.99	0.21
Japan - South Korea	0.85	0.17
Australia - South Korea	0.83	0.13


Notes: Estimated bargaining parameters for the first country in each pair.

According to our point estimates, Japan looks to be the strongest overall bargainer, the US and Australia appear to be in the weakest bargaining positions where we note that Australia’s bilateral with Japan is its most important bilateral, and the overall bargaining positions of South Korea and the EU seem to be somewhere in the middle.³¹ That said, it is important to keep in mind when interpreting these estimates that we are abstracting from political economy issues and focusing on non-agricultural tariff bargaining to the exclusion of the broader set of issues covered by the Uruguay Round negotiating agenda.

The factors that determine our bargaining power parameter estimates are both subtle and illuminating. On one level, our model tends to attribute lower bargaining power to a country in a given bilateral if that country’s tariffs under negotiation in the bilateral are reduced to a greater extent than those of its bargaining partner. This is because the division of the surplus will tend to tilt against a country that gives greater tariff cuts than it receives, due to adverse terms-of-trade movements. In the familiar setting of transferable utility this would be the end of the story as far as bargaining power goes: the bargaining frontier then has slope -1 when plotted in payoff space, so that there is a one-to-one mapping between the equilibrium division of surplus and the bargaining powers of the parties. But that mapping does not hold here, because utility is not transferable: the bargaining parties must rely on tariff changes both to create the surplus in their bargain and to divide the surplus between them. Thus, depending on the slope of the bargaining frontier in a given bilateral, an observed asymmetry across countries in the negotiated

³¹The relative importance to Australia of its bilateral with Japan as compared to its bilateral with South Korea is apparent from the numbers of products under negotiation in these bilaterals as displayed in the bottom panel of Table 4, and it is also true based on the bilateral trade volumes involved.

Figure 1: Japan and South Korea Bargaining Frontier



Notes: The solid curve represents the frontier of feasible welfare pairs for the Japan-South Korea bilateral negotiations holding the other pairs fixed at the equilibrium outcomes. The dotted line passes through the agreement point and has slope equal to -1 . The dashed line passes through the agreement point and has slope equal to the negative of the ratio of the Japan price level to the South Korea price level.

tariff cuts in their bilateral and the implied asymmetry in their division of the bargaining surplus need not translate neatly into implied asymmetries in bargaining power.

This is illustrated in Figure 1 for the case of the Japan-South Korea bilateral. With South Korea welfare on the vertical axis and Japan welfare on the horizontal axis and equilibrium disagreement welfares marked at the origin, the solid curve line in Figure 1 depicts the bargaining frontier between these two countries according to our model estimates. Figure 1 also depicts the equilibrium Japan-South Korea agreement point marked at its position on the frontier. This agreement point is defined where an iso-Nash-product contour (not pictured), in the family of such contours with the Japan bargaining power parameter set to its estimated value of 0.85, is tangent to the bargaining frontier. For comparison, the dotted line running through the equilibrium agreement point has slope -1 and confirms that the Japan-South Korea bargaining frontier is steeper at the agreement point than would be the case under transferable utility. We find that the slope of the Japan-South Korea bargaining frontier at this point is -1.96 , implying that it is relatively easy for the two countries to shift surplus to South Korea as compared to shifting surplus to Japan. As a consequence, Japan needs to have more bargaining power than it would need in a transferable utility world to achieve the division of surplus implied by the equilibrium Japan-South Korea agreement.

If direct income transfers were possible between countries, the slope of the bargaining frontier would be equal to the ratio of national price indices in the bargaining countries. To confirm that this is not the only determinant of bargaining frontier slopes in our setting, we include in Figure 1 a dashed line passing through the Japan-South Korea agreement point with slope equal to the negative of the ratio of the Japan price level to the South Korea price level, which according to our model estimates is -1.65. With the bargaining frontier slope at this point equal to -1.96 and corresponding to the solid curve in Figure 1, it is apparent that the slope of the Japan-South Korea bargaining frontier reflects more than simply the ratio of the price indices of the bargaining countries. To understand the forces that determine the slope of the bargaining frontier in our setting, we must consider the two channels through which tariff changes affect national welfare.

One channel runs through the terms of trade. To the extent that tariff changes can alter the terms of trade, the income effects of the induced terms-of-trade changes will redistribute income across countries in proportion to each country's affected trade volume, so the change in income that each country in the bilateral experiences will depend on the volume of that country's trade that is impacted by the induced terms-of-trade movements. Thus, if one country reduces its tariff on a given import that is principally supplied by its bargaining partner but also supplied in substantial quantities by third-party country suppliers, there will be a substantial spillover to those third parties from the terms-of-trade movements triggered by the country's tariff cut that limits the amount of income that is redistributed to its bargaining partner. And how the resulting international redistribution of income triggered by the terms-of-trade movements is valued by each country in the bilateral depends on how this income translates into purchasing power, which in turn depends on the level of each country's price indices.

The other channel runs through domestic prices: depending on the direction of the change, tariff changes can either reduce or exacerbate domestic price distortions and have additional implications for national surplus. For example, starting from its reaction-curve tariffs, a country that reduces its tariffs by a small amount can redistribute a small amount of income to its bargaining partner through the induced terms-of-trade changes without itself suffering any loss of surplus, because the cost to this country of the lost income suffered as a result of its diminished terms of trade is just offset by the benefit that the country enjoys from the reduction in the domestic price distortions.

In short, in addition to the levels of their price indices, the slope of the bargaining frontier between any two countries depends on the strength of market power that each

country yields on world markets, the degree of third-party spillovers from the terms-of-trade movements induced by changes in the tariffs under negotiation by these two countries, and the position of each country's initial tariffs relative to reaction-curve tariffs. Table 6 presents evidence on the magnitudes of these forces for each of the eight bilaterals.

Table 6: Spillover Benefits to Third Parties (MFN Negotiations)

Country 1	Country 2	Reducing Country	Δ Welfare	Δ Welfare	Δ Welfare	Δ Welfare	Δ Welfare
			Country 1 (1)	Country 2 (2)	3rd Parties (3)	Partner + 3rd Parties (4)	Partner / (Partner + 3rd Parties) (5)
US	EU	US	-1.00	0.59	0.98	1.57	0.38
US	EU	EU	0.82	-1.00	1.24	2.06	0.40
US	Japan	US	-1.00	0.95	0.55	1.49	0.64
US	Japan	Japan	0.45	-1.00	0.61	1.06	0.43
US	Korea	US	-1.00	2.71	5.21	7.92	0.34
US	Korea	Korea	0.35	-1.00	0.51	0.86	0.41
EU	Japan	EU	-1.00	1.49	1.16	2.65	0.56
EU	Japan	Japan	0.33	-1.00	0.80	1.13	0.29
Japan	Aus	Japan	-1.00	0.39	0.22	0.60	0.64
Japan	Aus	Aus	2.10	-1.00	0.46	2.56	0.82
Japan	Korea	Japan	-1.00	1.83	1.36	3.19	0.57
Japan	Korea	Korea	0.51	-1.00	0.42	0.93	0.55
Aus	Korea	Aus	-1.00	2.96	8.66	11.62	0.25
Aus	Korea	Korea	0.33	-1.00	0.13	0.46	0.72

Notes: Each row corresponds to a unilateral marginal decrease in tariffs by the “reducing country.” The reducing country reduces tariffs on all goods that it negotiates with the partner country in that row from the negotiated agreement level. The welfare changes are normalized so that the reducing country experiences a reduction in welfare equal to one.

The first three columns of Table 6 report, for each bilateral and beginning from the agreement levels of all tariffs as predicted by our model, the impact on bargaining-partner welfare and third-party country welfare when one of the bargaining partners reduces proportionately all of its tariffs under negotiation in that bilateral by an amount that reduces its welfare by 1 unit. Consider for example the first two rows of Table 6, which relate to the US-EU bilateral. In the first row, the US is the tariff reducing country, and the first three columns indicate that, beginning from agreement tariffs, when the US lowers its tariffs under negotiation in this bilateral by an amount that reduces its welfare by 1 unit, it increases EU welfare by 0.59 units and the welfare of third parties by 0.98 units. In the second row the EU is the tariff reducing country, and the first three columns indicate that, beginning from agreement tariffs, when the EU lowers its tariffs under negotiation in this bilateral by an amount that reduces its welfare by 1 unit, it

increases US welfare by 0.82 units and the welfare of third parties by 1.24 units.

In the fourth column of Table 6 we report the world-wide external effect – the sum of the bargaining-partner and third-party country welfare impacts – of these tariff reductions; this serves as an indication of both the degree of market power over world prices wielded by the tariff reducing country and the distance of this country’s tariffs from reaction curve levels. And in the fifth column we report the fraction of this world-wide external effect that is captured by the bargaining partner, as an indication of the importance of third-party spillovers. The first and second rows show that in the case of the US-EU bilateral, the US increases external world-wide welfare by 1.57 units when it reduces its tariffs under negotiation by an amount that lowers its welfare by 1 unit, while the EU increases external world-wide welfare by 2.06 units when it reduces its tariffs under negotiation by an amount that lowers its welfare by 1 unit. The fraction of the external world-wide effect that is captured by the bargaining partner when the US and EU reduce their tariffs in this fashion is 0.38 and 0.40, respectively.

As the rows of Table 6 reflect, there is significant variation across and within bilaterals in the degree to which tariff cuts can be used to shift surplus to bargaining partners, as well as in the extent of third-party spillovers. Notably, however, the third-party spillovers reported in column 3 are substantial and uniformly positive, a feature that is also reflected in the fact that the fraction of external world-wide effects captured by the bargaining partner as reported in column 5 are always substantially less than one.

While the overall impact on the bargaining partner and third parties (as reported in column 4 of Table 6) together may be expected to be positive when an importing country reduces its MFN import tariffs as part of a bilateral bargain due to the induced terms-of-trade effects, the impact on just the third parties as a group could be positive or negative depending on trade patterns. Intuitively, a given reduction in an MFN import tariff is expected to generate a terms-of-trade gain (loss) for third-party countries that export (import) the corresponding product. The sign of the impact on third parties as a group is thus an empirical question for which column 3 of Table 6 provides an answer.

7 MFN Tariff Bargaining in the Uruguay Round

We now present our quantification of MFN tariff bargaining in the Uruguay Round. The welfare impacts of the Uruguay Round’s MFN tariff bargaining as predicted by our model are presented in the first column of Table 7.

We find that for the tariffs under negotiation, the average tariff fell by 30.22% (the trade-weighted average tariff fell by 19.54%). The world real income gain of 0.06% indicates that the Uruguay Round achieved roughly one third of the world real income gains that could have been had if countries had found a way to negotiate to free trade on the tariffs that they were negotiating in the Uruguay Round, and it achieved a little under 20 percent of the gains compared to the world-real-income maximizing potential. From this perspective, our model indicates that the gains from tariff negotiations in the Uruguay Round were not insubstantial, but the Round still left a significant fraction of potential world-wide gains from negotiating over this set of tariffs as “unfinished business.”

Table 7: Estimated Uruguay Round and Counterfactual Outcomes

	Estimated Bargaining Parameters		All 0.5 Bargaining Parameters	
	MFN	No MFN	MFN	No MFN
	$\Delta\%$ 1990	$\Delta\%$ 1990	$\Delta\%$ 1990	$\Delta\%$ 1990
Δ Mean Tariff	-30.22%	-28.85%	-29.22%	-28.59%
Δ Trade Wgt'd Mean Tariff	-19.54%	-130.36%	-19.11%	-131.27%
Country Welfare				
US	0.01%	-0.13%	0.02%	0.11%
EU	0.02%	0.06%	0.03%	0.08%
Japan	0.10%	0.32%	0.06%	-0.24%
South Korea	0.17%	0.15%	0.32%	0.04%
Australia	0.06%	-1.71%	0.07%	0.18%
Canada	0.02%	-0.49%	0.02%	-0.59%
Africa NES	0.03%	-0.08%	0.02%	-0.09%
America NES	0.02%	-0.05%	0.02%	-0.04%
Asia NES	0.15%	-0.37%	0.17%	-0.33%
Europe NES	0.02%	-0.12%	0.02%	-0.16%
MENA NES	0.01%	-0.28%	0.00%	-0.49%
Mean	0.10%	-0.21%	0.11%	-0.19%
World Real Income	0.06%	-0.02%	0.06%	-0.07%

Notes: Each column represents changes in the row relative to the pre-Uruguay tariff levels. Tariff averages are computed among non-agriculture sectors for the bargaining countries. The mean across countries is weighted by population.

Column 1 also indicates that all countries and regional entities gained from the Uruguay Round. For the five countries who were engaged in the Round’s tariff bargaining, these gains reflect the combined effects of the efficiency benefits of their own MFN liberalization and the terms-of-trade movements triggered by the Round’s MFN tariff-cutting results. For Canada and for each of our five regional entities, none of whom are among our bargaining pairs, these gains came in the form of terms-of-trade improvements triggered by the MFN tariff liberalization of our seven bargaining pairs in the Round. In total, the

“free-rider” gains from MFN accruing to these non-participants amount to 35.38% of the total world real income gains generated by the Round.

We can also use our model to investigate how the bilaterals of the Uruguay Round interacted with each other in the presence of MFN. A simple way to do this is to compare the full equilibrium outcome of the Round according to our model, reproduced in column 1 of Table 8, with an alternative in which a single bargaining pair engages in MFN tariff bargaining over the tariffs that the bargaining pair negotiated in the Uruguay Round.

Table 8: MFN Bargaining Outcomes

	MFN Equil- ibrium (1)	Single Bargains							Sum of Single Bargains (9)
		US- EU (2)	US- Japan (3)	US- SK (4)	EU- Japan (5)	Japan- Aus (6)	Japan- SK (7)	Aus- SK (8)	
US	0.01%	0.00%	0.00%	0.00%	0.01%	0.00%	0.01%	0.00%	0.00%
EU	0.02%	0.01%	0.00%	0.00%	0.00%	0.00%	0.01%	0.00%	0.01%
Japan	0.10%	0.00%	0.00%	0.00%	0.04%	0.00%	0.05%	0.00%	0.09%
South Korea	0.17%	0.00%	0.00%	0.01%	0.01%	0.00%	0.10%	0.01%	0.11%
Australia	0.06%	0.00%	0.00%	0.00%	0.01%	0.00%	0.00%	0.05%	0.05%
Canada	0.02%	0.00%	0.00%	0.00%	0.01%	0.00%	0.00%	0.00%	
Africa NES	0.03%	0.00%	0.00%	0.00%	0.01%	0.00%	0.00%	0.00%	
America NES	0.02%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%	0.00%	
Asia NES	0.15%	0.00%	0.00%	0.01%	0.01%	0.00%	0.08%	0.01%	
Europe NES	0.02%	0.00%	0.00%	0.00%	0.01%	0.00%	0.00%	0.00%	
MENA NES	0.01%	0.00%	0.00%	0.00%	0.01%	-0.01%	0.00%	0.00%	
Mean	0.09%	0.00%	0.00%	0.00%	0.01%	0.00%	0.04%	0.01%	0.06%
World Real Income	0.06%	0.00%	0.00%	0.00%	0.01%	0.00%	0.03%	0.00%	0.05%

Notes: Each column represents changes in the row relative to the pre-Uruguay tariff levels. The first set of columns represents the Horn-Wolinsky MFN equilibrium at the estimated bargaining parameters. The next columns represent the outcomes of single pair MFN bargains holding the other pair’s at their 1990 levels. The mean across countries is weighted by population.

Columns 2 through 8 of Table 8 report the welfare changes associated with this single-bargain counterfactual taking each of our seven bargaining pairs in turn, while column 9 reports the welfare changes that each bargaining country would have experienced as a result of the sum of the welfare effects of its single-pair bilaterals. As a comparison of columns 1 and 9 reveals, between one quarter and two thirds of the welfare gains enjoyed by the EU, South Korea and Australia as a result of the Uruguay Round occur because of the interaction effects between bilaterals in the Round, while these interaction effects accounted for virtually all of the US gains and just 10 percent of Japan’s. In part the presence of those bilaterals impacts each of these bargaining countries directly through

the terms-of-trade benefits associated with the MFN tariff cuts negotiated by those bilaterals. But the presence of those bilaterals also impacts each bargaining country indirectly, through their impacts on the equilibrium outcomes of each country's own bilaterals. This indirect effect is on stark display in column 3 of Table 8, which indicates that if the US-Japan bilateral had been initiated in isolation, the US and Japan would have failed to find mutually beneficial tariff reductions. Finally, as the last row of Table 8 reveals, overall the interaction effects across bilaterals of the Uruguay Round accounted for roughly 20% of the increases in world real income that the Round generated.

Our bargaining protocol does not impose a reciprocity norm, but it is nevertheless of interest to assess the degree to which the equilibrium Horn-Wolinsky MFN tariff bargains predicted by our model conform to such a norm. As Bagwell et al. (2020a) observe, a key innovation of the GATT was that it multilateralized the reciprocity norm that had previously taken a bilateral form in the context of various bilateral trade agreements. Multilateral reciprocity in GATT negotiations can be understood as referring to an idealized agreement under which each country experiences an overall increase in export volume that is the same value as the overall increase in its import volume, where exports and imports are valued at pre-negotiation international prices. When this balance is achieved, terms-of-trade changes as a result of the agreement are thereby prevented.³²

To assess the degree to which the equilibrium Horn-Wolinsky MFN tariff bargains predicted by our model conform to multilateral reciprocity, we calculate the predicted change in each country's import and export volumes as a result of the tariff changes negotiated in the Round. The ratio of the change in export volume to change in import volume for each country then provides a measure of the degree of multilateral reciprocity achieved by that country, with a ratio that differs from one signaling deviations from reciprocity and a ratio less than (greater than) one associated with improving (worsening) terms of trade for that country as a result of the Round. We find that for our bargaining countries this ratio varies from a value of 0.99 for the US to a value of 1.16 for Australia, suggesting only modest deviations from multilateral reciprocity, with the pattern of deviations that do exist among these countries seemingly unrelated to bargaining power (the correlation between this ratio and our mean estimated bargaining power parameter for each country among its bargains is 0.09).³³ And we find that the ratio varies between 0.53 and 0.86

³²For a theoretical analysis of multilateral tariff negotiations when strict adherence to MFN and multilateral reciprocity is required, see Bagwell and Staiger (2018).

³³Using other measures to gauge conformity with reciprocity, Limão (2006), Limão (2007) and

for our non-bargaining countries, suggesting more significant deviations from multilateral reciprocity among non-bargaining countries that reflect their ability to free-ride on the MFN negotiations of others.

As constructed, our model of simultaneous bilateral bargaining does not allow for an explicit multilateral process for the implementation of multilateral reciprocity. To the extent that multilateral reciprocity was nevertheless achieved in the Uruguay Round, how, then, should its implementation be interpreted? An interpretation consistent with our model is possible provided that appropriate symmetry conditions, either by chance or by design, were met in the structure of bargaining that transpired in the Round. At the same time, our findings regarding multilateral reciprocity are also consistent with the occurrence of some degree of explicit multilateral coordination, particularly among bargaining countries, to ensure that a sufficient degree of multilateral reciprocity was a feature of the outcome of the Round.³⁴

8 Tariff Bargaining in the Absence of MFN

Given the potential drag on tariff liberalization generated by the positive third-party externalities associated with MFN tariff cuts as reported in Table 6, could the abandonment of MFN have allowed countries to achieve greater tariff liberalization than occurred under the MFN restriction, and in so doing have allowed the Uruguay Round to achieve greater gains in world welfare? And would the distribution of the gains from tariff liberalization in the Uruguay Round across countries have been impacted in a substantial way had the MFN requirement not been in place? We now compare the outcomes from the Uruguay Round with the outcomes that would be predicted by our model had the Uruguay Round negotiations occurred under a bargaining protocol that abandons the MFN requirement.

Recall that we represent the Uruguay Round bargaining protocol with three institutional constraints: countries (i) are restricted to bargain over MFN tariffs, (ii) must respect existing GATT tariff commitments and not raise their tariffs above these commitments, and (iii) must abide by the principal supplier rule. We now consider an alternative bargaining protocol under which the first and third of these constraints are removed. Our

Karacaovali and Limão (2008) also find evidence that the tariff negotiating outcomes of the Uruguay Round for bargaining countries (the US and the EU) were broadly consistent with reciprocity.

³⁴For example, countries might achieve such coordination by using a multilateral review. We discuss this possibility further in the Conclusion.

focus is on how relaxation of the MFN requirement impacts tariff bargaining, and as the principal supplier rule was introduced into the GATT bargaining protocols to facilitate bilateral MFN tariff bargaining, it is natural to remove these two constraints at the same time. Because the model does not perfectly predict tariffs under our representation of the Uruguay Round protocol, we compare simulated outcomes under the counterfactual protocol to the simulated outcomes under our representation of the Uruguay Round protocol rather than to the observed post-Uruguay tariffs.

To predict outcomes under discriminatory negotiations, we again solve for a bargaining equilibrium with our estimated bargaining parameters. In the discriminatory case, however, each pair negotiates only over tariffs that they apply to each other. Differently from MFN tariff cuts, discriminatory tariff reductions thus do not directly alter the tariffs faced by third parties. Instead, discriminatory bilateral tariff bargains affect third-party country welfare indirectly by altering the trade equilibrium. While as Table 6 confirmed, MFN tariff reductions generate positive spillovers, the third-party effect of a discriminatory tariff reduction agreed within a bilateral is expected to be negative, due to the less favorable terms of trade for third-party countries. The implied transfer of surplus from third parties to bargaining parties is then likely to drive down the levels of these negotiated discriminatory tariffs from what the negotiated levels of these tariffs would be under MFN.³⁵

To highlight the intensive-margin impact that the third-party effects of discriminatory tariff reductions have on tariff bargaining outcomes in our model, we consider a counterfactual in which, for each country, the set of its tariffs being negotiated is constrained to include only the sectors that were negotiated by that country in the Uruguay Round, and the set of countries negotiating on these tariffs is constrained to include only the countries that it negotiated with in the Uruguay Round. That is, if county i was negotiating an MFN tariff cut on sector k imports with the principal supplier of sector- k exports into its market, then in our counterfactual country i is allowed to negotiate a discriminatory tariff cut on sector- k imports with each of the countries that it bargained with in the Uruguay Round and that also export sector- k goods to its market. But we do not allow additional extensive margin effects on the pattern of bargaining.

³⁵Intuitively, negotiated tariffs are likely to be lower in the absence of the MFN constraint for the following reason: starting at their negotiated MFN tariffs, if two countries were allowed to use discriminatory tariffs, then they could jointly gain from exchanging (perhaps small) discriminatory tariff cuts, since they thereby could enjoy a transfer of surplus from third parties due to the implied world-price movements. See Bagwell et al. (2018) for confirmation of this reasoning in a three-country model.

The second column of Table 7 presents the results of this counterfactual. In the absence of MFN we find that for the tariffs under negotiation, the average tariff would have fallen by 28.85% as compared to the 30.22% number under MFN reported in the first column of Table 7. But the greater incentive to negotiate tariff cuts that abandoning MFN can engender is revealed when the No-MFN averages are calculated on a trade-weighted basis: as column 2 of Table 7 reports, the trade-weighted average tariff would have fallen by 130.36% if MFN were abandoned as compared to 19.54% under MFN.³⁶ Abandoning MFN can thus introduce a strong liberalizing force into tariff bargaining. But as the last line of Table 7 reveals, the liberalizing force unleashed by the abandonment of MFN is in fact *overly* strong, wiping out the world real income gains that MFN tariff bargaining would have produced and leading to a 0.02% reduction in world real income relative to the 1990 status quo itself.

We expect from the findings of Bagwell et al. (2020b) that in the absence of the MFN requirement Nash-in-Nash tariff bargaining would result in inefficient over-liberalization. But what a comparison across the findings reported in columns 1 and 2 of Table 7 indicates is that, according to our estimates and in the Uruguay Round context, the degree of over-liberalization is sufficiently important to outweigh the inefficient under-liberalization that arises under the MFN requirement, resulting in worse outcomes for world real income under discriminatory tariff bargaining than under MFN tariff bargaining.

To explore the driving force of this excessive liberalization, Table 9 provides analogous information to Table 6 for the counterfactual case of discriminatory tariff bargaining. The most striking difference is in the spillovers to third parties reported in column 3, which for MFN tariff bargaining are uniformly positive but which for discriminatory tariff bargaining are (with one exception) always negative. This difference is reflected as well in the fraction of the world-wide external effect of a country's tariff cuts that are captured by

³⁶As this level of tariff cuts indicates, we find under the No-MFN counterfactual that countries are led to adopt discriminatory import subsidies in a substantial number of sectors. The desirability of discriminatory import subsidies can be understood intuitively by noting that, beginning from free trade on all goods, two countries would have incentive to introduce small discriminatory subsidies on each other's imports in order to enjoy the third-party terms-of-trade benefits that such subsidies would induce without generating any first-order inefficiency costs. And while in practice explicit import subsidies are not common, excessive liberalization of non-tariff barriers that were especially restrictive to imports from a bargaining partner would serve the same purpose. We note, too, that for an extended model that includes political economy or distributional concerns, excessive liberalization might take the form of import tariffs that are positive but still lower than efficient from the perspective of government preferences.

its bargaining partner as reported in column 5, which for MFN are always below one but for discriminatory bargaining are (with one exception) always above one. This negative third-party externality drives down the levels of the negotiated tariffs in the absence of the MFN constraint from what the negotiated levels of these tariffs would be under MFN.

Table 9: Spillover Benefits to Third Parties (Discriminatory Negotiations)

Country 1	Country 2	Reducing Country	Δ Welfare Country 1 (1)	Δ Welfare Country 2 (2)	Δ Welfare 3rd Parties (3)	Δ Welfare Partner + 3rd Parties (4)	Δ Welfare Partner / (Partner + 3rd Parties) (5)
US	EU	US	-1.00	0.76	-0.44	0.33	2.34
US	EU	EU	1.30	-1.00	-0.02	1.28	1.01
US	Japan	US	-1.00	1.24	-0.11	1.13	1.10
US	Japan	Japan	0.78	-1.00	-0.37	0.41	1.90
US	Korea	US	-1.00	1.93	-0.40	1.53	1.26
US	Korea	Korea	0.52	-1.00	-0.31	0.21	2.50
EU	Japan	EU	-1.00	1.72	0.00	1.72	1.00
EU	Japan	Japan	0.57	-1.00	-0.39	0.19	3.07
Japan	Aus	Japan	-1.00	1.39	-0.54	0.84	1.65
Japan	Aus	Aus	0.60	-1.00	-0.18	0.42	1.44
Japan	Korea	Japan	-1.00	1.93	-0.23	1.70	1.14
Japan	Korea	Korea	0.48	-1.00	-0.12	0.36	1.33
Aus	Korea	Aus	-1.00	2.11	-0.72	1.39	1.51
Aus	Korea	Korea	0.47	-1.00	-0.38	0.09	5.50

Notes: Each row corresponds to a unilateral marginal decrease in tariffs by the “reducing country.” The reducing country reduces tariffs on all goods that it negotiates with the partner country in that row from the discriminatory agreement. The welfare changes are normalized so that the reducing country has an absolute welfare change equal to one.

In principle, both the excessive liberalization of tariffs and the additional tariff discrimination that results from discriminatory tariff bargaining could play a role in the poor performance of the No-MFN bargaining protocol when judged on the basis of world real income. To assess quantitatively the roles played by these two factors for the findings reported in Table 7, we recalculate the trade equilibrium, setting the discriminatory tariffs that each bargaining country applies to the other bargaining countries in each sector equal to that country’s average negotiated discriminatory tariff among the bargaining countries in that sector, and thereby eliminating the tariff discrimination among the negotiating countries that arose as a result of the discriminatory tariff bargaining. We find that world real income would have fallen by 0.01% relative to the 1990 status quo under this calculation, accounting for roughly 90% of the difference between the 0.06% rise in world real income relative to the 1990 status quo under MFN bargaining and the 0.02%

fall under No-MFN bargaining reported in Table 7. This suggests in turn that the bulk of the poor performance of the discriminatory tariff bargaining protocol is attributable to excessive liberalization, with only a relatively small part of the loss coming from the tariff discrimination among the negotiating countries that results from bargaining.

Turning to the impact of the abandonment of MFN on the distribution of welfare effects across countries, column 2 of Table 7 shows that most countries would have lost from the abandonment of MFN, with the average country experiencing a 0.21% reduction in welfare below its 1990 status quo level.

For comparison, columns 3 and 4 of Table 7 report the results of MFN and discriminatory bargaining respectively when all bargaining powers are set to one half. We include this alternative specification of bargaining powers for two reasons. The first reason is to explore robustness. As previously noted, the standard errors for many of our bargaining parameter estimates are large; hence, the specific country-level welfare changes identified in columns 1 and 2 must be interpreted in this light and thus with caution. Reassuringly, as a comparison of columns 1 and 2 with columns 3 and 4 confirms, our main findings are robust to the alternative specification for bargaining parameters. Specifically, under both specifications, most countries would have lost and average country welfare would have been lower in the described discriminatory tariff bargaining environment. Along the same lines, we evaluate the precision of the predicted difference in welfare between MFN and discriminatory bargaining using the bootstrapped parameter estimates. For world real income, the 95% confidence interval is [0.03%, 0.53%] for this difference, indicating that the finding that MFN leads to better aggregate outcomes than discriminatory bargaining is not sensitive to the imprecision in the bargaining parameter estimates.

The second reason to include the alternative specification for bargaining parameters in Table 7 is that, by comparing the magnitudes of welfare effects across the two specifications, we can provide support for the hypothesis that MFN mitigates the exercise of bargaining power. To this end, consider Japan, the US and Australia, the bargaining countries that are at the extremes of the distribution of welfare effects according to our estimates. As indicated in columns 1 and 2, according to our estimates, Japan would have experienced significant welfare gains while the US and Australia would have experienced significant welfare losses, were MFN abandoned. We interpret these predictions as indicating that under discriminatory tariff bargaining Japan would have been less constrained in exercising its strong bargaining power; and by the same token, the weak bargaining power of the US and Australia would then be a greater liability. The specification of equal

bargaining weights can be used to further develop this interpretation. As a comparison of column 3 to column 1 reveals, bargaining power has a modest impact on the bargaining outcomes under MFN; but comparing columns 4 and 2, it is clear that in the absence of MFN the countries that we estimate as having strong bargaining power such as Japan can more effectively exert their power at the expense of the countries which we estimate as having the weakest bargaining power such as the US and Australia.

To further aid in the interpretation of our findings, Table 10 reports the analogous information reported in Table 8, but for the No-MFN counterfactual. Column 1 of Table 10 reproduces the full equilibrium outcome of No-MFN bargaining according to our estimates, while columns 2 through 8 report the welfare changes associated with a single-bargain No-MFN counterfactual taking each of our seven bargaining pairs in turn, and column 9 reports the welfare changes that each bargaining country would have experienced as a result of the sum of the welfare effects of its single-pair No-MFN bilaterals.

Table 10: Discriminatory Bargaining Outcomes

	Discrim. Equil- ibrium (1)	Single Bargains US- EU (2)	US- Japan (3)	US- SK (4)	EU- Japan (5)	Japan- Aus (6)	Japan- SK (7)	Aus- SK (8)	Sum of Single Bargains (9)
US	-0.13%	0.00%	0.00%	0.00%	-0.03%	0.00%	-0.01%	0.00%	0.00%
EU	0.06%	0.10%	-0.03%	-0.01%	0.00%	0.00%	0.00%	0.00%	0.10%
Japan	0.32%	-0.02%	0.08%	-0.01%	0.14%	0.00%	0.06%	0.00%	0.27%
South Korea	0.27%	-0.14%	0.00%	0.46%	0.00%	0.00%	0.12%	0.01%	0.59%
Australia	-1.69%	-0.01%	-0.39%	0.00%	-0.24%	0.00%	-0.03%	0.08%	0.08%
Canada	-0.49%	-0.02%	-0.17%	-0.01%	-0.09%	0.00%	-0.01%	0.00%	
Africa NES	-0.08%	-0.03%	-0.03%	0.00%	0.00%	0.00%	0.00%	0.00%	
America NES	-0.05%	-0.02%	-0.01%	-0.01%	-0.01%	0.00%	0.00%	0.00%	
Asia NES	-0.37%	-0.09%	-0.07%	-0.07%	-0.03%	0.00%	-0.02%	-0.01%	
Europe NES	-0.12%	-0.03%	-0.03%	0.00%	-0.01%	0.00%	0.00%	0.00%	
MENA NES	-0.28%	-0.05%	-0.05%	0.01%	-0.06%	-0.01%	0.00%	0.00%	
Mean	-0.23%	-0.05%	-0.05%	-0.03%	-0.02%	0.00%	-0.01%	0.00%	-0.16%
World Real Income	-0.02%	0.00%	0.00%	0.01%	0.01%	0.00%	0.01%	0.00%	0.03%

Notes: Each column represents changes in the row relative to the pre-Uruguay tariff levels. The first set of columns represents the Horn-Wolinsky discriminatory equilibrium at the estimated bargaining parameters. The next columns represent the outcomes of single pair discriminatory bargains holding the other pair's at their 1990 levels. The mean across countries is weighted by population.

Strikingly, as column 9 of Table 10 reveals, if each of the bargaining countries were to consider its discriminatory bilaterals in isolation, that is, as a collection of single-pair discriminatory bargains, abandoning MFN would look like a winning proposition. The

EU, Japan, South Korea and even Australia would all expect to enjoy substantial gains relative to the gains they experience under MFN bargaining, with the US being essentially indifferent across the two protocols. It is only when the full equilibrium consequences of the abandonment of MFN are taken into account that the true losses suffered by the US, Australia, South Korea and the world as a whole become apparent. This is because under discriminatory tariff bargaining a beggar-thy-neighbor dynamic unfolds as the bilaterals compete against each other for market access, with each country in a given bilateral reducing its tariffs on its bargaining partner's imports in a discriminatory fashion in an attempt to give access to its bargaining partner's exporters at the expense of exporters from third-party countries.³⁷ As a comparison of columns 1 and 9 in Table ?? reveals, according to our estimates Japan is the only country that gains from the equilibrium interaction of other discriminatory bilaterals, and this helps account for its welfare gains in the No-MFN counterfactual relative to the Uruguay Round.³⁸

9 Conclusion

We embed a quantitative model of world trade into a model of bilateral bargaining over tariffs to examine the welfare effects of the most-favored-nation (MFN) requirement that characterizes negotiations at the GATT/WTO. We estimate the model using trade flows from 1990 and tariff outcomes from the Uruguay Round of GATT/WTO negotiations. In a trade model whose parameters are estimated to match observed trade flows, we quantify a free-rider issue that is created by the positive third-party externality from the MFN requirement and that leads to under liberalization relative to efficient outcomes. But our results indicate that the abandonment of MFN in tariff bargaining would create negative third-party externalities that are even more powerful and that would lead to

³⁷This competition in discriminatory tariff cutting is related to the “tariff complementarity effect” (see Bagwell and Staiger (1999b)), whereby a country that lowers in a discriminatory fashion a tariff on imports of a good from one partner has an incentive to lower the tariff as it applies to imports of that good from other partners as well.

³⁸The reason for this reflects an interaction between Japan's strong bargaining position in the Round and the tariff complementarity effect described in note 37. With its strong bargaining power, Japan is able to push down the discriminatory tariffs that its bargaining partners offer it, and through the tariff complementarity effect its bargaining partners are then softer bargaining partners for other countries, as they are more willing to lower their discriminatory tariffs for those other countries in those bilaterals; but this then makes them even softer bargaining partners for Japan.

substantial over liberalization. On balance we find for the Uruguay Round that MFN tariff negotiations are superior to counterfactual discriminatory tariff negotiations in terms of increasing world-wide welfare for this reason.

There are several promising avenues for future research. An obvious direction is to expand the current framework to allow for more products, that is, to handle a more disaggregated product classification. While this is essentially a computational challenge, it is an important extension, as actual tariff negotiations occur at a much more disaggregate level than the (essentially 2 digit) level that we have modeled here, and greater disaggregation could have important impacts on the principal supplier status that is central in shaping the bargaining patterns of the Round and associated externalities.

Also important is to consider the possibility that countries bargained under an additional constraint in the Uruguay Round, namely, that of reciprocity. We have presented here evidence on the degree to which our estimated bargaining outcomes conform to multilateral reciprocity. But we have not introduced bounds on permitted departures from multilateral reciprocity as a constraint into our bargaining protocol. How such an interval constraint would impact bargaining outcomes relative to those predicted here under the Horn-Wolinsky MFN tariff bargaining protocol is an open and important question.³⁹

Our modeling framework highlights bilateral tariff negotiations, which were a central feature of the Uruguay Round. The Round featured multilateral elements as well, however, such as “zero-for-zero” tariff negotiations in certain sectors. In addition, the negotiating countries agreed to a minimum target for overall tariff reductions to meet or exceed, and assessed progress relative to this target with multilateral reviews that supplemented the bilateral negotiations.⁴⁰ An interesting and important direction for future work is to build on the model of bilateral negotiations considered here to include additional multilateral elements. Related, modeling negotiations on non-tariff issue areas such as intellectual property rules would be of interest. And introducing political economy/distributional concerns into the objectives of governments should be high on the agenda.

³⁹As Bagwell and Staiger (2018) and Bagwell et al. (2020a) argue, the bargaining environment would change dramatically if multilateral reciprocity were strictly enforced so that no (or minimal) departures were permitted. We also note that there are several ways in which reciprocity finds representation in GATT/WTO rules and norms, one of which applies to renegotiations of tariff commitments and which can be interpreted as implying constraints on the tariff bargaining frontier which we have not considered here. See also Bagwell and Staiger (1999a).

⁴⁰For further discussion, see (Hoda, 2018, p 33-53).

References

- Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Bagwell, K. and R. W. Staiger (1999a). An economic theory of GATT. *American Economic Review* 89(1), 215–248.
- Bagwell, K. and R. W. Staiger (1999b). Regionalism and multilateral tariff co-operation. In J. Piggot and A. Woodland (Eds.), *International trade policy and the pacific rim*, pp. 157–190. Springer.
- Bagwell, K. and R. W. Staiger (2005). Multilateral trade negotiations, bilateral opportunism and the rules of GATT/WTO. *Journal of International Economics* 67(2), 268–94.
- Bagwell, K. and R. W. Staiger (2018). Multilateral trade bargaining and dominant strategies. *International Economic Review* 59(4), 1785–1824.
- Bagwell, K., R. W. Staiger, and A. Yurukoglu (2018). “Nash-in-Nash” tariff bargaining with and without mfn.
- Bagwell, K., R. W. Staiger, and A. Yurukoglu (2020a). Multilateral trade bargaining: A first look at the GATT bargaining records. *American Economic Journal: Applied Economics* 12(3), 72–105.
- Bagwell, K., R. W. Staiger, and A. Yurukoglu (2020b). “Nash-in-Nash” tariff bargaining. *Journal of International Economics* 122, 103263.
- Caliendo, L., R. C. Feenstra, J. Romalis, and A. M. Taylor (2017). Tariff reductions, entry, and welfare: Theory and evidence for the last two decades.
- Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies* 82(1), 1–44.
- Caner, M. and H. H. Zhang (2014). Adaptive elastic net for generalized methods of moments. *Journal of Business & Economic Statistics* 32(1), 30–47.
- Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019). “Nash-in-Nash” bargaining: a microfoundation for applied work. *Journal of Political Economy* 127(1), 163–195.

- Costinot, A., D. Donaldson, and I. Komunjer (2011). What goods do countries trade? a quantitative exploration of Ricardo's ideas. *The Review of Economic Studies* 79(2), 581–608.
- Crawford, G. S., R. S. Lee, M. D. Whinston, and A. Yurukoglu (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica* 86(3), 891–954.
- Crawford, G. S. and A. Yurukoglu (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review* 102(2), 643–685.
- Cremer, J. and M. H. Riordan (1987). On governing multilateral transactions with bilateral contracts. *The RAND Journal of Economics*, 436–451.
- Dagenais, M. G. (1983). Extension of the ridge regression technique to non-linear models with additive errors. *Economics Letters* 12(2), 169–174.
- Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: measuring the burden of adjustment. *IMF Economic Review* 55(3), 511–540.
- Dimaranan, B. and R. McDougall (2001). Global trade analysis and production: The gtap 5 data base. *Center for Global Trade Analysis, Purdue, University, West Lafayette, IN, USA. Available at: http://www.agecon.purdue.edu/GTAP/databases/v5/v5_doco.asp.*
- Dimaranan, B. V. and R. A. McDougall (2002). *Global Trade, Assistance, and Production: the GTAP5 Database*. Center for Global Trade Analysis, Purdue University.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Feenstra, R. C., R. E. Lipsey, and H. P. Bowen (1997). World trade flows, 1970-1992, with production and tariff data. Technical report, National Bureau of Economic Research.
- Feenstra, R. C., R. E. Lipsey, H. Deng, A. C. Ma, and H. Mo (2005). World trade flows: 1962-2000. Technical report, National Bureau of Economic Research.
- Fieleke, N. S. (1995). The uruguay round of trade negotiations: An overview. *New England Economic Review*, 3–15.

- Gandhi, A., Z. Lu, and X. Shi (2019). Estimating demand for differentiated products with zeroes in market share data. *Available at SSRN 3503565*.
- Gowrisankaran, G., A. Nevo, and R. Town (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review* 105(1), 172–203.
- Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review* 103(1), 145–177.
- Hart, O. and J. Tirole (1990). Vertical integration and market foreclosure. *Brookings Papers on Economic Activity. Microeconomics*, 205–286.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. *Handbook of International Economics* 4, 131.
- Ho, K. and R. S. Lee (2017). Insurer competition in health care markets. *Econometrica* 85(2), 379–417.
- Hoda, A. (2018). *Tariff Negotiations and Renegotiations under the GATT and the WTO: Procedures and Practices*. Cambridge University Press.
- Horn, H. and A. Wolinsky (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, 408–419.
- Karacaovali, B. and N. Limão (2008). The clash of liberalizations: Preferential vs. multilateral trade liberalization in the European Union. *Journal of International Economics* 74(2), 299 – 327.
- Lenzen, M., K. Kanemoto, D. Moran, and A. Geschke (2012). Mapping the structure of the world economy. *Environmental science & technology* 46(15), 8374–8381.
- Lenzen, M., D. Moran, K. Kanemoto, and A. Geschke (2013). Building eora: a global multi-region input–output database at high country and sector resolution. *Economic Systems Research* 25(1), 20–49.
- Limão, N. (2006). Preferential trade agreements as stumbling blocks for multilateral trade liberalization: Evidence for the united states. *American Economic Review* 96(3), 896–914.

- Limão, N. (2007). Are preferential trade agreements with non-trade objectives a stumbling block for multilateral liberalization? *The Review of Economic Studies* 74(3), 821.
- Mayer, T. and S. Zignago (2011). Notes on cepii's distances measures: The geodist database.
- McAfee, R. P. and M. Schwartz (1994). Opportunism in multilateral vertical contracting: Nondiscrimination, exclusivity, and uniformity. *American Economic Review*, 210–230.
- Novy, D. (2013). Gravity redux: measuring international trade costs with panel data. *Economic inquiry* 51(1), 101–121.
- Ossa, R. (2014). Trade wars and trade talks with data. *American Economic Review* 104(12), 4104–46.
- Ossa, R. (2015). Why trade matters after all. *Journal of International Economics* 97(2), 266–277.
- Ossa, R. (2016). Quantitative models of commercial policy. *Handbook of Commercial Policy* 1, 207–259.
- Silva, J. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics* 88(4), 641–658.
- Spearot, A. (2016). Unpacking the long-run effects of tariff shocks: New structural implications from firm heterogeneity models. *American Economic Journal: Microeconomics* 8(2), 128–167.

A Data Appendix

Trade Data

The main source of trade data is NBER-United Nations Trade Data. 1962-2000⁴¹ We supplement the 1995 Russian import data and the 2000 Indian import data with the Comtrade data. We aggregate the trade data up to the level of regional and product category used in the text. Our 49 traded product categories are defined in Table A1. Our first 13 traded product categories cover agriculture, with traded product categories 14-49 covering manufactures. Our 18 non-traded product categories are defined in Table A2.

Tariff Data

The tariff data is from the TRAINS data accessed through WITS⁴². We use the MFN applied ad valorem rate throughout the analysis. If the tariff data is not available for either 1990 or 2000, we borrow it from the closest year available. We then calculate the trade-weighted import tariff by the importing country (region) and the product category. For European countries, we calculate the euro-zone common import tariffs and apply to each country product-wise. For a given importing country (region) and a product category, if the import tariff is missing for a particular partner, we simply assume that the MFN tariff is applied to this partner.

Export Ratio

The ratio of exports to country level production is calculated using the GTAP 5 data (Dimaranan and McDougall, 2002), which provides the total production and the export for each country and sector in 1997. We then match the GTAP industries with our product classification to derive the export ratio by each product category.

Gravity Data & Preferential Trade Agreements

Gravity variables and the PTA relations between countries are from CEPII (Mayer and Zignago, 2011). For gravity variables, we use information on distance, GDP, and population. For distance between regions, we apply population weighted distance.

Intermediate share, non-traded sector, and final consumption

We download EORA purchaser-price measured input-output table for the USA in 1990 to construct our input-output share parameter γ 's. We match 428 reported industries into our industry classification that consists of 49 tradeable sectors and 18 non-tradeable sectors.

⁴¹<http://cid.econ.ucdavis.edu/nberus.html>

⁴²<http://wits.worldbank.org/>

Table A1: Traded Sector Classifications

Product Category	Corresponding SITC rev.2	Description
1	0	Live animals chiefly for food
2	1	Meat and meat preparations
3	2	Dairy products and birds'eggs
4	3	Fish,crustaceans,mollucs,preparations thereof
5	4	Cereals and cereal preparations
6	5,22	Vegetables and fruit; Oil seeds and oleaginous fruit
7	6	Sugar,sugar preparations and honey
8	7	Coffee,tea,cocoa,spices,manufactures thereof
9	8	Feeding stuff for animals,not incl.unmil.cereals
10	9	Miscel.edible products and preparations
11	11	Beverages
12	12	Tobacco and tobacco manufactures
13	21,61	Hides,skins and furskins,raw; Leather, leather manuf., n.e.s.and dressed furskisg
14	23	Crude rubber (including synthetic and reclaimed)
15	24	Cork and wood
16	25	Pulp and waste paper
17	26	Textile fibres (except wool tops) and their wastes
18	27,55,56,57	Crude materials; Essential oils & perfume mat.;toilet-cleansing mat; Fertilizers; Pyrotechnic products
19	28	Metalliferous ores and metal scrap
20	29	Crude animal and vegetable materials,n.e.s.
21	32	Coal,coke and briquettes
22	33,34	Petroleum,petroleum products and related; Gas,natural and manufactured materials
23	41,42,43	Animal oils and fats; Fixed vegetable oils and fats; Animal-vegetable oils-fats,processed,and waxes
24	51	Organic chemicals
25	52	Inorganic chemicals
26	53	Dyeing,tanning and colouring materials
27	54	Medicinal and pharmaceutical products
28	58	Artif.resins,plastic mat.,cellulose esters/ethers
29	59	Chemical materials and products,n.e.s.
30	62	Rubber manufactures,n.e.s.
31	63	Cork and wood manufactures (excl.furniture)
32	64	Paper,paperboard,artic.of paper,paper-pulp/board
33	65	Textile yarn,fabrics,made-upart.,related products
34	66	Non-metallic mineral manufactures,n.e.s.
35	67	Iron and steel
36	68,69	Non-ferrous metals; Manufactures of metal,n.e.s.
37	71	Power generating machinery and equipment
38	72,73,74	Machinery specialized for particular industries; Metalworking machinery; General industrial machinery & equipment,and parts
39	75,76	Office machines & automatic data processing; Telecommunications & sound recording apparatus equip.
40	77	Electrical machinery,apparatus & appliances n.e.s.
41	78	Road vehicles (incl. air cushion vehicles)
42	79	Other transport equipment
43	81	Sanitary,plumbing,heating and lighting fixtures
44	82	Furniture and parts thereof
45	83,84	Travel goods,handbags and similair containers; Articles of apparel and clothing accessories
46	85	Footwear
47	87,88	Professional,scientific & controling instruments ; Photographic apparatus,optical goods,watches
48	89	Miscellaneous manufactured articles,n.e.s.
49	90,91,93,94,95,96,97	Others

Table A2: List of Non-traded Sectors

Product Category	Description
50	Electricity, Gas and Water Supply
51	Construction
52	Sale, Maintenance and Repair of Motor Vehicles and Motorcycles; Retail Sale of Fuel
53	Wholesale Trade and Commission Trade, Except of Motor Vehicles and Motorcycles
54	Retail Trade, Except of Motor Vehicles and Motorcycles; Repair of Household Goods
55	Hotels and Restaurants
56	Inland Transport
57	Water Transport
58	Air Transport
59	Other Supporting and Auxiliary Transport Activities; Activities of Travel Agencies
60	Post and Telecommunications
61	Financial Intermediation
62	Real Estate Activities
63	Renting of M&Eq and Other Business Activities
64	Public Admin and Defence; Compulsory Social Security
65	Education
66	Health and Social Work
67	Other Community, Social and Personal Services

Using our constructed sectoral level trade, output (that we calculate from the NBER trade dataset and the GTAP dataset) and the sectoral output share that we calculate from the EORA dataset, we calculate sectoral gross expenditures by countries. Next we calculate the values of α 's that justify gross expenditures given the γ 's we found. Here we find that implied intermediate good expenditures in some industries exceed that of gross expenditures, leading to negative α 's. Therefore we slightly adjust the values of γ 's so that the calculated γ 's have the lowest sum of squared distance from the γ 's from the EORA data conditioning on that all α 's are positive. To make gains from trade comparable across countries, we assume that α 's of non-tradeable sectors are the same as the world average.