

MULTILATERAL TRADE BARGAINING AND DOMINANT STRATEGIES*

BY KYLE BAGWELL AND ROBERT W. STAIGER¹

Stanford University, U.S.A. and NBER, U.S.A.; Dartmouth College, U.S.A. and NBER, U.S.A.

Motivated by General Agreement on Tariffs and Trade bargaining behavior and renegotiation rules, we construct a three-country, two-good general-equilibrium model of trade and examine multilateral tariff bargaining under the constraints of nondiscrimination and multilateral reciprocity. For a general representation of government preferences, we identify the bargaining outcomes that can be achieved using dominant strategy proposals for all countries. In our analysis, dominant strategy outcomes emerge when tariff proposals are followed by multilateral rebalancing. The resulting bargaining outcome is efficient relative to government preferences if and only if the initial tariff vector positions the initial world price at its “politically optimal” level.

1. INTRODUCTION

Since 1947, the General Agreement on Tariffs and Trade (GATT) and its successor organization, the World Trade Organization (WTO), have provided the multilateral forum for bargaining over trade policy. The GATT was formed in 1947 among 23 original signatory countries and sponsored eight multilateral rounds of trade-policy negotiations. The final completed round, known as the Uruguay Round, resulted in the creation of the WTO on January 1, 1995. The WTO currently has more than 160 member countries and has struggled with its now-suspended Doha Round. But in combination, the GATT/WTO rounds surely represent one of the most important episodes of bargaining in economic history.

What accounts for the success of the GATT/WTO as a bargaining forum? We provide in this article a stylized model of multilateral tariff bargaining that embodies key institutional features of GATT/WTO practice. We argue that several of these features dramatically simplify the bargaining environment, and, in their presence, we show that all countries have dominant strategy proposals. We characterize the resulting bargaining outcomes, show that the associated sequence of initial proposals followed by “multilateral rebalancing” mimics stylized facts associated with GATT/WTO tariff bargaining, and describe conditions under which the bargaining outcomes are efficient.

The protocols for tariff negotiations in the GATT/WTO vary somewhat from round to round but have important common features. First, the negotiations are a form of barter: Each government makes commitments (offers “concessions”) on its own import tariffs in exchange for reciprocal commitments from its trading partners. Second, the negotiations are undertaken in the context of specific rules and norms. Under the principle of nondiscrimination as enshrined in GATT Article I, the tariff commitment that a country makes with respect to any given import good must be extended to all GATT/WTO countries.² Tariffs thus must satisfy a

*Manuscript received November 2016; revised December 2017.

¹ We thank Harold Cole and three anonymous referees for helpful comments. We also thank Matt Jackson, Gea Lee, and seminar participants at LSE, Stanford, and the 2017 Research Workshop on the Economics of International Trade Agreements in Villars for helpful discussions. We thank the NSF (Grant SES-1326940) for financial support. Bagwell thanks CASBS at Stanford for support and hospitality. Please address correspondence to: Kyle Bagwell, Department of Economics, Stanford University, Stanford, CA 94305, USA. E-mail: kbagwell@stanford.edu.

² GATT rules also include important exceptions to the principle of nondiscrimination; for example, GATT Article XXIV provides conditions for the formation of preferential trading agreements.

“most-favored nation” or MFN rule. In addition, reciprocity rules and norms shape the pattern of negotiation.

A primary expression of GATT rules concerning reciprocity is found in GATT Article XXVIII, which addresses rules for renegotiation. Under this article, after negotiations are completed, a country retains the right to withdraw a tariff commitment and reposition an import tariff at a higher level, with the understanding that its principal trading partners are then allowed to behave in a reciprocal manner and withdraw “substantially equivalent concessions.” Thus, GATT rules ensure that a form of reciprocity is followed when concessions are renegotiated, and this rule is, of course, known to participating countries at the time of negotiation. As well, when tariffs are originally negotiated within a round, a reciprocity norm shapes the expectations of negotiators and thus the negotiating outcome.

Following Bagwell and Staiger (1999, 2002), we regard the principle of reciprocity as being satisfied when two countries exchange tariff reductions (or tariff increases, in the case of renegotiation) such that each country experiences changes in the volume of its imports that are equivalent in magnitude to the changes in the volume of its exports, with the changes in trade volumes valued at existing world prices. In GATT parlance, reciprocal tariff liberalization then facilitates a “balance of concessions.”³ As Bagwell and Staiger (1999, 2002) show for a two-good general equilibrium model of trade, when two countries make tariff changes that satisfy the principle of reciprocity, the changes leave the world price that governs trade between the two countries unaltered. In this way, when applied in the context of original negotiations, reciprocity prevents countries from manipulating their terms of trade and thereby neutralizes the fundamental source of inefficiency in noncooperative tariff setting. Moreover, when applied in the context of renegotiations, reciprocity suggests that no country can be forced to accept a trade volume in excess of its preferred volume at the fixed terms of trade. Bagwell and Staiger (2005) show further that, in a three-country, two-good general-equilibrium setting, if two countries negotiate in a manner that satisfies both the principles of nondiscrimination and reciprocity, then the preservation of the terms of trade between the countries ensures as well the absence of any third-party externality.

The notion of reciprocity studied by Bagwell and Staiger (1999, 2002, 2005) may be understood as a form of bilateral reciprocity. We focus in this article on a related but distinct notion of multilateral reciprocity. For the three-country setting and under the MFN rule, we regard the principle of multilateral reciprocity as being satisfied when the three countries undertake any combination of tariff changes that leaves the world price unaltered. Multilateral reciprocity holds when two countries change their respective tariffs in a way that satisfies bilateral reciprocity. Multilateral reciprocity is a more general notion, however, since it can hold even if a country negotiates separately with each of two trading partners where each individual negotiation violates bilateral reciprocity. This happens when the individual negotiations push the world price in different directions, with one movement exactly offsetting the other so that the combined effect leaves the world price unaltered. As Bagwell et al. (2017) argue for the three-country model under the MFN rule, by fixing the terms of trade, multilateral reciprocity neutralizes the key source of inefficiency under noncooperative tariff setting.

The important role of multilateral as opposed to bilateral reciprocity was emphasized in early writings on GATT negotiations. The key point is that a bilateral exchange of tariff cuts is “multilateralized” under the MFN rule and may offer indirect benefits to other parties, where such indirect gains can be more effectively internalized in a multilateral bargaining forum.⁴ GATT commentators routinely express the view that the facilitation of multilateral as opposed to bilateral reciprocity made possible by GATT’s multilateral bargaining forum was a key

³ For further discussion of the principle of reciprocity in GATT/WTO and recent research related to this topic, see Bagwell and Staiger (2016).

⁴ Note that, according to the findings of Bagwell and Staiger (2005) just described, bilateral tariff cuts under the MFN rule offer potential indirect benefits to third parties only when those cuts fail to satisfy bilateral reciprocity (i.e., only when those cuts on their own would change world prices).

institutional innovation underlying GATT's success (see, for example, Interim Commission for the International Trade Organisation [ICITO], 1949, p. 10, and Curzon, 1966, pp. 75–77).

Utilizing recently declassified data from the GATT/WTO on tariff bargaining, Bagwell et al. (2017) study the pattern of tariff bargaining in the GATT Torquay (1950–51) Round. Negotiations in this round took a “request-offer” form, whereby the initial proposal of a country consisted of the tariff cuts it requested from its trading partner and the tariff cuts it offered in exchange. For our purposes here, we highlight three key findings from their study. First, the numbers of back-and-forth offers and counteroffers in any bargain were relatively small. Second, once the initial proposals were on the table, the focus of bargaining narrowed to each country's own tariff-cut offers. Countries responded to imbalances in the outstanding offers by adjusting their own offers instead of by adjusting their requests. Third, adjustments in offers typically took a simple and striking form. Offers for given import goods were rarely deepened as the round progressed, suggesting an absence of strategic screening behavior along this dimension.⁵ Instead, when adjustments in offers did occur, the adjustment typically involved a country reducing the depth of its offer.⁶ A potential interpretation of this pattern is that a country would propose for a given import good the tariff that generated its preferred trade volume for a fixed terms of trade, with the expectation that any subsequent “rebalancing” of offers necessary for multilateral reciprocity would arise later in the round after all offers had been recorded and might entail a reduction in the depth of its offer.

In this article, motivated by GATT bargaining behavior and renegotiation rules, we construct a general-equilibrium model of trade and examine multilateral tariff bargaining under the constraints of nondiscrimination and multilateral reciprocity. The model has two goods and three countries, corresponding to a home country and two foreign countries, where the foreign countries trade only with the home country. For a general representation of government preferences, we construct a simple bargaining game and identify bargaining outcomes that can be achieved using dominant strategy proposals for all countries. The resulting bargaining outcome is efficient relative to government preferences if and only if the initial tariff vector positions the initial world price at its “politically optimal” level.⁷ In symmetric settings, if the initial tariffs correspond to Nash tariffs, then the resulting bargaining outcome also ensures greater-than-Nash trade volumes and welfares for all countries.

To establish these results, we develop a bargaining game inspired by the request-offer structure that we described above in the context of the Torquay Round, a structure that has been commonly used across the GATT/WTO tariff bargaining rounds.⁸ In this game, countries simultaneously make proposals concerning their own tariffs (their “offers”) and the tariffs of their trading partners (their “requests”). For each country, we require that any proposed change in tariffs satisfy the principles of nondiscrimination and multilateral reciprocity. We capture GATT Article XXVIII renegotiation provisions in a short-hand way, by assuming also that the bargaining outcome must respect “voluntary exchange,” in the specific sense that no country is ever forced to accept a trade volume in excess of that implied by its proposal. A key issue is that the proposals may disagree and thus be imbalanced, with one side of the market seeking greater trade volume than the other. To address this issue, we construct a mechanism that maps

⁵ Strategic screening occurs, for instance, in theoretical models with one-sided uncertainty in which the uninformed party makes all the offers (e.g., Gul et al., 1986).

⁶ A country could reduce the depth of its offer on the intensive margin (by reducing the magnitude of its tariff cut on a given import good) or on the extensive margin (by reducing the range of import goods in a bargain). In the Torquay Round, rebalancing typically occurred on the extensive margin. In the model that we develop here, each country has one import good, and so we focus on intensive-margin adjustments.

⁷ The politically optimal tariffs are the tariffs that would be selected in the hypothetical situation in which governments are not motivated by the terms-of-trade implications of their unilateral trade-policy choices. The politically optimal world price is then the market-clearing world price that emerges when all governments select their politically optimal policies. See Bagwell and Staiger (1999, 2002) and Section 5 below for further discussion.

⁸ With the exception of the Kennedy and Tokyo Rounds, which relied primarily on tariff-cutting formulas, all of the GATT rounds and the now-suspended WTO Doha Round have made use of request-offer bargaining structures of the kind we describe here.

tariff proposals into assigned tariffs. When the proposals “agree,” as we define that term, the mechanism assigns the corresponding tariff vector upon which all parties agree. If the proposals disagree, the mechanism must introduce a rationing rule, since one side of the market is short relative to the other. The rationing rule that we employ uses randomization and ensures that foreign countries subjected to rationing are treated in an *ex ante* symmetric way. We construct a mechanism that maximizes the value of trade volume while satisfying the constraints of nondiscrimination, multilateral reciprocity, and voluntary exchange. We then show that the mechanism so constructed results in dominant strategies for each country, and we characterize the outcomes that can be achieved under dominant strategy proposals for each country.

In our model, dominant strategy proposals lead to assigned tariff vectors once adjustments are made that maintain multilateral reciprocity while not requiring any country to import more than is implied by its proposal. This simple sequence—initial tariff proposals followed by multilateral rebalancing—is broadly consistent with observed patterns identified by Bagwell et al. (2017) in the bargaining records for the GATT Torquay Round. As mentioned above, a notable feature of bargaining behavior in this round is the lack of back-and-forth haggling over the levels of proposed tariffs. In addition, subsequent to the initial proposals, our constructed mechanism orchestrates multilateral rebalancing by placing primary emphasis on the own-tariff offers in those proposals, much as did the narrowed focus of bargaining in the Torquay Round once initial proposals were on the table. Also, in line with the observed patterns from the Torquay Round, multilateral rebalancing is obtained through reductions in the depths of offers.⁹

Despite the importance of bargaining in GATT/WTO rounds, relatively little work has formally examined multilateral tariff bargaining behavior. As mentioned above, Bagwell and Staiger (1999, 2002, 2005) examine the purpose and design of trade agreements for a related setting. As we discuss in further detail in Section 2, some of the key implications of (bilateral) reciprocity and MFN that we utilize here are developed in their work. Our theoretical work is most closely related to that in Bagwell and Staiger (1999). In a related multicountry setting, Bagwell and Staiger (1999) study a tariff negotiation game in which countries make proposals consistent with reciprocity. Under disagreement, they assume that the assigned tariff vector maximizes the value of trade while not requiring any country to import a volume in excess of that implied by its proposal. Their analysis is at one level more general, since they also allow for discriminatory tariffs.¹⁰ At other levels, however, the model considered in this article offers several advantages.

In particular, Bagwell and Staiger (1999) assume that the home country proposes its own tariff policy as well as trade shares from foreign countries, and they then use this proposal and the reciprocity restriction to determine the home country’s implied proposals for foreign tariffs. Similarly, Bagwell and Staiger (1999) assume that each foreign country proposes a tariff for itself, and they then use all foreign tariff proposals and the reciprocity restriction to determine the implied proposal by foreign countries together for home tariffs. By contrast, the model that we present here relates more directly to the GATT request-offer tariff bargaining data, since we assume that each country directly proposes a tariff for itself and for each of its trading partners. Instead of endowing the home country with the power to choose trade shares directly, we propose an explicit rationing rule for settings in which one side of the market is short.¹¹ Another key difference is that the home country in the Bagwell–Staiger (1999) model adopts best-response instead of dominant proposals, whereas in the results featured here, all countries adopt dominant strategy proposals. We also show that one of Bagwell and Staiger’s (1999) main

⁹ As previously noted, however, in our two-good model, we are not able to include rebalancing that takes the form of changes in the range of products that are included in the bargain.

¹⁰ Bagwell and Staiger (1999) find that discriminatory policies are incompatible with efficiency for the game that they analyze.

¹¹ Although we focus on a probabilistic rationing rule in our main analysis, we consider other rationing rules in Section 7. As we discuss just below, one advantage of this approach is that we can relate our analysis to the literature on strategy-proof social choice functions, which considers alternative rationing rules that are compatible with dominant strategy implementation.

findings—that efficiency can be achieved if and only if the politically optimal MFN tariff vector is induced—continues to hold when we generalize their analysis so that all countries make direct tariff proposals and use dominant strategy proposals.

Ludema's (1991) work is also related. He offers an interesting model of tariff bargaining under the MFN rule in a dynamic setting. His focus is on whether an MFN-efficient bargaining outcome is possible when countries can reject an outcome and continue bargaining in the event that a country attempts to free ride and withholds its own tariff cuts. By comparison, our bargaining setup is essentially static, imposes the additional constraint of multilateral reciprocity, addresses associated rationing issues, and features dominant strategy outcomes.

Our work is also broadly related to previous research that examines strategy-proof rationing rules in other economic settings. In our model, a rationing issue arises at the given world price when the volume of trade proposed by the home country is lower than that proposed by the foreign countries in aggregate. We then require a rationing rule that is consistent with dominant strategy behavior. Our approach is to pick a foreign country at random, assign to that country its proposed volume so long as its proposed volume does not exceed the volume proposed by the home country, and then allocate any remaining volume to the other foreign country. Each country has a preferred volume of trade given the fixed world price, where preferences are single-peaked, and as noted no country is ever forced to accept a trade volume in excess of that implied by its proposal. A potentially attractive feature of this random priority rationing rule is that foreign countries are treated in an *ex ante* symmetric way. However, asymmetric rationing rules that use fixed (i.e., deterministic) priority schemes may also be attractive. For example, in settings where one trading partner is a primary supplier, the GATT principal supplier rule suggests that this partner receives priority. We consider such asymmetric rationing rules in Section 7 and argue that our results are robust to this extension.

Strategy-proof rationing rules are also examined by Benassy (1982) in the context of an analysis of fixed-price equilibria. Benassy assumes that each agent presents a preferred net demand volume in a given market, and he then considers rationing schemes that map from the set of net demand requests presented by all agents to actual transactions. He emphasizes deterministic rationing rules that are strategy-proof, efficient (in that there are not both rationed demanders and suppliers in a given market), and respect voluntary exchange (in that no agent can be forced to purchase more units than that agent demands or sell more units than that agent supplies). One scheme considered by Benassy that satisfies these properties is a priority or queuing scheme, where demanders (or suppliers) are ordered in a predetermined fashion. The setting that we consider is similar in some respects; however, we focus on a barter environment and use a random process to determine which foreign country is prioritized to be first in line. Given our tariff bargaining setting, we also assume that each country proposes tariffs for itself *and* its trading partner(s), subject to the institutional constraints of nondiscrimination and multilateral reciprocity. The relevant transaction price in our model is the world price, which is fixed under these institutional constraints.

Strategy-proof rationing rules are also used in the social choice literature that considers dominant strategy implementation when the domain of possible preferences is restricted in various ways.¹² By contrast, we adopt particular tariff assignment rules that are motivated by GATT/WTO practice, show that those rules are consistent with dominant strategy proposals, and consider the positive and normative implications of the resulting bargaining outcomes. Despite these differences, our analysis of dominant strategy outcomes and rationing rules has some interesting parallels in the social choice literature, wherein a related allotment problem is concerned with rationing volumes across agents with single-peaked preferences in a fixed-price setting with an exogenous total volume.¹³ In this context, the rationing rule that we

¹² See Barbera (2011) for an excellent survey.

¹³ It is important to note that differences exist as well. In the model considered here, the total volume that is allocated when the foreign countries are rationed is endogenously determined by the home-country proposal, and the world price is fixed as a consequence of the institution-based constraints (nondiscrimination, multilateral reciprocity) that we impose on the proposal strategy space.

employ when the home country is on the short side of the market is closely related to that used in the random priority mechanism.¹⁴ As we discuss further in Section 7, our findings also hold for other rationing rules featured in research on the allotment problem, including asymmetric rules that use fixed priority schemes and a rule based on the “uniform allocation rule.”¹⁵

Before proceeding to our formal analysis, we pause to comment on our general modeling approach. In our main analysis, we assume that government preferences are described by a fixed preference vector, focus on dominant strategy proposals, and impose a specific mechanism. Our assumption of a fixed preference vector follows Bagwell and Staiger (1999, 2002) and, as we argue in Section 7, is convenient but inessential. Given our constructed mechanism and dominant strategy solution concept, with some additional notation, our findings may be easily translated to a related private values setting in which the vector of government preferences is drawn from a space of possible preferences and where governments may be privately informed about their respective realized preferences.¹⁶ Second, we note that an attractive feature of the dominant strategy solution concept is that it does not require governments to correctly forecast each other’s strategies.¹⁷ Finally, we emphasize that our primary goal is not to construct an optimal mechanism; rather, we seek to explore the predictive and normative implications of a specific mechanism that is constructed to reflect important rules and norms of the GATT institution and bargaining process.¹⁸

The rest of the article is organized into seven remaining sections, which, respectively, (i) present the basic general-equilibrium trade model and its key properties, (ii) construct a mechanism that maps tariff proposals satisfying MFN and multilateral reciprocity into assigned tariffs, (iii) characterize dominant strategy proposals for each country, (iv) examine the efficiency properties of the resulting bargaining outcomes, (v) examine the resulting bargaining outcomes when the initial tariffs are Nash tariffs, (vi) discuss extensions concerning multiple countries, alternative rationing rules, additional constraints, and preference spaces and private information, and (vii) conclude.

2. FRAMEWORK

We begin by presenting a general-equilibrium model of trade. The model has two goods and three countries. After presenting the model, we provide a general representation of government preferences, describe Nash tariffs, and consider notions of reciprocity in trade negotiations.

2.1. The Trade Model. Three countries trade two goods, where the goods are normal in consumption and produced in perfectly competitive markets under conditions of increasing opportunity costs. The three countries are, respectively, referred to as the home country (or home), foreign country *1, and foreign country *2. Throughout, we denote foreign country variables with an asterisk. The home country imports good x from each of the two foreign

¹⁴ See Abdulkadiroglu and Sonmez (1998) and Bogomolnaia and Moulin (2001) for further discussion of the random priority mechanism (which the former paper refers to as a random serial dictator mechanism), as applied to settings in which indivisible objects are assigned to agents where each agent has use for only one unit.

¹⁵ In the social choice literature, priority schemes and other asymmetric rationing approaches are considered by Barbera et al. (1997) and Moulin (2000), for example, whereas the uniform allocation rule is analyzed by Sprumont (1991).

¹⁶ As we discuss in Section 7, with this extension, we can relate our results more directly to the class of problems studied in implementation theory.

¹⁷ When government preferences are drawn from a space of possible preferences, the dominant strategy solution concept also has the advantage of being robust with respect to the specification of the probability density over preference vectors and even the beliefs that governments hold about this specification.

¹⁸ Implicit in our approach is a view shared with a larger literature that institutions (or mechanisms) may be long-lived and embody persistent rules and norms that are costly to change. See, for example, Jackson (2001, pp. 661–2).

countries, and each foreign country imports good y from the home country.¹⁹ We assume that the foreign countries do not trade with one another.

Each country selects an ad valorem import tariff. Although the home country imports from two countries, we assume that it uses an MFN tariff and thus does not discriminate between imports from different foreign countries. With $t > -1$ denoting the home-country ad valorem import tariff and $t^{*i} > -1$ denoting the ad valorem import tariff of foreign country $*i$, we may define $\tau \equiv 1 + t > 0$ and $\tau^{*i} \equiv 1 + t^{*i} > 0$. We interpret $\tau > 1$ as an import tariff and similarly for τ^{*i} .

Our approach is to present the basic trade model under the assumption that the tariff vector $(\tau, \tau^{*1}, \tau^{*2})$ is such that the home country and each foreign country exchange a positive volume of trade. After describing the model under this assumption, we introduce notation to define the associated set of tariff vectors. Finally, we present an assumption under which we may extend the set of tariff vectors to include the possibility of a prohibitive tariff for one foreign country.

To present the basic trade model, we let $p \equiv p_x/p_y$ denote the local price faced by producers and consumers in the home country, and we similarly use $p^{*i} \equiv p_x^{*i}/p_y^{*i}$ to denote the local price in foreign country $*i$, $i = 1, 2$, respectively. The world (i.e., untaxed) price for trade between the home country and foreign country $*i$ is defined as $p^{wi} \equiv p_x^{*i}/p_y$. The local price in the home country is determined from the world price as $p = \tau p^{wi}$. Clearly, since the home country adopts an MFN tariff τ , a single world price must prevail: $p^w \equiv p^*/p_y$. The local price in foreign country $*i$ is now determined as $p^{*i} = p^w/\tau^{*i}$, where the local price varies across foreign countries if they select different import tariffs. Observe next that p^w represents foreign country $*i$'s terms of trade and that the home country's terms of trade are similarly given by $1/p^w$. Finally, it is convenient to introduce functional notation that summarizes the relationships between local and world prices: $p = \tau p^w \equiv p(\tau, p^w) > 0$ and $p^{*i} = p^w/\tau^{*i} \equiv p^{*i}(\tau^{*i}, p^w) > 0$.

In each country, the production levels of goods x and y are determined by the local price. Let $Q_g = Q_g(p)$ and $Q_g^{*i} = Q_g^{*i}(p^{*i})$ denote production levels for good g , where $g = x, y$, in the home country and in foreign country $*i$, respectively. Consumption in each country is determined by the local price in that country along with the world price: $C_g = C_g(p, p^w)$ and $C_g^{*i} = C_g^{*i}(p^{*i}, p^w)$ for $g = x, y$ and $i = 1, 2$, respectively. Within any country, the local price defines the trade-off that confronts consumers and also determines the level and distribution of factor income. Together, the local and world prices also determine tariff revenue for the country, which we assume is distributed lump sum to consumers.²⁰ Hence, for each good in each country, consumption is determined by the local price in that country along with the world price.

We now define import and export volumes. For the home country, imports of good x and exports of good y are, respectively, denoted as $M(p, p^w) \equiv C_x(p, p^w) - Q_x(p)$ and $E(p, p^w) \equiv Q_y(p) - C_y(p, p^w)$. Similarly, for foreign country $*i$, imports of good y and exports of good x are, respectively, denoted as $M^{*i}(p^{*i}, p^w) \equiv C_y^{*i}(p^{*i}, p^w) - Q_y^{*i}(p^{*i})$ and $E^{*i}(p^{*i}, p^w) \equiv Q_x^{*i}(p^{*i}) - C_x^{*i}(p^{*i}, p^w)$. We assume that these functions are all twice-continuously differentiable given positive trade volumes.

For a given world price, we also assume that each country imports less of its import good when the relative price of the import good rises in that country. Formally, for a given world price, we assume that $\partial M(p, p^w)/\partial p < 0$ and that $\partial M^{*i}(p^{*i}, p^w)/\partial p^{*i} > 0$. We assume further that, for a given world price, the import volume for a given country can be made arbitrarily small as the tariff of that country is made sufficiently high. Formally, for any p^w and values $\underline{M} > 0$ and $\underline{M}^{*i} > 0$, we assume that there exists $\tau(\underline{M}) > 0$ and $\tau^{*i}(\underline{M}^{*i}) > 0$ such that $M(p(\tau, p^w), p^w) < \underline{M}$ for all $\tau > \tau(\underline{M})$ and $M^{*i}(p^{*i}(\tau^{*i}, p^w), p^w) < \underline{M}^{*i}$ for all $\tau^{*i} > \tau^{*i}(\underline{M}^{*i})$.²¹ Notice that this assumption

¹⁹ We assume throughout that trade policies never reverse the “natural” direction of trade; thus, the home country never imports good y and the foreign countries never import good x .

²⁰ See Bagwell and Staiger (1999, 2002) for further details concerning the determination of tariff revenue.

²¹ The role of this assumption is to ensure that we can always assign tariffs to generate the import values that arise in our constructed mechanism in Section 3.

does not imply for any country the existence of a prohibitive tariff level that generates zero import volume.

Trade volumes are subject to two market relationships. First, for any world price, trade balance must be satisfied in each country. For the home country, the trade-balance condition is that, for any p^w ,

$$(1) \quad p^w M(p(\tau, p^w), p^w) = E(p(\tau, p^w), p^w).$$

Similarly, for foreign country $*i$, the trade-balance condition is that, for any p^w ,

$$(2) \quad M^{*i}(p^{*i}(\tau^{*i}, p^w), p^w) = p^w E^{*i}(p^{*i}(\tau^{*i}, p^w), p^w).$$

The trade-balance conditions are budget constraints that are captured as restrictions on the import and export functions in a given country.

The second relationship is market clearing. We thus define the equilibrium world price, $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2})$, as the unique value for p^w that satisfies market clearing in good y :

$$(3) \quad E(p(\tau, p^w), p^w) = M^{*1}(p^{*1}(\tau^{*1}, p^w), p^w) + M^{*2}(p^{*2}(\tau^{*2}, p^w), p^w).$$

As is standard, market clearing for good x is implied by the trade-balance conditions along with the requirement of market clearing for good y .

As noted, we have presented the basic trade model under the assumption that the tariff vector is such that a positive trade volume is exchanged between the home country and each foreign country. We now introduce the notation Υ to define the set of tariff vectors for which trade volumes are positive at the associated market-clearing world price:

$$\Upsilon \equiv \{(\tau, \tau^{*1}, \tau^{*2}) \in \mathfrak{R}_+^3 \mid M^{*i}(p^{*i}(\tau^{*i}, \tilde{p}^w), \tilde{p}^w) > 0, i = 1, 2\}.$$

Notice that, for $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$, we have from (1) and (3) that $M(p(\tau, \tilde{p}^w), \tilde{p}^w) > 0$ holds as well. Thus, for $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$, the assumptions presented above for the basic trade model all hold. We assume that Υ is a nonempty set.

Our presentation to this point assumes positive trade volumes between the home country and each foreign country, but our trade-balance and market-clearing conditions above can also be used to capture limiting cases where one foreign country has no trade volume.²² Although we do not assume that prohibitive tariffs exist as a general matter, we do assume the existence of prohibitive tariffs for individual foreign countries under one specific scenario.

To formalize this assumption, suppose we start with some initial tariff vector $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$ and associated initial world price $p_0^w \equiv \tilde{p}^w(\tau_0, \tau_0^{*1}, \tau_0^{*2})$. Suppose further that for some foreign country $*i$, it is possible to satisfy the market-clearing condition (3) at the world price p_0^w when the trade volume for foreign country $*j$ is set to 0; that is, suppose that there exists τ_1^{*i} such that $E(p(\tau_0, p_0^w), p_0^w) = M^{*i}(p^{*i}(\tau_1^{*i}, p_0^w), p_0^w)$. For such initial tariff vectors, our assumption is that there exists a lowest prohibitive tariff $\tau_1^{*j} > \tau_0^{*j}$ for foreign country $*j$ such that $M^{*j}(p^{*j}(\tau_1^{*j}, p_0^w), p_0^w) = 0$ with $M^{*j}(p^{*j}(\tau^{*j}, p_0^w), p_0^w) > 0$ for all $\tau^{*j} \in [\tau_0^{*j}, \tau_1^{*j}]$.²³ We then use Υ_+ to denote the extension of the set Υ to include limit-case tariff vectors $(\tau_0, \tau_1^{*1}, \tau_1^{*2})$ that can be constructed in this way from some initial tariff vector $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$. We also extend the definition of the market-clearing world price to include such limit-case vectors and thus write $p_0^w = \tilde{p}^w(\tau_0, \tau_1^{*1}, \tau_1^{*2})$ even though one foreign country has no trade volume.

²² Notice from (1) and (3) that if for foreign country $*1$, say, we have $M^{*1}(p^{*1}(\tau^{*1}, \tilde{p}^w), \tilde{p}^w) = 0$ for $\tilde{p}^w = \tilde{p}^w(\tau, \tau^{*1}, \tau^{*2})$, then we may still have that $M(p(\tau, \tilde{p}^w), \tilde{p}^w) > 0$ and $M^{*2}(p^{*2}(\tau^{*2}, \tilde{p}^w), \tilde{p}^w) > 0$.

²³ In our analysis below, this assumption ensures that we can assign tariff vectors under which one foreign country receives no trade volume.

We turn now to our assumptions on market-clearing prices. Our maintained assumptions are that Metzler and Lerner paradoxes are ruled out; thus, for $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$, we assume that

$$(4) \quad \begin{aligned} \frac{dp(\tau, \tilde{p}^w)}{d\tau} &> 0 > \frac{dp^{*i}(\tau^{*i}, \tilde{p}^w)}{d\tau^{*i}}, \\ \frac{\partial \tilde{p}^w(\tau, \tau^{*1}, \tau^{*2})}{\partial \tau} &< 0 < \frac{\partial \tilde{p}^w(\tau, \tau^{*1}, \tau^{*2})}{\partial \tau^{*i}}. \end{aligned}$$

The world-price inequalities in (4) ensure that each country is large (i.e., that each country can improve its terms of trade by raising its import tariff).

2.2. Government Preferences. We follow Bagwell and Staiger (1999, 2002) and assume that government preferences are fixed and represented in a general way that is consistent with the traditional case in which governments maximize national income as well as the case in which governments have distributional concerns. Specifically, for a given vector of tariffs, the home government welfare function is represented as $W(p(\tau, \tilde{p}^w), \tilde{p}^w)$, whereas the welfare function for the government of foreign country $*i$ is represented as $W^{*i}(p^{*i}(\tau^{*i}, \tilde{p}^w), \tilde{p}^w)$. For simplicity, we refer to W as the welfare of the home country and to W^{*i} as the welfare of foreign country $*i$.

Our primary assumption on welfare functions is that, holding its local price fixed, each government values an improvement in its terms of trade:

$$(5) \quad W_{\tilde{p}^w}(p(\tau, \tilde{p}^w), \tilde{p}^w) < 0 < W_{\tilde{p}^w}^{*i}(p^{*i}(\tau^{*i}, \tilde{p}^w), \tilde{p}^w)$$

for $i = 1, 2$ and for $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$.²⁴ As Bagwell and Staiger (1999, 2002) discuss in detail, the experiment considered here may be understood as corresponding to underlying tariff changes whereby a government raises its own tariff while a trading partner reduces its tariff, with the end result being that the government’s local price is unchanged while its terms of trade improve. This assumption is satisfied when governments maximize national income and also in the leading political-economy models of trade policy, as Bagwell and Staiger (1999, 2002) discuss.

2.3. Nash Tariffs. In the absence of a trade agreement, we assume that governments would each set their optimal unilateral policies, leading to a Nash equilibrium. To provide additional context for our analysis, we define and briefly characterize the Nash tariffs.

The optimal or best-response tariffs for the home country and foreign country $*i$ are, respectively, defined by

$$(6) \quad \begin{aligned} W_p \frac{dp}{d\tau} + W_{\tilde{p}^w} \frac{\partial \tilde{p}^w}{\partial \tau} &= 0, \\ W_{p^{*i}} \frac{dp^{*i}}{d\tau^{*i}} + W_{\tilde{p}^w}^{*i} \frac{\partial \tilde{p}^w}{\partial \tau^{*i}} &= 0, \end{aligned}$$

where $i = 1, 2$ and for simplicity, we suppress notation for functional dependencies. A *Nash equilibrium* is a tariff vector, $(\tau_N, \tau_N^{*1}, \tau_N^{*2}) \in \Upsilon$, satisfying these three first-order conditions. We assume that a unique Nash equilibrium exists, and we refer to the associated tariffs as *Nash tariffs*.²⁵

²⁴ For a given world price, we also impose an assumption that each country has single-peaked preferences with respect to its own tariff (or local price). We postpone a formal statement of this assumption until Section 4.

²⁵ As is well known (see Dixit, 1987), autarky Nash tariff equilibria typically exist as well. Our focus here is on characterizing interior Nash tariffs, which seems the natural focus for our purposes given that, in Section 6, we consider the possibility that these tariffs are the initial tariffs from which countries bargain.

Bagwell and Staiger (1999, 2002) show that the Nash tariffs are inefficient, where efficiency is measured relative to the preferences of governments (i.e., relative to the welfare functions, W , W^{*1} , and W^{*2}). The key intuition is that each government is motivated in part by the terms-of-trade implications of its trade-policy selection, and a terms-of-trade gain for the home country is a terms-of-trade loss for the foreign countries. As Bagwell and Staiger (1999, 2002) discuss, this means that starting from the Nash tariffs, each government would gain from a small increase in its trade volume (i.e., a small decrease in the relative price of its import good) if this could be attained without altering its terms of trade.

The formal argument is useful for our subsequent discussion and is as follows. Using (4)–(6), we may easily verify that $W_p < 0 < W_{p^{*i}}$ at the Nash tariffs. Now suppose that the home country and foreign country $*i$ exchange small reciprocal tariff cuts that preserve the terms of trade. Since $p = \tau^w$ and $p^{*i} = p^w/\tau^{*i}$, it then follows that p falls and p^{*i} rises. With $W_p < 0 < W_{p^{*i}}$ at the Nash tariffs, we thus conclude that the home country and foreign country $*i$ both enjoy welfare gains from exchanging small reciprocal tariff cuts that preserve the terms of trade.

A government cannot alter its local price and preserve its terms of trade with a unilateral trade-policy adjustment; however, under (4) and as we discuss in greater detail next, governments can attain trade-volume increases without altering the terms of trade if they liberalize trade through reciprocal adjustments in trade policies.

2.4. Reciprocity. We now explore notions of reciprocal adjustments in trade policies. Motivated by the preceding discussion, we are interested in the effects of reciprocal adjustments on negotiating partners and also third parties.

To begin our discussion, we follow Bagwell and Staiger (2005) and define a *welfare-preservation property* for the model. Specifically, for tariff vectors in Υ , if the home country and foreign country $*i$ negotiate changes in their respective tariffs that leave unaltered the world price, then the welfare of the government of foreign country $*j$, where $j \neq i$, is unaltered as well. This implication is easily verified. If the negotiation between the governments of home and foreign country i leaves the world price \tilde{p}^w fixed at some level p_0^w , and if the tariff τ^{*j} of foreign country $*j$ does not change, then $p^{*j} = \tilde{p}^w/\tau^{*j}$ is also unaltered, and so $W^{*j}(p^{*j}(\tau^{*j}, \tilde{p}^w), \tilde{p}^w)$ is unaltered as well.

When would a negotiated change in the tariffs of the home country and foreign country $*i$ leave the world price unaltered? Referring to (4), we see that a negotiated change in the tariffs of the home country and foreign country $*i$ can preserve the world price \tilde{p}^w only if the tariffs move in the same direction. Using foreign country $*i$'s trade-balance condition (2), Bagwell and Staiger (2005) go further and argue that the world price remains unaltered in this two-good setting if the home country and foreign country $*i$ negotiate tariff changes that satisfy the principle of bilateral reciprocity whereby any change in the volume of foreign country $*i$'s imports is of equal value to the change in the volume of its exports.²⁶ Given our maintained assumption that the home country's trade policy is nondiscriminatory, the principle of bilateral reciprocity thus completely insulates third parties from spillover effects associated with bilateral negotiations.

For our purposes here, it is convenient to formally define the principle of bilateral reciprocity directly in terms of its implication for the world price. To this end, we specify an initial tariff

²⁶ Formally, consider a negotiation between the home country and, say, foreign country $*1$ concerning changes in their respective tariffs. Let $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$ and $(\tau_1, \tau_1^{*1}, \tau_0^{*2}) \in \Upsilon$ represent the initial and negotiated tariff vectors, with associated world prices $p_0^w \equiv \tilde{p}^w(\tau_0, \tau_0^{*1}, \tau_0^{*2})$ and $p_1^w \equiv \tilde{p}^w(\tau_1, \tau_1^{*1}, \tau_0^{*2})$. Let the corresponding local prices in foreign country $*1$ be denoted as $p_0^{*1} = p_0^w/\tau_0^{*1}$ and $p_1^{*1} = p_1^w/\tau_1^{*1}$. Bagwell and Staiger (2005) define the principle of bilateral reciprocity as holding if and only if

$$M^{*1}(p_1^{*1}, p_1^w) - M^{*1}(p_0^{*1}, p_0^w) = p_0^w [E^{*1}(p_1^{*1}, p_1^w) - E^{*1}(p_0^{*1}, p_0^w)].$$

As Bagwell and Staiger (2005) argue, after applying (2) for $i = 1$ at both the initial and negotiated tariff vectors, it follows easily that the principle of bilateral reciprocity holds if and only if $p_0^w = p_1^w$.

vector, $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$, and associated world price, $p_0^w = \tilde{p}^w(\tau_0, \tau_0^{*1}, \tau_0^{*2})$, and consider a negotiation between home and foreign country *1, say. With τ^{*2} held fixed at the initial level τ_0^{*2} , the negotiated tariff vector may be represented as $(\tau_1, \tau_1^{*1}, \tau_0^{*2}) \in \Upsilon$, with an associated world price that is given as $p_1^w = \tilde{p}^w(\tau_1, \tau_1^{*1}, \tau_0^{*2})$. We now say that a negotiated tariff change between home and foreign country *1 satisfies the principle of *bilateral reciprocity* if the world price is unchanged: $p_0^w = p_1^w$. The definition extends in the obvious way if the bilateral negotiation is between the home country and foreign country *2.

Even though two countries are unable to alter the terms of trade when undertaking a negotiation that satisfies the principle of bilateral reciprocity, they are able to change their local prices and thereby experience welfare effects as a consequence of such a negotiation. Indeed, as noted above, starting at the Nash tariffs, both negotiating countries gain from exchanging small tariff cuts that satisfy the principle of bilateral reciprocity. Furthermore, as the welfare-preservation property indicates, the third country experiences no welfare effect as a consequence of a negotiated tariff change that is undertaken by the other countries and that satisfies bilateral reciprocity.

We turn now to multilateral negotiations in which the tariffs of *all* countries may change. One implication is that all three countries may experience welfare effects even when tariffs satisfy a notion of multilateral reciprocity. To make this argument precise, we first must define multilateral reciprocity.

To this end, we now consider a multilateral negotiation among all three countries. Let us specify an initial tariff vector, $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$, and associated world price, $p_0^w = \tilde{p}^w(\tau_0, \tau_0^{*1}, \tau_0^{*2})$, and also a negotiated tariff vector, $(\tau_1, \tau_1^{*1}, \tau_1^{*2}) \in \Upsilon$, and associated world price, $p_1^w = \tilde{p}^w(\tau_1, \tau_1^{*1}, \tau_1^{*2})$. We say that a negotiated tariff change satisfies the principle of *multilateral reciprocity* if the world price is unchanged: $p_0^w = p_1^w$. Of course, multilateral reciprocity is satisfied if the home country negotiates with only one foreign country and the negotiation conforms with bilateral reciprocity. Multilateral reciprocity is likewise satisfied if the home country negotiates separately with both foreign countries, with each individual negotiation satisfying bilateral reciprocity. Multilateral reciprocity is also consistent, however, with a scenario where individual negotiations, if viewed in isolation, would violate bilateral reciprocity. Finally, multilateral reciprocity is consistent as well with a scenario in which only the foreign tariffs are changed, where, under (4), the world price is maintained only if the foreign tariffs are changed in opposite directions.

Building directly on arguments made above, we can now report as well a simple welfare result for negotiated tariff cuts that satisfy the principle of multilateral reciprocity. Specifically, starting at the Nash equilibrium, if all three countries offer slight tariff cuts and the three tariff cuts in combination satisfy the principle of multilateral reciprocity, then all three countries experience a welfare gain.

3. THE CONSTRUCTED MECHANISM

With the basic model and some key properties now defined, we are prepared to analyze the manner in which multilateral tariffs may be determined through negotiation. Motivated by GATT tariff bargaining behavior, our approach is to require that any proposed change in tariffs satisfy the principles of nondiscrimination and multilateral reciprocity. In line with renegotiation provisions found in GATT Article XXVIII, we require further that no country ever be forced to accept a trade volume in excess of that implied by its proposal.²⁷ The remaining challenge is then to introduce any priority or rationing rules that may be needed for situations in which one side

²⁷ As Bagwell and Staiger (1999, 2002) discuss in detail, GATT Article XXVIII contains renegotiation provisions under which a country can initiate reciprocal tariff increases that preserve the terms of trade. If a negotiated tariff agreement were such that a country received greater trade volume than it preferred at the given terms of trade, then the country could subsequently utilize GATT Article XXVIII and renegotiate to obtain its preferred trade volume at the given terms of trade. Following Bagwell and Staiger (1999, 2002), we capture the implications of this constraint here by assuming that no country can be forced to accept a trade volume in excess of that implied by its proposal.

of the market is short relative to the other. We address these issues in this section and construct a mechanism that translates simultaneous tariff proposals from each country into a vector of assigned tariffs. In the following section, we then argue that the constructed mechanism results in dominant strategies for each country.

3.1. *Setup.* We assume that the three countries begin their negotiation with an exogenous initial tariff vector, $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$. The associated initial world price is represented as $p_0^w \equiv \tilde{p}^w(\tau_0, \tau_0^{*1}, \tau_0^{*2}) > 0$.

3.2. *Strategies.* The game form involves simultaneous proposals by all three countries. We impose two restrictions on the strategy set. First, we assume that each country can only make proposals concerning its own tariff and that of its trading partner(s). Since the two foreign countries do not trade with one another, this means that no foreign country proposes a change in the tariff of the other foreign country.²⁸ Second, we assume that each proposal if accepted must maintain the world price. The latter restriction means that we are restricting proposals to satisfy multilateral reciprocity, as that term is defined above.²⁹

Formally, the respective strategy spaces are:

Home country's strategy: A proposal $T_h \equiv (T_h^h, T_h^{*1}, T_h^{*2}) \in \Upsilon$ such that $p_0^w \equiv \tilde{p}^w(T_h^h, T_h^{*1}, T_h^{*2})$.

*Foreign country *1's strategy:* A proposal $T_{*1} \equiv (T_{*1}^h, T_{*1}^{*1}, T_{*1}^{*2}) \in \Upsilon$ such that $p_0^w \equiv \tilde{p}^w(T_{*1}^h, T_{*1}^{*1}, T_{*1}^{*2})$ and $T_{*1}^{*2} = \tau_0^{*2}$.

*Foreign country *2's strategy:* A proposal $T_{*2} \equiv (T_{*2}^h, T_{*2}^{*1}, T_{*2}^{*2}) \in \Upsilon$ such that $p_0^w \equiv \tilde{p}^w(T_{*2}^h, T_{*2}^{*1}, T_{*2}^{*2})$ and $T_{*2}^{*1} = \tau_0^{*1}$.

Using h to denote the home country and $*i$ to denote foreign country $*i$, we interpret the notation for proposals as follows: T_h is the tariff vector proposed by the home country, T_h^h is the tariff that the home country proposes for itself, T_h^{*i} is the tariff that the home country proposes for foreign country $*i$, T_{*i} is the tariff vector proposed by foreign country $*i$, T_{*i}^{*i} is the tariff proposed by foreign country $*i$ for itself, T_{*i}^h is the tariff that foreign country $*i$ proposes for the home country, and $T_{*i}^{*j} = \tau_0^{*j}$ is the (status quo) tariff that foreign country $*i$ proposes for foreign country $*j$, where $i \neq j$. Let S_h, S_{*1} , and S_{*2} denote the respective strategy spaces so defined, and let $S \equiv S_h \times S_{*1} \times S_{*2}$.

The proposal of the home country may entail tariff changes by some or all countries, and the requirement of multilateral reciprocity ensures that any such changes would preserve the world price at its initial level. By contrast, the proposal of foreign country $*i$ can only entail changes in the tariffs of the home country and foreign country $*i$, and so the requirement of multilateral reciprocity for the proposal of foreign country $*i$ amounts to a requirement of bilateral reciprocity.

Summarizing, once the simultaneous proposals are made, we have the following objects:

$$T_h \equiv (T_h^h, T_h^{*1}, T_h^{*2}), T_{*1} \equiv (T_{*1}^h, T_{*1}^{*1}, T_{*1}^{*2}), \text{ and } T_{*2} \equiv (T_{*2}^h, T_{*2}^{*1}, T_{*2}^{*2}), \text{ where}$$

$$p_0^w \equiv \tilde{p}^w(T_h^h, T_h^{*1}, T_h^{*2}) = \tilde{p}^w(T_{*1}^h, T_{*1}^{*1}, T_{*1}^{*2}) = \tilde{p}^w(T_{*2}^h, T_{*2}^{*1}, T_{*2}^{*2}),$$

$$T_{*1}^{*2} = \tau_0^{*2}, \text{ and } T_{*2}^{*1} = \tau_0^{*1}.$$

3.3. *Equivalence Class.* From the point of view of any single country, there exists a continuum of tariff adjustments by the other two countries that leave unaltered the world price. For any foreign country, tariff changes by the home country and the other foreign country that

²⁸ This assumption is in line with GATT bargaining practice, wherein tariff negotiations are generally aligned with trade patterns. The assumption is not essential for our analysis, however.

²⁹ Recall that the MFN requirement is built into the model itself, with the assumption that the home country has only one import tariff, τ .

preserve the world price correspond to changes that satisfy bilateral reciprocity. By the welfare-preservation property, such adjustments leave the former foreign country indifferent. We recall also that, by (4), a world-price-preserving tariff adjustment between the home country and any one foreign country requires that the tariffs be changed in the same direction. Foreign tariffs also may be adjusted in a way that maintains the world price and thus leaves the home country indifferent. In this case, as previously noted, we may use (4) to conclude that the adjusted foreign tariffs move in opposite directions.

In recognition of such policies, we are led to define for each country a class of tariffs for the other countries that are equivalent from the former country’s perspective. We begin by defining an equivalence class of tariffs for each foreign country.

DEFINITION 1. Given $T_{*1} \equiv (T_{*1}^h, T_{*1}^{*1}, T_{*1}^{*2})$ and $p_0^w \equiv \tilde{p}^w(T_{*1})$, we may define an equivalence class for foreign country *1 as a tariff set

$$EC_{*1}(T_{*1}) \equiv \{\hat{T}_{*1} \equiv (\hat{T}_{*1}^h, \hat{T}_{*1}^{*1}, \hat{T}_{*1}^{*2})\},$$

which satisfies the requirements that (i) $\hat{T}_{*1}^{*1} = T_{*1}^{*1}$ and (ii) $\tilde{p}^w(\hat{T}_{*1}) = p_0^w$. Likewise, given $T_{*2} \equiv (T_{*2}^h, T_{*2}^{*1}, T_{*2}^{*2})$ and $p_0^w \equiv \tilde{p}^w(T_{*2})$, we may define an equivalence class for foreign country *2 as a tariff set

$$EC_{*2}(T_{*2}) \equiv \{\hat{T}_{*2} \equiv (\hat{T}_{*2}^h, \hat{T}_{*2}^{*1}, \hat{T}_{*2}^{*2})\},$$

which satisfies the requirements that (i) $\hat{T}_{*2}^{*2} = T_{*2}^{*2}$ and (ii) $\tilde{p}^w(\hat{T}_{*2}) = p_0^w$.

Thus, an equivalence class for foreign country *i maintains *i’s proposed tariff for itself and allows for alternative tariffs for home and foreign country *j, where $j \neq i$, such that these alternative tariffs when joined with *i’s proposed tariff for itself serve to maintain the initial world price. Notice that an equivalence class is not empty, since $T_{*i} \in EC_{*i}(T_{*i})$. Next, we observe that we can reach any member of $EC_{*i}(T_{*i})$ by fixing $\hat{T}_{*i}^{*i} = T_{*i}^{*i}$ and then allowing changes from (T_{*i}^h, T_{*i}^{*j}) to $(\hat{T}_{*i}^h, \hat{T}_{*i}^{*j})$ that satisfy bilateral reciprocity between home and foreign country *j and that thus maintain the initial world price. Finally, let $e_{*i}(T_{*i})$ denote a representative member of $EC_{*i}(T_{*i})$. Since T_{*i} and any $e_{*i}(T_{*i}) \in EC_{*i}(T_{*i})$ generate the same world price p_0^w and local price in foreign country *i (as *i’s own tariff is unaltered by requirement (i)), it follows that T_{*i} and any $e_{*i}(T_{*i}) \in EC_{*i}(T_{*i})$ generate the same economic magnitudes (e.g., trade volumes) and government welfare for foreign country *i. The latter implication is a restatement of the welfare-preservation property.

We now define an equivalence class of tariffs for the home country.

DEFINITION 2. Given $T_h \equiv (T_h^h, T_h^{*1}, T_h^{*2})$ and $p_0^w \equiv \tilde{p}^w(T_h)$, we may define an equivalence class for the home country as a tariff set

$$EC_h(T_h) \equiv \{\hat{T}_h \equiv (\hat{T}_h^h, \hat{T}_h^{*1}, \hat{T}_h^{*2})\},$$

which satisfies the requirements that (i) $\hat{T}_h^h = T_h^h$ and (ii) $\tilde{p}^w(\hat{T}_h) = p_0^w$.

Thus, an equivalence class for the home country maintains home’s proposed tariff for itself and allows for alternative tariffs for the two foreign countries such that these alternative tariffs when joined with home’s proposed tariff for itself serve to maintain the initial world price. Given that $T_h \in EC_h(T_h)$, we know that $EC_h(T_h)$ is not empty. We say that we can reach any member of $EC_h(T_h)$ by fixing $\hat{T}_h^h = T_h^h$ and then allowing changes from (T_h^{*1}, T_h^{*2}) to $(\hat{T}_h^{*1}, \hat{T}_h^{*2})$ that maintain the initial world price. Finally, since T_h and any $e_h(T_h) \in EC_h(T_h)$ generate the same world price p_0^w and local price in the home country (as home’s own tariff is unaltered

by requirement (i)), it follows that T_h and any $e_h(T_h) \in EC_h(T_h)$ generate the same economic magnitudes (e.g., trade volumes) and government welfare for the home country.

3.4. *Implied Import Volumes.* Each country’s proposal can be associated with an implied import volume for itself. We now introduce some notation with which to represent for each country the import volume that is implied by its proposal.

DEFINITION 3. The home country’s proposal T_h is associated with an implied import volume for home defined as

$$M_h \equiv M(p(T_h^h, p_0^w), p_0^w).$$

Similarly, foreign country $*i$ ’s proposal T_{*i} is associated with an implied import volume for foreign country $*i$ defined as

$$M_{*i} \equiv M^{*i}(p^{*i}(T_{*i}^{*i}, p_0^w), p_0^w).$$

Notice that, given the initial world price, each country’s implied import volume depends only on that country’s proposed tariff for its own imports. Notice also that, for any given country and proposal by that country, any member of the resulting equivalence class entails the same tariff for that country and the same world price, and so generates as well the same implied import volume for that country. Finally, since proposals are members of Υ , we also observe that $M_h > 0$ and $M_{*i} > 0$ for $i = 1, 2$.

3.5. *Agreement.* We now define agreement between the three tariff proposals.

DEFINITION 4. The proposals $\{T_h, T_{*1}, T_{*2}\}$ agree if and only if there exist $e_h(T_h) \in EC_h(T_h)$, $e_{*1}(T_{*1}) \in EC_{*1}(T_{*1})$, and $e_{*2}(T_{*2}) \in EC_{*2}(T_{*2})$ such that $e_h(T_h) = e_{*1}(T_{*1}) = e_{*2}(T_{*2})$.

Thus, proposals agree when a common tariff vector is in the equivalence class for each country.

We observe that the common tariff vector under agreement must use the proposal that each country makes for its own tariff. We record this observation as follows.

LEMMA 1. *The proposals $\{T_h, T_{*1}, T_{*2}\}$ agree if and only if*

$$(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) \in EC_h(T_h) \cap EC_{*1}(T_{*1}) \cap EC_{*2}(T_{*2}).$$

We now illustrate the notion of agreement with two examples. For simplicity, we assume in each example that the initial tariffs are $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) = (15, 15, 15)$, the initial world price is $p_0^w = 1$, and the world price is preserved when one unit of home liberalization is balanced against a total of one unit of liberalization from the foreign countries. In the first example, the common tariff vector that lies in all three equivalence classes is simply the tariff proposal of the home country. The second example illustrates, however, that the common tariff vector that falls in all three equivalence classes may differ in part from all three proposals.

EXAMPLE 1. Suppose that home’s proposal is $T_h = (5, 10, 10)$, which means that home proposes to cut its tariff by 10 in exchange for cuts of 5 by both of its trading partners. Suppose that foreign country $*1$ ’s proposal is $T_{*1} = (10, 10, 15)$, so that foreign country $*1$ proposes to exchange tariff cuts of 5 with the home country while leaving foreign country $*2$ ’s tariff at its initial level. Assume that foreign country $*2$ makes a symmetric proposal, $T_{*2} = (10, 15, 10)$. The proposals of home and each foreign country $*i$ maintain the world price under our assumptions, since $p_0^w \equiv \tilde{p}^w(15, 15, 15) = \tilde{p}^w(5, 10, 10) = \tilde{p}^w(10, 10, 15) = \tilde{p}^w(10, 15, 10)$. We establish now

that the proposals agree since the home proposal $T_h = (5, 10, 10)$ is a member of all three equivalence classes. It is immediate that $T_h \in EC_h(T_h)$. Next, observe that given $T_{*1} = (10, 10, 15)$, we can reach $T_h = (5, 10, 10)$ by having home and foreign country $*2$ exchange tariff cuts of 5 units, which preserves the world price and leaves foreign country $*1$ indifferent. In other words, $T_h \in EC_{*1}(T_{*1})$. By a similar argument, $T_h \in EC_{*2}(T_{*2})$. We conclude that the proposals $\{T_h = (5, 10, 10), T_{*1} = (10, 10, 15), T_{*2} = (10, 15, 10)\}$ agree.

EXAMPLE 2. Suppose that home's proposal is $T_h = (5, 10, 10)$ but that $T_{*1} = (12, 12, 15)$ and $T_{*2} = (8, 15, 8)$. The proposals of home and each foreign country $*i$ maintain the world price under our assumptions, since $p_0^w \equiv \tilde{p}^w(15, 15, 15) = \tilde{p}^w(5, 10, 10) = \tilde{p}^w(12, 12, 15) = \tilde{p}^w(8, 15, 8)$. We establish now that the proposals agree, since the proposal $(5, 12, 8)$ is a member of all three equivalence classes. To show that $(5, 12, 8) \in EC_h(T_h)$, we note that given T_h , we can reach $(5, 12, 8)$ by having foreign country $*1$ raise its tariff by 2 units, while foreign country $*1$ cuts its tariff by 2 units, which preserves the world price and leaves the home country indifferent. Next, $(5, 12, 8) \in EC_{*1}(T_{*1})$ follows, since we can reach $(5, 12, 8)$ from T_{*1} by having home and foreign country $*2$ exchange tariff cuts of 7 units, which preserves the world price and leaves foreign country $*1$ indifferent. Likewise, starting at T_{*2} , we can reach $(5, 12, 8)$ by having home and foreign country $*1$ exchange tariff cuts of 3 units, which preserves the world price and leaves foreign country $*2$ indifferent, thus establishing that $(5, 12, 8) \in EC_{*2}(T_{*2})$. We conclude that the proposals $\{T_h = (5, 10, 10), T_{*1} = (12, 12, 15), T_{*2} = (8, 15, 8)\}$ agree.

We may now report the following implication of agreement.

LEMMA 2. *The proposals $\{T_h, T_{*1}, T_{*2}\}$ agree if and only if $p_0^w M_h = M_{*1} + M_{*2}$.*

PROOF. The home proposal T_h implies the market-clearing world price of $\tilde{p}^w(T_h) = p_0^w$ and an implied import volume for home of $M_h = M(p(T_h^h, p_0^w), p_0^w)$. For the home proposal T_h , the home trade-balance condition (1) may be stated as

$$(7) \quad p_0^w M_h = E(p(T_h^h, p_0^w), p_0^w).$$

Suppose first that the proposals $\{T_h, T_{*1}, T_{*2}\}$ agree. Then we know from Lemma 1 that $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) \in EC_h(T_h) \cap EC_{*1}(T_{*1}) \cap EC_{*2}(T_{*2})$. Since $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) \in EC_h(T_h)$, we have that $\tilde{p}^w(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w$. It follows that

$$\begin{aligned} E(p(T_h^h, p_0^w), p_0^w) &= M^{*1}(p^{*1}(T_{*1}^{*1}, p_0^w), p_0^w) + M^{*2}(p^{*2}(T_{*2}^{*2}, p_0^w), p_0^w) \\ &= M_{*1} + M_{*2}. \end{aligned}$$

We may now refer to (7) to conclude that $p_0^w M_h = M_{*1} + M_{*2}$.

Suppose next that the proposals $\{T_h, T_{*1}, T_{*2}\}$ are such that $p_0^w M_h = M_{*1} + M_{*2}$. Using (7), we may rewrite this equality as

$$E(p(T_h^h, p_0^w), p_0^w) = M^{*1}(p^{*1}(T_{*1}^{*1}, p_0^w), p_0^w) + M^{*2}(p^{*2}(T_{*2}^{*2}, p_0^w), p_0^w).$$

It follows from the market-clearing condition (3) that $\tilde{p}^w(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w$. Given that the proposal vector $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$ employs the proposal that each country makes for its own tariff and satisfies $\tilde{p}^w(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w$, we conclude that

$$(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) \in EC_h(T_h) \cap EC_{*1}(T_{*1}) \cap EC_{*2}(T_{*2}).$$

Thus, by Lemma 1, the proposals $\{T_h, T_{*1}, T_{*2}\}$ agree. ■

We include the proof of Lemma 2 in the main text, since the proof is direct and reveals the central role played by the home trade-balance condition (1) and the market-clearing condition (3). As the proof illustrates, these conditions together ensure that the equation $p_0^w M_h = M_{*1} + M_{*2}$ can be alternatively expressed as indicating that the common tariff vector $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$ composed of each country's proposed tariff for itself delivers p_0^w as the market-clearing world price.

We further note that, under agreement, the common tariff vector $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$ is a member of Υ . This follows from the strategy-space restriction that each country's own proposal vector belongs to Υ . Finally, since the common tariff vector under agreement delivers the world price $\tilde{p}^w(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w$, the common tariff vector also results in the same import volume for each country as does that country's own proposal vector.

3.6. Mechanism. We define a *mechanism* as a pair $(S, g(\cdot))$, where S is the strategy space as defined above and g is an outcome function that maps from a vector of tariff proposals, (T_h, T_{*1}, T_{*2}) , to a vector of tariffs, $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+$. Given a vector of tariff proposals, a mechanism thus *assigns* (or *selects*) a tariff vector for application.

We impose two *baseline rules or requirements* for the mechanism that we construct. First, if the tariff proposals agree, then we require that the mechanism assign the tariff vector constituted of each country's proposal for its own tariff: $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$. We know from Lemma 1 that this is the only tariff vector that is in the equivalence class for each country. We also know from Lemma 2 that the proposals agree if and only if the home country and foreign countries (in aggregate) propose the same value of implied import volumes ($p_0^w M_h = M_{*1} + M_{*2}$). Second, if the tariff proposals do not agree, then we require that the mechanism assign a tariff vector that maximizes trade volume valued at world prices while not forcing any country to import more than its implied import volume and while preserving the initial world price. As we will see, under disagreement, the baseline rules do not uniquely determine the outcome function, and so we will add further rules below to ensure a unique mapping for our constructed mechanism. For now, we note that there are two ways that disagreement may occur: The home country may be on the long side ($p_0^w M_h > M_{*1} + M_{*2}$), or the home country may be on the short side ($p_0^w M_h < M_{*1} + M_{*2}$). We address each of these two cases and define corresponding assignment rules. Then, in the next section, we characterize dominant strategies for countries when the resulting constructed mechanism is used.

3.7. Agreement. Our first requirement for the mechanism is associated with the case in which the tariff proposals agree. As just noted, in this case, we require that the mechanism assign the tariff vector constituted of each country's proposal for its own tariff: $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$. By Lemma 1, this is the unique tariff vector that is in the equivalence class for each country. We also know from Lemma 2 that the proposals agree if and only if $p_0^w M_h = M_{*1} + M_{*2}$.

3.8. Disagreement. Our second requirement for the mechanism is associated with the case in which the tariff proposals fail to agree. By Lemma 2, failure to agree, or disagreement, occurs if and only if $p_0^w M_h \neq M_{*1} + M_{*2}$. When disagreement occurs, we require that the mechanism assign a tariff vector that maximizes trade volume valued at world prices while not forcing any country to import more than its implied import volume and while preserving the initial world price. After formally stating the associated program, we consider its implications when home is on the long side ($p_0^w M_h > M_{*1} + M_{*2}$) and when home is on the short side ($p_0^w M_h < M_{*1} + M_{*2}$).

To state the program, we must define trade volume valued at world prices. Consider any vector of applied tariffs, $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+$, for which $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$. We define the associated value of trade volume as

$$(8) \quad TV(\tau, \tau^{*1}, \tau^{*2}) \equiv p_0^w M(p(\tau, p_0^w), p_0^w) + \sum_{i=1,2} M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w).$$

This definition is consistent with market clearing for tariffs such that $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$.

Consider now any vector of tariff proposals, (T_h, T_{*1}, T_{*2}) , for which $p_0^w M_h \neq M_{*1} + M_{*2}$. Given such a vector, our *disagreement program* is defined as follows:

$$\max_{(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+} TV(\tau, \tau^{*1}, \tau^{*2}) \equiv p_0^w M(p(\tau, p_0^w), p_0^w) + \sum_{i=1,2} M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w)$$

subject to

$$\begin{aligned} \tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) &= p_0^w, \\ M(p(\tau, p_0^w), p_0^w) &\leq M_h, \\ M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) &\leq M_{*i}, \text{ for } i = 1, 2. \end{aligned}$$

Our second requirement now may be succinctly stated: Whenever disagreement occurs, the mechanism must assign a tariff vector that is a solution to the disagreement program.

3.9. *Home Long.* We begin with the case in which the tariff proposals are such that the home country is on the long side:

$$(9) \quad p_0^w M_h > M_{*1} + M_{*2}.$$

For this case, an initial point is that, under our baseline rules, we cannot assign tariffs that deliver the implied import volume for the home country. To see why, suppose that we assign an applied tariff vector, $(\tau, \tau^{*1}, \tau^{*2})$, for which $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and such that $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \leq M_{*i}$, for $i = 1, 2$, as required by our baseline rules. We then obtain

$$\begin{aligned} M(p(\tau, p_0^w), p_0^w) &= [1/p_0^w] E(p(\tau, p_0^w), p_0^w) \\ &= [1/p_0^w] \sum_{i=1,2} M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \\ &\leq [1/p_0^w] [M_{*1} + M_{*2}] \\ &< M_h, \end{aligned}$$

where the first equality follows from the home trade-balance condition (1), the second equality uses $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and the market-clearing condition (3), the first inequality employs $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \leq M_{*i}$, for $i = 1, 2$, and the final inequality is a restatement of (9). The import volume for the home country under the assigned tariff vector thus must be below its implied import volume limit: $M(p(\tau, p_0^w), p_0^w) < M_h$. Put differently, the home country's assigned tariff must exceed the tariff that it proposed for itself.

Building from this point, we may anticipate that our assignment for this case must be such that each foreign country imports a volume that equals that implied by its proposal: $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = M_{*i}$, for $i = 1, 2$. Such an assignment, in turn, implies that each foreign country must be assigned the tariff that it proposes for itself. The home country's assigned tariff then must be raised above its proposed tariff for itself to restore balance and deliver the initial world price as the market-clearing world price.

Formally, for the case in which the home country is on the long side as captured by (9), we claim that a mechanism that satisfies our baseline rules must assign the tariff vector

$$(10) \quad (\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2}),$$

where $\tilde{\tau} = \tilde{\tau}(T_{*1}^{*1}, T_{*2}^{*2})$ is defined to satisfy

$$(11) \quad \tilde{p}^w(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w.$$

Notice that, conditional on (9) holding, the proposed assignment rule uses only the foreign proposals and assigns to each foreign country the tariff that it proposed for itself.

To establish this claim, we state and prove the following.

LEMMA 3. *For a vector of tariff proposals (T_h, T_{*1}, T_{*2}) such that (9) holds, the unique solution to the disagreement program is given by the tariff vector $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where $\tilde{\tau}$ is defined by (11).*

Lemma 3 is proved in the Appendix.

Thus, in the case of disagreement in which the home country is on the long side, our baseline rules or requirements deliver a unique assigned tariff vector $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where, conditional on (9), the assigned tariff vector uses only the foreign proposals and assigns to each foreign country the tariff that it proposed for itself. The home country’s assigned tariff is raised above its proposed tariff for itself.

3.10. *Home Short.* The other case of disagreement occurs when the tariff proposals are such that the home country is short:

$$(12) \quad p_0^w M_h < M_{*1} + M_{*2}.$$

We begin our analysis of this case with the following observation: Under our baseline rules, we can no longer assign tariffs that deliver the implied import volumes for the foreign countries, since to do so would violate the implied import volume limit for the home country. To confirm this observation, let us assume to the contrary that we assign an applied tariff vector, $(\tau, \tau^{*1}, \tau^{*2})$, for which $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and such that $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = M_{*i}$ for each i . Using the home trade-balance condition (1), $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and the market-clearing condition (3), $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = M_{*i}$ for each i , and (12), we then obtain

$$\begin{aligned} p_0^w M(p(\tau, p_0^w), p_0^w) &= E(p(\tau, p_0^w), p_0^w) \\ &= \sum_{i=1,2} M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \\ &= M_{*1} + M_{*2} \\ &> p_0^w M_h, \end{aligned}$$

which means that the import volume for home under this tariff vector must exceed the home country’s implied import volume limit: $M(p(\tau, p_0^w), p_0^w) > M_h$. Hence, the assigned tariff vector must be such that at least one foreign country imports a volume that is strictly lower than its implied import volume.

We may anticipate that our assignment for this case must be such that the home country imports a volume that is equal to that implied by its proposal: $M(p(\tau, p_0^w), p_0^w) = M_h$. The home country’s assigned tariff is then equal to the tariff that it proposed for itself. At least one foreign country, however, must import a volume that is strictly below its implied import volume. Thus, for at least one foreign country, the assigned tariff for that country must be raised above the tariff that it proposed for itself.

In fact, there is a continuum of possible ways to allocate a fixed value of trade volume, $p_0^w M_h$, across the two foreign countries, even while maintaining the initial world price and ensuring that no foreign country imports a volume that exceeds its implied import volume. We require additional rules therefore if we seek a basis for a unique assigned tariff vector when the tariff

proposals are such that the home country is short. The problem is essentially one of choosing a rationing rule for allocating the fixed value of trade volume, $p_0^w M_h$, across the two foreign countries.

One approach might be to construct a mechanism that assigns tariffs so that foreign countries split the difference, with both foreign countries importing less than the volumes implied by their respective proposals. Looking ahead toward our dominant strategy arguments, however, a potential danger with this approach is that a foreign country might strategically choose its implied import volume in order to diminish the extent to which its assigned import volume falls short of the import volume that it actually prefers.³⁰ We thus pursue a different approach here. We pick a foreign country at random and specify for that country a tariff that delivers its implied import volume, provided that the value of that volume is no greater than the value of the home country's implied import volume. Otherwise, the selected foreign country imports a volume equal in value to the home country's implied import volume. The tariff of the other foreign country is then set so as to import the remaining value, if any, of the home country's implied import volume.³¹ Finally, with these assignments in place, the home country's tariff is set so as to deliver the initial world price as the market-clearing world price. As we confirm below, the home country's proposed tariff for itself is then the home country tariff that delivers the initial world price.

To formalize this approach, suppose that the proposals are such that (12) holds and that foreign country $*i$ is randomly selected as the "first" country. In the assigned tariff vector, foreign country $*i$ then sets the tariff τ^{*i} such that

$$(13) \quad M^{*i} (p^{*i} (\tau^{*i}, p_0^w), p_0^w) = \min \{M_{*i}, p_0^w M_h\}.$$

There are thus two cases, which we consider in turn.

The first case arises if

$$(14) \quad M_{*i} \geq p_0^w M_h.$$

For this case, we claim that the mechanism satisfies our baseline rules by assigning the tariff vector $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2}) \in \Upsilon_+$, which is defined as follows. First, $\hat{\tau}^{*i}$ is set so as to satisfy

$$(15) \quad M^{*i} (p^{*i} (\tau^{*i}, p_0^w), p_0^w) = p_0^w M_h.$$

Next, $\hat{\tau}^{*j}$ is set at a prohibitive level, so that

$$(16) \quad M^{*j} (p^{*j} (\tau^{*j}, p_0^w), p_0^w) = 0.$$

Finally, we define $\hat{\tau}$ so that the world price is maintained, given these foreign tariffs:

$$(17) \quad \tilde{p}^w (\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2}) = p_0^w. \quad 32$$

³⁰ For example, if $p_0^w M_h < M_{*1} + M_{*2}$ with $M_{*1} = M_{*2}$, then one approach might assign tariffs such that foreign country $*i$ imports $p_0^w M_h/2 < M_{*i}$. If the mechanism further specifies that foreign country $*i$ obtain its implied import volume limit when it instead makes a proposal that implies the import volume limit of $M'_{*i} = M_{*i} + \varepsilon$, then foreign country $*i$ would have incentive to propose for itself a lower tariff that implies the import volume M'_{*i} , even if its preferred volume is M_{*i} , since $M'_{*i} = M_{*i} + \varepsilon$ is closer to M_{*i} than is $p_0^w M_h/2$. Hence, when the home country is short, dominant strategy arguments may fail to apply under some natural (Bertrand-like) assignment rules. Dominant strategy characterizations likewise fail to apply under proportional rationing schemes (see Benassy, 1982).

³¹ We could allow the "second" foreign country to choose the minimum of its implied import volume and any remaining value of the home country's implied import volume. The scenario we consider here, however, is one in which the home country is short, which is to say that the implied import volume of the second foreign country exceeds any remaining value of the country's implied import volume.

To establish this claim, we state and prove the following.

LEMMA 4. *For a vector of tariff proposals (T_h, T_{*1}, T_{*2}) such that (12) holds in the form of the first case (14), a solution to the disagreement program is given by the tariff vector $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2})$ where $\hat{\tau}^{*i}$ is defined by (15), $\hat{\tau}^{*j}$ is defined by (16), and $\hat{\tau} = T_h^h$ is defined by (17).*

We see from Lemma 4 that, as anticipated, the home country’s assigned tariff is the tariff that it proposed for itself, $\hat{\tau} = T_h^h$. Lemma 4 is proved in the Appendix.

We note in this first case that our baseline rules do not uniquely identify an assigned tariff vector; for example, we could deliver the same trade volume while satisfying the other constraints by slightly increasing (lowering) the assigned tariff for foreign country $*i$ (foreign country $*j$) in a fashion that maintains the aggregate implied import volume for the two foreign countries. Thus, when the proposals induce this first case, we impose some additional rules in constructing our mechanism so as to arrive at the assigned tariff vector, $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2})$.

The second case arises if

$$(18) \quad M_{*i} < p_0^w M_h.$$

For this case, we claim that the mechanism satisfies our baseline rules by assigning the tariff vector $(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2}) \in \Upsilon$, which is defined as follows. First, $\bar{\tau}^{*i}$ is set so as to satisfy

$$(19) \quad M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = M_{*i},$$

from which it follows that $\bar{\tau}^{*i} = T_{*i}^{*i}$. Next, $\bar{\tau}^{*j}$ is set so that

$$(20) \quad M^{*j}(p^{*j}(\tau^{*j}, p_0^w), p_0^w) = \min\{M_{*j}, p_0^w M_h - M_{*i}\} = p_0^w M_h - M_{*i},$$

where the second equality follows from (12). Finally, we define $\bar{\tau}$ so that the initial world price is maintained, given these foreign tariffs:

$$(21) \quad \tilde{p}^w(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2}) = p_0^w.$$

To establish this claim, we state and prove the following.

LEMMA 5. *For a vector of tariff proposals (T_h, T_{*1}, T_{*2}) such that (12) holds in the form of the second case (18), a solution to the disagreement program is given by the tariff vector $(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2})$, where $\bar{\tau}^{*i} = T_{*i}^{*i}$ is defined by (19), $\bar{\tau}^{*j}$ is defined by (20), and $\bar{\tau} = T_h^h$ is defined by (21).*

As noted above, in this second case foreign country $*i$, the randomly selected “first” country, is assigned the tariff that it proposed for itself, $\bar{\tau}^{*i} = T_{*i}^{*i}$. As anticipated, Lemma 5 also confirms that the home country’s assigned tariff is again the tariff that it proposed for itself, $\bar{\tau} = T_h^h$. Lemma 5 is proved in the Appendix.

Just as in the first case, we note in this second case that our baseline rules do not uniquely identify an assigned tariff vector. Hence, when the proposals induce this second case, we impose some additional rules in constructing our mechanism so as to arrive at the assigned tariff vector, $(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2})$.

³² Notice that we assume here that, given the initial world price, there exists a finite tariff for foreign country $*j$ at and above which import volume into foreign country $*j$ is 0. We show in the proof of Lemma 4 that $\hat{\tau} = T_h^h$, where $T_h \in \Upsilon$. Thus, our assumptions in Section 2 ensure the existence of $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2}) \in \Upsilon_+$ such that foreign country $*j$ receives no trade volume.

3.11. *The Constructed Mechanism.* Recall that a mechanism is defined by the strategy space S and an outcome function g that takes tariff proposals and assigns tariffs. We may now summarize the tariffs that our *constructed mechanism* assigns as a function of the tariff proposals:

- A. If the tariff proposals agree, then the assigned tariff vector is $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$. The tariff proposals can agree if and only if $p_0^w M_h = M_{*1} + M_{*2}$.
- B. If the tariff proposals do not agree and the home country is long, so that $p_0^w M_h > M_{*1} + M_{*2}$, then the assigned tariff vector is $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where $\tilde{\tau}$ satisfies $\tilde{p}^w(\tau, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w$.
- C. If the tariff proposals do not agree and the home country is short, so that $p_0^w M_h < M_{*1} + M_{*2}$, then there are two cases:
 - 1. In the first case, the randomly selected first country, foreign country $*i$, makes a proposal such that $M_{*i} \geq p_0^w M_h$. The assigned tariff vector is then $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2})$, where $\hat{\tau}^{*i}$ satisfies $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = p_0^w M_h$, $\hat{\tau}^{*j}$ satisfies $M^{*j}(p^{*j}(\tau^{*j}, p_0^w), p_0^w) = 0$, and $\hat{\tau} = T_h^h$ satisfies $\tilde{p}^w(\tau, \hat{\tau}^{*1}, \hat{\tau}^{*2}) = p_0^w$.
 - 2. In the second case, the randomly selected first country, foreign country $*i$, makes a proposal such that $M_{*i} < p_0^w M_h$. The assigned tariff vector is then $(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2})$, where $\bar{\tau}^{*i} = T_{*i}^{*i}$ satisfies $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = M_{*i}$, $\bar{\tau}^{*j}$ satisfies $M^{*j}(p^{*j}(\tau^{*j}, p_0^w), p_0^w) = p_0^w M_h - M_{*i}$, and $\bar{\tau} = T_h^h$ satisfies $\tilde{p}^w(\tau, \bar{\tau}^{*1}, \bar{\tau}^{*2}) = p_0^w$.

We may now summarize our findings for this section.

PROPOSITION 1. *The constructed mechanism satisfies the two baseline rules or requirements.*

PROOF. The constructed mechanism is defined above, and we may use Lemmas 1–5 to confirm that the mechanism so defined satisfies the two baseline rules or requirements. ■

To conclude this section, we relate our model and constructed mechanism to the “request-offer” language and practice of GATT tariff bargaining. In the context of this language, we may think of T_h^{*1} and T_h^{*2} as the home country’s *requests* of foreign countries $*1$ and $*2$, respectively, and T_h^h as the home country’s *offer*. In similar fashion, for $i = 1, 2$, we may think of T_{*i}^{*i} as foreign country $*i$ ’s *request* of the home country and T_{*i}^{*i} as foreign country $*i$ ’s *offer*. In the context of this language, therefore, we may think of the tariff that a country proposes for itself as that country’s offer. An interesting implication of our constructed mechanism is then that the offer in each proposal plays a central role, both when tariff proposals agree and when they do not. Indeed, when tariff proposals agree, the mechanism assigns the tariff vector that is constituted of each country’s offer. Under disagreement, the assigned tariff vector for foreign countries corresponds to their offers when the home country is long, and the assigned tariff vector for the home country corresponds to its offer when the home country is short.³³ When the proposals disagree, a country that is on the long side may need to reduce the depth of its offer in order to ensure that the assigned tariff vector satisfies multilateral reciprocity. The central role for offers resonates with patterns reported by Bagwell et al. (2017), wherein once initial proposals were on the table, tariff bargaining in the GATT Torquay Round focused on possible modifications to the offers in each proposal instead of requests, and such modifications typically entailed reductions in the depths of offers.

Finally, as we have noted, our mechanism requires that any proposed change in tariffs satisfy the principles of nondiscrimination and multilateral reciprocity, and we assume that proposals are made simultaneously. Here, we observe that simultaneous tariff bargaining subject to the constraint of multilateral reciprocity (and nondiscrimination) describes well what was viewed at the time as a hallmark of the GATT tariff bargaining forum. As one early GATT report put it (see also Curzon, 1966, pp. 75–77):

³³ In line with the discussion above relating to Lemma 5, if the home country is short and the second case prevails, then the constructed mechanism also assigns the tariff offer for one foreign country.

Multilateral tariff bargaining, as devised at the London Session of the Preparatory Committee in October 1946 and as worked out in practice at Geneva and Annecy, is one of the most remarkable developments in economic relations between nations that has occurred in our time. It has produced a technique whereby governments, in determining the concessions they are prepared to offer, are able to take into account the indirect benefits they may expect to gain as a result of simultaneous negotiations between other countries, and whereby world tariffs may be scaled down within a remarkably short time. . . . The multilateral character of the Agreement enabled the negotiators to offer more extensive concessions than they might have been prepared to grant if the concessions were to be incorporated in separate bilateral agreements. Before the Geneva negotiations a country would have aimed at striking a balance between the concessions granted to another country and the direct concessions obtained from it without taking into account indirect benefits which might accrue from other prospective trade agreements; it might even have been unwilling to grant an important concession if it had been obliged to extend that concession to third countries without compensation. (ICITO, 1949, p. 10)

In effect, the claim of the ICITO report is that governments could *exchange* externalities across bilaterals in a balanced way that allowed them to maintain multilateral reciprocity and that this exchange was made possible by the simultaneous negotiations facilitated in GATT rounds.

4. DOMINANT STRATEGIES

We now consider the endogenous determination of tariffs in the constructed mechanism when governments use only dominant strategies. In this section, we take the first step in this process and characterize the respective sets of dominant strategies for foreign country **i* and the home country when the constructed mechanism is used.

An initial observation is that, for any foreign country **i*, the proposal strategy is completely described by T_{*i}^{*i} . To see the point, consider foreign country **1*. Since foreign country **1* is restricted to set $T_{*1}^{*2} = \tau_0^{*2}$ and to set $T_{*1}^h = T_{*1}^h(T_{*1}^{*1})$ at the unique value given T_{*1}^{*1} that delivers $p_0^w \equiv \tilde{p}^w(T_{*1}^h, T_{*1}^{*1}, \tau_0^{*2})$, its proposal $(T_{*1}^h, T_{*1}^{*1}, T_{*1}^{*2}) = (T_{*1}^h(T_{*1}^{*1}), T_{*1}^{*1}, \tau_0^{*2})$ is completely determined by its selection of T_{*1}^{*1} . By contrast, as noted previously, the home country's proposal is not fully determined by its proposal for its own tariff, T_h^h , since the home country's proposal includes levels for both foreign tariffs and these tariffs can be combined in different ways to generate the initial world price.

To characterize dominant strategies, we must utilize the payoff functions. For a given initial world price p_0^w , the home-country payoff is defined by the welfare function $W(p(\tau, p_0^w), p_0^w)$, whereas the payoff for foreign country **i* is defined by the welfare function $W^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w)$. In (5), we describe how each welfare function varies with the world price, for a given local price. For our present purposes, however, the world price is fixed at its initial level, p_0^w , and the more relevant issue is how each welfare function then varies with the corresponding local price. We now impose some modest structure on the dependence of each welfare function on the corresponding local price.

To begin, we follow Bagwell and Staiger (1999, 2002) and define the *politically optimal reaction tariff for the home country* as the tariff that satisfies

$$(22) \quad W_p(p(\tau, p_0^w), p_0^w) = 0,$$

and the *politically optimal reaction tariff for foreign country *i* as the tariff that satisfies

$$(23) \quad W_{p^{*i}}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) = 0.$$

Let τ_{PO} and τ_{PO}^{*i} denote the respective politically optimal reaction tariffs for the home country and foreign country **i*.

Having defined politically optimal reaction tariffs for each country, we may now state three further assumptions. Our first assumption, which requires a couple of additional definitions, is

that the corresponding politically optimal proposals for each country generate positive trade volumes for all countries and thus reside in Υ . To state this assumption formally, let us define for foreign country $*i$ the proposal strategy $T_{*i,PO}$ where $T_{*i,PO}^{*i} = \tau_{PO}^{*i}$, $T_{*i,PO}^{*j} = \tau_0^{*j}$, and $T_{*i,PO}^h = T_{*i}^h(\tau_{PO}^{*i})$ is then uniquely specified to deliver p_0^w as the market-clearing world price. Similarly, for the home country, we define a set of home-country proposal strategies for which $T_h^h = \tau_{PO}$ with T_h^{*1} and T_h^{*2} then specified in any fashion so that $\tilde{p}^w(\tau_{PO}, T_h^{*1}, T_h^{*2}) = p_0^w$. With these definitions in place, our first assumption is that the initial tariff vector and associated world price are such that (i) $T_{*i,PO} \in \Upsilon$ and (ii) there exists (T_h^{*1}, T_h^{*2}) such that $\tilde{p}^w(\tau_{PO}, T_h^{*1}, T_h^{*2}) = p_0^w$ and $(\tau_{PO}, T_h^{*1}, T_h^{*2}) \in \Upsilon$.³⁴

Our second assumption is that, given the fixed initial world price p_0^w , each country has *single-peaked preferences* with respect to its own tariff (or equivalently, with respect to its local price).³⁵ Formally, for the given p_0^w , we assume that $W_p(p(\tau, p_0^w), p_0^w) > 0$ for $p(\tau, p_0^w) < p(\tau_{PO}, p_0^w)$ and $W_p(p(\tau, p_0^w), p_0^w) < 0$ for $p(\tau, p_0^w) > p(\tau_{PO}, p_0^w)$, where we recall that $p(\tau, p_0^w) = \tau p_0^w$. Similarly, for the given p_0^w , we assume that $W_{p^{*i}}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) > 0$ for $p^{*i}(\tau^{*i}, p_0^w) < p^{*i}(\tau_{PO}^{*i}, p_0^w)$ and $W_{p^{*i}}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) < 0$ for $p^{*i}(\tau^{*i}, p_0^w) > p(\tau_{PO}, p_0^w)$, where we recall that $p^{*i}(\tau^{*i}, p_0^w) = (1/\tau^{*i})p_0^w$. These assumptions are all understood to hold for tariffs such that the corresponding country has positive trade volume. Any scenario in which a foreign country has zero trade volume corresponds to a limiting case.

A final assumption that we add at this point serves to simplify the characterization of dominant strategies for the home country. Specifically, we assume that, for any foreign country $*i$ and tariff vector $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$ such that $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$, there exists an alternative tariff vector $(\tau', \tau'^{*1}, \tau'^{*2}) \in \Upsilon$ such that $\tilde{p}^w(\tau', \tau'^{*1}, \tau'^{*2}) = p_0^w$ and $\tau'^{*i} = \tau^{*i}$, $\tau'^{*j} = \tau_0^{*j}$. Thus, if it is feasible for foreign country $*i$ to obtain an import volume $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w)$ with some tariff vector $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$, then we assume that it is also feasible for foreign country $*i$ to obtain that same import volume with world-price-preserving modifications in τ and τ^{*j} that reposition the latter tariff to the initial level, τ_0^{*j} .³⁶

4.1. *Dominant Strategies for Foreign Country *i.* We characterize first the dominant strategies for foreign country $*i$. As the following proposition establishes, the characterization of foreign country $*i$'s dominant strategy set is quite simple. In the single member of this set, foreign country $*i$ proposes τ_{PO}^{*i} for itself and proposes a tariff for the home country that in combination with the initial tariff for foreign country $*j$ maintains the market-clearing world price at its initial level, p_0^w .

PROPOSITION 2. *Given the constructed mechanism, the set of dominant strategy proposals for foreign country *i is nonempty and, in fact, contains a singleton defined by $T_{*i,PO}$, where $T_{*i,PO}^{*i} = \tau_{PO}^{*i}$, $T_{*i,PO}^{*j} = \tau_0^{*j}$, and $T_{*i,PO}^h = T_{*i}^h(\tau_{PO}^{*i})$ is then uniquely set to deliver p_0^w as the market-clearing world price.*

The proof of Proposition 2 is in the Appendix.

³⁴ For the home country, our goal here is simply to ensure the existence of a tariff vector such that the home-country tariff is τ_{PO} , the market-clearing world price is p_0^w , and all countries receive positive import volumes. Tariff vectors for which the home-country tariff is τ_{PO} , the market-clearing world price is p_0^w , and some foreign country receives zero import volume may exist but are not in Υ and thus are not part of the set of dominant strategy proposals for the home country.

³⁵ Bagwell and Staiger (1999, 2002) capture this assumption with the assumption that global second-order conditions hold for all maximization problems. The global second-order condition associated with the maximization problem leading to home's politically optimal reaction tariff, for example, is that $W_{pp}(p, p_0^w) < 0$.

³⁶ Obtaining the import volume $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w)$ with τ^{*j} set equal to τ_0^{*j} would not be feasible if this could not be obtained with a home tariff level above -1 and below the level that would prohibit trade between the home country and foreign country $*j$. Hence, with this assumption, we rule out extreme size asymmetries between the home country and each of its foreign trading partners.

The structure of the proof takes the following form. To establish that $T_{*i,PO}$ is a dominant strategy, we must show that foreign country $*i$ cannot improve upon $T_{*i,PO}$ regardless of the proposals of the home country and foreign country $*j$. The argument is straightforward when $T_{*i,PO}$ and the proposals of the other countries result in agreement or in the home country being long, since under the constructed mechanism, foreign country $*i$ then enjoys its favorite local price for the given initial world price. The remaining possibility in which the home country is short requires consideration of more cases, as discussed in the Appendix. To establish that no other proposal strategy can be a dominant strategy for foreign country $*i$, we show that, for any alternative proposal strategy for foreign country $*i$, we can identify strategies for the home country and foreign country $*j$ such that the proposal strategy $T_{*i,PO}$ generates strictly higher welfare for foreign country $*i$ than does the given alternative. For a given alternative strategy, the key idea is to find proposal strategies for other countries such that, whether foreign country $*i$ selects the alternative strategy or $T_{*i,PO}$, the proposals result in agreement or in the home country being long, so that under the constructed mechanism, foreign country $*i$ is assigned the tariff that it proposed for itself.

4.2. *Dominant Strategies for the Home Country.* We characterize next the dominant strategies for the home country. As the following proposition establishes, the characterization of the home country’s dominant strategy set is quite simple, too. In any member of this set, the home country proposes the tariff τ_{PO} for itself and proposes a pair of tariffs for foreign countries such that the three tariffs together maintain the market-clearing world price at its initial level, p_0^w .

PROPOSITION 3. *Given the constructed mechanism, the set of dominant strategy proposals for the home country is a nonempty set whose members are the home-country proposal strategies for which $T_h^h = \tau_{PO}$, with T_h^{*1} and T_h^{*2} then specified in any fashion so that $\tilde{p}^w(\tau_{PO}, T_h^{*1}, T_h^{*2}) = p_0^w$.*

Proposition 3 is proved in the Appendix.

In the proof, we take an arbitrary proposal strategy, $T_{h,PO}$, from the set described in Proposition 3 and begin by showing that the home country cannot improve upon $T_{h,PO}$ regardless of the proposals of the foreign countries. Similar to the proof of Proposition 2, the argument is straightforward when $T_{h,PO}$ and the proposals of the foreign countries result in agreement or in the home country being short, since under the constructed mechanism, the home country then enjoys its favorite local price for the given initial world price. In the remaining possibility in which the home country is long, we show that the home country is also unable to improve upon $T_{h,PO}$, since an alternative strategy that alters its assigned tariff (by putting the home country on the short side) would only move its assigned tariff further above its politically optimal reaction tariff. Next, to establish that any alternative proposal strategy T_h with $T_h^h \neq \tau_{PO}$ is not a dominant strategy, we identify strategies for the foreign countries such that $T_{h,PO}$ generates strictly higher welfare for the home country than does T_h . The key idea is to utilize the final additional assumption listed above and find proposal strategies for the foreign countries such that, whether the home country selects T_h or $T_{h,PO}$, the proposals result in agreement or in the home country being short, so that under the constructed mechanism, the home country is assigned the tariff that it proposed for itself.

5. DOMINANT STRATEGY OUTCOMES

Having characterized the sets of dominant strategies for foreign country $*i$ and the home country, we now characterize the tariff vectors that can be *achieved* as assignments in the constructed mechanism when governments use only dominant strategies.³⁷ We refer to any

³⁷ A possible alternative terminology would be to refer here to the tariff vectors that can be “implemented” using dominant strategies and the constructed mechanism. We postpone use of this term, however, until Section 7, when we extend our framework and explicitly include a space of possible government preferences from which the constructed mechanism delivers a mapping to negotiated tariff outcomes. See also footnote 40.

negotiated tariff outcome that can be achieved in this sense as a *dominant strategy outcome*. Our discussion shows how the negotiated outcome actually emerges from dominant strategy proposals, and we also assess the efficiency of the negotiated outcome.

To begin, we require some further definitions. Recall from (22) and (23) that the politically optimal reaction tariffs, τ_{PO} , τ_{PO}^{*1} , and τ_{PO}^{*2} , are defined relative to an initial world price, p_0^w , which we have taken to be exogenous. Let us now define p_{PO}^w as the particular initial world price p_0^w such that if each country were to apply its politically optimal reaction tariff relative to $p_0^w = p_{PO}^w$, then the resulting market-clearing world price would preserve the initial world price: $\tilde{p}^w(\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2}) = p_{PO}^w$ when $p_0^w = p_{PO}^w$.

Next, following Bagwell and Staiger (1999, 2002), we define the *politically optimal tariffs* as the three tariffs that simultaneously solve the following three equations:

$$(24) \quad W_p(p(\tau, \tilde{p}^w), \tilde{p}^w) = 0 = W_{p^{*i}}(p^{*i}(\tau^{*i}, \tilde{p}^w), \tilde{p}^w), \quad i = 1, 2.$$

At the politically optimal tariffs, the associated market-clearing world price \tilde{p}^w takes the value p_{PO}^w . We may thus understand p_{PO}^w to be the politically optimal world price. For $p_0^w = p_{PO}^w$, we assume that the politically optimal tariffs reside in Υ .³⁸

Bagwell and Staiger (1999, 2002) show that politically optimal tariffs are efficient, where efficiency is measured relative to the preferences of governments and is assessed relative to the full set of tariff vectors $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon$.³⁹ Efficiency fails, however, at tariffs for which some but not all of the equations in (24) are satisfied. In terms of the analysis conducted here, the politically optimal tariffs that Bagwell and Staiger (1999, 2002) define are the politically optimal reaction tariffs when those tariffs are defined relative to the initial world price p_{PO}^w . An immediate implication is that the politically optimal reaction tariffs are efficient when they are defined relative to p_{PO}^w .

We seek now to characterize dominant strategy outcomes. For a given vector of government preference functions, our approach is to characterize how the dominant strategy outcome varies as different values for the exogenous world price p_0^w are considered. In particular, and motivated by the discussion above, we distinguish between two cases: $p_0^w \neq p_{PO}^w$ and $p_0^w = p_{PO}^w$.⁴⁰

Case 1. $p_0^w \neq p_{PO}^w$. In this case, a first point is that, when dominant strategy proposals are used, agreement does not occur. To see this, assume to the contrary that agreement occurs. Our constructed mechanism then requires that the assigned tariff vector be $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2})$. For dominant strategy proposals, we further have from Propositions 2 and 3 that $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) =$

³⁸ In terms of our assumptions above, for $p_0^w = p_{PO}^w$, we thus assume that the pair $(\tau_{PO}^{*1}, \tau_{PO}^{*2})$ is one of the pairs (T_h^{*1}, T_h^{*2}) such that $\tilde{p}^w(\tau_{PO}, T_h^{*1}, T_h^{*2}) = p_0^w$ and $(\tau_{PO}, T_h^{*1}, T_h^{*2}) \in \Upsilon$.

³⁹ We distinguish this notion of efficiency from that used in the social choice literature on strategy-proof rationing rules. In the social choice literature, the volume of supply and price are given exogenously, and the notion of efficiency refers to whether the rationing rule leads to an efficient allocation of this fixed volume across agents. (See Barbera, 2011, section 9.1, for additional discussion of that literature.) The counterpart of this situation in our analysis arises when proposals are such that the home country is short, and the rationing rule then concerns how the home-country implied trade volume is allocated across foreign countries. In this context, the rationing rule that we employ allocates volume across foreign countries in an efficient manner. Our notion of efficiency in this article, however, is defined in relation to the full set of tariff vectors, Υ , and thus recognizes that different tariff vectors give rise to different world prices and implied home-country trade volumes.

⁴⁰ Notice that the given preference vector implies values for the politically optimal reaction tariffs, world price, and implied import volumes, where the latter values are defined below just after (25). In Section 7, we consider a second and closely related approach in which p_0^w is fixed but government preferences are drawn from a space of possible preferences. For the parameterization that we feature, this approach also leads to two cases, corresponding to whether, under the realized preference vector, the politically optimal world price equals or differs from p_0^w . As discussed in Section 7, for this second approach, we may use the findings from this section to characterize the social choice functions that can be implemented by the constructed mechanism using dominant strategy proposals.

$(\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$, where as always the politically optimal reaction tariffs are defined relative to the initial world price, p_0^w . By Lemma 2, agreement is possible if any only if

$$(25) \quad p_0^w M_{h,PO} = M_{*1,PO} + M_{*2,PO},$$

where $M_{h,PO} \equiv M(p(\tau_{PO}, p_0^w), p_0^w)$ and $M_{*i,PO} \equiv M^{*i}(p^{*i}(\tau_{PO}^{*i}, p_0^w), p_0^w)$. Using (25), the home trade-balance condition (1), and the market-clearing condition (3), we then see that $\tilde{p}^w(\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2}) = p_0^w$, which contradicts the assumption that $p_0^w \neq p_{PO}^w$.

Thus, when $p_0^w \neq p_{PO}^w$, there are two possibilities to consider for our characterization of dominant strategy outcomes. One possibility is that under dominant strategy proposals, the home country is long. Using Propositions 2 and 3, we know that under dominant proposals $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$, and so this possibility occurs if and only if

$$(26) \quad p_0^w M_{h,PO} > M_{*1,PO} + M_{*2,PO}.$$

When home is long so that (26) obtains, the tariff vector that is achieved is $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where $\tilde{p}^w(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2}) = p_0^w$ defines $\tilde{\tau}$. Thus, with $T_{*i}^{*i} = \tau_{PO}^{*i}$ under dominant proposals, $\tilde{\tau}$ is defined to satisfy $\tilde{p}^w(\tilde{\tau}, \tau_{PO}^{*1}, \tau_{PO}^{*2}) = p_0^w$. Given $p_{PO}^w \neq p_0^w$, it follows that $\tilde{\tau} \neq \tau_{PO}$. Since for the given initial world price, the foreign countries obtain their favorite local price while the home country does not, an implication from Bagwell and Staiger (1999, 2002) is that the dominant strategy outcome is inefficient when the home country is long and $p_0^w \neq p_{PO}^w$.⁴¹

When $p_0^w \neq p_{PO}^w$, it is interesting to consider how the dominant strategy outcome emerges from the proposed tariffs if the home country is long. The home-country proposal entails $T_h^h = \tau_{PO}$ along with proposals for foreign tariffs that preserve the initial world price, and this proposal is clearly not achieved since $\tilde{\tau} \neq \tau_{PO}$ when $p_{PO}^w \neq p_0^w$. Provided that $\tau_0^{*j} \neq \tau_{PO}^{*j}$, foreign country $*i$'s proposal is also not exactly achieved. But the achieved tariff vector is in foreign country $*i$'s equivalence class given its proposal. Thus, for each foreign country $*i$, we can think in this case of its proposal being accepted, once world-price-preserving modifications are made to its proposed tariffs for the home country and for foreign country $*j$.

The other possibility when $p_0^w \neq p_{PO}^w$ is that, under dominant strategies, the home country is short. Since Propositions 2 and 3 ensure that under dominant proposals $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$, this possibility occurs if and only if

$$(27) \quad p_0^w M_{h,PO} < M_{*1,PO} + M_{*2,PO}.$$

If home is short so that (27) obtains, the home tariff that is achieved is $T_h^h = \tau_{PO}$. The foreign tariffs that are achieved depend on whether, for the randomly selected foreign country $*i$, we have $M_{*i,PO} \geq p_0^w M_{h,PO}$ or $M_{*i,PO} < p_0^w M_{h,PO}$. Under the first inequality, the achieved tariff, $\hat{\tau}^{*i}$, for the randomly selected foreign country $*i$ satisfies $\hat{\tau}^{*i} \geq \tau_{PO}^{*i}$, whereas foreign country $*j$ sets a prohibitive tariff $\hat{\tau}^{*j}$ that thus satisfies $\hat{\tau}^{*j} > \tau_{PO}^{*j}$. Under the second inequality, the achieved tariff, $\bar{\tau}^{*i}$, for the randomly selected foreign country $*i$ satisfies $\bar{\tau}^{*i} = \tau_{PO}^{*i}$, whereas foreign country $*j$ sets a tariff $\bar{\tau}^{*j}$ under which by (20) we have $M^{*j}(p^{*j}(\bar{\tau}^{*j}, p_0^w), p_0^w) < M_{*j,PO}$ and so $\bar{\tau}^{*j} > \tau_{PO}^{*j}$. Given that the home country applies the tariff τ_{PO} while at least one foreign

⁴¹ To see why such a tariff vector is inefficient, suppose that $W_p < 0 = W_{p^{*i}}$ at the tariffs $(\tilde{\tau}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$. From here, suppose all three tariffs are slightly decreased in a way that preserves the market-clearing world price at its initial level, p_0^w . The foreign countries experience only a second-order loss, since they initially enjoyed their favorite local prices, while the home country enjoys a first-order gain from the induced lower value for p . We can now make a further small adjustment (e.g., a small cut in τ) that slightly raises the world price, so as to give the foreign countries a first-order gain. Provided that the second change is small relative to the first, the home country continues to enjoy a first-order gain, and in this way, a Pareto improvement has been engineered. A related argument would apply if $W_p > 0 = W_{p^{*i}}$ at the tariffs $(\tilde{\tau}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$.

country applies a tariff that differs from its politically optimal reaction tariff, an implication from Bagwell and Staiger (1999, 2002) is that the dominant strategy outcome is inefficient if the home country is short and $p_0^w \neq p_{PO}^w$.⁴²

When $p_0^w \neq p_{PO}^w$, it is also interesting to consider how the dominant strategy outcome emerges from the proposed tariffs if the home country is short. The home-country proposal entails $T_h^h = \tau_{PO}$ along with proposals for foreign tariffs that preserve the initial world price. Since the tariff vector that is achieved specifies that the home country applies the tariff τ_{PO} , any difference between the home-country proposal and the achieved tariff vector must involve the associated foreign-country tariffs. Provided that $\widehat{\tau}^{*j}$ and $\bar{\tau}^{*j}$ differ from τ_0^{*j} , the randomly selected foreign country $*i$'s proposal is also not exactly achieved, even when (as under the second inequality above) its assigned tariff is τ_{PO}^{*i} . The dominant strategy proposal for foreign country $*j$ (i.e., the foreign country that is not randomly selected) sets $T_{*j}^{*j} = \tau_{*j}^{*j}$; thus, its proposal is clearly not achieved, since the tariff vector that is achieved specifies either $\widehat{\tau}^{*j} > \tau_{PO}^{*j}$ or $\bar{\tau}^{*j} > \tau_{PO}^{*j}$ for foreign country $*j$. It thus follows as well that the achieved tariff vector is not in foreign country $*j$'s equivalence class given its proposal. We can think in this case of the home-country proposal being accepted as is (if foreign country tariffs are specified appropriately), or after world-price-preserving modifications are made to the home-country proposal for foreign-country tariffs.⁴³

Case 2. $p_0^w = p_{PO}^w$. Since $p_0^w = p_{PO}^w$ in this case, we have that $\tilde{p}^w(\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2}) = p_{PO}^w$. As above, for dominant strategy proposals, Propositions 2 and 3 yield $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$. For this case, under dominant strategy proposals, we thus have that

$$(28) \quad p_0^w M_{h,PO} = M_{*1,PO} + M_{*2,PO},$$

where again $M_{h,PO} \equiv M(p(\tau_{PO}, p_0^w), p_0^w)$ and $M_{*i,PO} \equiv M^{*i}(p^{*i}(\tau_{PO}, p_0^w), p_0^w)$. By Lemma 2, the proposals thus agree. The tariff vector that is achieved is thus $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$. Since for the given initial world price, all countries obtain their favorite local price, the implemented tariff vector corresponds to the vector of politically optimal tariffs. Hence, an implication from Bagwell and Staiger (1999, 2002) is that the dominant strategy outcome is efficient when $p_0^w = p_{PO}^w$.

When $p_0^w = p_{PO}^w$, it is also interesting to consider how the dominant strategy outcome $(T_h^h, T_{*1}^{*1}, T_{*2}^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$ emerges from the proposed tariffs. The set of dominant strategy proposals for the home country includes $(T_h^h, T_h^{*1}, T_h^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$. This proposal is now feasible, since $\tilde{p}^w(\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2}) = p_{PO}^w = p_0^w$. The home country may also propose $(T_h^h, T_h^{*1}, T_h^{*2})$, where $T_h^h = \tau_{PO}$, $(T_h^{*1}, T_h^{*2}) \neq (\tau_{PO}^{*1}, \tau_{PO}^{*2})$, and the market-clearing world price under the home country's proposed tariffs equals $p_0^w = p_{PO}^w$. Relative to home's proposal, the dominant strategy outcome then entails a world-price-preserving adjustment in foreign tariffs and is thus equivalent from home's perspective. Next, recall that the dominant strategy set for foreign country $*i$ is a singleton defined by $T_{*i,PO}$, where $T_{*i,PO}^{*i} = \tau_{PO}^{*i}$, $T_{*i,PO}^{*j} = \tau_0^{*j}$, and $T_{*i,PO}^h = T_h^h(\tau_{PO}^{*i})$ is then uniquely set to deliver p_0^w as the market-clearing world price. Since $p_{PO}^w = p_0^w$, the dominant strategy outcome is in the equivalence class for foreign country $*i$ that is defined by its dominant strategy proposal.

If the home country proposes $(T_h^h, T_h^{*1}, T_h^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$, then we may understand that the home-country proposal is accepted as is. If instead the home country proposes

⁴² At the achieved tariffs, we have $W_p = 0$ and $W_{p^{*i}} \neq 0$ for at least one i . We may thus engineer small local price changes that offer a first-order gain to at least one foreign country while imposing only a second-order cost on the home country. In line with the discussion in footnote 41, from here we can engineer a Pareto improvement with a small change that generates a terms-of-trade gain for home.

⁴³ When making its proposal, the home country does not know which foreign country will be randomly selected to go first; thus, the home-country proposal cannot be accepted as is with probability 1.

$(T_h^h, T_h^{*1}, T_h^{*2})$, where $T_h^h = \tau_{PO}$ and $(T_h^{*1}, T_h^{*2}) \neq (\tau_{PO}^{*1}, \tau_{PO}^{*2})$, then we can think of the home-country proposal being accepted, once world-price-preserving adjustments are made to its proposed tariffs for the foreign countries. Whether or not the home-country proposal is accepted as is, we can think of each foreign country $*i$'s proposal as being accepted, once world-price-preserving modifications are made to its proposed tariffs for the home country and for foreign country $*j$.

5.1. *Summary.* We have established the following proposition.

PROPOSITION 4. *Given the constructed mechanism, under dominant strategy proposals, (i) the achieved tariff vector is efficient if and only if $p_0^w = p_{PO}^w$, and (ii) when $p_0^w = p_{PO}^w$, the efficient tariff vector that is achieved is the politically optimal tariff vector.*

Put differently, Proposition 4 states that the dominant strategy outcome is efficient if and only if $p_0^w = p_{PO}^w$, and in this case, the efficient outcome that obtains is the politically optimal tariff vector.

As illustrated above, we can also interpret the process through which the dominant strategy outcome actually emerges from tariff proposals. In some cases, the outcome entails accepting the home-country proposal as is or after world-price-preserving modifications are made to the home-country proposal for foreign-country tariffs. In other cases, we can think of each foreign country $*i$'s proposal as being accepted once world-price-preserving modifications are made to its proposed tariffs for the home country and for foreign country $*j$. When $p_0^w = p_{PO}^w$, the proposals actually agree. Then, the home-country proposal may be accepted as is, and the proposal of each foreign country $*i$ is accepted once world-price-preserving modifications are made to its proposed tariffs for the home country and for foreign country $*j$.

With Proposition 4 in hand, we now return to the “request-offer” language and practice of GATT tariff bargaining and again highlight the central role played by offers (i.e., T_h^h , T_{*1}^{*1} , and T_{*2}^{*2}). In the case where $p_0^w = p_{PO}^w$, for example, the tariff vector that is achieved consists of each country's offer from its initial proposal. The world-price-preserving modification of foreign country $*i$'s proposal reaches this tariff vector by replacing foreign country $*i$'s requests of the home country and foreign country $*j$ with their offers contained in their proposals. From a practical perspective, the countries could obtain this outcome by simply agreeing to the offers on the table. Similarly, when $p_0^w \neq p_{PO}^w$, the offers of foreign countries are achieved when the home country is long, and the offer of the home country (and perhaps one foreign country) is achieved when the home country is short. In this situation, a country that is on the long side may need to modify its offer in order to ensure that the achieved tariff vector satisfies multilateral reciprocity.

We also note that the lack of strategic behavior implied by our dominant strategy characterizations is consistent with the observations of GATT commentators regarding the early GATT rounds (as well as with the findings of Bagwell et al., 2017, as we discuss in the introduction). For example, in his discussion of the bargaining techniques used by countries over the first five GATT rounds of request-offer tariff negotiations, Curzon (1966) observes:

Their requests cannot be higher than their offers and negotiations start from this maximum position: if all requests are granted all the offers will be fulfilled. Similarly all other contracting parties are likely to make offers which match the requests they have made. As some of the requests are rejected, some of the offers are withdrawn. This procedure has been raised to a Gatt principle and is not laid down by any rule. It is a convention but one which creates a much better negotiating climate than the opposite trend which was a feature of the classical bilateral negotiations. Then, everyone put forward very low offers with the intention of increasing gradually if the bargaining proved profitable. A country never knew, however, when it had reached the maximum its partner was willing to concede. (p. 74)

According to our findings here, the constraints of nondiscrimination and multilateral reciprocity can be interpreted as the key GATT institutional features that make the convention described by Curzon a dominant bargaining strategy in this setting.

6. NASH BEGINNINGS

Throughout, we have assumed a setup under which the initial tariff vector, $(\tau_0, \tau_0^{*1}, \tau_0^{*2})$, and associated world price, $p_0^w \equiv \tilde{p}^w(\tau_0, \tau_0^{*1}, \tau_0^{*2})$, are exogenous. An interesting possibility is that the initial tariff vector corresponds to Nash tariffs, so that $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) = (\tau_N, \tau_N^{*1}, \tau_N^{*2})$ with $p_0^w = \tilde{p}^w(\tau_N, \tau_N^{*1}, \tau_N^{*2}) \equiv p_N^w$.⁴⁴ We now briefly discuss some implications associated with this possibility.

As we note in Section 2, Nash tariffs are inefficient; furthermore, starting at the Nash equilibrium, if all three countries offer slight tariff cuts that in combination satisfy the principle of multilateral reciprocity, then all three countries experience a welfare gain. The key ideas behind the latter finding are that multilateral reciprocity fixes the world price and $W_p < 0 < W_{p^{*i}}$ at the Nash tariffs.

Let us now suppose that countries start at the Nash equilibrium, with $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) = (\tau_N, \tau_N^{*1}, \tau_N^{*2})$ and $p_0^w = \tilde{p}^w(\tau_N, \tau_N^{*1}, \tau_N^{*2}) \equiv p_N^w$. If countries use dominant strategy proposals, then $(T_h^i, T_{*1}^{*1}, T_{*2}^{*2}) = (\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$, where $(\tau_{PO}, \tau_{PO}^{*1}, \tau_{PO}^{*2})$ is defined relative to the initial world price $p_0^w = p_N^w$. Since no country can be forced to import a volume greater than its implied import volume, the achieved tariff vector $(\tau, \tau^{*1}, \tau^{*2})$ satisfies $M(p(\tau, p_0^w), p_0^w) \leq M_{h,PO} \equiv M(p(\tau_{PO}, p_0^w), p_0^w)$ and also satisfies $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \leq M_{*i,PO} \equiv M^{*i}(p^{*i}(\tau_{PO}^{*i}, p_0^w), p_0^w)$, $i = 1, 2$.

Suppose that the proposals are such that agreement occurs or the home country is long. We then may conclude that $M(p(\tau, p_0^w), p_0^w) > M_{h,N} \equiv M(p(\tau_N, p_0^w), p_0^w)$ and $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) > M_{*i,N} \equiv M^{*i}(p^{*i}(\tau_N^{*i}, p_0^w), p_0^w)$, $i = 1, 2$, so that the achieved tariff vector gives each country greater-than-Nash trade volumes and welfares. Given the initial world price $p_0^w = p_N^w$, this conclusion follows under (4) from the assumption of single-peaked preferences, $W_p < 0 < W_{p^{*i}}$ at the Nash tariffs, $W_p = 0$ at τ_{PO} , and $W_{p^{*i}} = 0$ at τ_{PO}^{*i} . The result holds as well when the home country is short, provided that the residual volume left for foreign country $*j$ (i.e., the “second” foreign country) is positive and exceeds its Nash volume, $M_{*j,N}$. This case is expected if the home country is large relative to the individual foreign countries but is not guaranteed by our assumptions.

One case of particular interest occurs when the home country is “symmetric” relative to the two foreign countries in the specific sense that $p_N^w = p_{PO}^w$.⁴⁵ Hence, under symmetry with $p_N^w = p_{PO}^w$, if countries use dominant strategies, then agreement would occur and the politically optimal tariffs would be achieved. The outcome is then efficient, with all three countries enjoying trade volume and welfare gains relative to the Nash starting point.

We summarize these findings as follows.

PROPOSITION 5. *Suppose the initial tariff vector is the Nash tariff vector. Given the constructed mechanism, under dominant strategy proposals, the following obtains: (i) If agreement occurs or*

⁴⁴ Given our maintained assumption that $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$, if $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) = (\tau_N, \tau_N^{*1}, \tau_N^{*2})$, then it follows that $(\tau_N, \tau_N^{*1}, \tau_N^{*2}) \in \Upsilon$.

⁴⁵ The equality of Nash and politically optimal world prices would obtain in a two-country model if the model were structured in a symmetric way so that government preferences were identical and the countries’ economies took mirror-image forms. Our three-country model, by contrast, has a built-in asymmetry in that the Nash tariff for good x is set by the home country alone, whereas the Nash tariffs for good y are independently set by the two foreign countries. All else equal, this asymmetry implies that the home country has greater market power to push world prices in its favor in the Nash equilibrium. Hence, the equality of Nash and politically optimal world prices would not be expected in our three-country model if preferences were identical and the home country and foreign countries in aggregate were to have mirror-image economies. The specific notion of symmetry that we define is thus distinct from a symmetric structure in this sense.

the home country is long, then the achieved tariff vector gives each country greater-than-Nash trade volumes and welfares; and (ii) if the setting is symmetric with $p_0^w = p_N^w$, then agreement occurs and the achieved tariff vector is the politically optimal tariff vector, which is efficient and gives each country greater-than-Nash trade volumes and welfares.

7. EXTENSIONS

The analysis can be extended in multiple ways. We discuss here four possible extensions.

7.1. Multiple Countries. As a first example, we may extend the analysis to consider multiple foreign countries. Suppose, for example, that there are three foreign countries. When the home country is short, one of the foreign countries, say foreign country $*i$, is selected as the “first” foreign country. As above, in the constructed mechanism, foreign country $*i$ imports a volume that is the minimum of M_{*i} and $p_0^w M_h$. If $p_0^w M_h > M_{*i}$, then the value of the home country’s implied import volume is not exhausted by foreign country $*i$ ’s implied import volume. We can then treat this residual value, $p_0^w M_h - M_{*i}$, as the value of home trade volume that is allocated across the remaining two foreign countries (i.e., $p_0^w M_h - M_{*i}$ in the model with three foreign countries then plays the role of M_h in the constructed mechanism defined above for the model with two foreign countries). In this general way, we can proceed inductively to define the constructed mechanism for any number of foreign countries.

7.2. Additional Constraints. We place little restriction above on the initial tariff vector, $(\tau_0, \tau_0^{*1}, \tau_0^{*2})$. The achieved tariff vector may involve tariff cuts if the initial tariffs are high, as Proposition 5 illustrates. At the same time, if the initial tariff vector entails tariffs that are low, then an efficient outcome such as described in Proposition 4 could entail tariff increases. When the initial tariff vector is such that the home country is short under dominant strategy proposals, it is also possible that one foreign country (the “second” foreign country) is assigned a prohibitive tariff and receives zero trade volume under the constructed mechanism. This assignment is not fundamental to our arguments, however. For example, with some additional notation, we could impose an exogenous upper bound on the tariff that each foreign country can apply, in which case the prohibitive tariff would be replaced with this maximum tariff. A natural candidate for the exogenous upper bound would be the existing tariff level that each foreign country sets going into the negotiations, where as one example the existing tariff level could be the Nash tariff.⁴⁶ Another possible extension would be to include a minimum-welfare, or ex post participation, constraint for each country, where the threshold welfare level is exogenous.

7.3. Alternative Rationing Rules. In the constructed mechanism, when the home country is short, we specify that a randomly selected foreign country be given first priority in the allocation of trade volume. As we noted, alternative rationing rules could be used. A key task is to ensure that the any new rule is also consistent with dominant strategy behavior. We mention here two kinds of alternative rationing rules.

A first alternative rule allows that different foreign countries may be treated differently, even in an ex ante sense. Suppose, for example, that foreign country $*i$ is, for exogenous reasons, a more significant trading partner for the home country. The home country might then regard foreign country $*i$ as a “principal supplier” in the GATT/WTO context. An alternative rationing rule could then be specified in which, when the home country is short, foreign country $*i$ always assumes the first-priority position (i.e., conditional on home being short, foreign country $*i$ is always treated as if it were randomly selected to go “first” in the scheme above). The results

⁴⁶ More generally, the existing tariff level for each country could correspond to any tariff that lies between its politically optimal reaction and Nash tariffs. For the extended model with maximum tariffs, when the home country is short under dominant strategy proposals, the “first” foreign country could be assigned a tariff that reciprocated all the increased trade volume implied by the home country’s proposal; but its assigned tariff could be no lower than this value, as the “second” foreign country could never be assigned a tariff that was higher than its existing tariff level.

above would continue to hold under this asymmetric rationing rule that utilizes a fixed priority scheme. Indeed, our assumption above that each foreign country is selected to go first with equal probability generates an attractive symmetry property but plays no formal role in our analysis, and so the fixed (i.e., deterministic) priority scheme can be understood as a limiting case of the analysis already conducted.

The second rule is based on the “uniform allocation rule.” Sprumont (1991) defines this rule for an allotment problem in which shares of a fixed volume of a divisible good are to be allocated at a fixed price among N agents, each of whom has a continuous, single-peaked preference as to its preferred share. The rule gives each agent his preferred share, provided that the shares do not fall outside of upper and lower bounds that are determined so that shares add up to 1.⁴⁷ Sprumont shows that an allotment rule is efficient, strategy-proof, and anonymous if and only if it is the uniform allocation rule.⁴⁸

We could impose a similar rule, although we would need to modify the rule in the case of excess supply (i.e., when the home country is long).⁴⁹ Similar results would then obtain. An appealing feature of the uniform allocation rule is that both foreign countries receive positive trade volumes when the home country is short, even if the import volume implied by one foreign country’s proposal exceeds that implied by the home country’s proposal. At the same time, in a decentralized bargaining setting, the uniform allocation rule may require significant coordination among parties to ensure that the bounds are properly determined.

7.4. Preference Spaces and Private Information. In our analysis above, we follow Bagwell and Staiger (1999, 2002) and assume that governments’ preferences are fixed. Using our constructed mechanism, we then detail in Section 5 the dominant strategy outcomes that emerge depending on the value that the initial world price takes relative to the politically optimal world price that is implied by the given government preferences. We consider now a second and closely related approach in which the value of p_0^w is fixed but government preferences are drawn from a space of possible preferences in a private values setting. For the parameterization developed here, this approach similarly leads to two cases, corresponding to whether under the realized preference vector the politically optimal world price equals or differs from p_0^w . As we argue below, we can allow as well that each government privately observes its realized preferences.⁵⁰ We can then use the findings from Section 5 to characterize for the private information setting the social choice functions that can be *implemented* by the constructed mechanism using dominant strategy proposals.⁵¹

To this end, we now suppose that government preferences are given as $W(p(\tau, p_0^w), p_0^w; \theta)$ and $W^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w; \theta^{*i})$, where $\theta \in \Theta$ and $\theta^{*i} \in \Theta^{*i}$ are preference parameters or “types.” We may think of Θ and Θ^{*i} as intervals on the real line, for example. Thus, in this extension, we explicitly allow for a space of possible government preferences. Importantly, we assume that θ and θ^{*i} correspond to private-value preference types and do not directly impact the

⁴⁷ Sonmez (1994) provides a helpful algorithm for reaching the outcome prescribed by the uniform rule. See also Barbera (2011).

⁴⁸ Efficiency is defined in this context relative to the problem of allocating the fixed volume among privately informed agents in an efficient way (see also footnote 39). An anonymous rule treats agents symmetrically.

⁴⁹ Given our requirement that a country cannot be forced to accept a trade volume in excess of that implied by its proposal, we are not able to require a foreign country to import a reference volume that exceeds its proposed (i.e., preferred) volume. We could thus handle the case in which the home country is long as we do above and use the uniform allocation rule when the home country is short.

⁵⁰ Given our constructed mechanism and dominant strategy solution concept, our analysis can be similarly extended to a complete information setting in which the government preference vector is drawn from a space of possible preferences.

⁵¹ With this extension, we can thus relate our results more directly to the class of problems considered in the large literature on implementation theory. For surveys, see Barbera (2001) and Jackson (2001). In this context, we note that utility is not (perfectly) transferable in our trade model. Governments may value world-price changes differently, and even world-price-preserving adjustments in import tariffs cause local prices to change and thus induce local-price distortions as a by-product.

determination of economic variables (i.e., $\tau_0, \tau_0^{*1}, \tau_0^{*2}, p_0^w$, and the function \tilde{p}^w are all independent of θ and θ^{*i}).

For this extended model, we may define strategies for each government just as before, except that now each government’s strategy is a function of its privately observed type. Thus, the proposal strategy of the home country can be represented as a function $T_h(\theta)$ where for all $\theta \in \Theta$, we require $T_h(\theta) \equiv (T_h^h(\theta), T_h^{*1}(\theta), T_h^{*2}(\theta)) \in \Upsilon$ and $p_0^w \equiv \tilde{p}^w(T_h^h(\theta), T_h^{*1}(\theta), T_h^{*2}(\theta))$. The proposal strategy for foreign country $*i, T_{*i}(\theta^{*i})$, is analogously defined. We may similarly extend the definition of the implied import volume associated with a country’s proposal to reflect that country’s type, so that $M_h(\theta) \equiv M(p(T_h^h(\theta), p_0^w), p_0^w)$ and $M_{*i}(\theta) \equiv M^{*i}(p^{*i}(T_{*i}^{*i}(\theta^{*i}), p_0^w), p_0^w)$. The constructed mechanism maps from actual tariff proposals to assigned tariffs and is defined just as before.

The definition of politically optimal tariff reactions may be similarly extended. Thus, $\tau_{PO}(\theta)$ satisfies $W_p(p(\tau, p_0^w), p_0^w; \theta) = 0$ and $\tau_{PO}^{*i}(\theta^{*i})$ satisfies $W_{p^{*i}}(p^{*i}(\tau^{*i}, p_0^w), p_0^w; \theta^{*i}) = 0$. The assumptions imposed previously on the government welfare functions are now assumed to apply for every ex post state $(\theta, \theta^{*1}, \theta^{*2})$; in particular, the assumption of single-peaked preferences continues to hold for each ex post state and thus is defined for $W(p(\tau, p_0^w), p_0^w; \theta)$ relative to $\tau_{PO}(\theta)$ and for $W^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w; \theta^{*i})$ relative to $\tau_{PO}^{*i}(\theta^{*i})$, respectively. The politically optimal world price can be represented in the extended setting as $p_{PO}^w(\theta, \theta^{*1}, \theta^{*2}) \equiv \tilde{p}^w(\tau_{PO}(\theta), \tau_{PO}^{*1}(\theta^{*1}), \tau_{PO}^{*2}(\theta^{*2}))$, and the associated politically optimal import volumes can be defined as $M_{h,PO}(\theta) \equiv M(p(\tau_{PO}(\theta), p_0^w), p_0^w)$ and $M_{*i,PO}(\theta^{*i}) \equiv M^{*i}(p^{*i}(\tau_{PO}^{*i}(\theta^{*i}), p_0^w), p_0^w)$, for $i = 1, 2$.

With these notational adjustments in place, our results may be directly translated to the private information setting. Our characterizations of dominant strategies in Propositions 2 and 3 can be restated as ex post conditions that must hold for every type that a country may have, where τ_{PO}^{*i} is replaced by $\tau_{PO}^{*i}(\theta^{*i})$ in Proposition 2 and τ_{PO} is replaced by $\tau_{PO}(\theta)$ in Proposition 3.⁵² Using the arguments in Section 5 leading up to Proposition 4, we can also characterize for the private information setting the social choice functions that can be implemented by the constructed mechanism using dominant strategy proposals. For any given preference vector realization $(\theta, \theta^{*1}, \theta^{*2})$, the resulting values for the politically optimal world price and implied import volumes indicate whether case 1 or case 2 applies, and within case 1 whether (26) or (27) obtains. The two-case taxonomy in Section 5 can thus be used to provide a (random) mapping from preference vector realizations to negotiated tariff outcomes.

Proposition 4 extends as well, once we fix a state $(\theta, \theta^{*1}, \theta^{*2})$, replace p_{PO}^w with $p_{PO}^w(\theta, \theta^{*1}, \theta^{*2})$, and interpret an “achieved tariff vector” as the implemented outcome that obtains at this state. For a given state $(\theta, \theta^{*1}, \theta^{*2})$, the key point is that the implemented outcome is efficient if and only if $p_{PO}^w(\theta, \theta^{*1}, \theta^{*2}) = p_0^w$. Thus, ex post efficiency obtains if and only if the preference domain is restricted so that only preference vectors for which $p_{PO}^w(\theta, \theta^{*1}, \theta^{*2}) = p_0^w$ are possible. For more general preference domains, the implemented social choice function is not ex post efficient.

8. CONCLUSION

Motivated by GATT bargaining behavior and renegotiation rules, we construct a three-country, two-good general-equilibrium model of trade and examine multilateral tariff bargaining under the constraints of nondiscrimination and multilateral reciprocity. For a general representation of government preferences, we identify bargaining outcomes that can be achieved using dominant strategy proposals for all countries. The resulting bargaining outcome is

⁵² For example, Proposition 3 can be restated as follows: Given the constructed mechanism, the set of dominant strategy proposals for the home country is a nonempty set whose members are the home-country proposal strategies for which, for all $\theta \in \Theta$, we have that $T_h^h(\theta) = \tau_{PO}(\theta)$ with $T_h^{*1}(\theta)$ and $T_h^{*2}(\theta)$ then specified in any fashion so that $\tilde{p}^w(\tau_{PO}(\theta), T_h^{*1}(\theta), T_h^{*2}(\theta)) = p_0^w$. Proposition 2 can be restated for each $\theta^{*i} \in \Theta^{*i}$ in an analogous fashion.

efficient relative to government preferences if and only if the initial tariff vector positions the initial world price at its “politically optimal” level. In symmetric settings, if the initial tariffs correspond to Nash tariffs, then the initial world price is indeed positioned at its politically optimal level, and the resulting bargaining outcome is efficient and ensures greater-than-Nash trade volumes and welfares for all countries.⁵³

In our model, dominant strategy proposals lead to assigned tariff vectors once a round of adjustments are made to ensure that multilateral reciprocity is maintained while not requiring any country to import more than is implied by its proposal. This simple sequence—tariff proposals followed by multilateral rebalancing—is broadly consistent with observed patterns identified by Bagwell et al. (2017) in the bargaining records for the GATT Torquay Round. A notable feature of bargaining behavior in this round is the lack of back-and-forth haggling over the levels of proposed tariffs. Moreover, the own-tariff offers contained in the proposals—instead of the requested tariff changes by others that complete each proposal—take center stage in our constructed mechanism, consistent with the additional finding of Bagwell, Staiger and Yurukoglu (2017) that, once initial proposals were on the table, tariff bargaining at Torquay focused on modifications to the offers in each proposal instead of requests. Finally, and also in broad alignment with observed patterns from the Torquay Round, when offers are modified to obtain multilateral reciprocity in the constructed mechanism, the modifications take the form of shallower, instead of deeper, offers.

Our work also contributes to the theoretical literature on tariff bargaining under nondiscrimination and reciprocity. For the case of nondiscriminatory tariffs, we generalize prior work by Bagwell and Staiger (1999) by allowing each country to make direct proposals regarding the tariffs of its trading partner(s) and by characterizing bargaining outcomes that obtain when all countries use dominant strategy proposals. This generalization that we consider here offers a more direct path to the GATT tariff bargaining records. With this generalization, we also show that a key finding—that efficiency can be achieved if and only if the politically optimal MFN tariff vector is induced—extends to a setting where all countries make direct tariff proposals and use dominant strategy proposals. We leave a many-good generalization of our results as an important direction for future work.

Finally, our work contributes by drawing novel links between multilateral tariff bargaining and the literatures that feature strategy-proof rationing schemes. As we argue, when tariff negotiations are constrained to satisfy the MFN rule and multilateral reciprocity, trade is conducted at a fixed world price. The rationing problems that emerge are then broadly similar to those that have been studied in other research on strategy-proof rationing. We mention above several extensions and directions for future work that could help to develop this relationship more fully. Interesting future work might also explore the relationship between multilateral tariff bargaining and strategy-proof exchange. In light of our findings, we view the MFN rule and the norm of multilateral reciprocity as providing governments with a pragmatic solution to a complicated multiparty bargaining problem. Going further, and building on previous work on strategy-proof exchange by Barbera and Jackson (1995), it would be interesting to determine whether circumstances exist under which some form of multilateral reciprocity is, in fact, necessary for dominant strategy implementation of multilateral tariff bargaining outcomes.

⁵³ The resulting bargaining outcome is inefficient if countries are symmetric (in that $p_N^w = p_{PO}^w$) but initial tariffs are asymmetric (in that $p_0^w \neq p_N^w$); moreover, even if countries and initial tariffs are symmetric (in that $p_0^w = p_N^w = p_{PO}^w$), the gains from further liberalization may be small if, due to previous negotiation rounds, the initial tariffs are already close to efficient levels. Our analysis thus provides a partial perspective on the failure of the Doha Round, since, in broad terms, that round began with asymmetric and low tariff bindings for major industrialized countries relative to those for emerging countries. We note, however, that the Doha Round also included issues that extend beyond the market-access concerns featured in our analysis (for a more extensive discussion of issues facing the Doha Round, see Bagwell and Staiger, 2014).

APPENDIX

LEMMA 3. For a vector of tariff proposals (T_h, T_{*1}, T_{*2}) such that (9) holds, the unique solution to the disagreement program is given by the tariff vector $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where $\tilde{\tau}$ is defined by (11).

PROOF. The assigned tariff vector $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where $\tilde{\tau}$ is defined by (11) obviously maintains the initial world price. It also satisfies the constraint that implied import volumes are not exceeded. To see this, observe that at the assigned tariff vector, foreign country $*i$ imports

$$(A.1) \quad M^{*i}(p^{*i}(T_{*i}^{*i}, p_0^w), p_0^w) = M_{*i}.$$

Next, at the assigned tariff vector, the home country imports

$$(A.2) \quad \begin{aligned} M(p(\tilde{\tau}, p_0^w), p_0^w) &= [1/p_0^w] E(p(\tilde{\tau}, p_0^w), p_0^w) \\ &= [1/p_0^w] \sum_{i=1,2} M^{*i}(p^{*i}(T_{*i}^{*i}, p_0^w), p_0^w) \\ &= [1/p_0^w] [M_{*1} + M_{*2}] \\ &< M_h, \end{aligned}$$

where the first equality follows from the home trade-balance condition (1), the second equality follows from (11) and the market-clearing condition (3), the third equality uses (A.1), and the inequality employs (9).

Finally, we confirm that, given the implied import volume limits and world price, the assigned tariff vector also maximizes the value of trade volume. To see this, we consider an arbitrary vector of tariffs, $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+$, for which $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and such that no country imports a volume in excess of its implied import volume. In other words, we consider any tariff vector that satisfies the constraints of the disagreement program.

Using the definition of $TV(\tau, \tau^{*1}, \tau^{*2})$ given in (8) and the home trade-balance condition (1), $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and the market-clearing condition (3), and (A.1), we obtain

$$\begin{aligned} TV(\tau, \tau^{*1}, \tau^{*2}) &= E(p(\tau, p_0^w), p_0^w) + \sum_{i=1,2} M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \\ &= 2[M^{*1}(p^{*1}(\tau^{*1}, p_0^w), p_0^w) + M^{*2}(p^{*2}(\tau^{*2}, p_0^w), p_0^w)] \\ &\leq 2[M_{*1} + M_{*2}] \\ &= 2[M^{*1}(p^{*1}(T_{*1}^{*1}, p_0^w), p_0^w) + M^{*2}(p^{*2}(T_{*2}^{*2}, p_0^w), p_0^w)] \\ &= TV(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2}), \end{aligned}$$

where the inequality follows from the restriction that imported volumes for foreign countries do not exceed their respective implied import volume limits. Since by (A.1), the assigned tariff vector ensures that both foreign countries import volumes that equal their respective implied import volume limits, the assigned vector of tariffs thus obtains the maximum value for the value of trade volume.

In fact, given the proposals, any tariff vector $(\tau, \tau^{*1}, \tau^{*2}) \neq (\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$ such that $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+$, $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and no country imports a volume in excess of its implied import volume must deliver strictly less trade volume: $TV(\tau, \tau^{*1}, \tau^{*2}) < TV(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$. This is because the proposal can be distinct while maintaining the initial world price only if $(\tau^{*1}, \tau^{*2}) \neq (T_{*1}^{*1}, T_{*2}^{*2})$. It follows that $(\tau, \tau^{*1}, \tau^{*2})$ implies a distinct trade volume for at least one foreign country;

thus, since $M^{*i}(p^{*i}(T_{*i}^{*i}, p_0^w), p_0^w) = M_{*i}$ and $M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \leq M_{*i}$ for all $i = 1, 2$, we may conclude that $TV(\tau, \tau^{*1}, \tau^{*2}) < TV(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2}) = 2[M_{*1} + M_{*2}]$. Consequently, our baseline rules are sufficient to ensure the unique determination of the assigned tariff vector when the tariff proposals satisfy (9). In particular, for such tariff proposals, our constructed mechanism must assign the tariff vector $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$, where $\tilde{\tau}$ is defined by (11). ■

LEMMA 4. For a vector of tariff proposals (T_h, T_{*1}, T_{*2}) such that (12) holds in the form of the first case (14), a solution to the disagreement program is given by the tariff vector $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2})$, where $\hat{\tau}^{*i}$ is defined by (15), $\hat{\tau}^{*j}$ is defined by (16), and $\hat{\tau} = T_h^h$ is defined by (17).

PROOF. The assigned tariff vector, $(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2}) \in \Upsilon_+$, delivers the initial world price by (17) and is also such that no country imports a volume that exceeds its implied import volume. It is immediately clear from (14)–(16) that neither foreign country imports a volume that exceeds its implied import volume. The home country imports a volume that equals its implied import volume. To see this, we, respectively, use the home trade-balance condition (1), $\tilde{p}^w(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2}) = p_0^w$, and the market-clearing condition (3), (16), and (15) to obtain

$$\begin{aligned} \text{(A.3)} \quad p_0^w M(p(\hat{\tau}, p_0^w), p_0^w) &= E(p(\hat{\tau}, p_0^w), p_0^w) \\ &= M^{*i}(p^{*i}(\hat{\tau}^{*i}, p_0^w), p_0^w) + M^{*j}(p^{*j}(\hat{\tau}^{*j}, p_0^w), p_0^w) \\ &= M^{*i}(p^{*i}(\hat{\tau}^{*i}, p_0^w), p_0^w) \\ &= p_0^w M_h, \end{aligned}$$

and thus $M(p(\hat{\tau}, p_0^w), p_0^w) = M_h$. Given this equality, it now follows that $\hat{\tau} = T_h^h$; thus, in the first case, the tariff that is assigned to the home country is the home country’s proposed tariff for itself.

Finally, we confirm that it is not possible to find another tariff vector that generates a greater value for trade volume while also delivering the initial world price and ensuring that no country imports a volume in excess of its implied import volume. To see this, we consider an arbitrary vector of tariffs, $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+$, for which $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and such that no country imports a volume in excess of its implied import volume. In other words, we consider any tariff vector that satisfies the constraints of the disagreement program.

Using the definition of $TV(\tau, \tau^{*1}, \tau^{*2})$ as in (8), the foreign-country trade-balance condition (2), $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$, and the implied market-clearing condition for good x , the requirement that the import volume of the home country does not exceed its implied import volume, and (A.3), we obtain

$$\begin{aligned} TV(\tau, \tau^{*1}, \tau^{*2}) &= p_0^w M(p(\tau, p_0^w), p_0^w) + p_0^w \sum_{i=1,2} E^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \\ &= 2p_0^w M(p(\tau, p_0^w), p_0^w) \\ &\leq 2p_0^w M_h \\ &= 2p_0^w M(p(\hat{\tau}, p_0^w), p_0^w) \\ &= TV(\hat{\tau}, \hat{\tau}^{*1}, \hat{\tau}^{*2}). \end{aligned}$$

For the first case, the assigned tariff vector thus obtains the maximum value for the value of trade volume, given the initial world price and the restriction that no country imports a volume in excess of its implied import volume. ■

LEMMA 5. For a vector of tariff proposals (T_h, T_{*1}, T_{*2}) such that (12) holds in the form of the second case (18), a solution to the disagreement program is given by the tariff vector $(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2})$, where $\bar{\tau}^{*i} = T_{*i}^{*i}$ is defined by (19), $\bar{\tau}^{*j}$ is defined by (20), and $\bar{\tau} = T_h^h$ is defined by (21).

PROOF. This assigned tariff vector, $(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2}) \in \Upsilon$, delivers the initial world price and is also such that no country imports a volume that exceeds its implied import volume. It is immediately clear from (19) and (20) that neither foreign country imports a volume that exceeds its implied import volume. The home country imports a volume that equals its implied import volume. To see this, we respectively use the home trade-balance condition (7), $\tilde{p}^w(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2}) = p_0^w$, and the market-clearing condition (3), (19), and (20) to obtain

$$\begin{aligned}
 \text{(A.4)} \quad p_0^w M(p(\bar{\tau}, p_0^w), p_0^w) &= E(p(\bar{\tau}, p_0^w), p_0^w) \\
 &= M^{*i}(p^{*i}(\bar{\tau}^{*i}, p_0^w), p_0^w) + M^{*j}(p^{*j}(\bar{\tau}^{*j}, p_0^w), p_0^w) \\
 &= M_{*i} + M^{*j}(p^{*j}(\bar{\tau}^{*j}, p_0^w), p_0^w) \\
 &= M_{*i} + p_0^w M_h - M_{*i} \\
 &= p_0^w M_h,
 \end{aligned}$$

and thus $M(p(\bar{\tau}, p_0^w), p_0^w) = M_h$. Given this equality, it now follows that $\bar{\tau} = T_h^h$; thus, in the second case as well, the home country’s assigned tariff is the tariff that it proposed for itself.

Finally, we confirm that it is not possible to find another tariff vector that generates a greater value for trade volume while also delivering the initial world price and ensuring that no country imports a volume in excess of its implied import volume. To see this, we employ a argument similar to that above for the first case. Specifically, we consider an arbitrary vector of tariffs, $(\tau, \tau^{*1}, \tau^{*2}) \in \Upsilon_+$, for which $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$ and such that no country imports a volume in excess of its implied import volume. In other words, we consider any tariff vector that satisfies the constraints of the disagreement program.

Using the definition of $TV(\tau, \tau^{*1}, \tau^{*2})$ given in (8), the foreign-country trade-balance condition (2), $\tilde{p}^w(\tau, \tau^{*1}, \tau^{*2}) = p_0^w$, and the implied market-clearing condition for good x , the requirement that the import volume of the home country does not exceed its implied import volume, and (A.4), we obtain

$$\begin{aligned}
 TV(\tau, \tau^{*1}, \tau^{*2}) &= p_0^w M(p(\tau, p_0^w), p_0^w) + p_0^w \sum_{i=1,2} E^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w) \\
 &= 2p_0^w M(p(\tau, p_0^w), p_0^w) \\
 &\leq 2p_0^w M_h \\
 &= 2p_0^w M(p(\bar{\tau}, p_0^w), p_0^w) \\
 &= TV(\bar{\tau}, \bar{\tau}^{*1}, \bar{\tau}^{*2}).
 \end{aligned}$$

Thus, for the second case, the assigned tariff vector obtains the maximum value for the value of trade volume, given the initial world price and the restriction that no country import a volume in excess of its implied import volume. ■

PROPOSITION 2. Given the constructed mechanism, the set of dominant strategy proposals for foreign country $*i$ is nonempty and, in fact, contains a singleton defined by $T_{*i,PO}$, where $T_{*i,PO}^{*i} = \tau_{PO}^{*i}$, $T_{*i,PO}^{*j} = \tau_0^{*j}$, and $T_{*i,PO}^h = T_{*i}^h(\tau_{PO}^{*i})$ is then uniquely set to deliver p_0^w as the market-clearing world price.

PROOF. Consider foreign country $*i$. We wish to argue that foreign country $*i$'s dominant strategy is to propose $T_{*i,PO}$ defined by $T_{*i,PO}^{*i} = \tau_{PO}^{*i}$ with $T_{*i,PO}^{*j} = \tau_0^{*j}$ and $T_{*i,PO}^h = T_{*i}^h(\tau_{PO}^{*i})$ then set to deliver the initial world price, p_0^w , as the market-clearing world price under the proposal. By assumption, $T_{*i,PO} \in \Upsilon$ exists. We compare this proposal strategy to an alternative strategy T_{*i} associated with a different own tariff for foreign country $*i$, $T_{*i}^{*i} \equiv \tau^{*i} \neq T_{*i,PO}^{*i} = \tau_{PO}^{*i}$, where as described $T_{*i}^{*j} = \tau_0^{*j}$ and $T_{*i}^h = T_{*i}^h(\tau^{*i})$ then follow. Let $M_{*i,PO} \equiv M^{*i}(p^{*i}(\tau_{PO}^{*i}, p_0^w), p_0^w)$ and $M_{*i} = M^{*i}(p^{*i}(\tau^{*i}, p_0^w), p_0^w)$ denote the corresponding implied import volumes.

If the proposals of the home country and foreign country $*j$ are such that the tariff proposals agree when foreign country $*i$ proposes $T_{*i,PO}$, then the alternative proposal T_{*i} cannot possibly represent an improvement for foreign country $*i$. This follows since under proposal $T_{*i,PO}$, foreign country $*i$ enjoys its favorite local price for the given initial world price. By similar reasoning, if the proposals of the home country and foreign country $*j$ are such that the home country is long when foreign country $*i$ proposes $T_{*i,PO}$, then the proposal $T_{*i,PO}$ again results in foreign country $*i$'s favorite local price for the given initial world price, and so the alternative proposal T_{*i} cannot possibly represent an improvement for foreign country $*i$.⁵⁴ The remaining possibility is that the proposal $T_{*i,PO}$ and the proposals of other countries are such that the home country is short. To address this remaining possibility, we distinguish between two cases for the alternative proposal, T_{*i} .

The first case is that the alternative proposal entails a lower tariff for foreign country $*i$: $\tau^{*i} < \tau_{PO}^{*i}$ and thus $M_{*i} > M_{*i,PO}$. Given that the proposals of other countries are such that the home country is short when foreign country $*i$ proposes $T_{*i,PO}$, the home country will again be short in this first case under the alternative proposal, T_{*i} . If foreign country $*i$ is not randomly selected to go first, then its assigned tariff does not depend on its proposal (conditional on the home country being short), and so it is then indifferent between the two proposals. If foreign country $*i$ is randomly selected to go first, then it enjoys its favorite local price under the proposal $T_{*i,PO}$ if $M_{*i,PO} \leq p_0^w M_h$. The alternative proposal then cannot represent an improvement for foreign country $*i$. If foreign country $*i$ is randomly selected to go first and $M_{*i,PO} > p_0^w M_h$, then the assigned tariff for foreign country $*i$ is $\hat{\tau}^{*i}$, where $\hat{\tau}^{*i}$ satisfies $M^{*i}(p^{*i}(\hat{\tau}^{*i}, p_0^w), p_0^w) = p_0^w M_h$. The alternative proposal entails an even higher implied import volume, $M_{*i} > M_{*i,PO}$, and thus leads to the same assigned tariff for foreign country $*i$. Hence, when the home country is short, and in the first case where $\tau^{*i} < \tau_{PO}^{*i}$, we conclude that the politically optimal proposal $T_{*i,PO}$ is always at least weakly preferred to the alternative proposal, T_{*i} .

The second case is that the alternative proposal entails a higher tariff for foreign country $*i$: $\tau^{*i} > \tau_{PO}^{*i}$ and thus $M_{*i} < M_{*i,PO}$. Suppose first that the other proposals are such that the home country remains short under the alternative proposal (despite the fact that the alternative proposal implies a lower trade volume for foreign country $*i$). If foreign country $*i$ is not randomly selected to go first, then its assigned tariff does not depend on its proposal (conditional on the home country being short) and so it is then indifferent between the two proposals. If foreign country $*i$ is randomly selected to go first, then it enjoys its favorite local price under the proposal $T_{*i,PO}$ if $M_{*i,PO} \leq p_0^w M_h$. The alternative proposal then cannot represent an improvement for foreign country $*i$. If foreign country $*i$ is randomly selected to go first and $M_{*i,PO} > p_0^w M_h$, then the assigned tariff for foreign country $*i$ under the proposal $T_{*i,PO}$ is $\hat{\tau}^{*i}$, where $\hat{\tau}^{*i}$ satisfies $M^{*i}(p^{*i}(\hat{\tau}^{*i}, p_0^w), p_0^w) = p_0^w M_h$. If the alternative proposal satisfies $M_{*i} \geq p_0^w M_h$, then the same tariff is assigned for foreign country $*i$. If the alternative proposal satisfies $M_{*i} < p_0^w M_h$, then the assigned tariff for foreign country $*i$ under the alternative proposal T_{*i} is $\bar{\tau}^{*i}$, which satisfies $M^{*i}(p^{*i}(\bar{\tau}^{*i}, p_0^w), p_0^w) = M_{*i}$. Given $M_{*i} < p_0^w M_h < M_{*i,PO}$, we then see that $\bar{\tau}^{*i} > \hat{\tau}^{*i} > \tau_{PO}^{*i}$. We conclude that the politically optimal proposal $T_{*i,PO}$ is then preferred to the alternative proposal T_{*i} , since the assigned tariff for foreign country $*i$ is closer to its politically optimal reaction tariff under the proposal $T_{*i,PO}$.

⁵⁴ Recall that, under the constructed mechanism defined above, foreign country $*i$'s proposed tariff for itself is assigned under agreement and also under disagreement when the home country is long.

Continuing with the second case where $M_{*i} < M_{*i,PO}$, we suppose, second, that the other proposals are such that the home country is short under the proposal $T_{*i,PO}$ but is not short under the alternative proposal T_{*i} . Thus, we now focus on the scenario where $M_{*i} + M_{*j} \leq p_0^w M_h < M_{*i,PO} + M_{*j}$. Since $M_{*i} > 0$, the other proposals then must be such that $M_{*j} < p_0^w M_h$. Under the proposal $T_{*i,PO}$, if foreign country $*i$ is randomly selected, then it enjoys a trade volume of either $M_{*i,PO}$ (if $M_{*i,PO} \leq p_0^w M_h$) or $p_0^w M_h$ (if $M_{*i,PO} > p_0^w M_h$). Under the proposal $T_{*i,PO}$, if foreign country $*i$ is not randomly selected, then it enjoys a trade volume of $p_0^w M_h - M_{*j} > 0$. Under the alternative proposal, the home country is not short, and so foreign country $*i$ enjoys a trade volume of $M_{*i} \leq p_0^w M_h - M_{*j}$. Thus, foreign country $*i$ enjoys a trade volume closer to $M_{*i,PO}$ under the proposal $T_{*i,PO}$ than under the proposal T_{*i} .

Having now considered all possible trade-volume scenarios when foreign country $*i$ proposes $T_{*i,PO}$, and the corresponding possibilities were foreign country $*i$ instead to propose an alternative proposal T_{*i} , we now conclude that alternative proposals T_{*i} that entail $\tau^{*i} \neq \tau_{PO}^{*i}$ are dominated by the proposal $T_{*i,PO}$ that entails $\tau^{*i} = \tau_{PO}^{*i}$.

Finally, we argue that any tariff proposal for foreign country $*i$ such that $T_{*i} \neq T_{*i,PO}$ is not a dominant strategy. Let T_{*i} be any tariff proposal for foreign country $*i$ for which $T_{*i}^{*i} \neq \tau_{PO}^{*i}$. Proposals $T_{*i,PO}$ and T_{*i} are in Υ . Of these two proposals, the one with the lower proposed tariff for foreign country $*i$ must also have the lower proposed tariff for the home country. Denote this tariff proposal vector as $T_{*i,L}$. Suppose now that the home country proposes the vector $T_{*i,L}$ as well, and that foreign country $*j$ proposes the initial or status quo tariff vector, $(\tau_0, \tau_0^{*1}, \tau_0^{*2}) \in \Upsilon$. Given these tariff proposals for foreign country $*j$ and the home country, the home country is long or agrees whether foreign country $*i$ proposes T_{*i} or $T_{*i,PO}$; thus, in both cases, foreign country $*i$ imports a volume that equals its implied import volume, which is M_{*i} under the proposal T_{*i} and $M_{*i,PO}$ under the proposal $T_{*i,PO}$, respectively, with $M_{*i} \neq M_{*i,PO}$. The tariff proposal $T_{*i,PO}$ is then strictly better for foreign country $*i$ than is the tariff proposal T_{*i} . ■

PROPOSITION 3. *Given the constructed mechanism, the set of dominant strategy proposals for the home country is a nonempty set whose members are the home-country proposal strategies for which $T_h^h = \tau_{PO}$, with T_h^{*1} and T_h^{*2} then specified in any fashion so that $\tilde{p}^w(\tau_{PO}, T_h^{*1}, T_h^{*2}) = p_0^w$.*

PROOF. We wish to argue that the set of dominant strategies for the home country is defined by a set of proposals under which the home country proposes the tariff τ_{PO} for itself and tariffs for the foreign countries that in combination with τ_{PO} deliver the initial world price as the market-clearing world price. By assumption, we know that some such proposals exist in Υ . To fix ideas, let us select any proposal $T_{h,PO}$ from this set, where the proposal strategy $T_{h,PO}$ is thus defined by $T_{h,PO}^h = \tau_{PO}$ with $(T_{h,PO}^{*1}, T_{h,PO}^{*2})$ then satisfying $p_0^w \equiv \tilde{p}^w(\tau_{PO}, T_{h,PO}^{*1}, T_{h,PO}^{*2})$. We compare this proposal strategy to an alternative strategy T_h for which $T_h^h \equiv \tau \neq T_{h,PO}^h = \tau_{PO}$, with (T_h^{*1}, T_h^{*2}) then satisfying $p_0^w \equiv \tilde{p}^w(\tau, T_h^{*1}, T_h^{*2})$. Recall that, given the initial world price, the home country's implied import volume is determined by its proposed tariff for itself. Let $M_{h,PO} \equiv M(p(\tau_{PO}, p_0^w), p_0^w)$ and $M_h \equiv M(p(\tau, p_0^w), p_0^w)$ denote the corresponding implied import volumes.

If the proposals of the foreign countries are such that the tariff proposals agree when the home country proposes $T_{h,PO}$, then the alternative proposal T_h cannot possibly represent an improvement for the home country. This follows since, under proposal $T_{h,PO}$, the home country enjoys its favorite local price for the given initial world price. By similar reasoning, if the proposals of the foreign countries are such that the home country is short when the home country proposes $T_{h,PO}$, then the home country is again assigned the tariff τ_{PO} and thus enjoys its favorite local price for the given initial world price; hence, once again, the alternative proposal T_h cannot possibly represent an improvement for the home country.

The remaining possibility is that the proposal $T_{h,PO}$ and the proposals of the foreign countries are such that the home country is long so that $p_0^w M_{h,PO} > M_{*1} + M_{*2}$. Under the proposal $T_{h,PO}$,

the home country is then assigned the tariff $\tilde{\tau}$ defined given the foreign proposals $(T_{*1}^{*1}, T_{*2}^{*2})$ to deliver $p_0^w \equiv \tilde{p}^w(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$. Given the foreign proposals, if the home country remains long under the alternative proposal T_h , then the same tariff vector is assigned, and so the alternative proposal fails to offer an improvement for the home country. Suppose then the home country obtains agreement or is short under the alternative proposal T_h . For the given foreign proposals, under any of these cases for the alternative proposal T_h , the tariff that the home country proposes for itself, τ , is assigned, and so we have that $p_0^w M(p(\tau, p_0^w), p_0^w) = p_0^w M_h \leq M_{*1} + M_{*2} = p_0^w M(p(\tilde{\tau}, p_0^w), p_0^w) < p_0^w M_{h,PO}$, where the final equality follows from (A.2). It follows that $\tau_{PO} < \tilde{\tau} \leq \tau$, and so the alternative proposal T_h results in an assigned home-country tariff τ that is (weakly) further from τ_{PO} than is the assigned home-country tariff $\tilde{\tau}$ that results from the proposal $T_{h,PO}$.

Having now considered all possible trade-volume scenarios when the home country proposes $T_{h,PO}$, and the corresponding possibilities were the home country instead to propose an alternative proposal T_h , we conclude that alternative proposals T_h that entail $\tau \neq \tau_{PO}$ are dominated by the proposal $T_{h,PO}$ that specifies a home-country tariff of τ_{PO} . Since $T_{h,PO}$ is an arbitrary selection from the set of home proposals for which $T_{h,PO}^h = \tau_{PO}$, we conclude that the alternative proposal T_h is dominated by all such home-country proposal strategies that specify a home-country tariff of τ_{PO} .

Finally, we compare distinct home-country proposal strategies that both specify a home-country tariff of τ_{PO} , and we argue that for any given foreign proposals, such home-country proposal strategies must result in the same assigned tariff vector. A corollary is that one home-country strategy in this class cannot dominate another. To develop this argument, we use $T_{h,PO} = (\tau_{PO}, T_{h,PO}^{*1}, T_{h,PO}^{*2})$ to denote one such proposal, and we denote an alternative such proposal as $T'_{h,PO} = (\tau_{PO}, T_{h,PO}^{*1'}, T_{h,PO}^{*2'})$, where $(T_{h,PO}^{*1}, T_{h,PO}^{*2}) \neq (T_{h,PO}^{*1'}, T_{h,PO}^{*2'})$. For given foreign proposals, if $T_{h,PO}$ obtains agreement, then $T'_{h,PO}$ obtains agreement as well. Thus, the same tariff vector would be assigned under either home-country proposal. Given the foreign proposals, if the home country is short under $T_{h,PO}$, then the home country similarly is short under $T'_{h,PO}$. For this case, whether the home country proposes $T_{h,PO}$ or $T'_{h,PO}$, it is assigned the tariff τ_{PO} and imports a volume that equals its implied import volume, $M_{h,PO}$. Further, the assigned foreign tariff vectors are independent of whether the home country proposes $(T_{h,PO}^{*1}, T_{h,PO}^{*2})$ or $(T_{h,PO}^{*1'}, T_{h,PO}^{*2'})$ for the foreign countries. The same tariff vector thus would be assigned under either home-country proposal.⁵⁵ Finally, given the foreign proposals, if the home country is long under $T_{h,PO}$, then the home country is long as well under $T'_{h,PO}$. The assigned tariff vector is then $(\tilde{\tau}, T_{*1}^{*1}, T_{*2}^{*2})$ whether the home country proposes $T_{h,PO}$ or $T'_{h,PO}$.

We now summarize our arguments and complete the proof. Having considered all possible trade-volume scenarios when the home country proposes $T_{h,PO}$, and the corresponding possibilities were the home country instead to propose alternative proposal T_h , we conclude that alternative proposals T_h that entail $\tau \neq \tau_{PO}$ are dominated by the proposal $T_{h,PO}$ that entails $\tau = \tau_{PO}$. We have also argued that the home-country proposal $T_{h,PO}$ offers the same assigned tariff vector for any given foreign proposals as does the home-country proposal $T'_{h,PO}$, where $T_{h,PO}$ and $T'_{h,PO}$ both entail $\tau = \tau_{PO}$ but propose different foreign tariffs.

Finally, we argue that any tariff proposal for the home country such that $T_h^h \neq \tau_{PO}$ is not a dominant strategy. Let T_h be any tariff proposal for the home country for which $T_h^h \neq \tau_{PO}$. Proposals $T_{h,PO}$ and T_h are in Υ . Of these two proposals, and given that both must deliver market clearing at the world price p_{PO}^w , the one with the lower proposed tariff for the home country must be associated with a higher aggregate foreign import volume. Denote this tariff proposal vector as $T_{h,L}$. This proposal implies positive import volumes for each foreign country $*i$ when the tariff proposed for that country by the home country, $T_{h,L}^{*i}$, is imposed. Since

⁵⁵ This statement is understood to refer to expected values when the home country is short and a “first” foreign firm is selected at random. The key point is that the selection probability is random and thus independent of the specific home-country proposal.

$\tilde{p}^w(T_{h,L}^h, T_{h,L}^{*1}, T_{h,L}^{*2}) = p_0^w$, we know from our assumptions that the same positive import volume for foreign country $*i$ can be obtained with an alternative proposal in Υ by foreign country $*i$ such that $T_{*i}^{*i} = T_{h,L}^{*i}$ and $T_{*i}^{*j} = \tau_0^{*j}$, with T_{*i}^h then selected so that $\tilde{p}^w(T_{*i}^h, T_{*i}^{*1}, T_{*i}^{*2}) = p_0^w$. Suppose now that the foreign countries propose the alternative proposals in Υ just constructed. The implied import volumes, M_{*1} and M_{*2} , are then such that $M_{*1} + M_{*2} = p_0^w M_{h,L}$, where $M_{h,L} \equiv M(p(T_{h,L}^h, p_0^w), p_0^w)$. Given these constructed foreign proposals, the home country is short or agrees whether it proposes $T_{h,PO}$ or T_h ; thus, for both proposals, the home country imports a volume that is equal to its implied import volume, which is $M_{h,PO}$ under the proposal $T_{h,PO}$ and $M_h \neq M_{h,PO}$ under the proposal T_h , respectively. The tariff proposal $T_{h,PO}$ is then strictly better for the home country than is the proposal T_h . ■

REFERENCES

- ABDULKADIROGLU, A., AND T. SONMEZ, "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems," *Econometrica* 66 (1998), 689–701.
- BAGWELL, K., AND R. W. STAIGER, "An Economic Theory of GATT," *American Economic Review* 89(1) (1999), 215–48.
- , AND ———, *The Economics of the World Trading System* (Cambridge, MA: The MIT Press, 2002).
- , AND ———, "Multilateral Trade Negotiations, Bilateral Opportunism and the Rules of GATT," *The Journal of International Economics* 67(2) (2005), 268–94.
- , AND ———, "Can the Doha Round Be a Development Round? Setting a Place at the Table," in R. C. Feenstra and A. M. Taylor, eds., *Globalization in an Age of Crisis: Multilateral Economic Cooperation in the Twenty-First Century* (Chicago: University of Chicago Press, 2014), 91–124.
- , AND ———, "The Design of Trade Agreements," in K. Bagwell and R. W. Staiger, eds., *The Handbook of Commercial Policy*, Vol. 1A (Amsterdam: Elsevier, 2016), 435–529.
- , ———, AND A. YURUKOGLU, "Multilateral Trade Bargaining: A First Look at the GATT Bargaining Records," mimeo, July 2017.
- BARBERA, S., "Strategyproof Social Choice," in K. J. Arrow, A. Sen, and K. Suzumura, eds., *Handbook of Social Choice and Welfare*, Vol. 2 (Amsterdam: Elsevier, 2011), 731–831.
- , AND M. O. JACKSON, "Strategy-Proof Exchange," *Econometrica* 63(1) (1995), 51–87.
- , ———, AND A. NEME, "Strategy-Proof Allotment Rules," *Games and Economic Behavior* 18(1) (1997), 1–21.
- BENASSY, J.-P. *The Economics of Market Disequilibrium* (New York: Academic Press, 1982).
- BOGOMOLNAIA, A., AND H. MOULIN, "A New Solution to the Random Assignment Problem," *Journal of Economic Theory* 100 (2001), 295–328.
- CURZON, G., *Multilateral Commercial Diplomacy: The General Agreement on Tariffs and Trade and Its Impacts on National Commercial Policies and Techniques* (New York: Praeger, 1966).
- DIXIT, A., "Strategic Aspects of Trade Policy," in T. F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress* (New York, NY: Cambridge University Press, 1987), 329–62.
- GUL, F., H. SONNENSCHNEIN, AND R. WILSON, "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory* 39 (1986), 155–90.
- INTERIM COMMISSION FOR THE ITO (ICITO), *The Attack on Trade Barriers: A Progress Report on the Operation of the General Agreement on Tariffs and Trade* (Geneva: Interim Commission for the International Trade Organisation, August 1949).
- JACKSON, M. O., "A Crash Course in Implementation Theory," *Social Choice and Welfare* 18 (2001), 655–708.
- LUDEMA, R., "International Trade Bargaining and the Most-Favored Nation Clause," *Economics and Politics* 3(1) (1991), 1–20.
- MOULIN, H., "Priority Rules and Other Asymmetric Rationing Methods," *Econometrica* 68(3) (2000), 643–84.
- SONMEZ, T., "Consistency, Monotonicity, and the Uniform Rule," *Economic Letters* 46 (1994), 229–35.
- SPRUMONT, Y., "The Division Problem with Single-Peaked Preferences: A Characterization of the Uniform Allocation Rule," *Econometrica* 59(2) (1991), 509–19.