Submitted to *Econometrica*

WHY IS TRADE NOT FREE? A REVEALED PREFERENCE APPROACH

RODRIGO ADÃO
Booth School of Business, University of Chicago and NBER

JOHN STURM BECKO
Department of Economics, Princeton University

ARNAUD COSTINOT
Department of Economics, MIT and NBER

DAVE DONALDSON
Department of Economics, MIT and NBER

A prominent explanation for why trade is not free is politicians' desire to protect some of their constituents at the expense of others. In this paper we develop a methodology that can be used to reveal the welfare weights that a nation's import tariffs implicitly place on different groups of society. Applied in the context of the United States in 2017, this method implies that redistributive trade protection accounts for about a third of US tariff variation and that it causes large monetary transfers between US individuals, mostly driven by differences in welfare weights across sectors of employment. In terms of political economy mechanisms, the welfare weights that we estimate are consistent with lobbying being a significant driver of US trade protection. Perhaps surprisingly, swing states in US presidential elections play a much smaller role.

KEYWORDS: International trade policy, Political economy.

1. INTRODUCTION

International trade is decidedly not free. Countries around the world routinely impose tariffs and other barriers to trade. One prominent explanation for why they do so is redistributive politics. Even if trade restrictions reduce the size of the pie, they can increase the slice received by some at the expense of others.

The existing literature on the political economy of trade policy reviewed in Rodrik (1995), Gawande and Krishna (2003), and McLaren (2016) is rich with theories explaining why politicians may choose to favor particular constituents of society. Direct democracy may lead politicians to cater to the median voter; competition for electoral-college votes, as in a US presidential election, may bias their preferences towards swing-state voters; and lobbying activities may give a disproportionate weight to politically-organized sectors. In each of these theories, a fully-specified political process is combined with a typically stylized economic environment to generate predictions about the structure of trade protection. Empirical work, in turn, can test

Rodrigo Adão: rodrigo.adao@chicagobooth.edu
John Sturm Becko: becko@princeton.edu
Arnaud Costinot: ddonald@mit.edu
Dave Donaldson: ddonald@mit.edu

We are grateful to Keelan Beirne, Shotaro Beppu, Robin Li, Noah Siderhurst and Akash Thakkar for outstanding research assistance, to Thomas Chaney, Paola Conconi, Ben Faber, Pablo Fajgelbaum, Penny Goldberg, Gordon Hanson, John McLaren, Giovanni Maggi, Bob Staiger, Jon Vogel and seminar audience members for helpful comments, and to Chad Bown, Aksel Erbahar, Alessandro Ferrari, Kari Heerman, Ralph Ossa, and Justin Pierce for generous help with data. We also gratefully acknowledge the financial support of NSF grant number 2242367.

whether such specific predictions hold in practice. The influential tariff formula arising from Grossman and Helpman's (1994) protection for sale model and its subsequent test by Goldberg and Maggi (1999) embodies this canonical approach.[1]

In this paper, we propose an alternative approach, grounded in the logic of revealed preference. It is designed to flexibly identify the welfare weights that a nation's import tariffs implicitly place on different members of society, i.e., who the politically-favored are, without imposing any a priori restrictions on the reasons for such favoritism. When implemented in the context of the United States in 2017, our analysis implies that the redistributive motive for trade protection accounts for a significant fraction of tariff variation and causes large monetary transfers between US individuals, mostly driven by differences in welfare weights across sectors of employment. Perhaps surprisingly, differences in welfare weights across states play a much smaller role.

The starting point of our revealed-preference approach is a general tariff formula that only relies on the assumption that trade taxes are set via some political process that is constrained Pareto efficient. That is, we assume that, given available policy instruments, politicians' incentives in the country of interest ("Home") are such that there does not exist a change in trade taxes that could strictly increase the utility of some of their domestic constituents without strictly decreasing the utility of others—a requirement satisfied by leading political economy models, as we discuss further below. For any such political process, our general tariff formula states that, for any good g, trade taxes t_g must satisfy

$$t_g = -\sum_n \beta(n) \times \frac{\partial[\omega(n) - \bar{\omega}]}{\partial m_g} + \text{Residual}_g, \tag{1}$$

where the derivative $\partial[\omega(n) - \bar{\omega}]/\partial m_g$ denotes the marginal change in the real earnings of a given individual n, relative to the average earnings change in the population, associated with a marginal increase in the (net) imports m_g of good g. The term $\beta(n)$ denotes the social marginal return of a transfer to individual n that Home's trade policymaking process arrives at, whatever is the underlying process that causes it to "choose" this particular point on the economy's Pareto frontier. Finally, Residual_g captures the effects of non-redistributive motives for protection, namely terms-of-trade manipulation and second-best corrections for distortions, which we also fully characterize.

Empirically, we propose to treat our general tariff formula (1) as a regression equation in which the dependent variable is the trade tax t_g, the regressors are $\{\partial[\omega(n)-\bar{\omega}]/\partial m_g\}$, and the coefficients of interest are (minus) the vector of welfare weights $\{\beta(n)\}$. Intuitively, the choice of a higher tariff on a given good g reveals a stronger preference of society for the individuals whose real income would have been more negatively affected by a marginal increase in good g's imports. Because the estimated welfare weights are valid regardless of the underlying political process that gives rise to them, our revealed-preference approach highlights a natural division of labor in the study of how distributional forces can lead to protectionism. First, we can draw on the vast body of recent work on the general-equilibrium impact of trade policy on earnings and prices in order to construct empirically credible measures of $\{\partial[\omega(n) - \bar{\omega}]/\partial m_g\}$. And second, given the estimated welfare weights, we can go on to evaluate the importance of the redistributive motive for trade policy and explore its political determinants.

We apply our general formula to study the redistributive motive embedded in the trade policy of the United States in 2017—that is, before the changes introduced in 2018. To measure

[1]A non-exhaustive list of empirical papers testing the predictions of Grossman and Helpman (1994) and various extensions of the original "protection for sale" model includes Gawande and Bandyopadhyay (2000), Mitra et al. (2002), Matschke and Sherlund (2006), Bombardini (2008), and Gawande et al. (2009).

$\partial[\omega(n) - \bar{\omega}]/\partial m_g$, we develop a model of the US economy that features heterogeneous exposure to international trade across US regions and sectors, both directly via exports and imports and indirectly via input-output and domestic trade linkages. Following Fajgelbaum et al. (2020), we model the rest of the world as a series of import demand and export supply curves whose elasticities are estimated from the variation in the prices and quantities of US exports and imports induced by the 2018 tariff changes. We calibrate other model parameters to match available data on trade and production across sectors and states in 2017. This economic model yields estimates of $\{\partial[\omega(n) - \bar{\omega}]/\partial m_g\}$ for all US individuals based on their region and sector of employment, and for thousands of country-product varieties g. Critically, these predictions are consistent with the differential changes in labor market outcomes across US regions and sectors observed during the US-China trade war, as we establish using the testing procedure of Adao et al. (2025).

Armed with estimates of $\{\partial[\omega(n) - \bar{\omega}]/\partial m_g\}$ as well as measures of import tariffs t_g for the United States in 2017, we turn to the estimation of the welfare weights using equation (1). The key assumption in our implementation is that after controlling for a subset of non-redistributive motives for trade protection, such as terms-of-trade manipulation, our measures of the economic return from imports $\{\partial[\omega(n) - \bar{\omega}]/\partial m_g\}$ are orthogonal to the other motives left in Residual_g, such as the correction of domestic distortions. This rules out, for instance, the possibility that import restrictions on goods that favor a subset of individuals also systematically alleviate or worsen externalities due to carbon emissions, since measures of such motives will not be included in our control set.

Our baseline estimates reveal that US individuals employed in different sectors and states differ substantially in their welfare weights, with a long upper tail. They imply that, from society's perspective, a hypothetical \$1 received by an individual at the 90th, 95th, and 99th percentiles of our estimates are equivalent to \$1.17, \$1.74, and \$1.91, respectively, received by an individual at the 10th percentile. These differences are mostly driven by a large dispersion in sector-specific welfare weights, with the dispersion in state-specific welfare weights playing only a minor role. Individuals employed in the three sectors with highest welfare weights, Apparel, Vehicles, and Metals, have a social marginal utility of income that is 138% higher than the US average. In contrast, the social marginal utility of income enjoyed by individuals in the three states with the highest welfare weights, Florida, Vermont, and Wyoming, is only 7.5% higher.

Reassuringly, our baseline estimates are highly correlated with those obtained from a battery of alternative specifications meant to deal, in a theory-consistent way, with the existence of non-tariff measures, such as antidumping and countervailing duties; production subsidies that may vary either across states or sectors; and income taxes that are imposed both at the state and federal level. They are also highly correlated with estimates that account for exogenous constraints on US tariffs, such as those due to WTO membership; misspecification of terms-of-trade controls, perhaps due to negotiated trade taxes; and reverse causality between tariffs and imports.

We conclude by using our estimated welfare weights to quantify the importance of redistributive trade policy in the United States and shed light on its primitive determinants. A number of novel insights emerge. In terms of magnitude, the redistributive motive accounts for a significant share of the cross-sectional variation in US tariffs: 32% in 2017, out of which 28% derives from sector-specific welfare weights and 4% from state-specific ones. Likewise, the monetary transfers associated with redistributive trade protection are large, with a long upper tail. Transitioning to a counterfactual US economy with welfare weights equalized across individuals—and hence with tariffs purged of their redistributive component—would reduce the real earnings of individuals at the 90th, 95th, and 99th percentiles of the gains from redistributive trade protection by \$538, \$1,727, and \$2,861, respectively. In contrast, individuals at the 10th percentile would gain \$621.

In terms of political economy mechanisms, we use our estimates of welfare weights to discriminate between two leading explanations for the existence of tariffs: sectors' abilities to lobby, as formalized by Grossman and Helpman (1994), and states' abilities to swing presidential elections, as formalized by Ma and McLaren (2018). We show that the cross-sectional variation in our welfare weights is consistent with lobbying being a significant driver of US trade protection, with a lobbying dummy explaining about a fifth of the overall variation in welfare weights across sectors and regions. In contrast, we do not find that individuals from swing states receive significantly higher weights.

Related Literature

This paper combines central ideas from the fields of political economy, public finance, and international trade to shed new light on the nature and importance of redistributive trade protection.

The literature on redistributive politics is rich and varied. Persson and Tabellini (2002) describe its focus to be on "how the interplay between democratic institutions and self-seeking individuals, lobby groups, and parties determines the degree of redistribution." Numerous theoretical models of the previous interactions have been developed and applied to the analysis of trade policy. Classic examples include Mayer (1984), Grossman and Helpman (1994), Dixit and Londregan (1996), and Grossman and Helpman (2005). This line of research has been complemented by detailed empirical studies highlighting different dimensions of these interactions, from lobbying expenditure on trade policy, as in Bombardini and Trebbi (2012) and Kim (2017), to the votes of elected officials on trade bills and trade agreements, as in Conconi et al. (2014) and Conconi et al. (2024).

To evaluate the importance of redistributive trade protection and identify the individuals who benefit from it, this paper leverages a key implication of the previous models: competition, whether between political candidates or lobbies, ensures that trade policies are constrained Pareto efficient. Our approach, which relies on this high-level property of the political process and abstracts from the specifics, has both costs and benefits. The main cost is in terms of scope. Since we do not specify a structural model of the political process, we cannot say anything about how changes in this process, for instance new limits on political contributions or changes in electoral competition, may endogenously affect welfare weights and, in turn, redistributive trade protection. We can only estimate the welfare weights associated with current political institutions. The main benefit of our approach is in terms of robustness. The upside of not relying on a specific structural model of the political process is, of course, that the validity of our empirical estimates does not rely on it either.[2]

Although our empirical strategy is agnostic about the political process through which tariffs come about, our estimates of welfare weights do have implications for what this underlying

[2] In abstracting from the details of the political process and focusing on the associated welfare weights, our analysis also relates to Baldwin (1987) who stresses the equivalence between tariffs chosen by lobbying-influenced policy makers and those maximizing a social welfare function with extra weight on profits in the context of a partial-equilibrium model. In the same spirit, Helpman (1995) points out that the political support function pioneered by Hillman (1982) can be viewed as a reduced-form of the influence-driven contribution approach in Grossman and Helpman (1994), though the emphasis is on the limits of this equivalence and the distinct predictions of different political economy models.

political process may or may not be. To take a well-known example, Grossman and Helpman's (1994) protection for sale model predicts that the welfare weights of different individuals only depend on whether they are employed in politically-organized sectors or not. In Ma and McLaren (2018), individuals instead only receive positive welfare weights if employed in a swing state. These are stark predictions that one can formally test using our empirical estimates, as we will do in Section 6. More generally, one can view the gaps in the social marginal utility of income of different individuals as the analog of the "wedges" in the misallocation literature (see e.g. Chari et al., 2007 and Hsieh and Klenow, 2009). In this sense, our estimates can be used to discriminate between existing political economy models in the same way that wedges can be used to discriminate between different sources of distortions.

From the public finance literature, our paper borrows the general idea of using observed taxes to flexibly estimate welfare weights, as in, for instance, Werning (2007), Bourguignon and Spadaro (2012), and Jacobs et al. (2017).[3] Our general tariff formula also builds on the type of necessary first-order condition that is common in optimal commodity taxation, e.g. Diamond and Mirrlees (1971) and Greenwald and Stiglitz (1986). The structure of our formula is most closely related to Costinot and Werning (2023) who also express the optimal tariff on each good as a function of the marginal impact of its imports on a few key sufficient statistics. This provides a Pigouvian perspective on the determinants of trade protection. Pigouvian taxation calls for taxes on any economic activity whose effect on social welfare is not internalized by those directly involved in that activity. We apply this general principle to importing activities, regardless of whether what fails to be internalized is an increase in pollution (as in Markusen, 1975, Kortum and Weisbach, 2021 and Hsiao, 2026), psychosocial costs (as in Grossman and Helpman, 2021), an aggravation of output distortions under imperfect competition (as in Helpman and Krugman, 1989), a change in international prices leading to deteriorated terms of trade (as in de V. Graaff, 1949-1950, Grossman and Helpman, 1995, and Bagwell and Staiger, 1999), or, as is the focus of our empirical analysis, a change in domestic prices redistributing income away from individuals with higher social marginal utility.

Within the trade literature, others have attempted to estimate the welfare weights that governments may assign to different sectors, including Gardner (1987), Tyers (1990), Francois and Nelson (2014), and Ossa (2014), and to a lesser extent different regions, as in Ma and McLaren (2018) and Gawande et al. (2024). Previous work typically focuses on environments that are partial equilibrium or feature a small number of tariffs. Ossa (2014), for instance, combines a multi-sector gravity model with data on the average tariffs for thirty-three industries in order to identify the thirty-three welfare weights that they each receive, up to normalization. Our paper offers a "shovel-ready" way to estimate welfare weights in general economic environments and without having to restrict the granularity of observed trade policy. This is critical to identify welfare weights that may simultaneously vary across sectors and regions of employment.

A striking feature of previous work on the political economy of trade policy is the limited extent to which it draws on advances in the empirical estimation of causal responses of labor market outcomes to trade policy, as reviewed in Goldberg and Pavcnik (2016). Despite the fact that the existence of heterogeneous causal impacts of changes in imports on earnings is the primary rationale for trade protection in the political economy literature, modern understanding of such impacts is not actually used when attempting to infer the reasons for protectionism. Theoretical tractability rather than empirical credibility drives the way earnings are implicitly assumed to respond to trade protection. In contrast, we infer these responses using a model

[3] All three papers focus on income taxes to infer welfare weights at different quantiles of the income distribution. In related work, Fajgelbaum et al. (2023) use proposals for California's High-Speed Rail system to infer policymakers' preferences for location-based redistribution.

designed to harness the substantial heterogeneity in exposure to trade across the US population, and, importantly, we validate our model's predictions against the estimated causal response of US labor markets to the US-China trade war.

2. A GENERAL TARIFF FORMULA

The goal of this section is to characterize the structure of Pareto efficient trade taxes. We do so via a general tariff formula that features three generic motives for trade policy: (i) redistribution, which will be the main focus of our empirical analysis; (ii) terms-of-trade manipulation, which will be controlled for in our regressions; and (iii) distortions, which will be treated as a structural residual.

2.1. *Environment*

We focus on a single country, Home, that can trade with the rest of the world subject to its preferred trade taxes. Home comprises many firms $f \in \mathcal{F}$ and individuals $n \in \mathcal{N}$. Firms and individuals can produce and consume goods $g \in \mathcal{G}$. Goods encompass final goods and intermediate inputs, as well as labor and other primary factors. Both production and consumption may be subject to externalities $z \equiv \{z_k\}$ to be described further below.

Domestic Technology. Firm f's technology is described by a production set $\Upsilon(z; f)$. A production plan consists of a net output vector $y(f) \equiv \{y_g(f)\}$. It is feasible if

$$y(f) \in \Upsilon(z; f).$$

Domestic Preferences. A consumption plan for individual n consists of a vector of goods demanded $c(n) \equiv \{c_g(n)\}$. It delivers utility

$$u(n) = u(c(n), z; n).$$

Prices, Taxes, and Transfers. International transactions are subject to specific trade taxes $t \equiv \{t_g\} \in \mathcal{T}$. Trade taxes create a wedge between the prices $p \equiv \{p_g\}$ faced by domestic firms and individuals and the prices $p^w \equiv \{p_g^w\}$ in the rest of the world. For any good g that is traded between Home and the rest of the world,

$$p_g = p_g^w + t_g. \tag{2}$$

If a good g is imported, $t_g \geq 0$ corresponds to an import tariff, while $t_g \leq 0$ corresponds to an import subsidy. If good g is exported, $t_g \geq 0$ corresponds to an export subsidy, while $t_g \leq 0$ corresponds to an export tax. Trade taxes on a given good are either unrestricted, $t_g \in \mathbb{R}$, or restricted to be zero, $t_g \in \{0\}$. For instance, Home's government may be unable to tax imports of services, for technological reasons, or prohibited from imposing export taxes, for constitutional reasons. We let $\mathcal{G}^T$ denote the set of goods that can be taxed.[4] Tax revenues are rebated to domestic individuals through a uniform lump-sum transfer τ.

[4]Although the choice of numeraire never appears explicitly in our analysis, the numeraire good, whose trade tax can be normalized to zero, is always implicitly excluded from $\mathcal{G}^T$. This convention explains why indeterminacy of trade taxes due to Lerner symmetry plays no role in Proposition 1 below.

Foreign Offer Curve. We summarize trade with the rest of the world by an offer curve $\Omega(p^w, z)$. For given foreign prices p^w, it describes the vector of Home's net imports $m \equiv \{m_g\}$ that the rest of the world is willing to export. A vector of net imports is feasible if

$$m \in \Omega(p^w, z). \tag{3}$$

Externalities. For a given domestic allocation $\{y(f), c(n)\}$, a vector of net imports m, and a vector of domestic and foreign prices (p, p^w), the vector of externalities satisfies

$$z \in \mathcal{Z}(\{y(f), c(n)\}, m, p, p^w). \tag{4}$$

This accommodates financial frictions and knowledge spillovers that affect firms' production sets $\Upsilon(z; f)$, carbon emissions that may affect both firms' technologies and individuals' utilities $u(c(n), z; n)$, as well as psychosocial costs that may only affect the latter, as in recent models of identity politics.

2.2. *Competitive Equilibrium*

Profit Maximization. Each firm f chooses $y(f)$ to solve

$$\max_{y \in \Upsilon(z;f)} p \cdot y, \tag{5}$$

where the dot product $\cdot$ refers to the inner product, $p \cdot y = \sum_g p_g y_g$. We let $\pi(p, z; f)$ denote the value function associated with (5), i.e., the profits of firm f, expressed as a function of the domestic prices p and the externalities z.

Utility Maximization. Each individual n chooses $c(n)$ to solve

$$\max_c u(c, z; n) \tag{6}$$
$$\text{subject to: } p \cdot c = \pi \cdot \phi(n) + \tau,$$

where $\pi \equiv \{\pi(p, z; f)\}$ is the vector of firms' profits and $\phi(n) \equiv \{\phi(f, n)\}$ is the vector of firms' shares held by individual n, with $\sum_{n \in \mathcal{N}} \phi(f, n) = 1$ for all $f \in \mathcal{F}$. Endowments of goods or factors by individual n correspond to her fully owning simple firms with production sets given by a singleton, as will be the case in the next section. Below we let $y(n) \equiv \{\sum_{f \in \mathcal{F}} y_g(f)\phi(f, n)\}$ denote the vector of output associated with individual n, $\mu(n)$ denote the Lagrange multiplier associated with her budget constraint, and $e(p, z, u; n) \equiv \min_c\{p \cdot c | u(c, z; n) \geq u\}$ denote her expenditure function.

Market Clearing and Government's Budget Balance. Total demand by domestic individuals equals total supply by domestic firms and total exports from the rest of the world,

$$\sum_{n \in \mathcal{N}} c(n) = \sum_{f \in \mathcal{F}} y(f) + m. \tag{7}$$

Finally, the budget constraint of the domestic government is

$$t \cdot m = N\tau, \tag{8}$$

where N is the total number of individuals at Home.

Competitive Equilibrium. We are now ready to define a competitive equilibrium.

DEFINITION 1: *A competitive equilibrium with trade taxes $t \in \mathcal{T}$ is a vector of domestic and foreign prices (p, p^w), a vector of net imports m, a vector of externalities z, a domestic allocation $\{y(f), c(n)\}$, and a transfer τ such that: (i) (p, p^w) satisfy* (2)*; (ii) m satisfies* (3)*; (iii) z satisfies* (4)*; (iv) $y(f)$ solves* (5) *for all $f \in \mathcal{F}$; (v) $c(n)$ solves* (6) *for all $n \in \mathcal{N}$; (vi) all markets clear, as described in* (7)*; and (vii) the government's budget is balanced, as described in* (8)*.*

2.3. *Pareto-Efficient Trade Taxes*

It is standard in the literature on the political economy of trade policy to model explicitly various features of the political process. We propose instead to remain agnostic about these considerations and only require that trade taxes be (constrained) Pareto-efficient.

DEFINITION 2: *A vector of trade taxes t^* is Pareto-efficient if, for any individual n_0, there exists a vector of utility $\{\underline{u}(n)\}_{n \neq n_0}$ such that t^* solves*

$$\max_{t \in \mathcal{T}} \max_{\{u(n)\}} u(n_0)$$
$$\text{subject to: } u(n) \geq \underline{u}(n) \text{ for } n \neq n_0,$$
$$\{u(n)\} \in \mathcal{U}(t),$$

where $\mathcal{U}(t)$ is the set of utility profiles attainable in a competitive equilibrium with trade taxes t.

Examples. Trade taxes that satisfy Definition 2 may arise under various political systems. Mayer (1984) focuses on direct democracy. Under majority voting, he shows that tariffs in his model maximize the utility of the median voter.[5] The same median voter outcome would occur in a representative democracy in which two political parties compete for office by announcing (and committing) to a policy as in Downs (1957). If individuals themselves run for office (and lack the ability to commit to a policy), as in Besley and Coate (1997), then the median voter outcome may not arise, but the utility of whoever wins the election will continue to be maximized. More generally, electoral competition for votes between political parties with different ideologies may lead policy choices to maximize a weighted sum of individuals' utility, with higher weights given to "non-partisan" or "swing voters" whose votes are more likely to be swayed by changes in policy, as in Dixit and Londregan (1996). This is the approach followed by Ma and McLaren (2018) to analyze the role of "swing states" for US trade policy. Other potential determinants of utility weights across groups of individuals include turnout and knowledge, as discussed in Grossman and Helpman (2001). In each case, one can view the utility weights as the Lagrange multipliers associated with the utility constraints in Definition 2, with the different utility levels $\{\underline{u}(n)\}$ implicitly revealing the importance of these political forces.

The efficiency property embodied in Definition 2 may not only arise from political parties competing for votes, but also from special interests competing for influence. Grossman and Helpman's (1994) influential protection for sale model offers one such efficiency result for

[5] Mayer (1984) focuses on the case of a small open economy. Dhingra (2014) extends Mayer's (1984) optimal tariff formula to the case of a large country.

trade policy. In their model, a government that cares both about (utilitarian) welfare and political contributions ends up choosing trade taxes that maximize a weighted sum of individuals' utility, with higher weights being given to those in politically-organized sectors that are able to make such contributions.[6] Becker (1983) and Wittman (1989) offer early formalizations of the argument for efficiency.[7]

Finally, we note that Definition 2 is fully consistent with what Corden (1984) describes as a "conservative social welfare function" and views as "implicit in much policy-making." Such a function says that in response to market disturbances, e.g. a surge in Chinese imports, "one (or the) objective of policy is to prevent significant falls in real incomes of any significant sector of the economy." One can view market disturbances as shifts in the Pareto frontier and, whenever feasible, a conservative social welfare function simply calls for selecting allocations on the new frontier that Pareto dominate the original one. In this case, the utility levels $\{\underline{u}(n)\}$ entering the utility constraints in Definition 2 would be those associated with the pre-shock equilibrium.[8]

Counter-examples. If competition— whether for votes or for influence—is expected to lead to efficiency, what forces then may lead to departures from Definition 2? A key force may arise in dynamic environments if a lack of commitment by policy makers leads to inefficiencies. Acemoglu (2003) gives central stage to this consideration to explain the absence of a "Political Coase Theorem." In Acemoglu and Robinson (2001), the inability to commit to future policies is used to explain inefficient redistribution towards farmers. If future transfers depend on the number of farmers in the economy, they may favor trade taxes, which keep the agricultural sector large, to lump-sum transfers, which do not. In Coate and Morris (1995), it is also politicians' desire to influence future electoral outcomes that may lead them to choose inefficient, non-lump-sum transfers today in order to preserve their reputation and be in charge of policymaking tomorrow.

Despite these examples, it is important to note that for our purposes the key issue is not whether inferior forms of redistribution, like trade taxes, will be deployed instead of lump-sum transfers. We have already discussed in footnote 7 a number of reasons why they may be. The key issue is whether, conditional on trade taxes being used, they will be used in a way that violates Definition 2. In the case of reputational concerns, for instance, this would require not only that the choice of trade taxes versus lump-sum transfers has differential effects on voters' beliefs about the type of the politicians that they face. It would further require different trade policy choices to have differential effects on voters' beliefs, perhaps because some tariffs are more likely to reflect a legitimate second-best motive, say those targeting goods with high levels of carbon emissions.

[6]Since political contributions are real resources that are taken away from individuals in politically-organized sectors, constrained Pareto efficiency in Definition 2 should be interpreted as applying conditional on a given level of lobbying activities. In Grossman and Helpman (1994), the previous qualification has no bearing on the choice of trade taxes due to the assumption of quasi-linear preferences. It would more generally, as in Dixit et al. (1997).

[7]Left open is the question of why competition for office or influence would lead to trade taxes that are constrained Pareto-efficient rather than, even more radically, non-distortionary forms of redistribution like targeted lump-sum transfers. Although we remain mostly agnostic about this issue, two observations are worth making at this point. First, our approach need not be inconsistent with the existence of more efficient forms of redistribution. As we discuss in Section 2.4, if lump-sum transfers were indeed available, the tariff formula presented in Proposition 1 would continue to hold. Second, there are potential reasons why either lobbies or the government may want to commit to using inferior instruments for redistribution. Lobbies may prefer tariffs to lump-sum transfers because it allows them to extract a higher share of total surplus, as in Dixit et al. (1997). Benevolent governments, in turn, may prefer tariffs because lobbying activities, and in turn the overall level of distortions, may be lower for inefficient instruments, as in Rodrik (1986).

[8]Deardorff (1987) formalizes the conservative social welfare function along these lines.

The general point is that trade taxes may fail to satisfy Definition 2 in dynamic environments where choices over today's trade policy may also affect the policy implemented tomorrow. In such situations, the "welfare weights" that we recover from our empirical analysis will reflect both the social marginal utility of income of a given group, in line with Definition 2, and the shadow values of manipulating future redistribution.[9]

Characterization. Let $\nu(n)$ denote the Lagrange multiplier associated with the utility constraint of individual n, with the convention $\nu(n_0) = 1$. Hence the social marginal utility of n's income is $\lambda(n) \equiv \mu(n)\nu(n)$, the average social marginal utility of income is $\bar{\lambda} \equiv \sum_{n\in\mathcal{N}} \lambda(n)/N$, and the social marginal return of a hypothetical lump-sum transfer to individual n is $\beta(n) \equiv \lambda(n)/\bar{\lambda}$.[10] Figure 1 illustrates how the choice of Pareto-efficient trade taxes $t^* \in \mathcal{T}$ implicitly reveals those social marginal returns.

FIGURE 1.—Pareto-efficient trade taxes

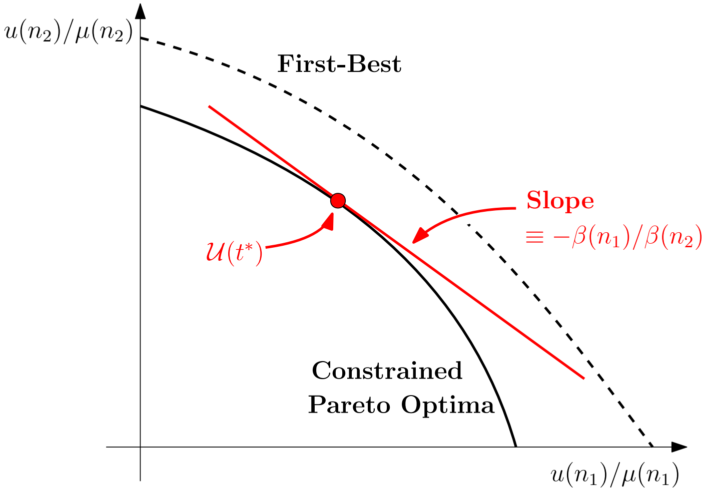

Notes: This figure plots the constrained Pareto frontier (solid line) between two individuals n_1 and n_2 that obtains as one varies the trade taxes $t \in \mathcal{T}$ applied in a competitive equilibrium. The slope of the constrained Pareto frontier at the chosen trade taxes t^* reveals the ratio of social marginal returns $\beta(n_1)/\beta(n_2)$. The first-best frontier (dashed line) is the set of Pareto optima that would arise if only technological and resource constraints applied. Due to the presence of externalities, the two frontiers may not be tangent at any point.

To characterize Pareto efficient trade taxes, it is convenient to treat equilibrium variables as functions of the vector of taxable imports $m^T \equiv \{m_g\}_{g\in\mathcal{G}^T}$ rather than the vector of trade taxes $t \in \mathcal{T}$.[11] Under this convention, the vectors of partial derivatives $\partial p/\partial m_g$, $\partial p^w/\partial m_g$,

[9]Appendix A.3 presents a two-period example that illustrates this discussion.

[10]Since all tax revenues are assumed to be rebated via a uniform lump-sum transfer, the average social marginal utility of income $\bar{\lambda}$ is also equal to the social marginal value of fiscal revenues λ_{gov}. If one were to relax the assumption of uniform lump-sum transfers and allow for more general forms of government spending and social programs, then our characterization of tariffs would be unchanged, except for the definition of $\beta(n)$. Namely, the social marginal return of a hypothetical lump-sum transfer to individual n should be defined more generally as $\beta(n) \equiv \lambda(n)/\lambda_{gov}$.

[11]Formally, if $\tilde{x}(t)$ denotes the equilibrium value of a variable x as a function t, then the function of imports $x(m)$ that we consider is defined as $x(m^T) \equiv \tilde{x}(t^{-1}(m^T))$, with $t^{-1}(m^T)$ the vector t that solves: $\tilde{m}_g(t) = m_g$ for all $g \in \mathcal{G}^T$. Throughout we assume that, local to the observed equilibrium, the inverse $t^{-1}(m^T)$ exists and is unique. We view this as a mild requirement that rules out extreme environments, such as those where preferences over net imports are Leontief and so multiple vectors of trade taxes t may be associated with the same import vector m^T.

and $\partial z/\partial m_g$ then refer to the changes in domestic prices, foreign prices, and externalities, respectively, associated with whatever change in trade taxes induces a marginal increase in the net imports of any given good $g \in \mathcal{G}^T$, holding fixed the imports of all other goods in $\mathcal{G}^T$. In particular, let $\partial\omega(n)/\partial m_g \equiv [y(n) - c(n)] \cdot (\partial p/\partial m_g)$ denote the change in individual n's real income caused by the increase in net imports of good g via its impact on domestic prices p; let $\partial\bar{\omega}/\partial m_g \equiv \sum_{n\in\mathcal{N}}[\partial\omega(n)/\partial m_g]/N$ denote its average across the population; and let $\partial(\omega - \bar{\omega})/\partial m_g \equiv \{\partial(\omega(n) - \bar{\omega})/\partial m_g\}$ denote the vector of deviations from the average.

Our main proposition shows how the previous statistic, together with the changes in foreign prices $\partial p^w/\partial m_g$ and externalities $\partial z/\partial m_g$, shapes Pareto efficient trade taxes.

PROPOSITION 1: *Pareto efficient trade taxes satisfy*

$$t_g^* = \underbrace{-\beta \cdot \frac{\partial(\omega - \bar{\omega})}{\partial m_g}}_{\textit{redistribution}} + \underbrace{m \cdot \frac{\partial p^w}{\partial m_g}}_{\textit{terms-of-trade}} + \underbrace{\epsilon \cdot \frac{\partial z}{\partial m_g}}_{\textit{distortions}} \;\; \textit{for all } g \in \mathcal{G}^T, \tag{9}$$

where $\beta \equiv \{\beta(n)\}$ denotes the social marginal returns of transfers to different individuals; and $\epsilon \equiv \sum_{n\in\mathcal{N}}\beta(n)[e_z(n) - \pi_z(n)]$ denotes the social marginal cost of externalities, with $e_z(n) \equiv \{\partial e(p, z, u(n); n)/\partial z_k\}$ and $\pi_z(n) \equiv \{\sum_{f\in\mathcal{F}}\phi(f, n)\partial\pi(p, z; f)/\partial z_k\}$.

Proposition 1 derives from the following necessary first-order condition,

$$-t^* \cdot dm = \beta \cdot d(\omega - \bar{\omega}) - m \cdot dp^w - \epsilon \cdot dz.$$

It states that at a constrained Pareto optimum, the marginal cost of any tax change in terms of fiscal revenues, $-t^* \cdot dm$, should be equal to its marginal benefit in terms of redistribution, $\beta \cdot d(\omega - \bar{\omega})$, terms-of-trade manipulation, $-m \cdot dp^w$, and correcting domestic distortions, $-\epsilon \cdot dz$. Equation (9) then specializes this condition to changes in trade taxes that only affect the imports of a single good $g \in \mathcal{G}^T$, as shown in Appendix A.1. Our general tariff formula highlights the three broad reasons why Home's government may want to tax the net imports of a given good g.

First, restricting net imports may affect real incomes via changes in domestic prices. Thus, a government may engineer as-if transfers from individuals with low social marginal return (i.e., a low $\beta(n)$) towards individuals with high social marginal return (i.e., a high $\beta(n)$). This is the redistributive motive captured by the first term, $-\beta \cdot \partial(\omega - \bar{\omega})/\partial m_g$, which will be at the core of our empirical analysis. Note that the redistributive motive is zero if $\partial(\omega - \bar{\omega})/\partial m_g = 0$, which occurs if changes in imports do not differentially affect real earnings in the population, or if $\beta(n) = 1$ for all n, which occurs if individuals have identical quasi-homothetic preferences and Home's government is utilitarian, a standard benchmark in the trade literature.

Second, restricting net imports may lower Home's import prices and increase its export prices. This is the terms-of-trade motive captured by the second term, $m \cdot \partial p^w/\partial m_g$, which we will use as a control in our main specification. As usual, this second term is zero in the case of a small open economy that may manipulate domestic prices p, but not foreign prices p^w. Note that given our change of variables, the terms-of-trade motive takes a particularly simple form in equation (9). It is akin to the classical optimal tariff formula that obtains in a two-good environment—in which the optimal tariff is equal to the inverse of the elasticity of the foreign export supply curve—despite the fact that we impose no restrictions on the number of goods (nor on preferences and technology).[12]

[12]The terms-of-trade effect being equal to $m \cdot \partial p^w/\partial m_g$ does rely on our assumption that all domestic firms are owned by domestic individuals, i.e., $\sum_{n\in\mathcal{N}}\phi(f, n) = 1$ for all $f \in \mathcal{F}$. If there was foreign ownership of domestic

Third, restricting net imports may reduce negative externalities or raise positive ones. This is the typical second-best motive for trade protection captured by the third term, $\epsilon \cdot \partial z / \partial m_g$. Again, due to our change of variables, this third motive can be expressed in an intuitive manner as the sum of the marginal change in distortionary activities caused by one extra imported unit of good g, each multiplied by the social cost of that activity.

In a competitive equilibrium, domestic individuals and firms do not internalize any of the three previous considerations. Following a general Pigouvian logic, the optimal trade tax on a given good g requires them to pay, at the margin, for the potential negative impact of that good's imports on social welfare, a perspective emphasized in Costinot and Werning (2023). This is true regardless of whether import restrictions may affect social welfare via redistribution or efficiency considerations.

2.4. *Extensions*

Below we will use Proposition 1 to estimate the role played by the redistributive motive for trade protection. Before we do so, we discuss its robustness to a number of considerations from which we have abstracted. Formal proofs can be found in Appendix A.2.

Other Policy Instruments. While the economic environment considered in Section 2.1 is general along many dimensions, it restricts the policy instruments available to the domestic government to specific trade taxes. Beginning with other trade policy instruments, it is well-known that the restriction to specific rather than ad-valorem trade taxes is without loss of generality under perfect competition. The critical assumption is that the government can create a wedge between foreign prices p^w—which affect the decision of foreigners via (3)—and domestic prices p—which affect the decisions of domestic firms and individuals via (5) and (6). The specific or ad-valorem nature of the trade tax through which the wedge comes about is irrelevant.

In practice, a government may also choose to restrict trade flows via various non-tariff measures, from product standards to anti-dumping duties. If such barriers do not generate fiscal revenues, as would be the case for product standards, then Proposition 1 is unchanged. That is, the existence of standards may affect the particular values of the sufficient statistics entering equation (9), but not the fact that equation (9) must continue to hold.[13] If instead non-tariff measures act as another trade tax, as would be the case for anti-dumping duties, then the associated fiscal externalities should be accounted for, as shown in Proposition A.1. Intuitively, (9) reflects the first-order conditions for the choice of t, holding all other policy choices constant. For the purposes of choosing t the only new consideration created by other policies is therefore the way that the fiscal revenue they generate is affected when t is adjusted.

The same observations apply to the case of domestic policy instruments. When these instruments do not generate fiscal externalities, as in the case of lump-sum transfers, even individual-specific ones, the first-order condition underlying Proposition 1 remains valid and our analysis goes through without qualification. The value of β observed in a competitive equilibrium may, of course, depend on whether these additional instruments exist or not, but their existence does

---

firms, then the terms-of-trade effect in equation (9) would need to be adjusted to $m_{FDI} \cdot \partial p^w / \partial m_g$, with $m_{FDI} \equiv m + \sum_{f \in \mathcal{F}} \phi^w(f) y(f)$ and $\phi^w(f)$ the share of firm f held by foreigners. Domestic ownership of foreign firms could be dealt with in a similar manner, provided that the economic environment also explicitly describes foreign firms and their profits.

[13] Similarly, tariffs that are conditional on the use of production techniques, such as the rules-of-origin restrictions that often appear in trade agreements (Conconi et al., 2018), can be handled by defining goods on the basis of such techniques.

not change the fact that β can be estimated using (9). In contrast, when the government uses distortionary taxes, such as producer subsidies, the fiscal externalities associated with these taxes should again be explicitly controlled for in our regressions, as shown in Proposition A.2.[14] In the presence of income taxes, one should further take into account how domestic taxes affect the ultimate incidence of transfers engineered by trade restrictions. Specifically, changes in real earnings should be adjusted by marginal income tax rates, as also established in Proposition A.2.

Constrained Trade Taxes. Our baseline analysis abstracts from WTO rules and trade agreements. Hence it treats a country's decision to abide by them as a choice that can be used to reveal social welfare weights. An alternative view is that these arrangements impose constraints on the choice of the government. Tariffs on some goods may be fixed at some exogenous level, due to prior trade agreements, or they may be constrained to be constant across subsets of goods—for instance, they may be prohibited from varying across goods from different origin countries. In the former case, the existence of non-zero, but fixed trade taxes implies another source of distortions due to fiscal externalities, as changes in the subset of trade taxes controlled by the government may now also affect the fiscal revenues generated by exogenous trade taxes on other goods, as shown in Proposition A.3. In the latter case, Proposition 1 continues to hold provided that marginal changes in imports are aggregated at the level at which trade taxes can vary, e.g. total imports of a product from all WTO countries in the case of the most-favored-nation (MFN) clause, as also shown in Proposition A.3.

Negotiated Trade Taxes. The tariff formula in Proposition 1 hinges on a strict dichotomy between domestic individuals, whose utility the domestic government takes into account when setting trade taxes, and foreigners, who are absent from the government's problem in Definition 2. In practice, various rounds of trade negotiations and bargaining may lead governments to, at least partly, internalize the impact of their preferred trade taxes on foreigners' welfare. Accordingly, it is common in the trade literature to model negotiated tariffs, such as those arising from GATT negotiations, as Pareto efficient from a world standpoint (e.g. Bagwell and Staiger, 2002).

From the point of view of Home, the only difference between the structure of Pareto efficient trade taxes that we consider and those that would arise from "trade talks" derives from the treatment of the terms-of-trade motive, $m \cdot \partial p^w / \partial m_g$ in equation (9). Under the assumption that the domestic government takes into account the aggregate real income of foreigners, the coefficient in front of $m \cdot \partial p^w / \partial m_g$ is now equal to $1 - \lambda_F / \bar{\lambda}$ instead of 1, with λ_F the social marginal utility (still from the point of view of Home's government) of foreigners' income. The general logic behind our formula is unchanged. The key observation is that Home's government now not only values redistribution towards various domestic individuals, as reflected in $-\beta \cdot \partial(\omega - \bar{\omega})/\partial m_g = \sum_{n \in \mathcal{N}} (1 - \lambda(n)/\bar{\lambda})(\partial \omega(n)/\partial m_g)$, but also redistribution towards foreigners, as reflected in $(1 - \lambda_F/\bar{\lambda})(\partial \omega_F / \partial m_g)$, with $\partial \omega_F / \partial m_g \equiv m \cdot \partial p^w / \partial m_g$ the change in foreigners' real income. Since many tariffs that we consider in our empirical analysis have been

[14]In theory, not controlling for such considerations may lead to severe omitted variable bias. Consider the case of a small open economy in which individuals have identical quasi-homothetic preferences and there are no distortions. Suppose that the domestic government has a linear social welfare function, so that $\beta > 0$ is exogenously given. If the domestic government can freely choose producer subsidies $\{s^y(f)\}$ in this environment, then it will set the fiscal externality associated with producer subsidies, $\sum_{f \in \mathcal{F}} s^y(f) \cdot \partial y(f)/\partial m_g$, to be exactly equal to the redistributive term, $\beta \cdot \partial(\omega - \bar{\omega})/\partial m_g$. According to Proposition A.2, one would therefore observe $t^* = 0$ and always infer from equation (9) (when failing to account for $\sum_{f \in \mathcal{F}} s^y(f) \cdot \partial y(f)/\partial m_g$) that $\beta = \{1\}$, regardless of whether or not this is true. We thank Bob Staiger for suggesting this example.

negotiated, we will allow the coefficient in front of the terms-of-trade motive, $m \cdot \partial p^w/\partial m_g$, to differ from 1 in our baseline regressions. More generally, if Home places different weights on different foreign countries (perhaps due to preferential trade agreements) as well as on different individuals n from these countries (perhaps due to political forces in i influencing its trade negotiations with Home, as in Grossman and Helpman, 1995), this can be allowed for by controlling more flexibly for $m(i,n) \cdot \partial p^w/\partial m_g$, with $m(i,n) \equiv c(i,n) - y(i,n)$ the net imports of individual n in country i, as shown in Proposition A.4.[15]

Other Distortions. We conclude by noting that Proposition 1 assumes that the only source of distortions is externalities in an otherwise perfectly competitive environment. This abstracts from firm-delocation effects, as in Venables (1987) and Ossa (2011), or profit-shifting effects, as in Brander and Spencer (1984) and Mrazova (2023). In Proposition A.5, we show how general distortions due to imperfect competition may be incorporated in our tax formula. Since this formula reflects a necessary first-order condition, this new source of distortions enters additively, as the extra social cost of firms' output distortions. In the case where firms only produce a single good and the social marginal utility of income is equalized across individuals, this is equal to the change in the final output of the firm multiplied by the difference between its price and marginal cost, as is standard in the literature on misallocation. For our purposes, it is enough to note that such extra considerations would appear as part of our structural residual and would only matter if they are systematically correlated with changes in real earnings $\partial(\omega - \bar{\omega})/\partial m_g$.

3. MEASURING THE SENSITIVITY OF REAL EARNINGS TO IMPORTS

The goal of our empirical analysis is to use equation (9) to go from observed US trade taxes t to the welfare weights β and, in turn, to explore the importance and nature of redistributive trade protection. Doing so requires measures of the sensitivity of real earnings to the imports of any given good g, $\partial(\omega - \bar{\omega})/\partial m_g$, holding constant the imports of all other goods. Direct estimation without a priori restrictions would require estimating as many derivatives $\partial(\omega(n) - \bar{\omega})/\partial m_g$ as there are individual-good pairs (n, g) in the US, which is infeasible. To arrive at such estimates, we therefore build a model of the US economy using a specific version of the general environment from Section 2. We then show that this model replicates the causal impact of observed tariff shocks on relative earnings, which raises confidence that it provides an accurate measure of $\partial(\omega - \bar{\omega})/\partial m_g$.

3.1. *A Model of the US Economy*

The specific environment that we rely on for the rest of our analysis is a variation of the model in Fajgelbaum et al. (2020) (FGKK), which we will calibrate using data from 2017.

Regions, Sectors, Products, and Trade Partners. A domestic individual n may live in one of many regions $r \in \mathcal{R}_H$ and work in one of many sectors $s \in \mathcal{S}$. Given data availability, we take the set of domestic regions $\mathcal{R}_H$ to be the 50 US states, plus the District of Columbia, and the set of sectors $\mathcal{S}$ to be 21 tradable industries based on 3-digit NAICS industries, plus one services sector that is not internationally traded. To these 51×22 groups of individuals, we add

[15]In the case of a representative agent in each foreign country i, and hence no concern for foreign redistribution, the controls reduce to $m(i) \cdot \partial p^w/\partial m_g$, with $m(i)$ the net imports of country i.

a "residual" individual whose pattern of (net) expenditure will allow our model to match data from 2017 and whose behavior we hold fixed throughout our analysis.[16]

We let N_{rs} denote the fixed number of individuals living in region r and working in sector s.[17] Each individual is endowed with ϕ_{rs} units of equipped labor, which she sells to firms f in that region and sector at a wage w_{rs}.[18] Firms hire labor and buy intermediate goods from other domestic firms and foreigners in order to produce differentiated products $h \in \mathcal{H}_s$, which they sell to foreigners, other domestic firms, and individuals. The set of all products $\mathcal{H} \equiv \cup_{s\in\mathcal{S}} H_s$ is based on the 6-digit HS system, resulting in 5,299 products with positive trade in 2017.

Foreigners may be located in one of many countries $i \in \mathcal{R}_F$. We take the set of foreign countries $\mathcal{R}_F$ to be the top 100 US trade partners, plus the rest of the world treated as a single country; the top 100 partners account for 99.0% of US exports and 99.6% of its imports. A good g in the general notation of Section 2 either corresponds to labor from a given region-sector pair (r, s) or to an origin-destination-product triplet (o, d, h), where each origin o and destination d is either a domestic region r or a foreign country i. We let $\mathcal{R} \equiv \mathcal{R}_H \cup \mathcal{R}_F$ denote the set of all locations.

Trade Taxes. In terms of policy instruments, we assume that there are no export taxes or subsidies. The only US trade taxes are specific import tariffs t_{ih} that may vary across foreign origins $i \in \mathcal{R}_F$ and products $h \in \mathcal{H}$.[19] All tariff revenues are rebated uniformly across individuals. Note that since a tradable good g is an origin-destination-product triplet, trade taxes are constrained to be equal across all domestic destinations, i.e., different US regions cannot impose different tariffs. As discussed in Section 2.4, our general formula still holds in this case provided that marginal changes in imports now refer to total changes in Home's imports of product h from country i, $m_{ih} \equiv \sum_{r\in\mathcal{R}_H} m_{irh}$, where m_{irh} denotes bilateral imports to each region $r \in \mathcal{R}_H$, as described in equation A.14

Sensitivity of Real Earnings to Imports. For any individual n endowed with ϕ_{rs} units of labor from region r and sector s, the change in real earnings associated with an increase in imports of product h from a foreign country i reduces to

$$\frac{\partial \omega(n)}{\partial m_{ih}} = \phi_{rs} \frac{\partial w_{rs}}{\partial m_{ih}} - \sum_{o\in\mathcal{R}, v\in\mathcal{H}} c_{orv}(n) \frac{\partial p_{orv}}{\partial m_{ih}}, \tag{10}$$

where $c_{orv}(n)$ is the consumption of product v from origin o by individual n from region r in the competitive equilibrium with trade taxes t and p_{orv} denotes the domestic price of that

[16] We also assume that the residual individual's social marginal utility of income is equal to the US average. Hence, its existence does not create any further motive for redistribution. Like exogenous lump-sum transfers in quantitative trade and spatial models, the only role of the residual individual is to rationalize imbalances in the data (e.g. between countries or states). This specific modeling choice has little effect on the welfare weights that we estimate, as discussed further in our working paper (Adao et al., 2023).

[17] Following FGKK, this specification does not allow for mobility across sectors and regions. As such, it should be thought of an approximation for the short-run impact of tariffs on earnings.

[18] As already discussed in Section 2.1, factor endowments can be interpreted in the context of our general environment as full ownership of simple firms that produce factor services without any additional inputs. Under perfect competition and constant returns to scale, as we assume below, factor endowments will be individuals' only source of income other than government transfers.

[19] In practice, the vast majority of import tariffs imposed by the United States are ad-valorem rather than specific. As already discussed in Section 2, there is no loss of generality in focusing on an environment where all import tariffs are assumed to be specific instead. This is because one can always go from an observed equilibrium with ad-valorem tariffs to an equivalent one with specific tariffs by letting the latter be equal to the ad-valorem ones times the price of US imports (pre-tariff) in the observed equilibrium, without any further consequences for our analysis.

good. The first term on the right-hand side, $\phi_{rs}(\partial w_{rs}/\partial m_{ih})$, is the change in the individual's earnings, whereas the second term, $\sum_{o\in\mathcal{R},v\in\mathcal{H}} c_{orv}(n)(\partial p_{orv}/\partial m_{ih})$, is the change in her cost of living. Next, we describe the parametric restrictions that we impose on domestic technology, domestic preferences, and foreign offer curves to compute the sensitivity of real earnings to imports in (10), with further details about analytical derivations relegated to the Online Appendix.

3.2. *Parametric Restrictions*

Domestic Technology. For each region $r\in\mathcal{R}_H$, destination $d\in\mathcal{R}$, and product $h\in\mathcal{H}_s$ from sector $s\in\mathcal{S}$, there is a representative firm f whose gross output $q(f)$ is

$$q(f)=\theta_{rds}[\ell_{rs}(f)]^{\alpha_s}\prod_{k\in\mathcal{S}}[Q_{rk}(f)]^{\alpha_{ks}}, \tag{11}$$

$$Q_{rk}(f)=\left[\sum_{c=H,F}\sum_{o\in\mathcal{R}_c}\sum_{v\in\mathcal{H}_k}(\theta^c_{orkv})^{\frac{1}{\sigma}}[q_{orv}(f)]^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{12}$$

where $\ell_{rs}(f)$ denotes labor from region r and sector s used by firm f, $q_{orv}(f)$ denotes its use of intermediate inputs of product v from origin o delivered to region r, and $\sigma\geq 0$ denotes the elasticity of substitution between different inputs.[20] Since we only observe product-level trade flows between domestic regions and foreign countries, but not between pairs of domestic regions, we impose $\theta^H_{orkv}=\bar{\theta}^H_{ork}$. Finally, we normalize input demand shifters so that $\alpha_s+\sum_{k\in\mathcal{K}}\alpha_{ks}=\sum_{c=H,F}\sum_{o\in\mathcal{R}_c}\sum_{v\in\mathcal{H}_k}\theta^c_{orkv}=1$. Note that trade costs, of the standard iceberg form, are implicitly embedded in demand shifters. For instance, if a product v from sector k is non-tradable from an origin o to region r, then $\theta^c_{orkv}=0$.

Domestic Preferences. In each region $r\in\mathcal{R}_H$, the utility $U(n)$ of any individual n is

$$U(n)=E(z,n)\prod_{s\in\mathcal{S}}[C_{rs}(n)]^{\gamma_s}, \tag{13}$$

$$C_{rs}(n)=\left[\sum_{c=H,F}\sum_{o\in\mathcal{R}_c}\sum_{h\in\mathcal{H}_s}(\theta^c_{orsh})^{\frac{1}{\sigma}}[c_{orh}(n)]^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{14}$$

where $E(z,n)$ denotes the impact of externalities on the utility of individual n. Except for the Cobb-Douglas parameters $\{\gamma_s\}$ that may differ from $\{\alpha_{ks}\}$ in equation (11), note that all other demand shifters as well as elasticities in equation (14) are the same as in equation (12). That is, both domestic firms and individuals demand the same "sector composite," a standard restriction in quantitative trade models that is driven by data availability (or lack thereof). In line with our treatment of technology, we impose the normalization $\sum_{s\in\mathcal{S}}\gamma_s=1$. Lastly, we let ξ_{rs} denote the residual individual's net spending on all products from a sector s and domestic region $r\in\mathcal{R}_H$, which we treat as an exogenous preference shifter.

[20] In terms of the general notation of Section 2, the associated vector of net output $y(f)$ is obtained by entering gross output with a positive sign for good $g=(r,d,h)$ and entering all inputs with a negative sign. This vector is then feasible, $y(f)\in\Upsilon(z;f)$, if (11) and (12) hold. Note that there are no externalities in production in our model, a point we return to below.

Foreign Offer Curve. For each foreign country $i \in \mathcal{R}_F$, domestic region $r \in \mathcal{R}_H$, and product $h \in \mathcal{H}$, i's gross exports $q_{irh}^{X,F}$ and gross imports $q_{rih}^{M,F}$ satisfy

$$p_{irh}^{X,F} = \theta_{irh}^{X,F} (q_{irh}^{X,F})^{\psi^{X,F}}, \tag{15}$$

$$p_{rih}^{M,F} = \theta_{rih}^{M,F} (q_{rih}^{M,F})^{-\psi^{M,F}}, \tag{16}$$

where $p_{irh}^{X,F}$ is the price received by foreign sellers of product h in country i that serves region r and $p_{rih}^{M,F}$ is the price paid by foreign buyers of product h from region r in country i. Since products are differentiated by origin, foreign gross exports are also equal to domestic net imports of these goods, $q_{irh}^{X,F} = m_{irh}$, whereas foreign gross imports are equal to minus domestic net imports, $q_{rih}^{M,F} = -m_{rih}$.[21] The first elasticity $\psi^{X,F} \geq 0$ denotes the inverse of foreigners' export supply elasticity, whereas the second $\psi^{M,F} \geq 0$ denotes the inverse of their import demand elasticity. Provided that either of these two elasticities is different from zero, then Home may affect world prices $p^w \equiv \{p_{irh}^{X,F}, p_{rih}^{M,F}\}$.

Externalities. Externalities only affect the utility of US individuals, and they do so leaving individuals' marginal rates of substitution unchanged, as can be seen from the impact of $E(z,n)$ in (13). This implies that the only role of externalities in the rest of our analysis will be to provide a rationale for, and interpretation of, the structural residual in our regressions. Accordingly, we do not impose further restrictions on the externalities included in the vector z nor on their determinants in equation (4).

3.3. *Baseline Calibration*

The last piece of information needed to measure the sensitivity of real earnings to imports in equation (10) consists of the values of the structural parameters that determine the competitive equilibrium of our model in 2017. These parameters comprise: the three elasticities, $\{\sigma, \psi^{X,F}, \psi^{M,F}\}$; the technology shifters, preference shifters, and labor endowments, $\{\alpha_s, \alpha_{ks}, \gamma_s, \xi_{rs}, \theta_{rds}, \theta_{orsh}^c, \theta_{irh}^{X,F}, \theta_{rih}^{M,F}, \phi_{rs}, N_{rs}\}$; and the US import tariffs, $\{t_{ih}\}$. We now describe how we calibrate each of them.

Elasticities. We set the values of the three elasticities $\{\sigma, \psi^{X,F}, \psi^{M,F}\}$ equal to FGKK's estimates, since their estimating equations for each of these parameters are consistent with the parametric assumptions imposed in Section 3.2. FGKK's empirical strategy uses the US-China trade war in 2018-2019 as an exogenous source of variation that allows identification of: the elasticity of US import demand σ, by using the changes in US tariffs as a shifter of the prices paid by US importers; the elasticity of foreign export supply $\psi^{X,F}$, by using the change in US tariffs as a shifter of the prices received by foreign exporters; and the elasticity of foreign import demand $\psi^{M,F}$, by using the foreign retaliatory tariffs as a shifter of the prices paid by foreign importers. The point estimates are $\sigma = 2.53$, $\psi^{X,F} = 0.00$, and $\psi^{M,F} = 0.96$.

Technology Shifters, Preference Shifters, and Labor Endowments. We set the values of $\{\alpha_s, \alpha_{ks}, \gamma_s, \xi_{rs}, \theta_{rds}, \theta_{orsh}^c, \theta_{irh}^{X,F}, \theta_{rih}^{M,F}, \phi_{rs}, N_{rs}\}$ to match US data from 2017 on: value added and employment by US region and sector; domestic trade flows by US region and sector; and international trade flows by US region, foreign country, and product. In our baseline calibration, we further normalize all prices and wages to one. This amounts to a

[21]The vector of net imports $m = \{q_{irh}^{X,F}, -q_{rih}^{M,F}\}$ is feasible, $m \in \Omega(p^w, z)$, if equations (15) and (16) hold.

choice of units of account for goods and factors that ultimately pins down the levels of $\{\theta_{rds}, \theta^{c}_{orsh}, \theta^{X,F}_{irh}, \theta^{M,F}_{rih}, \phi_{rs}\}$, without further implications for the rest of our analysis. We briefly describe the various data sources used in this procedure below and offer further details about data construction and calibration in the Online Appendix.

Value-added and employment by US region and sector. We combine the BEA's national and regional accounts to obtain value-added at the region-sector level. We begin with nationwide data on value-added by sector, which are available from the BEA's make-use tables (Bureau of Economic Analysis, 2019) at the (equivalent of) 3-digit NAICS level. Within each sector, we then assign these national value-added amounts to each region in proportion to its share of sectoral value-added in the BEA regional accounts (Bureau of Economic Analysis, 2023).[22] We directly obtain total employment by region-sector from the BEA regional accounts.

International trade flows by US region, foreign country, and product. We obtain foreign imports and exports of products by US region from the US Census (U.S. Census Bureau, 2020a,b). These flows by foreign country are available at the 6-digit HS level, which we concord to our sector classification. Due to lack of data, we assume that the services sector has zero international trade flows. For each of the 21 tradable sectors, we rescale regional trade flows to match aggregate imports and exports in the national accounts.[23]

Domestic trade flows by US region and sector. To measure the value of flows from any domestic region-sector to any other, we first compute domestic sales as gross output minus exports for each producing region and sector. We then use Commodity Flow Survey data (U.S. Census Bureau, 2019) to apportion domestic sales across purchasing regions and use national input-output tables to apportion regional purchases across different sectors as well as final demand.

US Import Tariffs. We use USITC data (U.S. International Trade Commission, 2018) to calculate the applied ad valorem tariff charged by the United States on each 6-digit HS product h from each foreign country i in 2017. We take the ratio of calculated duties to the FOB import value, which we denote t^{av}_{ih}, as the ad valorem tariff for a given product-country pair.[24] Under our price normalization, specific and ad-valorem tariffs coincide: $t_{ih} = t^{\text{av}}_{ih}$. For a given product h, t_{ih} may differ across origin countries due to country i being part of a preferential or regional trade agreement with the United States—e.g. the Generalized System of Preferences or NAFTA—and non-MFN ("column two") treatment of non-WTO members.

3.4. *Model-Implied Sensitivity of Real Earnings to Imports*

Given the parametric restrictions from Section 3.2 and the calibration from Section 3.3, we use equation (10) to compute the changes in real earnings associated with imports. For each

[22] It is worth emphasizing that we do not decompose value-added into labor and capital payments, hence the reference to "equipped labor" in the model of Section 3.1. Throughout our analysis, earnings of individuals from a given region and sector always refers to total value added, not just the wage bill.

[23] Data availability also explains why we assume that each US region is a state rather than a smaller spatial unit, like a US county. To the best of our knowledge, there is no data on foreign imports and exports of products by US county.

[24] In practice, US tariffs may vary at the HS-10 level. As mentioned above, the US Census, however, only records state-level exports and imports at the 6-digit level, hence our choice to aggregate tariffs at the same level. Allowing tariffs to vary at the 6-digit level still makes our analysis much more granular than previous ones (e.g. Goldberg and Maggi (1999) and Ossa (2014) consider 107 and 57 sectors, respectively).

FIGURE 2.—Sensitivity of real earnings to imports

(a) Across sectors, $\partial(\omega_s - \bar{\omega})/\partial m_k$ (b) Across regions, $\partial(\omega_r - \bar{\omega})/\partial m_k$

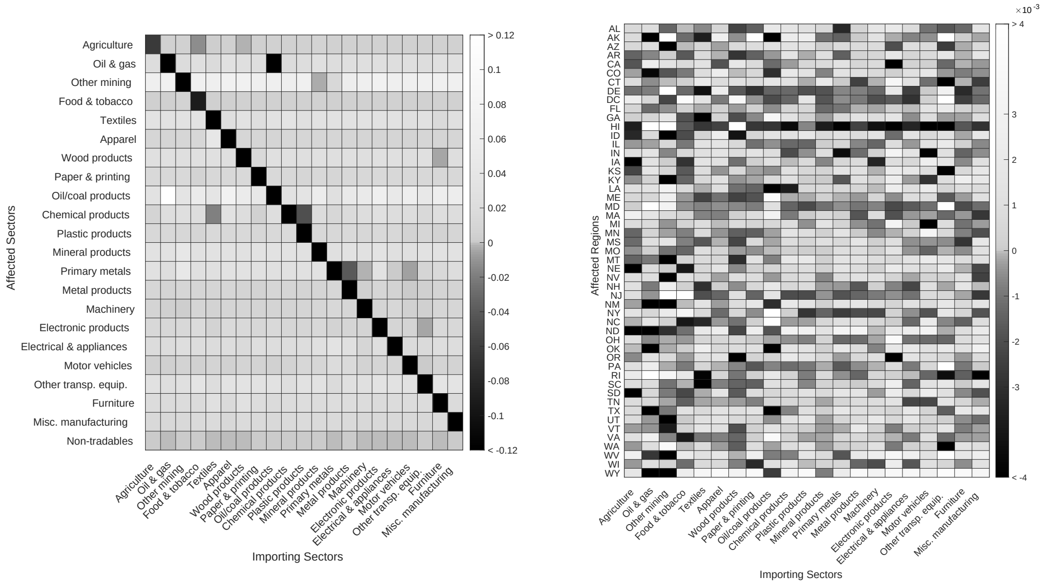

Notes: Figure 2a plots estimates of how a marginal change in imports of goods in the sector shown on the x-axis affects the difference between the average real earnings of individuals employed in each of the sectors shown on the y-axis and the average real earnings of all US individuals. Figure 2b plots the same estimates for the average real earnings of individuals living in each of the regions shown on the y-axis. Import units are chosen so that each cell reports the 2017 dollar average change in real earnings associated with a one million 2017 dollar increase in US import values. All values reported are de-meaned within each row.

region-sector (r,s) and each country-product (i,h), we let $\partial\omega_{rs}/\partial m_{ih}$ denote the change in real earnings associated with imports of product h from country i for all individuals living in region r and working in sector s. The resulting Jacobian matrix $\{\partial\omega_{rs}/\partial m_{ih}\}$ has $51 \times 22 = 1{,}122$ rows, one for each region-sector pair $(r,s) \in \mathcal{R}_H \times \mathcal{S}$, and $5{,}299 \times 101 = 535{,}199$ columns, one for each country-product pair $(i,h) \in \mathcal{R}_F \times \mathcal{H}$.

We will use the entries of the Jacobian matrix $\{\partial\omega_{rs}/\partial m_{ih}\}$ to construct the right-hand side variables in our regressions. To help visualize some of the variation that will allow us to identify welfare weights in Section 4, we report in Figure 2 the average impact of import restrictions from different tradable sectors k on the real earnings of workers from different sectors s or different regions r relative to the US average, which we denote $\partial(\omega_s - \bar{\omega})/\partial m_k$ and $\partial(\omega_r - \bar{\omega})/\partial m_k$, respectively. Changes in imports are measured in millions of dollars, whereas changes in real earnings per worker are measured in dollars.

Figure 2a focuses on sectors. For each cell (s,k), it displays $\partial(\omega_s - \bar{\omega})/\partial m_k$, further de-meaned by the average impact $\partial(\omega_s - \bar{\omega})/\partial m$ across import restrictions from all sectors k to reflect the fact that our regressions will include a constant.[25] Dark shades, which indicate negative entries, are primarily on display when $s = k$. Light shades, which indicate that individuals

[25]When averaging across imported products from any given sector k, we restrict ourselves to country-pair products (i,h) whose value of US imports in 2017 is greater than \$100,000, which leaves us with a total of 71,688 country-product pairs across all sectors. These are the same country-product pairs that we use in the empirical analysis of Section 4.

gain from imports, can be seen when $s \neq k$. The dark diagonal is the result of the natural force of protection: individuals tend to gain less from an increase in imports in their own sectors, relative to the US average, since the firms employing them also have to compete directly against foreign goods. For individuals employed in the non-tradable sector (in the bottom row), we see little difference across imports from different sectors.

Figure 2b turns to analogous effects across states. Each cell (r, k) now reports $\partial(\omega_r - \bar{\omega})/\partial m_k$, again de-meaned by $\partial(\omega_r - \bar{\omega})/\partial m$. The range of entries is an order of magnitude smaller than that for sectors in Figure 2a, indicating that as we vary import restrictions from different sectors k, there is typically less dispersion in the changes in real earnings experienced by workers from a given region r than by workers from a given sector s. This reflects the fact that most workers in all regions are employed in the non-tradable sector as well as the fact that diversification of employment across tradable sectors within each region partly smoothens the heterogeneity documented in Figure 2a.

The previous observations notwithstanding, it is important to note that Figure 2 abstracts from substantial heterogeneity in the responses to the changes in the imports of different origin-product pairs (i, h) that belong to the same sector k. This heterogeneity, which we display in Appendix Figure D.1, implies that the cross-regional variation in $\partial(\omega_r - \bar{\omega})/\partial m_{ih}$ captures more than differences in sectoral composition because, for example, different regions import different products from different foreign origins.[26] We will leverage this variation to identify separately welfare weights across sectors and regions below.

3.5. *Validating Model-Implied Sensitivity of Real Earnings to Imports*

Before turning to the identification of welfare weights by combining our tariff formula with the values of $\{\partial(\omega_{rs} - \bar{\omega})/\partial m_{ih}\}$ implied by our model, we propose to validate our model's predictions. Although one cannot directly estimate the Jacobian matrix $\{\partial(\omega_{rs} - \bar{\omega})/\partial m_{ih}\}$—which would amount to separately identifying $1,173 \times 535,199$ local causal effects—one can focus on a subset of exogenous changes in imports that have been observed in the data and ask whether, for these changes, the causal responses of earnings predicted by our model are "close" to observed ones.

Given our objective to identify US welfare weights in 2017, the ideal experiment would focus on plausibly exogenous changes in US import tariffs around that time. As a proxy for such an experiment, we return to the tariff changes implemented between 2017 and 2019. The logic is that these tariff changes derive from exogenous changes in the welfare weights of US policymakers rather than changes to the economic fundamentals of the US economy. Formally, we estimate the following linear regression via OLS,

$$\Delta \log w_{rs}^{\text{obs.}} = \alpha \Delta \log w_{rs}^{\text{pred.}} + \text{Controls}_{rs} + \varepsilon_{rs}, \tag{17}$$

where $\Delta \log w_{rs}^{\text{obs.}}$ is the log-change in earnings per worker observed in region r and sector s between 2017 and 2019, measured as value-added per worker in the BEA regional accounts, and $\Delta \log w_{rs}^{\text{pred.}} \equiv \sum_{i,h} (\partial \log \tilde{w}_{rs}/\partial t_{ih}^{\text{av}}) \Delta t_{ih}^{\text{av}}$ is the predicted response of the same outcome to the changes in US import tariffs $\Delta t_{ih}^{\text{av}}$ observed over this period, computed from our model up to a first-order approximation. Following the logic of Borusyak et al. (2022), we include

[26] This general pattern is illustrated in Appendix Figure D.2. For each region r, it reports the R^2 of a regression of $\partial(\omega_r - \bar{\omega})/\partial m_{ih}$ on the set of variables $\{\partial(\omega_s - \bar{\omega})/\partial m_{ih}\}_{s \in \mathcal{S}}$. For most regions, we find a low R^2, consistent with our region- and sector-level measures of real earnings sensitivity to imports capturing different sources of variation across origin-product pairs (i, h).

TABLE I

RESPONSES OF EARNINGS PER WORKER TO THE 2017-2019 CHANGES IN US TARIFFS: A TEST

	Dependent variable: $\Delta \log w_{rs}^{\text{obs.}}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log w_{rs}^{\text{pred.}}$	0.982	0.982	1.048	0.934	0.931	1.338
Std. error	(0.456)	(0.456)	(0.456)	(0.515)	(0.520)	(0.313)

Notes: OLS regressions of the 2017-19 change in log observed value-added per worker, across region-sector pairs, on the model-predicted equivalent arising from the 2017-19 change in US import tariffs. The sample consists of 1,000 region-sector pairs with positive value-added per worker in 2017 and 2019. All specifications include a constant and control for the overall exposure of each region-sector pair to US tariff changes. Column (2) includes controls for model-predicted response to foreign retaliatory tariffs. Column (3) includes controls for an indicator for the sector being tradable. Column (4) includes controls for log employment and log value-added in 2016. Column (5) includes all the controls from columns (1) through (4) simultaneously. And column (6) reports the specification in column (1) but weighted by 2016 employment. Standard errors in parentheses computed with ACD's version of inference for shift-share specifications.

in Controls$_{rs}$ a constant as well as the overall exposure of each region-sector pair to US tariff changes, measured by $\sum_{i,h} \partial \log \tilde{w}_{rs} / \partial t_{ih}^{\text{av}}$.

The previous specification corresponds to a special case of the "IV-based test" in Adao et al. (2025) (ACD) when the model predictions and instrumental variables (IV) coincide. Under the joint assumption that our model's predictions are correct, and that tariff changes have a common mean and are independent of the other non-tariff shocks that enter the error term ε_{rs} in (17), the regression coefficient α should be one in expectation.[27] We use standard errors computed via the procedure in ACD in order to test whether estimates of α are statistically different from this value, implying that the null of no misspecification should be rejected.

Table I reports our estimates. Column (1) presents the simplest form of our test. The point estimate of $\widehat{\alpha} = 0.98$ is close to one, which implies, alongside the standard error of 0.46, that we do not reject the null at standard levels ($p = 0.97$). The next four columns of this table vary the set of variables included in Controls$_{rs}$. Column (2) adds the predicted effect of foreign retaliatory tariffs $\sum_{i,h} (\partial \log \tilde{w}_{rs} / \partial t_{ih}^{F,\text{av}}) \Delta t_{ih}^{F,\text{av}}$, again up to a first-order approximation, with $\Delta t_{ih}^{F,\text{av}}$ the observed change in the ad-valorem tariff imposed by a foreign country i on US exports of product h. The stability of $\widehat{\alpha}$ relative to column (1) shows that the effects of such retaliatory tariffs are orthogonal to those of US tariffs, our primary focus. Column (3) adds to Controls$_{rs}$ an indicator variable for whether sector s is tradable or not, with the reported coefficient indicating that our estimates are not being driven by the differential change between tradable or non-tradable segments of the economy. Column (4) adds controls for initial size (log employment and log value-added) and column (5) adds all of the controls used in columns (1) through (4) simultaneously. All of these estimates demonstrate that the results of our test are robust to allowing for asymmetric shocks to economic fundamentals across different segments of the US economy. As a final check, column (6) returns to the baseline of column (1) but uses a weighted regression, weighting by initial employment, to confirm that the observed and predicted responses to tariff shocks are similarly aligned for small and large regions and sectors.

[27] Since the estimates of elasticities $\{\sigma, \psi^{X,F}, \psi^{M,F}\}$ already leverage the same tariff changes, one may wonder whether additional testing can be conducted. As discussed in ACD, the answer is yes. The reason is that our model relies on a large number of untested assumptions, from the structure of domestic input-output linkages to a lack of factor mobility. ACD's test implicitly sheds light on the overall credibility of those assumptions by using extra moment conditions, distinct from those already used in estimation. We also note that our test relies on responses for outcomes that are not used in estimation; namely, earnings per worker across regions and sectors.

4. REVEALING WELFARE WEIGHTS

4.1. *Empirical Specification*

We now return to the empirical specification suggested by Proposition 1: a regression of tariffs on a measure of the sensitivity of individuals' real earnings to imports in order to reveal the social marginal return of transfers to individuals, and hence the strength and nature of redistributive motives for protectionism.

For empirical purposes, we assume that welfare weights are an additively separable function of the "socioeconomic groups" $j \in \mathcal{J}$ to which individuals $n \in \mathcal{N}$ may belong,

$$\beta(n) = \sum_{j \in \mathcal{J}} \text{Dummy}_j(n) \times \beta_j, \tag{18}$$

where $\text{Dummy}_j(n)$ is an indicator variable that takes the value of one if individual n is a member of group j and zero otherwise, and β_j is the social marginal return of transfers to members of group j. Note that any individual can be a member of multiple groups. Note also that since $\frac{1}{N}\sum_{n\in\mathcal{N}} \beta(n) = 1$, as established in Section 2, we must also have $\frac{1}{N}\sum_{j\in\mathcal{J}} N_j\beta_j = 1$, with $N_j \equiv \sum_{n\in\mathcal{N}} \text{Dummy}_j(n)$ the number of individuals in group j.

We focus on the scope for redistribution based on individuals' sectors and regions. In particular, we model socio-economic groups that are defined according to two considerations: "working in sector s," with welfare weights $\{\beta_s\}_{s\in\mathcal{S}}$, and "residing in region r," with welfare weights $\{\beta_r\}_{r\in\mathcal{R}_H}$, respectively.[28] Thus, using the notation from Section 3, we can write our general tariff formula (9) as

$$t_{ih} = -\sum_{s\in\mathcal{S}} \beta_s N_s \frac{\partial(\omega_s - \bar{\omega})}{\partial m_{ih}} - \sum_{r\in\mathcal{R}_H} \beta_r N_r \frac{\partial(\omega_r - \bar{\omega})}{\partial m_{ih}} + \text{Controls}_{ih} + \varepsilon_{ih}, \tag{19}$$

where t_{ih} denotes the tariff on the 6-digit HS product h from foreign country i applied by the US in 2017, as described in Section 3.3. We restrict our sample to the 71,688 country-product pairs with US imports in 2017 above \$100,000.[29]

According to (19), the tariff t_{ih} is a function of four terms. The first two terms are our objects of primary interest. They capture redistribution across individuals according to their sectors of employment and states of residence. The unknown sets of parameters, β_s and β_r, represent the social marginal return of transfers to the individuals in any given sector s and region r. The corresponding regressors $\partial(\omega_s - \bar{\omega})/\partial m_{ih}$ and $\partial(\omega_r - \bar{\omega})/\partial m_{ih}$ capture the sensitivity of the average real earnings of individuals working in sector s of region r (relative to the US

[28]This interest is motivated, in part, by the contrasting predictions of models emphasizing sectoral mechanisms, such as Grossman and Helpman (1994), and regional mechanisms, such as Ma and McLaren (2018). We thus pursue a version with sector- and region-specific effects for reasons of parsimony, i.e., $|\mathcal{R}_H| + |\mathcal{S}|$ parameters to estimate rather than $|\mathcal{R}_H| \times |\mathcal{S}|$, even though separate values of β_j for each region-sector combination are formally identified. For the interested reader, Appendix Figure D.5 reports results from an augmented specification that also includes eight socioeconomic groups formed from the interactions between three binary categories associated with education (college-educated and not), gender, and race (white or non-white) available in region-sector ACS data. This version yields values of β_s and β_r that are extremely similar to those in our baseline, as further discussed in our working paper (Adao et al., 2023).

[29]For computational reasons, when implementing (19), we only consider variations of the tariffs for the 71,688 country-product pairs in our sample, while keeping all other tariffs fixed. Thus the regressors $\{\partial(\omega_s - \bar{\omega})/\partial m_{ih}\}$ and $\{\partial(\omega_r - \bar{\omega})/\partial m_{ih}\}$ formally correspond to the changes in real earnings caused by a marginal change in tariffs such that dm_{ih} is non-zero, all other changes in imports are zero within our sample, and all changes in tariffs are zero outside our sample.

average $\bar{\omega}$) to a change in the quantity of imports m_{ih} in product h from foreign country i. The measurement of such sensitivities was the focus of Section 3, with Figure 2 summarizing the variation in these regressors.[30] The third term in (19) refers to additional factors that we control for, beyond sector- and region-based redistributive motives. Our baseline analysis populates this set with an intercept, the terms-of-trade motive for trade protection $m \cdot (\partial p^w / \partial m_{ih})$, and, for computational reasons, the fiscal externalities associated with country-product pairs outside our sample.[31] Finally, ε_{ih} captures the impact of trade protection on distortions as well as any measurement error in trade taxes or misspecification. It is the unobserved error in our regression.

We begin by estimating equation (19) via OLS. This requires $\partial(\omega_s - \bar{\omega})/\partial m_{ih}$ and $\partial(\omega_r - \bar{\omega})/\partial m_{ih}$ to be uncorrelated with the residual ε_{ih} given our control set. Potential violations from such an orthogonality requirement may arise from the existence of other trade policy instruments, production subsidies, income taxes, exogenous constraints on US tariffs, misspecification of terms-of-trade controls due to negotiated trade taxes, and reverse causality between tariffs and imports. In Section 4.3, we show that alternative specifications meant to deal, in a theory-consistent way, with such concerns deliver estimates that are highly correlated with those obtained from our baseline specification.

A distinct challenge arises from the component of tariffs that may derive from attempts to correct externalities—the term $\epsilon \cdot (\partial z / \partial m_{ih})$ in equation (9). The plausibility of our orthogonality assumption inherently depends on the types of externalities that are thought to be empirically important. For instance, in the case where the consumption of various imported goods may generate different health hazards, so that $z = \{m_{ih}\}$ and the externality experienced by each individual is $E(\{m_{ih}\}, n) = \sum_{i,h} E_{ih} m_{ih}$—as is often considered in the product standards literature—our exclusion restriction requires no systematic correlation between health damage E_{ih} and the sensitivity of real earnings with respect to imports across country-product pairs. A second type of externality that has featured prominently in the study of distortion-correcting tariff policy concerns foreign externalities due to carbon emissions (e.g. Kortum and Weisbach, 2021 and Hsiao, 2026), which are a function of total production abroad and work through the world price p^w. While the vector $\partial p^w / \partial m_{ih}$ is too high-dimensional to control for directly, it seems likely that our flexible approach to controlling for terms-of-trade motives, discussed below, may do much to mitigate this concern. Finally, in the case of social identity in Grossman and Helpman (2021), the psychosocial cost externality is assumed to be a linear function of changes in others' real earnings. In this case, our estimates of welfare weights

[30] Since the value of these regressors depend on the estimates of the elasticities $\{\sigma, \psi^{X,F}, \psi^{M,F}\}$, the confidence intervals for the estimates of β_s and β_r that we report below take into account the sampling uncertainty in these three elasticities. To do so, we first augment FGKK's replication code to write their estimation problem as a just-identified GMM model (so as to obtain the variance-covariance matrix for $\{\sigma, \psi^{X,F}, \psi^{M,F}\}$). We then treat the regressors $\partial(\omega_s - \bar{\omega})/\partial m_{ih}$ and $\partial(\omega_r - \bar{\omega})/\partial m_{ih}$ as generated regressors whose standard errors are augmented using the delta method.

[31] Like in FGKK's original model, foreign export supply curves are (almost) perfectly elastic in our analysis. It does not follow, however, that the terms-of-trade motive, $m \cdot (\partial p^w / \partial m_{ih})$, is zero. While the prices of US imports do not respond to import restrictions, the prices of US exports do, as US tariffs tend to increase the demand for US products and raise US wages. In turn, from the point of view of manipulating the US terms of trade, the optimal US tariff is uniform, but non-zero. Due to the computational considerations referred to in footnote 29, $m \cdot (\partial p^w / \partial m_{ih})$ differs slightly from the optimal tariff on ih, which would have to be computed holding fixed the imports of all goods, not just those in our sample. For this reason, $m \cdot (\partial p^w / \partial m_{ih})$ exhibits a small amount of variance so we control for it. For the same reason, we also include in Controls_{ih} the sum of the fiscal externalities associated with goods outside our sample, consistent with Proposition A.3. In our baseline specification, we require the coefficient on these fiscal externalities to be minus one, also consistent with theory. Since the value of US imports on goods outside our sample is small, relaxing this restriction, or ignoring this adjustment altogether, has very little impact on our results.

FIGURE 3.—Distribution of estimated welfare weights

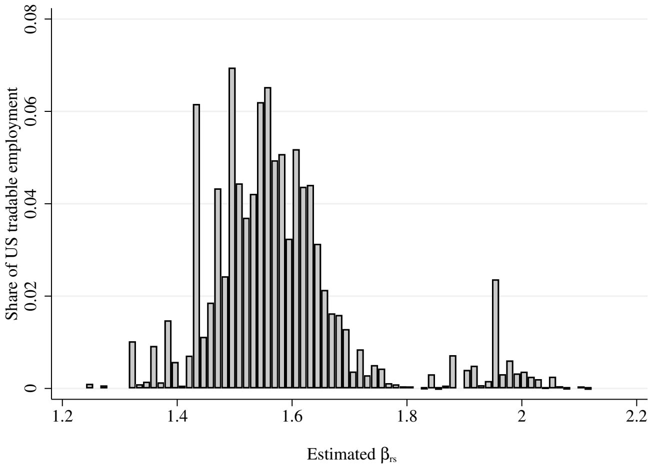

Notes: This figure displays the histogram of estimates of welfare weights across regions and tradable sectors, $\hat{\beta}_{rs} = \hat{\beta}_s + \hat{\beta}_r$, weighted by employment N_{rs}, where $\hat{\beta}_s$ and $\hat{\beta}_r$ denote the OLS estimates of β_s and β_r in (19). We omit 51 values, all corresponding to the Apparel sector in different regions, with $\hat{\beta}_{rs} \geq 2.5$.

would recover the total impact of a transfer to an individual, both direct and indirect via its effects on other individuals' welfare.

4.2. *Baseline Estimates*

Using the OLS estimates of β_s and β_r in equation (19), we can compute the welfare weights $\hat{\beta}_{rs} \equiv \hat{\beta}_s + \hat{\beta}_r$ for any of the N_{rs} individuals working in a given sector s and region r. Figure 3 reports the distribution of welfare weights across US individuals employed in tradable sectors. Three features are immediately apparent.

First, the fact that all welfare weights are above one implies that individuals in tradable sectors, who directly receive trade protection, tend to have higher social marginal utilities of income than individuals outside these sectors. Second, there is substantial dispersion in welfare weights among individuals employed in tradable sectors, with a long upper tail. Compared to individuals at the 10th percentile of our estimates, those located at the 90th, 95th, and 99th percentiles enjoy social marginal utilities of income that are 20%, 37%, and 239% higher, respectively. This underscores a clear sense in which trade policy is far from redistribution-neutral, since a world in which trade protection is not used to achieve redistributive goals would display no estimated welfare weight dispersion. Third, despite the wide range of these estimates, all estimated welfare weights $\hat{\beta}_{rs}$ are statistically larger than zero at standard significance levels. This lends credence to the constrained Pareto efficiency assumption underlying our methodology, which requires non-negative welfare weights for all individuals.

To explore further the determinants of welfare weights, Figure 4 reports separately our sector- and region-based estimates, $\hat{\beta}_s$ and $\hat{\beta}_r$, each normalized to have an employment-weighted mean of zero.[32] 95% confidence intervals for these estimates are also shown. Across sectors, estimates

[32]Our theory requires $\frac{1}{N}\sum_{r,s}\beta_{rs}N_{rs} = 1$. This restriction implies that $\{\beta_{rs}\}$ are identified. In contrast, $\{\beta_r\}$ and $\{\beta_s\}$, are only identified up to a constant: starting from $\beta_{rs} = \beta_r + \beta_s$, one can always construct $\beta'_r = \beta_r + c$ and $\beta'_s = \beta_s - c$ that satisfy $\beta_{rs} = \beta'_r + \beta'_s$, hence the need for a normalization in Figure 4.

FIGURE 4.—Estimates of welfare weights across sectors and regions

(a) $\hat{\beta}_s$ estimates (b) $\hat{\beta}_r$ estimates

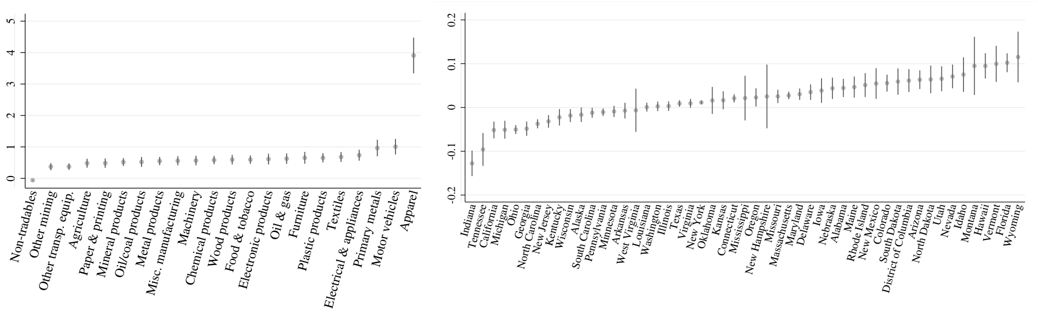

Notes: Figure 4a displays estimates of the sectoral component of welfare weights, $\hat{\beta}_s$, for each sector s, as obtained from the OLS estimation of (19) and normalized such that the mean of $\hat{\beta}_s$ across s is zero. Figure 4b does the same but for the regional component of welfare weight, $\hat{\beta}_r$, for each region r, normalized such that the mean of $\hat{\beta}_r$ across r is zero. Circles correspond to point estimates and bars denote 95% confidence intervals. Standard errors are clustered at the product-level and adjusted for sampling uncertainty in the estimates of the elasticities $\{\sigma, \psi^{X,F}, \psi^{M,F}\}$ that enter our regressors.

of $\hat{\beta}_s$ range from -0.06 for Services to 3.91 for Apparel, and while Apparel is a clear outlier, which was omitted from Figure 3, even the second- and third-largest values, 1.01 for Motor vehicles and 0.97 for Metals, are considerably higher than, and statistically different from, that of the Other mining sector, the tradable sector with lowest welfare weight.[33] As can be seen from Figure 4b, the region-level estimates $\hat{\beta}_r$ are also precisely estimated—for the reasons discussed at the end of Section 3.4—and reveal significant dispersion across regions, but these estimates are of a strikingly smaller scale than the sector-specific ones, ranging from -0.13 for Indiana to 0.12 for Wyoming.

One way to evaluate the economic magnitude of the previous estimates is to consider the average welfare weights within each sector and region, $\bar{\beta}_s \equiv \sum_r (N_{rs}/N_s)\hat{\beta}_{rs}$ and $\bar{\beta}_r \equiv \sum_s (N_{rs}/N_r)\hat{\beta}_{rs}$, which we report in Appendix Figure D.4. By construction, these welfare weights adjust for compositional differences in how individuals employed in different sectors are distributed across regions as well as how individuals from different states are distributed across sectors. Quantitatively, our estimates imply that individuals employed in Apparel, Motor Vehicles, and Metals, which combine to comprise 0.7% of total employment and 8.0% of tradable sector employment, enjoy an average social marginal utility of income that is 138% higher than the average across all US individuals. This means that, from society's perspective, giving \$1 to an individual in one of these three sectors is equivalent to giving \$2.38 to any randomly-chosen individual in the United States. In contrast, giving \$1 to an individual in the state with the highest social marginal utility of income is only equivalent to \$1.20 given to someone in the bottom state. This asymmetry between sectors and states foreshadows the greater importance of sector- rather than region-based considerations in accounting for US trade protection, as we demonstrate in Section 5 below.

[33]Not surprisingly, Apparel has the largest average ad-valorem tariffs of any sector by a wide margin. Appendix Figure D.3 contains a version of Figure 4a without Apparel for greater clarity.

TABLE II

SENSITIVITY ANALYSIS

	Specification	Correlation of $\widehat{\beta}_{rs}$ with baseline (1)	Standard deviation of $\widehat{\beta}_{rs}$ across sectors (2)	Standard deviation of $\widehat{\beta}_{rs}$ across states (3)
1	Baseline	–	0.74	0.05
2	Control for fiscal externality of NTMs	1.00	0.74	0.05
3	Drop observations with NTMs	0.81	1.08	0.08
4	Add AD and CV rates to dependent variable	1.00	0.75	0.05
5	Add fiscal externality of subsidies to dep. var.	1.00	0.75	0.05
6	Control for fiscal externality of subsidies	1.00	0.77	0.05
7	Control flexibly for fiscal externality of subsidies	0.98	1.41	0.11
8	Measure income post-income taxes	1.00	0.98	0.06
9	Drop constrained tariff lines	0.95	0.89	0.17
10	Add MFN clause constraint	1.00	0.79	0.07
11	Control for country terms-of-trade (ToT) effects	0.99	0.90	0.06
12	Control for country-sector ToT effects	0.95	1.05	0.07
13	Use alternative model as IV	0.88	0.33	0.05

Notes: Rows 2-13 of this table report results from specifications that introduce alternative modeling assumptions to those in the baseline (row 1). See text for details. "AD" and "CV" refer to antidumping and countervailing duties, respectively.

4.3. *Sensitivity Analysis*

The baseline estimates of $\hat{\beta}_s$ and $\hat{\beta}_r$ reported in Section 4.2 were obtained under a number of assumptions that we now probe further. Table II describes results from 12 different approaches to estimating these parameters under a range of alternative specifications, each displayed in a separate row. For each of them, column 1 reports how alternative estimates of $\hat{\beta}_s$ and $\hat{\beta}_r$ correlate with the baseline values described in Section 4.2 and columns 2 and 3 report the dispersion of the estimates of welfare weights across sectors and regions. We summarize each of these 12 alternative procedures here, and provide further details about methods and auxiliary data sources in the Online Appendix.

- **Other trade policy instruments:** To deal with non-tariff measures (NTMs), we implement the tariff formula from Proposition A.1. First, in row 2 we add controls for the fiscal externality associated with the revenue potentially generated by the six types of NTMs available from TRAINS (UNCTAD, 2023) as well as those reported in the Temporary Trade Barriers database (Bown et al., 2020). Second, in row 3 we restrict ourselves to the subsample of 26% of observations in which there is no NTM in place at all, and the new and old formula coincide. Finally, in row 4 we add to our measure of tariffs t_{ih} an estimate of antidumping and countervailing duties.
- **Production subsidies:** To deal with production subsidies, we implement the tariff formula from Proposition A.2. It highlights an extra fiscal externality equal to $\sum_{f\in\mathcal{F}} s^y(f)\cdot \partial y(f)/\partial m_{ih}$. We account for it by measuring (following Ferrari and Ossa, 2023) state and local subsidies $s^y(f)$ (W.E. Upjohn Institute for Employment Research, 2023, Good Jobs First, 2023, Environmental Working Group, 2023) and either subtracting $\sum_{f\in\mathcal{F}} s^y(f)\cdot \partial y(f)/\partial m_{ih}$ from the dependent variable in row 5 or controlling for this effect (which relaxes the assumption that its coefficient is one) in row 6. Alternatively, in row 7 we control for the interaction of $\partial y(f)/\partial m_{ih}$ with sector and region fixed-effects, which dis-

penses with the need to measure $s^y(f)$ (under the assumption that subsidies take the form $s^y(f) = s^y_s + s^y_r$).

- **Income taxes:** We again use Proposition A.2. Our baseline analysis omits income taxation. More generally, one needs to distinguish between pre-tax and post-tax real incomes. In row 8 we do so by using data on state and federal income tax schedules from TaxSim (Feenberg and Coutts, 1993) and microdata from the American Community Survey (ACS) (Ruggles et al., 2023), which allows us to compute an income-weighted average of marginal tax rates faced by individuals within each state-sector pair.
- **Constrained trade taxes:** To deal with constraints on a country's ability to choose its trade taxes, we use the tariff formula from Proposition A.3. In row 9 we use WTO data (World Trade Organization, 2023) in order to confine attention to the 3.6% of observations that appear to be unconstrained (that is, where no trade agreement applies and the MFN rate exhibits "overhang") and add controls for the fiscal externalities associated with potentially constrained trade taxes (that is, those in the complementary 96.4% of observations). As a second exercise, row 10 focuses on the constraint imposed by MFN itself. This calls for using as regressors the marginal changes with respect to $M_{\text{WTO}h} \equiv \sum_{i\in\text{WTO}} m_{ih}$ rather than m_{ih}, since $M_{\text{WTO}h}$ is what the government controls.[34]
- **Negotiated trade taxes:** Our baseline specification, which controls for the terms-of-trade motive via $m \cdot (\partial p^w/\partial m_{ih})$, is valid if Home places the same welfare weight on all foreign individuals, including zero weight, as is commonly assumed. However, preferential trade agreements may lead US policy to internalize foreign countries' welfare in heterogeneous ways. To allow for this possibility, we implement the tariff formula from Proposition A.4. Row 11 first includes controls for 101 separate terms-of-trade motives, one for each foreign country. Going further, row 12 controls for 303 separate terms-of-trade effects within each combination of country and 3 broad sectors, which allows for sector-level political economy considerations within foreign countries to affect trade negotiations with the US.[35]
- **Instrumental variable specification:** Our baseline analysis measures the sensitivity of imports to real earnings via a model in which tariff revenue is redistributed uniformly across all agents. One way to assess the potential for misspecification due to this assumption is to use an instrumental variable (IV) that measures sensitivity in an alternative model that does not redistribute tariff revenue. Further, if this alternative model is calibrated to the case of a counterfactual economy with zero tariffs (as in row 13) then this IV also probes concerns about simultaneity bias, as in prior work (Trefler, 1993, Goldberg and Maggi, 1999).

As Table II makes clear, the resulting estimates of $\hat{\beta}_s$ and $\hat{\beta}_r$ that we obtain across these 12 specifications remain similar to those of our baseline analysis. In particular, the correlation between baseline and alternatives (column 1) is never lower than 0.81. And the core finding of Section 4.2—that the variance of welfare weights across sectors (column 2) is considerably larger than that across states (column 3)—remains true in all specifications.

[34] An additional potential constraint on trade taxes is that they are non-negative—import subsidies are not used by the US. As in analogous procedures used by Trefler (1993) and Goldberg and Maggi (1999), one can explore the role of such censoring with the use of a Tobit model under the extra assumption that latent tariffs are normally distributed. The resulting estimates of welfare weights are extremely similar in such an analysis, as discussed in Adao et al. (2023).

[35] Within each foreign country, this specification controls separately for the terms-of-trade effects associated with imports and exports of goods in each of three broad sector groups: agriculture, manufacturing, or services.

FIGURE 5.—Variance decomposition of US tariffs in 2017

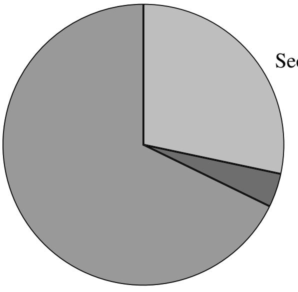

Notes: This figure plots the share of variance in US tariffs t_{ih} in 2017 that can be explained, according to estimates of equation (19), due to each of three components: redistribution based on individuals' sector of employment; redistribution based on individuals' state of residence; and other factors. The decomposition of variance reported is computed using the Owen-Shapley method.

5. HOW IMPORTANT IS REDISTRIBUTIVE TRADE PROTECTION?

5.1. *Cross-Sectional Variation in Tariffs*

A first way to evaluate the importance of the redistributive motive for trade protection is to ask how much of the observed cross-sectional variation in US tariffs in 2017 can be explained by the combination of the sector- and region-based redistributive motives.

To answer this question, we carry out an Owen-Shapley regression decomposition that returns the share of the variance explained by the first two sets of regressors in (19). The results of this decomposition are displayed in Figure 5. Two findings are evident. First, although we aim to explain the variation in 71,688 tariff lines using changes in real earnings across only 51 regions and 22 sectors, the redistributive motives, either due to sector- or region-based considerations, account for about one third (32.1%) of the total variation in US trade policy. Second, sector-based motives for redistribution explain the lion's share of total redistributive motives (28.2%), implying that region-based considerations are indeed relatively minor (3.9%).

5.2. *Gains from Redistributive Trade Protection*

Another way to evaluate the importance of redistributive protection is to ask, from an economic rather than a statistical standpoint: how large are the monetary transfers caused by redistributive trade protection?

To provide estimates of the causal impact of redistributive tariffs on the real income of different individuals, we return to the model from Section 3 to construct a counterfactual US economy with trade taxes purged of the redistributive component that we have estimated in Section 4. Formally, we set counterfactual trade taxes t'_{ih} equal to

$$t'_{ih} = t_{ih} - (t_{R,ih} - t_R) \tag{20}$$

with t_{ih} the observed US tariff, $t_{R,ih} \equiv -\sum_{s\in\mathcal{S}} \hat{\beta}_s N_s \frac{\partial(\omega_s - \bar{\omega})}{\partial m_{ih}} - \sum_{r\in\mathcal{R}_H} \hat{\beta}_r N_r \frac{\partial(\omega_r - \bar{\omega})}{\partial m_{ih}}$ the redistributive component associated with our baseline estimates of welfare weights, and t_R the average value of $t_{R,ih}$. This is equivalent to considering a counterfactual US economy in which social marginal returns to different individuals have been equalized, holding fixed the average level of trade protection. We then calculate the counterfactual changes in real income of all individuals in the economy. Gains from redistributive trade protection are equal to minus these real income changes. They can be interpreted as the as-if transfer, either positive or negative,

FIGURE 6.—Distribution of gains from redistributive trade protection

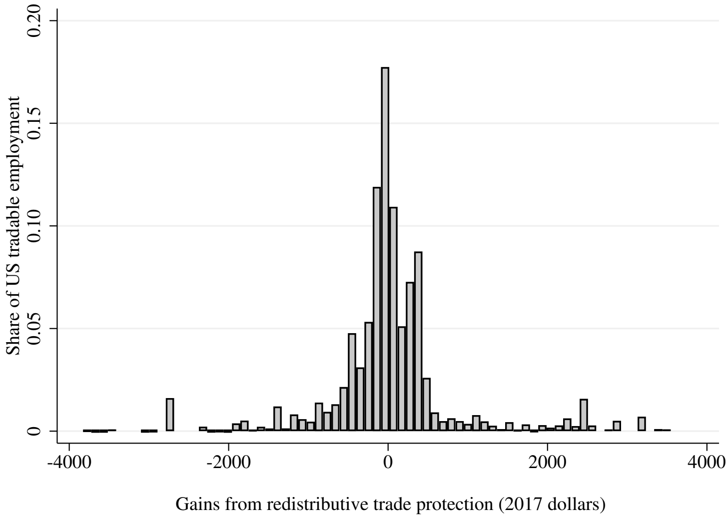

Notes: This figure reports the histogram of the gains from redistributive trade protection across regions and tradable sectors, weighted by employment N_{rs}. Gains are defined as minus the change in real earnings that results from a counterfacual US economy in which US tariffs are taken from their observed 2017 values to the value that would obtain in the absence of redistributive motives, as described in equation (20). We omit 20 region-sector pairs with either gains above or losses below $4,000.

that each individual receives as a result of the redistributive motive embedded in the US tariff schedule of 2017.

Figure 6 displays the distribution of the gains from redistributive trade protection across individuals employed in tradable sectors. In line with the substantial dispersion in welfare weights seen in Figure 4, some individuals experience sizable gains from redistributive trade protection. For individuals located at the 90th, 95th, and 99th percentiles of the distribution, the changes in real income caused by the redistributive component of the US tariffs are $538, $1,727, and $2,861, respectively. In contrast, those at the 10th percentile lose $621.

6. POLITICAL ECONOMY MECHANISMS

Up to this point, we have remained agnostic about the specific dimensions of the US political process that may be driving redistributive tariffs. All that matters for our estimates of the welfare weights β_s and β_r, as well as for the associated gains from redistributive tariffs, is that this process arrives at some Pareto-efficient outcome. To conclude our analysis, we illustrate how to use our estimates in order to discriminate between two leading explanations for the existence of tariffs in the previous political economy literature: sectors' abilities to lobby, as formalized by Grossman and Helpman (1994), and states' abilities to swing presidential elections, as formalized by Ma and McLaren (2018).

6.1. *Structurally Interpreting Welfare Weights*

In Grossman and Helpman (1994), lobbying is modeled as a menu auction à la Bernheim and Whinston (1986). A subset of sectors, those that are politically organized, can make political contributions that are contingent on the government's choice of trade policy, whereas other sectors cannot. The government's objective function is assumed to be a weighted sum of total political contributions and the welfare of the economy's representative agent. For our purposes, the key implication of Grossman and Helpman's (1994) protection for sale model is that welfare weights take the form:

$$\beta_{GH}(n) = \begin{cases} \frac{a+1}{a+\alpha_L} & \text{if individual } n \text{ is employed in a politically organized sector,} \\ \frac{a}{a+\alpha_L} & \text{otherwise,} \end{cases} \tag{21}$$

where $a \geq 0$ denotes the weight that the government assigns to aggregate welfare relative to political contributions and α_L denotes the share of the population employed in politically organized sectors. Hence welfare weights only depend on whether individuals are employed in a politically organized sector, in which case they receive higher weights.

In contrast, Ma and McLaren (2018) focus on a political environment in which two national parties compete à la Lindbeck and Weibull (1993). Citizens are located in different states, with different ideological biases. Swing states correspond to neutral states. National parties choose trade policies to maximize the number of states in which they receive a majority of votes. The key implication of Ma and McLaren (2018) is that welfare weights take the form:

$$\beta_{MM}(n) = \begin{cases} \beta_r & \text{if individual } n \text{ is employed in a region } r \text{ that is a swing state,} \\ 0 & \text{otherwise,} \end{cases} \tag{22}$$

where $\beta_r \geq 0$ denotes the weight that the government assigns to any voter in region r. Hence welfare weights only depend on whether individuals are located in a swing state, in which case they receive higher weights.

6.2. *Testing Existing Political Economy Models*

In both models, welfare weights, as described in equations (21) and (22), are sufficient statistics for the impact of the political process on the structure of trade policy. To test these models' core political economy predictions—while abstracting from their simplifying assumptions about the economic environment—we can therefore check whether our estimates of welfare weights are systematically higher for politically organized sectors and/or for swing states.[36]

In their original work, Goldberg and Maggi (1999) defined a sector as politically organized if its campaign contributions were above a certain threshold, with the threshold suggested by a natural break in the data. In the same spirit, we define sectors as politically organized based on LobbyView data (Kim, 2018) that allow us to compute total lobbying expenditures during 2000-2016 on filings that cite trade policy as their primary issue, as further described in Appendix B.2.[37] To find a natural break in the data, we then use a k-means clustering procedure, with

[36] Such a test implicitly relies on the assumption that the government, in the case of Grossman and Helpman (1994), or political parties, in the case of Ma and McLaren (2018), have the same beliefs as households and firms about the economic environment, as described in Section 3. This symmetric treatment of policymakers, households, and firms is standard in the existing literature, including the numerous empirical tests of Grossman and Helpman's (1994) protection for sale model. It should be clear, however, that if policymakers have different beliefs about the economic environment, then identifying their preferences would also require taking a stand on what these beliefs are. While we do not attempt to contribute to the literature along this dimension, recent empirical work by Bombardini et al. (2023) suggests a potential way forward.

[37] Our decision to focus on lobbying expenditure data, which were first used to study trade policy in Bombardini and Trebbi (2012), reflects the fact that such expenditures can be linked to specific policy issues as well as the arguments about the long-standing monetary dominance of this form of political influence voiced in Bombardini and Trebbi (2020).

TABLE III

POLITICAL ECONOMY MECHANISMS: A TEST

	Dependent variable: welfare weight ($\widehat{\beta}_{rs}$)		
	(1)	(2)	(3)
Politically organized sector	0.903		0.904
	(0.087)		(0.087)
Swing state		0.030	0.032
		(0.034)	(0.031)
Constant	0.988	0.993	0.981
	(0.015)	(0.019)	(0.017)
R^2	0.209	0.003	0.212

Notes: Sample of 1,033 region-sector pairs with positive employment in 2017. All specifications are weighted by employment in 2017. The dependent variable is the estimated welfare weight of each region-sector, $\hat{\beta}_{rs} = \hat{\beta}_s + \hat{\beta}_r$ with $\hat{\beta}_s$ and $\hat{\beta}_r$ the OLS estimates of β_s and β_r in (19). Robust standard errors in parentheses.

$k = 2$, based on the log of trade-lobbying spending per worker. The sectors in the politically organized group are: Apparel, Chemical products, Electronic products, Oil and coal products, and Primary metals.

Likewise, we divide US regions into swing states and non-swing states based on margins of victory in past elections as reported by the MIT Election Data and Science Lab (MIT Election Data and Science Lab, 2017). Specifically, we focus on presidential elections between 2000 and 2016 and compute the absolute difference between the average voting shares of Democrat and Republican candidates in each state over that time period, as further described in Appendix B.2. Following the earlier work of Ma and McLaren (2018) and Bown et al. (2023), we then define a US region as a swing state if the margin of victory is below 5%. The swing states are: Colorado, Florida, Iowa, Nevada, New Hampshire, Ohio, Pennsylvania, Virginia, and Wisconsin.

Table III reports the results of OLS regressions that project the welfare weights (estimated in Section 4.2) on two dummy variables, one for whether or not individuals are employed in politically organized (i.e., high trade-lobbying) sectors and one for whether or not they are located in a swing state. Column (1) focuses on the politically organized dummy. We see that those in organized sectors have statistically significantly higher welfare weights, with the organized dummy explaining about a fifth of the overall variation in welfare weights across sectors and regions. The ratio of the slope (0.903) to the intercept (0.988) identifies $1/\hat{a} = 0.91$. This implies that workers employed in politically organized sectors enjoy an average social marginal utility that is about twice as large as that for workers employed in other sectors. Column (2) turns to the swing state dummy. In contrast, consistent with our earlier variance decomposition highlighting the importance of sector- relative to region-based characteristics, we find no statistically significant effect of being in a swing state. The same pattern holds when both politically organized sector and swing state dummies are included, as can be seen from Column (3).[38]

[38]Although neither Grossman and Helpman (1994) nor Ma and McLaren (2018) study the interaction between sectors' abilities to lobby and states' abilities to swing presidential elections, an interesting possibility is that politically organized sectors may systematically receive more trade protection when they are located in swing states. Our baseline analysis also abstracts from such interaction by assuming that welfare weights β_{rs} take an additive form $\beta_{rs} = \beta_r + \beta_s$. A simple way to extend our empirical strategy to shed light on this interaction is to assume instead that $\beta_{rs} = \beta_r + \beta_s + \gamma \times \text{Dummy}(r = \text{swing state}) \times \text{Dummy}(s = \text{politically organized})$. When we do so, and otherwise follow the exact same procedure as in Section 4.1, we obtain estimates of β_r and β_s that are almost identical, but the interaction term is positive and statistically significant $\hat{\gamma} = 0.24$ (0.06), consistent with a role for swing states to shape US trade policy. We thank an anonymous referee for suggesting this exercise.

7. CONCLUDING REMARKS

Why is trade not free? A prominent answer to this question is redistributive politics. In this paper we have developed a revealed-preference approach to identify who the politically-favored are and, in turn, to quantify the importance of redistributive tariffs.

Our approach builds on a general tariff formula that emerges from any political process provided that such a process results in policies that are constrained Pareto efficient. It highlights a simple sense in which the redistributive motives behind the tariff observed on any good should be the product of two considerations: the as-if welfare weights of different constituents of society and the marginal impact of that good's imports on the real income of these constituents. Inverting this logic, a simple regression of tariffs on estimates of the marginal welfare impact of imports can reveal the point on society's Pareto frontier that the political process arrives at.

We have applied the previous methodology to US trade policy in 2017 and estimated welfare weights across individuals from 50 states, plus Washington DC, and 22 sectors. Our estimates imply that redistributive trade protection among these broad groups accounts for almost one third of the variance in US tariffs observed across thousands of products and origin countries. The monetary transfers caused by redistributive tariffs are large as well, with the estimated difference between gains at the 90th percentile and losses at the 10th percentile equal to $1,159 per worker annually. The previous conclusions are mainly driven by differences in welfare weights across sectors, with differences across states only playing a minor role. In line with this observation, the welfare weights that we estimate are consistent with lobbying being a significant driver of US trade protection. Perhaps surprisingly, swing states in US presidential elections play a much smaller role.

While trade is decidedly not free, import tariffs are by no means the only policy tools available to governments seeking to help some of their constituents at the expense of others. Environmental policy, competition policy, and financial regulation are all areas to which the approach developed in this article would be straightforward to apply. In all such cases, we hope that our analysis can offer a blueprint for identifying who the politically-favored are and for evaluating the economic importance of the political favors they receive.

REFERENCES

ACEMOGLU, DARON (2003): "Why not a Political Coase theorem? Social conflict, commitment, and politics," *Journal of Comparative Economics*, 31, 620–652. [9]

ACEMOGLU, DARON AND JAMES A. ROBINSON (2001): "Inefficient Redistribution," *The American Political Science Review*, 95 (3), 649–661. [9]

ADAO, RODRIGO, ARNAUD COSTINOT, AND DAVE DONALDSON (2025): "Putting Quantitative Models to the Test: An Application to the US-China Trade War," *Quarterly Journal of Economics*, 140 (2), 1471–1524. [3, 21]

ADAO, RODRIGO, ARNAUD COSTINOT, DAVE DONALDSON, AND JOHN STURM (2023): "Why is Trade Not Free? A Revealed Preference Approach," NBER working paper number 31798. [15, 22, 27]

BAGWELL, KYLE AND ROBERT W. STAIGER (1999): "An Economic Theory of GATT," *American Economic Review*, 89 (1), 215–248. [5]

——— (2002): *The Economics of the World Trading System*, MIT press. [13]

BALDWIN, RICHARD (1987): "Politically Realistic Objective Functions and Trade Policy," *Economics Letters*, 24, 287–290. [4]

BECKER, GARY S. (1983): "A Theory of Competition Among Pressure Groups for Political Influence," *Quarterly Journal of Economics*, 98 (3), 371–400. [9]

BERNHEIM, B. DOUGLAS AND MICHAEL D. WHINSTON (1986): "Menu Auctions, Resource Allocation, and Economic Influence," *Quarterly Journal of Economics*, 101 (1), 1–31. [29]

BESLEY, TIMOTHY AND STEPHEN COATE (1997): "An Economic Model of Representative Democracy," *Quarterly Journal of Economics*, 112 (1), 85–114. [8]

BOMBARDINI, MATHILDE (2008): "Firm heterogeneity and lobby participation," *Journal of International Economics*, 75 (2), 329–348. [2]

BOMBARDINI, MATILDE, BINGJING LI, AND FRANCESCO TREBBI (2023): "Did US politicians expect the China shock?" *American Economic Review*, 113 (1), 174–209. [30]

BOMBARDINI, MATILDE AND FRANCESCO TREBBI (2012): "Competition and political organization: Together or alone in lobbying for trade policy?" *Journal of International Economics*, 97 (1), 18–26. [4, 30]

——— (2020): "Empirical models of lobbying," *Annual Review of Economics*, 12, 391–413. [30]

BORUSYAK, KIRILL, PETER HULL, AND XAVIER JARAVEL (2022): "Quasi-experimental shift-share research designs," *Review of Economic Studies*, 89 (1), 181–213. [20]

BOURGUIGNON, FRANCOIS AND AMEDEO SPADARO (2012): "Tax-benefit revealed social preferences," *The Journal of Economic Inequality*, 10 (75-108). [5]

BOWN, CHAT, MILLA CIESZKOWSKY, AKSEL ERBAHAR, AND JOSE SIGNORET (2020): *Temporary Trade Barriers Database*, Washington, D.C.: World Bank. [26]

BOWN, CHAD, PAOLA CONCONI, AKSEL ERBAHAR, AND LORENZO TRIMARCHI (2023): "Politically Motivated Trade Protection," Unpublished manuscript. [31]

BRANDER, JAMES AND BARBARA SPENCER (1984): "Tariff Protection and Imperfect Competition," in *Monopolistic Competition and International Trade*, ed. by H. Kierzkowski, Clarendon Press, Oxford. [14]

BUREAU OF ECONOMIC ANALYSIS (2019): *2017 Benchmark Input-Output Accounts*, U.S. Department of Commerce. [18]

——— (2023): *State Annual GDP (SAGDP)*, U.S. Department of Commerce. [18]

CHARI, V. V., PATRICK J. KEHOE, AND ELLEN C. MCGRATTAN (2007): "Business Cycle Accounting," *Econometrica*, 75, 781–836. [5]

COATE, STEPHEN AND STEPHEN MORRIS (1995): "On the Form of Transfers to Special Interests," *Journal of Political Economy*, 103 (6), 1210–1235. [9]

CONCONI, PAOLA, FLORIN CUCU, FEDERICO GALLINA, AND MATTIA NARDOTTO (2024): "A Political Disconnect? Evidence from Voting on EU Trade Agreements," *CEPR Discussion Paper DP19609*. [4]

CONCONI, PAOLA, GIOVANNI FACHINI, AND MAURIZIO ZANARDI (2014): "Policymakers' horizon and trade reforms: The protectionist effect of elections," *Journal of International Economics*, 94, 102–118. [4]

CONCONI, PAOLA, MANUEL GARCÍA SANTANA, LAURA PUCCIO, AND ROBERTO VENTURINI (2018): "From Final Goods to Inputs: The Protectionist Effect of Preferential Rules of Origin," *American Economic Review*, 108 (8), 2335–2365. [12]

CORDEN, MAX (1984): "Market disturbances and protection: efficiency versus the conservative social welfare function," *Discussion paper no. 92, Centre for Economic Policy Research, Australian National University*. [9]

COSTINOT, ARNAUD AND IVAN WERNING (2023): "Robots, Trade, and Luddism: A Sufficient Statistics to Optimal Technology Regulation," *Review of Economic Studies*, 90 (5), 2261–2291. [5, 12]

DE V. GRAAFF, J. (1949-1950): "On Optimum Tariff Structures," *The Review of Economic Studies*, 17 (1), 47–59. [5]

DEARDORFF, ALAN V. (1987): "Safeguard Policy and the Conservative Social Welfare Function," in *Protection and Competition in International Trade: Essays in Honor of W. M. Corden*, ed. by Henryk Kierzkowski, Oxford: Blackwell, 22–40. [9]

DHINGRA, SWATI (2014): "Reconciling observed tariffs and the median voter model," *Economics and Politics*, 26 (3), 483–504. [8]

DIAMOND, PETER A. AND JAMES A. MIRRLEES (1971): "Optimal Taxation and Public Production II: Tax Rules," *American Economic Review*, 61 (3), 261–278. [5]

DIXIT, AVINASH, GENE M. GROSSMAN, AND ELHANAN HELPMAN (1997): "Common Agency and Coordination: General Theory and Application to Government Policy Making," *Journal of Political Economy*, 105 (4), 752–769. [9]

DIXIT, AVINASH AND JOHN LONDREGAN (1996): "The Determinants of Success of Special Interests in Redistributive Politics," *The Journal of Politics*, 58 (4), 1132–1155. [4, 8]

DOWNS, ANTHONY (1957): *An Economic Theory of Democracy*, New York: Harper Collins. [8]

ENVIRONMENTAL WORKING GROUP (2023): *EWG Farm Subsidy Database*. [26]

FAJGELBAUM, PABLO D., CECILE GAUBERT, NICOLE GORTON, EDUARDO MORALES, AND EDOUARD SCHAAL (2023): "Political Preferences and the Spatial Distribution of Infrastructure: Evidence from California's High-Speed Rail," NBER working paper no. 31438. [5]

FAJGELBAUM, PABLO D., PINELOPI K. GOLDBERG, PATRICK J KENNEDY, AND AMIT K KHANDELWAL (2020): "The Return to Protectionism," *Quarterly Journal of Economics*, 135 (1), 1–55. [3, 14]

FEENBERG, DANIEL R. AND ELISABETH COUTTS (1993): "An Introduction to the TAXSIM Model," *Journal of Policy Analysis and Management*, 12 (1), 189–194. [27]

FERRARI, ALESSANDRO AND RALPH OSSA (2023): "A quantitative analysis of subsidy competition in the US," *Journal of Public Economics*, 224, 104919. [26]

FRANCOIS, JOSEPH AND DOUGLAS R. NELSON (2014): "Political Support for Trade Policy in the European Union," *European Journal of Political Economy*, 36, 243–253. [5]

GARDNER, BRUCE L. (1987): "Causes of U.S. Farm Commodity Programs," *Journal of Political Economy*, 95 (2), 290–310. [5]

GAWANDE, KISHORE AND USREE BANDYOPADHYAY (2000): "Is protection for sale? Evidence on the Grossman-Helpman theory of endogenous protection," *Review of Economics and Statistics*, 82 (1), 139–152. [2]

GAWANDE, KISHORE AND PRAVIN KRISHNA (2003): "The Political Economy of Trade Policy: Empirical Approaches," in *Handbook of International Trade*, ed. by E. Kwan Choi and James Harrigan, Blackwell Publishing. [1]

GAWANDE, KISHORE, PRAVIN KRISHNA, AND MARCELO OLARREAGA (2009): "What Governments Maximize and Why: The View from Trade," *International Organization*, 63, 491–532. [2]

GAWANDE, KISHORE, PABLO M. PINTO, AND SANTIAGO M. PINTO (2024): "Heterogeneous Districts, Interests, and Trade Policy," *mimeo*. [5]

GOLDBERG, PENNY AND GIOVANNI MAGGI (1999): "Protection for Sale: An Empirical Investigation," *American Economic Review*, 89 (5), 1135–1155. [2, 18, 27, 30]

GOLDBERG, PINELOPI K. AND NINA PAVCNIK (2016): "The Effects of Trade Policy," in *Handbook of Commercial Policy*, ed. by Kyle Bagwell and Robert W. Staiger, Elsevier, vol. 1. [5]

GOOD JOBS FIRST (2023): *Subsidy Tracker*. [26]

GREENWALD, BRUCE C. AND JOSEPH E. STIGLITZ (1986): "Externalities in Economies with Imperfect Information and Incomplete Markets," *Quarterly Journal of Economics*, 101 (2), 229–264. [5]

GROSSMAN, GENE M. AND ELHANAN HELPMAN (1994): "Protection for Sale," *American Economic Review*, 84 (4), 833–850. [2, 4, 5, 8, 9, 22, 29, 30, 31]

——— (1995): "Trade Wars and Trade Talks," *Journal of Political Economy*, 103 (4), 675–708. [5, 14]

——— (2001): *Special Interest Politics*, MIT Press. [8]

——— (2005): "A Protectionist Bias in Majoritarian Politics," *Quarterly Journal of Economics*, 120 (4), 1239–1282. [4]

——— (2021): "Identity politics and trade policy," *The Review of Economic Studies*, 88 (3), 1101–1126. [5, 23]

HELPMAN, ELHANAN (1995): "Politics and Trade Policy," *NBER Working Paper 5309*. [4]

HELPMAN, ELHANAN AND PAUL R. KRUGMAN (1989): *Trade Policy and Market Structure*, Cambridge, Massachussetts: MIT Press. [5]

HILLMAN, ARYE L. (1982): "Declining Industries and Political-Support Protectionist Motives," *American Economic Review*, 72, 1180–1187. [4]

HSIAO, ALLAN (2026): "Coordination and Commitment in International Climate Action: Evidence from Palm Oil," *Econometrica*, 94 (1), 1–33. [5, 23]

HSIEH, CHANG-TAI AND PETER J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (4), 1403–1448. [5]

JACOBS, BAS, EGBERT L.W. JONGEN, AND FLORIS T. ZOUTMAN (2017): "Revealed social preferences of Dutch political parties," *Journal of Public Economics*, 156, 81–100. [5]

KIM, IN SONG (2017): "Political cleavages within industry: firm-level lobbying for trade liberalization," *American Political Science Review*, 111 (1), 1–20. [4]

——— (2018): *LobbyView: Firm-Level Lobbying & Congressional Bills Database*, Working Paper, MIT. [30]

KORTUM, SAMUEL AND DAVID WEISBACH (2021): "Optimal Unilateral Carbon Policy," Cowles Foundation Discussion Paper No. 2311. [5, 23]

LINDBECK, ASSAR AND JORGEN W. WEIBULL (1993): "A model of political equilibrium in a representative democracy," *Journal of Public Economics*, 51 (2), 195–209. [30]

MA, XIANGJUN AND JOHN MCLAREN (2018): "A swing-state theorem, with evidence," NBER working paper no. w24425. [4, 5, 8, 22, 29, 30, 31]

MARKUSEN, JAMES R. (1975): "International Externalities and Optimal Tax Structures," *Journal of International Economics*, 5, 15–29. [5]

MATSCHKE, XENIA AND SHANE M. SHERLUND (2006): "Do labor issues matter in the determination of US trade policy? An empirical reevaluation," *American Economic Review*, 96, 405–421. [2]

MAYER, WOLFGANG (1984): "Endogenous Tariff Formation," *American Economic Review*, 74, 970–985. [4, 8]

MCLAREN, JOHN (2016): "The Political Economy of Commercial Policy," in *Handbook of Commercial Policy*, ed. by Kyle Bagwell and Robert W. Staiger, Elsevier. [1]

MIT ELECTION DATA AND SCIENCE LAB (2017): *U.S. President 1976–2020*, Harvard Dataverse, Version 6. [31]

MITRA, DEVASHISH, DIMITROS D. THOMAKOS, AND MEHMET A. ULUBASOGLU (2002): ""Protection for sale" in a developing country: democracy vs. dictatorship," *Review of Economics and Statistics*, 84 (3), 497–508. [2]

MRAZOVA, MONIKA (2023): "Trade Agreements when Profits Matter," *mimeo*. [14]

OSSA, RALPH (2011): “A "New Trade" Theory of GATT/WTO Negotiations,” *Journal of Political Economy*, 119 (1), 112–152. [14]

——— (2014): “Trade Wars and Trade Talks with Data,” *American Economic Review*, 114 (12), 4104–46. [5, 18]

PERSSON, TORSTEN AND GUIDO TABELLINI (2002): *Political Economics: Explaining Economic Policy*, MIT Press. [4]

RODRIK, DANI (1986): “Tariffs, Subsidies, and Welfare with Endogenous Policy,” *Journal of International Economics*, 21, 285–299. [9]

——— (1995): “Political Economy of Trade Policy,” in *Handbook of International Economics*, ed. by Gene M. Grossman and Kenneth Rogoff, Elsevier, vol. 3. [1]

RUGGLES, STEVEN, SARAH FLOOD, MATTHEW SOBEK, ET AL. (2023): *IPUMS USA: Version 13.0*, Minneapolis, MN: IPUMS, [dataset]. [27]

TREFLER, DANIEL (1993): “Trade liberalization and the theory of endogenous protection: an econometric study of US import policy,” *Journal of Political Economy*, 101 (1), 138–160. [27]

TYERS, ROD (1990): “Implicit Policy Preferences and the Assessment of Negotiable Trade Policy Reforms,” *European Economic Review*, 34, 1399–1426. [5]

UNCTAD (2023): *Trade Analysis Information System (TRAINS)*, United Nations Conference on Trade and Development. [26]

U.S. CENSUS BUREAU (2019): *2017 Commodity Flow Survey Public Use Microdata*, U.S. Department of Commerce. [18]

——— (2020a): *USA Trade Online — National-Level Trade Data*, U.S. Department of Commerce. [18]

——— (2020b): *USA Trade Online — State Trade Data*, U.S. Department of Commerce. [18]

U.S. INTERNATIONAL TRADE COMMISSION (2018): *HTS Annual Tariff Schedule, 2017*, Washington, DC: USITC. [18]

VENABLES, ANTHONY J. (1987): “Trade and Trade Policy with Differentiated Products: A Chamberlinian-Ricardian Model,” *The Economic Journal*, 97 (387), 700–717. [14]

W.E. UPJOHN INSTITUTE FOR EMPLOYMENT RESEARCH (2023): *State Business Incentive Database*. [26]

WERNING, IVAN (2007): “Pareto Efficient Income Taxation,” Unpublished manuscript. [5]

WITTMAN, DONALD (1989): “Why Democracies Produce Efficient Results,” *Journal of Political Economy*, 97 (6), 1395–1424. [9]

WORLD TRADE ORGANIZATION (2023): *Consolidated Tariff Schedules (CTS) Database*, Geneva: WTO. [27]