Optimal Trade and Industrial Policies in the Global Economy: A Gradient-Based Learning Approach[*]

Zi Wang[†1], Xingcheng Xu[‡2], Yanqing Yang[§3], and Xiaodong Zhu[¶4]

[1]Hong Kong Baptist University
[2]Shanghai Artificial Intelligence Laboratory
[3]Fudan University
[4]The University of Hong Kong

February 19, 2024

Abstract

We propose a gradient-based learning framework for solving optimal (unilaterally and mutually) policies in multi-country-multi-sector general equilibrium models. It is computationally efficient in characterizing global competitions in which multiple major economies compete with each other using high-dimensional policy tools. We apply this framework for solving optimal tariffs and industrial policies in the presence of sectoral scale economies, as in Lashkaripour and Lugovskyy (2023). We find that (i) the "internal cooperation" assumption for the analytical optimal policies in Lashkaripour and Lugovskyy (2023) does not hold in our fully optimal policies; (ii) Nash tariffs and industrial subsidies vary substantially across sectors and countries and increase with respect to sectoral scale economies; (iii) Nash tariffs lead to significant welfare losses in all economies, whereas Nash industrial policies could result in welfare gains for all economies if there is no distortion in specifying and implementing industrial subsidies; and (iv) the global social planner would choose close-to-zero tariffs and substantial industrial subsidies that increase with sectoral scale economies to maximize the global welfare.

JEL classification: F12; F51; C61; C63

Keywords: Tariff Wars; Industrial Policies; Nash Equilibrium; Gradient-based Learning; Best-response with Gradients; Machine Learning

[*]The ideas expressed herein are solely our own and do not represent the views or opinions of any institution or organization.

[†]wangzi@hkbu.edu.hk

[‡]xingcheng.xu18@gmail.com

[§]yanqingyang@fudan.edu.cn

[¶]xdzhu@hku.hk

1 Introduction

In recent years, economic conflicts have emerged among major world economies, making it important to understand the incentives and consequences of these conflicts for both economic studies and policy making. In this agenda, two trends have become increasingly pronounced. First, in addition to trade policies, countries rely on a combination of policies, including but not limited to industrial policies and innovation policies, to compete with each other. Second, in many real-world scenarios, countries act as if they are playing non-cooperative games. Therefore, when thinking about one country's optimal policies, we must take into account other countries' retaliation. Given these trends in practice, to evaluate the costs of economic conflicts and the benefits of potential economic cooperation, it is essential to solve for Nash equilibria in multi-country-multi-sector quantitative general equilibrium models that incorporate a combination of policies.

However, it is computationally challenging to solve for high-dimensional optimal policies and the corresponding Nash equilibria in multi-country-multi-sector quantitative general equilibrium models. There are three main challenges. First, conditional on a set of policies, equilibrium outcomes are determined by solving a high-dimensional nonlinear system. Second, we often consider a high-dimensional policy space consisting of continuous policy tools. Third, solving for Nash equilibria requires repeatedly solving for unilateral optimal policies in each country.

There are two main approaches in the literature to address the challenges above. First, Judd and Su (2012) develop a constrained optimization approach that can compute optimal policies taking equilibrium conditions as constraints. Combined with the standard Newton-based algorithms for constrained optimization, this approach has speed advantage since it does not repeatedly solve the equilibrium system at each guess of policies. However, if the dimension of equilibrium conditions is too high, this approach will lose its speed advantage due to expanding dimensionality of the solution space. Based on Judd and Su (2012), Ossa (2014) takes four days to compute optimal tariffs (unilaterally and mutually) in a GE model with 7 economies and 33 industries but without input-output linkages, using "a high-end desktop computer and standard MATLAB software". It is computationally infeasible to include more economies, industries, and policies as well as incorporate input-output linkages.

Second, Lashkaripour and Lugovskyy (2023) and Bartelme, Costinot, Donaldson, and Rodriguez-Clare (2021) derive sufficient statistics of optimal policies without solving the high-dimensional optimization problem. This approach is appealing since it does not require computation power but can get optimal policies by calculating several simple and intuitive sufficient statistics. However, this approach has to make several simplification assumptions to get the exact form of sufficient statistics. For example, the analytical forms of unilaterally optimal policies in Lashkaripour and Lugovskyy (2023) rely on the "*internal cooperation*" assumption: they assume that the relative wages in other economies remain unchanged under one economy's optimal policies. In real-world scenarios where such assumptions are unlikely to hold, we still need to solve the high-dimensional optimization problem.

In this paper, we propose a gradient-based learning framework for solving optimal (unilaterally and mutually) policies in multi-country-multi-sector general equilibrium models. Unlike Judd and Su (2012), we do not include equilibrium outcomes into the solution space and take equilibrium conditions as constraints. As a result, we avoid the exploding dimensionality of our solution space. Instead, we solve the equilibrium system by iteration at each guess of policies and focus on improving the efficiency of policy updating. To this end, we employ the widely-used gradient-based learning algorithm that can efficiently update the parameters based on the gradient of the objective function with respect to the parameters. Moreover, we implement the gradient-based algorithm using a machine learning (ML) open-source framework, PyTorch. This framework provides automatic differentiation, which is crucial for efficiently computing gradients.

We utilize our gradient-based learning framework to solve optimal trade and industrial policies in a multi-country-multi-sector GE model with trade, input-output linkages, and sectoral scale economies. Our model is an generalization of the model in Lashkaripour and Lugovskyy (2023). We consider the world with 6 major economies plus the rest of the world (ROW) and 44 sectors (22 of them are tradable sectors). Using one CPU in a laptop, we compute Nash tariffs and industrial policies among 7 regions (6 major economies + ROW) in this model by about 10 hours. Our solution is shown to be precise and our convergence process is stable. We further show that parallelization can further improve the efficiency of our method.

Our quantitative results shed light on how countries would behave in a non-cooperative game of tariffs and industrial subsidies. First, we find that the "*international cooperation*" assumption in Lashkaripour and Lugovskyy (2023) does not hold in our fully optimal policies, *i.e.* the optimal policies in one economy do affect the relative wages in other economies. This indicates that the analytical optimal policies in Lashkaripour and Lugovskyy (2023) are not fully optimal and thereby numerical methods are essential for solving the fully optimal policies.

Second, we find that Nash tariffs and industrial subsidies vary substantially across sectors and countries and increase with respect to sectoral scale economies. This finding confirms the incentives for a country to manipulate trade and industrial policies to gain from terms-of-trade and home market effects. This result is similar to Theorem 1 in Lashkaripour and Lugovskyy (2023) but, as we have mentioned, their analytical results are based on the "*international cooperation*" assumption that does not hold in our fully optimal policies.

Third, we find that Nash tariffs reduce welfare in all participating countries, whereas Nash subsidies lead to substantial welfare gains in all countries if they are efficiently implemented. Intuitively, subsidizing sectors with strong scale economies would lower the exporting prices and thus benefit foreign countries. This result indicates that in the era of global competition different forms of competition have very different welfare implications. However, welfare gains from industrial subsidies tend to rely on proper specification and implementation. We further show that given other countries set their industrial subsidies optimally, a country is likely to lose if its industrial policies deviate from the optimal levels.

Finally, we consider the global social planner choosing trade and industrial policies in all economies to maximize the global welfare.[1] We find that the global optimal policies consist of close-to-zero tariffs and substantial industrial subsidies that increase with sectoral scale economies. These policies would result in considerable welfare gains in most of the economies.

Related Literature. This paper is the first attempt to use gradient-based learning algorithm and machine-learning implementation to solve for optimal policies in multi-country-multi-sector

[1] We measure the global welfare as the population-weighted average real income across economies.

GE models. Our computation framework outperforms previous attempts, *e.g.* Ossa (2014), based on Judd and Su (2012) in terms of efficiency. This framework is widely applicable for evaluating optimal policies and policy competition in quantitative trade and spatial models.

This paper also relates to recent quantitative analysis on trade and industrial policies. Ju, Ma, Wang, and Zhu (2024) characterize optimal trade and industrial policies, in the U.S. and China, that are uniform across certain industries. This paper, in contrast, computes the fully optimal policies. Lashkaripour and Lugovskyy (2023) and Bartelme et al. (2021) characterize optimal trade and industrial policies using sufficient statistic approach. But as we have discussed, these theoretical explorations require simplification assumptions that may not hold in practice.

The rest of the paper is organized as follows. In Section 2, we propose the general problem and solution framework. We then apply our general framework to solving optimal trade and industrial policies in a multi-country-multi-sector GE model. In Section 3, we set the model, characterize the equilibrium, and calibrate the model. In Section 4, we compute optimal trade and industrial policies and characterize their features. We conclude in Section 5.

2 The General Problem and Solution Framework

2.1 Optimal Policies and Nash equilibria

Consider N countries in the world indexed by $i = 1,2,\ldots,N$. We consider a vector of equilibrium outcomes, x, in the world, and a vector of policies a_i in each country i. The collection of policies in all countries but i is denoted as a_{-i}. Let $a \equiv (a_i, a_{-i}) \in \mathcal{A}$. The equilibrium outcomes x are determined by the following (nonlinear) equilibrium conditions:

$$G_i(x, a_i, a_{-i}) = 0, \quad \forall i = 1,2,\ldots,N. \tag{1}$$

The government in country i aims to maximize $W_i(a_i, a_{-i}; x)$, an objective depending on equilibrium outcomes and (potentially) policies in all countries. More formally, given a_{-i}, country i

solves:

$$\begin{aligned} &\max_{(a_i;x)} W_i(a_i, a_{-i}; x) \\ &\text{s.t. } G_i(x, a_i, a_{-i}) = 0, \quad \forall i = 1, 2, \ldots, N. \end{aligned} \tag{2}$$

The set of solutions to the problem above is denoted as $\mathbf{b}_i^*(a_{-i})$. Then a Nash equilibrium is a set of policy choices for all countries, $(a_i^*)_{i=1}^N$, such that for all $i = 1, 2, \ldots, N$

$$a_i^* \in \mathbf{b}_i^*(a_{-i}^*). \tag{3}$$

2.2 Best-Response Dynamics

To solve for the Nash equilibrium $(a_i^*)_{i=1}^N$, we transform the static problem into a system of differential equations that characterizes best-response dynamics:

$$\dot{a}_i = \mathbf{b}_i^*(a_{-i}) - a_i, \quad \forall i = 1, 2, \ldots, N, \quad \forall a \in \mathcal{A}^N, \tag{4}$$

where we assume that $\mathbf{b}_i^*(a_{-i})$ has a unique element for all $a_{-i} \in \mathcal{A}_{-i}$.

A discrete-time scheme for the best-response dynamics is as follows:

$$a_i^{t+1} = \eta_t \mathbf{b}_i^*(a_{-i}^t) + (1 - \eta_t) a_i^t, \quad \forall i = 1, 2, \ldots, N, \quad \forall a \in \mathcal{A}^N, \tag{5}$$

where $\eta_t \in (0, 1]$.

The best-response dynamics capture the idea of players adjusting their strategies iteratively in response to the strategies chosen by others. Researchers have extensively studied these dynamics in different game-theoretic contexts to gain insights into stability, convergence rates, computational complexity, and other important properties of strategic decision-making. Matsui (1992) explores stability concepts derived from best response dynamics and their equivalence to static concepts. It shows the existence of set-valued versions of these concepts and provides examples illustrating their usefulness in analyzing forward induction and preplay communication. Hofbauer and

Sorin (2006) studies best response dynamics in continuous time for continuous concave-convex zero-sum games. They prove convergence of the dynamics to the set of saddle points, providing a dynamical proof of the minmax theorem. Barron, Goebel, and Jensen (2010) extends a convergence result for best response dynamics in continuous concave-convex zero-sum games to nonconcave, nonconvex payoff functions. Leslie, Perkins, and Xu (2020) analyzes three learning dynamics for two-player zero-sum discounted-payoff stochastic games. They prove convergence of a continuous-time best-response dynamic and a fictitious-play-like process to the set of Nash equilibrium strategies. Dindoš and Mezzetti (2006) considers n-person games with quasi-concave payoffs and analyzes the better-reply dynamics. Durand and Gaujal (2016) analyzes the worst-case and average execution time of the best-response algorithm for computing pure Nash equilibria in finite potential games. Swenson, Murray, and Kar (2018) studies the convergence properties of best-response dynamics in potential games. It shows that in most potential games, the dynamics have a unique solution and converge to pure-strategy Nash equilibria exponentially. Other notable books on best response dynamics include, for example, Hofbauer and Sigmund (1998), Fudenberg and Levine (1998), and Nisan, Roughgarden, Tardos, and Vazirani (2007).

2.3 Gradient-based Learning

To find the best response $\mathbf{b}_i^*\left(a_{-i}^t\right)$ at each iteration t of the best-response dynamics (5), gradient-based learning algorithms can be used.

Gradient ascent (descent) is an iterative optimization method that aims to find the local maximum (minimum) of a function. It updates the parameters based on the gradient of the objective function with respect to the parameters. Since $\mathbf{b}_i^*\left(a_{-i}^t\right) = \arg\max_{a_i \in \mathcal{A}} W_i(a_i, a_{-i}^t; x)$, finding the maximum of $W_i(a_i, a_{-i}^t; x)$ is equivalent to finding the minimum of $O_i(a_i) \equiv -W_i(a_i, a_{-i}^t; x)$. The general update rule is:

$$a_i^{\tau+1} = a_i^{\tau} - \gamma\omega(\nabla O_i|_{a_i=a_i^{\tau}}), \tag{6}$$

where a_i^τ is the value at timestep τ, γ is the learning rate (step size), $\nabla O_i|_{a_i=a_i^\tau}$ is the gradient of the objective function with respect to a_i at timestep τ, and $\omega(\cdot)$ is a functional of the gradient.

The gradient provides information about the direction and magnitude of descent. Through the iterative updates in (6), the values are driven towards a minimum of the objective function $O_i(a_i)$, i.e. the maximum of $W_i(a_i, a_{-i}^t; x)$.

Gradient-based learning is studied in multiagent learning algorithms. Several papers have investigated the convergence properties of these algorithms in different game settings. Singh, Kearns, and Mansour (2000) analyzes agents' strategy adaptation through gradient ascent in two-player, two-action, iterated general-sum games. It shows convergence to a Nash equilibrium or convergence of average payoffs to Nash equilibrium payoffs. Chasnov, Ratliff, Calderone, Mazumdar, and Burden (2019) focuses on convergence guarantees for gradient-based learning algorithms in non-cooperative multi-agent settings. The analysis utilizes game Hessian singular values to obtain finite-time convergence bounds to an ε-differential Nash equilibrium. Raghunathan, Cherian, and Jha (2019) introduces the Gradient-based Nikaido-Isoda (GNI) function, which efficiently computes equilibrium. It shows that gradient descent converges to a stationary Nash point and provides error bounds. Mazumdar, Ratliff, and Sastry (2020) presents a framework for competitive gradient-based learning and analyzes its effectiveness in general-sum and potential games by studying the limiting behavior and local Nash equilibria. Varma, Veetaseveera, Postoyan, and Morărescu (2021) proposes an algorithm for optimizing ordinal potential games using a "better response" approach, reducing computational time compared to traditional methods.

2.4 Automatic Differentiation for The Gradient Estimate

To implement gradient-based learning, automatic differentiation (see e.g. Baydin, Pearlmutter, Radul, and Siskind (2018)), commonly referred to as backpropagation in machine learning, is used to estimate gradients. This process, known as the reverse mode of automatic differentiation, is crucial for training neural network models. It allows for the calculation of gradients, which optimization algorithms then use to update the network parameters.

The process of reaching equilibrium outcomes, denoted as x, is governed by nonlinear equilibrium conditions $G_i(x, a_i, a_{-i}) = 0$ for all participants $i = 1, 2, \ldots, N$. This process defines an iterative approach to solving the system, transitioning through a series of intermediate variables from an initial value to the equilibrium outcomes:

$$x_0 \xrightarrow{(a_i, a_{-i})} x_1 \xrightarrow{(a_i, a_{-i})} \cdots \xrightarrow{(a_i, a_{-i})} x_T.$$

This sequence continues until the difference $|x_T - x_{T-1}|$ is less than a predetermined threshold ε.

Drawing parallels to neural networks, each iteration step to address the nonlinear equilibrium conditions can be viewed as a hidden layer in a neural network. The intermediate variables, x_t, act as nodes within this layer, contributing to the final equilibrium outcomes x_T. These nodes are interconnected through the iteration relation $G = (G_i)_{i=1}^N$, which dictates the equilibrium conditions. In this framework, policies a_i are analogous to parameters that optimize the welfare of country i, with welfare W_i serving as the loss function in a neural network. The gradient computation mirrors that of forward-backward propagation in neural networks, utilizing the chain rule:

$$\frac{\partial W_i}{\partial a_i} = \frac{\partial W_i}{\partial x_T} \frac{\partial x_T}{\partial x_{T-1}} \cdots \frac{\partial x_2}{\partial x_1} \frac{\partial x_1}{\partial a_i}.$$

This analogy underscores the similarity between solving nonlinear economic equilibrium systems and optimizing neural networks, as depicted in the Figure 1 illustrating the neural network analogy.

For practical implementation of the best-response dynamics with gradients, we implement the gradient-based algorithm using the machine learning (ML) open-source framework such as PyTorch (Paszke, Gross, Massa, Lerer, Bradbury, Chanan, Killeen, Lin, Gimelshein, Antiga et al. (2019)), TensorFlow (Abadi, Agarwal, Barham, Brevdo, Chen, Citro, Corrado, Davis, Dean, Devin et al. (2016)) or Google JAX (Frostig, Johnson, and Leary (2018)). These frameworks provide automatic differentiation, which allows efficient computation of gradients through forward and backward propagation. This feature is crucial for obtaining gradients with respect to the action in our problem efficiently. Additionally, ML frameworks offer support for various machine learning optimization methods, including SGD, ADAM, Adagrad, RMSprop, and more. This versatility

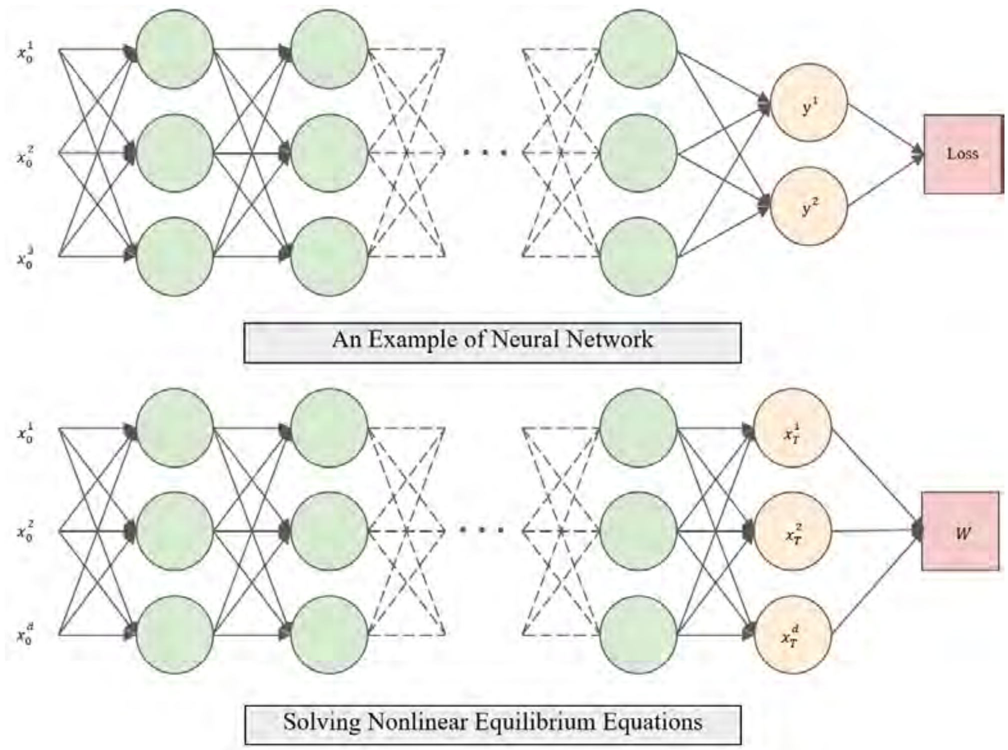

Figure 1: Analogy between nonlinear economic equilibrium systems and neural networks

allows us to choose the most appropriate gradient-based learning method for the specific task.

3 Model and Calibration

To shed light on what specific problems our gradient-based learning framework can solve, we present a multi-country multi-sector general equilibrium model with trade, input-output linkages, and sectoral scale economies and solve the optimal (unilaterally and mutually) trade and industrial policies in this model. This model, a la Bartelme et al. (2021), Lashkaripour and Lugovskyy (2023), and Ju et al. (2024), provides a laboratory for quantifying global impacts of tariffs and industrial subsidies. In this section, we first introduce the model and its equilibrium and then discuss the model's calibration. In the next section, we discuss computation and characterize the optimal policies obtained using our gradient-based learning framework.

3.1 Model Setup

In a world consisting of N countries denoted by i and n, there exists a population of L_i workers in each country i. J sectors, indexed by j and s, are also present. Workers cannot relocate between countries, but are free to move across sectors. Each sector j consists of a unit mass of varieties.

The representative consumer of country i has a two-tiered preference:

$$U_i = \sum_{j=1}^{J} \alpha_i^j \log \left[\left(\int_0^1 \left[C_i^j(\omega) \right]^{\frac{\sigma_j - 1}{\sigma_j}} d\omega \right)^{\frac{\sigma_j}{\sigma_j - 1}} \right] \tag{7}$$

where α_i^j is the share of sector j in country i and σ_j is the elasticity of substitution across varieties within sector j. We assume that each variety is produced under perfect competition using labor and composite intermediates.

Trade between countries is hindered by three costs. First, there is an iceberg trade cost $\tau_{in}^j \geq 1$ for shipping goods from country i to country n, where $\tau_{ii}^j = 1$. Second, an ad valorem tariff $t_{in}^j \geq 0$ is imposed by importing country n on goods j imported from country i, with $t_{ii}^j = 0$. Lastly, an ad valorem output tax $e_{in}^j \geq -1$ is imposed by production country i on goods j destined for country i. Notably, this tax takes the form of industrial subsidies when negative and uniform for all destination countries, i.e. $e_i^j = e_{in}^j$ for all n.

The production technology is characterized by the following unit cost function: the unit cost of variety ω of sector j in country i is $c_i^j(\omega) = \frac{c_i^j}{z_i^j(\omega)}$, where c_i^j is defined as

$$c_i^j = \frac{1}{\left(L_i^j\right)^{\psi_j}} w_i^{\beta_i^j} \left[\prod_{s=1}^{J} (P_i^s)^{\gamma_i^{sj}} \right]^{1-\beta_i^j}. \tag{8}$$

Here, P_i^s is the price index of good s in country i, L_i^j is the labor allocated to sector j of country i, and $\psi_j \geq 0$ characterizes the external economies of scale in sector j. The parameter γ_i^{sj} represents the share of good s in producing good j in sector j in country i, satisfying $\sum_{s=1}^{J} \gamma_i^{sj} = 1$. Moreover, w_i is the wage in country i and β_i^j represents the value-added share in sector j of country i.

Following Eaton and Kortum (2002), we assume that the Hicks-neutral productivity $z_i^j(\omega)$ is drawn independently from a Frechét distribution. The probability that $z_i^j(\omega) \leq z$ is given by

$$\mathbb{P}\left(z_i^j(\omega) \leq z\right) = \exp\left(-T_i^j z^{-\theta_j}\right) \tag{9}$$

where $z > 0$ and $\theta_j > \max\{\sigma_j - 1, 1\}$. Here, T_i^j characterizes the average productivity of sector j in country i and θ_j characterizes the dispersion of productivities in sector j.

3.2 Equilibrium

We proceed by aggregating the individual decisions and delivering equilibrium conditions. Based on the property of Frechét distribution and the ideal price index of CES preferences, the sectoral price index can be expressed as

$$P_n^j = \left(\sum_{i=1}^{N} T_i^j \left[c_i^j \tau_{in}^j (1+t_{in}^j)(1+e_{in}^j)\right]^{-\theta_j}\right)^{-\frac{1}{\theta_j}}. \text{[2]} \tag{10}$$

The expenditure share of country n on good j from country i is given by

$$\pi_{in}^j = \frac{X_{in}^j}{X_n^j} = \frac{T_i^j \left[c_i^j \tau_{in}^j (1+t_{in}^j)(1+e_{in}^j)\right]^{-\theta_j}}{(P_n^j)^{-\theta_j}}, \tag{11}$$

where X_{in}^j is the expenditure of country n on good j from country i and X_n^j is the total expenditure on good j in country n.

Sectoral employment satisfies

$$w_i L_i^j = \beta_i^j \sum_{n=1}^{N} \frac{X_{in}^j}{(1+t_{in}^j)(1+e_{in}^j)}. \tag{12}$$

[2] To save notations, we ignore the constant related to the Gamma function.

The wage is determined by labor market clearing

$$\sum_{j=1}^{J} L_i^j = L_i. \tag{13}$$

Assuming that export tariffs (if any) are collected before import tariffs and all tax revenues or subsidy expenditures are transferred to workers as lump-sum payments, we can express the total final income as

$$Y_i = w_i L_i + \sum_{j=1}^{J}\sum_{n=1}^{N} \frac{e_{in}^j}{1+e_{in}^j} X_{in}^j + \sum_{j=1}^{J}\sum_{k=1}^{N} \frac{t_{ki}^j}{(1+t_{ki}^j)(1+e_{ki}^j)} X_{ki}^j. \tag{14}$$

The sectoral expenditure can be expressed by

$$X_i^j = \alpha_i^j Y_i + \sum_{s=1}^{J} (1-\beta_i^s)\gamma_i^{js} \sum_{n=1}^{N} \frac{X_{in}^s}{(1+t_{in}^s)(1+e_{in}^s)}. \tag{15}$$

Definition 1 (Equilibrium)**.** *Given parameters* $\left\{\theta_j, \psi_j, \alpha_i^j, \beta_i^j, \gamma_i^{sj}, L_i, e_{in}^j, t_{in}^j, T_i^j, \tau_{in}^j\right\}$*, the equilibrium consists of* $\left(w_i, L_i^j, P_i^j, X_i^j\right)$ *such that:*

1. *Equation* (10) *defines the price indices* P_n^j.
2. *Equation* (12) *determines the allocation of labor across sectors.*
3. *Equation* (13) *sets the wage level.*
4. *Equation* (15) *ensures sectoral market clearing.*

The welfare in country n can be measured by its real income, $W_n \equiv \frac{Y_n}{P_n}$, where the aggregate price index for final consumption goods can be expressed as

$$P_n = \prod_{j=1}^{J} \left(P_n^j\right)^{\alpha_n^j}. \tag{16}$$

The equilibrium system in Definition 1 consists $3NJ+N$ nonlinear equations that relate to $3NJ+N$ unknown variables, which can be solved when given a numeraire. However, this system poses a challenge as it depends on a complex set of parameters, including T_i^j, τ_{in}^j, which are difficult to calibrate. To address this issue, we use the "exact-hat" algebra developed by Dekle, Eaton, and Kortum (2008) to compute changes in equilibrium outcomes relative to changes in exogenous shocks. We denote the value of any variable after change as Z' and use the notation $\widehat{Z}=\frac{Z'}{Z}$.

Given the values of $\left(\alpha_i^j,\beta_i^j,\gamma_i^{sj};\psi_j,\theta_j\right)$, as well as data on $\left(X_{in}^j,t_{in}^j,e_{in}^j\right)$, we can compute the changes in equilibrium outcomes, $\left(\widehat{w}_i^j,\widehat{L}_i^j,\widehat{P}_i^j,\widehat{X}_i^j\right)$, by solving a system of $3NJ+N$ nonlinear equations. The details are presented in Appendix A.1.

The optimization problem faced by each country i involves maximizing its welfare by adjusting its import tariffs and industry subsidies:

$$\begin{aligned} &\max_{\left\{t_{in}^j,e_{ni}^j;w_i,L_i^j,P_i^j,X_i^j\right\}_{i,j}} W_n \equiv \frac{Y_n}{P_n}, \quad \forall n=1,2,\cdots,N \\ &\text{s.t. Equation (10), (12),(13), and (15)} \end{aligned} \tag{17}$$

In practice, we express this problem in terms of relative changes in equilibrium outcomes. The details are also presented in Appendix A.1.

3.3 Model's Calibration

To solve for the changes in equilibrium outcomes with respect to policy changes, we need data on bilateral trade shares $\left(\pi_{in}^j\right)$, sectoral consumption shares $\left(\alpha_i^j\right)$, sectoral value-added shares $\left(\beta_i^j\right)$, sectoral expenditure $\left(X_n^j\right)$, input expenditure shares γ_i^{js}, and initial policies $\left(t_{in}^j,e_{in}^j\right)$. We also need the values of parameters (ψ_j,θ_j).

We consider 6 major economies (US, China, Japan, EU, Brazil, and India) and the rest of world (ROW).[3] We utilize the OECD Inter-Country Input-Output database (ICIO) for 2017 to extract

[3] European Union (EU) includes 28 countries including the UK.

internationally comparable data on country-sectoral production, value-added, bilateral trade flows, and input-output linkages. The ICIO table includes 22 tradable sectors and 22 nontradables.[4] We get the initial import tariffs from the World Integrated Trade System (WITS) for 2017 and assume that initially $e_{in}^{j}=0$ for all (i,n,j).

We calibrate (ψ_j,θ_j) from Lashkaripour and Lugovskyy (2023). The calibrated values of (ψ_j,θ_j) in Appendix A.2. Lashkaripour and Lugovskyy (2023) recover ψ_j from the effects of variation in sector size on equilibrium quantities, exploiting variation in countries' population and preferences as instruments.

4 Quantitative Application

4.1 Best-Response Dynamics in Practice

Let the action taken by country n be denoted as $a_n=(t_{in}^{j},e_{in}^{j})_{i\in[N],j\in[J]}\in\mathcal{A}\subset\mathbb{R}^{(N-1)\times J+J}$, where $[N]$ represents the set of all countries and $[J]$ represents the set of industries. The action space, denoted as $\mathcal{A}$, is assumed to be identical for all countries and defined as $\mathcal{A}=[0,1]^{(N-1)\times J}\times[-1,0]^{J}$.

Each country has the ability to impose a tariff on each industry for all other countries except itself, resulting in $(N-1)\times J$ possible tariff components. Additionally, each country can offer a subsidy for each industry, resulting in J possible subsidy components. Hence, there are a total of $N\times J$ possible action components for each country.

We define the relative change in welfare for country n as:

$$\widehat{\mathrm{W}}_n=F_n\left(a_n,a_{-n};a^{(0)},X^{(0)}\right),\ n\in[N]. \tag{18}$$

[4] The ICIO has 45 industries. We disregard the last one, which is "Activities of households as employers; undifferentiated goods- and services-producing activities of households for own use", because it has many zeros. See OECD. (2021) OECD Inter-Country Input-Output Database, http://oe.cd/icio.

Here, a_{-n} represents the actions taken by all countries except country n, $a^{(0)}$ denotes the initial action profile of all countries, and $X^{(0)} = (X_{in}^{j})$ represents the initial state of the economic environment.

To compute a Nash equilibrium for equation (18), we employ the method of best-response dynamics (refer to e.g. Matsui (1992), Hofbauer and Sigmund (1998), Fudenberg and Levine (1998), Nisan et al. (2007), Hofbauer and Sorin (2006), Swenson et al. (2018), and Heinrich, Jang, Mungo, Pangallo, Scott, Tarbush, and Wiese (2023)), which is an iterative algorithm utilized to find a Nash equilibrium in a non-cooperative game. The algorithm assumes that each player selects a strategy from their available options and, in each subsequent round, players adjust their strategies based on the best response to the strategies chosen by other players in the previous round. The algorithm converges to a stable state where no player can unilaterally improve their payoffs by deviating from their current strategies.

In Heinrich et al. (2023), the authors explored two specific playing sequences for determining the order in which players update their actions during the game. One approach is a fixed cyclic order, where players take turns adjusting their strategies in a predetermined order. This playing sequence assumes that players move sequentially, and the order of play remains constant throughout the game, often referred to as the clockwork playing sequence. Another approach is the random playing sequence, where, at each time step, a player is selected uniformly at random from among all players to update their strategy.

A random shuffle playing sequence combines the characteristics of both playing sequences. This sequence introduces stochasticity into the playing order, which helps to reduce positional advantages that may arise in a clockwork playing sequence. By shuffling the playing order randomly at the beginning of each round, this ensures that each player has an equal opportunity to adjust their actions in each round. The best-response dynamics with random shuffle playing sequence is described in Algorithm 1.

Algorithm 1 Best-response dynamics with random shuffle playing sequence

1. Give the initial action profile of all countries $a^{(0)}$, and the initial state $X^{(0)} = (X_{in}^{j})$ of the economic environment.

2. Draw a uniform starting action profile $(a_n^0)_{n\in[N]}$ at random from $\mathcal{A}^N$, or start from a chosen point in $\mathcal{A}^N$.

3. Repeat until converge:

   - random shuffle playing sequence for each round
   - for the n-th player $s(n)$ in the shuffled sequence s

     (a) best response search for the player $s(n)$ in sequence

$$a_{s(n)}^{t} = \arg\max_{a\in\mathcal{A}} F_{s(n)}\left(a, a_{-s(n)}^{t-1}; a^{(0)}, X^{(0)}\right) \tag{19}$$

     (b) set $a_{-s(n)}^{t} = a_{-s(n)}^{t-1}$

4.2 Gradient-based Learning

To optimize the best response for player n, denoted as $a_n = \arg\max_{a\in\mathcal{A}} F_n(a, a_{-n}; a^{(0)}, X^{(0)})$, gradient-based learning algorithms are commonly employed, see e.g. Singh et al. (2000), Chasnov et al. (2019), Raghunathan et al. (2019), Mazumdar et al. (2020), and Varma et al. (2021).

Gradient descent is an iterative optimization algorithm that aims to find the local minimum of a function. It updates the parameters based on the negative gradient of the objective function with respect to the parameters. Denote the objective function as $L_n(a) = -F_n(a, a_{-n}; a^{(0)}, X^{(0)})$. The update step in gradient descent is mathematically represented as

$$a^{t+1} = a^t - \gamma \nabla L_n|_{a=a^t}, \tag{20}$$

where a^t is the parameter value at iteration t, γ is the learning rate (step size), and $\nabla L_n|_{a=a^t}$ represents the gradient of the objective function with respect to a at iteration t.

ADAM (Adaptive Moment Estimation, see Kingma and Ba (2014)) is an extension of gradient descent that combines adaptive learning rates and momentum. It utilizes the first and second moments of the gradients to adaptively adjust the learning rate for each parameter. ADAM is widely employed in machine learning optimization, such as large language models (LLMs) GPT3 (175B) with 175 billions parameters, BLOOM (176B), MT-NLG (530B), Gopher (280B), ERNIE 3.0 Titan (260B), and so on (see e.g. the survey Zhao, Zhou, Li, Tang, Wang, Hou, Min, Zhang, Zhang, Dong et al. (2023)).

The update step in ADAM involves the following steps:

1. Compute the gradient of the objective function with respect to the parameters: $\nabla L_n|_{a=a^t}$.

2. Calculate the first moment estimate of the gradients, $m_t = \beta_1 m_{t-1} + (1-\beta_1)\nabla L_n|_{a=a^t}$, where β_1 is the first moment decay rate and $m_0 = 0$.

3. Calculate the second moment estimate of the gradients $v_t = \beta_2 v_{t-1} + (1-\beta_2)(\nabla L_n|_{a=a^t})^{\odot 2}$, where β_2 is the second moment decay rate, $v_0 = 0$, and $(\nabla L_n|_{a=a^t})^{\odot 2}$ is the element-wise square.

4. Bias-correct the first and second moment estimates, $\hat{m}_t = \frac{m_t}{1-\beta_1^t}$, $\hat{v}_t = \frac{v_t}{1-\beta_2^t}$, where β_1^t and β_2^t are values β_1 and β_2 to the power t.

5. Update the parameters using the bias-corrected estimates:

$$a^{t+1} = a^t - \gamma \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \varepsilon}, \tag{21}$$

   where ε is a small constant to avoid division by zero.

In the above steps, a^t represents the parameter value at iteration t, γ is the learning rate, β_1 and β_2 are the decay rates for the first and second moments, and m_t and v_t are the first and second moment estimates at iteration t, respectively. For our best-response search in each round, we utilize the ADAM optimization method, which offers faster convergence compared to naive gradient descent.

Figure 2 illustrates the entire calculation process, which is a summary of all steps above.

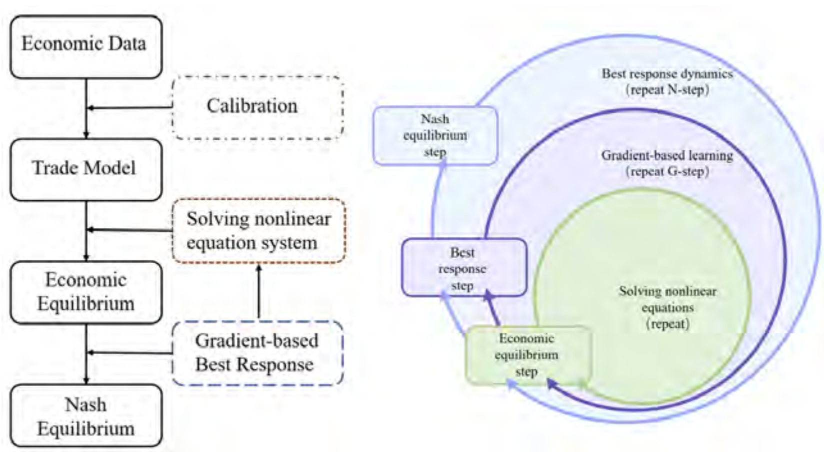

Figure 2: The Flow Chart of the Calculation

4.3 Machine Learning Implementation and Computational Performances

Our stochastic best-response dynamics is implemented using the PyTorch machine learning open-source framework (see Paszke et al. (2019)). PyTorch allows us to take advantage of automatic differentiation, enabling us to efficiently compute gradients through forward and backward propagation. This feature is crucial for obtaining gradients with respect to the action in our problem effectively.

Additionally, the ML framework offers support for various machine learning optimization methods, including SGD, ADAM, Adagrad, and RMSprop. This versatility allows us to choose the most appropriate gradient-based learning method for our specific task.

Due to the capabilities of the PyTorch framework, our stochastic best-response dynamic program is easily implemented. The combination of PyTorch's flexibility and efficient optimization methods contributes to the rapid convergence of our stochastic best-response dynamics. It's worth mentioning that other machine learning frameworks such as TensorFlow (Abadi et al. (2016)) or Google JAX (Frostig et al. (2018)) could also be utilized for similar purposes.

Our baseline exercise involves computing Nash tariffs and industrial subsidies for all 7 economies in our calibrated model (US, China, Japan, EU, Brazil, India, and the rest of the world). We refer this baseline scenario as "**global dual policy competition**". The computation statistics of our

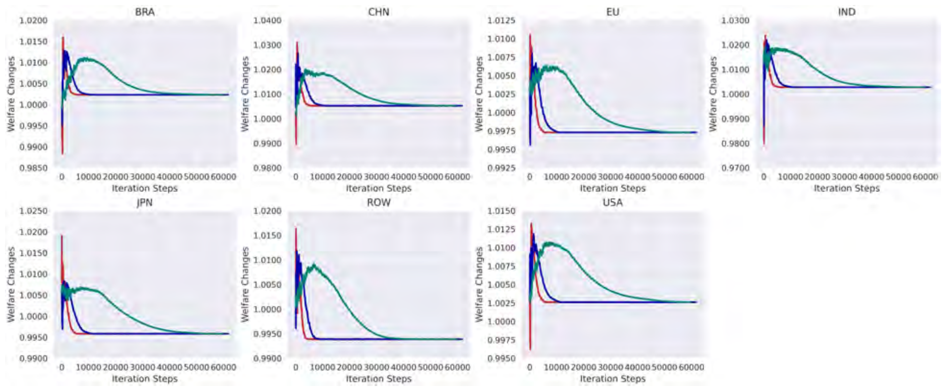

Figure 3: Iteration Curve for Nash Equilibrium of Global Dual Policy Competition with Scale Economies

framework in this baseline scenario is summarized in Table 1.

Table 1: Computation statistics of Solving for Global Dual Policy Competition

Information	**Value**
ML Framework	PyTorch
Device	CPU
Num. epochs	~ 20
Iterations per player/epoch	~ 50
Num. players	7
Playing sequence	Random Shuffle
Optimizer	ADAM
Anneal learning rate	False
Learning rate	$10^{-4} \sim 10^{-3}$
Is clipping grad	{True, False}
Max grad norm	10.0
Max Computation Time	5h

Figure 3 illustrates the convergence path of our algorithm in each country. It shows that even under relatively strong scale economies, our algorithm converges steadily after about 10,000 iterations.

Finally, we examine the optimality of our solutions. Figure 4 illustrates landscape near Nash equilibrium for global dual policy competition, suggesting the local optimality of our solutions.

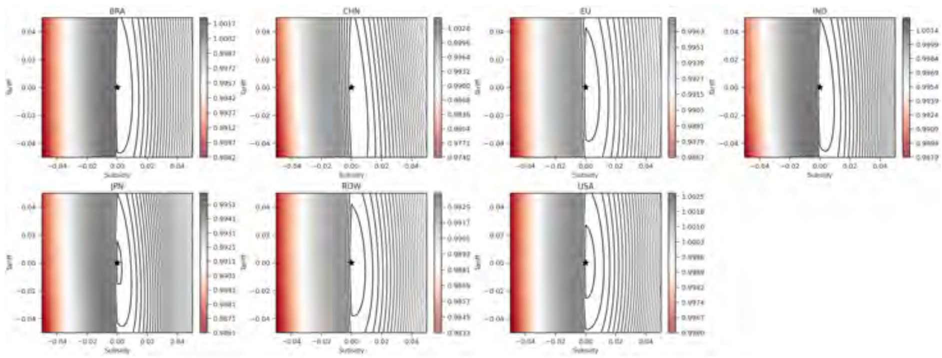

Figure 4: Landscape Near Nash Equilibrium for Global Dual Policy Competition with Scale Economies

4.4 Quantitative Results

4.4.1 Unilaterally Optimal Policies

In this section, we solve for the unilaterally optimal policies in China. More specifically, We consider China's optimal (i) industrial subsidies, (ii) import tariffs, (iii) dual policies, *i.e.* industrial subsidies + import tariffs, and (iv) full policies, *i.e.* industrial subsidies + import and export tariffs.

The welfare consequences of these policies are summarized in Table 2. We find that China gains significantly from unilaterally optimal policies, in particularly from industrial subsidies. The Chinese welfare increases by 3.71% under optimal industrial subsidies and by 2.22% under optimal import tariffs. Moreover, the Chinese welfare gain from the optimal dual policies is 5.78%, slightly lower than the sum of the gain from industrial subsidies and the gain from import tariffs (3.71% + 2.22%). This result suggests that in our application industrial subsidies and import tariffs are substitutes for welfare improvement. Finally, if three policies, industrial subsidies, import tariffs, and export tariffs, are optimally implemented, the Chinese welfare would increase by 6.30%. This numerical result indicates substantial incentives for a large country like China to manipulate its trade and industrial policies.

In the meanwhile, China's unilaterally optimal policies generally result in welfare losses in other major economies. Japan suffers most from China's optimal policies, reflecting the tough competition between China and Japan in many industries. In contrast, the United States only loses modestly under China's optimal policies.

Table 2: Welfare Changes Under China Unilateral Optimization

	China (Δ%)			
	Subsidy	Import Tariff	Dual	Full
	(1)	(2)	(3)	(4)
China	3.71	2.22	5.78	6.30
United States	-0.09	0.08	-0.03	-0.16
European Union	-0.17	-0.11	-0.28	-0.26
Japan	-0.36	-0.42	-0.67	-1.07
India	-0.01	1.20	-0.47	-0.63
Brazil	0.17	-0.56	-0.44	-0.36
Rest of the World	-0.17	-0.93	-1.19	-1.32

We then characterize the Chinese optimal dual policies. The left panel of Figure 5 suggests that the optimal import tariffs vary substantially across importing countries and sectors, which is mainly due to the sectoral heterogeneity in scale economies. Lashkaripour and Lugovskyy (2023) show that in the absence of scale economies the optimal import tariffs are uniform across importing countries and sectors.[5]

The middle and right panels of Figure 5 suggest that the unilaterally optimal import tariffs and industrial subsidies are strongly increasing with respect to sectoral scale elasticities, consistent with Theorem 1 in Lashkaripour and Lugovskyy (2023).[6]

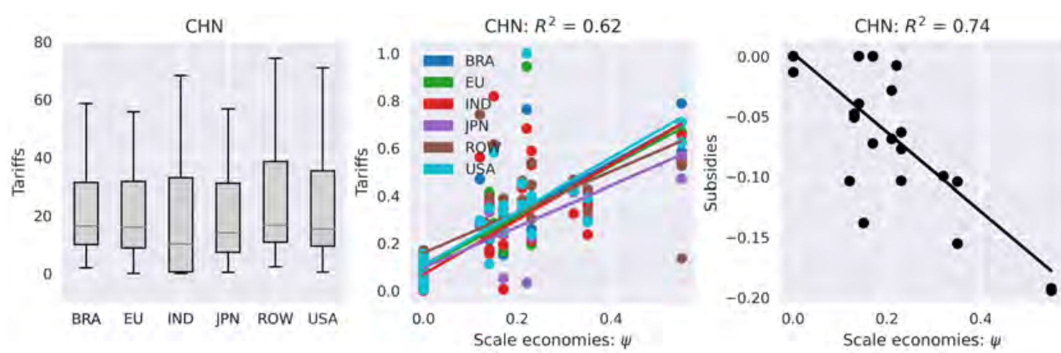

Figure 5: Import Tariffs and Industrial Subsidies in China's Unilaterally Optimal Dual Policies

[5]In the absence of scale economies, we find that the optimal import tariffs are almost uniform across sectors. Please see Appendix B.2 for the details.

[6]Lashkaripour and Lugovskyy (2023) primarily consider the optimal full policies. In Appendix B.3, we characterize the Chinese optimal full policies. The main patterns in Figure 5 preserve under the optimal full policies.

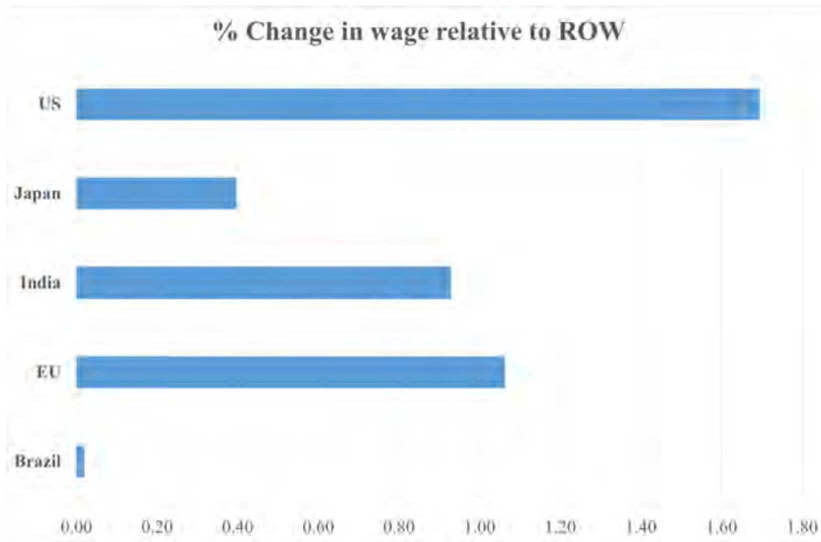

Figure 6: Changes in Relative Wage under China's Optimal Full Policies (in percentage)
(Notes: Full policies refer to industrial subsidies, import tariffs, and export tariffs. We normalize the change in nominal wage in the rest of the world (ROW) as 0.)

We have mentioned that one of the most restricted assumptions made by Lashkaripour and Lugovskyy (2023) in analytically characterizing the optimal policies is the other countries' internal cooperation, *i.e.* they require the relative wages among other countries remain unchanged under the optimal policies. In contrast, the optimal full policies in our numerical exercises are not subject to this "internal cooperation" restriction.

Figure 6 illustrates changes in relative wages under China's full policies. It suggests that the "internal cooperation" assumption made by Lashkaripour and Lugovskyy (2023) does not hold in our fully optimized solution: the relative wages among other major economies do change under China's optimal policies. The changes in relative wage are moderate here because the economies considered in our numerical exercises are very large-6 major economies plus the rest of the world. In this case, the extraterritorial terms-of-trade effects are small. In a world with more, smaller, and more heterogeneous economies, the extraterritorial terms-of-trade effects of optimal policies would be larger, resulting in larger gaps between the theoretical optimal policies in Lashkaripour and Lugovskyy (2023) and the fully optimal policies in our numerical exercises.

4.4.2 Nash Policies

To understand policy competitions among major economies, we utilize our model and gradient-based learning framework to characterize Nash equilibria in trade and industrial policy competitions. In particular, We characterize global dual policy competition in which every country chooses import tariffs and industrial subsidies to maximize its own welfare.

Figure 7 shows that Nash subsidies increase with scale economies. Intuitively, a country can address misallocation across sectors led by external economies of scale by subsidizing higher-return-to-scale sectors. This intuition is consistent with analytical results developed by Lashkaripour and Lugovskyy (2023) under "internal cooperation" assumptions.

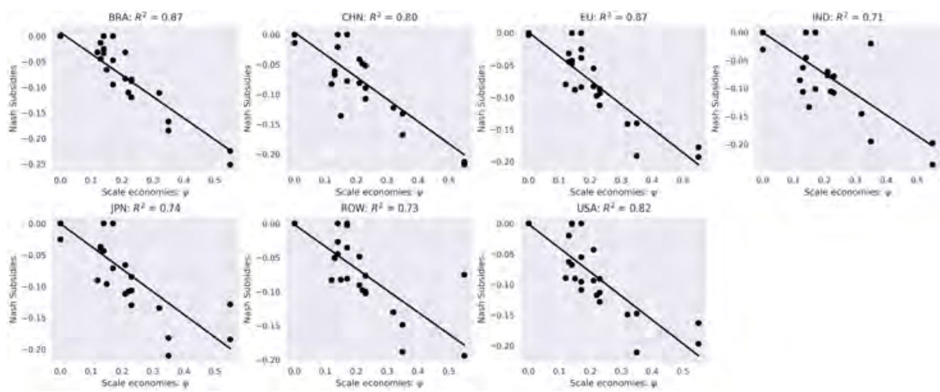

Figure 7: Nash Subsidies for Global Dual Policy Competition with Scale Economies

How are the optimal industrial policies and import tariffs correlated with each other under dual policy competition? Figure 8 suggests that with scale economies, Nash tariffs are positively correlated with Nash subsidies, indicating that with scale economies, two policy instruments, import tariffs and industrial subsidies, are close to complements in global dual policy competition.

What drives this complementarity between the optimal import tariffs and industrial policies? It is possible that this complementarity is due to the negative correlation between trade and scale economies in the calibration of Lashkaripour and Lugovskyy (2023). To verify this possibility, we rearrange the scale economies so that they are in the same rank across sectors with trade elasticities.[7] Figure 9 suggests that in this case industrial and trade policies are substitutes.

[7] We multiply ψ_j in Lashkaripour and Lugovskyy (2023) by 0.3 to avoid that $\theta_j \psi_j > 1$.

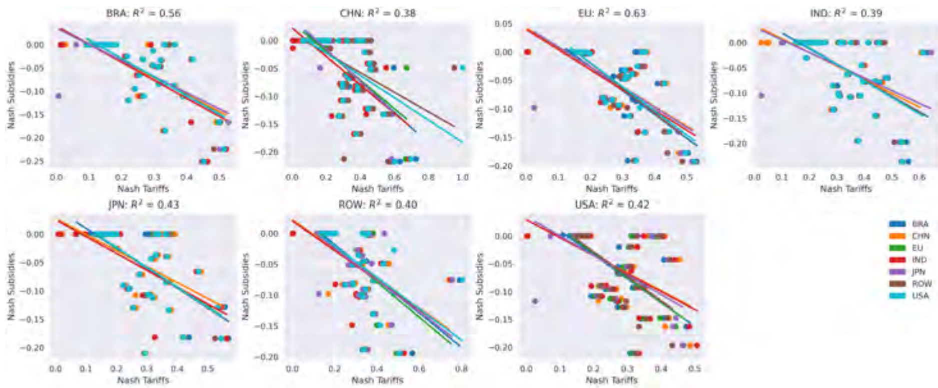

Figure 8: Nash Tariffs and Subsidies for Global Dual Policy Competition with Scale Economies

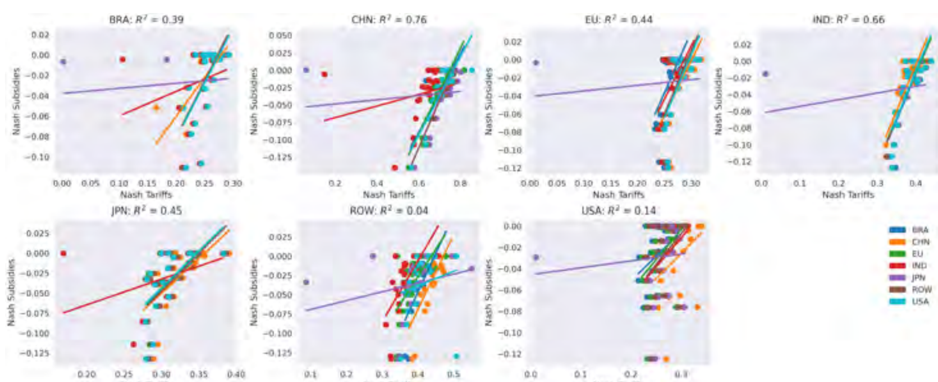

Figure 9: Nash Tariffs and Subsidies for Global Dual Policy Competition with Scale Economies under Positive Correlation

We proceed by investigating the welfare effects of global competition.

Table 3: Welfare Changes at Nash Equilibrium with Scale Economies

	China and US (Δ%)			**World** (Δ%)			
	Subsidy	Tariff	Dual	Subsidy	Subsidy-Uni	Tariff	Dual
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
United States	1.28	-0.07	1.23	1.31	1.31	-1.01	0.26
China	3.31	-0.33	3.17	2.42	1.80	-2.56	0.52
European Union	-0.51	0.02	-0.49	1.03	0.53	-1.80	-0.27
Japan	-0.61	0.03	-0.57	1.14	0.30	-2.21	-0.42
India	-0.50	0.01	-0.51	2.56	0.60	-1.84	0.27
Brazil	-0.05	0.00	-0.05	1.97	1.71	-2.08	0.23
Rest of the World	-0.44	0.05	-0.41	1.47	1.76	-2.33	-0.62

Note: "Subsidy" refers cases where players can adjust their industry subsidies only. "Subsidy-uni" refers to the cases where each player can only choose a uniform subsidy rate for all manufacturing sectors (sector 6-22 in Table A.1). "Tariff" refers to cases where players can modify their import tariffs solely. "Dual" refers to cases where players have the flexibility to adjust both their industry subsidies and import tariffs. "China and US" refers to cases where only China and the US are allowed to adjust their policies, whereas "World" refers to cases where all economies can adjust their policies.

We first consider the case where only China and the U.S. can adjust their policies. In the case of industrial policy competition (Column (1) of Table 3), China and the U.S. gain substantially at the expense of other economies. Intuitively, industrial policies increase the sizes of high-return-to-scale sectors in China and the U.S. and decrease the scale of these sectors in other countries. In contrast, if the competition is focused solely on import tariff policy (Column (2) of Table 3), both China and the U.S. will encounter a significant decrease in welfare, while most other economies will experience slight negative impacts. Finally, in the dual policy competition (Column (3) of Table 3), both China and the U.S. will obtain a sizable increase in welfare, whereas other economies will again bear the brunt of tariff wars between China and the U.S., leading to welfare losses.

We then consider global competition in trade and industrial policies. Column (4) of Table 3 suggests that all countries substantially benefit from the non-cooperative industrial subsidies, if these subsidies are specified and implemented correctly. In contrast, a global trade war would result in substantial welfare losses of all economies (Column (6) of Table 3). The overall welfare effects of two policy competitions, as shown in Column (7) of Table 3, are therefore mixed.

We further consider constraints for governments in implementing industrial subsidies. In particular, we consider the case where each country can only impose a uniform subsidy on all of its manufacturing sectors (sector 6-22 in Table A.1). The welfare effects of the optimal uniform subsidies are shown in Column (5) of Table 3. Comparing with the sector-specific optimal subsidies in each country, the optimal uniform subsidies lead to much smaller welfare gains from most economies.

Finally, we investigate the interactions of trade and industrial policy competitions. In particular, we compare Nash tariffs under tariff competition and under dual competition. Figure 10 shows the result in global competition. It suggests that countries tend to impose higher import tariffs in the global competition if they are not allowed to compete via industrial subsidies, which is consistent with the quantitative results in Ju et al. (2024). Figure 11 shows the similar results in the US-China competition.

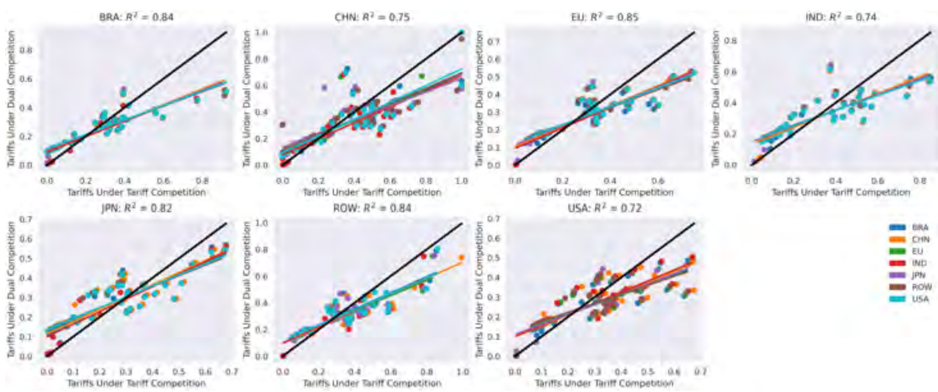

Figure 10: Optimal Tariffs Under Global Tariff Competition and Dual Competition

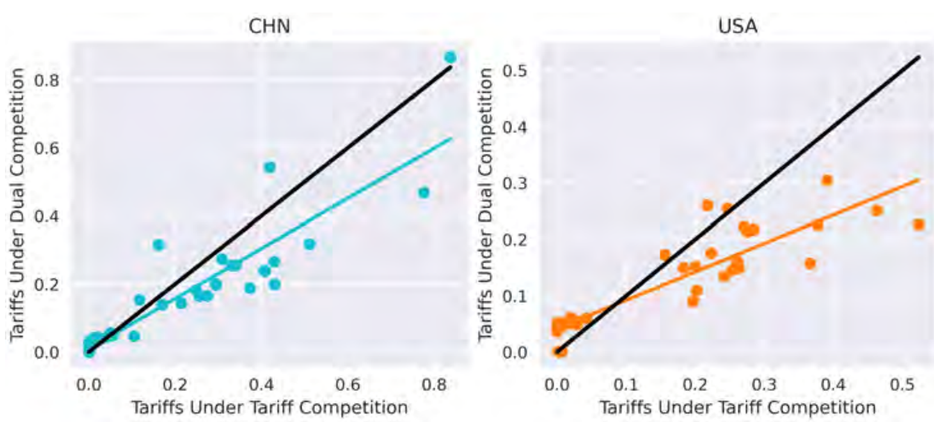

Figure 11: Optimal Tariffs Under China and US Tariff Competition and Dual Competition

In sum, our quantitative results for Nash policies suggest that (i) tariff competitions lead to substantial losses for all countries; (ii) non-cooperative industrial policies, if correctly specified and implemented, could result in welfare gains; and (iii) countries tend to impose lower tariffs if they are allowed to implement industrial subsidies.

4.4.3 Imperfect Implementation of Industrial Subsidies

In section 4.4.2, we have shown that China gains considerably from non-cooperative industrial subsidies. However, this welfare gain hinges on accurately specifying and implementing industrial subsidies. In this subsection, we consider the case in which China cannot perfectly implement its optimal industrial subsidies in the global industrial policy competition. Instead, China chooses from $\left[0.1e_{CHN}^{j^*}, 1.9e_{CHN}^{j^*}\right]$, following a uniform distribution, for all tradable sector j.

We randomly draw China's subsidies for 1000 times. Figure 12 is the histogram of the corresponding welfare changes in China. It suggests that given other countries set their industrial subsidies optimally, China has to precisely specify its industrial subsidies to gain as in Column (4) of 3 (on average 2.42%). Once deviate from its optimal values, China's industrial subsidies would only lead to small welfare gains, or in many cases, welfare losses.

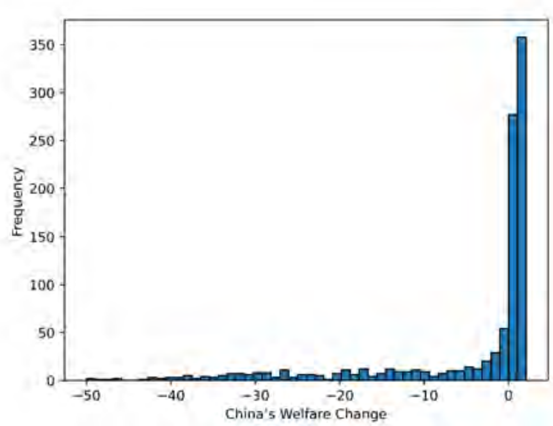

Figure 12: China's Welfare Change (%) under Randomly Drawn Industrial Subsidies

(Note: Each e_{CHN}^{j} is drawn uniformly from $\left[0.1e_{CHN}^{j^*}, 1.9e_{CHN}^{j^*}\right]$. We draw $\left(e_{CHN}^{j}\right)_{j=1}^{22}$ for 1000 times.)

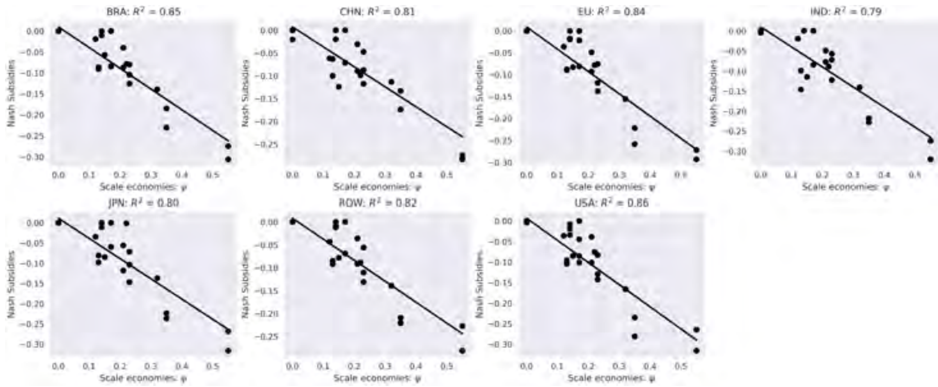

Figure 13: Subsidies in Global Dual Policy Cooperation

4.4.4 Cooperative Polices

Having characterized non-cooperative policy competitions across major economies, we turn to consider cooperative policies aiming to maximize the global welfare. In particular, we consider a world social planner that chooses trade and industrial policies in all economies to maximize the weighted-average welfare change defined as follows

$$\hat{W} \equiv \sum_{n=1}^{N} \frac{Y_n}{\sum_{k=1}^{N} Y_k} \left(\frac{\hat{Y}_n}{\hat{P}_n} \right). \tag{22}$$

Notice that (i) the global cooperative policies may lead to welfare reduction in certain countries since the global social planner cares about weighted-average welfare changes in the global economy; and (ii) since cooperative policies have to be determined simultaneously in all countries and sectors, solving for these policies are much more computationally challenging than solving for the Nash policies. It takes about 24 hours to solve for the global dual cooperative policies in which the global social planner chooses tariffs and subsidies in all countries to maximize the welfare expressed by Equation (22).

Figure 13 links global cooperative subsidies with sectoral scale economies. It shows that (i) global cooperative subsidies are substantial, indicating strong motives for the global social planner to address sectoral misallocation by subsidies; and (ii) like Nash subsidies, global cooperative subsidies are strongly increasing with sectoral scale economies.

However, tariffs in global dual policy cooperation are close to zero and uniform across countries and sectors, as shown in Figure 14.

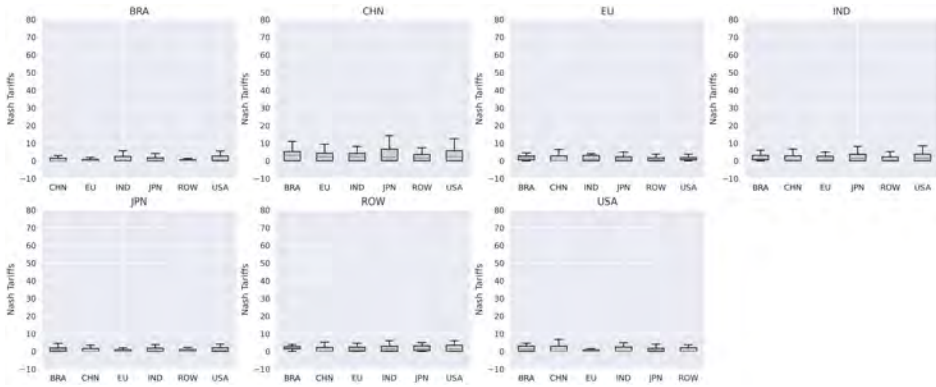

Figure 14: Tariffs in Global Dual Policy Cooperation

In sum, the global social planner tends to set all tariffs close to zero and utilize industrial subsidies to address sectoral misallocation.

Table 4 summarizes the welfare effects of global policy cooperation. We find that global dual policy cooperation increases welfare in all major economies except India (Column (3) of Table 4). India suffers welfare losses due to its low share in the world pre-change final income. These welfare gains, as shown in Column (1) and (2) in Table 4, mainly come from global subsidy cooperation. Global trade cooperation only modestly increases welfare in major economies.

Table 4: Welfare Changes of Global Cooperation with Scale Economies (%)

	Subsidy	Tariff	Dual
	(1)	(2)	(3)
United States	2.7853	0.1899	2.8218
China	3.5491	1.1299	4.2923
European Union	1.8508	0.1850	1.7828
Japan	1.4536	0.3724	1.3948
India	-0.0603	2.3386	-0.2739
Brazil	3.3496	0.0790	3.4655
Rest of the World	1.9264	0.0594	1.8776

5 Conclusion

In this paper, we propose a gradient-based learning framework to solve optimal policies in quantitative trade and spatial models. This framework utilizes advanced techniques on gradient descending and we implement it using the machine learning (ML) open-source framework such as PyTorch, TensorFlow, and Google JAX. Our method is efficient in computing high-dimensional optimal policies in trade models with very high-dimensional equilibrium systems.

We use our framework to compute optimal trade and industrial policies in a multi-country-multi-sector GE model with trade, input-output linkages, and sectoral scale economies. Our quantitative results suggest that (i) the analytical optimal policies characterized by Lashkaripour and Lugovskyy (2023) rely on the “internal cooperation” assumption which does not hold in our fully optimal policies; (ii) Nash tariffs and industrial subsidies strongly increase with sectoral scale economies; (iii) countries tend to gain from Nash industrial subsidies but lose from Nash tariffs; and (iv) the global optimal policies consist of close-to-zero tariffs and substantial industrial subsidies that increase with sectoral scale economies.

Our method is widely applicable in trade and spatial economics. It does not suffer from the curse of dimensionality. Therefore, we can use this method to compute high-dimensional optimal (continuous) policies in models where a large number of agents (possibly regions) interact with each other. The applications include but are not limited to optimal carbon emissions, optimal corporate taxes, and optimal innovation subsidies over space.

References

ABADI, M., A. AGARWAL, P. BARHAM, E. BREVDO, Z. CHEN, C. CITRO, G. S. CORRADO, A. DAVIS, J. DEAN, M. DEVIN, ET AL. (2016): "TensorFlow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*.

BARRON, E., R. GOEBEL, AND R. JENSEN (2010): "Best response dynamics for continuous games," *Proceedings of the American Mathematical Society*, 138, 1069–1083.

BARTELME, D., A. COSTINOT, D. DONALDSON, AND A. RODRIGUEZ-CLARE (2021): "The Textbook Case for Industrial Policy: Theory Meets Data," *Mimeo*.

BAYDIN, A. G., B. A. PEARLMUTTER, A. A. RADUL, AND J. M. SISKIND (2018): "Automatic differentiation in machine learning: a survey," *Journal of Marchine Learning Research*, 18, 1–43.

CHASNOV, B., L. J. RATLIFF, D. CALDERONE, E. MAZUMDAR, AND S. A. BURDEN (2019): "Finite-Time Convergence of Gradient-Based Learning in Continuous Games," *Mimeo*.

DEKLE, R., J. EATON, AND S. KORTUM (2008): "Global rebalancing with gravity: Measuring the burden of adjustment," *IMF Staff Papers*, 55, 511–540.

DINDOŠ, M. AND C. MEZZETTI (2006): "Better-reply dynamics and global convergence to Nash equilibrium in aggregative games," *Games and Economic Behavior*, 54, 261–292.

DURAND, S. AND B. GAUJAL (2016): "Complexity and optimality of the best response algorithm in random potential games," in *International Symposium on Algorithmic Game Theory*, Springer, 40–51.

EATON, J. AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

FROSTIG, R., M. J. JOHNSON, AND C. LEARY (2018): "Compiling machine learning programs via high-level tracing," *Systems for Machine Learning*, 4.

FUDENBERG, D. AND D. K. LEVINE (1998): *The theory of learning in games*, vol. 2, MIT press.

HEINRICH, T., Y. JANG, L. MUNGO, M. PANGALLO, A. SCOTT, B. TARBUSH, AND S. WIESE (2023): "Best-response dynamics, playing sequences, and convergence to equilibrium in random games," *International Journal of Game Theory*, 1–33.

HOFBAUER, J. AND K. SIGMUND (1998): *Evolutionary games and population dynamics*, Cambridge university press.

HOFBAUER, J. AND S. SORIN (2006): "Best response dynamics for continuous zero-sum games," *Discrete and Continuous Dynamical Systems Series B*, 6, 215.

JU, J., H. MA, Z. WANG, AND X. ZHU (2024): "Trade Wars and Industrial Policy Competitions: Understanding the US-China economic conflicts," *Journal of Monetary Economics*, 141.

JUDD, K. L. AND C.-L. SU (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80.

KINGMA, D. P. AND J. BA (2014): "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*.

LASHKARIPOUR, A. AND V. LUGOVSKYY (2023): "Profits, scale economies, and the gains from trade and industrial policy," *American Economic Review*, 113, 2759–2808.

LESLIE, D. S., S. PERKINS, AND Z. XU (2020): "Best-response dynamics in zero-sum stochastic games," *Journal of Economic Theory*, 189, 105095.

MATSUI, A. (1992): "Best response dynamics and socially stable strategies," *Journal of Economic Theory*, 57, 343–362.

MAZUMDAR, E., L. J. RATLIFF, AND S. S. SASTRY (2020): "On gradient-based learning in continuous games," *SIAM Journal on Mathematics of Data Science*, 2, 103–131.

NISAN, N., T. ROUGHGARDEN, E. TARDOS, AND V. V. VAZIRANI (2007): *Algorithmic Game Theory*, Cambridge University Press.

OSSA, R. (2014): "Trade Wars and Trade Talks with Data," *American Economic Review*, 104, 4104–4146.

PASZKE, A., S. GROSS, F. MASSA, A. LERER, J. BRADBURY, G. CHANAN, T. KILLEEN, Z. LIN, N. GIMELSHEIN, L. ANTIGA, ET AL. (2019): "PyTorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, 32.

RAGHUNATHAN, A., A. CHERIAN, AND D. JHA (2019): "Game Theoretic Optimization via Gradient-based Nikaido-Isoda Function," in *Proceedings of the 36th International Conference on Machine Learning*, ed. by K. Chaudhuri and R. Salakhutdinov, PMLR, vol. 97 of *Proceedings of Machine Learning Research*, 5291–5300.

SINGH, S., M. J. KEARNS, AND Y. MANSOUR (2000): "Nash Convergence of Gradient Dynamics in General-Sum Games." in *UAI*, 541–548.

SWENSON, B., R. MURRAY, AND S. KAR (2018): "On best-response dynamics in potential games," *SIAM Journal on Control and Optimization*, 56, 2734–2767.

VARMA, V. S., J. VEETASEVEERA, R. POSTOYAN, AND I.-C. MORĂRESCU (2021): "Distributed gradient methods to reach a Nash equilibrium in potential games," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 3098–3103.

ZHAO, W. X., K. ZHOU, J. LI, T. TANG, X. WANG, Y. HOU, Y. MIN, B. ZHANG, J. ZHANG, Z. DONG, ET AL. (2023): "A survey of large language models," *arXiv preprint arXiv:2303.18223*.

A Model and Calibration

A.1 "Exact-hat" algebra

Changes in unit costs can be expressed as

$$\hat{c}_i^j = \frac{1}{\left(\hat{L}_i^j\right)^{\psi_j}} \hat{w}_i{}^{\beta_i^j} \left[\prod_{s=1}^{J} \left(\hat{P}_i^s\right)^{\gamma_i^{sj}}\right]^{1-\beta_i^j}. \tag{A.1}$$

Changes in trade share:

$$\hat{\pi}_{in}^j = \frac{\left[\hat{c}_i^j \widehat{1+t_{in}^j} \widehat{1+e_{in}^j}\right]^{-\theta_j}}{\left(\hat{P}_n^j\right)^{-\theta_j}}. \tag{A.2}$$

Changes in price indices:

$$\hat{P}_n^j = \left[\sum_{i=1}^{N} \pi_{in}^j \left[\hat{c}_i^j \widehat{1+t_{in}^j} \widehat{1+e_{in}^j}\right]^{-\theta_j}\right]^{-\frac{1}{\theta_j}}. \tag{A.3}$$

Changes in sectoral wage incomes:

$$\hat{w}_i \hat{L}_i^j w_i L_i^j = \beta_i^j \sum_{n=1}^{N} \frac{\hat{\pi}_{in}^j \hat{X}_n^j X_{in}^j}{\left(1+t_{in}^j\right)' \left(1+e_{in}^j\right)'}. \tag{A.4}$$

Changes in sectoral labor allocation satisfy:

$$\sum_{j=1}^{J} \hat{L}_i^j L_i^j = \bar{L}_i. \tag{A.5}$$

Changes in the total income:

$$\hat{Y}_i Y_i = \hat{w}_i w_i \bar{L}_i + \sum_{j=1}^{J} \sum_{n=1}^{N} \frac{\left(e_{in}^j\right)'}{\left(1+e_{in}^j\right)'} \left(X_{in}^j\right)' + \sum_{j=1}^{J} \sum_{k=1}^{N} \frac{\left(t_{ki}^j\right)'}{\left(1+t_{ki}^j\right)' \left(1+e_{ki}^j\right)'} \left(X_{ki}^j\right)'. \tag{A.6}$$

Changes in sectoral expenditure:

$$\hat{X}_i^j X_i^j = \alpha_i^j \hat{Y}_i Y_i + \sum_{s=1}^{J} \left(1-\beta_i^s\right) \gamma_i^{js} \sum_{n=1}^{N} \frac{\left(X_{in}^s\right)'}{\left(1+t_{in}^s\right)' \left(1+e_{in}^s\right)'}. \tag{A.7}$$

Changes in aggregate price indices:

$$\hat{P}_n = \prod_{j=1}^{J} \left(\hat{P}_n^j\right)^{\alpha_n^j}. \tag{A.8}$$

Optimal policies are solved by

$$\begin{aligned} &\max_{\left\{t_{in}^{j\prime}, e_{ni}^{j\prime}; \hat{w}_i, \hat{L}_i^j, \hat{P}_i^j, \hat{X}_i^j\right\}_{i,j}} \hat{W}_n \equiv \frac{\hat{Y}_n}{\hat{P}_n}, \quad \forall n = 1, 2, \cdots, N \\ &\text{s.t. Equation (A.3), (A.5),(A.4), and (A.7)} \end{aligned} \tag{A.9}$$

A.2 Calibration of (ψ_j, θ_j)

The sector-specific trade and scale elasticities, (θ_j, ψ_j), are calibrated from Lashkaripour and Lugovskyy (2023). The calibrated values are summarized in Table A.1.

Table A.1: Calibration of (ψ_j, θ_j) from Lashkaripour and Lugovskyy (2023)

Industry	ICIO code	Description	θ_j	ψ_j
1	D01T02	Agriculture	6.23	0.14
2	D03	Fishing	6.23	0.14
3	D05T06	Mining, energy	5.28	0.17
4	D07T08	Mining, non-energy	5.28	0.17
5	D09	Mining support	5.28	0.17
6	D10T12	Food	2.30	0.35
7	D13T15	Textiles	3.36	0.22
8	D16	Wood	3.90	0.23
9	D17T18	Paper	2.65	0.32
10	D19	Petroleum	0.64	0.35
11	D20	Chemical	3.97	0.23
12	D21	Pharmaceutical	3.97	0.23
13	D22	Rubber	5.16	0.14
14	D23	Non-metallic	5.28	0.17
15	D24	Basic metals	3.00	0.21
16	D25	Fabricated metal	3.00	0.21
17	D26	Computer	1.24	0.55
18	D27	Electrical equipment	1.24	0.55
19	D28	Machinery nec	7.75	0.12
20	D29	Motor vehicles	2.81	0.13
21	D30	Other transport equipment	2.81	0.13
22	D31T33	Manufacturing nec	6.17	0.15

Notes: We set $\theta_j = 10$ and $\psi_j = 0$ for non-tradable sectors.

B Supplementary Quantitative Results

B.1 Convergence Graphs

The computation process of finding a Nash equilibrium involves iterative algorithms that aim to converge towards a stable solution. These iterations refine the strategies of each player until an equilibrium is reached. To demonstrate convergence, we track the welfare change metric at each iteration. The welfare change metric represents how much the overall welfare of the players changes from one iteration to the next. We plot the number of iterations on the x-axis and the corresponding metric value on the y-axis. As the iterations progress, the values of the welfare change metric eventually approach a stable point. This indicates that the algorithm is converging towards a Nash equilibrium, where no player has an incentive to unilaterally change their strategy. This visual representation helps demonstrate the iterative refinement of strategies and the eventual attainment of equilibrium, which provides confidence that the algorithm is finding a stable solution.

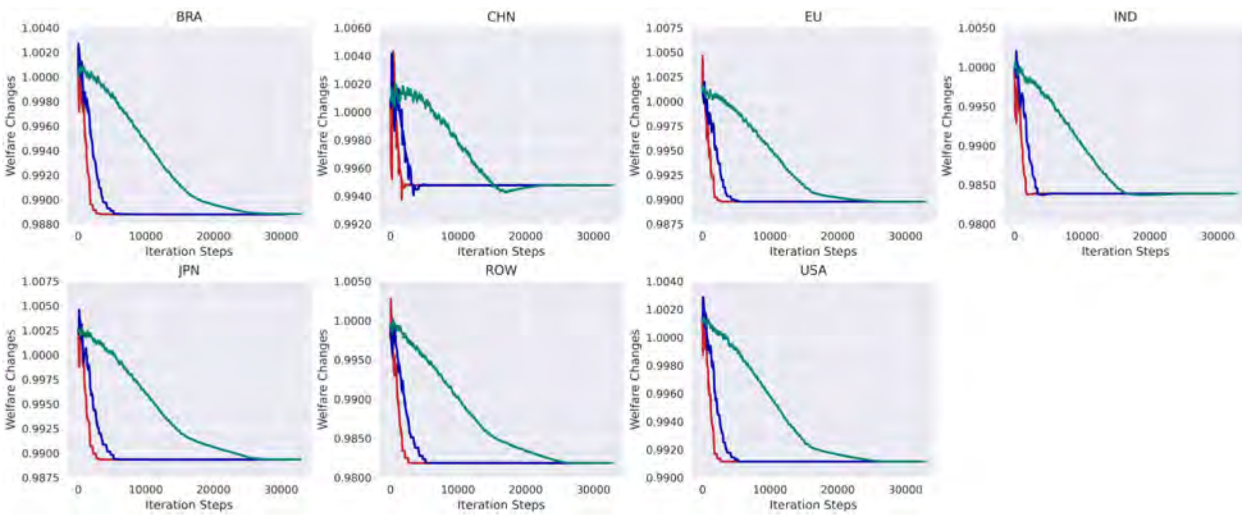

Figure B.1: Computation of Nash Equilibrium for Global Dual Policy Competition without Scale Economies ($\psi = 0$)

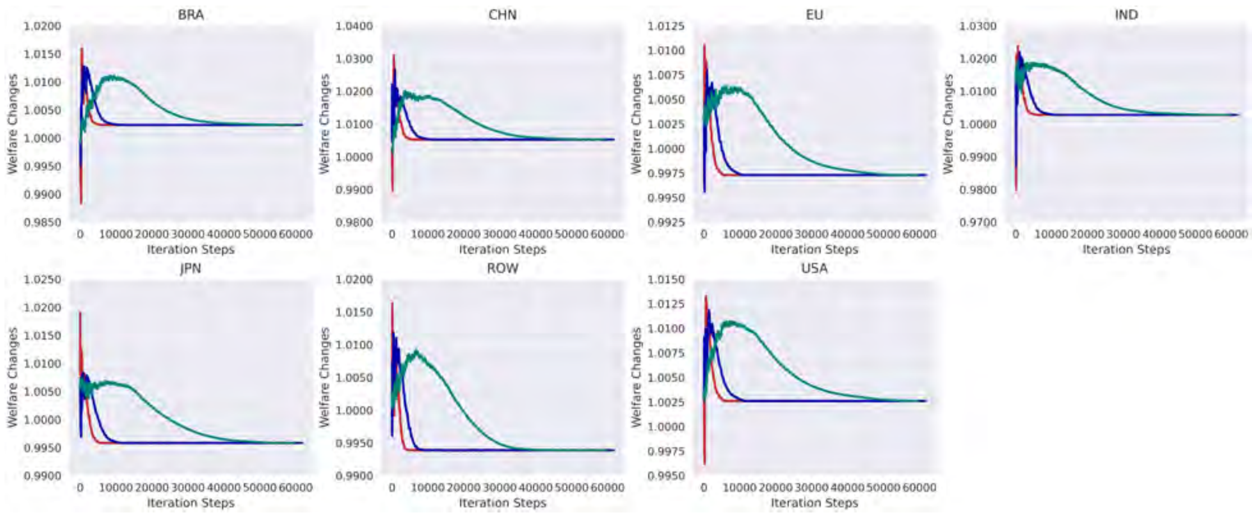

Figure B.2: Computation of Nash Equilibrium for Global Dual Policy Competition with Scale Economies ($\psi \neq 0$)

B.2 Optimal Trade and Industrial Policies without Scale Economies

In this subsection, we consider the neoclassical model without scale economies, i.e. $\psi_j = 0$, as a special case. Figure B.3 illustrates the import tariffs in the Nash equilibrium in which each country decides its tariffs and industrial subsidies to maximize its own real income. It suggests that in a global competition without scale economies, each economy tends to apply almost identical tariffs to all other economies across all industries.

Similar to Nash tariffs in Figure B.3, import tariffs under unilaterally optimal policies, as illustrated by Figure B.4, are also almost identical across industries and economies.

Table B.1 presents welfare changes under various Nash equilibria. The "subsidy" column denotes cases where players can only adjust their industry subsidies. The "tariff" column signifies situations where players can solely modify their import tariffs. Finally, the "dual" column indicates scenarios where players have the flexibility to adjust both their industry subsidies and import tariffs.

When considering a competition solely between China and the U.S. without scale economies, the outcomes vary depending on the policy pursued. However, it is important to note that the effects of all three scenarios are relatively small.

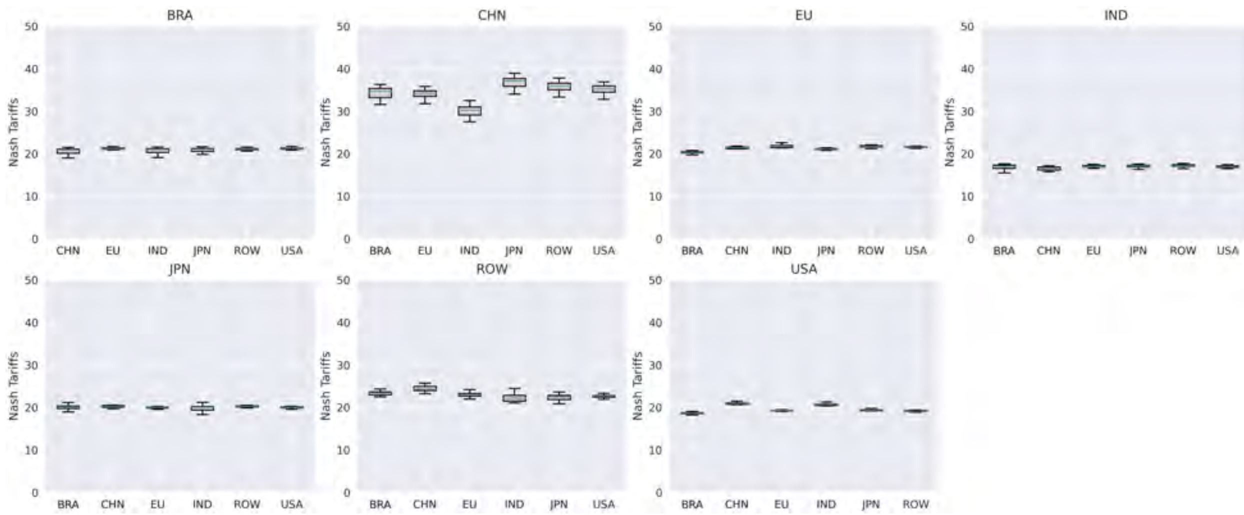

Figure B.3: Nash Tariffs for Global Dual Policy Competition without Scale Economies

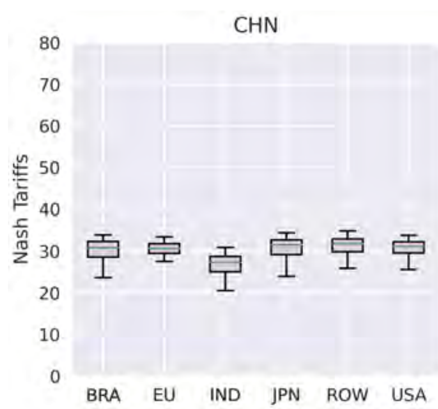

Figure B.4: Import Tariffs in China's Unilaterally Optimal Dual Policies Without Scale Economies

In the absence of scale economies, the imposition of tariffs and the implementation of dual policies (Column 5 and 6, respectively) have adverse effects on all seven economies. These competition scenarios give rise to Prisoner's Dilemma situations. When there is only subsidy competition (Column 5), Brazil and China experience slight benefits from the competition, while the other economies are negatively impacted.

Economies	**China and US** ($\Delta\%$)			**World** ($\Delta\%$)		
	Subsidy	Tariff	Dual	Subsidy	Tariff	Dual
	(1)	(2)	(3)	(4)	(5)	(6)
BRA	-0.0264	0.0014	-0.0248	0.0102	-1.1504	-1.1277
CHN	0.1262	-0.0488	0.0713	0.0445	-0.5856	-0.5247
EU	-0.0228	0.0022	-0.019	-0.0492	-1.0517	-1.0265
IND	-0.0152	0.0063	-0.0057	-0.0951	-1.6609	-1.6261
JPN	-0.0385	0.0012	-0.0368	-0.0757	-1.1136	-1.0731
ROW	-0.1035	-0.0004	-0.1039	-0.0794	-1.8126	-1.832
USA	-0.013	-0.0362	-0.0467	-0.0502	-0.9101	-0.8904

Table B.1: Welfare Changes at Nash Equilibrium without Scale Economies ($\psi = 0$)

B.3 Unilaterally Optimal Full Policies

Figure B.5 illustrates the unilaterally optimal full policies in China. The main patterns in Figure 5 preserve.

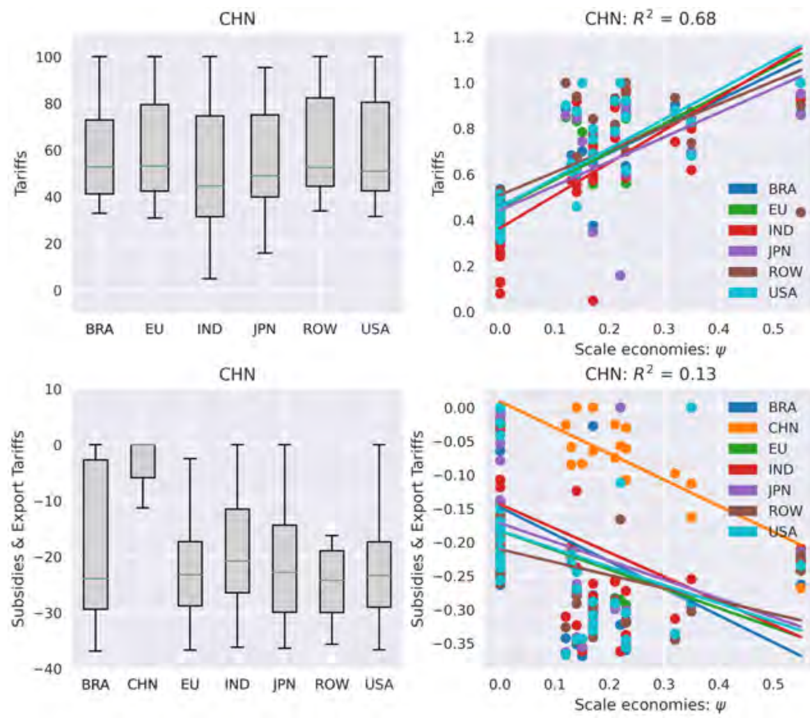

Figure B.5: Import Tariffs, Industry Subsidies and Export Tariffs in China's Unilaterally Optimal Full Policies