

Optimal Trade and Industrial Policies in the Global Economy: A Deep Learning Framework*

Zi Wang^{†1}, Xingcheng Xu^{‡2}, Yanqing Yang^{§3}, and Xiaodong Zhu^{¶4}

¹Hong Kong Baptist University

²Shanghai Artificial Intelligence Laboratory

³Fudan University

⁴The University of Hong Kong

July 21, 2024

Abstract

We propose a deep learning framework, *DL-opt*, designed to efficiently solve for optimal policies in quantifiable general equilibrium trade models. *DL-opt* integrates (i) a nested fixed point (NFXP) formulation of the optimization problem, (ii) automatic implicit differentiation to enhance gradient descent for solving unilateral optimal policies, and (iii) a best-response dynamics approach for finding Nash equilibria. Utilizing *DL-opt*, we solve for non-cooperative tariffs and industrial subsidies across 7 economies and 44 sectors, incorporating sectoral external economies of scale. Our quantitative analysis reveals significant sectoral heterogeneity in Nash policies: Nash industrial subsidies increase with scale elasticities, whereas Nash tariffs decrease with trade elasticities. Moreover, we show that global dual competition, involving both tariffs and industrial subsidies, results in lower tariffs and higher welfare outcomes compared to a global tariff war. These findings highlight the importance of considering sectoral heterogeneity and policy combinations in understanding global economic competition.

JEL classification: F12; F51; C61; C63

Keywords: Deep Learning; Tariff Wars; Industrial Policies; Optimal Policies; Nash Equilibria; Best-response dynamics

*We are grateful to Yucheng Yang, Jeffrey Sun, Costas Arkolakis, Daniel Xu, Dan Lu, and numerous seminar audiences for helpful suggestions and comments. 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 among leading economies have intensified, highlighting the need to understand the incentives and consequences of these conflicts for both economic studies and policymaking. Two notable trends have emerged within this context. First, countries increasingly rely on a combination of trade and industrial policies to enhance their global competitiveness. Second, in many real-world scenarios, countries behave as if engaged in non-cooperative games, implementing policies that benefit themselves at the expense of others. Given these trends, it is essential to evaluate the costs of economic conflicts and the benefits of potential economic cooperation by solving for optimal trade and industrial policies in the global economy.

However, it is computationally challenging to solve for high-dimensional continuous optimal policies and the corresponding Nash equilibria in quantitative general equilibrium models. Three primary difficulties arise. First, conditional on a set of policies, equilibrium outcomes must be determined by solving a high-dimensional nonlinear system. Second, the policy space often includes numerous continuous policy tools. Third, solving for Nash equilibria necessitates iteratively determining unilateral optimal policies for each country.

There are two main approaches in the literature to address the challenges above. First, [Judd and Su \(2012\)](#) develop the mathematical programming with equilibrium constraints (MPEC) framework. When combined with standard algorithms for constrained optimization, this approach offers a speed advantage because it avoids repeatedly solving the equilibrium system for each policy guess. However, this speed advantage diminishes when the dimension of equilibrium conditions is high, due to the exploding dimensionality of the solution space.¹

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

¹For example, based on [Judd and Su \(2012\)](#), [Ossa \(2014\)](#) required four days to compute optimal tariffs (both unilaterally and mutually) in a general equilibrium model with 7 economies and 33 industries, but without input-output linkages, using a high-end desktop computer and standard MATLAB software. Incorporating more economies, industries, and policies, as well as input-output linkages, presents significant computational challenges.

can get optimal policies by calculating several simple and intuitive sufficient statistics. However, this approach has to make simplification assumptions to get the exact form of sufficient statistics.² As a result, their optimal policies are in fact constrained optimal policies under simplification assumptions. To obtain fully optimal policies without these assumptions, we still need to solve the high-dimensional nonlinear optimization problem.

In this paper, we propose a deep learning framework, *DL-opt*, for solving optimal (unilaterally and mutually) policies in quantifiable general equilibrium trade models. To avoid the exploding dimensionality of solution space in MPEC, we do not include equilibrium outcomes into the solution space and take equilibrium conditions as constraints. Instead, we employ the nested fixed point (NFXP) algorithm developed by Rust (2000) that solves the equilibrium system by iteration at each guess of policies. Utilizing NFXP, our focus is on improving the efficiency of policy updating. To this end, we employ the widely-used deep learning algorithm that can efficiently update the policies based on the gradient of the objective function with respect to the parameters. The key for gradient descent is automatic implicit differentiation, which efficiently compute the gradient of equilibrium outcomes with respect to policies via the implicit function defined by the equilibrium system.

Armed with our deep learning method to solve for unilaterally optimal policies, we employ state-of-the-art best-response dynamics method for solving Nash equilibria. In particular, we use a random shuffle playing sequence to improve algorithm efficiency and to avoid being trapped in local solutions. In sum, by combining the NFXP algorithm, automatic implicit differentiation, and best-response dynamics, our *DL-opt* framework is able to solve for high-dimensional non-cooperative policies in general equilibrium trade models.

We then apply our *DL-opt* framework to solve optimal trade and industrial policies in a multi-country-multi-sector general equilibrium trade model with input-output linkages and sectoral scale economies, à la the models of Lashkaripour and Lugovskyy (2023), Bartelme et al. (2021), and Ju, Ma, Wang, and Zhu (2024). We consider the world with 6 major economies plus the rest of

²For 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.

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 5 hours. In comparable computational exercises, our method demonstrates an efficiency at least ten times greater than that of widely used nonlinear solvers, such as Knitro, particularly when analytical Jacobian and Hessian matrices are unavailable. Our solutions are shown to be precise, with a stable convergence process. Additionally, we demonstrate that parallelization can further enhance the efficiency of our method.

Our quantitative explorations focus on *global dual competition* in which each economy chooses import tariffs and industrial subsidies to maximize its own real income. To isolate the impacts of trade wars and understand the independence of trade and industrial policies, we also consider the *global trade war* in which each economy chooses import tariffs to maximize its own real income. Nash policies computed using our *DL-opt* framework shed light on global competition in the following aspects:

First, Nash industrial subsidies increase with sectoral scale elasticities but are not correlated with trade elasticities. Consistent with previous studies such as [Lashkaripour and Lugovskyy \(2023\)](#), these results suggest that industrial subsidies can be used to address misallocation across sectors led by external economies of scale. The terms-of-trade effect, decreasing with trade elasticities, is not a primary concern in determining optimal industrial subsidies.

Second, Nash tariffs under global dual competition increase with scale elasticities but decrease with trade elasticities. Moreover, we find that the positive correlation between Nash tariffs and scale elasticities is much stronger under trade wars than that under dual competition. This result suggests that when industrial policies are not allowed, tariffs have to be used to address misallocation across sectors and thereby be prohibitively high on sectors with strong increasing returns to scale. In contrast, when industrial policies are employed, tariffs are mainly used to manipulate terms-of-trade and thereby decrease with trade elasticities.

Third, we find that Nash tariffs under trade war is much higher than those under dual competition. The average Nash tariffs under global trade war is 42%, whereas that under global dual competition is 35%. Intuitively, if a country is able to efficiently subsidize its high-tech industries

with strong economies of scale, it has lower incentives to impose high import tariffs on these industries. Consequently, comparing with global trade war, global dual competition via the combination of tariffs and industrial policies leads to substantially larger welfare gains (or lower welfare losses).

Fourth, we extend our framework to consider global cooperative policies decided by a global social planner that aims to maximize average real income across countries. We find that cooperative tariffs are close to zero, whereas cooperative industrial subsidies are strongly increasing with scale elasticities. Moreover, comparing with global cooperation on tariffs, global cooperation on tariffs and industrial subsidies results in larger welfare gains for all major economies except India.

Related Literature. This paper is the first attempt to use deep learning algorithm and machine-learning implementation to solve for optimal policies in general equilibrium 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 advances in employing deep learning methods to solve high-dimensional macroeconomic models. [Han, E, and Yang \(2021\)](#) and [Sun \(2023\)](#) propose deep learning methods to solve for heterogeneous-agent models with aggregate shocks. [Feng, Han, and Zhu \(2023\)](#) explores optimal taxation with incomplete market. The central objective of this literature is to approximate high-dimensional value functions or policy functions using deep learning techniques. In contrast, our *DL-opt* framework aims at identifying the optimal points within high-dimensional objective functions derived from high-dimensional equilibrium systems. Although our model is static, it often encompasses higher dimensions than those typically addressed in dynamic models within this field.

This paper contributes to recent quantitative analysis on trade and industrial policies. [Ju et al. \(2024\)](#) 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. Additionally, [Lashkaripour and Lugovsky \(2023\)](#) and [Bartelme et al. \(2021\)](#) utilize a sufficient statistic approach to characterize optimal trade and industrial policies. However, their reliance on simplification assumptions results in constrained optimal policies, as previously discussed.

The remainder of the paper is organized as follows. In Section 2, we propose the general problem and solution framework of *DL-opt*. We then apply this framework to solve for optimal trade and industrial policies in a multi-country-multi-sector trade model. In Section 3, we set the model, define the equilibrium, calibrate the model, and compute the optimal policies using our *DL-opt* framework. Section 4 characterizes the features of Nash tariffs and industrial subsidies solved in Section 3. We conclude in Section 5.

2 DL-opt: General Framework

Our *DL-opt* framework consists of three parts. In this section, we first introduce (i) the nested fixed point (NFXP) algorithm and (ii) automatic implicit differentiation to solve for unilaterally optimal policies. We then discuss (iii) best-response dynamics to solve for Nash policies.

2.1 Unilaterally Optimal Policies

We start from one granular player, called the government, that decides an action vector \mathbf{a} . All atomic players, including households and firms, take \mathbf{a} as given. The behaviors of all atomic players and the granular player determine the equilibrium outcome vector \mathbf{X} . The general equilibrium is characterized by the following nonlinear system:

$$\mathbf{X} = G(\mathbf{X}, \mathbf{a}). \quad (1)$$

The government aims to maximize $W(\mathbf{X}, \mathbf{a}) \in \mathbb{R}$ that depends on both the action \mathbf{a} and the equilibrium outcome \mathbf{X} . We can express the government's problem in a MPEC framework:

$$\begin{aligned} & \max_{(\mathbf{a}, \mathbf{X})} W(\mathbf{X}, \mathbf{a}) \\ & \text{s.t. } \mathbf{X} = G(\mathbf{X}, \mathbf{a}) \end{aligned} \quad (2)$$

However, in many applications \mathbf{X} has much higher dimensions than \mathbf{a} . Therefore, the MPEC framework in Equation (2) incurs the curse of dimensionality and is computationally challenging.

Nested Fixed Point (NFPX) Algorithm. Suppose that we have calculated \mathbf{X} by solving the fixed point in Equation (1) and obtain $\mathbf{X} = \mathbf{X}(\mathbf{a})$ for any given \mathbf{a} . Then the government’s problem (2) can be re-expressed as

$$\max_{\mathbf{a}} W(\mathbf{X}(\mathbf{a}), \mathbf{a}). \quad (3)$$

This is the nested fixed point (NFXP) algorithm developed by Rust (2000). It consists of two tiers of loops.

1. Inner loop: Given \mathbf{a} , solve the fixed point \mathbf{X} in Equation (1).
2. Outer loop: Update \mathbf{a} by the following general rule:

$$\mathbf{a}_{t+1} = \mathbf{a}_t + \gamma \omega(\nabla_{\mathbf{a}} W(\mathbf{X}(\mathbf{a}_t), \mathbf{a}_t)), \quad (4)$$

where $\nabla_{\mathbf{a}} W$ is the gradient of the objective function with respect to \mathbf{a} , γ is the learning rate, and $\omega(\cdot)$ is a function of the gradient specified by the detailed algorithms.

Fixed Point Iterations. Given policy \mathbf{a} , we calculate \mathbf{X} by the following contraction iterations:

$$\mathbf{X}_{t+1} = G(\mathbf{X}_t, \mathbf{a}). \quad (5)$$

Rust (2000) argues that the calculation of fixed point \mathbf{X} can be accelerated by utilizing Newton-Kantorovich iterations after contraction iterations. The Newton-Kantorovich iterations have the following form:

$$\mathbf{X}_{t+1} = \mathbf{X}_t - [\mathbf{I} - \nabla_{\mathbf{X}} G(\mathbf{X}_t, \mathbf{a})]^{-1} [\mathbf{X}_t - G(\mathbf{X}_t, \mathbf{a})], \quad (6)$$

where \mathbf{I} is the identity matrix and $\nabla_{\mathbf{X}} G$ is the gradient of equilibrium system with respect to \mathbf{X} . However, this method is computationally burdensome when G and \mathbf{X} are high-dimensional since it requires computing the gradient $\nabla_{\mathbf{X}} G$ for each \mathbf{X}_t . In the *DL-opt* framework, we prefer using contraction iterations in Equation (5) to obtain the fixed point \mathbf{X}_T .

Automatic Implicit Differentiation. The key challenge to solve the problem in Equation (3) is to compute $\nabla_{\mathbf{a}}W(\mathbf{X}(\mathbf{a}), \mathbf{a})$ in Equation (4). By the chain rule, it can be expressed as

$$\nabla_{\mathbf{a}}W(\mathbf{X}_T, \mathbf{a}) = [\nabla_{\mathbf{a}}W(\mathbf{X}_T, \mathbf{a}) + \nabla_{\mathbf{X}}W(\mathbf{X}_T, \mathbf{a})\nabla_{\mathbf{a}}\mathbf{X}_T], \quad (7)$$

where \mathbf{X}_T is the fixed point in Equation (1) under \mathbf{a} .

Due to high-dimensionality of both \mathbf{X} and \mathbf{a} , it is generally difficult to compute $\nabla_{\mathbf{a}}\mathbf{X}_T$. However, by the implicit function theorem, we have

$$\nabla_{\mathbf{a}}\mathbf{X}_T = -[\nabla_{\mathbf{X}}G(\mathbf{X}_T, \mathbf{a}) - \mathbf{I}]^{-1} \nabla_{\mathbf{a}}G(\mathbf{X}_T, \mathbf{a}). \quad (8)$$

We utilize automatic differentiation in deep learning to efficiently calculate $\nabla_{\mathbf{X}}G(\mathbf{X}_T, \mathbf{a})$ and $\nabla_{\mathbf{a}}G(\mathbf{X}_T, \mathbf{a})$ at the point \mathbf{X}_T . Equation (8) combines the implicit function theorem with automatic differentiation. Therefore, it is called *automatic implicit differentiation*.³

Now we briefly introduce the idea of automatic differentiation. Automatic differentiation is in particular efficient when $G(.,.)$ is high-dimensional and consists of many elementary operations and functions. Suppose that \mathbf{X} has K dimensions. Then to calculate $\nabla_{\mathbf{X}}G(\mathbf{X}_T, \mathbf{a})$ using finite difference, we need to evaluate $G(.,.)$ for $K + 1$ times. This can be very challenging when K is large.

Automatic differentiation, in contrast, decomposes $G(.,.)$ into many intermediate variables. Each intermediate variable can be computed by either (i) the elements of \mathbf{X} , or (ii) other intermediate variables through elementary operations and functions. These intermediate variables formulate a computation graph that connects $G(.,.)$ with \mathbf{X} .

The key advantage of automatic differentiation is that many intermediate variables are common in calculating $\frac{\partial G}{\partial X_i}$ and $\frac{\partial G}{\partial X_j}$ for $i \neq j$. Therefore, we do not need to evaluate the high-dimensional $G(.,.)$ for K times to compute $\nabla_{\mathbf{X}}G(\mathbf{X}_T, \mathbf{a})$. In Appendix A.1, we provide a toy example to illustrate the advantage of automatic differentiation in detail.

³The idea of automatic implicit differentiation is first proposed in Jeffrey Sun's computation notes. See https://jeffreyesun.com/auto_implicit_diff.pdf.

Repeated application of automatic differentiation to Equation (8) and (7) establishes a connection between $W(\cdot, \cdot)$ and \mathbf{a} through multiple layers of intermediate variables and elementary operations. Figure 1 draws a parallel between policy optimization and neural networks. In this analogy, policies \mathbf{a} are like neural network parameters, intermediate variables correspond to nodes in hidden layers, and the objective function W is analogous to the loss function. Policy optimization adjusts policies \mathbf{a} to maximize W via layers of intermediate variables (Panel (a) of Figure 1), similar to how neural networks adjust parameters to approximate outputs and minimize the loss function (Panel (b) of Figure 1). Consequently, gradient computation through automatic implicit differentiation mirrors the forward-backward propagation in neural networks depicted in Panel (b). This powerful analogy suggests that well-developed deep learning tools can effectively implement automatic differentiation for any well-defined function G .




Figure 1: Analogy between Policy Optimization and Neural Networks

2.2 Nash Policies

We proceed by considering multiple granular players ($i = 1, \dots, N$) and Nash policies. In particular, play i decides action \mathbf{a}_i , taking others' actions \mathbf{a}_{-i} as given. The best response function of player i can be given by

$$b_i^*(\mathbf{a}_{-i}) \equiv \arg \max_{\mathbf{a}_i} W(\mathbf{X}(\mathbf{a}_i, \mathbf{a}_{-i}), \mathbf{a}_i, \mathbf{a}_{-i}), \quad (9)$$

where \mathbf{X} is the fixed point in $\mathbf{X} = G(\mathbf{X}, \mathbf{a}_i, \mathbf{a}_{-i})$. Given \mathbf{a}_{-i} , $b_i^*(\mathbf{a}_{-i})$ can be computed using the NFXP algorithm and automatic implicit differentiation in our *DL-opt* framework described in Section 2.1.

With the best-response function, $b_i^*(\mathbf{a}_{-i})$, in hand, we solve for Nash equilibria, $(\mathbf{a}_i^*)_{i=1}^N$, where $\mathbf{a}_i^* \in b_i^*(\mathbf{a}_{-i})$ for $i = 1, \dots, N$. We transform the static problem into a system of best-response dynamics:

$$\mathbf{a}_i^{t+1} = \eta_t b_i^*(\mathbf{a}_{-i}^t) + (1 - \eta_t) \mathbf{a}_i^t, \quad \forall i = 1, 2, \dots, N, \quad (10)$$

where $\eta_t \in (0, 1]$ is the stepsize.

The best-response dynamics capture the idea of players adjusting their strategies iteratively in response to the strategies chosen by others. In the iterative algorithm expressed by Equation (10), players subsequently chooses their best responses, given others' strategies. The key for this algorithm is the sequence of players' actions. [Heinrich, Jang, Mungo, Pangallo, Scott, Tarbush, and Wiese \(2023\)](#) consider 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. The other 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.

Our *DL-opt* framework utilizes a random shuffle playing sequence that combines the features of two sequences above. This sequence introduces randomness 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.

Figure 2 summarizes the calculation flows of our DL-opt framework. First, the general equilibrium outcome \mathbf{X} is determined using fixed point iteration within the inner loop of the NFXP

Algorithm 1 Best-response dynamics with random shuffle playing sequence

1. Guess the initial action profile of all players \mathbf{a}^0 .
2. Repeat until converge:
 - Random shuffle playing sequence for each round
 - For the i -th player $s(i)$ in the shuffled sequence s
 - (a) Compute the best response of the player $s(i)$ in sequence s

$$\mathbf{a}_{s(i)}^{t+1} = \eta_t b_{s(i)}^* \left(\mathbf{a}_{-s(i)}^t \right) + (1 - \eta_t) \mathbf{a}_{s(i)}^t. \quad (11)$$

- (b) Update $\mathbf{a}_{-n}^{t+1} = \mathbf{a}_{-n}^t$ for all $n \neq s(i)$.
-




Figure 2: Calculation Flows of the DL-opt Framework

algorithm. Second, the best response function is computed via gradient descent, supported by automatic implicit differentiation. Finally, the Nash equilibrium is achieved through best response dynamics with a randomly shuffled playing sequence.

3 Model, Calibration, and Computation

In this section, we present a multi-country-multi-sector general equilibrium trade model with 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 trade and industrial policy competitions. We first introduce the model and its equilibrium, then discuss the model’s calibration, and finally utilize our *DL-opt* framework to compute Nash tariffs and industrial subsidies.

3.1 Model and Equilibrium

Consider N countries with labor endowments $(L_i)_{i=1}^N$. There are J sectors. 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 in country i has a Cobb-Douglas preference over J sectors, with share α_i^j for sector j , and a CES preference over varieties within in sector, with the elasticity of substitution $\sigma_j > 1$.

Each variety is produced under perfect competition using labor and composite intermediates. 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

$$c_i^j = \frac{1}{(L_i^j)^{\psi_j} w_i^{\beta_i^j}} \left[\prod_{s=1}^J (P_i^s)^{\gamma_i^{sj}} \right]^{1-\beta_i^j}, \quad (12)$$

where w_i is the wage in country i and β_i^j represents the value-added share, P_i^s is the price index

of good s in country i and $\gamma_i^{s_j}$ represents the share of good s in producing good j in sector j in country i , and 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 .

Following Eaton and Kortum (2002), we assume that $z_i^j(\omega)$ is drawn independently from a Fréchet distribution:

$$\mathbb{P}\left(z_i^j(\omega) \leq z\right) = \exp\left(-T_i^j z^{-\theta_j}\right) \quad (13)$$

where $z > 0$ and $\theta_j > \max\{\sigma_j - 1, 1\}$.

Exporting good j from country i to n incurs an iceberg trade cost $\tau_{in}^j \geq 1$. Moreover, we consider two policies. First, an ad valorem tariff $t_{in}^j \geq 0$ is imposed by importing country n on goods j imported from country i , with $t_{ii}^j = 0$. Second, an ad valorem output tax $s_i^j \in [0, 1)$ is imposed by production country i on goods j .

Based on the property of Fréchet distribution and the ideal price index of CES preferences, the sectoral price index can be expressed as

$$P_n^j = \left(\sum_{i=1}^N T_i^j \left[c_i^j \tau_{in}^j (1 + t_{in}^j) (1 - s_i^j) \right]^{-\theta_j} \right)^{-\frac{1}{\theta_j}}. \quad (14)$$

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

$$\pi_{in}^j = \frac{X_{in}^j}{X_n^j} = \frac{T_i^j \left[c_i^j \tau_{in}^j (1 + t_{in}^j) (1 - s_i^j) \right]^{-\theta_j}}{\left(P_n^j \right)^{-\theta_j}}, \quad (15)$$

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

$$w_i L_i^j = \beta_i^j \sum_{n=1}^N \frac{X_{in}^j}{(1 + t_{in}^j)(1 + e_{in}^j)}. \quad (16)$$

The wage is determined by labor market clearing

$$\sum_{j=1}^J L_i^j = L_i. \quad (17)$$

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

$$Y_i = w_i L_i + \sum_{j=1}^J \sum_{n=1}^N \frac{e_{in}^j}{1 + e_{in}^j} X_{in}^j + \sum_{j=1}^J \sum_{k=1}^N \frac{t_{ki}^j}{(1 + t_{ki}^j)(1 + e_{ki}^j)} X_{ki}^j. \quad (18)$$

The sectoral expenditure can be expressed by

$$X_i^j = \alpha_i^j Y_i + \sum_{s=1}^J (1 - \beta_i^s) \gamma_i^{js} \sum_{n=1}^N \frac{X_{in}^s}{(1 + t_{in}^s)(1 + e_{in}^s)}. \quad (19)$$

Definition 1 (Equilibrium). *Given parameters $\{\theta_j, \psi_j, \alpha_i^j, \beta_i^j, \gamma_i^{sj}, L_i, s_i^j, t_{in}^j, T_i^j, \tau_{in}^j\}$, the equilibrium consists of $(w_i, L_i^j, P_i^j, X_i^j)$ such that (i) (L_i^j) satisfy Equation (16); (ii) (w_i) are determined by Equation (17); (iii) (P_i^j) satisfy Equation (14); (iv) (X_i^j) satisfy Equation (19).*

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 $P_n = \prod_{j=1}^J (P_n^j)^{\alpha_n^j}$.

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 $\widehat{Z} = Z'/Z$. Given the values of $(\alpha_i^j, \beta_i^j, \gamma_i^{sj}; \psi_j, \theta_j)$, as well as data on $(X_{in}^j, t_{in}^j, s_i^j)$, we can compute the changes in equilibrium outcomes, $(\widehat{w}_i^j, \widehat{L}_i^j, \widehat{P}_i^j, \widehat{X}_i^j)$, by solving a system of $3NJ + N$ nonlinear

equations. In our application, this system is the equilibrium system $G(.,.)$ described in Section 2.1. The details of this system are presented in Appendix A.2.

The optimization problem faced by each country i involves maximizing the change in its welfare, \hat{W}_i , by choosing (t_{ki}^j, s_i^j) for all $k \neq i$ and j . The details of this problem are also presented in Appendix A.2.

3.2 Calibration

The “exact-hat” algebra requires data on bilateral trade shares (π_{in}^j) , sectoral consumption shares (α_i^j) , sectoral value-added shares (β_i^j) , sectoral expenditure (X_n^j) , input expenditure shares γ_i^{js} , and initial policies (t_{in}^j, s_i^j) . 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).⁴ We utilize the OECD Inter-Country Input-Output database (ICIO) for 2017 to extract 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.⁵ We get the initial import tariffs from the World Integrated Trade System (WITS) for 2017 and assume that initially $s_i^j = 0$ for all (i, j) .

We calibrate (ψ_j, θ_j) from Lashkaripour and Lugovskyy (2023). The calibrated values of (ψ_j, θ_j) in Appendix A.3. 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.

3.3 Computation

We utilize our *DL-opt* framework to compute (t_{ki}^j, s_i^j) that maximize \hat{W}_i for each country i . We select the Adaptive Moment Estimation (ADAM) algorithm to implement gradient descent in Equa-

⁴European Union (EU) includes 28 countries including the UK.

⁵The 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>.

tion (4). The procedures and advantages of the ADAM algorithm are presented in Appendix A.4.

Our stochastic best-response dynamics is implemented using the PyTorch machine learning open-source framework.⁶ 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, PyTorch 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 or Google JAX⁷ 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 competition*. The computation statistics of our framework in this baseline scenario is summarized in Table 1. It takes about 5 hours in a single-core personal computer to compute Nash tariffs and industrial policies for all 7 economies and 44 industries.

We compared the computational efficiency of our *DL-opt* framework with widely-used non-linear solvers such as Knitro. The *DL-opt* framework solves for unilaterally optimal industrial subsidies in China in less than **one hour**, whereas Knitro, using interior-point methods without analytical Jacobian and Hessian matrices, takes **seven hours** to achieve the same solution. For this small-scale problem, our *DL-opt* framework is approximately ten times more efficient than Knitro under these conditions.

Figure 3 illustrates the convergence path of our algorithm in each country. The three lines

⁶See Paszke, Gross, Massa, Lerer, Bradbury, Chanan, Killeen, Lin, Gimelshein, Antiga et al. (2019) for the introduction of PyTorch.

⁷See Abadi, Agarwal, Barham, Brevdo, Chen, Citro, Corrado, Davis, Dean, Devin et al. (2016) and Frostig, Johnson, and Leary (2018) for the details of these frameworks.

Table 1: Computation statistics of Solving for Global Dual 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: Iteration Curve for Nash Equilibrium of Global Dual Competition

(Note: The three lines represent distinct computational experiments. The red line corresponds to a learning rate of 0.001, the blue line to a learning rate of 0.0005, and the green line to a learning rate of 0.0001.)

denote separate computational experiments, each conducted under different learning rates. Despite these variations, the outcomes across all experiments remain consistent. It also 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. Moreover, welfare decreases dramatically when subsidy rates exceed their optimal levels but decrease slowly when they are below the optimal. This is mainly because excessive subsidies require additional subsidy expenses and thereby lead to greater welfare losses.




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

4 Optimal Trade and Industrial Policies

Our *DL-opt* framework is efficient in solving for high-dimensional continuous optimal policies in multi-country-multi-sector general equilibrium models. Comparing with the existing methods, our framework takes into account (i) rich heterogeneity across countries and sectors and (ii) combinations of policy instruments. In this section, we focus on global dual competition in which each country decides its import tariffs and sector-specific production subsidies to maximize its real income, given other countries' tariffs and industrial subsidies. To understand the interdependence of trade and industrial policies, we further compare tariffs under global dual competition with those under the global trade war.

We discuss our quantitative results as follows. First, we depict sectoral heterogeneity of non-cooperative tariffs and industrial subsidies solved using our *DL-opt* framework. Second, we explore how non-cooperative tariffs will change if countries are allowed to compete via industrial subsidies. Third, we extend our framework to solve for global cooperative tariffs and industrial subsidies. Finally, we show that in global competition, industrial subsidies have to be correctly specified to generate welfare gains.

4.1 Sectoral Heterogeneity of Optimal Policies

We use our *DL-opt* framework to solve for (unilaterally and mutually) optimal tariffs and industrial subsidies, taking into account rich heterogeneity across countries and sectors. [Lashkaripour and Lugovskyy \(2023\)](#) has provided analytical characterizations of unilaterally optimal tariffs and industrial policies. However, their analytical results are built on the “internal cooperation” assumption: the relative wages among other countries remain unchanged under the optimal policies. Utilizing our *DL-opt* framework, we investigate changes in relative wages under China’s fully optimal policies to see whether this assumption is restrictive in solving for the fully optimal policies. The result suggests that the “internal cooperation” assumption is restrictive and therefore the analytical unilaterally optimal tariffs and industrial policies in [Lashkaripour and Lugovskyy \(2023\)](#) are constrained optimal. The detailed results are presented in Appendix [B.1](#).

We proceed by depicting optimal policies computed using our *DL-opt* framework. Figure 5 compares the unilaterally optimal and Nash policies in China, suggesting that Nash tariffs and industrial subsidies are considerably higher than unilaterally optimal policies. Intuitively, other countries’ optimal trade and industrial policies would shrink China’s production scale in increasing-return-to-scale industries and thereby strengthen China’s incentives to promote production scale by tariffs and industrial subsidies. This result highlights the importance of characterizing mutually optimal (Nash) policies in understanding global competition.




Figure 5: Unilaterally Optimal vs. Nash Policies in China

(Notes: The dash line is the 45-degree line.)

We then illustrate how our Nash policies relate to sectoral trade elasticities, (θ_j) , and scale elasticities (ψ_j) . To this end, we regress $\log(t_i^j)$ (under trade war and under dual competition)

on $\log(\theta_j)$ and $\log(\psi_j)$, controlling for importer and exporter fixed effects. Moreover, we regress $\log\left(\frac{1}{1-s_i^j}\right)$ on $\log(\theta_j)$ and $\log(\psi_j)$, controlling for country fixed effects. The results are reported in Table 2.

Table 2: Nash Policies and Trade and Scale Elasticities

Dep. Var:	$\log\left(\frac{1}{1-s_i^j}\right)$ -dual	$\log(t_{in}^j)$ -dual	$\log(t_{in}^j)$ -trade war
	(1)	(2)	(3)
$\log(\theta_j)$	-0.01 (0.01)	-0.15*** (0.01)	0.09*** (0.03)
$\log(\psi_j)$	0.12*** (0.01)	0.20*** (0.02)	0.64*** (0.03)
Exp. FE	✓	✓	✓
Imp. FE		✓	✓
# obs	154	924	923
R^2	0.63	0.30	0.31

Note: “trade war” 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.

First, Nash industrial subsidies under dual competition are positively correlated with scale elasticities (ψ_j) but not correlated with trade elasticities (θ_j). As discussed in [Lashkaripour and Lugovskyy \(2023\)](#), industrial subsidies are mainly used to address misallocation across sectors: subsidizing sectors with larger ψ_j could lead to greater gains from the increase in production scale, whereas the terms-of-trade effect is not a concern in determining Nash industrial subsidies. Our numerical result about Nash industrial subsidies is consistent with theoretical results in [Lashkaripour and Lugovskyy \(2023\)](#).

Second, Nash tariffs (under trade war and under dual competition) are positively correlated with scale elasticities, (ψ_j). This is due to the home-market effect: the increase in import tariffs would induce more domestic production, the benefit of which is larger in sectors that have stronger increasing returns to scale. Notice that the positive correlation between Nash tariffs and scale elasticities is much stronger under trade wars than that under dual competition (comparing coefficients in Column (2) and (3) in Table 2). When countries can compete via industrial policies, they mainly rely on industrial subsidies to address misallocation of sectoral production led by external economies of scale. In contrast, when they can only manipulate tariffs, tariffs have to be

prohibitively high on sectors with strong increasing returns to scale.

Third, Nash tariffs under dual competition are negatively correlated with trade elasticities, (θ_j) . As figured out by Eaton and Kortum (2002), smaller θ_j indicates stronger within-sector comparative advantage and thereby larger terms-of-trade effects of tariffs. Under dual competition, tariffs are used to address terms-of-trade effects and thereby negatively correlated with (θ_j) . Notably, Nash tariffs under trade war is slightly positively correlated with θ_j . This is because when industrial subsidies are absent, tariffs have to be used to address misallocation across sectors. In this case, the terms-of-trade effect is not the primary concern in determining Nash tariffs.

How do Nash tariffs and dual policies, summarized in Table 2, affect labor allocation across sectors? Figure 6 compares changes in production shares under global trade war and under dual competition. The results indicate that dual competition significantly impacts the production shares of most tradable sectors more than the trade war. This implies that, compared to dual policies, tariffs are a less effective tool for shifting labor towards sectors with large economies of scale.

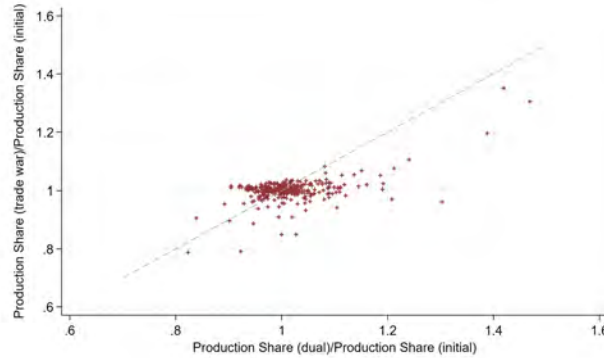


Figure 6: Effects of Nash Tariffs and Dual Policies on Sectoral Production Shares

(Notes: Sectoral production share is defined as $\frac{\sum_{n=1}^N X_{in}^j}{\sum_{s=1}^J \sum_{n=1}^N X_{in}^s}$. Non-tradable sectors are excluded. The dash line is the 45-degree line.)

Table 2 also indicates connections between Nash tariffs and industrial subsidies. To see it more clearly, we depict Nash tariffs and industrial subsidies in Figure 7 and find significantly positive correlation between these two policy tools. There are two driving forces behind this positive correlation. First, scale elasticities push Nash tariffs and industrial policies under dual competition towards the same direction (comparing column (1) with (2) in Table 2). Second, θ_j and ψ_j is

negatively correlated under the calibration in [Lashkaripour and Lugovsky \(2023\)](#) (see Appendix Table A.1). Smaller θ_j results in larger Nash tariffs and thereby the positive correlation between Nash tariffs and industrial subsidies.

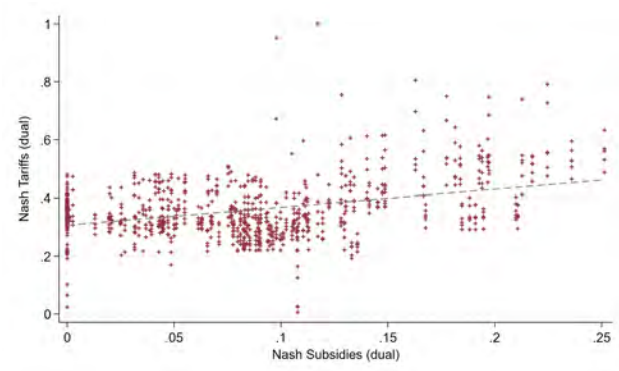


Figure 7: Nash Tariffs and Subsidies under Dual Competition

(Notes: Regressing Nash tariffs on Nash subsidies (controlling for importer fixed effects) leads to a coefficient 0.61 with standard error 0.05.)

4.2 Trade and Industrial Policy Competitions

In this section, we investigate the interactions of trade and industrial policies in the global competition. To this end, we compare Nash tariffs under trade war with those under dual competition. The results are depicted by Figure 8. It suggests that when industrial policy competition is allowed, countries have much weaker incentives to impose prohibitive non-cooperative tariffs: the simple average of Nash tariffs under dual competition is 35%, much lower than 42% under trade war. As shown in Table 2, Nash tariffs have a much larger positive correlation with scale elasticity ψ_j under trade war than that under dual competition. Figure 8 confirms this pattern: Countries impose prohibitive Nash tariffs on industries with ψ_j above the median under trade war. A policy implication of this result is that when industrial policies are not considered, the global competition tends to result in high tariffs on industries with strong increasing returns to scale. Recent examples in the US-China trade war starting from 2018 include semiconductors, electric vehicles, and photovoltaic.

What are the welfare consequences of global competition when tariffs and industrial policies




Figure 8: Nash Tariffs under Trade War vs. under Dual Competition

(Notes: Simple average of Nash tariffs under trade war is 42%, whereas simple average of Nash tariffs under dual competition is 35%.)

are combined in the policy consideration? Table 3 reports welfare effects of the US-China Nash policies and of the global Nash policies. The results suggest that comparing with trade wars, dual competition via the combination tariffs and industrial policies leads to larger welfare gains (or lower welfare losses). If a country is able to efficiently subsidize its high-tech industries with strong economies of scale, it has lower incentives to impose high import tariffs on these industries, as shown in Figure 8. As a result, dual competition tends to result in fewer distortions and thereby better welfare outcomes than tariff wars (Column (4) in Table 3).

Table 3: Welfare Effects of Nash Policies

	China and US ($\Delta\%$)		World ($\Delta\%$)	
	Trade war (1)	Dual (2)	Trade war (3)	Dual (4)
United States	-0.07	1.23	-1.01	0.26
China	-0.33	3.17	-2.56	0.52
European Union	0.02	-0.49	-1.80	-0.27
Japan	0.03	-0.57	-2.21	-0.42
India	0.01	-0.51	-1.84	0.27
Brazil	0.00	-0.05	-2.08	0.23
Rest of the World	0.05	-0.41	-2.33	-0.62

Note: “Trade war” 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.

Comparing Column (3) and (4) in Table 3, we observe that industrial policy competition generally leads to welfare gains for participants.⁸ The positive welfare impact of non-cooperative industrial policies, achieved by directing labor into high-economy-of-scale sectors, outweighs the negative welfare impact caused by terms-of-trade deterioration, *i.e.* the “immiserizing effect” of non-cooperative industrial policies highlighted by [Lashkaripour and Lugovsky \(2023\)](#).

4.3 Extension: Cooperative Policies

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

$$\hat{W} \equiv \sum_{n=1}^N \bar{\omega}_n \left(\frac{\hat{Y}_n}{\hat{P}_n} \right), \quad \bar{\omega}_n \equiv \frac{Y_n}{\sum_{k=1}^N Y_k}. \quad (20)$$

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 (20).

Our first finding is that global cooperation leads to very low tariffs: the average tariff under trade cooperation is 3.31%; and it will reduce to 2.23% if both tariffs and industrial policies are considered in global cooperation. Cooperative tariffs on many industries are close to zero. In sum, global cooperation tends to make the world tariff-free.

⁸Appendix Table B.1 reports welfare effects of industrial subsidy competitions. The results show that all major economies gain from global industrial subsidy competition.

We then investigate the correlation between cooperative policies and trade and scale elasticities. The regression results are reported in Table 4. We find that if tariff is the only policy tool for global cooperation, tariffs are lower in industries with higher scale elasticities (ψ_j) and lower trade elasticities (θ_j). Cooperative tariffs thus reallocate factors of production from weak-increasing-return-to-scale industries to strong-increasing-return-to-scale industries and result in welfare gains. When industrial policies are used for global cooperation, this reallocation is mainly achieved by industrial subsidies which are much larger in industries with higher ψ_j and lower θ_j .

Table 4: Cooperative Policies and Trade and Scale Elasticities

Dep. Var:	$\log(t_{in}^j)$ -coop	$\log(t_{in}^j)$ -dual coop	$\log\left(\frac{1}{1-s_i^j}\right)$
	(1)	(2)	(3)
$\log(\theta_j)$	1.83*** (0.50)	-0.33*** (0.10)	-0.07*** (0.01)
$\log(\psi_j)$	0.99 (0.60)	0.28** (0.13)	0.13*** (0.01)
Exp. FE	✓	✓	✓
Imp. FE	✓	✓	
# obs	229	525	154
R^2	0.24	0.21	0.80

Finally, We quantify welfare consequences of cooperative policies and report the results in Table 5. Cooperative policies improve welfare in most of the major economies, except for India under dual cooperation. For all major economies except India, dual cooperation results in larger welfare gains than tariff cooperation. Comparing the welfare effects reported in Table 3 for Nash policies and in Table 5 for cooperative policies, we can figure out large welfare gains by pushing the world from non-cooperative competition to global cooperation.

4.4 Extension: Imperfect Implementation of Industrial Subsidies

In section 4.2, we have shown that most of the major economies, including China, gain 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

Table 5: Welfare Effects of Cooperative Policies ($\Delta\%$)

	Tariff	Dual
	(1)	(2)
United States	0.19	2.82
China	1.13	4.29
European Union	0.19	1.78
Japan	0.37	1.39
India	2.34	-0.27
Brazil	0.08	3.46
Rest of the World	0.06	1.88

competition. Instead, China chooses from $[0.1s_{CHN}^{j*}, 1.9s_{CHN}^{j*}]$, following a uniform distribution, for all tradable sector j .

We randomly draw China's subsidies for 1000 times. Figure 9 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 from global industrial policy competition.⁹ Once deviate from its optimal values, China's industrial subsidies tend to result in small welfare gains, or in many cases, welfare losses.




Figure 9: China's Welfare Change (%) under Randomly Drawn Industrial Subsidies

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

⁹See Appendix Table B.1 for welfare effects of global industrial policy competition.

5 Conclusion

In this paper, we introduce a deep learning framework, DL-opt, designed to determine optimal policies within quantitative trade and spatial models. This framework integrates the nested fixed point (NFXP) algorithm, automatic implicit differentiation, and best-response dynamics. Our approach efficiently computes high-dimensional optimal policies in trade models with intricate equilibrium systems, making it particularly suitable for analyzing the incentives and consequences of contemporary global trade and industrial policy competitions.

We apply our framework to calculate optimal trade and industrial policies within a multi-country-multi-sector general equilibrium trade model that incorporate input-output linkages and sectoral scale economies. Our counterfactual analysis highlights the importance of sectoral heterogeneity and policy combinations in understanding global economic competition.

Our method is broadly applicable within trade and spatial economics. It effectively overcomes the curse of dimensionality, enabling the computation of high-dimensional optimal (continuous) policies in models where numerous agents, potentially regions, interact. Applications of this method include, but are not limited to, optimizing carbon emissions, corporate taxes, and innovation subsidies across different spatial contexts.

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*.
- BARTELME, D., A. COSTINOT, D. DONALDSON, AND A. RODRIGUEZ-CLARE (2021): “The Textbook Case for Industrial Policy: Theory Meets Data,” *Mimeo*.
- DEKLE, R., J. EATON, AND S. KORTUM (2008): “Global rebalancing with gravity: Measuring the burden of adjustment,” *IMF Staff Papers*, 55, 511–540.

- EATON, J. AND S. KORTUM (2002): “Technology, Geography, and Trade,” *Econometrica*, 70, 1741–1779.
- FENG, Z., J. HAN, AND S. ZHU (2023): “Optimal Taxation with Incomplete Markets: an Exploration via Reinforcement Learning,” *Mimeo*.
- FROSTIG, R., M. J. JOHNSON, AND C. LEARY (2018): “Compiling machine learning programs via high-level tracing,” *Systems for Machine Learning*, 4.
- HAN, J., W. E. AND Y. YANG (2021): “DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks,” *Mimeo*.
- 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.
- 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.
- 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.

RUST, J. (2000): “Nested Fixed Point Algorithm,” *mimeo*.

SUN, J. (2023): “Continuation Value Is All You Need: A Deep Learning Method for Solving Heterogeneous-Agent Models with Aggregate Uncertainty,” *Mimeo*.

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, Calibration, and Computation

A.1 Automatic Differentiation: A Toy Example

We consider function $y = f(x_1, x_2) = [x_1^2 + x_1/x_2 - \exp(x_2)] [x_1/x_2 - \exp(x_2)]$, aiming to compute $\frac{\partial y}{\partial x_i}$ at point (\bar{x}_1, \bar{x}_2) .

We can use finite difference to approximate $\frac{\partial y}{\partial x_1}$: $\frac{\partial y}{\partial x_1} \simeq \frac{f(\bar{x}_1+h, \bar{x}_2) - f(\bar{x}_1, \bar{x}_2)}{h}$. Notice that to compute $\frac{\partial y}{\partial x_1}$ and $\frac{\partial y}{\partial x_2}$, we need to evaluate function $f(x_1, x_2)$ at least three times, including initial evaluation of $f(\bar{x}_1, \bar{x}_2)$.

Automatic differentiation, instead, introduces a series of intermediate variables: $v_{-1} = x_1$, $v_0 = x_2$, $v_1 = v_{-1}/v_0$, $v_2 = v_{-1}^2$, $v_3 = \exp(v_0)$, $v_4 = v_1 - v_3$, $v_5 = v_2 + v_4$, and $v_6 = v_4 \cdot v_5 = y$. The relationship between (x_1, x_2) and y , through these intermediate variables, can be illustrated by the computation graph in Figure A.1:




Figure A.1: Computation Graph of Automatic Differentiation: Toy Example

By the chain rule, we have

$$\begin{aligned} \frac{\partial y}{\partial x_1} &= \frac{\partial y}{\partial v_6} \left(\frac{\partial v_6}{\partial v_4} \frac{\partial v_4}{\partial v_1} \frac{\partial v_1}{\partial v_{-1}} + \frac{\partial v_6}{\partial v_5} \left(\frac{\partial v_5}{\partial v_4} \frac{\partial v_4}{\partial v_1} \frac{\partial v_1}{\partial v_{-1}} + \frac{\partial v_5}{\partial v_2} \frac{\partial v_2}{\partial v_{-1}} \right) \right) \frac{\partial v_{-1}}{\partial x_1}, \\ \frac{\partial y}{\partial x_2} &= \frac{\partial y}{\partial v_6} \left(\frac{\partial v_6}{\partial v_4} \left(\frac{\partial v_4}{\partial v_1} \frac{\partial v_1}{\partial v_0} + \frac{\partial v_4}{\partial v_3} \frac{\partial v_3}{\partial v_0} \right) + \frac{\partial v_6}{\partial v_5} \frac{\partial v_5}{\partial v_4} \frac{\partial v_4}{\partial v_1} \frac{\partial v_1}{\partial v_0} \right) \frac{\partial v_0}{\partial x_2}. \end{aligned} \quad (\text{A.1})$$

Notice that there are many common terms in computing $\frac{\partial y}{\partial x_1}$ and $\frac{\partial y}{\partial x_2}$. Therefore, we do not need to evaluate $f(x_1, x_2)$ twice. We only need to compute the following terms: $\frac{\partial y}{\partial v_6}$, $\frac{\partial v_{-1}}{\partial x_1}$, $\frac{\partial v_0}{\partial x_2}$, $\frac{\partial v_6}{\partial v_4}$, $\frac{\partial v_6}{\partial v_5}$, $\frac{\partial v_5}{\partial v_4}$, $\frac{\partial v_5}{\partial v_2}$, $\frac{\partial v_4}{\partial v_3}$, $\frac{\partial v_4}{\partial v_1}$, $\frac{\partial v_3}{\partial v_0}$, $\frac{\partial v_2}{\partial v_{-1}}$, $\frac{\partial v_1}{\partial v_0}$, $\frac{\partial v_1}{\partial v_{-1}}$. Calculating the gradients for these elementary functions is much more efficient than evaluating the entire function multiple times.

A.2 “Exact-hat” algebra

Changes in unit costs can be expressed as

$$\hat{c}_i^j = \frac{1}{(\hat{L}_i^j)^{\psi_j}} \hat{w}_i^{\beta_i^j} \left[\prod_{s=1}^J (\hat{P}_i^s)^{\gamma_i^j} \right]^{1-\beta_i^j}. \quad (\text{A.2})$$

Changes in trade share:

$$\hat{\pi}_{in}^j = \frac{\left[\hat{c}_i^j \widehat{1 + t_{in}^j} \widehat{1 + e_{in}^j} \right]^{-\theta_j}}{(\hat{P}_n^j)^{-\theta_j}}. \quad (\text{A.3})$$

Changes in price indices:

$$\hat{P}_n^j = \left[\sum_{i=1}^N \pi_{in}^j \left[\hat{c}_i^j \widehat{1 + t_{in}^j} \widehat{1 + e_{in}^j} \right]^{-\theta_j} \right]^{-\frac{1}{\theta_j}}. \quad (\text{A.4})$$

Changes in sectoral wage incomes:

$$\hat{w}_i \hat{L}_i^j w_i L_i^j = \beta_i^j \sum_{n=1}^N \frac{\hat{\pi}_{in}^j \hat{X}_n^j X_{in}^j}{(1 + t_{in}^j)' (1 + e_{in}^j)'}. \quad (\text{A.5})$$

Changes in sectoral labor allocation satisfy:

$$\sum_{j=1}^J \hat{L}_i^j L_i^j = \bar{L}_i. \quad (\text{A.6})$$

Changes in the total income:

$$\hat{Y}_i Y_i = \hat{w}_i w_i \bar{L}_i + \sum_{j=1}^J \sum_{n=1}^N \frac{(e_{in}^j)'}{(1 + e_{in}^j)'} (X_{in}^j)' + \sum_{j=1}^J \sum_{k=1}^N \frac{(t_{ki}^j)'}{(1 + t_{ki}^j)' (1 + e_{ki}^j)'} (X_{ki}^j)'. \quad (\text{A.7})$$

Changes in sectoral expenditure:

$$\hat{X}_i^j X_i^j = \alpha_i^j \hat{Y}_i Y_i + \sum_{s=1}^J (1 - \beta_i^s) \gamma_i^{js} \sum_{n=1}^N \frac{(X_{in}^s)'}{(1 + t_{in}^s)' (1 + e_{in}^s)'}. \quad (\text{A.8})$$

Changes in aggregate price indices:

$$\hat{P}_n = \prod_{j=1}^J (\hat{P}_n^j)^{\alpha_n^j}. \quad (\text{A.9})$$

Optimal policies are solved by

$$\max_{\{t_{in}^j, e_{in}^j, \hat{W}_i, \hat{L}_i^j, \hat{P}_i^j, \hat{X}_i^j\}_{i,j}} \hat{W}_n \equiv \frac{\hat{Y}_n}{\hat{P}_n}, \quad \forall n = 1, 2, \dots, N \quad (\text{A.10})$$

s.t. Equation (A.4), (A.6), (A.5), and (A.8)

A.3 Calibration of (ψ_j, θ_j)

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

A.4 Adaptive Moment Estimation (ADAM) Algorithm

The Adaptive Moment Estimation (ADAM, see [Kingma and Ba \(2014\)](#)) algorithm 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:

Table A.1: Calibration of (ψ_j, θ_j) from [Lashkaripour and Lugovsky \(2023\)](#)

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.

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:

$$a^{t+1} = a^t - \gamma \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \varepsilon}, \quad (\text{A.11})$$

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.

B Quantification

B.1 The “Internal Cooperation” Assumption in [Lashkaripour and Lugovskyy \(2023\)](#)

Figure B.1 suggests that the “internal cooperation” assumption made by [Lashkaripour and Lugovskyy \(2023\)](#) is restrictive in solving for the fully optimal policies: the relative wages among other major economies do change under China’s unilaterally 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.




Figure B.1: Changes in Relative Wage under China’s Unilaterally Optimal Policies (in percentage)

(Notes: Here, to compare with the results in [Lashkaripour and Lugovskyy \(2023\)](#), we consider China’s unilaterally optimal import tariffs, export tariffs, and industrial subsidies. We normalize the change in nominal wage in the rest of the world (ROW) as 0.)

B.2 Welfare Effects of Industrial Subsidy Competition

This section reports welfare effects of industrial subsidy competitions, as a supplement to the welfare results in Table 3.

Table B.1: Welfare Effects of Industrial Subsidy Competitions

	China and US ($\Delta\%$)	World ($\Delta\%$)	
	Subsidy (1)	Subsidy (2)	Subsidy-Uni (3)
United States	1.28	1.31	1.31
China	3.31	2.42	1.80
European Union	-0.51	1.03	0.53
Japan	-0.61	1.14	0.30
India	-0.50	2.56	0.60
Brazil	-0.05	1.97	1.71
Rest of the World	-0.44	1.47	1.76

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). “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.

B.3 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.2 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.2, import tariffs under unilaterally optimal policies, as illustrated by Figure B.3, are also almost identical across industries and economies.

Table B.2 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.

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 com-

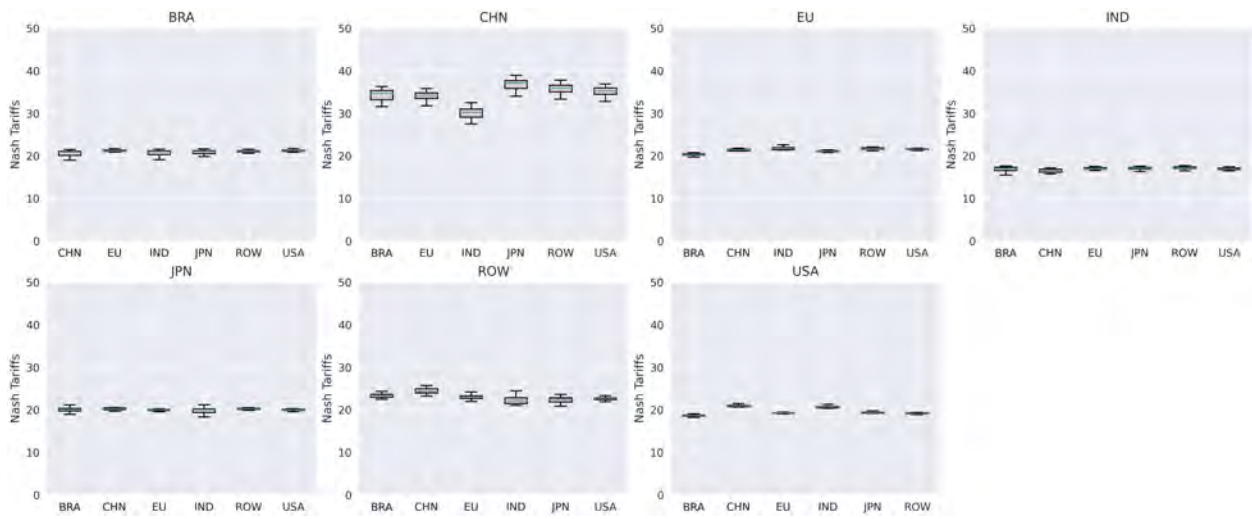


Figure B.2: Nash Tariffs for Global Dual Policy Competition without Scale Economies




Figure B.3: Import Tariffs in China's Unilaterally Optimal Dual Policies Without Scale Economies

petition 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.

Table B.2: Welfare Effects of Nash Policies: $\psi = 0$

	China and US ($\Delta\%$)			World ($\Delta\%$)		
	Subsidy (1)	Tariff (2)	Dual (3)	Subsidy (4)	Tariff (5)	Dual (6)
Brazil	-0.0264	0.0014	-0.0248	0.0102	-1.1504	-1.1277
China	0.1262	-0.0488	0.0713	0.0445	-0.5856	-0.5247
European Union	-0.0228	0.0022	-0.019	-0.0492	-1.0517	-1.0265
India	-0.0152	0.0063	-0.0057	-0.0951	-1.6609	-1.6261
Japan	-0.0385	0.0012	-0.0368	-0.0757	-1.1136	-1.0731
Rest of the World	-0.1035	-0.0004	-0.1039	-0.0794	-1.8126	-1.832
United States	-0.013	-0.0362	-0.0467	-0.0502	-0.9101	-0.8904