Import Competition, Innovation, and the Cost of Protectionism

Deniz Atalar*
UNIVERSITY OF CAMBRIDGE

October 29, 2025

Most Recent Draft Here

Abstract

How does import protectionism affect catch-up innovation and welfare? I develop a small open-economy model in which trade costs shape buyer–supplier relationships, and suppliers' incentives to innovate. When trade costs rise, domestic buyers become more inclined to shift sourcing from foreign to domestic suppliers. The possibility of this shift triggers heterogeneous innovation responses among domestic suppliers: those suppliers that are less technologically advanced than their foreign competitors increase their innovation, while other, more advanced suppliers reduce it. I verify and quantify the buyers' shift in sourcing as well as the suppliers' heterogeneous innovation responses to this shift using novel Turkish firm–product-level data on firm-to-firm transactions, imports, production, and innovation expenditures. I build on these firm level responses to assess the impact of import protectionism on welfare. If domestic suppliers in the aggregate respond to a rise in trade costs by increasing innovation, the welfare losses of such a policy are partially mitigated. If, on the other hand, they respond by decreasing innovation, the welfare losses are amplified. Calibrating the model to Turkish microdata, I find that a 10 percent rise in trade costs in Turkiye triggers an increase in aggregate innovation, which mitigates roughly one-quarter of the welfare loss relative to a no-innovation benchmark.

Keywords: *Import Competition, Innovation, Trade, Technological Upgrading*
JEL Classification: *F14; F61; F62; O31.*

*Email: da497@cam.ac.uk. I am deeply grateful to my supervisor Vasco M. Carvalho and my advisor Meredith A. Crowley for their guidance and support throughout this project. I am further grateful to Banu Demir for her many helpful suggestions. I thank Treb Allen, Giammario Impullitti, Joris Hoste, Sarur Chaudhary, Fred Maeng, Ramya Raghavan, Chuan-Han Cheng and Megan Haasbroek for their suggestions and feedback on this project. This paper has benefited from comments at the University of Cambridge Empirical Micro Workshop, Macro Workshop, and Trade Workshop, as well as at the Dartmouth Globalization Seminar and the CEP–Warwick Junior Trade Workshop. I gratefully acknowledge financial support from two separate Keynes Fund grants.

1 Introduction

Import protection is once again on the rise. Over the past five years, more than one-third of countries have significantly increased protectionist measures (Financial Times, 2024). The number of new discriminatory policies rose sharply from 951 in 2020 to 1,869 in 2024 (Global Trade Alert, 2025). Wary of the pitfalls of protectionism, proponents of conventional trade theory have long cautioned that protectionist policies raise import prices and shift production toward less efficient firms, generating welfare losses (Dornbusch et al., 1977). But the efficiency of domestic firms is not fixed. Import protectionism prompts importers to rewire their supply chains toward domestic suppliers, creating new opportunities and competitive pressures. This, in turn, reshapes domestic suppliers' incentives to innovate, thereby altering their efficiency. The real welfare cost of protectionism, then, not only arises from the static misallocation of production that such measures create, but also from how domestic suppliers adjust their innovation in response.

To assess this cost, it is thus necessary to address the question of how protectionist trade policies shape firms' innovation incentives. Two opposing views prevail. Proponents of protectionism argue that trade costs can promote innovation by giving domestic firms the opportunity to catch up (Turkish Presidency, 2024). Whereas opponents argue that it is openness that promotes innovation by adding competitive pressures that force firms to improve (World Trade Organization, 2023). Faced with these contradictory views, policymakers are left with considerable uncertainty. Does protectionism foster or hinder innovative activity, and by how much?

In this paper, I answer this question and use the results to assess the welfare impact of protectionism. I begin by developing a small open-economy model in which trade costs rewire buyer–supplier relationships and, in turn, shape suppliers' incentives to innovate. The model shows that when trade costs rise, domestic buyers shift sourcing from foreign to domestic suppliers. Anticipating this shift, domestic suppliers respond heterogeneously: technological laggards increase innovation to capture new market opportunities, while leaders reduce it as competitive pressure diminishes. I then test the model's predictions using novel Turkish firm–product–level records that link firm-to-firm transactions, imports, production, and innovation expenditures. These data provide a uniquely detailed view of how trade costs reshape sourcing patterns and trigger heterogeneous innovation responses—the micro-level evidence needed to quantify welfare implications. Finally, I calibrate the model to these data and find that a 10 percent increase in trade costs in Turkiye raises aggregate innovation, which offsets roughly one-quarter of the welfare loss relative to a counterfactual without innovation responses.

To formalize these ideas, I develop a small open-economy model that links trade costs, downstream buyers' sourcing decisions, and upstream suppliers' incentives to innovate. A downstream buyer produces a final good using a continuum of differentiated inputs. For each input variety, the buyer sources from either a domestic or foreign supplier, with trade costs shaping the relative attractiveness of each option. This creates a winner-takes-all environment: for each input, the supplier offering the lowest quality-adjusted price captures the buyer's demand for that variety. Domestic suppliers can lower their relative prices through technological upgrading, knowing that their ability to attract or retain downstream demand depends on their position vis-à-vis foreign rivals. This exclusivity generates strong incentives to innovate and provides the microfoundation for how downstream demand reallocations translate into incentives to innovate by upstream suppliers.

A central theoretical prediction of the model is an inverted-U relationship between a domestic supplier's initial productivity and its investment in innovation—a pattern consistent with earlier evidence on innovation and competition in closed-economy settings (Aghion et al., 2005) and recent extensions to international environments (Akcigit et al., 2025). Low-productivity suppliers recognize that their chances of attracting downstream buyers away from foreign competitors are small, so the expected gains from upgrading are limited and their incentives to innovate are weak. High-productivity suppliers, already preferred by downstream buyers, gain little from further improvement and therefore have similarly weak incentives to innovate. The strongest incentives arise for suppliers in the middle of the productivity distribution—those close to their foreign competitors—because a marginal improvement in quality-adjusted prices can determine whether they secure downstream demand. This demand-based mechanism concentrates innovation among suppliers that are "neck-and-neck" with their foreign rivals, linking trade exposure to heterogeneous innovation responses.

Changes in trade costs alter domestic suppliers' incentives to innovate. When trade costs rise, domestic suppliers that were initially more productive than their foreign rivals become even more attractive to downstream buyers. This makes it easier for them to secure demand, reducing their incentive to further improve. In contrast, less productive suppliers find themselves closer—in effective terms—to their foreign competitors, which strengthens their incentives to upgrade in order to attract downstream demand. Thus, higher trade costs reallocate innovation incentives across the domestic productivity distribution, with opposing effects depending on a firm's initial position.

This duality helps reconcile opposing views on trade protection. The decline in innovation among leading firms aligns with the concerns of opponents of protectionism, who argue that reduced competition weakens incentives to improve. Meanwhile, the increase in innovation among lagging firms reflects the argument of proponents, who see protection as an opportunity for catch-up. This framework, therefore, generalizes both perspectives, showing that the effect of protectionism on innovation depends on where firms stand in the productivity distribution.

The welfare impact of trade costs thus depends on how they reshape innovation incentives across the distribution of domestic suppliers relative to their foreign competitors. When most domestic suppliers lag behind, higher trade costs create opportunities to capture market share from weakened foreign rivals, spurring catch-up innovation that partly offsets the static welfare losses from distorted resource allocation. Conversely, when most suppliers are technological leaders, higher trade costs reduce competitive pressure, dampening innovation and amplifying welfare losses. The net welfare effect therefore hinges on whether the innovation gains among laggards outweigh the innovation losses among leaders.

To credibly quantify the net welfare effect it is first necessary to test whether the model's mechanisms hold in the data and estimate their magnitudes across the supplier distribution. This requires first validating the model's foundational assumptions—particularly the buyer–supplier structure that generates innovation incentives—and then quantifying the innovation responses across heterogeneous suppliers. Doing so requires exceptionally detailed microdata linking import costs to buyer-supplier transactions and firm-level innovation expenditures—a combination I am able to construct using novel Turkish administrative data. Using these data, I pursue three empirical goals. First, I document that downstream buyers source each product from a single upstream supplier, verifying the competitive environment that shapes innovation decisions. Second, I show that rising trade costs induce buyers to reallocate sourcing from foreign to domestic suppliers, confirming the demand reallocation mechanism. Third, I quantify how domestic suppliers' innovation responds to foreign competition, testing the model's prediction of an inverted-U relationship.

I begin by documenting that downstream buyers in Turkiye typically source each product from a single upstream supplier—either domestic or foreign—consistent with the model's buyer–supplier structure. Using administrative data on firm-to-firm transactions covering the universe of formal firms, merged with customs import records, I find that in more than 70 percent of cases, a downstream buyer purchases a given input from a single supplier, indicating little diversification of input sourcing.[1] Yet these relationships

[1]The VAT data identify the sector of the supplier but not the specific product transacted. Since products are typically more disaggregated than sectors, this estimate likely understates the true share of single-supplier relationships at the product level.

are far from stable: on average, buyers purchase from around twelve manufacturing suppliers and reshuffle about half of them each year—dropping roughly six and adding a similar number.[2] This high turnover underscores that buyer-supplier relationships are short-lived and contestable, validating a central assumption of the model: suppliers can realistically expect innovation to influence buyers' sourcing decisions.

I then show that rising trade costs induce downstream buyers to reallocate sourcing from foreign to domestic suppliers, as predicted by the model. When import prices increase, importing firms reduce purchases from foreign suppliers along both the intensive and extensive margins and expand purchases from domestic suppliers. These adjustments indicate that buyers not only substitute across existing relationships but also reconfigure their supplier networks more fundamentally, consistent with the model's mechanism through which trade shocks shift demand toward domestic suppliers.

Having shown that domestic buyer–supplier relationships adjust in response to import prices, I next turn to suppliers' innovation behavior—the key channel through which these reallocations affect welfare. In the model, the formation of buyer–supplier relationships is the pillar of innovation incentives: suppliers compete to attract downstream demand, and their incentive to innovate depends on their technological proximity to foreign competitors. When this proximity is large—whether because a domestic supplier is far behind or already far ahead—suppliers see little prospect of altering downstream demand in their favor and invest little in innovation. When they are close in proximity, suppliers perceive innovation as a viable way to capture demand from importers, and innovation intensifies. This mechanism generates an inverted-U relationship between innovation and proximity to foreign competitors. The model therefore predicts that this relationship can be tested empirically by relating firms' observed innovation activity to the quality-adjusted price gap between a domestic product and its foreign counterpart.

I therefore construct this measure of proximity and innovation at the firm–product level using administrative and customs data from Turkiye. The data combine the Turkish version of the EU's Prodcom survey (PRODTR), which records firms' output by ten-digit product codes, with detailed customs data on imports. This merge yields a harmonized firm–product panel that captures both domestic and foreign suppliers within narrowly defined product markets, allowing me to compare quality-adjusted prices across them. I measure competition using the quality-adjusted price gap between domestic and foreign suppliers, instrumented with Bartik-style shocks to global export prices that are exogenous to Turkish firms.

To capture firms' investment in innovation, I use administrative employer–employee

[2] Bernard et al. (2022) and Antras et al. (2017) report related findings.

matched data reporting wage payments by occupation. From these, I construct firm-year spending on technology-related labor—wages paid to scientists, engineers, and technical workers—which captures broader forms of catch-up innovation.[3]

Consistent with the model's predictions, I find that innovation expenditure follows an inverted-U relationship with a domestic suppliers proximity to foreign competitors. This pattern emerges visually in binscatter plots and formally in quadratic IV regressions with the peak located near parity between domestic and foreign prices. No comparable pattern appears for non-technical occupations: their wage share declines monotonically with the price gap, consistent with firms pricing below foreign rivals expanding production employment and those pricing above reducing it. This contrast supports the model's prediction that the inverted-U relationship reflects technological upgrading incentives rather than general labor-demand dynamics.

Having established the theoretical mechanism and documented empirical patterns consistent with its predictions, I next quantify its aggregate implications by calibrating the model to rich Turkish firm-level data. The calibration links each structural parameter to directly observed moments in the data, including product-level price dispersion, domestic–import market shares, and firms' employment of technical and production workers. This mapping ensures that the model reproduces key features of the Turkish economy—such as the 43% domestic share in tradable manufacturing and the 14.5% share of R&D labor in total employment—allowing the quantitative exercise to trace how changes in trade costs affect innovation, aggregate productivity, and welfare.

The counterfactual analysis reveals that trade costs always reduce welfare, but the magnitude of the loss depends on how innovation responds to competition. Moderate trade barriers raise aggregate productivity by stimulating catch-up innovation among firms close to the world technology, partially offsetting the static efficiency losses from misallocation. In the Turkish calibration, a 10% increase in all intermediate import prices—relative to free trade— reduces welfare by 3.1%—compared with 4.1% in a no-innovation benchmark—implying that innovation mitigates roughly one-quarter of the static loss. Yet as trade barriers rise further, competitive pressure fades, innovation incentives collapse, and the dynamic gains vanish. Thus, in an economy where firms endogenously innovate, trade protection never improves welfare but moderate measures could cushion its decline.

[3]This measure offers advantages over conventional indicators such as patents or R&D expenditures. Patents primarily capture frontier innovation, which is relatively rare in Turkiye's manufacturing sector, while R&D data from balance sheets can be affected by reporting incentives and subsidy eligibility. Moreover, an extension using R&D spending confirms the same inverted-U pattern.

Literature Review. This paper relates to a growing literature documenting non-monotonic firm responses to trade and competition (e.g., Lim et al., 2022; Fieler and Harrison, 2023; Akcigit et al., 2025). Lim et al. (2022) and Fieler and Harrison (2023) show that trade liberalization generates heterogeneous innovation responses across firms, with competition stimulating innovation for some firms but discouraging it for others. My paper embeds this mechanism in a general-equilibrium setting to study its welfare consequences, showing that the aggregate effects of protectionism depend on the distribution of firms across the technological spectrum.

This paper is most closely related to Akcigit et al. (2025) (hereafter AAI), who integrate the Aghion et al. (2005) inverted-U relationship between competition and innovation into an open-economy framework. Their two-country quantitative model, calibrated to the U.S. economy, provides important insights into how R&D subsidies and protectionist policies shape long-run growth and welfare in a frontier economy.

My paper departs from AAI along three dimensions. First, in terms of results, I show that the welfare effects of protectionism depend on the distribution of firm productivities across domestic and foreign competitors. The model yields closed-form expressions demonstrating that trade protection can either magnify or attenuate welfare losses depending on where domestic firms lie along this distribution.[4] Second, I build a direct bridge between theory and empirics: rather than analyzing long-run growth, my model is constructed to be taken to firm-level data, linking its key mechanisms—competition, reallocation, and innovation—to observable quantities. This design allows me to quantify the strength of these channels empirically using detailed firm–product information. Third, the context differs fundamentally. I study a small open economy where innovation takes the form of within-firm, catch-up upgrading rather than frontier R&D. Using administrative data from Turkiye, I measure firms' spending on technology-related labor, which offers a broader and more appropriate proxy for innovation in this setting than patents or balance-sheet R&D. Together, these departures extend AAI's insights beyond frontier economies and provide the first empirical validation of competition-driven innovation within a unified theoretical and quantitative framework.[5]

My paper also contributes to a large body of work that studies the link between trade policy and innovation. The evidence put forward in this literature often points in oppo-

[4]In Appendix F, AAI explore numerically how alternative technological-gap distributions affect aggregate innovation, but they do not connect this variation to welfare or provide analytical or empirical validation.

[5]Trade may also affect innovation through market size (Lileeva and Trefler, 2010; Coelli, 2022), knowledge spillovers (Bai et al., 2020; Bergeaud et al., 2024; Alvarez and López, 2008; Alvarez et al., 2015; Buera and Oberfield, 2020), input costs (Amiti and Konings, 2007), and comparative advantage (Grossman and Helpman, 1991). For broader reviews, see Melitz and Redding (2021) and Akcigit and Melitz (2022).

site directions. One strand shows that trade restrictions can spur domestic innovation and productivity, particularly when they shelter sectors with learning externalities. Historical episodes such as the Napoleonic Blockade (Juhász, 2018; Juhász et al., 2023) and South Korea's industrial policies (Lane, 2025), as well as modern restrictions like China's Great Firewall (Zhou, 2025) and geopolitical trade barriers (Flynn et al., 2025), demonstrate how limiting foreign competition can encourage R&D and technological upgrading.[6] In contrast, an extensive literature finds that trade liberalization promotes innovation and quality upgrading by intensifying competition and expanding market access (Amiti and Konings, 2007; Bloom et al., 2016; Khandelwal, 2010; Amiti and Khandelwal, 2013; Dang, 2017; Medina, 2017; Lim et al., 2022; Fieler and Harrison, 2023; Piveteau and Smagghue, 2022; Crowley et al., 2024).[7] My paper reconciles these seemingly conflicting findings both empirically and theoretically, and show how these opposing responses can result in different welfare costs for protectionism.[8]

Lastly, my work contributes to the literature on how trade policy reshapes domestic production networks (Arkolakis et al., 2023; Dhyne et al., 2021; Lim et al., 2017; Amiti et al., 2024; Demir et al., 2024; Magerman and Palazzolo, 2024). While this literature focuses on how trade shocks alter the structure and composition of domestic buyer-supplier relationships, I extend it by showing that trade-induced demand reallocation also influences the innovation behavior of upstream suppliers.[9]

Outline. The remainder of the paper is structured as follows. Section 2 develops the theoretical framework linking trade and innovation. Section 3 presents the empirical analysis. Section 4 calibrates the model to Turkish firm-level data and conducts counterfactual simulations. Section 5 concludes.

[6] Other studies consist of but are not limited to Autor et al. (2020); Hashmi (2013); Hombert and Matray (2018)

[7] Other studies consist of but are not limited to Feenstra and Hanson (1996); Bernard et al. (2006); Teshima (2008); Goldberg et al. (2010); Kueng et al. (2017); Fernandes and Paunov (2013); Chen and Juvenal (2022); Medina (2024).

[8] Only a few studies provide direct empirical evidence of such non-monotonic responses. Aghion et al. (2009, 2018, 2024a,b) document inverted-U effects of competition and trade on innovation using patent data, but these measures capture frontier innovation rather than the within-firm technological upgrading that dominates developing economies. The exception, Amiti and Khandelwal (2013), studies process innovation via imported inputs but does not link these effects to welfare. My paper complements this literature by providing both micro-level evidence on non-frontier innovation and a tractable framework that connects these responses to aggregate welfare. In contrast to their focus on export-oriented firms, I use rich Turkish firm-level data that capture a broader set of producers competing in the domestic market. I measure adjustments in technical employment, reflecting within-firm innovation activity, rather than in patent counts or product quality.

[9] This work builds on a broader literature on production networks, which examines how shocks propagate through supply chains; see Carvalho et al. (2021) for a review. Relatedly, McNerney et al. (2022) show that technological improvements can also propagate along production chains, amplifying growth and structural change.

2 Theory

I develop a small open-economy model to evaluate how protectionist policies affect welfare when domestic suppliers can undertake cost-reducing innovation.[10] The economy features a downstream buyer that assemble a differentiated final good using a constant elasticity of substitution (CES) aggregator of intermediate varieties. For each variety, a domestic and a foreign supplier compete 'a la Bertrand to sell to downstream firms. Domestic intermediate producers can invest in process innovation to lower their marginal costs and improve their chances of being chosen as suppliers. Changes in trade costs shift relative prices, thereby altering the intensity of foreign competition and reshaping the sourcing decisions of downstream buyers. These sourcing reallocations, in turn, determine suppliers' incentives to innovate. Aggregating across firms yields predictions for productivity, innovation, and welfare. The subsections that follow present the setup, characterize equilibrium behavior, and derive aggregate and welfare implications.[11]

2.1 Setup

2.1.1 Representative Household

In this model, there is a Home country H that is infinitesimal relative to the rest of the world W. The analysis focuses on the behavior of Home agents and firms, taking world prices and wages as given.

The Home economy is populated by a representative household endowed with one unit of labor, a fraction ℓ_H^P of which is employed in production at wage w_H^P, and the remaining share $\ell_H^R = 1 - \ell_H^P$ engaged in technological upgrading at wage w_H^R.

The household consumes two goods, both freely traded: a homogeneous good Z_H, priced at p_z, and a differentiated good $\mathcal{D}_H$, with its price normalized to one. Preferences are Cobb—Douglas, so a share α of income is spent on Z_H and the remaining share $1-\alpha$ on $\mathcal{D}_H$, such that:

$$Z_H = \frac{\alpha I_H}{p_z}, \qquad \mathcal{D}_H = (1-\alpha) I_H,$$

where I_H denotes income (see Appendix A.1 for details).

The homogeneous good is produced solely with labor in both Home and the rest of the world and is freely traded. Let o_W denote labor productivity in the world and w_W^p the world wage. Since Home is infinitesimal, the world determines the price of the homoge-

[10]Related models include Akcigit et al. (2025), Bustos (2011), and Helpman (2024).

[11]In the main text, I focus on process innovation (reductions in unit costs). Appendix B shows that a product/quality innovation version is isomorphic, so the results carry over.

neous good, which pins Home's production wage:

$$w_H^p = \frac{1}{o_W} w_W^p, \qquad p_z = w_H^p.$$

Throughout, w_W^p and o_W are taken as exogenous.

2.1.2 Representative Downstream Producer

Home hosts a representative downstream producer that uses a country-specific fixed factor F_H (e.g., land) and a continuum of intermediate inputs indexed by $j \in [0,1]$ in production. For each variety, a domestic and a foreign supplier compete to serve this buyer, and the downstream producer sources from whichever offers the lower quality-adjusted delivered price.

Production follows a CES aggregator:

$$Y_H \;=\; \frac{F_H^{\beta}}{1-\beta} \int_0^1 \left(k_H^{\frac{\beta}{1-\beta}} \, q_{Hj} + k_W^{\frac{\beta}{1-\beta}} \, q_{Wj} \right)^{1-\beta} dj, \qquad 0 < \beta < 1, \tag{1}$$

where q_{Hj} and q_{Wj} denote the quantities of variety j purchased from the Home and World suppliers, respectively. The parameters k_H and k_W represent country-level quality shifters that scale the effective productivity of intermediates. The fixed factor F_H is normalized to 1 for simplicity. The parameter β governs both the Cobb–Douglas share of the fixed factor and the CES elasticity of substitution $1/\beta$ across intermediate inputs.[12]

Downstream demand for inputs. The downstream producer rents the fixed factor F_H at rate f_H and chooses inputs from domestic and foreign suppliers to maximize profits. Profit maximization delivers two conditions.

First, the rental rate of the fixed factor is proportional to output:

$$f_H = \frac{\beta Y_H}{F_H}. \tag{2}$$

Second, demand for each intermediate variety j is iso-elastic in its quality-adjusted delivered price. Let p_{Hj} and p_{Wj} denote the pre-trade-cost prices offered by Home and World suppliers, respectively, and let $\tau_H \geq 0$ denote the iceberg trade cost on imported

[12] Following Akcigit et al. (2025), I equate these parameters to simplify notation; relaxing this assumption does not alter the qualitative results.

intermediates. Then,

$$q_j = \begin{cases} k_H \left(p_{Hj}\right)^{-1/\beta}, & \text{if } \dfrac{p_{Hj}}{k_H^{\frac{\beta}{1-\beta}}} \leq (1+\tau_H)\,\dfrac{p_{Wj}}{k_W^{\frac{\beta}{1-\beta}}}, \\ k_W \left((1+\tau_H)p_{Wj}\right)^{-1/\beta}, & \text{otherwise.} \end{cases} \tag{3}$$

That is, the downstream producer sources each input from the supplier offering the lowest quality-adjusted delivered price, with trade costs raising the effective price of imported intermediates.

To capture the empirical setting where Home (e.g., Turkiye) primarily exports homogeneous goods rather than differentiated intermediates, I assume that the World imposes a prohibitive tariff on differentiated goods, $\tau_W = \infty$. This implies that World downstream producers purchase only from World intermediates, while Home downstream producers choose between domestic and imported inputs. Consequently, Home intermediate firms face import competition but do not export, consistent with the stylized trade structure motivating the analysis. In contrast, the homogeneous good is freely traded and constitutes Home's only export to the World.

2.1.3 Upstream Supplier's Problem

A Home upstream supplier is characterized by its *post-innovation* labor productivity $\phi_H \in (0, \infty)$.[13] Output is linear in efficiency units of labor, so the marginal cost is

$$mc_H(\phi_H) = \frac{w_H^p}{\phi_H},$$

where w_H^p is the Home production wage.

Market structure. For each variety j, a Home supplier with productivity ϕ_H competes with a World supplier with productivity ϕ_W to serve the Home downstream producer's demand. Competition follows a two-stage Bertrand game. In stage 1, each supplier observes its own and its foreign competitor's productivity and decides whether to remain active by paying a small fixed cost $\varepsilon > 0$. In stage 2, the active suppliers set prices, and the downstream producer sources from the one offering the lowest quality-adjusted delivered price. Since each supplier's marginal cost and quality are common knowledge, the supplier with the higher delivered cost (after trade costs) anticipates being undercut

[13] Each supplier's initial productivity is drawn from a Fréchet distribution and may be upgraded through innovation; ϕ_H here denotes productivity *after* innovation. See Section 2.1.4 for details.

and exits in stage 1, so only the cost leader survives.[14]

Pricing. Delivered marginal costs are w_H^p/ϕ_H for Home suppliers and $(1+\tau_H)w_W^p/\phi_W$ for World suppliers. Under CES demand, an active supplier charges a constant markup over marginal cost:

$$p^*(\phi_H) = \frac{1}{1-\beta}\,\frac{w_H^p}{\phi_H}. \tag{4}$$

An analogous expression holds for a World supplier with wage w_W^p and productivity ϕ_W.

Profits. To characterize supplier profits, it is convenient to introduce a few composite terms that simplify the notation. The constant

$$\kappa = \frac{\beta}{1-\beta}\,(1-\beta)^{1/\beta}$$

collects demand-side parameters from CES aggregation. The parameter

$$\delta = \frac{1}{\beta} - 1 > 0$$

captures the elasticity of substitution across intermediate varieties: a higher δ (or a lower β) implies stronger substitutability and thus tougher price competition. Finally, define the *quality-adjusted wage*

$$\tilde{w}_i^p = \frac{w_i^p}{k_i^{1/\delta}}, \quad i \in \{H, W\},$$

which rescales wages by the country's efficiency level k_i. In effect, wages in more efficient environments are "discounted," reflecting that higher quality lowers the effective cost of supplying a variety.

Consider now a Home upstream supplier with productivity ϕ_H competing with a World supplier with productivity ϕ_W in the same variety. The Home supplier's profit is

$$\pi_H(\phi_H, \phi_W) = \kappa\,\phi_H^{\delta}\,(\tilde{w}_H^p)^{-\delta}\,\mathbf{1}\left\{\frac{\phi_H}{\tilde{w}_H^p} \geq \frac{\phi_W}{(1+\tau_H)\tilde{w}_W^p}\right\}. \tag{5}$$

Profits are positive only if the Home supplier's productivity, scaled by its wage and quality, exceeds its World rival's once the trade cost is applied to the rival's delivered cost.

[14] Similarly to Akcigit et al. (2025), I model competition as a two-stage game where the higher-cost supplier exits. An alternative approach would be to let the more productive supplier set a limit price equal to its rival's marginal cost, as De Ridder (2024) does. Both formulations yield the same main implication of the model: an inverted-U-shaped relationship between foreign competition and innovation.

Formally, the condition is

$$\frac{\phi_H}{\tilde{w}_H^p} \quad \geq \quad \frac{\phi_W}{(1+\tau_H)\tilde{w}_W^p}.$$

When this inequality holds, the Home supplier captures the entire demand for that variety. Its profits are then proportional to ϕ_H^δ, decreasing in $\tilde{w}_H^p$, and scaled by κ. The exponent $\delta = 1/\beta - 1$ governs the elasticity of substitution across varieties. Although competition within a variety is perfectly competitive, with only the lowest-cost supplier serving the market, suppliers still face monopolistic competition across varieties through the CES aggregator. A higher δ means that varieties are closer substitutes, so a productivity advantage ϕ_H translates more strongly into demand shifting toward that supplier, amplifying profits for the more productive firm. If instead the World supplier has a lower quality- and tariff-adjusted cost (whether because of higher productivity ϕ_W or a lower $\tilde{w}_W^p$), the Home supplier exits and earns zero.

2.1.4 Innovation by Home Upstream Suppliers

In the previous subsection, I described profits and sourcing decisions conditional on suppliers' *post-innovation* productivities (ϕ_H, ϕ_W). I now explain how these productivities are determined.

Reflecting the small open economy setting, Home suppliers endogenously choose whether to invest in innovation in response to tariffs and foreign competition, while World suppliers are treated as exogenous and non-responsive to Home conditions. In particular, World suppliers are assumed to operate with a fixed post-innovation productivity distribution, which already reflects the outcomes of their own innovation processes. That is, their productivities ϕ_W are drawn directly from a given distribution.

In contrast, Home upstream suppliers endogenously invest in cost-reducing innovation. After drawing a baseline productivity $\tilde{\phi}_H$, a supplier chooses an innovation effort $u \in [0,1]$, which determines the probability of upgrading to $(1+a)\tilde{\phi}_H$, where $a > 0$ is a fixed step size. With probability $1-u$, the supplier retains its baseline productivity. This mechanism captures a discrete improvement in technology driven by endogenous effort.

Assumption 1 (Productivity distributions)**.** *Baseline productivities in Home are drawn from a Fréchet distribution:*

$$F_H(x) \equiv \Pr(\tilde{\phi}_H \leq x) = \exp(-T_H\, x^{-\theta}), \qquad x > 0.$$

World productivities are drawn directly from a post-innovation Fréchet distribution:

$$G_W(x) \equiv \Pr(\phi_W \le x) = \exp(-T_W\, x^{-\theta}), \qquad x > 0.$$

The scale parameters T_H and T_W govern the average productivity in each country: a higher T implies a rightward shift of the distribution and, hence, a more competitive pool of suppliers. The shape parameter $\theta > 1$ controls the thickness of the right tail and determines the degree of heterogeneity across suppliers, with lower values of θ implying greater dispersion. By assuming that both countries share the same shape parameter θ, I impose that the extent of supplier-level heterogeneity is identical across countries, while allowing average productivity (via T_i) to differ between Home and World.

Before describing how u is determined, it is useful to lay out the timing of events:

1. Each upstream supplier in Home draws a baseline productivity $\tilde{\phi}_H$ from the distribution $F_H(\cdot)$ (Assumption 1).

2. Home suppliers decide on an innovation probability $u \in [0, 1]$, trading off the cost against expected profits. Importantly, at this stage a Home supplier does not observe the productivity of its specific World rival in its variety. Instead, it only knows the distribution of World productivities, denoted $G_W(\cdot)$, and therefore bases its choice on expected profits.[15]

3. Innovation outcomes are realized: with probability u, the Home supplier's productivity upgrades to $\phi_H = (1 + a)\tilde{\phi}_H$; with probability $1 - u$, it remains at $\phi_H = \tilde{\phi}_H$.

4. Each variety is then populated by one Home supplier and one World competitor. At this stage, suppliers observe each other's realized productivities (ϕ_H, ϕ_W) and engage in the two-stage Bertrand pricing game described in Section 2.1.3.

This setup reflects the fact that Home suppliers make innovation choices under uncertainty, responding to the distribution of foreign competition rather than to the productivity of a specific rival. Now that the timing of events is established, we can describe the Home supplier's profit maximization problem, which determines the optimal innovation probability u.

[14]One can interpret T_W as the outcome of an innovation process by World suppliers, starting from some baseline $\tilde{T}_W$ and upgrading to $T_W = \tilde{T}_W(1 + a_W)^{-\theta}$. Because this is a small open economy model, I treat T_W as fixed and exogenous from the perspective of Home. This abstraction is without loss of generality, since World suppliers are not assumed to respond to conditions in Home.

[15]Qualitatively identical results could be obtained if suppliers are matched with their specific World competitor before choosing u. I adopt the ex ante formulation to better align the model's innovation choice with the empirical setting, where suppliers respond to the distribution of foreign competition rather than to individual rivals.

Home supplier's maximization problem. Given its baseline productivity $\tilde{\phi}_H$, the Home supplier chooses an innovation probability $u \in [0,1]$—the likelihood of successfully upgrading from $\tilde{\phi}_H$ to $(1+a)\tilde{\phi}_H$—to maximize expected profits net of the innovation cost:

$$\begin{aligned}\max_{u\in[0,1]} \;& (1-u)\; \underbrace{\eta\,\kappa\,\tilde{\phi}_H^{\delta}\,(\tilde{w}_H^{p})^{-\delta}}_{\text{operating profit at } \tilde{\phi}_H}\; \underbrace{G_W\big((1+\tau_H)\tilde{o}_W\,\tilde{\phi}_H\big)}_{\text{win prob. at } \tilde{\phi}_H} \\ & + \; u\; \underbrace{\eta\,\kappa\,\big((1+a)\tilde{\phi}_H\big)^{\delta}\,(\tilde{w}_H^{p})^{-\delta}}_{\text{operating profit at } (1+a)\tilde{\phi}_H}\; \underbrace{G_W\big((1+\tau_H)\tilde{o}_W\,(1+a)\tilde{\phi}_H\big)}_{\text{win prob. at } (1+a)\tilde{\phi}_H} \\ & - \; \underbrace{R(u,\tilde{\phi}_H)}_{\text{cost of } u}.\end{aligned} \tag{6}$$

The maximization problem in (6) has three elements. The first term reflects the operating profit the Home supplier would earn if it serves the market with its baseline productivity ($\eta\;\kappa\;\tilde{\phi}_H^{\delta}\;(\tilde{w}_H^{p})^{-\delta}$), weighted by the probability of winning against a randomly drawn World rival, $G_W\left((1+\tau_H)\,\tilde{o}_W\tilde{\phi}_H\right)$. The second term captures the case in which innovation succeeds: productivity rises to $(1+a)\tilde{\phi}_H$, operating profits increase accordingly, and the probability of winning is recalculated under the World distribution $G_W(\cdot)$. Finally, the last term subtracts the cost of innovation, $R(u,\tilde{\phi}_H)$. Because the Home supplier does not observe the realized productivity of its specific World competitor, it cannot condition directly on ϕ_W. Instead, it evaluates expected profits using the distribution $G_W(\cdot)$: in equilibrium, the supplier either earns the monopoly profit if it undercuts its rival, or zero if it is undercut.

Innovation cost. Innovation is implemented by hiring domestic *technically skilled workers* paid wage w_R. The cost of raising u is assumed to be convex and increasing in the supplier's baseline productivity $\tilde{\phi}_H$: when productivity is already high, further improvements require disproportionately more resources. The parameter $f>0$ captures the baseline difficulty or fixed intensity of innovation, so that the cost function is

$$R(u,\tilde{\phi}_H) \;=\; \tfrac{f}{2}\,u^2\,\tilde{\phi}_H^{\delta}\,w_R.$$

The quadratic term in u captures the convexity of the cost of raising the innovation probability. More generally, any specification of the form u^{γ} with $\gamma>1$ would capture convex costs. Following the literature on innovation (see Acemoglu et al. (2018) for discussion), I set $\gamma=2$. This choice is standard and does not affect the model's qualitative results, particularly the inverted-U shape of innovation incentives.[16]

Optimal innovation. From (6), the FOC equates marginal benefits with marginal costs of innovation. Innovation generates gains along two distinct dimensions. First, keeping constant the probability of capturing the market, upgrading by a factor $(1+a)$ raises its profits proportionally—this is the *scale effect*. Second, innovation also increases the probability of overtaking the World rival and capturing the market—this is the *competition effect*.

Formally, the expected profit gain net of the supplier-level cost multiplier can be written as:

$$\begin{aligned} S(\tilde{\phi}_H) =& \underbrace{G_W\big((1+\tau_H)\tilde{o}_W\tilde{\phi}_H\big)\left[(1+a)^\delta - 1\right]}_{\text{Scale Effect}} \\ & + \underbrace{\Big[G_W\big((1+\tau_H)\tilde{o}_W(1+a)\tilde{\phi}_H\big) - G_W\big((1+\tau_H)\tilde{o}_W\tilde{\phi}_H\big)\Big](1+a)^\delta}_{\text{Competition Effect}} \\ =& (1+a)^\delta\, G_W\big((1+\tau_H)\tilde{o}_W\,(1+a)\tilde{\phi}_H\big) \;-\; G_W\big((1+\tau_H)\tilde{o}_W\,\tilde{\phi}_H\big) \end{aligned}$$

where the first term reflects the *scale effect* (higher profits conditional on winning) and the second term reflects the *competition effect* (a higher chance of winning the market).

Given this definition, the supplier's optimal innovation intensity is

$$u^*(\tilde{\phi}_H) \;=\; \min\left\{1,\; \frac{\kappa\,(\tilde{w}_H^p)^{-\delta}}{f\,w_R}\; S(\tilde{\phi}_H)\right\}. \tag{7}$$

Here κ and f are parameters, $\tilde{w}_H^p$ is taken as exogenous, and w_R is determined in general equilibrium. As such, when analyzing how innovation probability u^* varies with baseline productivity $\tilde{\phi}_H$, the relevant object is $S(\tilde{\phi}_H)$: the constants simply scale the level of u^* but do not affect its shape.

Intuitively, suppliers innovate more when the net expected payoff $S(\tilde{\phi}_H)$ is high. The central question, however, is which suppliers innovate most intensively—that is, at what

[16] The exponent on productivity, $\tilde{\phi}_H^\delta$, links the cost of innovation to the supplier's baseline efficiency. Intuitively, more productive suppliers face steeper challenges in pushing the frontier further. I assume here that the elasticity of cost with respect to productivity equals δ, the same parameter that governs the demand-side elasticity of substitution across varieties. More generally, the cost could be specified as $\tilde{\phi}_H^\eta$ for any $\eta \geq \delta$. If $\eta < \delta$, the profit gains from higher productivity would increasingly outweigh the rise in innovation costs, implying that more productive suppliers should innovate disproportionately more—a pattern not supported by the data, as shown in Section 3.3.1. By contrast, when $\eta \geq \delta$, the model preserves its qualitative implications. I adopt the normalization $\eta = \delta$ for tractability, following Akcigit et al. (2025), who make a similar assumption equating the elasticity of innovation costs with the elasticity of quality improvements.

point in the productivity distribution the innovation policy $u^*(\tilde{\phi}_H)$ reaches its maximum. Proposition 1 characterizes this function and establishes the inverted-U result.

Proposition 1 (Limits, peak, and comparative statics of innovation incentives)**.** *Fix $a > 0$, $\tau_H \geq 0$, $T_W > 0$, and $0 < \delta < \theta$. The innovation probability $u^*(\tilde{\phi}_H)$ satisfies*

$$\lim_{\tilde{\phi}_H \to 0} u^*(\tilde{\phi}_H) = 0, \qquad \lim_{\tilde{\phi}_H \to \infty} u^*(\tilde{\phi}_H) = \frac{\kappa\,(\tilde{w}_H^p)^{-\delta}}{f\,w_R}\left[(1+a)^{\delta} - 1\right].$$

It attains a unique interior maximum at

$$\tilde{\phi}_H^* = \frac{1}{(1+\tau_H)\tilde{o}_W}\left(\frac{T_W\left[1-(1+a)^{-\theta}\right]}{(\theta-\delta)\,\ln(1+a)}\right)^{1/\theta},$$

with $u''(\tilde{\phi}_H^*) < 0$, so that $u^*(\tilde{\phi}_H)$ is single-peaked (inverted-U shaped).[17]

The location of the peak satisfies the following comparative statics:

1. $\frac{\partial \tilde{\phi}_H^*}{\partial T_W} > 0$*: stronger World technology shifts the peak innovation to higher Home productivity levels.*
2. $\frac{\partial \tilde{\phi}_H^*}{\partial \tau_H} < 0$*: higher tariffs reduce the effective competitiveness of World rivals, shifting the peak innovation threshold toward lower-productivity Home suppliers.*
3. $\frac{\partial \tilde{\phi}_H^*}{\partial a} < 0$*: larger innovation steps make it worthwhile for laggard suppliers to invest in upgrading.*
4. $\frac{\partial \tilde{\phi}_H^*}{\partial \delta} > 0$*: when varieties are closer substitutes, upgrading is especially valuable for already-productive suppliers.*
5. $\lim_{a\to 0}\lim_{\theta\to\infty}\tilde{\phi}_H^* = \frac{1}{(1+\tau_H)\tilde{o}_W}\left(\frac{\theta}{\theta+1}T_W\right)^{1/\theta} = \frac{1}{(1+\tau_H)\tilde{o}_W}\phi_{W,\textit{mode}}$*: The peak aligns with the tariff-adjusted modal World productivity in the limit of vanishing World productivity dispersion ($\theta \to \infty$) and small innovation steps ($a \to 0^+$).*

Proof. See Appendix A.2. □

Proposition 1 is a key theoretical result. It formalizes why innovation is highest for suppliers in the middle of the productivity distribution—those for whom an upgrade substantially raises the probability of winning against World competitors—and lower for suppliers that are far behind or far ahead.

[17] Formally, $u^*(\tilde{\phi}_H) = \min\{u(\tilde{\phi}_H), 1\}$. The expression above assumes parameters such that $u(\tilde{\phi}_H) < 1$ for all $\tilde{\phi}_H$, which holds if f is sufficiently large to make $\frac{\kappa(\tilde{w}_H^p)^{-\delta}}{f\,w_R}\left[(1+a)^{\delta}-1\right] < 1$. If the cap binds, the profile remains inverted-U but is truncated at 1 around $\tilde{\phi}_H^*$.

The expected payoff from innovation is captured by $S(\tilde{\phi}_H)$. Since equilibrium innovation $u^*(\tilde{\phi}_H)$ is a scaled version of $S(\tilde{\phi}_H)$, its profile inherits the same inverted-U shape. Figure 1 therefore plots $S(\tilde{\phi}_H)$ to illustrate the hump-shaped incentives that characterize $u^*(\tilde{\phi}_H)$. The shape of $S(\tilde{\phi}_H)$ itself is governed by two forces: the *competition effect*, whereby an upgrade raises the probability of overtaking rivals, and the *scale effect*, the profit gain from higher productivity when the supplier already captures the market.[18]

Very low-productivity suppliers innovate little. They are too far behind their competitors: even after a successful step to $(1+a)\tilde{\phi}_H$, the *competition effect* is small because the probability of overtaking a competitor remains negligible. The *scale effect* is also limited, since such suppliers seldom win the market to begin with. With little upside and convex innovation costs, optimal effort is low.

Very high-productivity suppliers also innovate little. They already win with probability close to one, so the *competition effect* is negligible: there is almost no scope to raise the win probability further. In this limit, the innovation payoff converges to the *scale effect* alone, $(1+a)^\delta - 1$, which is simply the proportional profit gain from raising productivity when the supplier already captures the market.

By contrast, mid-productivity suppliers innovate the most. These suppliers are in the range where an upgrade meaningfully increases the likelihood of winning the market in expectation. The *competition effect* is strongest here: a small productivity jump flips a significant share of head-to-head contests against potential competitors. Combined with a nontrivial *scale effect*, this produces the highest payoff-to-cost ratio and thus the peak of $u^*(\tilde{\phi}_H)$. Figure 1 illustrates this inverted-U shape.

[18] I plot $S(\tilde{\phi}_H)$ rather than $u^*(\tilde{\phi}_H)$ because the latter depends on general-equilibrium objects such as w_R, which will be solved in the next subsection.

Figure 1: Expected innovation gain $S(\phi_H)$.

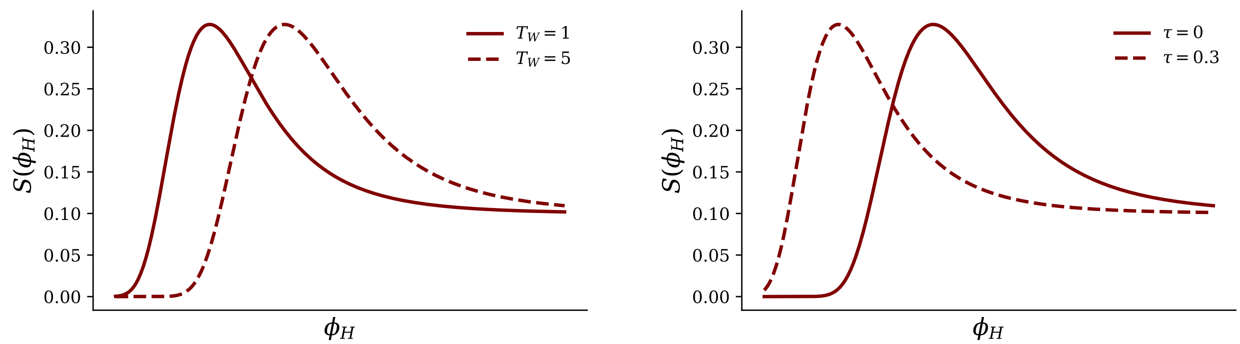

(a) Different T_W values ($\tau_H = 0$ fixed).

(b) Different τ_H values ($T_W = 5$ fixed).

Notes: These figures plot the expected gain from innovation, $S(\phi_H)$, as a function of a Home supplier's productivity, ϕ_H. $S(\phi_H)$ represents the supplier's expected increase in profit from upgrading its technology, given the distribution of foreign competitors' productivity. The relationship follows an inverted-U shape: innovation incentives are low for very unproductive suppliers (upgrading does little to close the gap) and for very productive suppliers (already dominant), and peak for intermediate suppliers that face the highest expected gain from overtaking competitors. Panel (a) varies the scale parameter of the foreign productivity distribution, T_W, while holding the domestic trade cost, τ_H, fixed. Panel (b) varies τ_H while holding T_W fixed. Parameters: $\beta = 0.5$, $\delta = 1$, $\theta = 8$, $a = 0.1$.

The position of the peak, $\tilde{\phi}_H^*$, shifts systematically with model parameters. When World technology T_W rises, the distribution of competitor productivities shifts rightward, so Home suppliers must reach a higher baseline productivity before innovation can meaningfully raise their probability of winning; the peak therefore moves right, as shown in panel (a) of Figure 1. By contrast, higher tariffs τ_H inflate the delivered costs of World competitors, shifting their effective distribution leftward in tariff-adjusted terms and making it easier for less productive Home suppliers to catch up; in this case, the peak moves left, as shown in panel (b) of Figure 1. Larger innovation steps a also shift the peak left, since even laggard suppliers can cover more of the gap in a single jump. Greater substitutability δ has the opposite effect: when varieties are closer substitutes, demand is highly sensitive to relative costs (across varieties), so upgrading is especially valuable for suppliers already near the frontier, pushing the peak toward higher productivity. Finally, when the shape parameter θ is large and when the jump size a is small, the distribution of World productivities becomes tightly concentrated, and $\tilde{\phi}_H^*$ aligns closely with the tariff-adjusted modal World productivity. In other words, the point at which domestic suppliers innovate the most corresponds to the productivity level where they face the highest probability of being close enough to competitors for innovation to meaningfully affect their chances of winning the market.

From theory to observables. Neither the innovation probability rate u^* nor baseline productivities $\tilde{\phi}_H$ are observed directly in the data, so Proposition 1 cannot be tested in its raw form. I therefore map the theory into two observables:

(i) Share of wage bill on technical occupations. The first observable is the share of a supplier's wage bill allocated to R&D workers (scientists, engineers, and related technical staff).

$$s^{\text{tech}}(\tilde{\phi}_H) \;\equiv\; \frac{w_R \ell_R(\tilde{\phi}_H)}{w_H \ell_P(\tilde{\phi}_H)}.$$

Under the quadratic innovation cost, the numerator $w_R \ell_R$ equals the cost of innovation effort $R(u, \tilde{\phi}_H)$, so $s^{\text{tech}}(\tilde{\phi}_H)$ is (up to a positive multiplicative constant) monotone in $u(\tilde{\phi}_H)^2$ and inherits the inverted-U prediction.[19]

(ii) Quality-adjusted price gap. The second observable measures the distance between a Home supplier's quality-adjusted price and that of a representative World competitor. A convenient benchmark is the *modal* World price, since the Fréchet distribution delivers a closed-form expression for the modal productivity (and hence modal quality-adjusted price) of World varieties. This makes the mode both analytically tractable in theory and empirically meaningful: in the data, where many products coexist with heterogeneous unit values, the mode serves as a natural reference point for the "typical" World competitor faced by a Home supplier.

Formally, define the quality-adjusted price of a Home supplier with baseline productivity $\tilde{\phi}_H$ as

$$\widehat{p}_H(\tilde{\phi}_H) \;\equiv\; k_H^{\frac{1}{1-\beta}}\, p_H(\tilde{\phi}_H).$$

Likewise, define the quality-adjusted price of the *modal* (i.e. typical) World competitor as

$$\widehat{p}_W^{\text{mode}} \;\equiv\; (1+\tau_H)\, k_W^{\frac{1}{1-\beta}}\, p_W^{\text{mode}}.$$

The relative gap in quality-adjusted prices can then be expressed as the ratio

$$D \;\equiv\; \frac{\widehat{p}_H}{\widehat{p}_W^{\text{mode}}}.$$

Because $\widehat{p}_H \propto \tilde{\phi}_H^{-1}$ and $\widehat{p}_W^{\text{mode}} \propto (\tilde{\phi}_W^{\text{mode}})^{-1}$, the gap D provides a one-to-one proxy for productivity differences between Home and World. The Fréchet mapping between $\tilde{\phi}_H$

[19]The constant of proportionality collects $\frac{f}{2}\tilde{\phi}_H^{\delta}$ and the denominator $w_H \ell_P(\tilde{\phi}_H)$. Its explicit form is given in Appendix A.3.

and D is derived in Appendix A.3.

Proposition 2 (Peak and comparative statics in the share of wage bill on technical occupations)**.** *Fix $a > 0$, $\tau_H \geq 0$, $T_W > 0$, and $0 < \delta < \theta$.*
The share of wage bill on technical occupations, s^{tech}, reaches a unique interior maximum in the quality-adjusted price gap D between a Home supplier and the modal World competitor.

The maximizer is given by

$$D^* = \left[\frac{\theta(\theta - \delta) \ln(1 + a)}{(\theta - 1)\left(1 - (1 + a)^{-\theta}\right)} \right]^{1/\theta},$$

for $\theta > 1$ and $0 < \delta < \theta$, with $s^{\text{tech}\,\prime\prime}(D^*) < 0$.

Moreover, the location of the peak satisfies similar comparative statics as in Proposition 1:

1. $\frac{\partial D^*}{\partial a} > 0$*: larger innovation steps make it worthwhile for laggard suppliers to invest in upgrading, shifting the peak rightward.*

2. $\frac{\partial D^*}{\partial \delta} < 0$*: when varieties are closer substitutes, the peak occurs at lower price gaps.*

3. $\lim_{a \to 0^+} \lim_{\theta \to \infty} D^* = 1$*: the peak aligns with the tariff-adjusted modal World price as $\theta \to \infty$ and $a \to 0^+$.*

Note: *The expression for D^* follows directly from the mapping between innovation $u^*(\tilde{\phi}_H)$ and the observable $s^{\text{tech}}(\tilde{\phi}_H)$ derived in Appendix A.3.*

To interpret this result, note that the price gap D moves inversely with Home productivity: as a supplier's productivity $\tilde{\phi}_H$ falls, D rises. By construction, $D = 1$ when the Home supplier's quality-adjusted price equals the delivered (tariff-adjusted) quality-adjusted price of the modal World competitor. Hence $D > 1$ indicates that the Home supplier is relatively expensive (less productive) than the typical World competitor, while $D < 1$ indicates the opposite.

The inverted-U in $s^{\text{tech}}(D)$ then follows directly from Proposition 1. Suppliers with very high quality-adjusted prices (large D, low productivity) innovate little: even after upgrading, the probability of catching up remains negligible. Symmetrically, suppliers with very low quality-adjusted prices (small D, high productivity) also innovate little: they already dominate, so further upgrades yield only limited additional gain. It is the intermediate suppliers—those close to parity with the typical World competitor—that devote the largest share of their wage bill to technical occupations, as innovation most increases their probability of capturing the market.

As in Proposition 1, the position of the peak shifts systematically with model parameters. Larger innovation steps a shift the peak D^* to the right, since higher-priced (large D, low productivity) suppliers can cover more of the gap in a single jump. Greater substitutability δ has the opposite effect: when varieties are closer substitutes, demand is highly sensitive to relative costs, so upgrading is especially valuable for suppliers that already price low (small D, high productivity), pushing the peak toward lower D.

Finally, when the innovation step is low ($a \to 0$) and the dispersion of World productivities vanishes ($\theta \to \infty$), the peak occurs precisely at parity ($D^* = 1$). In this limit, the suppliers that are neck-and-neck with the typical World competitor and have a chance to meaningfully win the market are those that innovate the most.

2.2 Equilibrium

In equilibrium, the World production wage w_W^p is solved separately from World labor–market clearing, and the Home production wage w_H^p is tied to it through the homogeneous–good arbitrage condition. Therefore, the remaining endogenous variables to be determined are the Home wage on technical occupations w_R and the Home fixed–factor rent f_H.

Production wages. Production labor in World is employed in both the homogeneous good and the differentiated sector. With Cobb–Douglas preferences, demand for the homogeneous good pins down labor use in that sector, while CES demand and markup pricing yield closed forms for labor demand in intermediates. Combining the labor market clearing condition with the World income identity gives the production wage:

$$w_W^p \;=\; \left\{ M_W \,(1-\beta)^{\frac{1}{\beta}-2} \left((1-\beta)^2 + \frac{\alpha}{1-\alpha} \right) \right\}^{\beta}, \qquad M_W := T_W^{\delta/\theta}\, \Gamma\!\Big(1 - \tfrac{\delta}{\theta}\Big). \tag{8}$$

Since the homogeneous good is freely traded, Home and World wages are tied by relative productivity:

$$w_H^p = \frac{1}{o_W}\, w_W^p. \tag{9}$$

Appendix A.4 provides the full derivation.

Wage of technical occupations. The second equilibrium condition ensures that labor demand for technical occupations equals its supply.

For a Home supplier with baseline productivity $\tilde{\phi}_H$, the equilibrium innovation rate is $u^*(\tilde{\phi}_H)$, as defined in Equation (7). The innovation cost function implies that the labor required to support this innovation rate is

$$\frac{f}{2}\, u^*(\tilde{\phi}_H)^2 \tilde{\phi}_H^{\delta}.$$

Total tech labor demand in Home is then

$$L_H^R = \int_0^{\infty} \frac{f}{2}\, u^*(\tilde{\phi}_H)^2 \tilde{\phi}_H^{\delta}\; dF_H(\tilde{\phi}_H), \tag{10}$$

which must equal the labor endowment ℓ_H^R, where $F_H(\cdot)$ is the distribution of pre–innovation productivities.

Solving this condition yields a closed form for the wage of technical occupations:

$$w_R \;=\; \frac{\kappa\,(\tilde{w}_H^p)^{-\delta}}{\sqrt{2f\ell_H^R}}\,\left(\int_0^{\infty} \tilde{\phi}_H^{\delta} S(\tilde{\phi}_H)^2 f_H(\tilde{\phi}_H)\, d\tilde{\phi}_H\right)^{1/2}. \tag{11}$$

Equation (28) highlights the determinants of the researcher wage w_R. The parameter κ captures overall profitability in the differentiated sector: higher profitability increases suppliers' willingness to pay for innovation, raising the demand for researchers and thus their wage. The factor $(\tilde{w}_H^p)^{-\delta}$ reflects the cost of production labor in Home: as production wages rise, profitability falls, discouraging innovation and depressing w_R. The denominator $\sqrt{2f\ell_H^R}$ shows that a higher fixed cost of innovation f lowers demand for research labor, while a larger supply of research workers ℓ_H^R puts downward pressure on their wage. Finally, the bracketed integral aggregates the effective gains from innovation across the productivity distribution. Here $\tilde{\phi}_H^{\delta}$ captures that innovation requires more efficiency units of labor at higher productivity levels, while $S(\tilde{\phi}_H)^2$ reflects the intensity of firms' innovation incentives at each point in the distribution. Together, these terms summarize how the wage for technical occupations adjusts to equate the supply and demand for research labor.

Fixed factor rents. As shown in the first–order condition (2), fixed–factor rents are proportional to output: $f_H = \beta Y_H$, ensuring that the fixed factor receives its equilibrium share of income.

2.3 Aggregate and Welfare Implications

2.3.1 Post-Innovation Productivities

I begin by characterizing how innovation reshapes the distribution of Home suppliers' productivities. A supplier with initial productivity $\tilde{\phi}_H$ ends up at one of two possible outcomes: it either remains at $\tilde{\phi}_H$ if it does not innovate, or jumps to $(1+a)\tilde{\phi}_H$ if it successfully innovates. The probability of upgrading is given by $u^*(\tilde{\phi}_H)$ from Equation (7). Hence the expected productivity of such a supplier is

$$\mathbb{E}[\phi_H^{\text{new}} \mid \tilde{\phi}_H] = (1 - u^*(\tilde{\phi}_H)) \cdot \tilde{\phi}_H + u^*(\tilde{\phi}_H) \cdot (1+a)\tilde{\phi}_H.$$

Aggregating across all suppliers delivers the post–innovation density of Home productivities:

$$g_H(\phi_H) = \underbrace{(1 - u^*(\phi_H))\, f_H(\phi_H)}_{\text{suppliers that do not upgrade}} + \underbrace{\tfrac{1}{1+a}\, u^*\big(\tfrac{\phi_H}{1+a}\big)\, f_H\big(\tfrac{\phi_H}{1+a}\big)}_{\text{suppliers that upgrade}}, \tag{12}$$

where $f_H(\cdot)$ is the initial Fréchet density of Home productivities. The first term represents suppliers that stay at their baseline ϕ_H, while the second term reflects those that successfully move up from $\phi_H/(1+a)$ to ϕ_H.[20]

2.3.2 Aggregate Productivity Gain from Innovation

To evaluate the aggregate implications, I follow Melitz (2003) and define aggregate Home productivity of *active* suppliers as:

$$\mathbf{\Phi}_H \;=\; \int_0^\infty \phi_H^\delta \, G_W\big((1+\tau_H)\tilde{o}_W\, \phi_H\big)\, \upsilon_H(\phi_H)\, d\phi_H,$$

where $G_W(\cdot)$ is the probability that a Home supplier with productivity ϕ_H is active in the domestic market (i.e. it outcompetes its World rival), and $\upsilon_H(\cdot)$ is the density of Home productivities. Here $\upsilon_H = f_H$ in the pre–innovation case and $\upsilon_H = g_H$ in the post–innovation case.

The *aggregate productivity gain from innovation* is then defined as the difference between post– and pre–innovation aggregate productivity:

$$\mathcal{G}(T_H, T_W, \tau_H) = \mathbf{\Phi}_H^{post} - \mathbf{\Phi}_H^{pre}.$$

In other words, $\mathcal{G}$ measures how much innovation shifts the productivity distribution of

[20]See Appendix A.2 for the general equilibrium derivation of $u^*(\phi_H)$.

active Home suppliers upward, relative to the counterfactual in which all suppliers were frozen at their baseline productivities.

Proposition 3 (Aggregate productivity gain)**.** *The aggregate productivity gain from innovation yields the closed form*

$$\mathcal{G}(T_H, T_W, \tau_H) = \left(\frac{2\ell_H^R}{f}\right)^{1/2} \left\{\Gamma\, T_H \left[\frac{(1+a)^{2\delta}}{(T_H+2b)^{1-s}} - \frac{2(1+a)^{\delta}}{(T_H+b+c)^{1-s}} + \frac{1}{(T_H+2c)^{1-s}}\right]\right\}^{1/2},$$

with $s = \delta/\theta$, $\widetilde{T}_W = T_W(1+\tau_H)^{-\theta}\tilde{o}_W^{-\theta}$, $b = \widetilde{T}_W(1+a)^{-\theta}$, *and* $c = \widetilde{T}_W$.

For any $T_H > 0$, $T_W > 0$, $\tau_H \geq 0$, $a > 0$, *and* $\delta \in (0, \theta)$*:*

1. *The aggregate productivity gain is strictly positive,*

$$\mathcal{G}(T_H, T_W, \tau_H) > 0.$$

2. *(Hump shape) Under* $\delta \in (0,\theta)$ *and* $a \in (0,1)$,

$$\left.\frac{\partial \mathcal{G}}{\partial \widetilde{T}_W}\right|_{\widetilde{T}_W = 0^+} > 0, \qquad \exists\, M < \infty : \frac{\partial \mathcal{G}}{\partial \widetilde{T}_W} < 0 \text{ for all } \widetilde{T}_W \geq M, \quad \text{and} \quad \lim_{\widetilde{T}_W \to \infty} \frac{\partial \mathcal{G}}{\partial \widetilde{T}_W} = 0.$$

By continuity, there exists at least one $\widetilde{T}_W^* \in (0, M)$ *with* $\frac{\partial \mathcal{G}}{\partial \widetilde{T}_W}(\widetilde{T}_W^*) = 0$, *so* $\mathcal{G}$ *is hump-shaped in* $\widetilde{T}_W$.

Hence $\mathcal{G}$ *increases for sufficiently small* $\widetilde{T}_W$ *and decreases for sufficiently large* $\widetilde{T}_W$.

Proof. See Appendix A.5. □

Proposition 3 delivers two messages. After providing a closed-form expression, the proposition shows that (i) the aggregate productivity gain from innovation $\mathcal{G}$ is *always positive* for all admissible parameter values; and (ii) it is *hump-shaped* in foreign competition.

To organize the intuition, define the *effective World productivity*

$$\widetilde{T}_W \equiv \frac{T_W}{\big((1+\tau_H)\tilde{o}_W\big)^{\theta}},$$

which summarizes how World aggregate productivity (T_W) is perceived in the Home market after adjusting for trade frictions (the tariff τ_H) and relative unit–labor productivity ($\tilde{o}_W$). Differentiating the closed form shows that the aggregate productivity gain $\mathcal{G}$ *rises* *when effective World productivity $\widetilde{T}_W$ is low* but *falls when $\widetilde{T}_W$ is high*; thus $\mathcal{G}$ is hump–shaped in $\widetilde{T}_W$.

The intuition is simple. When effective World productivity is very low, Home suppliers face little competitive pressure; the expected payoff from upgrading is small, so aggregate innovation is limited. When effective World productivity is very high, most Home suppliers are far behind; even a successful upgrade barely changes their chance of winning, so innovation again remains low. In the limit $\widetilde{T}_W \to \infty$, innovation vanishes because Home suppliers never expect to outcompete foreign rivals. Innovation is therefore most intense—and $\mathcal{G}$ is maximized—when Home and World are *neck–and–neck*, i.e., when effective productivities are close.

Figure 2 visualizes these effects. Panel (a) plots the aggregate productivity gain against World aggregate productivity, holding tariffs constant at zero, and shows three curves corresponding to low, medium, and high levels of Home productivity (0.5, 1, and 2). Panel (b) plots the same gain against the tariff rate, holding Home productivity fixed at 1, and varies the World productivity level (0.5, 1, and 2). Since effective World competitiveness increases with higher World productivity and decreases with higher tariffs, both panels capture the same underlying hump-shaped relationship—just traced out from different directions.

Figure 2: Aggregate productivity gain

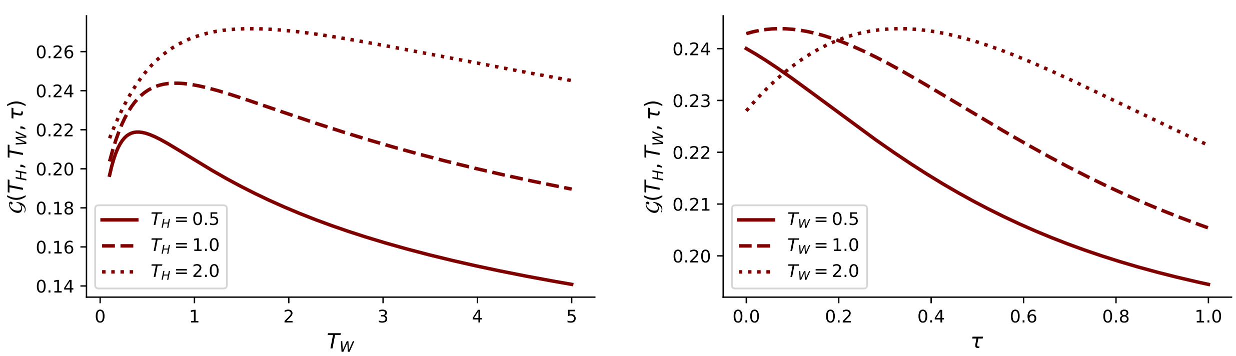

(a) Varying T_W at $\tau_H = 0$ for $T_H \in \{0.5, 1, 2\}$.

(b) Varying τ_H at $T_H = 1$ for $T_W \in \{0.5, 1, 2\}$.

Notes: This figure plots the aggregate productivity gain from innovation, $\mathcal{G}$, computed from the model under parameter values $\beta = 0.5$, $\delta = 1$, $\theta = 3$, and $a = 0.1$. Panel (a) varies the scale parameter of the foreign productivity distribution, T_W, while holding the domestic trade cost, τ_H, fixed. Panel (b) varies τ_H while holding T_W fixed. The figure shows that the aggregate productivity gain from innovation is inverted-U shaped when the Home economy is technologically behind (lower T_H relative to T_W), but becomes decreasing when it is technologically ahead. β governs demand curvature (with price elasticity equal to $1/\beta$) and thus affects markups; $\delta \equiv 1/\beta - 1$ is the return exponent in innovation payoffs; θ is the Fréchet tail parameter shaping productivity dispersion; and a is the innovation step size, such that a successful innovator's productivity scales by $(1 + a)$.

Panel (a) shows how the aggregate gain from innovation is affected by changes in the level of World productivity. The plot shows a clear "neck-and-neck" pattern across different levels of Home productivity. The aggregate productivity gain from innovation is largest when domestic and foreign technologies are similar. In contrast, when the World is far ahead, most domestic suppliers are too far behind to catch up and have little incentive to innovate. When the Home economy is far ahead, most suppliers already dominate and also face little pressure to improve. In both cases, fewer suppliers face meaningful competition, so innovation and aggregate productivity gains are limited.

Panel (b) shows how the aggregate gain from innovation is affected by tariffs. When the World is behind, increasing tariffs simply pushes Home and World suppliers further apart and lowers innovation. When the World and Home are close, raising tariffs breaks this balance and again reduces incentives to innovate. But when the World is far ahead, a moderate tariff actually brings more suppliers into neck-and-neck competition by narrowing the gap, which increases innovation. However, this effect eventually reverses—too much protection again fragments competition, reducing the gains from innovation.

Overall, both panels highlight that relative productivities and trade policy shape the distribution of neck-and-neck competitors, which in turn governs innovation incentives and the aggregate productivity gain.

Welfare decomposition. Household utility follows Cobb–Douglas preferences over a homogeneous good z_H and a differentiated composite D_H, yielding

$$U_H \;=\; \alpha^{\alpha}(1-\alpha)^{1-\alpha}\; p_z^{-\alpha}\; I_H, \tag{13}$$

where α is the share of income spent on the homogeneous good, p_z is the price of the homogeneous good, and I_H is income.[21]

Household income is

$$I_H \;=\; w_H^p\,\ell_H^P \;+\; w_H^R\,\ell_H^R \;+\; f_H\,F_H \;+\; \Pi_H, \tag{14}$$

where ℓ_H^P and ℓ_H^R denote the labor endowment in production and technical occupations. Income has four sources: (i) the production wage w_H^p (fixed by world markets), (ii) the research wage w_H^R (endogenous to innovation), (iii) suppliers' profits Π_H, and (iv) the rental income from the fixed factor F_H.

[21] The price of the differentiated good is normalized to one.

Welfare could be decomposed as

$$U_H \;=\; \alpha^\alpha(1-\alpha)^{1-\alpha}\,\big(U_H^{\text{bench}}(T_H,T_W,\tau_H)\;+\;U_H^{\text{innov}}(T_H,T_W,\tau_H)\big), \tag{15}$$

where the first term collects income generated at *baseline* productivity (non-innovating suppliers) and the second term is the *additional* income due to innovation.

Benchmark welfare (static effect of tariffs). The benchmark component is

$$\begin{aligned} U^{\text{bench}} = (w_H^p)^{-\alpha}\Bigg(&\ell_P w_H^p + (w_H^p)^{-\delta}\beta(1-\beta)^{\delta-1}\,\Gamma\left(T_H + T_W\big((1+\tau_H)\tilde{o}_W\big)^{-\theta}\right)^{s-1} \\ &\times\left[(2-\beta)\,T_H + T_W\big((1+\tau_H)\tilde{o}_W\big)^{-\theta}\right]\Bigg), \end{aligned} \tag{16}$$

with $s \equiv \delta/\theta \in (0,1)$ and $\Gamma \equiv \Gamma(1-s) > 0$. For any $T_H > 0$ and $T_W > 0$, and for parameter values

$$\frac{1-\beta}{\beta} \;<\; \theta \;<\; \frac{2-\beta}{\beta},$$

U^{bench} is strictly *decreasing* in τ_H.[22] Intuitively, tariffs raise the cost of imported varieties and reduce allocative efficiency (lower output), while shifting market share and profits toward relatively low-productivity Home suppliers. Under the above (standard) restriction, the efficiency loss dominates the profit-shifting gain, so the benchmark welfare effect of tariffs is negative for all (T_H, T_W).

Innovation welfare (dynamic effect of tariffs). Proposition 4 characterizes the dynamic welfare contribution of innovation in closed form and establishes its shape with respect to the effective World productivity $\widetilde{T}_W$. As a reminder, this aggregate productivity summarizes the competitive pressure that Home suppliers face once World technology, trade costs, and relative labor productivities are all taken into account. The proposition shows two key results: (i) the innovation component of welfare is *always positive*, and (ii) its dependence on foreign competition is *inverted–U shaped*, increasing when the effective World scale is low but declining once foreign competition becomes very strong.

Proposition 4 (Welfare Gain from Innovation)**.** *The innovation component of welfare yields*

[22]Full derivations are in Appendix A.7. Also, please refer to Appendix C for the full derivations of this model with no innovation.

the closed form

$$U^{\text{innov}}(T_H,\widetilde{T}_W) = (w_H^p)^{-\alpha-\delta}\,\beta\,(1-\beta)^{\delta-1}\Bigg[\Big(\tfrac{\ell_H^R}{f}\,I_0(\widetilde{T}_W)\Big)^{1/2}\Big(2^{-1/2}+2^{1/2}(1-\beta)\Big)$$
$$+\Big(\tfrac{2\,\ell_H^R}{f\,I_0(\widetilde{T}_W)}\Big)^{1/2}T_H\Gamma\Bigg((1+a)^{2\delta}\frac{A(\widetilde{T}_W)^s}{P(\widetilde{T}_W)}-(1+a)^{\delta}B(\widetilde{T}_W)^s\Big(\tfrac{1}{P(\widetilde{T}_W)}+\tfrac{1}{R(\widetilde{T}_W)}\Big)+\frac{C(\widetilde{T}_W)^s}{R(\widetilde{T}_W)}\Bigg)\Bigg],$$

where $\alpha=(1+a)^{-\theta}$, $s=\delta/\theta\in(0,1)$, *and* $\Gamma=\Gamma(1-s)>0$. *The auxiliary functions* $A(\cdot)$, $B(\cdot)$, $C(\cdot)$, $P(\cdot)$, $R(\cdot)$, *and* $I_0(\cdot)$ *are positive additive functions of* T_H *and* $\widetilde{T}_W$, *defined in Appendix A.6.*

For any $T_H>0$, $\widetilde{T}_W\geq 0$, $a>0$, $\delta\in(0,\theta)$, $\beta\in(0,1)$, $\ell_H^R>0$, $f>0$, *and* $w_H^p>0$:

1. *The innovation objective is strictly positive,*

$$U^{\text{innov}}(T_H,\widetilde{T}_W)>0.$$

2. *(Hump shape) Under* $\delta\in(0,\theta)$ *and* $a\in(0,1)$,

$$\frac{\partial U^{\text{innov}}}{\partial\widetilde{T}_W}\bigg|_{\widetilde{T}_W=0^+}>0,\qquad \exists\,M<\infty:\ \frac{\partial U^{\text{innov}}}{\partial\widetilde{T}_W}<0 \text{ for all } \widetilde{T}_W\geq M,$$

$$\text{and}\quad \lim_{\widetilde{T}_W\to\infty}\frac{\partial U^{\text{innov}}}{\partial\widetilde{T}_W}=0^-.$$

By continuity, there exists at least one interior maximizer $\widetilde{T}_W^*\in(0,M)$ *with* $\frac{\partial U^{\text{innov}}}{\partial\widetilde{T}_W}(\widetilde{T}_W^*)=0$, *so* U^{innov} *is hump-shaped in* $\widetilde{T}_W$.

Proof. See Appendix A.6. □

The proposition formalizes the intuition developed earlier. When effective World productivity is very low, Home suppliers face little external pressure. Innovation raises productivity but contributes little to welfare, since most domestic suppliers already dominate their markets. As effective World productivity increases, more suppliers are drawn into neck-and-neck competition, where upgrading has the largest payoff. Innovation then rises, and so does its welfare contribution. Eventually, however, if effective World productivity becomes too large, Home suppliers are pushed too far behind: even if they innovate, they cannot realistically overtake their rivals. In that regime, innovation effort declines and the welfare gains from innovation taper off, converging to zero in the limit.

In short, Proposition 4 shows that the dynamic welfare effect of tariffs is non-monotonic. Tariffs shift the effective World scale $\widetilde{T}_W$: raising tariffs compresses the foreign productivity distribution and can stimulate innovation when Home is behind World, but the same policy reduces innovation when Home is ahead. This inverted-U dynamic sets the stage for the full welfare analysis below, where I combine the benchmark and innovation components.

Total welfare and interaction of margins. To evaluate the full effect of tariffs on aggregate welfare, both the static benchmark term and the dynamic innovation term must be considered. The benchmark component declines monotonically with tariffs: making imported varieties more expensive reduces allocative efficiency and output, while the profit shift toward lower-productivity Home suppliers only partially offsets this loss. By contrast, the innovation component responds non-monotonically to tariffs, depending on the relative technological gap between Home and World.

This interaction yields three distinct regimes. First, when Home is *technologically ahead* of World—i.e. $T_H > T_W \tilde{o}_W$—a tariff reduces both benchmark and innovation welfare. Tariffs lower efficiency through distorted input prices and also reduce innovation by relaxing competitive pressure. Fewer Home suppliers are neck-and-neck with their foreign competitors, and thus fewer find it worthwhile to upgrade. In this case, tariffs are *more harmful than what standard trade models would suggest*, as they erode both static efficiency and dynamic productivity gains.

Second, when Home is *far behind*—that is, $T_H \ll T_W \tilde{o}_W$—a small tariff still lowers benchmark welfare significantly, but has negligible impact on innovation. Because the technology gap is too large, tariff-induced changes in relative prices fail to move any domestic suppliers into the responsive region where innovation incentives are high. Hence, the dynamic margin is effectively dormant, and overall welfare closely tracks the benchmark term.

Third, when Home is *slightly behind*—approaching the neck-and-neck region from below—a small tariff can have an offsetting effect. By narrowing the effective productivity gap, it brings more suppliers into the high-return zone for innovation. As these suppliers innovate more, aggregate dynamic welfare rises. In this case, the welfare gain from innovation can partially or even fully offset the benchmark loss. Under certain parameter values, a modest tariff can *increase* total welfare—despite worsening static efficiency—by reallocating effort toward productivity-enhancing innovation. Under most circumstances, however—particularly when innovation steps a are small, fixed innovation costs f are high, or the share of research labor ℓ_H^R is low—the dynamic gain will be too weak to offset the static efficiency loss. In such cases, the total effect of tariffs remains negative.

To visualize these regimes, Figure 3 plots welfare as a function of tariffs for different (T_H, T_W) combinations (with $\tilde{o}_W = 1$ for simplicity). The left panel holds T_W fixed and varies T_H, while the right panel does the reverse. The figure illustrates how the total welfare effect depends crucially on the relative position of Home and World in the global technology distribution.

Figure 3: Home Welfare and Tariffs

(a) Home ahead of Foreign (b) Home slightly behind (c) Home lagging Foreign

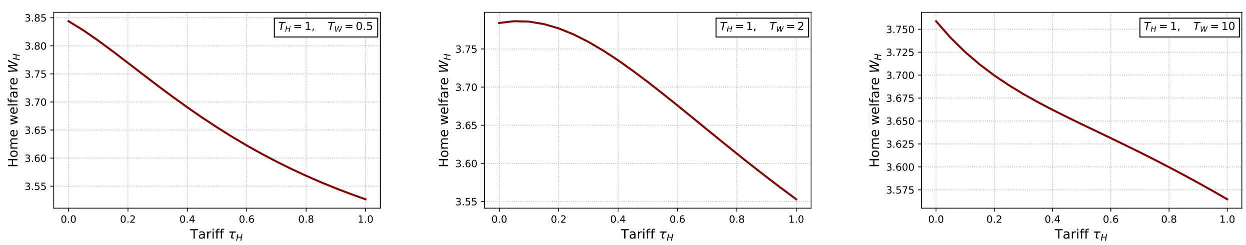

Note: Each panel plots Home welfare U_H as a function of tariffs τ under different relative productivity regimes. Panel (a): Home is ahead of Foreign ($T_H = 1 > T_W = 0.5$); tariffs reduce both efficiency and innovation incentives. Panel (b): Home is slightly behind ($T_H = 1 \lesssim T_W = 2$); tariffs compress price gaps, stimulate innovation, and may raise total welfare. Panel (c): Home is far behind ($T_H = 1 \ll T_W = 10$); tariffs lower welfare without generating meaningful innovation gains.
Parameters are set to $\beta = 0.5$, $\delta = 1$, $\theta = 3$, and $a = 0.1$. Here β governs demand curvature (markup elasticity), $\delta \equiv 1/\beta - 1$ is the implied return exponent in innovation payoffs, θ is the Fréchet shape parameter governing productivity dispersion, and a is the step size of innovation (a successful upgrade scales productivity by $(1 + a)$). See Appendix A.7 for derivations.

3 Empirical analysis of buyers' sourcing decisions and suppliers innovation expenditures

With the theory in place and its firm-level predictions made explicit, I now take the model to the data. The empirical analysis is designed to map the model's objects to observables and to test its core mechanisms using administrative firm–product data from Turkiye. The purpose of the empirical analysis is threefold. First, it documents that downstream buyers in Turkiye typically source each product from a single upstream supplier, consistent with the model's buyer–supplier structure. Second, it examines how changes in import prices affect sourcing patterns, showing that rising trade costs induce buyers to reallocate demand from foreign to domestic suppliers. Third, it tests the model's central prediction that domestic suppliers' innovation effort follows an inverted-U relationship with the intensity of foreign competition.

3.1 The structure and persistence of buyer-supplier relationships in Turkiye

A key feature of the theoretical framework is that each firm sources a given input from a single supplier and that these sourcing relationships are fluid. This assumption is crucial: when suppliers compete à la Bertrand within narrowly defined input markets, small changes in relative prices can trigger discrete switches in sourcing, giving rise to the model's inverted-U relationship between competition and innovation.

To examine whether the model's microfoundations are empirically relevant, I turn to detailed firm-level data from Turkiye. Turkiye provides an ideal setting for this analysis: it is a small open economy where import competition plays a central role in shaping firms' sourcing and innovation decisions, and where administrative records capture the full network of firm-to-firm transactions and import flows.[23]

In this section, I use these data to validate the model's key features. Specifically, I document two stylized facts about buyer–supplier relationships in the Turkish manufacturing sector: (1) within each sector, purchases are highly concentrated in a single supplier; and (2) supplier relationships are short-lived and fluid.[24]

To establish these facts, I combine two administrative datasets covering all formal firms in Turkiye between 2009 and 2019. The first is the universe of firm-to-firm transactions from the value-added tax (VAT) registry, which records every domestic sale between registered firms and identifies both the buyer and the seller. The second is the universe of customs declarations, which record import transactions by firm, product, and origin country. Merging these datasets yields a comprehensive view of each firm's sourcing structure—both domestic and foreign—on an annual basis.

For the purposes of this study, I restrict attention to domestic purchases from suppliers engaged in manufacturing or wholesale trade (NACE Rev. 2 codes 10–32 and 45–46), focusing on intermediate input linkages relevant for production rather than services. After these filters, the resulting dataset contains 838,904 firm–year observations.

Fact 1. Concentrated sourcing within sectors. Firms typically rely on a single supplier—either a domestic firm or a foreign country—within each *sector*, defined at the four-digit NACE level, in a given year. Because VAT transactions are not recorded at the

[23]Turkiye relies heavily on imported inputs—imports accounted for 30.1% of GDP in 2019—so shifts in import prices directly shape the competitive environment faced by domestic producers.

[24]While firms typically single-source within a sector, they maintain multiple supplier relationships overall—on average, twelve per firm per year—spanning a range of sectors and countries. This reflects the multi-input nature of production but does not affect the model's assumption of single-sourcing at the product level. Bacilieri et al. (2023); Bernard et al. (2022) also document similar facts with many sourcing.

product level, I measure input sourcing at the supplier's sector level.[25] Figure 4 shows that more than 70% of buyer–sector–year observations involve a single supplier.[26] This high degree of concentration implies that sourcing decisions are inherently discrete. Such lumpiness is central to the model's mechanism, as it allows small price changes to reallocate demand across suppliers and create strong, state-dependent incentives to innovate.

Figure 4: Number of Suppliers Per Sector

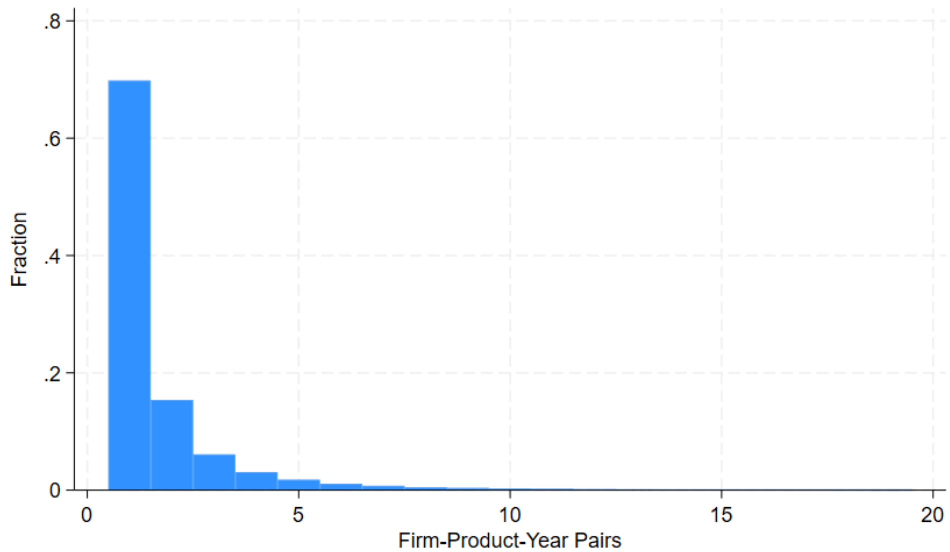

Notes: The figure plots the distribution of the number of suppliers per buyer–sector–year observation in the Turkish VAT dataset. Each bar shows how many distinct suppliers a buyer sources from within a given sector (four-digit NACE) in a given year. The data cover the period 2009–2019. This measure is defined at the sector level because VAT transaction records do not include product identifiers.

Fact 2. Short-lived and fluid relationships. Supplier links are brief and continuously reshuffled. The average duration of a buyer–supplier relationship is just over two years (median = 1), and firms add and drop suppliers at nearly identical rates—six (median = 3) per year on average—so that the total number of suppliers per firm remains stable despite high turnover. This churn indicates that buyer–supplier relationships are highly contestable: domestic firms anticipate substitution as a realistic possibility and therefore perceive innovation as a means of attracting and retaining demand. Such fluidity reinforces the model's assumption that sourcing decisions respond quickly to relative price changes.

[25]A four-digit NACE sector is more aggregated than a product category (defined later using ten-digit PRODTR codes). Hence, the 70% figure reported below likely understates the true extent of single-sourcing at the product level.

[26]This complements the findings of Antras et al. (2017).

Table 1: Supplier Relationship Dynamics

	Mean	P25	Median	P75	Obs.
Duration of buyer–supplier relationship (years)	2.10	1	1	2	4,631,072
Number of dissolved relationships per firm–year	6.13	1	3	6	838,904
Number of new relationships per firm–year	6.14	1	3	7	838,904
Total number of suppliers per firm–year	11.57	2	5	11	838,904

Notes: Statistics are computed at the firm–year level using the buyer–supplier panel of Turkish VAT data, 2009–2019. Supplier relationships are short-lived on average (median duration of one year). Firms add and drop a similar number of suppliers each year, so the overall number of suppliers per firm–year remains relatively stable.

Because supplier relationships are frequently dissolved and re-formed, small shifts in relative prices can trigger discrete switches in sourcing. This discreteness is precisely what creates strong incentives for domestic suppliers: by lowering their relative prices through productivity improvements, they can trigger a switch in their favour and capture new demand. This reallocation channel provides the microeconomic foundation for the inverted-U innovation response that I test in the next section.

3.2 How trade cost shocks impact buyers' choice of suppliers in Turkiye

I examine how changes in quality-adjusted import prices affect buyers' sourcing decisions. The analysis is carried out in two steps. First, I estimate how changes in foreign input prices, adjusted for quality differences, influence the firms' import decision at both the intensive and extensive margins. Second, I examine whether changes in quality-adjusted prices lead buyers to substitute toward domestic suppliers.

Constructing and instrumenting quality-adjusted import prices A key variable in the analysis is the firm–product–country–year measure of *quality-adjusted import prices*. Following Khandelwal et al. (2013), I recover country-level product quality from observed import values and unit prices under a CES demand framework. Specifically, I estimate

$$\ln(val_{fcpt}) + \frac{1-\beta}{\beta}\ln(p_{fcpt}) = \gamma_{pt} + \eta_{ft} + \nu_{ct} + \epsilon_{fcpt}, \tag{17}$$

where val_{fcpt} denotes the value of firm f's imports of product p from country c in year t, and p_{fcpt} is the corresponding average unit value. The fixed effects γ_{pt}, η_{ft}, and ν_{ct} absorb product–year, importer–year, and exporter–year variation, respectively. The residual ϵ_{fcpt} captures exporter-specific product quality under the assumption that higher-quality goods command higher prices and larger market shares.

The implied quality-adjusted log price is then

$$\ln(qadjp_{fcpt}) = \ln(p_{fcpt}) - \frac{\beta}{1-\beta}\epsilon_{fcpt},$$

which nets out the price premium attributable to quality, isolating the effective price conditional on quality differences.

Instrumenting for exogenous variation. Quality-adjusted import prices may still reflect Turkish demand shocks—for instance, if Turkish firms upgrade quality preferences or expand production in response to local shocks. To isolate variation that is plausibly exogenous to Turkish conditions, I instrument changes in firm–product–country quality-adjusted prices with corresponding changes in export prices to other destinations, in particular the European Union, using UN Comtrade data. The idea is that global shocks to exporters' costs or technologies affect prices faced by all destinations, including Turkiye, but are orthogonal to Turkish demand. This strategy follows the logic of related shift–share instruments used in the trade literature. The resulting instrument has strong predictive power for Turkish import prices but is unlikely to capture domestic demand or innovation responses.

Import margins. I begin by estimating a two-stage least squares (2SLS) specification linking firm–product–country–year import values to quality-adjusted import prices:

$$\ln(\text{Value}_{fpct}) = \beta \ln(qadjp_{fpct}) + \alpha_{fpc} + \alpha_{fpt} + \alpha_{ct} + \epsilon_{fpct}, \tag{18}$$

where α_{fpc}, α_{fpt}, and α_{ct} are firm–product–country, firm–product–year, and country–year fixed effects, respectively.

The first stage instruments Turkish import prices with global quality-adjusted export prices (excluding Turkiye):

$$\ln(qadjp_{fpct}) = \ln(qadjp^{exp}_{pct}) + \alpha_{fpc} + \alpha_{fpt} + \alpha_{ct} + \nu_{fpct}.$$

Column (1) of Table 2 confirms that the instrument is highly relevant: when global export prices increase, Turkish import prices co-move strongly in the same direction. The exclusion restriction is plausible since Turkiye accounts for a negligible share of world trade, so domestic shocks cannot influence global prices.

Turning to the second stage, Column (2) shows that higher import prices significantly reduce import values, with an implied elasticity of substitution of about -4.7. This estimate is well within the range reported in the international trade literature and indicates a strong intensive-margin response. Column (3) examines the extensive margin by regressing a firm–product import dummy on import prices.The coefficient of -0.71 implies that a one–log–point increase in a foreign supplier's quality-adjusted price reduces the probability that the buyer continues importing that product by 71 percentage points. This is a large and economically meaningful effect, indicating that import relationships are highly contestable: when a supplier's relative price rises substantially, buyers are very likely to switch away, consistent with the short-lived and fluid buyer–supplier relationships documented earlier.[27]

Table 2: How import prices impact imports of Turkish buyers

	FS	(2)	(3)
	$\ln(\text{Quality Adj. Price}_{fpct})$	$\ln(\text{Imports}_{fpct})$	Import Dummy$_{fpct}$
$\ln(\text{Quality Adj. Price}^{EU}_{pct})$	0.0511***		
	(0.00397)		
$\ln(\text{Quality Adj. Price}_{fpct})$		-3.739***	-0.712***
		(0.582)	(0.0861)
Observations	382,554	382,554	705,504
Firm-Product-Country FE	Yes	Yes	Yes
Firm-Product-Year FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes
F-Stat	-	352.5	796.2

Notes: Column (1) reports the first-stage regression, where the instrumented variable is $\ln q^{EU}_{adjp}$. Columns (2) and (3) report the corresponding second-stage (IV) regressions, in which the log of imports and the import dummy are regressed on the predicted quality-adjusted import price. All regressions include firm–product–country, firm–product–year, and country–year fixed effects. Standard errors are clustered at the product level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Domestic margins. Because product-level information is unavailable for domestic purchases, I aggregate import prices into a firm-level index that captures the average quality-adjusted cost of imported inputs. The (log) index combines product-level quality-adjusted

[27] Because the regression is estimated on firm–product–country panels with importer fixed effects, the dependent variable reflects the continuation of an existing import relationship—taking the value one when the firm continues importing and zero when it ceases to do so.

prices using two-year-lagged sales weights:

$$\ln(qadjp_{ft}) \equiv \sum_{p \in P_{f,t-2}} w_{fp,t-2}\, x_{fpt}, \qquad w_{fp,t-2} = \frac{\text{sales}_{fp,t-2}}{\sum_{p' \in P_{f,t-2}} \text{sales}_{fp',t-2}},$$

so that the weights are predetermined with respect to contemporaneous import shocks.

To understand how trade shocks affect domestic sourcing, I estimate at the buyer–supplier level:

$$\ln(\text{Value}_{ff't}) = \beta \ln(qadjp_{ft}) + \alpha_W + \alpha_{f't} + \alpha_{st} + \epsilon_{ff't},$$

where $\text{Value}_{ff't}$ is buyer f's purchases from domestic supplier f' in year t, and α_W, $\alpha_{f't}$, and α_{st} denote buyer, supplier–year, and sector–year fixed effects. Column (1) of Table 3 reports a coefficient of 0.053, indicating that a one–log–point increase in a buyer's import price index raises purchases from a given domestic supplier by 0.053%. Hence, when import prices rise, buyers expand orders from their existing domestic suppliers, reflecting an intensive-margin adjustment that reallocates demand toward local firms.

I then examine whether higher import prices lead firms to form new domestic relationships. At the firm–year level, I estimate

$$\ln(\#\text{domestic suppliers}_{ft}) = \beta \ln(qadjp_{ft}) + \alpha_W + \alpha_{st} + \epsilon_{ft},$$

where α_W and α_{st} are firm and sector–year fixed effects. Column (2) of Table 3 reports a coefficient of 0.185, indicating that a one–log–point increase in a firm's weighted import price index is associated with a 0.185% increase in the number of its domestic suppliers. Thus, when foreign inputs become costlier, firms expand their domestic supplier base.

Combining the extensive–margin results for imports and domestic purchases reveals a coherent substitution pattern. A one–log–point increase in import prices reduces the probability of importing by 0.71% while increasing the number of domestic suppliers by 0.185%. Given that the median firm purchases from five suppliers, the implied change of $0.185 \times 5 \approx 0.93\%$ closely mirrors the 0.71% decline in import participation. This near one-for-one relationship suggests that firms typically replace foreign suppliers with domestic ones when relative import costs rise, consistent with the model's single-sourcing assumption.

Firm sales. Replacing the dependent variable with $\ln(\text{sales}_{ft})$ yields a coefficient of -0.085, indicating that higher import prices are associated with lower firm output. A one–log–point increase in the weighted import price index reduces sales by about 0.085%, suggesting that substitution toward domestic inputs only partially offsets the loss of cheaper imported inputs.

Taken together, these results show that foreign and domestic suppliers are close substitutes: higher import prices reduce both the intensive and extensive margins of importing while simultaneously increasing reliance on domestic suppliers. However, the substitution is incomplete, as firms' total sales still decline when faced with higher import costs.

Table 3: How import prices impact domestic sourcing of Turkish buyers

	(1) $\ln(\text{Domestic Purchases}_{ff't})$	(2) $\ln(\text{No. Domestic Suppliers})_{ft}$	(3) $\ln(\text{Sales})_{ft}$
$\ln(\text{Quality Adj. Price}_{ft})$	0.0532***	0.1835***	-0.0852***
	(0.0011)	(0.0028)	(0.0077)
Observations	20,775,958	359,995	359,995
Firm FE	Yes	Yes	Yes
Seller-Year FE	Yes	-	-
Year FE	-	Yes	Yes
F-Stat	9.8e+05	1.5e+04	1.5e+04

Notes: All columns report instrumental-variable (IV) estimates where the endogenous regressor is the quality-adjusted import price. Column (1) reports results where the outcome is the log value of purchases of importer f from domestic supplier f' at time t; the specification includes importer fixed effects and supplier–year fixed effects. Columns (2) and (3) report results where the outcomes are, respectively, the number of distinct domestic suppliers from which firm f sources at time t, and the log value of its total sales; the specification includes firm and year fixed effects. Standard errors, clustered at the firm level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These patterns highlight two key mechanisms. First, substitution occurs along both extensive and intensive margins—whether to source domestically at all, and how much to purchase from each source. Second, because firms typically rely on a single supplier per input, these adjustments often take the form of discrete switches rather than gradual reallocations. This discreteness underpins the incentives for domestic firms: by lowering their relative prices through innovation, they can trigger a switch in their favor and capture new demand. This reallocation channel provides the microeconomic foundation for the inverted-U innovation response explored in the next subsection.

3.2.1 Additional checks

I also explore adjustment margins not explicitly modeled in the theory. Table J10 in the Appendix examines product-level responses to relative price changes. As domestic prices rise relative to foreign competitors, firms reduce the number of products they produce and are less likely to enter new product lines. This pattern suggests that intensified competition induces firms to drop their least productive products rather than diversify into new ones. While product-scope adjustments lie outside the model's single-product framework, they provide complementary evidence that competition reallocates activity toward more efficient producers.

3.3 Turkish suppliers' incentives to innovate

Having established that buyers make discrete sourcing choices and substitute toward domestic suppliers when trade costs rise, I now test the model's central prediction: domestic firms' innovation effort is hump–shaped in their proximity to foreign competitors. I study how the share of the wage bill paid to R&D workers varies with a firm–year measure of proximity, D_{ft}, constructed from quality–adjusted price gaps. I first provide a visual benchmark and then estimate a quadratic IV specification that instruments D_{ft} with global price movements.

Data sources. Domestic production data come from the Turkish *Prodcom* survey (PRODTR) compiled by TurkStat, which reports the value and quantity of every product produced by firms between 2009 and 2019. Products are defined at the 10-digit PRODTR level—two digits more disaggregated than the standard EU *Prodcom* classification— providing highly granular product coverage.[28] Import transactions are drawn from the Turkish *Customs* dataset, which records all imports at the 12-digit *GTIP* (Customs Tariff Statistics Position) level. Since PRODTR and GTIP codes are not directly comparable, I use TurkStat's official concordance to harmonize the two systems; most cases are many-to-one mappings from GTIP to PRODTR.[29] Prices are measured as unit values—the ratio of reported value to quantity—at the firm-product-year level for domestic data and the country-product-year level for imports. About 85% of PRODTR products share identical units with their GTIP counterparts, allowing direct comparison; products with mismatched units are excluded.

Measuring proximity to foreign competitors. A central variable in the analysis is firm-level proximity to foreign competitors, derived directly from the model. Proposition 2 predicts that innovation follows an inverted-U in this measure: firms that are "neck-and-neck" with foreign rivals invest most in upgrading, while laggards and leaders invest less.

I measure proximity using a quality-adjusted price gap that compares each firm's output prices to those of imported alternatives at the product level. Formally, for firm f in

[28] Appendix D describes the survey design and provides examples of product definitions.
[29] Appendix E documents the concordance procedure and presents descriptive statistics.

year t,

$$D_{ft} = \sum_{p \in P_{ft}} w_{fp,t-2} \left(x_{fpt} - y_{pt} \right),$$

where x_{fpt} is the firm's quality-adjusted log unit value for product p, y_{pt} is the modal quality-adjusted log import price for the same product, and $w_{fp,t-2}$ are two-year-lagged sales weights capturing the importance of each product in the firm's historical portfolio.

Quality adjustment follows Khandelwal et al. (2013), who recover product quality from a CES demand system; Appendix F details the estimation and construction. Intuitively, D_{ft} is positive when a firm's effective price exceeds that of its foreign rivals (a disadvantage) and negative when the firm is more competitive.

Figure F1 shows the distribution of firm-product price differences after adjusting for quality. The distribution is centered above zero: the median gap is 0.10, implying that half of Turkish firms charge prices at least 10 percent higher than foreign competitors, even after adjusting for quality. This highlights the competitive pressure Turkish producers face in globalized product markets.

Outcome variable: R&D wage bill. The main outcome of interest is the firm-level wage bill allocated to R&D workers. In the model, R&D expenditures correspond to wages paid to technical labor, which I proxy using the Turkish *4A Active Insured Administrative Records Micro* dataset—an employer–employee panel covering all insured workers in formal firms between 2009 and 2019. The dataset reports each worker's firm identifier, occupation (ISCO–08), and monthly earnings.

I link these employment records to the *Prodcom* sample using firm identifiers, yielding approximately 10,500 unique firms and 96,000 firm–year observations. Following Bernard et al. (2017), I classify occupations into broad categories and define R&D workers (Group R) as scientists, engineers, and technicians (ISCO 211X, 214X, 31XX).

For each firm–year, the R&D wage bill share is the fraction of total annual wages paid to Group R employees:

$$s_{ft}^{\text{tech}} = \frac{\text{wages to R\&D workers}_{ft}}{\text{total wage bill}_{ft}}.$$

More than 60 percent of firms employ at least one R&D worker during the sample period. Appendix G describes the occupational classification in detail and provides summary statistics.

Instrumenting for the price gap. To estimate the inverted-U relationship between innovation and proximity to foreign competitors, I require exogenous variation in the firm-level price gap D_{ft}. Direct OLS estimates would be biased if firms that expect stronger competitive pressure adjust their markups or innovate in ways that affect their own prices.

I therefore instrument D_{ft} with a Bartik-style measure based on global movements in quality-adjusted export prices that exclude Turkiye. Specifically, I interact world price shocks for each product with firms' historical (two-year lagged) product shares. This design captures how international price shifts—originating entirely outside the Turkish economy—alter the relative competitiveness of domestic firms depending on their exposure to particular products.

The instrument is both *relevant*, since Turkish import prices co-move with global export prices for the same products, and satisfies the *exclusion restriction*, as it is constructed from trade flows excluding Turkiye and weighted by lagged product shares. Because Turkiye represents less than 2% of global import demand and less than 5% of export supply in almost all product–country pairs, it is implausible that Turkish firms' behavior feeds back into the world prices used in the instrument.[30]

3.3.1 Inverted-U

To test the predicted inverted-U relationship, I estimate the following firm-level panel regression:

$$s_{ft}^{tech} = \alpha_W + \gamma_{st} + g(D_{ft}) + \epsilon_{ft}, \tag{19}$$

where s_{ft}^{tech} denotes the share of the wage bill allocated to R&D workers by firm f in year t, α_W is a firm fixed effect capturing time-invariant heterogeneity, and γ_{st} is a sector-year fixed effect accounting for common shocks. The term D_{ft} is the firm-level quality-adjusted price gap with foreign competitors, constructed as described above.

I begin with a binned scatter of the R&D wage bill on D_{ft} for visual reference. I then approximate $g(D_{ft})$ with a quadratic, $g(D_{ft}) = \beta_1 D_{ft} + \beta_2 D_{ft}^2$, and estimate this specification using instrumental variables (IV) to address the endogeneity of D_{ft}.

Binned Scatterplot. Figure 5 provides a visual benchmark for the relationship between firms' R&D wage bill and their proximity to foreign competitors (D_{ft}), controlling for firm and sector–year fixed effects. The binscatter offers an intuitive illustration of the inverted-U predicted by the model.

[30]Appendix H provides the full derivation of the instrument, including the estimation of global quality adjustments and the construction of firm-level exposure shares.

Panel (a) shows the conditional correlation using observed D_{ft}, while Panel (b) replaces D_{ft} with fitted values from the first-stage regression on the global price instrument. Both plots reveal a hump-shaped relationship: R&D spending peaks when firms are closest to their foreign competitors, and the curvature is more pronounced in the IV-adjusted plot, consistent with endogeneity bias attenuating the raw relationship.

Figure 5: Binned Scatterplots of Log R&D Wage Bill and Firms' Proximity to Foreign Competitors

(a) Raw (b) IV-Adjusted

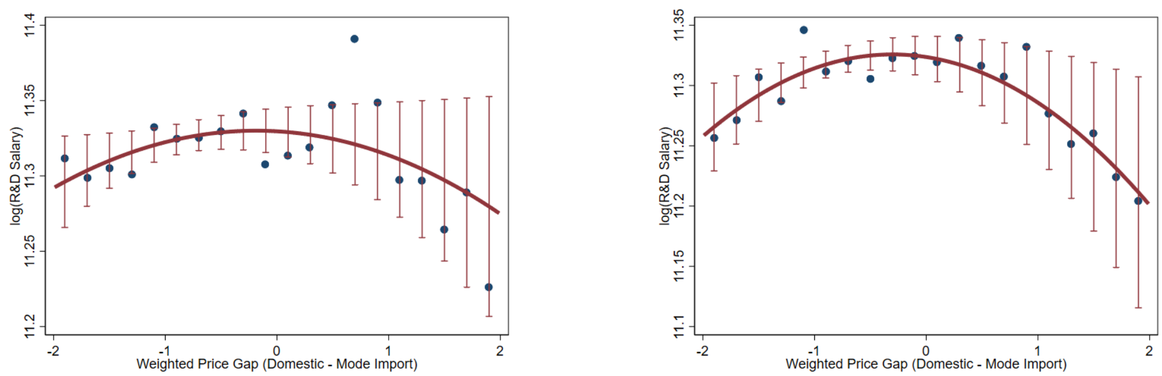

Notes: The figure plots binscatters of the log R&D wage bill against firms' proximity to foreign competitors (D_{ft}) using 20 bins, with firm and sector-year fixed effects. Panel (b) instruments D_{ft} with global price shocks, as described above. Robust bias-corrected confidence intervals are shown, following Cattaneo et al. (2019).

Quadratic Parametric Regressions. Having visually demonstrated the relationship between R&D spending and proximity to foreign competitors using binned scatterplots, I now turn to a formal regression analysis. Specifically, I estimate a quadratic IV model to test for an inverted U-shaped relationship. The second-stage regression is given by

$$s_{ft}^{\text{tech}} = \alpha_W + \gamma_{st} + \beta_1 D_{ft} + \beta_2 D_{ft}^2 + \varepsilon_{ft}, \tag{20}$$

where s_{ft}^{tech} denotes the share of the wage bill devoted to technical workers, α_W and γ_{st} are firm and sector-year fixed effects.

Because both D_{ft} and D_{ft}^2 are endogenous, they are instrumented using the weighted average of global export prices, z_{ft}, and its square, z_{ft}^2, as described above. The corresponding first-stage equations are

$$D_{ft} = \alpha_W + \gamma_{st} + \pi_{11} z_{ft} + \pi_{12} z_{ft}^2 + \nu_{1,ft}, \tag{21}$$

$$D_{ft}^2 = \alpha_W + \gamma_{st} + \pi_{21} z_{ft} + \pi_{22} z_{ft}^2 + \nu_{2,ft}. \tag{22}$$

Columns FS1–FS2 of Table 4 report the first-stage estimates. As expected from the construction, z_{ft} is negatively correlated with D_{ft} (a global increase in import prices reduces the price gap between domestic and imported products within Turkiye), while z_{ft}^2 is positively correlated with D_{ft}^2. Both coefficients have the predicted signs. The F-statistic for the joint significance of the excluded instruments is 80.21, comfortably above conventional weak-instrument thresholds (e.g., 20), indicating strong first-stage relevance.

The second-stage IV estimation results are presented in Column (1) of Table 4.[31] The estimates provide evidence of an inverted U-shaped relationship between the share of tech wage bill and the competitive price gap: the coefficient on D_{ft}^2 is negative and statistically significant. The implied peak of the quadratic curve is at $-\beta_1/(2\beta_2) = 0/(2\cdot 0.0538) = 0$. This suggests that the share of tech wage bill is maximized when the domestic firm's quality adjusted price matches that of its foreign competitor. Interpreted through the lens of the theoretical model, this corresponds to Turkish firms that are neck-and-neck with the majority of their foreign competitors, and thus have the strongest incentive to invest in innovation to gain market access.

To fix ideas, consider a firm whose quality-adjusted price is one standard deviation above its foreign competitor, i.e. $D_{ft} = 0.73$ (107% higher). Suppose a foreign shock closes this gap so that the firm moves to $D_{ft} = 0$ (neck-and-neck). Under the quadratic specification, the implied change in the share of the wage bill allocated to technical workers is

$$\Delta s^{\text{tech}} = \left[\beta_1 D + \beta_2 D^2\right]_{D=0} - \left[\beta_1 D + \beta_2 D^2\right]_{D=0.73} = \beta_1(0 - 0.73) + \beta_2(0^2 - 0.73^2).$$

Using the estimated coefficients ($\beta_1 = 0$, $\beta_2 = -0.0538$) gives

$$\Delta s^{\text{tech}} = 0.0538 \times 0.73^2 \ = \ 0.0286,$$

that is, a *2.86 percentage-point* increase in s^{tech}. Given that the median firm allocates 4.5% of its wage bill to technical workers, this corresponds to an increase of roughly $0.0286/0.045 \approx 63\%$ relative to the median.

For comparison, I re-estimate the same quadratic specification with alternative depen-

[31]OLS estimates of the same specification are reported in Table J1.

Table 4: Quadratic Regression Results: R&D vs. Other Employment Payments

	(FS1) D_{ft}	(FS2) D^2_{ft}	(1) s^{tech}_{ft}	(2) $\ln(sales_{ft})$	(3) Prod Emp$_{ft}$	(4) Mrk Shr$_{ft}$	(5) ln(profit$_{ft}$)	(6) Exporter$_{ft}$
D_{ft}			0.0418	-0.238**	-0.500**	-0.00578*	-0.361*	0.0216
			(0.0952)	(0.102)	(0.206)	(0.00303)	(0.20)	(0.0343)
D^2_{ft}			-0.0538***	0.0645***	0.0516	-0.00000745	-0.0940	-0.000545
			(0.0198)	(0.0213)	(0.0429)	(0.000630)	(0.060)	(0.00714)
z_{ft}	-0.0581***	0.0498*						
	(0.00627)	(0.0270)						
z^2_{ft}	0.0303***	0.111***						
	(0.00162)	(0.00697)						
Observations	95,851	95,851	95,851	95,851	95,851	95,851	73,760	95,851
R^2	0.845	0.789	-	-	-	-	-	-
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-Stat			80.21	80.21	80.21	80.21	43.48	80.21

All regressions include firm and sector-year fixed effects. Columns *FS1* and *FS2* report the first-stage results of the 2SLS estimation. Columns (1)–(6) present the second-stage IV regressions, all measured at the firm f at time t: the share of tech. wage bill to production wage bill in column (1); log firm sales in column (2); Production wage bill in column (3); market share in column (4); log profit levels in column (5); and a binary indicator for whether the firm is an exporter in column (6). Standard errors are reported in parentheses and are clustered at the sector level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

dent variables—log sales, the production wage bill (normalized by its sample median), market share, log profits, and a binary export indicator:

$$Y_{ft} = \alpha_W + \gamma_{st} + \beta_1 D_{ft} + \beta_2 D_{ft}^2 + X_{ft}'\delta + \varepsilon_{ft}. \tag{23}$$

Results are reported in Columns 2–6 of Table 4. A one–standard deviation decrease in proximity (from $D_{ft} = 0.73$ to $D_{ft} = 0$) is associated with: (i) sales rising by 15%; (ii) the production wage bill increasing by 36% (relative to its median); (iii) market share rising by 0.44 percentage points (equivalent to a 6% increase in the median firm's market share); and (iv) profits increasing by 31%. By contrast, the probability of exporting does not change significantly, consistent with exporters typically being far ahead of the foreign mode and thus less sensitive to movements in D_{ft}.[32] Unlike the wage bill on tech workers, the results in Columns (2)–(6) do not exhibit a clear inverted U-shaped pattern.

3.3.2 Extensions and Robustness checks

Nonparametric Estimation. As a complementary exercise, I estimate the innovation response to proximity using a nonparametric IV (NPIV) approach. This method flexibly recovers the shape of $g(D_{ft})$ without imposing a quadratic form, while instrumenting for the endogeneity of D_{ft}. Because NPIV estimation is technically demanding and sensitive to ill-posedness, I follow Chetverikov and Wilhelm (2017) in imposing monotonicity within subsamples, which stabilizes the estimates. The results, reported in Appendix Table I and Appendix Figure I1, show a rising relationship when domestic firms price below the modal foreign competitor and a declining relationship when they price above, consistent with the inverted-U documented in the quadratic regressions.

Robustness to trade dynamics and pre-trends. One concern is that the inverted U-shape may reflect changes in firms' own trade activities rather than shifts in domestic competitive pressure. To address this, I augment the baseline regression with exporter–sector–year and importer–sector–year fixed effects, which absorb time-varying shocks specific to trading firms within each sector. As reported in Appendix Table J5, the inverted U-shaped response remains robust.

I also test for pre-trends by examining whether R&D wage bill shares in $t-1$ respond to changes in proximity at t. As shown in Appendix Table J6, there are no significant effects, ruling out differential pre-trends.

[32] Percent changes are computed from the fitted quadratic by evaluating the outcome at $D_{ft} = 0.73$ and $D_{ft} = 0$.

Alternative outcome measures. As an alternative measure of R&D activity, I replace the wage bill of R&D workers with balance sheet expenditures that include both labor and non-labor components. Appendix Table J7 shows that the inverted U-shaped response remains robust, confirming that results are not driven by the specific R&D measure. For completeness, Appendix Figure J4 also presents binscatter evidence using the share of the wage bill as the outcome, which displays a similar non-monotonic pattern.

Additional checks. To ensure the inverted U-shape is specific to innovation activities, Appendix Figure J3 repeats the analysis for non-R&D workers. In contrast to R&D labor, spending on non-R&D workers declines monotonically with D_{ft}, consistent with competitive pressure affecting scale but not inducing technological upgrading. This pattern accords with the theory: as foreign competitors become more productive, demand shifts away from domestic firms, reducing their sales and thereby lowering the demand for production labor.

I also consider firms in markets with minimal import competition—defined as those where the domestic mode price is at least half the imported mode price. As shown in Appendix Figure J6, the inverted U-shape disappears in this subsample, with R&D labor spending declining monotonically in D_{ft}. This supports the interpretation that the hump-shaped response arises only when foreign competition is meaningful.

Finally, Appendix Table J2 shows that the inverted U-shape remains robust when including additional firm-level controls (age, cash ratio, sales, and the gap from the domestic mode price).[33]

Heterogeneity. I next explore whether the strength of the inverted U-shaped response varies across firms and sectors. Splitting the sample by sector-level profitability, I find that the response is statistically significant only in above-median profit sectors (Appendix Table J8, Columns 5–8). This suggests that in industries where the payoff to capturing market share is high, firms invest more aggressively in technology when close to foreign competitors.

Liquidity is another margin of heterogeneity. Splitting firms by their cash ratio, I find that liquidity-constrained firms exhibit a sharper inverted U-shaped response (Appendix Table J9, Column 4), while firms with higher liquidity display a flatter pattern (Column 8). This is consistent with the idea that constrained firms allocate scarce resources more selectively, focusing R&D investment where expected returns are highest.

[33]The results are also robust when using the log R&D wage bill as the dependent variable (Appendix Table J3), and when restricting the sample to firms with positive R&D employment (Appendix Table J4).

4 Calibration and Quantification

Having established both theoretically and empirically that firms' incentives to innovate depend non-monotonically on their technological distance from foreign competitors, I now return to the model to quantify how this mechanism shapes aggregate welfare. The central question of the paper—how tariffs affect welfare when firms can innovate—requires combining the model's closed-form expressions for welfare with parameter values disciplined by data. While the theoretical analysis shows that tariffs simultaneously alter static misallocation and dynamic innovation incentives, the relative strength of these forces depends on the parameters governing demand, productivity dispersion, and innovation behavior. To quantify these effects, I calibrate the model to Turkish microdata, using exactly identified equilibrium relationships that tie each parameter to a distinct empirical moment. With the calibrated parameter vector in hand, I then compute equilibrium outcomes and simulate how tariffs reshape production, innovation, and welfare.

Each structural parameter is mapped to a specific moment in the data that reflects its economic role. Parameters are organized into three groups according to their function and source: (i) preferences and market structure, (ii) technology and heterogeneity, and (iii) innovation and labor allocation.

The first group governs household behavior and the elasticity of substitution across varieties. These include the expenditure share on the homogeneous good (α), the price elasticity parameter (β), and the relative World wage (o_W). Each is drawn from aggregate or directly estimated data. The elasticity parameter β is estimated using customs data following Equation (18), with the resulting elasticity reported in Table 2. The expenditure share α is calculated from TurkStat national accounts as the share of Turkish consumption in agriculture and mining. These sectors are chosen because they are largely homogeneous and therefore provide a clean measure of the non-differentiated component of consumption. Finally, the relative World wage o_W is determined using the ratio of Turkiye's average wage to that of its trading partners, weighted by bilateral trade shares using World Bank data.

The second group captures productivity dispersion and relative competitiveness between domestic and foreign firms. The Fréchet shape parameter (θ) is estimated from within-product price dispersion across firm-level transactions. The model implies that the standard deviation of log prices is inversely related to θ:

$$\text{sd}[\ln p] = \frac{\pi}{\sqrt{6} \cdot \theta},$$

which can be inverted to yield

$$\hat{\theta} = \frac{\pi}{\sqrt{6} \cdot \text{sd}[\ln p]}, \tag{24}$$

where $\text{sd}[\ln p]$ is the empirical standard deviation of log prices across firms at the product-year level.[34]

The quality aggregates (k_H, k_W) are computed from median firm–product quality residuals $\exp(\varepsilon_{fpt})$ estimated from the following regression:

$$\ln(val_{f(c)pt}) + \left(\frac{1}{\beta} - 1\right) \ln(p_{f(c)pt}) = \gamma_{pt} + \varepsilon_{f(c)pt},$$

where $\ln(val_{f(c)pt})$ denotes the log of firm f's total sales (price times quantity) in product p at time t, and γ_{pt} captures product–time fixed effects. The residual $\varepsilon_{f(c)pt}$ thus represents a firm–product-specific quality component, after netting out price and aggregate demand effects. The aggregate quality index for each country c is then obtained as:

$$k_c = \text{median}_{f(c)} \left[\exp(\varepsilon_{f(c)pt})\right], \tag{25}$$

where c = TR for Home and c = W for World. Intuitively, a higher k_c reflects higher average quality (or technology-adjusted performance) among producers from country c.

The productivity scale ratio T_H/T_W is inferred from the ratio of domestic to imported market shares implied by the model. Aggregating sourcing decisions across downstream buyers, the share of total inputs sourced domestically depends on the relative productivity of domestic suppliers, adjusted for differences in product quality and wages. Formally,

$$\frac{T_H}{T_W} = \text{shr}_d \cdot \left(\frac{k_W}{k_H}\right)^{\theta/\delta} \cdot o_W^{-\theta}, \tag{26}$$

where shr_d is the observed domestic share of total purchases, $\delta = 1/\beta - 1$ reflects the elasticity of substitution across varieties, and $o_W = w_W/w_H$ denotes the relative foreign-to-domestic wage. The quality aggregates (k_H, k_W) enter directly here, capturing average differences in technological performance. Intuitively, a higher domestic share shr_d or a lower relative wage o_W raises the inferred T_H/T_W, indicating that domestic suppliers are relatively more competitive after accounting for differences in quality and input costs.

The third group governs firms' endogenous innovation behavior. The innovation step

[34] Under the assumption that firm productivities follow a Fréchet distribution, $F(\phi) = \exp[-T\phi^{-\theta}]$, and that prices are proportional to the inverse of productivity ($p \propto 1/\phi$), the distribution of log prices inherits the same dispersion as log productivities. Because the logarithm of a Fréchet-distributed variable is Gumbel-distributed with variance $\pi^2/(6\theta^2)$, the model implies $\text{sd}[\ln p] = \pi/(\sqrt{6}\theta)$.

size (a) determines the magnitude of productivity gains from a successful innovation. In the model, the position at which innovation intensity peaks depends on firms' technological proximity to their foreign competitors, as shown in Proposition 2. While the peak occurs when domestic and foreign suppliers are neck-and-neck—i.e., when their quality-adjusted prices coincide—the exact location of this peak depends on three key parameters: the elasticity of substitution (β), the Fréchet dispersion parameter (θ), and the innovation step size (a).

Having already estimated β and θ from price and quantity data, and having empirically identified the location of the innovation peak using Equation (20), I can use the model-implied relationship in Proposition 2 to back out the implied innovation step size a. Empirically, Equation (20) confirmed the model's prediction of an inverted-U relationship, with innovation intensity peaking when domestic and foreign suppliers' quality-adjusted prices coincide. Consistent with this, the calibration sets the model's implied productivity gain such that it reproduces the observed price ratio between domestic innovators and their modal imported competitors at this point of equality.

Formally, the calibration relies on the model-implied relationship between the observed productivity gain among upgrading firms and the innovation premium:

$$\frac{\tilde{p}_{fpt}^{\text{peak}}}{\tilde{p}_{pt}^{\text{mode, imp}}} = \left(\frac{\theta(\theta - \delta)\ln(1+a)}{(\theta-1)\left[1-(1+a)^{-\theta}\right]} \right)^{1/\theta}, \tag{27}$$

where θ is the Fréchet dispersion parameter estimated above from within-product price dispersion, and $\delta = 1/\beta - 1$ reflects the elasticity of substitution across varieties. Because innovation peaks when the relative quality-adjusted price ratio equals one, I solve this expression at $\tilde{p}_{fpt}^{\text{peak}} / \tilde{p}_{pt}^{\text{mode, imp}} = 1$ to obtain the implied innovation step size a. This ensures that the magnitude of productivity improvement captured by a is consistent with both the model's theoretical prediction and the empirically observed location of the innovation peak.

I identify the fixed cost of innovation by matching the model's implied R&D effort to observed R&D intensity in the data. Starting from the equilibrium research wage condition derived in the theory,

$$w_R = \frac{\kappa\,(\tilde{w}_H^p)^{-\delta}}{\sqrt{2f\,\ell_R}} \left(\int_0^\infty \tilde{\phi}_H^\delta\, S(\tilde{\phi}_H)^2\, f_H(\tilde{\phi}_H)\, d\tilde{\phi}_H \right)^{1/2}, \tag{28}$$

where $\delta = 1/\beta - 1$, I scale both sides by the production wage bill $w_P \ell_P$ and rearrange in terms of observables. Define the empirical R&D wage-bill share

$$s_R \;\equiv\; \frac{w_R \ell_R}{w_H^p \ell_P}.$$

Multiplying (28) by ℓ_R, dividing by $w_H^P \ell_P$, and solving for f yields

$$f \;=\; \frac{1}{2}\left(\frac{\kappa\,(\tilde{w}_H^p)^{-\delta}\;\left[\int_0^\infty \tilde{\phi}_H^\delta\; S(\tilde{\phi}_H)^2\; f_H(\tilde{\phi}_H)\, d\tilde{\phi}_H\right]^{1/2}\;\sqrt{\ell_R}}{s_R\, w_H^p \ell_P}\right)^2. \tag{29}$$

In (29), s_R is measured directly from matched employer–employee data (R&D vs. production occupations distinguished via ISCO-08); ℓ_R and ℓ_P are the corresponding employment counts; w_H^p is the Home wage for production workers implied from the small-open economy equilibrium condition using Equation (9); and the integral term summarizes aggregate innovation. Equation (29) pins down f so that the model reproduces the observed R&D intensity s_R. A larger observed R&D share s_R implies a lower fixed cost f is needed to rationalize the data, since firms are devoting substantial resources to research. Conversely, when the aggregate innovation effort or the constant on profits κ is larger, a higher f is required to match the same s_R; otherwise the model would over-predict R&D effort.

Together, these mappings form a fully determined system that ties the model's equilibrium conditions to observed economic magnitudes. Table 5 summarizes each parameter's interpretation, empirical value, and identifying equation.

The calibrated parameters are chosen such that the model reproduces key empirical moments of the Turkish economy. Most domestic producers operate below the World productivity frontier, with median domestic prices roughly 10% higher than World prices within narrowly defined product categories. By construction, the model matches the observed 43% share of domestic sales relative to imports across manufacturing products (excluding agriculture, food, and beverages, as well as products that are not both imported and produced domestically). It also reproduces the 14.5% share of R&D labor in total employment and the 22.4% share of R&D wages in total wage expenditure. These features ensure that the baseline equilibrium provides a realistic representation of the technological and trade structure of the Turkish economy.

With this benchmark in place, I quantify how trade costs affect aggregate productivity through innovation. Figure 6 shows that the relationship is non-monotonic. Aggregate productivity gains are maximized when trade costs increase by about 9.2%, at which point

Table 5: Calibrated Parameters, Economic Interpretation, and Empirical Sources

Parameter	Description	Value	Estimation or Data Source
(i) Preferences and Market Structure			
α	Expenditure share on homogeneous good	0.33	Calculated from TurkStat national accounts as the share of Turkish consumption in agriculture and mining.
β	Price elasticity of demand $(1/\beta)$	0.21	Estimated using customs data following Eq. (18).
o_W	Relative World wage (World vs. Home)	1.96	Ratio of Turkiye's average wage to that of its trading partners, weighted by bilateral trade shares (WITS, World Bank).
(ii) Technology and Heterogeneity			
θ	Fréchet shape parameter (technology dispersion)	8.81	Estimated from within-product price dispersion across firm-level transactions using Eq. (24).
T_H/T_W	Productivity scale ratio (Home vs. Foreign)	0.295	Inferred from the ratio of domestic to imported market shares using Eq. (26).
k_H	Domestic aggregate quality index	0.332	Median of domestic firm–product quality residuals $\exp(\varepsilon_{f(c)pt})$, Eq. (25).
k_W	World aggregate quality index	3.217	Median of foreign firm–product quality residuals $\exp(\varepsilon_{f(c)pt})$, Eq. (25).
(iii) Innovation and Labor Allocation			
a	Productivity gain from successful innovation	0.128	Estimated using Eq. (27), linking observed productivity gains among upgrading firms to the model-implied innovation premium.
f	Fixed cost of innovation	0.011	Inferred from Eq. (29), equating the model-implied innovation labor share to the observed technical labor share.
ℓ_R	Share of technical labor in Home	0.145	Employer–Employee data (2013); technical workers identified using ISCO-08 occupation codes.
ℓ_P	Share of production labor in Home	0.855	Employer–Employee data (2013); production workers identified using ISCO-08 occupation codes.

Note: Parameters are grouped by their economic role in the model. Group (i) governs household preferences and market structure: α determines the share of expenditure on homogeneous goods, β captures the elasticity of substitution across varieties, and o_W reflects the relative wage between Foreign and Home. Group (ii) characterizes technological heterogeneity: θ determines productivity dispersion, T_H/T_W measures relative efficiency between Home and Foreign, and (k_H, k_W) summarize average quality levels. Group (iii) governs innovation and labor allocation: a and f describe the step size and cost of innovation, while (ℓ_R, ℓ_P) denote the shares of technical and production labor used in Home. All parameters are obtained through exactly identified equilibrium relationships linking the model to Turkish microdata.

Figure 6: Aggregate Productivity Gain from Innovation under Tariffs

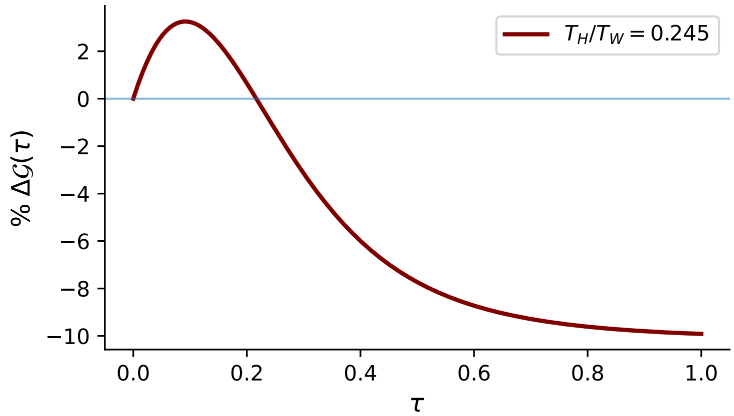

Note: The figure plots the percentage change in the aggregate productivity gain $\mathcal{G}(T_H, T_W, \tau)$ relative to the free-trade baseline ($\tau = 0$), computed from the calibrated model using the parameter values in Table 5. The curve is non-monotonic: moderate tariffs increase aggregate productivity by stimulating innovation among neck-and-neck firms, peaking at $\tau \approx 9.2\%$ with a gain of 3.3%. Higher tariffs reduce competitive pressure and innovation, leading to a decline in aggregate productivity that reaches about -10% under prohibitive protection.

aggregate productivity is 3.3% higher than under free trade. Up to this level, higher trade costs narrow effective cost gaps between domestic and foreign suppliers, bringing a larger mass of domestic firms into the neck-and-neck region where innovation incentives are strongest. These firms invest more in innovation to tip downstream demand in their favor. Beyond this point, however, additional protection shifts the competitive balance too far toward domestic suppliers, reducing the pressure to improve. As trade costs rise by 100%, aggregate productivity falls by roughly 10% relative to the free-trade baseline.

These results highlight that moderate trade barriers can stimulate innovation by intensifying competition among domestic suppliers, but excessive protection weakens that competition and diminishes incentives to upgrade. While openness is essential for sustaining innovation, policy can influence which firms face the most competition—and therefore which firms contribute most to aggregate technological progress.

Having established how trade costs shape innovation and productivity, I next use the model's closed-form welfare expressions to examine how these mechanisms translate into aggregate welfare, distinguishing between the benchmark (static) and innovation (dynamic) channels. Figure 7 summarizes the results. In Panel 7a, the innovation component U_H^{innov} initially rises with τ, because most Home producers lag behind the World frontier. A small increase in trade costs compresses effective cost gaps, brings more domestic firms into the neck-and-neck region, and raises the aggregate probability of successful innovation. As a result, aggregate productivity increases for moderate increases in trade costs,

Figure 7: Percent change in Home welfare components and comparison with no-innovation model

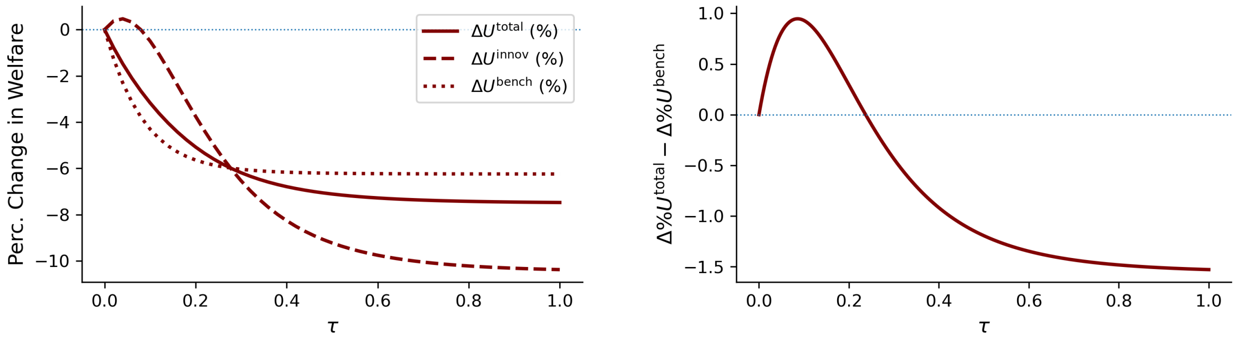

(a) Components: $\%\Delta U_H^{\text{bench}}(\tau)$, $\%\Delta U_H^{\text{innov}}(\tau)$, and $\%\Delta U_H^{\text{total}}(\tau)$

(b) With-innovation vs. no-innovation: $\%\Delta U_H^{\text{total}}(\tau)$

Note: Percent changes are own-baseline normalizations $[U(\tau) - U(0)]/U(0) \times 100$. Calibration given in Table 5. Welfare components are computed from the closed-form expressions in Appendix A.6. Panel 7a plots the benchmark, innovation, and total components; Panel 7b compares total welfare in the innovation economy to a no-innovation counterfactual.

generating a positive dynamic effect on welfare. This gain peaks around $\tau \approx 9.2\%$, after which additional protection widens effective gaps, weakens innovation incentives, and reduces productivity growth.

By contrast, the benchmark (static) component U_H^{bench} declines monotonically with τ, reflecting standard efficiency losses from protection: tariffs raise import prices and shift demand toward less efficient domestic suppliers, while non-switching firms face higher input costs. The total welfare effect combines these two forces. Panel 7b shows that total welfare U_H^{total} falls less steeply than in the no-innovation benchmark up to $\tau \approx 0.24$, reflecting the innovation stimulus for small increases in trade costs. Beyond this threshold, however, the innovation channel reverses and tariffs become more costly than in the static model. Quantitatively, a 10% tariff increase reduces total welfare by 3.1% with innovation, compared to 4.1% without innovation.[35]

These results show that any level of trade protection ultimately reduces welfare, but the magnitude of the loss depends on how innovation responds to competition. Tariffs generate static efficiency losses through misallocation while triggering dynamic responses that can offset part of these losses. In the Turkish calibration, moderate tariffs around 9.2% reduce the welfare loss by roughly one-third relative to a no-innovation

[35]In the model, as $\tau \to \infty$ (autarky), the welfare loss is about -7.5% with innovation and -6% in the benchmark. These magnitudes are in the ballpark of Costinot and Rodríguez-Clare (2014) under perfect competition (-11.9%). Because the homogeneous-good sector here is freely traded, the comparable number scales by its expenditure share: $0.66 \times (-11.9\%) \approx -7.8\%$.

benchmark, as firms closer to the frontier intensify upgrading. However, as protection becomes prohibitive, innovation incentives collapse and welfare losses exceed those in the static economy by about one-fourth. Thus, while trade exposure remains essential for aggregate productivity, its benefits depend on the underlying distribution of firms' technological positions. The key policy insight is that promoting innovation requires sustaining competitive pressure, not shielding domestic producers from it.

5 Conclusion

This paper shows that the welfare effects of protectionism depend not only on how tariffs distort prices and quantities, but also on how they rewire buyer–supplier relationships and thereby reshape suppliers' incentives to innovate. In a small open-economy model, domestic suppliers compete to capture downstream demand from buyers who source each input from a single supplier—either domestic or foreign. The exclusivity of these sourcing relationships means that small changes in relative prices can trigger discrete reallocations of demand, creating strong, state-dependent incentives to innovate. The model formalizes this mechanism and shows that suppliers' innovation follows an inverted-U with respect to their proximity to foreign competitors: firms that are "neck-and-neck" have the most to gain from upgrading, while laggards and leaders have the least.

Empirically, detailed Turkish firm-level data confirm these predictions. When import prices rise, downstream buyers substitute away from foreign suppliers and toward domestic ones, validating the model's microfoundation of contestable buyer–supplier links. Among the domestic suppliers that gain demand, innovation intensifies for those previously lagging behind but weakens for those already ahead, generating the same inverted-U pattern observed in theory. These heterogeneous responses highlight how trade shocks operate not through aggregate exposure alone, but through the micro-level reallocation of sourcing relationships that determine which firms face neck-and-neck competition.

Calibrating the model to Turkish microdata shows that trade protection remains costly even when firms respond endogenously by upgrading technology. A ten-percent increase in trade costs reduces welfare by 3.1 percent—compared with 4.1 percent in a no-innovation benchmark—implying that tariff-induced innovation softens, but does not overturn, the static efficiency losses from protection. Moderate increases in trade costs can compress technological gaps and spur catch-up innovation, yet higher barriers insulate leading suppliers, suppress competitive pressure, and amplify welfare losses. The overall welfare effect thus depends on how domestic suppliers are distributed between laggards and leaders when trade costs change.

More broadly, the findings underscore that protectionism reshapes not only the allocation of production but also the network through which competition and innovation propagate. Buyer–supplier linkages are the transmission channel connecting trade shocks to technological upgrading. Even when innovation responses cushion the welfare decline, tariffs remain costly in aggregate. A natural next step is to extend this framework dynamically—allowing suppliers' innovation to accumulate over time and feed back into the evolution of buyer–supplier networks—to assess how trade policy shapes an economy's long-run innovation trajectory and growth path.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, reallocation, and growth. *American Economic Review*, 108(11):3450–3491.

Aghion, P., Bergeaud, A., Lequien, M., and Melitz, M. J. (2024a). The heterogeneous impact of market size on innovation: Evidence from french firm-level exports. *Review of Economics and Statistics*, 106(3):608–626.

Aghion, P., Bergeaud, A., Lequien, M., Melitz, M. J., et al. (2018). *The impact of exports on innovation: Theory and evidence*. National Bureau of Economic Research Cambridge, MA.

Aghion, P., Bergeaud, A., Lequien, M., Melitz, M. J., and Zuber, T. (2024b). Opposing firm-level responses to the china shock: Output competition versus input supply. *American Economic Journal: Economic Policy*, 16(2):249–269.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and innovation: An inverted-u relationship. *The quarterly journal of economics*, 120(2):701–728.

Aghion, P., Blundell, R., Griffith, R., Howitt, P., and Prantl, S. (2009). The effects of entry on incumbent innovation and productivity. *The review of economics and statistics*, 91(1):20–32.

Akcigit, U., Ates, S. T., and Impullitti, G. (2025). Innovation and trade policy in a globalized world. *Journal of Political Economy Macroeconomics*. Forthcoming.

Akcigit, U. and Melitz, M. (2022). International trade and innovation. In *Handbook of international economics*, volume 5, pages 377–404. Elsevier.

Alvarez, I., Marín, R., and Santos-Arteaga, F. J. (2015). Foreign direct investment entry modes, development and technological spillovers. *The Manchester School*, 83(5):568–603.

Alvarez, R. and López, R. A. (2008). Is exporting a source of productivity spillovers? *Review of world economics*, 144(4):723–749.

Amiti, M., Duprez, C., Konings, J., and Van Reenen, J. (2024). Fdi and superstar spillovers: Evidence from firm-to-firm transactions. *Journal of International Economics*, 152:103972.

Amiti, M. and Khandelwal, A. K. (2013). Import competition and quality upgrading. *Review of Economics and Statistics*, 95(2):476–490.

Amiti, M. and Konings, J. (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American economic review*, 97(5):1611–1638.

Antras, P., Fort, T. C., and Tintelnot, F. (2017). The margins of global sourcing: Theory and evidence from us firms. *American Economic Review*, 107(9):2514–64.

Arkolakis, C., Huneeus, F., and Miyauchi, Y. (2023). Spatial production networks. Technical report, National Bureau of Economic Research.

Autor, D., Dorn, D., Hanson, G. H., Pisano, G., and Shu, P. (2020). Foreign competition and domestic innovation: Evidence from us patents. *American Economic Review: Insights*, 2(3):357–374.

Bacilieri, A., Borsos, A., Astudillo-Estevez, P., and Lafond, F. (2023). Firm-level production networks: what do we (really) know. *INET Oxford Working Paper*, 2023.

Bai, J., Barwick, P. J., Cao, S., and Li, S. (2020). Quid pro quo, knowledge spillover, and industrial quality upgrading: Evidence from the chinese auto industry. Technical report, National Bureau of Economic Research.

Bergeaud, A., Aghion, P., Gigout-Magiorani, T., Lequien, M., and Melitz, M. J. (2024). Exporting ideas: Knowledge flows from expanding trade in goods. Technical report, CEPR Discussion Papers.

Bernard, A. B., Dhyne, E., Magerman, G., Manova, K., and Moxnes, A. (2022). The origins of firm heterogeneity: A production network approach. *Journal of Political Economy*, 130(7):1765–1804.

Bernard, A. B., Jensen, J. B., and Schott, P. K. (2006). Survival of the best fit: Exposure to low-wage countries and the (uneven) growth of us manufacturing plants. *Journal of international Economics*, 68(1):219–237.

Bernard, A. B., Smeets, V., and Warzynski, F. (2017). Rethinking deindustrialization. *Economic Policy*, 32(89):5–38.

Bloom, N., Draca, M., and Van Reenen, J. (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The review of economic studies*, 83(1):87–117.

Buera, F. J. and Oberfield, E. (2020). The global diffusion of ideas. *Econometrica*, 88(1):83–114.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms. *American Economic Review*, 101(1):304–340.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.

Cattaneo, M. D., Crump, R. K., Farrell, M. H., and Feng, Y. (2019). On binscatter. *arXiv preprint arXiv:1902.09608*.

Chen, N. and Juvenal, L. (2022). Markups, quality, and trade costs. *Journal of International Economics*, page 103627.

Chetverikov, D. and Wilhelm, D. (2017). Nonparametric instrumental variable estimation under monotonicity. *Econometrica*, 85(4):1303–1320.

Coelli, F. (2022). Trade policy uncertainty and innovation: Evidence from china. *Available at SSRN 4169514*.

Costinot, A. and Rodríguez-Clare, A. (2014). Trade theory with numbers: Quantifying the consequences of globalization. In *Handbook of international economics*, volume 4, pages 197–261. Elsevier.

Crowley, M. A., Han, L., and Prayer, T. (2024). The pro-competitive effects of trade agreements. *Journal of International Economics*, 150:103936.

Dang, D. A. (2017). The effects of Chinese import penetration on firm innovation: Evidence from the Vietnamese manufacturing sector. Technical report, WIDER Working Paper.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, 114(1):199–251.

Demir, B., Javorcik, B., Michalski, T. K., and Ors, E. (2024). Financial constraints and propagation of shocks in production networks. *Review of Economics and Statistics*, 106(2):437–454.

Dhyne, E., Kikkawa, A. K., Mogstad, M., and Tintelnot, F. (2021). Trade and domestic production networks. *The Review of Economic Studies*, 88(2):643–668.

Dornbusch, R., Fischer, S., and Samuelson, P. A. (1977). Comparative advantage, trade, and payments in a ricardian model with a continuum of goods. *The American Economic Review*, 67(5):823–839.

Feenstra, R. C. and Hanson, G. H. (1996). Globalization, outsourcing, and wage inequality.

Fernandes, A. M. and Paunov, C. (2013). Does trade stimulate product quality upgrading? *Canadian Journal of Economics/Revue canadienne d'économique*, 46(4):1232–1264.

Fieler, A. C. and Harrison, A. E. (2023). Escaping import competition in china. *Journal of International Economics*, 145:103835.

Financial Times (2024). Western democracies drive a global rise in resource nationalism. *Financial Times*. Accessed July 2025.

Flynn, J. P., Levy, A. B., Moscona, J., and Wo, M. (2025). Foreign political risk and technological change. Technical report, National Bureau of Economic Research.

Global Trade Alert (2025). Global Trade Alert: Monitoring Policy Changes That Affect Global Commerce. https://globaltradealert.org/. Data accessed July 2025; independent repository of trade-policy interventions.

Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., and Topalova, P. (2010). Imported intermediate inputs and domestic product growth: Evidence from India. *The Quarterly Journal of Economics*, 125(4):1727–1767.

Grossman, G. M. and Helpman, E. (1991). Trade, knowledge spillovers, and growth. *European economic review*, 35(2-3):517–526.

Hashmi, A. R. (2013). Competition and innovation: The inverted-u relationship revisited. *Review of Economics and Statistics*, 95(5):1653–1668.

Helpman, E. (2024). Foreign competition and innovation. *Journal of International Economics*, 152:104007.

Hombert, J. and Matray, A. (2018). Can innovation help us manufacturing firms escape import competition from china? *The Journal of Finance*, 73(5):2003–2039.

Juhász, R. (2018). Temporary protection and technology adoption: Evidence from the napoleonic blockade. *American Economic Review*, 108(11):3339–3376.

Juhász, R., Lane, N., and Rodrik, D. (2023). The new economics of industrial policy. *Annual Review of Economics*, 16.

Khandelwal, A. (2010). The long and short (of) quality ladders. *The Review of Economic Studies*, 77(4):1450–1476.

Khandelwal, A. K., Schott, P. K., and Wei, S.-J. (2013). Trade liberalization and embedded institutional reform: Evidence from Chinese exporters. *American Economic Review*, 103(6):2169–95.

Kueng, L., Li, N., and Yang, M.-J. (2017). *The Impact of Emerging Market Competition on Innovation and Business Strategy: Evidence from Canada*. Swiss Finance Institute.

Lane, N. (2025). Manufacturing revolutions: Industrial policy and industrialization in south korea. *The Quarterly Journal of Economics*, page qjaf025.

Lileeva, A. and Trefler, D. (2010). Improved access to foreign markets raises plant-level productivity… for some plants. *The Quarterly journal of economics*, 125(3):1051–1099.

Lim, K. et al. (2017). Firm-to-firm trade in sticky production networks. In *2017 Meeting Papers*, volume 280. Society for Economic Dynamics.

Lim, K., Trefler, D., and Yu, M. (2022). Looking backward, innovating forward: A theory of competitive cascades. Technical report, National Bureau of Economic Research.

Magerman, G. and Palazzolo, A. (2024). Deglobalization and the reorganization of supply chains: Effects on regional inequalities in the eu. Technical report, NBB Working Paper.

McNerney, J., Savoie, C., Caravelli, F., Carvalho, V. M., and Farmer, J. D. (2022). How production networks amplify economic growth. *Proceedings of the National Academy of Sciences*, 119(1):e2106031118.

Medina, P. (2017). Import competition, quality upgrading and exporting: Evidence from the Peruvian apparel industry. *University of Toronto mimeo*.

Medina, P. (2024). Import competition, quality upgrading, and exporting: Evidence from the peruvian apparel industry. *Review of Economics and Statistics*, 106(5):1285–1300.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica*, 71(6):1695–1725.

Melitz, M. J. and Redding, S. J. (2021). Trade and innovation. Technical report, National bureau of economic research.

Newey, W. K. and Powell, J. L. (2003). Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578.

Piveteau, P. and Smagghue, G. (2022). Foreign competition along the quality ladder. In *Technical Report, Working Paper*.

Teshima, K. (2008). Import competition and innovation at the plant level: evidence from Mexico. *Unpublished paper, Columbia University*.

Turkish Presidency (2024). 2024 presidential annual program. Technical report, Presidency of Strategy and Budget, Presidency of the Republic of Türkiye. Section 4.2.1.2, p. 121.

World Trade Organization (2023). Trade in services for development: Chapter 2. Technical report, World Trade Organization. Accessed: 2025-09-07.

Zhou, J. (2025). Firewall for innovation. Job Market Paper, Massachusetts Institute of Technology. Available at: https://www.dropbox.com/s/hcdq7qmfq8yfubt/gfw_jmp.pdf?dl=1.

A Theoretical Derivations

A.1 Representative Household

Preferences and budget. The representative household in $Home$ has Cobb–Douglas preferences

$$U_H = Z_H^{\alpha} * \mathcal{D}_H^{(1-\alpha)}, \qquad 0 < \alpha < 1,$$

and faces the budget constraint

$$p_z \, Z_H + P \, \mathcal{D}_H \; = \; I_H,$$

where p_z is the price of the homogeneous good, P is the price of the differentiated final good (I later normalize $P \equiv 1$), and I_H is household income. Per-household income includes labor earnings from production and research, fixed-factor rents, and rebated profits:

$$I_H \; = \; w_H^p \, \ell_H^P \; + \; w^r \, \ell_H^R \; + \; f_H F_H + \Pi_H.$$

Demands. Utility maximization gives per-household demands

$$z_H = \frac{\alpha\, I_H}{p_z}, \qquad \mathcal{D}_c = \frac{(1-\alpha)\, I_c}{P}.$$

Homogeneous Good and Wage Equalization. The homogeneous good is freely traded. With unit labor productivity in Home and $o_W > 0$ in Foreign's homogeneous-good sector, unit costs are

$$p_z^H = w_H^p, \qquad p_z^W = \frac{1}{o_W}\, w_W^p.$$

Free trade and perfect competition imply $p_z^H = p_z^W \equiv p_z$, hence

$$w_H^p = \frac{1}{o_W}\, w_W^p, \qquad p_z = w_H^p = \frac{1}{o_W}\, w_W^p.$$

Because Home is small, w_W^p is pinned by Foreign and is taken as exogenous by Home.

A.2 Derivations for Proposition 1: Peak of innovation effort $u^*(\tilde{\phi}_H)$

Proposition 1 (Peak of u^*)**.** *Let*

$$u^*(\tilde{\phi}_H) = \frac{\kappa\, (\tilde{w}_H^p)^{-\delta}}{f\, w_R}\, S(\tilde{\phi}_H), \qquad S(\tilde{\phi}_H) = (1+a)^{\delta} G_W\big((1+\tau_H)\tilde{o}_W(1+a)\tilde{\phi}_H\big) - G_W\big((1+\tau_H)\tilde{o}_W\tilde{\phi}_H\big),$$

where $G_W(x) = \exp(-T_W x^{-\theta})$ with $T_W > 0$, $\theta > 0$, $a > 0$, and $(1+\tau_H)\tilde{o}_W > 0$. Assume $0 < \delta < \theta$. Then:

1. $\lim_{\tilde{\phi}_H \to 0^+} u^*(\tilde{\phi}_H) = 0.$

2. $\lim_{\tilde{\phi}_H \to +\infty} u^*(\tilde{\phi}_H) = \dfrac{\kappa(\tilde{w}_H^p)^{-\delta}}{f\, w_R}\big[(1+a)^{\delta} - 1\big].$

3. *u^* has a unique interior maximizer at*

$$\tilde{\phi}_H^* = \frac{1}{(1+\tau_H)\tilde{o}_W} \left(\frac{T_W\,[1-(1+a)^{-\theta}]}{(\theta-\delta)\ln(1+a)} \right)^{1/\theta},$$

and $u^{*\prime\prime}(\tilde{\phi}_H^*) < 0$.

Proof. Set $C \equiv \kappa(\tilde{w}_H^p)^{-\delta}/(f\, w_R) > 0$ and $b \equiv (1+\tau_H)\tilde{o}_W > 0$, and define

$$t(\tilde{\phi}_H) \equiv T_W\, b^{-\theta}\, \tilde{\phi}_H^{-\theta} > 0.$$

Using $G_W(x) = \exp(-T_W x^{-\theta})$, we can rewrite

$$u^*(\tilde{\phi}_H) = C\left[(1+a)^{\delta}\, e^{-(1+a)^{-\theta} t(\tilde{\phi}_H)} - e^{-t(\tilde{\phi}_H)}\right].$$

Limits. As $\tilde{\phi}_H \to 0^+$, $t(\tilde{\phi}_H) \to +\infty$ and both exponentials vanish, yielding $\lim_{\tilde{\phi}_H \to 0^+} u^*(\tilde{\phi}_H) = 0$. As $\tilde{\phi}_H \to +\infty$, $t(\tilde{\phi}_H) \to 0^+$ and both exponentials converge to 1, so $u^*(\tilde{\phi}_H) \to C\big[(1+a)^{\delta} - 1\big]$, establishing (i)–(ii).

Interior critical point. Differentiate using $t'(\tilde{\phi}_H) = -\theta\, T_W b^{-\theta} \tilde{\phi}_H^{-\theta-1} < 0$:

$$u^{*\prime}(\tilde{\phi}_H) = C\,\theta\, T_W b^{-\theta}\, \tilde{\phi}_H^{-\theta-1}\left[(1+a)^{\delta-\theta}\, e^{-(1+a)^{-\theta} t(\tilde{\phi}_H)} - e^{-t(\tilde{\phi}_H)}\right].$$

Since the front factor is positive, $u^{*\prime}(\tilde{\phi}_H) = 0$ iff

$$(1+a)^{\delta-\theta}\, e^{-(1+a)^{-\theta} t} = e^{-t} \quad\iff\quad (\delta-\theta)\ln(1+a) = -t\left[1-(1+a)^{-\theta}\right].$$

With $t = T_W b^{-\theta} \tilde{\phi}_H^{-\theta}$, this gives

$$T_W b^{-\theta} \tilde{\phi}_H^{-\theta} = \frac{(\theta-\delta)\ln(1+a)}{1-(1+a)^{-\theta}} > 0,$$

which (since $0 < \delta < \theta$) has the unique solution

$$\tilde{\phi}_H^* = b^{-1}\left(\frac{T_W\left[1-(1+a)^{-\theta}\right]}{(\theta-\delta)\ln(1+a)}\right)^{1/\theta}.$$

Second-order condition and uniqueness. Let $g(t) \equiv (1+a)^{\delta} e^{-(1+a)^{-\theta} t} - e^{-t}$ so that $u^*(\tilde{\phi}_H) = C\, g(t(\tilde{\phi}_H))$. We have

$$g'(t) = -(1+a)^{\delta-\theta} e^{-(1+a)^{-\theta} t} + e^{-t}, \qquad g''(t) = (1+a)^{\delta-2\theta} e^{-(1+a)^{-\theta} t} - e^{-t}.$$

At the critical point t^*, the first-order condition implies $e^{-t^*} = (1+a)^{\delta-\theta} e^{-(1+a)^{-\theta} t^*}$, hence

$$g''(t^*) = e^{-(1+a)^{-\theta} t^*}\left[(1+a)^{\delta-2\theta} - (1+a)^{\delta-\theta}\right] = -(1+a)^{\delta-2\theta}\big[(1+a)^{\theta} - 1\big]\, e^{-(1+a)^{-\theta} t^*} < 0.$$

Because $dt/d\tilde{\phi}_H < 0$, the chain rule yields $u^{*\prime\prime}(\tilde{\phi}_H^*) < 0$. Moreover, the equation for t above is linear in $\tilde{\phi}_H^{-\theta}$ and thus has a unique positive solution, so the maximizer is unique. □

We can also compare $\tilde{\phi}_H^*$ to the tariff–adjusted foreign modal productivity,

$$\phi_{W,\text{mode}} = \left(\tfrac{\theta}{\theta+1} T_W\right)^{1/\theta}, \qquad \bar{\phi}_{W,\text{mode}} \equiv \frac{\phi_{W,\text{mode}}}{(1+\tau_H)\tilde{o}_W}.$$

Then

$$\frac{\tilde{\phi}_H^*}{\bar{\phi}_{W,\text{mode}}} = \left[\frac{\frac{\theta+1}{\theta}\left(1-(1+a)^{-\theta}\right)}{(\theta-\delta)\,\ln(1+a)}\right]^{1/\theta},$$

which depends only on (a, δ, θ).

For small innovation steps $a \to 0^+$,

$$\frac{\tilde{\phi}_H^*}{\bar{\phi}_{W,\text{mode}}} \;\longrightarrow\; \left(\frac{\theta+1}{\theta-\delta}\right)^{1/\theta}.$$

A larger δ (equivalently, a higher elasticity of substitution) shifts the peak toward higher productivity. The intuition is that when inputs are close substitutes, the gains from an upgrade are not only about tipping the head-to-head contest against a Foreign rival but also about stealing demand from other varieties. This intensifies competition and implies that only firms already close to the frontier find innovation worthwhile. Conversely, when inputs are less substitutable (smaller δ), even firms further behind can benefit, and the peak occurs at lower productivity.

As a rises, the ratio $\tilde{\phi}_H^*/\bar{\phi}_{W,\text{mode}}$ falls, shifting the peak toward lower $\tilde{\phi}_H$: large innovation steps allow laggards to catch up more easily, so the firms with the strongest incentives are those farther from the frontier. Finally, as $\theta \to \infty$ the ratio tends to 1. When the foreign productivity distribution is very tight, competition comes overwhelmingly from firms near the mode, and the peak of Home innovation aligns exactly with the tariff–adjusted foreign modal productivity.

A.3 Share of Wage Bill on Tech Workers Inverted-U

Neither $u^*(\phi_H)$ nor ϕ_H is observed, so Proposition 1 cannot be taken directly to the data. I therefore use two observables. First, under the quadratic innovation cost, a firm's spending on technical labor is proportional to $u(\phi_H)^2$, and thus provides a monotone proxy for innovation intensity. Define the technical wage share

$$s^{\text{tech}}(\phi_H) \;\equiv\; \frac{w_R\,\ell_R(\phi_H)}{w_H\,\ell_P(\phi_H)},$$

which inherits the inverted-U prediction. Second, while productivities are unobserved, the model implies a one-to-one mapping between productivity and a *quality-adjusted price*. Define

$$\widehat{p}_H(\phi_H) \;\equiv\; p_H(\phi_H)\, k_H^{\frac{1}{1-\beta}}, \qquad \widehat{p}_W^{\,\text{mode}} \;\equiv\; (1+\tau)\, p_W^{\text{mode}}\, k_W^{\frac{1}{1-\beta}},$$

so that $\widehat{p}_H \propto \phi_H^{-1}$ and $\widehat{p}_W^{\,\text{mode}} \propto (\phi_W^{\text{mode}})^{-1}$. Our testable object is the quality-adjusted price gap

$$\widehat{D} \;\equiv\; \frac{\widehat{p}_H}{\widehat{p}_W^{\,\text{mode}}}, \qquad \widehat{v} \;\equiv\; \ln \widehat{D}.$$

Proposition 2 (Peak of s^{tech} in quality-adjusted price gap)**.** *Let $G(x) = \exp(-T_W x^{-\theta})$, $b \equiv \tilde{o}_W(1+\tau)$, and define the quality-adjusted prices as above, with ratio $\widehat{D} \equiv \widehat{p}_H/\widehat{p}_W^{\,\text{mode}}$. For $\theta > 1$ the mapping satisfies*

$$T_W\, b^{-\theta}\, \phi_H^{-\theta} = \frac{\theta-1}{\theta}\, \widehat{D}^{\theta}.$$

Consider

$$s^{\text{tech}}(\phi_H) = S_0\Big[(1+a)^{\delta} G\big(b(1+a)\phi_H\big) - G\big(b\phi_H\big)\Big]^2,$$

where $S_0 \equiv \frac{\eta\, w_H^{-\delta} k_H}{2 f w_R}\, \beta^2(1-\beta)^{1-2\beta/\beta}$. Then s^{tech} as a function of $\widehat{D}$ is

$$s^{\text{tech}}(\widehat{D}) = S_0\left[(1+a)^{\delta}\exp\left(-\frac{\theta-1}{\theta}\frac{\widehat{D}^{\theta}}{(1+a)^{\theta}}\right) - \exp\left(-\frac{\theta-1}{\theta}\widehat{D}^{\theta}\right)\right]^2.$$

If $0 < \delta < \theta$ (and $\theta > 1$), $s^{\text{tech}}(\widehat{D})$ is single-peaked with a unique interior maximizer $D^* > 0$ given by

$$D^{*\,\theta} = \frac{\theta(\theta-\delta)\,\ln(1+a)}{(\theta-1)\big(1-(1+a)^{-\theta}\big)}, \qquad D^* = \left[\frac{\theta(\theta-\delta)\,\ln(1+a)}{(\theta-1)\big(1-(1+a)^{-\theta}\big)}\right]^{1/\theta}.$$

Proposition 2 shows that the technical wage share is inverted-U in the observable price gap D: innovation effort rises as domestic and foreign prices converge, peaks at D^*, and falls again when the gap widens. In the small-step limit $a \to 0^+$,

$$D^* \;\to\; \left(\frac{\theta-\delta}{\theta-1}\right)^{1/\theta}.$$

Thus, when $\delta = 1$ (i.e., $\beta = 0.5$) and innovation steps are small, the peak occurs near $D^* = 1$, i.e. when the domestic quality-adjusted price is at parity with the foreign modal price.[36] Larger innovation steps a push the peak slightly above one—firms just behind

[36] In practice, estimates of β are typically closer to 0.2 (corresponding to an elasticity of substitution of 5).

innovate most—while tighter foreign dispersion (larger θ) pulls the peak back toward parity.

Economically, this structure reflects how innovation incentives hinge on competitive position. Firms far behind the frontier innovate little, as even a successful upgrade will not make them competitive. Firms far ahead also innovate little, since they already dominate and gain little market access from further upgrades. The strongest incentives arise for mid-range firms: when a marginal innovation flips the sourcing decision in their favor, it transforms zero profits into positive profits. This competitive margin generates the inverted-U pattern in innovation effort, and provides a simple empirical benchmark: innovation intensity should peak when domestic and foreign prices are closest.

A.4 Derivations for Production Wages

Foreign production wage. Foreign production labor is employed in the homogeneous good and in differentiated intermediates. Producing Z_W units of the homogeneous good requires Z_W/o_W labor, with $Z_W = \alpha I_W/p_z$ and $p_z = w_W^p/o_W$. Hence

$$L_z^W = \frac{Z_W}{o_W} = \frac{\alpha I_W}{w_W^p}.$$

For intermediates, CES demand $q = p^{-1/\beta}$ and markup pricing $p^*(\phi) = \frac{1}{1-\beta}\frac{w_W^p}{\phi}$ imply that a Foreign variety of productivity ϕ uses

$$\ell_W(\phi) = \frac{q(\phi)}{\phi} = (1-\beta)^{1/\beta}\,(w_W^p)^{-1/\beta}\,\phi^{\delta}, \qquad \delta = \tfrac{1}{\beta} - 1 \in (0,\theta).$$

Aggregating over the Fréchet density $g_W(\phi)$ gives

$$L_p^W(w_W^p) = \int_0^\infty \ell_W(\phi)\,g_W(\phi)\,d\phi = (1-\beta)^{1/\beta}\,(w_W^p)^{-1/\beta}\,M_W, \qquad M_W := T_W^{\delta/\theta}\,\Gamma\Big(1-\tfrac{\delta}{\theta}\Big).$$

The Foreign labor market clearing condition is therefore

$$1 \;=\; L_z^W + L_p^W \;=\; \frac{\alpha I_W}{w_W^p} \;+\; (1-\beta)^{1/\beta}\,(w_W^p)^{-1/\beta}\,M_W. \tag{30}$$

Using the Foreign income identity

$$I_W = w_W^p\Big(\tfrac{Z_W}{o_W} + L_p^W\Big) + f_W + \Pi_W^F + \Pi_W^H,$$

I use $\beta = 0.5$ here purely for expositional clarity, since it simplifies the algebra in this example.

with $f_W = \beta Y_W$ and $\Pi_W^H = o(\eta)$ in the SOE limit, we obtain

$$(1-\alpha) I_W = (1-\beta)^{\frac{1}{\beta}-2} \, (w_W^p)^{-\delta} \, M_W.$$

Substituting back into (30) yields

$$1 = \frac{\alpha}{1-\alpha}(1-\beta)^{\frac{1}{\beta}-2} \, M_W \, (w_W^p)^{-1/\beta} + (1-\beta)^{1/\beta} \, M_W \, (w_W^p)^{-1/\beta},$$

which implies

$$\left(w_W^p\right)^{1/\beta} = M_W \, (1-\beta)^{\frac{1}{\beta}-2} \, \left((1-\beta)^2 + \tfrac{\alpha}{1-\alpha}\right). \tag{31}$$

Raising both sides to the power β yields the closed form in (8).

Home production wage. Since the homogeneous good is freely traded, its price must equalize:

$$p_z^H = w_H^p, \qquad p_z^F = \tfrac{1}{o_W} w_W^p, \qquad p_z^H = p_z^F.$$

Hence

$$w_H^p = \frac{1}{o_W} \, w_W^p. \tag{32}$$

A.5 Gain in Aggregate Productivity From Innovation

I define the aggregate productivity gain from innovation as the change in the change in aggregate productivities of active firms:

$$\begin{aligned}
\mathcal{G}(T_H, T_W, \tau) &= \int_0^\infty \phi_H^\delta \, G_W\big((1+\tau)\tilde{o}_W \, \phi_H\big) \, g_H(\phi_H) \, d\phi_H \; - \int_0^\infty \phi_H^\delta \, G_W\big((1+\tau)\tilde{o}_W \, \phi_H\big) \, f_H(\phi_H) \, d\phi_H \\
&= \int_0^\infty \phi_H^\delta \, G_W\big((1+\tau)\tilde{o}_W \, \phi_H\big) \, \Big[\tfrac{1}{1+a} u^*(\tfrac{\phi_H}{1+a}) f_H(\tfrac{\phi_H}{1+a}) - u^*(\phi_H) f_H(\phi_H)\Big] \, d\phi_H \\
&= \left(\frac{2L_H^R}{f}\right)^{1/2} \left(\int_0^\infty \phi_H^\delta \, S(\phi_H, \tau)^2 \, f_H(\phi_H) \, d\phi_H\right)^{1/2},
\end{aligned}$$

where $S(\phi, \tau) = (1+a)^\delta \exp\big(-T_W(1+\tau)^{-\theta}\tilde{o}_W^{-\theta}(1+a)^{-\theta}\phi^{-\theta}\big) - \exp\big(-T_W(1+\tau)^{-\theta}\tilde{o}_W^{-\theta}\phi^{-\theta}\big)$. With $s \equiv \delta/\theta \in (0,1)$, $\Gamma \equiv \Gamma(1-s)$, and the Fréchet algebra,

$$\begin{aligned}
&\int_0^\infty \phi_H^\delta S(\phi_H, \tau)^2 f_H(\phi_H) \, d\phi_H \\
&= T_H \, \Gamma \Big[(1+a)^{2\delta}\big(T_H + 2b\big)^{s-1} - 2(1+a)^\delta\big(T_H + b + c\big)^{s-1} + \big(T_H + 2c\big)^{s-1}\Big],
\end{aligned}$$

where we have defined the *effective* foreign scale

$$\widetilde{T}_W \equiv \frac{T_W}{\big((1+\tau)\tilde{o}_W\big)^{\theta}}, \qquad b \equiv \widetilde{T}_W(1+a)^{-\theta}, \qquad c \equiv \widetilde{T}_W.$$

Hence the closed form used below:

$$\mathcal{G}(T_H, T_W, \tau) = \left(\frac{2L_H^R}{f}\right)^{1/2} \left\{ T_H\, \Gamma\left[\frac{(1+a)^{2\delta}}{(T_H+2b)^{1-s}} - \frac{2(1+a)^{\delta}}{(T_H+b+c)^{1-s}} + \frac{1}{(T_H+2c)^{1-s}} \right] \right\}^{1/2}.$$

Sign of $\partial\mathcal{G}/\partial\widetilde{T}_W$ at the extremes. Write the inner bracket in Proposition 3 as

$$F(\widetilde{T}_W) = K_1\,(T_H + 2\alpha\widetilde{T}_W)^{-(1-s)} - K_2\,(T_H + (\alpha+1)\widetilde{T}_W)^{-(1-s)} + K_3\,(T_H + 2\widetilde{T}_W)^{-(1-s)},$$

with

$$s = \frac{\delta}{\theta} \in (0,1), \quad \alpha = (1+a)^{-\theta} \in (0,1), \quad K_1 = (1+a)^{2\delta},\ K_2 = 2(1+a)^{\delta},\ K_3 = 1.$$

All multiplicative prefactors outside $F(\cdot)$ in $\mathcal{G}$ are strictly positive, so $\operatorname{sign} \partial\mathcal{G}/\partial\widetilde{T}_W = \operatorname{sign} F'(\widetilde{T}_W)$.

(i) Small $\widetilde{T}_W$. Differentiate and evaluate at 0:

$$F'(\widetilde{T}_W) = -(1-s)\Big[K_1{\cdot}2\alpha\,(T_H+2\alpha\widetilde{T}_W)^{-(2-s)} - K_2{\cdot}(\alpha+1)\,(T_H+(\alpha+1)\widetilde{T}_W)^{-(2-s)} + K_3{\cdot}2\,(T_H+2\widetilde{T}_W)^{-(2-s)}\Big].$$

At $\widetilde{T}_W = 0$ this reduces to

$$F'(0) = -(1-s)T_H^{-(2-s)} \underbrace{\big[2\alpha K_1 - (\alpha+1)K_2 + 2K_3\big]}_{B}.$$

Using $X := (1+a)^{\delta} > 1$, we can factor the bracket as

$$B = 2\big(\alpha X^2 - X(\alpha+1) + 1\big) = 2(X-1)(\alpha X - 1).$$

Because $\alpha X = (1+a)^{\delta-\theta} = (1+a)^{-(\theta-\delta)} < 1$ (since $\theta > \delta$), we have $(X-1) > 0$ and $(\alpha X - 1) < 0$, hence $B < 0$. Therefore $F'(0) > 0$, so $\partial\mathcal{G}/\partial\widetilde{T}_W > 0$ for all sufficiently small $\widetilde{T}_W$.

(ii) Large $\widetilde{T}_W$. For $\widetilde{T}_W \to \infty$,

$$F(\widetilde{T}_W) = \widetilde{T}_W^{-(1-s)}\Big[K_1(2\alpha)^{-(1-s)} - K_2(\alpha+1)^{-(1-s)} + K_3 2^{-(1-s)}\Big] + o\big(\widetilde{T}_W^{-(1-s)}\big).$$

Let $p := 1 - s \in (0,1)$ and define $f(t) := t^{-p}$ (decreasing and convex on $\mathbb{R}_+$). Set $d_1 = 2\alpha$, $d_2 = \alpha + 1 = \frac{d_1+d_3}{2}$, $d_3 = 2$ and note

$$\Sigma \;:=\; K_1 f(d_1) - K_2 f(d_2) + K_3 f(d_3) \;=\; f(d_1)X^2 - 2f(d_2)X + f(d_3), \quad X = (1+a)^{\delta}.$$

This is a quadratic in X with positive leading coefficient $f(d_1) > 0$. Its minimum equals

$$\min_X \Sigma \;=\; f(d_3) - \frac{f(d_2)^2}{f(d_1)}.$$

By convexity,

$$f(d_2) = f\Big(\tfrac{d_1+d_3}{2}\Big) \;\le\; \frac{f(d_1)+f(d_3)}{2} \;\le\; \sqrt{f(d_1)f(d_3)},$$

with strict inequalities since $d_1 \neq d_3$ when $a > 0$ (i.e., $\alpha < 1$). Hence $f(d_2)^2 < f(d_1)f(d_3)$ and therefore $\min_X \Sigma > 0$, implying $\Sigma > 0$ for all $X \ge 1$. It follows that

$$F'(\widetilde{T}_W) = -(1-s)\,\Sigma\,\widetilde{T}_W^{-(2-s)} + o\big(\widetilde{T}_W^{-(2-s)}\big) \;<\; 0 \quad \text{for sufficiently large } \widetilde{T}_W.$$

Combining (i) and (ii), $\partial\mathcal{G}/\partial\widetilde{T}_W > 0$ near 0 and $\partial\mathcal{G}/\partial\widetilde{T}_W < 0$ for large $\widetilde{T}_W$, proving the hump shape in $\widetilde{T}_W$.

A.6 Welfare

Welfare is

$$U_H \;=\; \alpha\log\alpha \;+\; (1-\alpha)\log(1-\alpha) \;+\; \log i_H(T_H, T_W, \tau) \;-\; \alpha\log p_z(T_W), \tag{33}$$

where the first two terms come from Cobb–Douglas preferences over the homogeneous good z_H and differentiated composite D_H, and $\log i_H - \alpha \log p_z$ is log real income.

Per household income is:

$$i_H(T_H, T_W, \tau) \;=\; w_H^p(T_W)\,\ell_H^P \;+\; w_H^r(T_H, T_W, \tau)\,\ell_H^R \;+\; \frac{f_H(T_H, T_W, \tau)\,F_H \;+\; \Pi_H(T_H, T_W, \tau)}{L_H}, \tag{34}$$

with $\ell_H^P \equiv L_H^P/L_H$ and $\ell_H^R \equiv L_H^R/L_H$ exogenous per-capita labor allocations.

The components entering welfare are given as:

$$w_H^p(T_W) = \frac{1}{o_W}\Big[T_W^s\,\Gamma(1-s)\,(1-\beta)^2\Big((1-\beta)^2 + \frac{\alpha}{1-\alpha}\Big)\Big]^{\beta}, \tag{35}$$

$$p_z(T_W) = w_H^p(T_W), \tag{36}$$

$$w_H^r(T_H, T_W, \tau) = \beta\,(1-\beta)^{\delta-1}\,\big(w_H^p(T_W)\big)^{-\delta}\left(\frac{\eta}{2f\,\ell_H^R}\,I_0(\tau)\right)^{1/2}, \tag{37}$$

$$\frac{f_H(T_H, T_W, \tau)\,F_H}{L_H} = \beta\,(1-\beta)^{\delta-1}\Big[\,I_1(\tau) + (1+\tau)^{-\delta} I_2(\tau)\Big], \tag{38}$$

$$\frac{\Pi_H(T_H, T_W, \tau)}{L_H} = \beta\,(1-\beta)^{\delta}\,\big(w_H^p(T_W)\big)^{-\delta}\,I_1(\tau). \tag{39}$$

where

$$D(\phi, T_W, \tau) \;=\; (1+a)^{\delta}\,G_W\big(\phi(1+\tau)(1+a)\big) \;-\; G_W\big(\phi(1+a)\big) \tag{40}$$

$$I_0(T_H, T_W, \tau) = \int_0^{\infty} \phi^{\delta}\,D(\phi, T_W, \tau)^2\,f_H(\phi)\,d\phi, \tag{41}$$

$$I_1(T_H, T_W, \tau) = \int_0^{\infty} \phi^{\delta}\,G_W\big(\phi(1+\tau)o_H\big)\,g_H(\phi)\,d\phi, \tag{42}$$

$$I_2(T_H, T_W, \tau) = \int_0^{\infty} \phi^{\delta}\,G_H\Big(\tfrac{\phi}{(1+\tau)o_H}\Big)\,g_W(\phi)\,d\phi, \tag{43}$$

The updated CDF and the PDF of the Foreign productivity (after exogenous upgrading) is taken as:

$$G_W(\phi) = e^{-T_W\phi^{-\theta}}, \qquad g_W(\phi) = T_W\theta\,\phi^{-(\theta+1)}e^{-T_W\phi^{-\theta}}.$$

Whereas, the initial CDF and PDF of the productivities at Home are taken as:

$$F_H(\phi) = e^{-T_H\phi^{-\theta}}, \qquad f_H(\phi) = T_H\theta\,\phi^{-(\theta+1)}e^{-T_H\phi^{-\theta}}.$$

Home (post-innovation): letting $u(\cdot)$ capture the R&D-induced upward shift in mass,

$$g_H(\phi) = f_H(\phi) \;+\; \frac{1}{1+a}\,u\Big(\frac{\phi}{1+a}\Big)\,f_H\Big(\frac{\phi}{1+a}\Big) \;-\; u(\phi)f_H(\phi),$$

$$G_H(\phi) = F_H(\phi) \;+\; \int_0^{\phi}\Big[\frac{1}{1+a}\,u\Big(\frac{x}{1+a}\Big)\,f_H\Big(\frac{x}{1+a}\Big) \;-\; u(x)f_H(x)\Big]dx.$$

T_W directly shifts w_H^p and p_z (world prices) and, via G_W, g_W, propagates to w_H^r, rents, and profits through $I_k(\cdot)$. T_H enters w_H^r through $I_0(\cdot)$ (via f_H) and affects rents/profits via

g_H, G_H in I_1, I_2. The tariff τ loads through $D(\phi,\tau)$ in I_0 and through $(1+\tau)o_W$ in I_1, I_2.

A.7 Decomposing Welfare into Benchmark and Innovation Components

In this subsection, I decompose welfare in Home into two main components: welfare from baseline and welfare gain from firms innovating:

$$\begin{aligned} U_H &= \alpha^\alpha(1-\alpha)^{1-\alpha}\Big[U_H^{\text{bench}}(T_H,T_W,\tau) + U_H^{\text{innov}}(T_H,T_W,\tau)\Big] \\ &= \alpha^\alpha(1-\alpha)^{1-\alpha}\Big[\underbrace{(w_H^p)^{-\alpha}\,(l_P w_H^p + f_H^n + \pi_H^n)}_{\text{Benchmark}} + \underbrace{(w_H^p)^{-\alpha}\,\big(l_R w_R + f_H^i + \pi_H^i\big)}_{\text{Gain from Innovation}}\Big] \end{aligned} \tag{44}$$

where the benchmark uses (f_H, F_H) and the innovation part captures the $u(\cdot)$-driven adjustments (g_H, G_H).

The research wage, w_H^r depends on I_0 computed with the baseline density f_H in (41) and is only present if innovation occurs; so it exists in the innovation component. For rents and profits,

$$f_H F_H = \underbrace{\beta(1-\beta)^{\delta-1}(\tilde{w}_H^p)^{-\delta}\Big[I_1^n + (1+\tau)^{-\delta} I_2^n\Big]}_{f_H^n} + \underbrace{\beta(1-\beta)^{\delta-1}(\tilde{w}_H^p)^{-\delta}\Big[I_1^i + (1+\tau)^{-\delta} I_2^i\Big]}_{f_H^i},$$

$$\Pi_H = \underbrace{\beta(1-\beta)^{\delta}\big(\tilde{w}_H^p\big)^{-\delta} I_1^n}_{\pi_H^n} + \underbrace{\beta(1-\beta)^{\delta}\big(\tilde{w}_H^p\big)^{-\delta} I_1^i}_{\pi_H^i},$$

with benchmark (no-innovation) terms

$$I_1^n = \int_0^\infty \phi^\delta G_W\big(\phi(1+\tau)\tilde{o}_W\big) f_H(\phi)\, d\phi, \quad I_2^n = \int_0^\infty \phi^\delta F_H\Big(\tfrac{\phi}{(1+\tau)\tilde{o}_W}\Big) g_W(\phi)\, d\phi,$$

and innovation components

$$\begin{aligned} I_1^i &= \int_0^\infty \phi^\delta G_W\big(\phi(1+\tau)o_H\big)\Big[\tfrac{1}{1+a} u\Big(\tfrac{\phi}{1+a}\Big) f_H\Big(\tfrac{\phi}{1+a}\Big) - u(\phi) f_H(\phi)\Big]\, d\phi, \\ I_2^i &= \int_0^\infty \phi^\delta \left(\int_0^{\phi/((1+\tau)o_H)} \Big[\tfrac{1}{1+a} u\Big(\tfrac{x}{1+a}\Big) f_H\Big(\tfrac{x}{1+a}\Big) - u(x) f_H(x)\Big] dx\right) g_W(\phi)\, d\phi. \end{aligned}$$

Thus U_H^{bench} reflects static trade and initial productivities, while U_H^{innov} adds the R&D-driven mass shifts toward higher ϕ through $u(\cdot)$, operating via I_0, I_1^i, and I_2^i.

To understand how tariffs affect Home welfare, I derive an expression for benchmark utility under the assumption that the quality adjusted wages in both countries are equal to each other, i.e., $\tilde{o}_H = 1$. Home welfare is given by:

$$
\begin{aligned}
U_H^{bench}(T_H, T_W, \tau) &= (w_H^p)^{-\alpha} \left(\ell_P w_H^p + f_H^n + \pi_H^n\right) \\
&= (w_H^p)^{-\alpha} \left(\ell_P w_H^p + (w_H^p)^{-\delta} \beta (1-\beta)^{\delta-1} \left[(2-\beta) I_1^n + (1+\tau)^{-\delta} I_2^n\right]\right),
\end{aligned} \tag{45}
$$

where $\delta = \frac{1-\beta}{\beta}$, and the terms I_1^n and I_2^n capture aggregate innovative output targeted toward domestic and foreign markets, respectively.

Using the closed-form expressions:

$$
I_1^n = \Gamma T_H \left(T_H + T_W((1+\tau)\tilde{o}_W)^{-\theta}\right)^{s-1}, \quad I_2^n = \Gamma T_W \left(T_W + T_H((1+\tau)\tilde{o}_W)^{\theta}\right)^{s-1},
$$

with $s = \delta/\theta \in (0,1)$ and $\Gamma = \Gamma(1-s) > 0$, equation (45) simplifies to:

$$
\begin{aligned}
U_H^{bench}(T_H, T_W, \tau) = (w_H^p)^{-\alpha} &\left(\ell_P w_H^p \right. \\
&+ (\tilde{w}_H^p)^{-\delta} \beta (1-\beta)^{\delta-1} \Gamma \left(T_H + T_W((1+\tau)\tilde{o}_W)^{-\theta}\right)^{s-1} \left[(2-\beta) T_H + T_W((1+\tau)\tilde{o}_W)^{-\theta}\right].
\end{aligned}
$$

(46)

This expression makes it transparent that the only channel through which τ affects welfare is via the effective contribution of Foreign technology, scaled by the iceberg trade cost $(1+\tau)^{-\theta}$. All other terms—such as $w_H^p = w_H^p(T_W)$—are independent of τ under the small open economy assumption.

To assess the effect of a marginal increase in the tariff τ, define $x = (1+\tau)^{-\theta}$, so that $x \in (0,1]$ and $\frac{dx}{d\tau} = -\theta(1+\tau)^{-\theta-1} < 0$. Then we can write the welfare-relevant term as:

$$
F(x) = (T_H + T_W x)^{s-1} \left[(2-\beta) T_H + T_W x\right].
$$

We aim to show that $F(x)$ is increasing in x, so that a higher τ (which lowers x) reduces welfare.

Differentiate:

$$
\begin{aligned}
\frac{dF}{dx} &= (s-1)(T_H + T_W x)^{s-2} T_W \left[(2-\beta) T_H + T_W x\right] + (T_H + T_W x)^{s-1} T_W \\
&= T_W (T_H + T_W x)^{s-2} \left[(s-1)\left((2-\beta) T_H + T_W x\right) + (T_H + T_W x)\right] \\
&= T_W (T_H + T_W x)^{s-2} \left[c(\beta, \theta) T_H + s T_W x\right],
\end{aligned}
$$

where $c(\beta, \theta) = 1 - (2 - \beta)(1 - s)$. This coefficient governs whether Foreign technology amplifies or dampens welfare as τ rises.

The derivative $\frac{dF}{dx} > 0$ for all $T_H, T_W > 0$ if and only if:

$$c(\beta, \theta) = 1 - (2 - \beta)(1 - s) \geq 0.$$

Solving this gives:

$$s \geq \frac{1 - \beta}{2 - \beta} \quad \iff \quad \theta \leq \frac{2 - \beta}{\beta}.$$

Combined with $s = \delta/\theta < 1$, we obtain the full condition:

$$\frac{1 - \beta}{\beta} < \theta < \frac{2 - \beta}{\beta}.$$

Under this empirically relevant range, benchmark welfare is strictly decreasing in τ for all $T_H, T_W > 0$.

The negative effect of tariffs arises because Foreign technology enhances the profitability and scalability of Home R&D. Tariffs reduce the effective spillover from Foreign technology by dampening its contribution to both the innovation scale and the output expansion terms. When θ lies in the intermediate range, these effects reinforce each other. Hence, tariffs reduce welfare not only by raising consumer prices, but also by stifling productivity growth through weaker global linkages.

The sensitivity of benchmark welfare to the tariff also depends on the relative size of T_W. For any positive T_H, the derivative $\frac{\partial B}{\partial \tau}$ approaches zero as $T_W \to 0$, and tends to negative infinity as $T_W \to \infty$. This is consistent with intuition. When Foreign technology is negligible ($T_W \to 0$), Home does not import intermediates, so tariff changes have little to no effect on welfare. By contrast, when Foreign technology dominates ($T_W \to \infty$), Home sources nearly all differentiated intermediates from abroad and produces only the homogeneous good. In this case, the economy becomes highly exposed to tariff distortions, and the welfare loss from increased trade costs becomes arbitrarily large.

A.8 Derivation of Welfare from Innovation

I now solve for the component of welfare from innovation:

Let $s \equiv \delta/\theta \in (0, 1)$ and $\Gamma \equiv \Gamma(1 - s) > 0$. Define

$$A \;=\; T_H + 2T_W\big((1 + a)(1 + \tau)\tilde{o}_W\big)^{-\theta}, \quad B \;=\; T_H + T_W((1 + \tau)\tilde{o}_W)^{-\theta}\big(1 + (1 + a)^{-\theta}\big),$$

$$C \;=\; T_H{+}2T_W((1{+}\tau)\tilde{o}_W)^{-\theta}, \quad P \;=\; T_H{+}T_W\big((1{+}a)(1{+}\tau)\tilde{o}_W\big)^{-\theta}, \quad R \;=\; T_H{+}T_W((1{+}\tau)\tilde{o}_W)^{-\theta}.$$

With this notation,

$$U^{innov}(T_H,T_W,\tau) = (w_H^p)^{-\alpha-\delta}\,\beta\,(1-\beta)^{\delta-1}\Bigg[\Big(\tfrac{\eta l_R}{f}I_0\Big)^{1/2}\Big(2^{-1/2}+2^{1/2}(1-\beta)\Big) + \Big(\tfrac{2\eta l_R}{fI_0}\Big)^{1/2}T_H\Gamma\left((1+a)^{2\delta}\frac{A^s}{P}-(1+a)^{\delta}B^s\Big(\tfrac{1}{P}+\tfrac{1}{R}\Big)+\frac{C^s}{R}\right)\Bigg],$$

and

$$I_0 \;=\; \Gamma T_H\Big[(1+a)^{2\delta}A^{s-1} - 2(1+a)^{\delta}B^{s-1} + C^{s-1}\Big].$$

Assumption 2 (Parameter domain)**.** *$T_H > 0$, $T_W \geq 0$, $\tau \geq 0$, $a > 0$, $\theta > 0$, $\delta \in (0,\theta)$, $\eta > 0$, $l_R > 0$, $f > 0$, $w_H^p > 0$, and $\beta \in (0,1)$.*

Proposition 3 (Positivity)**.** *Under Assumption 2, the innovation component of welfare satisfies $N(T_H,T_W,\tau) > 0$ for all $T_W \geq 0$ and $\tau \geq 0$.*

Proof. Since $w_H^p > 0$ and $\beta \in (0,1)$, the scalar prefactor $(w_H^p)^{-\alpha-\delta}\beta(1-\beta)^{\delta-1}$ is strictly positive.

By construction, $I_0 = \int \phi^\delta D(\phi,\tau)^2 f_H(\phi)\,d\phi > 0$ (as an integral of a square times a density). The closed form shown above is algebraically equal to this integral, hence $I_0 > 0$ and the square-root terms are well-defined.

$$\Big(2^{-1/2}+2^{1/2}(1-\beta)\Big) \;=\; \frac{3-2\beta}{\sqrt{2}} \;>\; 0 \quad \text{for } \beta \in (0,1.5),$$

hence, a fortiori, for $\beta \in (0,1)$.

Write the square-bracketed "correction" as

$$\Phi \;\equiv\; (1+a)^{2\delta}\frac{A^s}{P} - (1+a)^{\delta}B^s\Big(\frac{1}{P}+\frac{1}{R}\Big) + \frac{C^s}{R}.$$

Note that $C = R$, hence the last term equals C^{s-1}. Subtract and add $I_0/(T_H\Gamma) = (1+a)^{2\delta}A^{s-1} - 2(1+a)^{\delta}B^{s-1} + C^{s-1}$ to get

$$\Phi - \frac{I_0}{T_H\Gamma} = (1+a)^{2\delta}\Big(\frac{A^s}{P}-A^{s-1}\Big) \;-\; (1+a)^{\delta}B^{s-1}\Big(B\big(\tfrac{1}{P}+\tfrac{1}{R}\big)-2\Big).$$

Because $A > P$ and $s \in (0,1)$, $\frac{A^s}{P} - A^{s-1} = A^{s-1}\big(\frac{A}{P}-1\big) > 0$. Moreover, with $P < T_H + T_W(1+\tau)^{-\theta}(1+a)^{-\theta} < R$ and B lying between P and R, convexity of $x \mapsto B/x$ implies $B(\frac{1}{P}+\frac{1}{R}) - 2 \geq 0$. Hence the first term is nonnegative and the second is nonpositive, so $\Phi \geq I_0/(T_H\Gamma)$. Since $I_0 > 0$ and $\Gamma > 0$, we have $\Phi > 0$, with equality only at $T_W = 0$.

This implies the outer prefactor is strictly positive, the first inner term is strictly positive, and the curvature correction is nonnegative (indeed, strictly positive for $T_W > 0$). Therefore $N(T_H, T_W, \tau) > 0$. $\square$

Proof of Proposition 4. Let $s \equiv \delta/\theta \in (0,1)$, $\Gamma \equiv \Gamma(1-s) > 0$, and $\alpha \equiv (1+a)^{-\theta} \in (0,1)$. Define the effective World scale $\widetilde{T}_W := T_W/((1+\tau_H)\tilde{o}_W)^\theta$ and the shorthand

$$A := T_H{+}2\alpha\,\widetilde{T}_W, \quad B := T_H{+}(1{+}\alpha)\,\widetilde{T}_W, \quad C := T_H{+}2\,\widetilde{T}_W, \quad P := T_H{+}\alpha\,\widetilde{T}_W, \quad R := T_H{+}\widetilde{T}_W.$$

Then

$$I_0(\widetilde{T}_W) = \Gamma T_H\big[(1+a)^{2\delta}A^{\,s-1} - 2(1+a)^{\delta}B^{\,s-1} + C^{\,s-1}\big]\,,$$

and

$$U^{\text{innov}}(T_H,\widetilde{T}_W) = K_0\Bigg[\big(\tfrac{\eta\ell_R}{f}\,I_0(\widetilde{T}_W)\big)^{1/2}\,\underbrace{\big(2^{-1/2}+2^{1/2}(1-\beta)\big)}_{=:K_1>0} + \big(\tfrac{2\eta\ell_R}{f\,I_0(\widetilde{T}_W)}\big)^{1/2}\,T_H\Gamma\left((1+a)^{2\delta}\frac{A^{\,s}}{P} - (1+a)^{\delta}B^{\,s}\Big(\tfrac{1}{P}+\tfrac{1}{R}\Big) + \frac{C^{\,s}}{R}\right)\Bigg],$$

with $K_0 := (w_H^p)^{-\alpha-\delta}\,\beta\,(1-\beta)^{\delta-1} > 0$. For the sign of $U^{\text{innov}\,\prime}$ we analyze the two building blocks

$$\sqrt{I_0(\widetilde{T}_W)} \qquad\text{and}\qquad \frac{H(\widetilde{T}_W)}{\sqrt{I_0(\widetilde{T}_W)}},\quad H(\widetilde{T}_W) := T_H\Gamma\left((1+a)^{2\delta}\frac{A^{\,s}}{P} - (1+a)^{\delta}B^{\,s}\Big(\tfrac{1}{P}+\tfrac{1}{R}\Big) + \frac{C^{\,s}}{R}\right).$$

Step 1: behavior near $\widetilde{T}_W = 0$. At $\widetilde{T}_W = 0$ we have $A = B = C = P = R = T_H$ and

$$I_0(0) = \Gamma T_H\,T_H^{s-1}\big[(1+a)^{2\delta} - 2(1+a)^{\delta} + 1\big] = \Gamma T_H^{s}\big[(X-1)^2\big] > 0, \quad X := (1+a)^{\delta} > 1,$$

$$H(0) = T_H\Gamma\,T_H^{s-1}\big[(1+a)^{2\delta} - 2(1+a)^{\delta} + 1\big] = \Gamma T_H^{s}(X-1)^2 > 0.$$

Differentiate I_0 at 0. Using $\frac{d}{dx}(T_H+cx)^{s-1} = (s-1)c(T_H+cx)^{s-2}$,

$$\begin{aligned} I_0'(0) &= \Gamma T_H(s-1)T_H^{s-2}\Big[2\alpha(1+a)^{2\delta} - (1+\alpha)2(1+a)^{\delta} + 2\Big] \\ &= \Gamma T_H(s-1)T_H^{s-2}\cdot 2\big(X-1\big)\big(\alpha X-1\big). \end{aligned}$$

Since $X > 1$ and $\alpha X = (1{+}a)^{\delta-\theta} < 1$ (because $\theta > \delta$), the bracket is negative; as $(s{-}1) < 0$, it follows that $I_0'(0) > 0$, hence $\frac{d}{d\widetilde{T}_W}\sqrt{I_0(\widetilde{T}_W)}\big|_{0^+} > 0$. A parallel linearization of H (same coefficients and exponents) yields $H'(0) > 0$; together with $I_0(0) > 0$ this implies

$$\frac{d}{d\widetilde{T}_W}\frac{H(\widetilde{T}_W)}{\sqrt{I_0(\widetilde{T}_W)}}\Bigg|_{0^+} = \frac{H'(0)}{\sqrt{I_0(0)}} - \frac{H(0)I_0'(0)}{2I_0(0)^{3/2}} > 0 \quad \text{for all } a > 0,\ \delta \in (0,\theta).$$

Therefore $U^{\text{innov}\,\prime}(0^+) > 0$. By continuity, there exists $\varepsilon > 0$ such that $U^{\text{innov}\,\prime}(\widetilde{T}_W) > 0$ on $(0,\varepsilon)$.

Step 2: tail behavior as $\widetilde{T}_W \to \infty$. Use the standard asymptotics $(T_H + k\widetilde{T}_W)^{s-1} \sim k^{\,s-1}\widetilde{T}_W^{\,s-1}$ and $\frac{(T_H+k\widetilde{T}_W)^s}{T_H+m\widetilde{T}_W} \sim k^{\,s}m^{-1}\widetilde{T}_W^{\,s-1}$. Then

$$I_0(\widetilde{T}_W) = \Gamma T_H\,\Sigma_0\,\widetilde{T}_W^{\,s-1}\big(1+o(1)\big), \qquad H(\widetilde{T}_W) = T_H\Gamma\,\Sigma_1\,\widetilde{T}_W^{\,s-1}\big(1+o(1)\big),$$

where

$$\Sigma_0 := (1+a)^{2\delta}(2\alpha)^{\,s-1} - 2(1+a)^{\delta}(\alpha+1)^{\,s-1} + 2^{\,s-1},$$
$$\Sigma_1 := (1+a)^{2\delta}\frac{(2\alpha)^{\,s}}{\alpha} - (1+a)^{\delta}(\alpha+1)^{\,s}\Big(\tfrac{1}{\alpha}+1\Big) + \frac{2^{\,s}}{1}.$$

We claim $\Sigma_0 > 0$ and $\Sigma_1 > 0$. Set $p := 1-s \in (0,1)$ and $f(t) := t^{-p}$ (decreasing and convex). With $d_1 := 2\alpha$, $d_2 := \alpha+1 = \frac{d_1+d_3}{2}$, $d_3 := 2$, write

$$\Sigma_0 = f(d_1)X^2 - 2f(d_2)X + f(d_3), \quad \Sigma_1 = \frac{X^2}{\alpha}\,f(d_1) - \Big(\tfrac{1}{\alpha}+1\Big)X\,f(d_2) + f(d_3), \quad X := (1+a)^{\delta} > 1.$$

Both are quadratics in X with positive leading coefficients. Their minima occur at $X^\star = \frac{f(d_2)}{f(d_1)}$ (for Σ_0) and $X^\star = \frac{\alpha f(d_2)}{f(d_1)}$ (for Σ_1) and equal, respectively,

$$\min\Sigma_0 = f(d_3) - \frac{f(d_2)^2}{f(d_1)}, \qquad \min\Sigma_1 = f(d_3) - \frac{\alpha\,f(d_2)^2}{f(d_1)}.$$

By convexity (Jensen) and AM–GM,

$$f(d_2) = f\Big(\tfrac{d_1+d_3}{2}\Big) \le \tfrac{f(d_1)+f(d_3)}{2} \le \sqrt{f(d_1)f(d_3)},$$

with strict inequalities since $d_1 \ne d_3$ when $a > 0$ (i.e., $\alpha < 1$). Hence $f(d_2)^2 < f(d_1)f(d_3)$, which implies $\min\Sigma_0 > 0$ and $\min\Sigma_1 > 0$, and thus $\Sigma_0, \Sigma_1 > 0$ for all $X \ge 1$. Therefore,

$$\sqrt{I_0(\widetilde{T}_W)} \sim c_0\,\widetilde{T}_W^{\frac{s-1}{2}}, \qquad \frac{H(\widetilde{T}_W)}{\sqrt{I_0(\widetilde{T}_W)}} \sim c_1\,\widetilde{T}_W^{\frac{s-1}{2}},$$

with $c_0, c_1 > 0$, and

$$U^{\text{innov}}(T_H, \widetilde{T}_W) \sim C\, \widetilde{T}_W^{\frac{s-1}{2}} \xrightarrow[\widetilde{T}_W \to \infty]{} 0, \qquad \frac{d}{d\widetilde{T}_W} U^{\text{innov}}(T_H, \widetilde{T}_W) \sim \frac{s-1}{2}\, C\, \widetilde{T}_W^{\frac{s-3}{2}} \xrightarrow[\widetilde{T}_W \to \infty]{} 0^-,$$

since $s \in (0, 1)$.

Step 3: hump shape. From Steps 1–2 we have $U^{\text{innov}\prime}(\widetilde{T}_W) > 0$ on $(0, \varepsilon)$ and $U^{\text{innov}\prime}(\widetilde{T}_W) < 0$ for all $\widetilde{T}_W \geq M$ for some $0 < \varepsilon < M < \infty$, with U^{innov} continuous on $[0, \infty)$. By the intermediate value theorem, there exists at least one interior maximizer $\widetilde{T}_W^* \in (0, M)$ with $U^{\text{innov}\prime}(\widetilde{T}_W^*) = 0$. Thus U^{innov} increases for small effective World scale and decreases for sufficiently large effective World scale — an inverted-U (hump-shaped) response in $\widetilde{T}_W$. □

B Quality Improving Innovation

This appendix shows that if innovation upgrades *quality* rather than *productivity*, the innovation policy $u^*(\tilde{\phi}_H)$ remains single-peaked (inverted-U) with a unique interior maximizer.

Assumption 3 (Quality in profits and competition)**.** *Quality $k_H > 0$ enters per-period operating profits linearly and the pricing cutoff via $\tilde{w}_i^p = w_i^p / k_i^{1/\delta}$:*

$$\pi_H = \kappa\, \phi_H^\delta (\tilde{w}_H^p)^{-\delta} = \kappa\, \phi_H^\delta (w_H^p)^{-\delta} k_H, \qquad \frac{\phi_H}{\tilde{w}_H^p} = \frac{\phi_H\, k_H^{1/\delta}}{w_H^p}.$$

Quality-improving innovation raises k_H by a factor $(1 + a)$ with probability $u \in [0, 1]$ at cost $R(u, \tilde{\phi}_H) = \frac{f}{2} u^2\, \tilde{\phi}_H^\delta w_R$. World productivities are drawn from $G_W(x) = \exp(-T_W x^{-\theta})$. Assume $a > 0$, $T_W > 0$, $\tau_H \geq 0$, and $0 < \delta < \theta$.

Lemma 1 (Inverted-U under quality upgrading)**.** *Normalize baseline quality to $k_H \equiv 1$. When innovation succeeds, (k_H, ϕ_H) becomes $\big((1 + a), \tilde{\phi}_H\big)$; otherwise it remains $(1, \tilde{\phi}_H)$. Then the expected net gain from a successful innovation at baseline $\tilde{\phi}_H$ can be written as*

$$S_q(\tilde{\phi}_H) = (1 + a)\, G_W\big((1 + \tau_H)\tilde{o}_W\, (1 + a)^{1/\delta}\, \tilde{\phi}_H\big) \; - \; G_W\big((1 + \tau_H)\tilde{o}_W\, \tilde{\phi}_H\big),$$

so that

$$u^*(\tilde{\phi}_H) = \min\left\{1,\; \frac{\kappa\, (\tilde{w}_H^p)^{-\delta}}{f\, w_R}\, S_q(\tilde{\phi}_H)\right\}.$$

Moreover, $u^*(\tilde{\phi}_H)$ *is single-peaked with a unique interior maximizer*

$$\tilde{\phi}_H^* = \frac{1}{(1+\tau_H)\tilde{o}_W}\left(\frac{T_W\left[1-(1+a)^{-\theta/\delta}\right]}{\left(\frac{\theta}{\delta}-1\right)\,\ln(1+a)}\right)^{1/\theta}, \qquad u^{*\prime\prime}(\tilde{\phi}_H^*) < 0.$$

Proof. Write $b \equiv (1+\tau_H)\tilde{o}_W$ and $t(\tilde{\phi}_H) \equiv T_W b^{-\theta}\tilde{\phi}_H^{-\theta} > 0$. Under Assumption 3, a successful quality upgrade scales profits by $(1+a)$ and shifts the win cutoff as if $\tilde{\phi}_H$ were multiplied by $(1+a)^{1/\delta}$. Since $G_W(bx) = \exp(-T_W(bx)^{-\theta}) = \exp(-x^{-\theta}t)$, we obtain

$$S_q(\tilde{\phi}_H) = (1+a)\,e^{-(1+a)^{-\theta/\delta}t(\tilde{\phi}_H)} - e^{-t(\tilde{\phi}_H)}.$$

Differentiating as in the main text,

$$u^{*\prime}(\tilde{\phi}_H) \propto \left[(1+a)^{1-\theta/\delta}\,e^{-(1+a)^{-\theta/\delta}t(\tilde{\phi}_H)} - e^{-t(\tilde{\phi}_H)}\right].$$

Setting $u^{*\prime}(\tilde{\phi}_H) = 0$ yields

$$\left(1-\tfrac{\theta}{\delta}\right)\ln(1+a) = -\,t(\tilde{\phi}_H)\left[1-(1+a)^{-\theta/\delta}\right].$$

Because $0 < \delta < \theta$, the left-hand side is negative and the bracket on the right is positive, implying a unique solution for $t(\tilde{\phi}_H)$ and hence a unique $\tilde{\phi}_H^*$. Solving using $t = T_W b^{-\theta}\tilde{\phi}_H^{-\theta}$ gives the stated expression. The second derivative at $\tilde{\phi}_H^*$ is negative by the same argument as in Proposition 1, since $g_q(t) \equiv (1+a)e^{-(1+a)^{-\theta/\delta}t} - e^{-t}$ satisfies $g_q''(t^*) < 0$ at the critical point implied by $g_q'(t^*) = 0$. □

Comparing to the productivity-upgrading case. If we benchmark to the tariff-adjusted World mode $\bar{\phi}_{W,\text{mode}} = \frac{1}{(1+\tau_H)\tilde{o}_W}\left(\frac{\theta}{\theta+1}T_W\right)^{1/\theta}$, then

$$\frac{\tilde{\phi}_H^*}{\bar{\phi}_{W,\text{mode}}} = \left[\frac{\frac{\theta+1}{\theta}\left(1-(1+a)^{-\theta/\delta}\right)}{\left(\frac{\theta}{\delta}-1\right)\,\ln(1+a)}\right]^{1/\theta}, \qquad \frac{\tilde{\phi}_H^*}{\bar{\phi}_{W,\text{mode}}} \xrightarrow[a\to 0^+]{} \left(\frac{\theta+1}{\theta-\delta}\right)^{1/\theta}.$$

Thus the small-step limit coincides with the productivity-upgrading case, and the inverted-U shape is preserved.

C Benchmark Theory with No Innovation

This appendix solves the model when Home is a small economy. Home's population and fixed factor are scaled down by a factor $\eta \in (0,1)$ relative to Foreign. I shut down innovation and let intermediate productivities be exogenous Fréchet draws.

Sizes, preferences, and demand

There is one Foreign household and η identical Home households. Each household $c \in \{H,F\}$ has

$$U_c = \alpha \log z_c + (1-\alpha)\log D_c, \qquad I_c = p_z z_c + P D_c, \qquad 0 < \alpha < 1,$$

with $P \equiv 1$ by free trade in the differentiated final good. Per-household (per-capita) demands are

$$z_c = \frac{\alpha i_c}{p_z}, \qquad D_c = (1-\alpha) i_c,$$

where i_c is per-household income. Aggregates scale with population:

$$Z_H = \eta\, z_H, \quad \mathcal{D}_H = \eta\, D_H, \quad I_H = \eta\, i_H; \qquad Z_W = z_W,\ \mathcal{D}_W = D_W,\ I_W = i_W.$$

Home's stock of the fixed factor is $F_H = \eta$; Foreign's is $F_W = 1$.

Homogeneous good and production wages

The homogeneous good uses only production labor with unit productivity in Home and $o_W > 0$ in Foreign. Unit costs are

$$p_z^H = w_H^p, \qquad p_z^F = \frac{1}{o_W}\, w_W^p.$$

With free trade and perfect competition, $p_z^H = p_z^F \equiv p_z$, hence

$$w_H^p = \frac{1}{o_W}\, w_W^p, \qquad p_z = w_H^p = \frac{1}{o_W} w_W^p.$$

As Home is small, w_W^p (and thus p_z and w_H^p) are taken as exogenous from the world market.

Final producers, sourcing, and CES demands

The Foreign final producer sources only from Foreign intermediates:

$$\max_{F_W, \{q^F_{Fj}\}} \frac{F_W^\beta}{1-\beta} \int_0^1 \big(q^F_{Fj}\big)^{1-\beta} dj \ - \ f_W F_W \ - \ \int_0^1 p_{Fj}\, q^F_{Fj}\, dj,$$

with $F_W \equiv 1$. The Home final producer can source from Home and import from Foreign subject to a tariff $\tau \geq 0$ on intermediates:

$$\max_{F_H, \{q^H_{Hj}, q^F_{Hj}\}} \frac{F_H^\beta}{1-\beta} \int_0^1 \big(q^H_{Hj} + q^F_{Hj}\big)^{1-\beta} dj \ - \ f_H F_H \ - \ \int_0^1 p_{Hj} \big(q^H_{Hj} + q^F_{Hj}\big) dj,$$

with $F_H = \eta$. First-order conditions imply

$$f_W = \frac{\beta Y_W}{F_W} = \beta Y_W, \qquad f_H = \frac{\beta Y_H}{F_H} = \frac{\beta}{\eta} Y_H,$$

and (Marshallian) CES demands

$$q^F_{Fj} = (p_{Fj})^{-1/\beta}, \qquad q_{Hj} = \begin{cases} \eta\, (p^H_{Hj})^{-1/\beta}, & \text{if } p^H_{Hj} \leq (1+\tau) p^F_{Hj}, \\ \eta\, \big((1+\tau) p^F_{Hj}\big)^{-1/\beta}, & \text{if } (1+\tau) p^F_{Hj} < p^H_{Hj}, \end{cases}$$

where the delivered price in Home is $p_{Hj} = \min\{p^H_{Hj}, (1+\tau) p^F_{Hj}\}$.

Intermediates, pricing, and profits

An intermediate producer with productivity ϕ and production wage w^p_c has marginal cost $mc = w^p_c/\phi$ and charges the CES markup

$$p^*_c(\phi) = \frac{1}{1-\beta} \frac{w^p_c}{\phi}.$$

Given one Foreign and η Home consumers, per-variety operating profits are

$$\pi^F_W(\phi) = \kappa_W\, \phi^\delta, \qquad \pi^H_W(\phi) = \kappa_W\, (1+\tau)^{-1/\beta}\, \phi^\delta, \qquad \pi^H_H(\phi) = \eta\, \kappa_H\, \phi^\delta,$$

where $\delta = \dfrac{1}{\beta} - 1 > 0$ and $\kappa_c = \beta (1-\beta)^\delta\, (w^p_c)^{-\delta}$ for $c \in \{H, F\}$.

I assume no Home exports of intermediates by assigning tariff rates of infinity. Productivities are Fréchet:

$$G_H(\phi) = \exp(-T_H\phi^{-\theta}), \qquad G_W(\phi) = \exp(-T_W\phi^{-\theta}), \qquad \theta > 1,\ 0 < \delta < \theta.$$

Let $s \equiv \delta/\theta \in (0,1)$, $d \equiv (1+\tau)^{-\theta}$. I repeatedly use, for $K, T > 0$,

$$\int_0^\infty \phi^\delta e^{-K\phi^{-\theta}}\, T\theta\, \phi^{-(\theta+1)}\, d\phi = T\,\Gamma(1-s)\, K^{\,s-1}.$$

Aggregate outputs and profits Define $B(\tau) = T_W d + T_H$. Then

$$Y_H = \eta\,(1-\beta)^{\delta-1}\,(w_H^p)^{-\delta}\, B(\tau)^s\,\Gamma(1-s),$$
$$Y_W = (1-\beta)^{\delta-1}\,(w_W^p)^{-\delta}\, T_W^s\,\Gamma(1-s),$$

$$\Pi_H^H = \eta\,\kappa_H\, T_H\, B(\tau)^{\,s-1}\,\Gamma(1-s),$$
$$\Pi_W^F = \kappa_W\, T_W^s\,\Gamma(1-s),$$
$$\Pi_W^H = \eta\,\kappa_W\,(1+\tau)^{-\theta}\, T_W\, B(\tau)^{\,s-1}\,\Gamma(1-s).$$

Fixed-factor rents follow from the FOCs:

$$f_H = \frac{\beta}{\eta} Y_H = \beta(1-\beta)^{\delta-1}(w_H^p)^{-\delta}\, B(\tau)^s\,\Gamma(1-s), \qquad f_W = \beta Y_W = \beta(1-\beta)^{\delta-1}(w_W^p)^{-\delta}\, T_W^s\,\Gamma(1-s).$$

Labor market clearing

Let Z_c denote labor used in the homogeneous sector in country $c \in \{H, F\}$, and let L_H^H be Home labor used to produce intermediates for the Home market, L_W^H Foreign labor used to produce intermediates exported to Home, and L_W^F Foreign labor used to produce intermediates for the Foreign market.

With linear technologies for the homogeneous good, one unit of Home labor produces one unit, while one unit of Foreign labor produces o_W units. Hence, writing z_c for homogeneous-good quantities,

$$Z_H = z_H, \qquad Z_W = \frac{z_W}{o_W}.$$

Using the demand system $z_c = \alpha I_c/p_z$ and the free-trade pricing $p_z = w_H^p = \frac{1}{o_W} w_W^p$, this can be equivalently written as

$$Z_H = \frac{\alpha I_H}{w_H^p}, \qquad Z_W = \frac{\alpha I_W}{w_W^p}.$$

World labor market clearing. Total labor employed in the world equals the sum of homogeneous-good labor and intermediate-goods labor:

$$Z_H \;+\; Z_W \;+\; L_H^H \;+\; L_W^H \;+\; L_W^F \;=\; (L_H^P + L_W^P) \;=\; \eta + 1, \tag{47}$$

where I normalize Foreign production labor to $L_W^P = 1$ and scale Home by $\eta \in (0,1)$, i.e. $L_H^P = \eta$.

Small open economy limit. Home's final demand and all Home-market intermediate quantities scale linearly with η (recall $q_{Hj} = \eta\, p_{Hj}^{-1/\beta}$), so

$$Z_H = \mathcal{O}(\eta), \qquad L_H^H = \mathcal{O}(\eta), \qquad L_W^H = \mathcal{O}(\eta).$$

Taking $\eta \to 0$ in (47) yields

$$1 \;=\; Z_W \;+\; L_W^F, \tag{48}$$

i.e. world labor market clearing collapses to Foreign's own labor market. Therefore the Foreign production wage w_W^p is pinned down entirely by Foreign fundamentals (productivities, tariff-independent objects), and the Home production wage is determined residually via homogeneous-good arbitrage:

$$w_H^p = \frac{1}{o_W}\, w_W^p.$$

In the SOE benchmark, w_W^p (hence p_z and w_H^p) is thus *exogenous* to Home.

I pin down the Foreign production wage w_W^p from Foreign labor market clearing in intermediates. For a Foreign intermediate with productivity ϕ,

$$p_W^*(\phi) = \frac{1}{1-\beta}\,\frac{w_W^p}{\phi}, \qquad q_W(\phi) = \big(p_W^*(\phi)\big)^{-1/\beta}, \qquad \ell_W(\phi) = \frac{q_W(\phi)}{\phi}.$$

Hence, using $\delta = \frac{1}{\beta} - 1$,

$$\ell_W(\phi) = \frac{1}{\phi}\left(\frac{1}{1-\beta}\frac{w_W^p}{\phi}\right)^{-1/\beta} = (1-\beta)^{1/\beta}\,(w_W^p)^{-1/\beta}\,\phi^{\delta}.$$

With Fréchet productivity in Foreign,

$$g_W(\phi) = e^{-T_W\phi^{-\theta}}\,T_W\theta\,\phi^{-(\theta+1)}, \qquad s \equiv \frac{\delta}{\theta} \in (0,1),$$

the Foreign labor used in intermediates is

$$L_W^F = \int_0^\infty \ell_W(\phi)\, g_W(\phi)\, d\phi = (1-\beta)^{1/\beta}\, (w_W^p)^{-1/\beta} \underbrace{\int_0^\infty \phi^\delta\, g_W(\phi)\, d\phi}_{=T_W^s \Gamma(1-s)}.$$

Imposing the SOE resource condition $L_W^F = 1$ gives

$$1 = (1-\beta)^{1/\beta}\, (w_W^p)^{-1/\beta}\, T_W^s \Gamma(1-s) \implies (w_W^p)^{1/\beta} = (1-\beta)^{1/\beta}\, T_W^s \Gamma(1-s).$$

Therefore

$$w_W^p = \Big[(1-\beta)^{1/\beta}\, T_W^{\delta/\theta}\, \Gamma\big(1-\tfrac{\delta}{\theta}\big)\Big]^\beta = (1-\beta)\, T_W^{\frac{1-\beta}{\theta}} \Big[\Gamma\big(1-\tfrac{\delta}{\theta}\big)\Big]^\beta.$$

Free trade in the homogeneous good implies $p_z^H = p_z^F$ and, with unit Home productivity and Foreign productivity o_W in that sector,

$$p_z^H = w_H^p \qquad \text{and} \qquad p_z^F = \frac{1}{o_W}\, w_W^p,$$

so wage arbitrage gives

$$w_H^p = \frac{1}{o_W}\, w_W^p.$$

Income and welfare

Aggregate Home and Foreign incomes (no R&D) are

$$I_H = w_H^p \eta \;+\; f_H \eta \;+\; \Pi_H^H, \qquad I_W = w_W^p \;+\; f_W \;+\; \Pi_W^F \;+\; \Pi_W^H.$$

Per-household Home welfare (with $P \equiv 1$ and $p_z = w_H^p$) is

$$U_H = \alpha \log\!\left(\frac{\alpha i_H}{w_H^p}\right) + (1-\alpha)\log\big((1-\alpha) i_H\big) = \text{const} + \log i_H - \alpha \log w_H^p,$$

and aggregate (across η identical households) is ηU_H. Under the SOE closure, w_H^p is exogenous; thus

$$\frac{\partial U_H}{\partial \tau} = \frac{i_H'(\tau)}{i_H(\tau)} \quad \text{(or, in levels, } \frac{\partial I_H}{\partial \tau}\text{).}$$

Differentiating the above gives the usual decomposition: a negative output term through Y_H (since $B'(\tau) < 0$) and a positive rent-reallocation term through Π_H^H.

Per-household income in Home (with η Home households) is

$$i_H(\tau) \;=\; w_H^p \;+\; \frac{1}{\eta}\Pi_H^H \;+\; f_H.$$

Using the Fréchet algebra and the CES markup, this can be written as

$$i_H(\tau) \;=\; w_H^p + \left(w_H^p\right)^{1-\frac{1}{\beta}} \beta\,(1-\beta)^{\frac{1}{\beta}-2} \underbrace{\Big[(2-\beta)\,T_H \;+\; (1+\tau)^{-\theta}\,o_H^{\delta-\theta}\,T_W\Big]}_{\equiv\, S(\tau)} \underbrace{\Big[T_W\,o_H^{-\theta}(1+\tau)^{-\theta}+T_H\Big]^{s-1}}_{\equiv\, B(\tau)^{\,s-1}}, \tag{49}$$

where

$$s \equiv \frac{\delta}{\theta} \in (0,1), \qquad \delta = \frac{1}{\beta}-1, \qquad o_H > 0 \text{ (Home's homogeneous-good productivity).}$$

For notational convenience, define

$$B(\tau) = T_H + T_W\,o_H^{-\theta}(1+\tau)^{-\theta}, \qquad S(\tau) = (2-\beta)T_H + T_W\,o_H^{\delta-\theta}(1+\tau)^{-\theta}.$$

Monotonicity in τ. Since the prefactor in (49) is constant in τ, the sign of $i_H'(\tau)$ is the sign of

$$\frac{d}{d\tau}\Big\{\,B(\tau)^{\,s-1}\,S(\tau)\,\Big\} \;=\; B(\tau)^{\,s-2}\Big[(s-1)S(\tau)\,B'(\tau)+B(\tau)\,S'(\tau)\Big].$$

Using $B'(\tau) = -\theta T_W o_H^{-\theta}(1+\tau)^{-(\theta+1)} < 0$ and $S'(\tau) = -\theta T_W o_H^{\delta-\theta}(1+\tau)^{-(\theta+1)} < 0$, I can write

$$(s-1)S(\tau)\,B'(\tau)+B(\tau)\,S'(\tau) = B'(\tau)\Big[(s-1)S(\tau)+o_H^{\delta}\,B(\tau)\Big].$$

Hence $i_H'(\tau) < 0$ if and only if

$$(s-1)S(\tau)+o_H^{\delta}\,B(\tau) \;>\; 0\,.$$

Because $S(\tau)$ and $B(\tau)$ are affine in $(1+\tau)^{-\theta}$ with nonnegative slopes, the left-hand side is *increasing* in $(1+\tau)^{-\theta}$ and therefore minimized at $\tau\to\infty$, where

$$\lim_{\tau\to\infty}\Big[(s-1)S(\tau)+o_H^{\delta}B(\tau)\Big] = T_H\Big[o_H^{\delta}-(2-\beta)(1-s)\Big].$$

I conclude:

Proposition 4 (Per-household income falls with tariffs)**.** *Let $s=\delta/\theta\in(0,1)$ and $o_H>0$.*

Then $i_H(\tau)$ in (49) is strictly decreasing in τ for all $\tau \geq 0$ if and only if

$$o_H^{\delta} \;\geq\; (2-\beta)\left(1-\tfrac{\delta}{\theta}\right). \tag{50}$$

In the common normalization $o_H = 1$, this reduces to

$$\frac{\delta}{\theta} \;\geq\; \frac{1-\beta}{2-\beta}\,.$$

D PRODTR Data and Product Definitions

The Turkish *Prodcom* survey (*PRODTR*) is compiled annually by the Turkish Statistical Institute (TurkStat) and covers all establishments with at least 20 employees in the manufacturing sector (NACE Rev. 2 Section C) and in mining and quarrying (Section B). The dataset spans 2009–2019 and provides, for each firm and year, the value and physical quantity of every product produced.

Products are defined at the 10-digit PRODTR level, which extends the 8-digit European Union *Prodcom* classification by two additional digits. These final two digits introduce additional product differentiation specific to the Turkish production structure. The extra disaggregation greatly increases measurement precision relative to standard EU data, particularly in industries with substantial quality or material heterogeneity.

Example of additional disaggregation. The PRODTR classification subdivides many 8-digit *Prodcom* codes into finer product categories. For instance, EU *Prodcom* code 13.93.13.00—"Tufted carpets and other tufted textile floor coverings"—is split in the Turkish system into eight separate PRODTR subcategories:

1. 13.93.13.00.01 — Carpets and other floor coverings, of wool or fine animal hair, tufted.
2. 13.93.13.00.02 — Carpets and other floor coverings, of nylon or other polyamides, printed, area $<\leq 1$ m^2, tile, tufted.
3. 13.93.13.00.03 — Carpets and other floor coverings, of nylon or other polyamides, printed, area > 1 m^2, tufted.
4. 13.93.13.00.04 — Carpets and other floor coverings, of nylon or other polyamides, unprinted, area ≤ 1 m^2, tile, tufted.

5. 13.93.13.00.05 — Carpets and other floor coverings, of nylon or other polyamides, unprinted, area > 1 m^2, tufted.

6. 13.93.13.00.06 — Carpets and other floor coverings, of polypropylene, area ≤ 1 m^2, tile, tufted.

7. 13.93.13.00.07 — Carpets and other floor coverings, of polypropylene, area > 1 m^2, tile, tufted.

8. 13.93.13.00.99 — Tiles and other floor coverings of other synthetic substances suitable for weaving (tufted).

This level of granularity enables within-industry analysis of product differentiation, quality, and pricing behavior that is not possible with standard *Prodcom* data.

Coverage and reliability. The PRODTR survey is mandatory and forms part of the official industrial statistics program. Firms report production quantities and values based on administrative accounting records, which ensures high coverage and internal consistency. TurkStat cross-checks submitted values against VAT and export declarations to minimize reporting errors.

Quantities are reported in physical units (e.g., kilograms, square meters, pieces) following the same conventions as the EU *Prodcom* system, which facilitates matching with import data expressed under the GTIP classification. About 85 percent of product codes share identical unit definitions with their GTIP counterparts, allowing direct price comparisons after aggregation.

Usage in this paper. In the main text, I use PRODTR data to measure firm-product-year output values, quantities, and unit values (prices). These form the basis for estimating firm-level product quality and constructing the quality-adjusted price gap variable D_{ft}. Detailed documentation of the mapping between PRODTR and GTIP codes, and descriptive statistics on the overlap, are provided in Appendix E.

E PRODTR–GTIP Concordance

This appendix provides additional detail on the concordance between the Turkish industrial production survey (PRODTR, 10-digit codes) and the customs dataset (GTIP, 12-digit codes).

Table E1 summarizes the mapping structure. Out of 10,258 GTIP codes observed in the customs data, 9,383 map into PRODTR categories via many-to-one merges, 791 map one-to-one, and only 14 involve one-to-many merges. This confirms that harmonization primarily aggregates detailed GTIP classifications into broader—but still highly disaggregated—PRODTR categories.

Table E1: Product Classification and Mapping Statistics

Variable	**Value**
Number of PRODTR products	2,391
Number of GTIP products	10,258
Number of GTIP products with 1:1 merge	791
Number of GTIP products with 1:m merge	14
Number of GTIP products with m:1 merge	9,383
Number of firms producing each PRODTR product	
Mean	32.13
Median	15
25th percentile	9
75th percentile	31

To illustrate, consider two GTIP codes that both refer to woven fabrics used in the manufacture of medical bandages and gauze. In the official concordance, these GTIP codes are merged into a single PRODTR category that classifies cotton fabrics for medical use:

- **GTIP 520811100000:** Woven fabrics (≥85% cotton, unbleached, ≤100 g/m^2).
- **GTIP 520821100000:** Woven fabrics (≥85% cotton, bleached, ≤100 g/m^2).
- **PRODTR 13.20.20.20.00:** Woven cotton fabrics for medical gauze, bandages, and dressings (≤100 g/m^2).

This example highlights how harmonization aggregates detailed GTIP codes into consistent PRODTR categories, enabling direct comparison of domestic and imported goods at the product level.

F Construction of the Firm-Level Proximity Measure

This appendix provides full details on the construction of the proximity measure D_{ft} used in the main analysis.

Quality adjustment. Following Khandelwal et al. (2013), firm-level product quality is recovered from the CES demand restriction

$$\ln(\text{val}_{fpt}) + \frac{1-\beta}{\beta}\ln(p_{fpt}) = \gamma_{pt} + \epsilon_{fpt},$$

where val_{fpt} is the value of sales, p_{fpt} is the unit value, and γ_{pt} are product-year fixed effects. The residual ϵ_{fpt} captures firm-specific product quality under elasticity parameter $\beta \in (0,1)$. Quality-adjusted log prices are then

$$x_{fpt} = \ln(p_{fpt}) - \frac{\beta}{1-\beta}\epsilon_{fpt}.$$

This procedure is applied analogously to import prices at the country-product-year level to obtain quality-adjusted import prices.

Aggregation to the firm-year level. The firm-level proximity measure aggregates product-level quality-adjusted price gaps using lagged sales weights:

$$D_{ft} = \sum_{p \in P_{ft}} w_{fp,t-2}(x_{fpt} - y_{pt}), \quad w_{fp,t-2} = \frac{\text{sales}_{fp,t-2}}{\sum_{p' \in P_{f,t-2}} \text{sales}_{fp',t-2}}.$$

Using two-year-lagged shares mitigates endogeneity from contemporaneous demand or pricing shocks that affect both sales and innovation decisions.[37] In practice, more than 60 percent of firms are single-product producers, so the aggregate measure closely mirrors the product-level gap.

Modal import price. For each product p and year t, I compute y_{pt} as the mode of the kernel-smoothed distribution of quality-adjusted import prices across origin countries. The mode reflects the most typical foreign supplier and mitigates the influence of outliers or extreme price observations.

Descriptive statistics. The mean of the firm-product-level quality-adjusted price gap is 0.08 (standard deviation = 0.73); the median is 0.10. Thus, Turkish firms' effective prices are, on average, roughly 8–10 percent higher than those of their foreign competitors.

Figure F1 plots the kernel density of the firm-product-level gaps, while Figure J1 below shows the unadjusted counterpart using raw unit values for comparison.

[37]For example, if import prices fall for a product and a firm reduces its sales of that product, contemporaneous weighting would mechanically lower its contribution to D_{ft}. Lagged weights preserve the firm's ex-ante product importance.

Figure F1: Density plot of Firm-level Quality Adjusted Price Differences

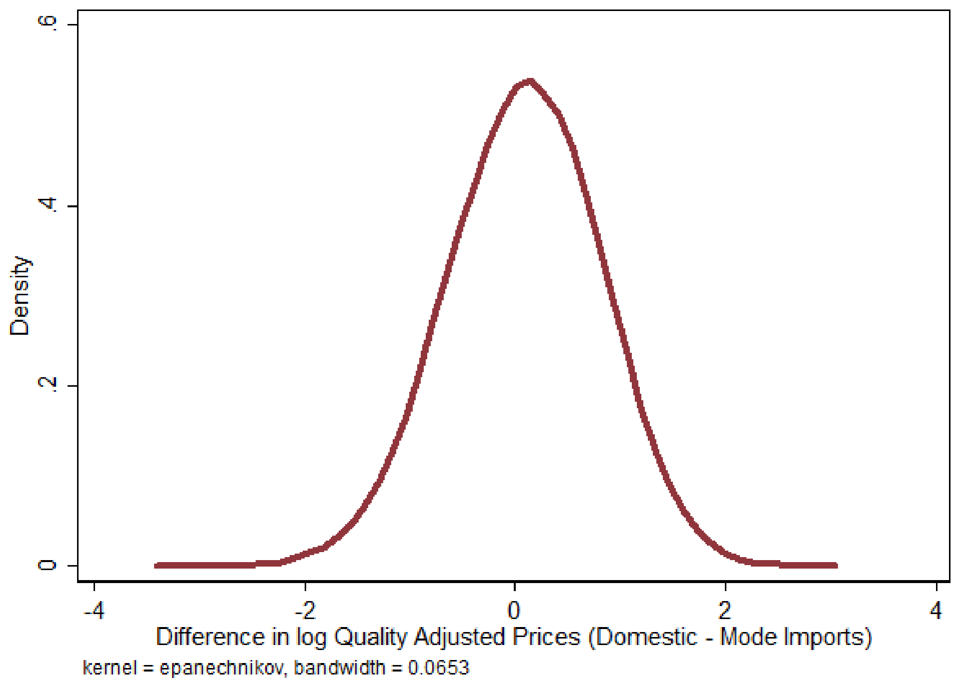

Notes: The figure plots the kernel density of quality-adjusted price differences at the firm–product level. A value above zero indicates that the domestic firm charges a higher price than its foreign competitors, while a value below zero indicates the opposite.

G Employer–Employee Data and Occupational Classification

The *4A Active Insured Administrative Records Micro* dataset is maintained by the Turkish Social Security Institution (SGK) and covers all employees in the formal sector. For each insured worker, it records the employing firm's unique tax identifier, the worker's occupation (ISCO–08, four-digit), and monthly earnings subject to social insurance contributions. The data span 2009–2019 and are reported directly by employers, ensuring full coverage of firms in the *Prodcom* sample.

Linkage to production data. I merge the employer–employee panel to the *Prodcom* (PRODTR) data using firm identifiers harmonized across datasets. The resulting linked sample contains 10,467 unique firms and 95,851 firm–year observations.

Occupational classification. Following Bernard et al. (2017), I group ISCO–08 occupations into six broad categories:

- **M:** Managers (ISCO 11XX–14XX)
- **R:** R&D-intensive occupations and technicians (ISCO 211X, 214X, 31XX)
- **S:** Support activities (clerical and administrative)
- **SA:** Sales occupations
- **L1:** Line workers in logistics
- **L2:** Line workers in other production activities

Group R represents high-skill technical labor responsible for research, development, and process engineering activities. For each firm–year, I compute both the total R&D wage bill and its share in the total wage bill:

$$s_{ft}^{\text{tech}} = \frac{\sum_{i \in R} w_{ift}}{\sum_i w_{ift}},$$

where w_{ift} denotes the annualized wage of worker i in firm f and year t. Analogously, I define the production wage bill share using Groups L1 and L2.

Summary statistics. Table G1 reports summary statistics for the R&D and production wage bills. The mean share of R&D wages in total pay is 12 percent, with a median of 4.5 percent. The corresponding mean share for production workers is 93 percent, reflecting the predominance of non-R&D labor in Turkish manufacturing.

Table G1: Summary Statistics for Employment Outcomes (2009–2019)

	R&D Wage Bill Share	**Log R&D Wage Bill**
Mean	0.1197	11.1513
Median	0.0453	11.1020
Standard deviation	0.1909	1.8834
25th percentile	0.0000	10.0097
75th percentile	0.1514	12.2939
Number of firm-year obs.	104,170	67,979
	Share of Production Wage Bill	**Log Production Wage Bill**
Mean	0.9314	12.8933
Median	0.9574	12.9748
Standard deviation	0.0920	1.7239
25th percentile	0.9158	12.1048
75th percentile	0.9815	13.8558
Number of firm-year obs.	104,170	103,337

Notes: The table reports summary statistics for the R&D wage bill and production wage bill at the firm–year level. The R&D wage bill is defined as total annual wages paid to ISCO-08 Group R (scientists, engineers, technicians). The share variables are relative to the total firm wage bill, and log values are in Turkish Lira.

H Construction of the Global Price Instrument

This section details the construction of the firm-level instrument z_{ft} used to identify exogenous variation in firms' proximity to foreign competitors (D_{ft}).

Step 1: Estimating global quality-adjusted export prices. Following Khandelwal et al. (2013), I first estimate origin–product–year quality using UN Comtrade data for all exporting countries excluding Turkiye:

$$\ln(\text{val}_{opt}) + \frac{1-\beta}{\beta}\ln(p_{opt}) = \alpha_{pt} + \epsilon_{opt}, \qquad o \neq \text{Turkiye}.$$

Here, val_{opt} and p_{opt} denote the export value and unit value of product p shipped from origin o in year t. The residual ϵ_{opt} captures origin-specific product quality, conditional on product–year fixed effects α_{pt}. Quality-adjusted log prices are then computed as

$$\tilde{y}_{opt} = \ln(p_{opt}) - \frac{\beta}{1-\beta}\epsilon_{opt}.$$

Step 2: Aggregating to the product level. For each product p and year t, I calculate the mode of $\tilde{y}_{opt}$ across exporting origins o that sold the product to Turkiye in year $t-2$. This mode captures the central tendency of global quality-adjusted export prices among relevant trading partners.

Step 3: Constructing the firm-level instrument. The firm-level instrument is a weighted average of these product-level global price trends:

$$z_{ft} = \sum_{p \in P_{ft}} w_{fp,t-2} \cdot \text{mode}_o(\tilde{y}_{opt}), \quad w_{fp,t-2} = \frac{\text{sales}_{fp,t-2}}{\sum_{p' \in P_{f,t-2}} \text{sales}_{fp',t-2}}.$$

Weights are lagged two years to ensure they are predetermined with respect to contemporaneous price shocks. A lower global export price (for a product heavily used by firm f) reduces z_{ft}, thereby tightening the firm's proximity to foreign competitors through a decline in D_{ft}.

Step 4: Validity. By construction, z_{ft} satisfies the two main IV conditions:

[nosep]*Relevance:* Turkish import prices strongly co-move with global export prices in the same product category. *Exclusion:* Global export prices are constructed excluding Turkish trade flows, and lagged product weights prevent contemporaneous feedback between Turkish innovation and global prices.

Turkiye accounts for less than 2% of global import demand and less than 5% of export supply in over 95% of origin–product–year observations (UN Comtrade, 2012–2021), supporting the assumption that Turkish shocks cannot influence world prices.

I Nonparametric Estimation of Foreign Proximity on Innovation

Having formally documented a quadratic relationship between the share of the tech wage bill and proximity to foreign competitors using parametric estimation, I now turn to a nonparametric approach. The goal is to estimate the function $g(D_{ft})$ flexibly, without imposing a specific functional form, while still addressing the endogeneity of D_{ft}.

The largest challenge in this context is that the proximity to foreign competitors D_{ft} is endogenous to the share of wage bill on technical workers s_{ft}^{tech}, which necessitates the use of instrumental variables in a nonparametric setting. This gives rise to the so-called *ill-posed inverse problem*: recovering the structural function $g(\cdot)$ from the conditional moment restriction $\mathbb{E}[s_{ft}^{tech} \mid \mathbf{z}_{ft}] = \mathbb{E}[g(D_{ft}) \mid \mathbf{z}_{ft}]$, where $\mathbf{z}_{ft}$ is the vector of instruments. As emphasized by Newey and Powell (2003) and Chetverikov and Wilhelm (2017), ill-posedness arises because the conditional expectation operator smooths g, so that inverting it amplifies small sampling errors in the estimated conditional moments into large fluctuations in $\hat{g}$. This instability is intrinsic to NPIV but does not arise in parametric IV regressions (such as the quadratic specification considered above), where one estimates only a finite number of coefficients.

In addition to this fundamental challenge, the inclusion of firm and sector-year fixed effects—entering the model linearly—raises the dimensionality of the problem and increases the difficulty of the estimation. The combination of high-dimensional fixed effects, instrumental variables, and nonparametric estimation makes this exercise particularly demanding.

To this end, I rely on the approach developed by Chetverikov and Wilhelm (2017), who propose imposing a monotonicity constraint on the estimated function to mitigate the ill-posedness of NPIV. By restricting $g(\cdot)$ to be monotone, their estimator eliminates artificial oscillations and greatly reduces variance, while remaining consistent under the monotonicity assumption.

Since my earlier results—both visual and quadratic parametric—suggest that the share of the tech wage bill follows a non-monotonic (inverted U-shaped) pattern, I adopt a subsample strategy to make the monotonicity restriction plausible within each subsample. Specifically, the quadratic estimates indicate that the share rises up to $D_{ft} \approx 0$ and falls thereafter. I therefore split the sample at $D_{ft} = 0$ and apply the monotone NPIV estimator separately to the two subsamples.

Formally, the estimated regression is

$$s_{ft}^{tech} = \alpha_W + \gamma_{st} + g(D_{ft}) + \epsilon_{ft},$$

where α_W and γ_{st} denote firm and sector-year fixed effects, and $g(\cdot)$ is estimated nonparametrically using z_{ft} as an instrument.

Figure I1: Nonparametric IV Estimation of Share of Tech Wage Bill and Proximity to Foreign Competitor

(a) Quality Adjusted Price Gap Below 0 | (b) Quality Adjusted Price Gap Above 0

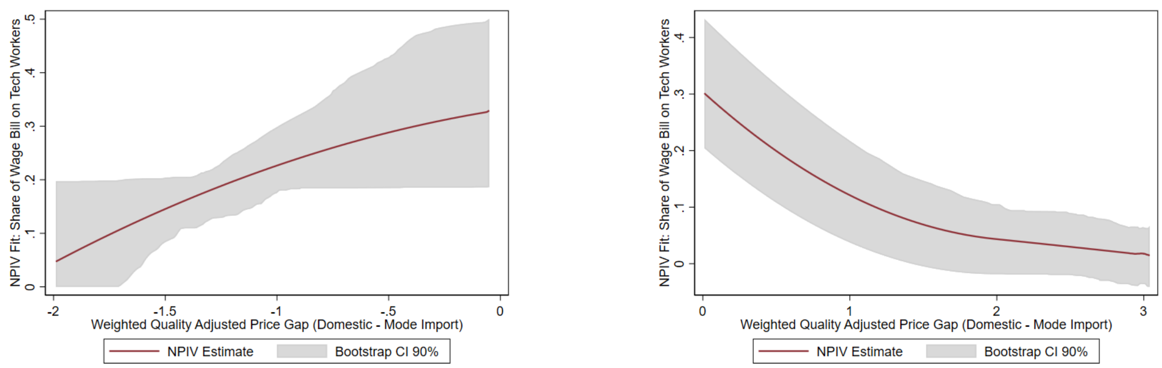

- *Note:* Monotone NPIV estimates are reported separately for the two subsamples ($D_{ft} < 0$ in panel (a) and $D_{ft} > 0$ in panel (b). Confidence intervals are constructed by bootstrapping 250 replications with resamples of size 10,000.

The results are shown in Figure I1. Panel (a) reports the nonparametric estimate when domestic firms charge less than the modal foreign competitor ($D_{ft} < 0$). The relationship is monotonically increasing: as domestic prices converge toward those of the foreign competitor, the share of the wage bill devoted to technical workers rises. Panel (b) reports results for $D_{ft} > 0$, where domestic firms charge more than the modal foreign competitor. Here the relationship is monotonically decreasing: the further domestic prices lie above those of foreign rivals, the lower the share of the wage bill allocated to technical workers. All estimates include firm and sector-year fixed effects, so identification comes from within-firm variation.

J Additional Figures and Tables

Figure J1: Density plot of Firm-level Price Differences

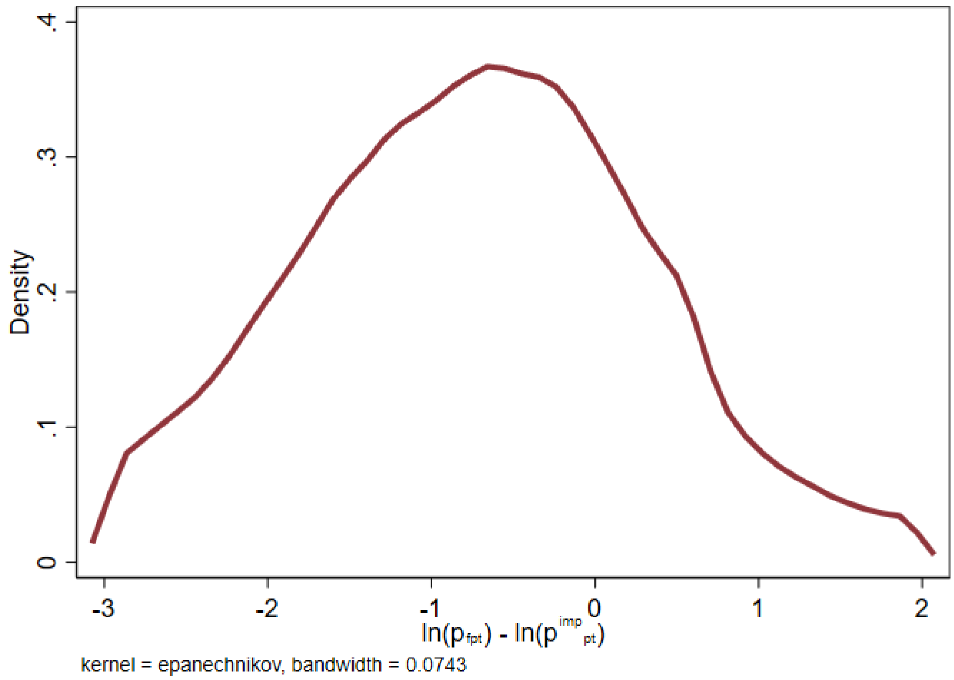

Note: This shows the k-density of the price difference at the firm-product level.

Table J1: Quadratic OLS Regression Results: R&D vs. Other Employment Payments

	(1) s_{ft}^{tech}	(2) $\ln(sales_{ft})$	(3) Prod Emp_{ft}	(4) Mrk Shr_{ft}	(5) $\ln(\text{profit}_{ft})$	(6) Exporter_{ft}
D_{ft}	0.00468	0.0736***	0.00758	0.000530***	0.0250***	0.00157
	(0.00362)	(0.00365)	(0.00776)	(0.000116)	(0.00517)	(0.00134)
D_{ft}^2	0.000781	-0.00269***	-0.00456**	0.0000786***	-0.00246**	-0.000215
	(0.000840)	(0.000848)	(0.00180)	(0.0000270)	(0.00119)	(0.000311)
Observations	95,851	95,851	95,851	95,851	73,355	95,851
R^2	0.709	0.873	0.979	0.912	0.839	0.754
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

All regressions include firm and sector-year fixed effects. Columns (1)–(6) present the OLS regressions, all measured at the firm f at time t: the share of tech. wage bill to production wage bill in column (1); log firm sales in column (2); Production wage bill in column (3); market share in column (4); log profit levels in column (5); and a binary indicator for whether the firm is an exporter in column (6). Standard errors are reported in parentheses and are clustered at the sector level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J2: Quadratic IV with Controls: Tech Wage Share vs. Competitive Price Gap

	(1) OLS s_{ft}^{tech}	(2) FS D_{ft}	(3) FS D_{ft}^2	(4) IV s_{ft}^{tech}
D_{ft}	0.00946**			0.0749
	(0.00446)			(0.120)
D_{ft}^2	0.000422			-0.0360***
	(0.000889)			(0.00988)
Cash ratio	0.000664	0.0163*	-0.00342	0.000772
	(0.00397)	(0.00879)	(0.0172)	(0.00405)
$\log(\text{Sales})$	-0.0105***	0.319***	-0.0628***	-0.0170*
	(0.00395)	(0.0194)	(0.0171)	(0.00893)
Distance from Domestic Mode	-0.0126**	0.300	-0.684***	-0.0811
	(0.00610)	(0.183)	(0.0228)	(0.0844)
z_{ft}		-0.0555***	0.0498*	
		(0.00572)	(0.0286)	
z_{ft}^2		0.00913***	0.132***	
		(0.00146)	(0.00729)	
Observations	85,650	85,650	85,650	85,650
R^2	0.715	0.268	0.0196	-0.0286
Firm FE	Yes	Yes	Yes	Yes
Sector–Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes

Notes: Column (1) reports OLS estimates of s_{ft}^{tech} on $\{D_{ft}, D_{ft}^2\}$ with controls (cash ratio, $\ln$ sales, and distance from the domestic mode) and firm, sector–year, and age fixed effects. Columns (2)–(3) report first-stage regressions for D_{ft} and D_{ft}^2 on instruments $\{z_{ft}, z_{ft}^2\}$ and the same controls. Column (4) reports the IV estimates. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J3: Quadratic Regression Results with Controls: R&D vs. Other Employment Payments

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln(R\&D_{ft})$	$\ln(\text{other emp.}_{ft})$	D_{ft}	D^2_{ft}	$\ln(R\&D_{ft})$	$\ln(\text{other emp.}_{ft})$
D_{ft}	0.00452	-0.000773			0.0144	-0.0118
	(0.00447)	(0.00361)			(0.0779)	(0.0621)
D^2_{ft}	-0.00302	-0.000559			-0.0390**	-0.00721
	(0.00213)	(0.000904)			(0.0192)	(0.0109)
$\log(\text{Age})$	0.0682**	0.178***	-0.0277	0.0790	0.0709**	0.178***
	(0.0282)	(0.0283)	(0.0230)	(0.0923)	(0.0290)	(0.0273)
$\log(\text{Liquidity})$	-0.00222	-0.00379	-0.00195	0.0174	-0.00157	-0.00370
	(0.00637)	(0.00288)	(0.00399)	(0.0134)	(0.00648)	(0.00285)
K^2_{ft}	-0.00193	0.00122	0.119***	0.126***	0.00143	0.00337
	(0.00142)	(0.000964)	(0.0144)	(0.0283)	(0.00892)	(0.00757)
$\log(\text{Sales})$	0.112***	0.125***	0.0436***	-0.0658***	0.109***	0.125***
	(0.0162)	(0.0176)	(0.00827)	(0.0197)	(0.0151)	(0.0194)
Exporter	0.0736***	0.0678***	0.00788	-0.0214	0.0727***	0.0677***
	(0.0101)	(0.00651)	(0.0119)	(0.0344)	(0.0103)	(0.00623)
z_{ft}			-0.0718**	0.194**		
			(0.0348)	(0.0871)		
z^2_{ft}			0.0217***	0.0459*		
			(0.00663)	(0.0252)		
Observations	68,747	68,741	68,747	68,747	68,747	68,741
R^2	0.882	0.956	0.108	0.0202	0.00538	0.0919
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
F-Stat					94.93	94.97

All regressions include firm and sector-year fixed effects. Columns (1)–(2) report OLS estimates; columns (3)–(4) report first-stage results of the 2SLS estimation; columns (5)–(6) report second-stage IV estimates. All specifications include controls for firm age, liquidity, sales, distance to mode domestic price K^2_{ft}, and exporter status. Standard errors are reported in parentheses and are clustered two-way at the firm and sector level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J4: Quadratic Specification with Positive Tech Workers: Tech Wage Share vs. Competitive Price Gap

	(1) OLS s_{ft}^{tech}	(2) FS D_{ft}	(3) FS D_{ft}^2	(4) IV s_{ft}^{tech}
D_{ft}	0.00538			0.0573
	(0.00553)			(0.125)
D_{ft}^2	0.00131			-0.0673***
	(0.00121)			(0.0260)
z_{ft}		-0.0734***	0.0579	
		(0.00843)	(0.0387)	
z_{ft}^2		0.0349***	0.123***	
		(0.00193)	(0.00885)	
Observations	61,710	61,710	61,710	61,710
R^2	0.719	0.00719	0.00790	-0.0703
Firm FE	Yes	Yes	Yes	Yes
Sector–Year FE	Yes	Yes	Yes	Yes

Notes: Column (1) reports the OLS relation between the tech wage bill share s_{ft}^{tech} and the competitive price gap D_{ft} (and its square). Columns (2) and (3) are first stages with D_{ft} and D_{ft}^2 as dependent variables and instruments z_{ft} and z_{ft}^2 (constructed from the price-gap shifters). Column (4) reports the second-stage IV estimates of s_{ft}^{tech} on $\{D_{ft}, D_{ft}^2\}$. All specifications include firm and sector–year fixed effects; standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J5: Robustness: Firms with Importer and Exporter FE

	(1)
	s_{ft}^{tech}
D_{ft}	0.0410
	(0.0954)
D_{ft}^2	-0.0539***
	(0.0199)
Observations	95,851
R^2	-0.0601
Firm FE	Yes
Sector-Year FE	Yes
Exporter-Year FE	Yes
Importer-Year FE	Yes
F-Stat	79.74

Sample restricted to firms that neither imported nor exported during the study period. Column (1) uses R&D wage bill; Column (2) uses balance-sheet R&D expenditures. Both columns include firm and sector-year fixed effects. Standard errors clustered at the firm and product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J6: Robustness: Pre-trends

	(1)
	s_{ft-1}^{tech}
D_{ft}	0.0648
	(0.917)
D_{ft}^2	-0.0769
	(0.153)
Observations	69,486
R^2	-0.00306
Firm FE	Yes
Sector-Year FE	Yes
F-Stat	36.23

The dependent variable is the share of wage bill on tech. workers in period $t-1$, to check for pre-trends. Standard errors clustered at the firm and sector level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure J2: Binned Scatterplots of R&D Spending and the Competitive Price Gap - Year FE only

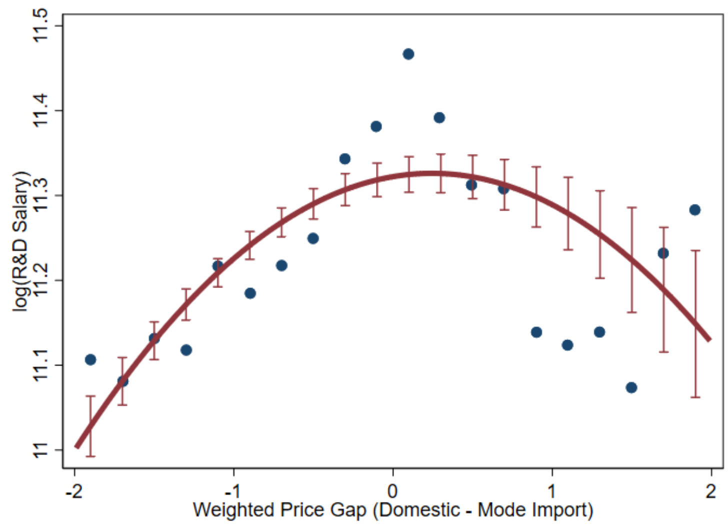

K Korea-Turkiye FTA

As stated above, the Bartik instrument is ideal for identifying the inverted U-shaped relationship between R&D spending and competitive pressure, thanks to the rich, continuous, and exogenous variation in price gaps it provides within firms and over time. However, the Bartik strategy is less suitable for studying pre-trends, dynamic adjustment paths, or clean before-and-after comparisons of firm outcomes. For this type of analysis, I turn to the Korea–Turkiye Free Trade Agreement (FTA), which provides a sharp, one-time, exogenous reduction in import prices for affected products. This policy shock is particularly well-suited for an event-study or difference-in-differences framework, allowing me to assess how sales, R&D, and productivity evolved in the periods before and after the FTA.

In this section, I first describe the details of the Korea–Turkiye FTA and explain how it changed import prices in the Turkish economy. I then document how this policy affected the Turkish manufacturing sector more broadly and, in particular, the R&D spending of domestic producers in the product categories exposed to the FTA. This analysis complements the earlier Bartik-based results by providing an alternative identification strategy, focused on a single, well-defined policy shock, to study how competitive pressure from foreign imports shapes domestic firms' innovation and performance.

Figure J3: Binned Scatterplots of Other employment spending and the Competitive Price Gap

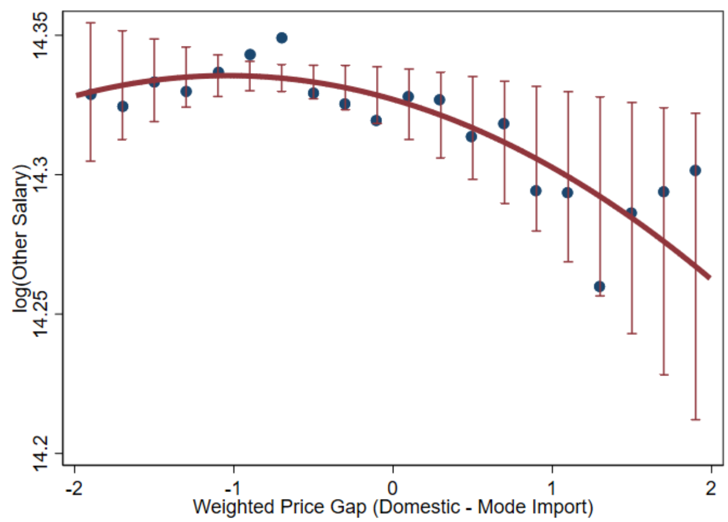

K.1 Context and Policy

Turkiye has been part of the European Union (EU)–Turkiye Customs Union since 1995. As part of this agreement, Turkiye committed to aligning its trade policy with that of the EU in order to avoid trade diversion and maintain the integrity of the customs union. Article 13 of the Customs Union Decision explicitly states:

> *1. Upon the date of entry into force of this Decision, Turkiye shall, in relation to countries which are not members of the Community, align itself on the Common Customs Tariff.*
> *2. Turkiye shall adjust its customs tariff whenever necessary to take account of changes in the Common Customs Tariff.*

These provisions require Turkiye to adopt the same tariff reductions that the EU agrees to with third parties, even though Turkiye does not participate in the negotiation of these agreements and has no formal say in their content. In effect, Turkiye is expected to adjust its trade policies exogenously in response to agreements over which it has no influence.

During the period covered by this study, there were relatively few major changes in Turkiye's trade policy. However, one significant event was the Free Trade Agreement (FTA) between the EU and South Korea, signed in 2011 and entering into force in July 2011. The EU–Korea FTA committed both parties to progressively abolish customs duties on over 98% of goods traded between them (as listed in Annex 2-A) within five years. This agreement had a substantial impact on bilateral trade: by 2024, total bilateral trade

Figure J4: Binned Scatterplots of Share of Wage Bill on Tech workers and the Firm-Level Price Gap

Figure J5: IV-Adjusted

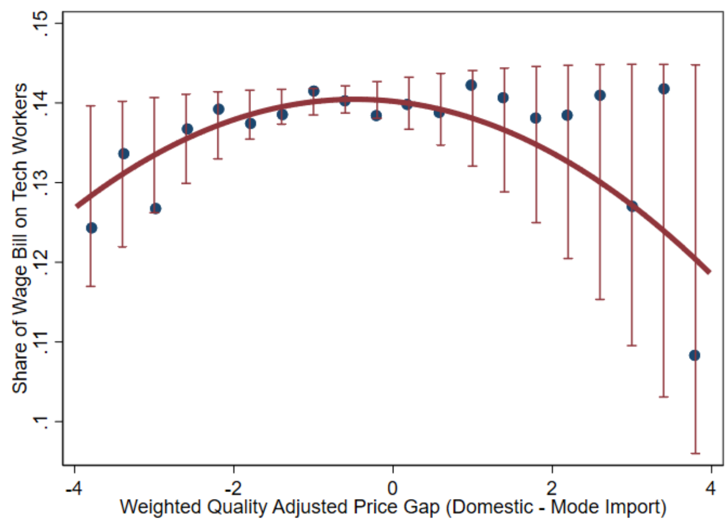

Note: The figure plots the kernel-weighted binscatter of the share wage bill on technical workers against the quality-adjusted price gap with foreign competitors, with 20 bins, firm and sector-year fixed effects. The right panel instruments D_{ft} with lagged global export prices.

in goods between the EU and Korea reached approximately €123.8 billion, representing a 96.5% increase compared to 2011 (€63 billion).

In accordance with the Customs Union rules, Turkiye adopted similar tariff reductions with Korea, signing its own bilateral trade agreement in 2013. This agreement applied to the same set of products for which tariffs were eliminated under the EU–Korea FTA. As a result, Turkish imports from Korea increased significantly in the products where tariff barriers declined. This exogenous policy shock—driven by the EU–Korea agreement and Turkiye's obligations under the Customs Union—provides a natural experiment for studying how Turkish firms responded to increased foreign competition.

Figure K1 presents a histogram of the number of products that previously faced high tariff rates but experienced tariff elimination after 2013 as a result of Turkiye's compliance with the Customs Union rules.

I focus on products that experienced more than a 3 percentage point reduction in tariffs following the 2013 Korea–Turkiye FTA. Figure K2 illustrates the evolution of imports from Korea for these products compared to products with tariff reductions below 3 percentage points. Panel (a) shows the value of imports from Korea, measured in USD and normalized to their 2009 level. The solid maroon line represents products with tariff re-

Figure J6: Binned Scatterplots of R&D Spending and the Competitive Price Gap for Firms Facing Low Import Competition

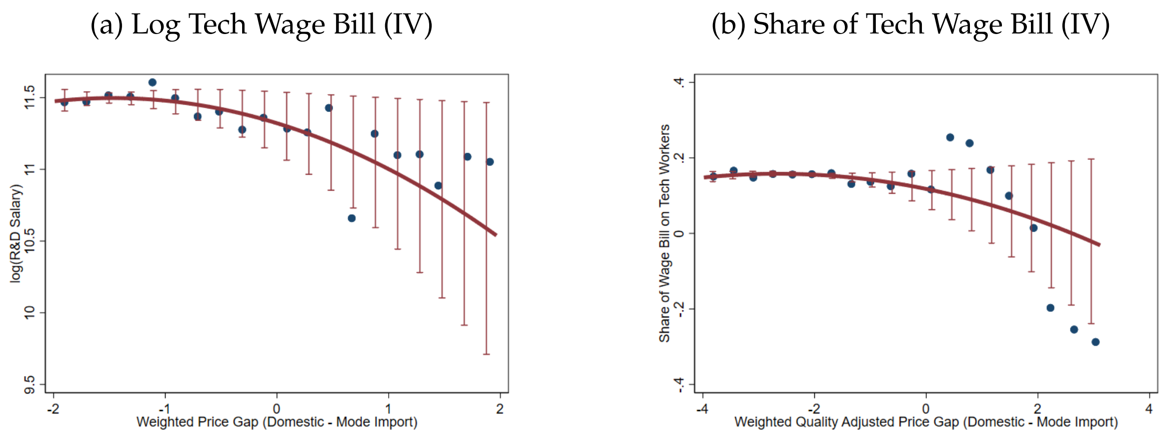

Table J7: Quadratic IV with R&D from Balance Sheet

	(1) IV $\ln(R\&D_{ft})$
D_{ft}	0.125
	(0.326)
D_{ft}^2	-0.148*
	(0.0961)
Observations	14,214
Firm FE	Yes
First-stage F	29.40

Notes: Column (1) reports IV estimates of $\ln(R\&D_{ft})$ on $\{D_{ft}, D_{ft}^2\}$ using balance sheet data. All specifications include firm fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ductions exceeding 3 percentage points, while the dashed navy line represents products with smaller reductions. Panel (b) shows the corresponding import quantities, measured in kilograms and similarly normalized. Both panels show a sharp increase in imports for products that experienced larger tariff reductions, while imports of other products remain relatively stable over the period.

While the primary focus of this study is on how the Korea–Turkiye FTA affected Turkish firms through increased import competition in the domestic market, it is important to acknowledge that the FTA also reduced tariffs on Turkish exports to Korea. In principle, these tariff reductions could generate new export opportunities for Turkish firms, potentially influencing their R&D and production decisions. To assess the relevance of this channel, I examine whether the set of products liberalized by Turkiye (i.e., those facing

Table J8: Quadratic IV by Sectoral Profitability: Tech Wage Share vs. Competitive Price Gap

	Below-Median Profit Sectors				Above-Median Profit Sectors			
	(1) OLS s_{ft}^{tech}	(2) FS D_{ft}	(3) FS D_{ft}^2	(4) IV s_{ft}^{tech}	(5) OLS s_{ft}^{tech}	(6) FS D_{ft}	(7) FS D_{ft}^2	(8) IV s_{ft}^{tech}
D_{ft}	0.0103			0.120	-0.000892			-0.0181
	(0.0112)			(0.134)	(0.0166)			(0.334)
D_{ft}^2	0.00113			-0.0185	-0.000827			-0.163***
	(0.00210)			(0.0127)	(0.00249)			(0.0351)
Liquidity	-0.00306	-0.00262	-0.0174	-0.00287	0.0140	-0.00663	0.0103	0.0152
	(0.00342)	(0.00606)	(0.0271)	(0.00499)	(0.0101)	(0.00666)	(0.0376)	(0.00999)
$\log$(Sales)	-0.0114**	0.0692***	-0.0252	-0.0185*	-0.0108	0.0499***	-0.155***	-0.0376*
	(0.00541)	(0.0139)	(0.0923)	(0.0106)	(0.0113)	(0.00811)	(0.0462)	(0.0216)
Distance to Dom. Mode	-0.00856	0.685***	-0.354	-0.0892	-0.0210	0.702***	-1.128**	-0.188
	(0.00765)	(0.0331)	(0.663)	(0.0929)	(0.0149)	(0.0589)	(0.521)	(0.254)
z_{ft}		-0.0583**	0.0618			-0.0496***	-0.0142	
		(0.0253)	(0.114)			(0.0139)	(0.0906)	
z_{ft}^2		0.0106**	0.112***			0.0127**	0.127**	
		(0.00454)	(0.0250)			(0.00565)	(0.0552)	
Observations	51,062	51,062	51,062	51,049	29,333	29,333	29,333	29,312
R^2	0.726	0.877	0.787	-0.0193	0.741	0.890	0.811	-0.300
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector–Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F				81.64				28.02

Notes: Columns (1)–(4) report estimates for sectors with below-median profits; Columns (5)–(8) for above-median profits. Columns (1) and (5) show OLS results for s_{ft}^{tech}. Columns (2)–(3) and (6)–(7) report first stages with D_{ft} and D_{ft}^2 as dependent variables. Columns (4) and (8) present IV estimates for s_{ft}^{tech}. All regressions include firm, sector–year, and age fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J9: Quadratic IV by Cash Ratio: Tech Wage Share vs. Competitive Price Gap

	Below-Median Cash Ratio Sectors				Above-Median Cash Ratio Sectors			
	(1) OLS s_{ft}^{tech}	(2) FS D_{ft}	(3) FS D_{ft}^2	(4) IV s_{ft}^{tech}	(5) OLS s_{ft}^{tech}	(6) FS D_{ft}	(7) FS D_{ft}^2	(8) IV s_{ft}^{tech}
D_{ft}	0.0141			-0.264	0.00720			0.294*
	(0.0174)			(0.213)	(0.00687)			(0.176)
D_{ft}^2	0.000853			-0.0949***	-0.000667			-0.0124
	(0.00482)			(0.0230)	(0.00119)			(0.0137)
Liquidity	-0.00704	0.00581	-0.0694	-0.0130	0.00567	-0.00340	0.0670*	0.00720
	(0.00913)	(0.00635)	(0.0619)	(0.0112)	(0.00589)	(0.00581)	(0.0338)	(0.00607)
$\log$(Sales)	-0.00900	0.0687***	-0.0583	0.00563	-0.0102	0.0601***	-0.0312	-0.0282**
	(0.0111)	(0.0138)	(0.0728)	(0.0163)	(0.00841)	(0.00868)	(0.0604)	(0.0119)
Distance to Dom. Mode	-0.0238	0.662***	-0.667	0.101	-0.0165	0.723***	-0.712	-0.231*
	(0.0232)	(0.0475)	(0.575)	(0.139)	(0.0176)	(0.0340)	(0.561)	(0.126)
z_{ft}		-0.0636***	0.0772			-0.0520**	0.0508	
		(0.0220)	(0.101)			(0.0193)	(0.103)	
z_{ft}^2		0.00914**	0.104***			0.00549	0.151***	
		(0.00382)	(0.0354)			(0.00532)	(0.0255)	
Observations	37,123	37,123	37,123	37,108	42,820	42,820	42,820	42,807
R^2	0.772	0.891	0.816	-0.174	0.719	0.879	0.796	-0.0778
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector–Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F				58.89				53.36

Notes: Columns (1)–(4) report estimates for sectors with below-median cash ratio; Columns (5)–(8) for above-median cash ratio. Columns (1) and (5) show OLS results for s_{ft}^{tech}. Columns (2)–(3) and (6)–(7) report first stages with D_{ft} and D_{ft}^2 as dependent variables. Columns (4) and (8) present IV estimates for s_{ft}^{tech}. All regressions include firm, sector–year, and age fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table J10: Adjustment Margins: Product Entry/Exit

	(1) ln(no. products)	(2) New Product Dummy
D_{ft}	-0.493***	-0.193***
	(0.118)	(0.0719)
D_{ft}^2	0.0339	0.0223
	(0.0265)	(0.0150)
Observations	96,807	96,807
Firm FE	Yes	Yes
Sector-Year FE	Yes	Yes
F-Stat	18.53	18.53

Each column uses a different dependent variable: log number of products for each firm-year and new product entry. Estimates capture firm responses to increased import prices. Standard errors clustered at the sector-level.

Table K1: Overlap of Turkiye and Korea Tariff Reductions (> 3pp)

	Korea tariff reduction > 3pp		
Turkiye tariff reduction > 3pp	No	Yes	Total
No	6,953 (57.3%)	5,184 (42.7%)	12,137 (100%)
Yes	2,456 (78.0%)	690 (22.0%)	3,146 (100%)
Total	9,409	5,874	15,283

Note: Table reports the cross-tabulation of products with tariff reductions greater than 3 percentage points by Turkiye and Korea under the FTA. Row percentages are reported in parentheses.

Figure K1: Histogram of tariff reductions

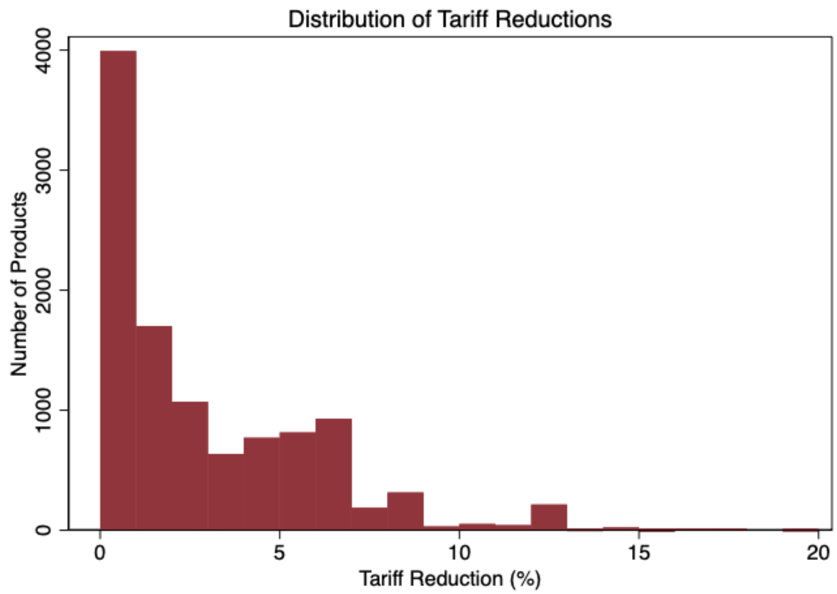

Note: The figure shows the distribution of tariff reductions following the Korea–Turkiye FTA.

Figure K2: Imports from Korea of products with tariff reductions above and below 3pp.

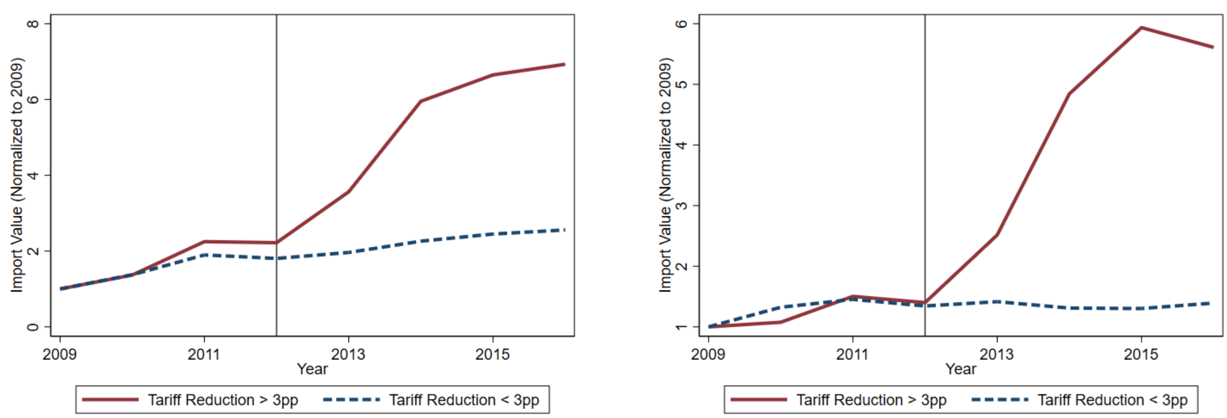

(a) Import value (USD, normalized to 2009) (b) Import quantity (kg, normalized to 2009)

Note: The solid maroon line depicts imports of products with tariff reductions greater than 3 percentage points; the dashed navy line depicts imports of products with smaller reductions. Both series are normalized to their respective 2009 levels.

increased competition from Korean imports) substantially overlaps with the set of products liberalized by Korea (i.e., those offering Turkish firms improved access to the Korean market). Appendix Table K1 reports the cross-tabulation of products that experienced more than a 3 percentage point tariff reduction on each side. Among the 3,146 products where Turkiye reduced tariffs, only 690 (22%) were also subject to tariff reductions by Korea. Conversely, 78% of products liberalized by Turkiye did not coincide with Korea's liberalization. This suggests that the export opportunities created by the FTA were concentrated in a different set of products than those affected by increased import competition, reinforcing the interpretation that the observed R&D responses were driven by heightened domestic competition rather than by improved access to the Korean market.

K.2 Imports of FTA-Liberalized Products

To more formally test the impact of the Korea–Turkiye FTA on Turkish imports from Korea, I estimate an event-study regression of the following form:

$$\ln(qty_{pct}) = \sum_{l=2009}^{2016} \beta_l \cdot 1(t = l) \cdot 1(c = \text{KOR}) \cdot 1(p \in P_{3pp}) + \alpha_{ct} + \gamma_{pt} + \eta_{pc} + \epsilon_{pct},$$

where qty_{pct} is the quantity of imports (in kilograms) of product p from country c in year t, and P_{3pp} denotes the set of products that experienced a tariff reduction of more than 3 percentage points under the Korea–Turkiye FTA. The coefficients β_l trace the differential change in imports of treated products from Korea relative to other products and trading partners in each year, normalized to their 2012 level. The specification includes country-year fixed effects (α_{ct}), product-year fixed effects (γ_{pt}), and product-country fixed effects (η_{pc}) to control for unobserved heterogeneity and time trends at each of these levels.

Figure K3 plots the estimated δ_l coefficients and their confidence intervals. The estimates indicate that imports of treated products from Korea remained stable prior to the FTA, suggesting no differential pre-trends, and then rose sharply following the implementation of the agreement. This pattern is consistent with the view that the Korea–Turkiye FTA substantially increased imports of products that saw meaningful tariff reductions, particularly on the extensive margin, while imports of other products remained relatively unchanged.

Figure K3: South Korea-Turkiye FTA Imports Event Study

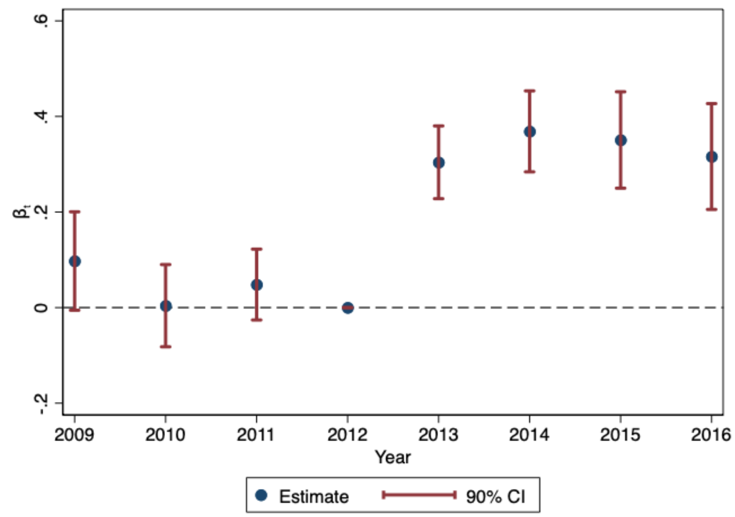

K.3 Domestic Producers of FTA-Liberalized Products

In this section, I examine how Turkish domestic producers of products affected by the Korea–Turkiye FTA adjusted their sales and R&D investments in response to increased competition from Korean imports. Similar to the previous section, I concord the GTIP import product codes with the corresponding PRODTR codes to match imported products that experienced a tariff reduction with their domestic counterparts. This allows me to track the outcomes of Turkish producers supplying the same products that became more exposed to Korean imports after the FTA.

To assess the dynamics of domestic sales, I estimate an event-study specification at the firm–product–year level, focusing on products that experienced more than a 3 percentage point tariff reduction and had a pre-FTA Korean import market share above 5%. The estimated specification is:

$$\ln q_{fpt} = \sum_{\ell \neq 2012} \beta_\ell \cdot 1\{p \in \text{TariffDrop} > 3 \ \wedge \ \text{MS}_p^{\text{KOR}} > 5\%\} \cdot 1\{t = \ell\} + \alpha_{fp} + \gamma_{st} + \epsilon_{fpt}.$$

Figure K4 shows that domestic producers of these products experienced a sharp decline in sales following the FTA, consistent with increased competition from Korean imports.

I further examine heterogeneity based on the firm's initial price position relative to Korean competitors. Firms that, prior to the FTA, priced their products higher than the Korean mode experienced more severe sales declines, while those that initially priced

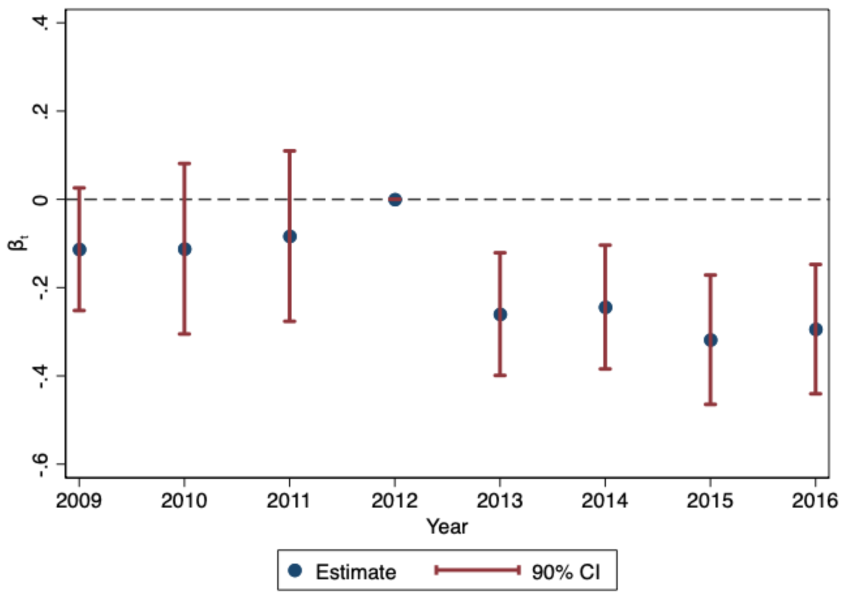

Figure K4: Event-study estimates of log sales of domestic producers of products affected by the Korea–Turkiye FTA.

below Korean competitors were less affected (Figure K5).

Finally, I analyze the dynamic response of R&D spending among domestic producers of affected products, weighted by their 2011 sales exposure to treated products. The estimated specification is:

$$\ln R\&D_{ft} = \sum_{\ell \neq 2012} \beta_{\ell}^{(g)} \cdot D_W \cdot 1\{t = \ell\} + \alpha_f + \gamma_{st} + \epsilon_{ft},$$

where D_W is the firm-level treatment intensity, computed as the weighted sum of sales shares of treated products in 2011, and g indicates the firm's initial price position relative to Korea.

The results, shown in Figure K6, reveal that firms with initially higher prices exhibit a noticeable increase in R&D spending after the FTA, consistent with efforts to improve competitiveness, while firms already pricing below Korean imports show a much more muted R&D response. This pattern aligns with the inverted-U relationship between competitive pressure and R&D documented in earlier sections.

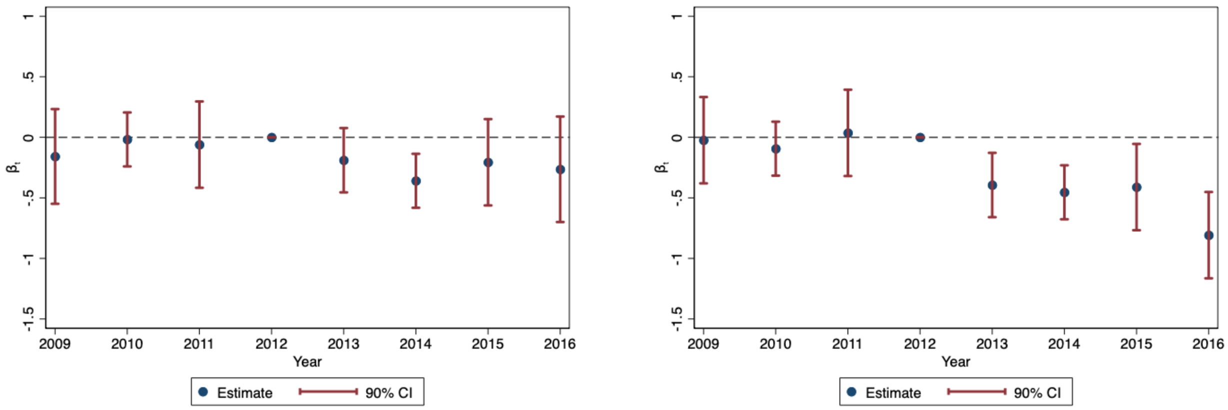

(a) Firms pricing higher than Korea (b) Firms pricing lower than Korea

Figure K5: Event-study estimates of log sales of domestic producers, by initial price position relative to Korea.

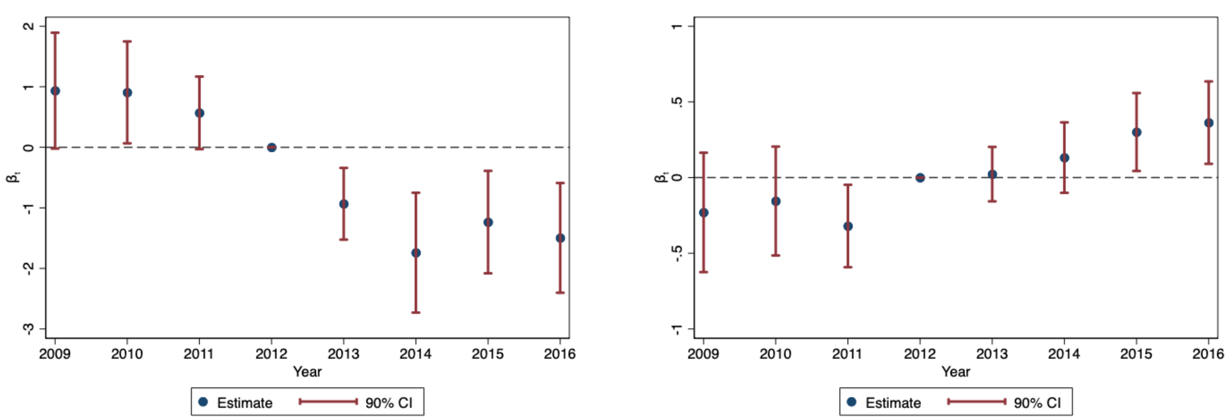

(a) Firms pricing higher than Korea (b) Firms pricing lower than Korea

Figure K6: Event-study estimates of R&D spending of domestic producers, by initial price position relative to Korea.

L Calibration Equations

L.1 Estimating θ from Price Dispersion

To estimate the shape parameter θ of the Fréchet distribution, I exploit the link between the dispersion of log prices and the underlying heterogeneity in productivity across firms implied by the Fréchet assumption. The model implies that log price gaps across firms follow a distribution whose standard deviation is inversely related to θ. Specifically, the standard deviation of log prices is given by:

$$\text{sd}[\ln p] = \frac{\pi}{\sqrt{6} \cdot \theta},$$

which can be inverted to yield an estimator:

$$\hat{\theta} = \frac{\pi}{\sqrt{6} \cdot \text{sd}[\ln p]},$$

where $\text{sd}[\ln p]$ is the empirical standard deviation of log prices across firms at the product-year level.

L.2 Domestic Production to Import Ratio

To back out the relative productivity shifter T_H/T_W, I use the model-implied expression for the domestic share of total input sourcing:

$$\frac{T_H}{T_W} = \frac{\text{shr}_d \cdot \left(\frac{k_W}{k_H}\right)^{\theta/\delta} \cdot o_W^{-\theta}}{T_W},$$

where shr_d is the domestic share of purchases, $\delta = 1/\beta - 1$, and o_W is the relative wage.

L.3 Quality Estimation for Home and Foreign

Firm-level prices and quantities are used to estimate the quality residuals $\varepsilon_{f(c)pt}$ from the following regression:

$$\ln(qp_{f(c)pt}) + \left(\frac{1}{\beta} - 1\right) \ln(p_{f(c)pt}) = \gamma_{pt} + \varepsilon_{f(c)pt},$$

where γ_{pt} captures product-time fixed effects. The aggregate quality index is then computed as the median of the exponentiated residuals:

$$k_c = \text{median}_{f(c)} \left[\exp(\varepsilon_{f(c)pt}) \right],$$

where $c = \text{TR}$ for Home and $c = \text{W}$ for World.

L.4 Innovation Productivity Gain

The productivity gain a from a successful innovation is disciplined by the ratio of observed prices at the peak R&D employment share and the modal imported price. From the model:

$$\frac{\tilde{p}_{fpt}^{\text{peak}}}{\tilde{p}_{pt}^{\text{mode, imp}}} = \left(\frac{t(t-\delta)\ln(1+a)}{(t-1)\left(1-(1+a)^{-t}\right)} \right)^{\frac{1}{t}},$$

where $t = \theta$, and $\delta = 1/\beta - 1$.

L.5 Fixed Cost of Innovation

The fixed cost of innovation f is backed out from the observed share of technical labor (R&D workers) and the model's equilibrium condition for innovation labor demand:

$$\frac{\eta \ell_R}{\ell_P} = \left(\frac{\eta \ell_R}{2f} \right)^{1/2} \cdot \frac{1}{\ell_P} \cdot \beta(1-\beta)^{1/\beta} \cdot \frac{\tilde{w}_H^{-\delta}}{w_H} \cdot \left[\int D(\phi)^2 \phi^\delta f_H(\phi) d\phi \right]^{1/2},$$

which implicitly solves for f, holding the integral and wage terms fixed by earlier calibration.