

Hémous, David; Lepot, Simon; Sampson, Thomas; Schärer, Julian

Working Paper

Trade, innovation and optimal patent protection

Working Paper, No. 456

Provided in Cooperation with:

Department of Economics, University of Zurich

Suggested Citation: Hémous, David; Lepot, Simon; Sampson, Thomas; Schärer, Julian (2024) : Trade, innovation and optimal patent protection, Working Paper, No. 456, University of Zurich, Department of Economics, Zurich,
<https://doi.org/10.5167/uzh-238643>

This Version is available at:

<https://hdl.handle.net/10419/306548>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series

ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 456

Trade, Innovation and Optimal Patent Protection

David Hémous, Simon Lepot, Thomas Sampson and Julian Schäfer

First version: November 2023
This version: September 2024

Trade, Innovation and Optimal Patent Protection*

David Hémous (U Zürich)

Simon Lepot (U Zürich)

Thomas Sampson (LSE)

Julian Schärer (U Zürich)

September 2024

Abstract

Intellectual property rights are a recurrent source of tensions between developed and developing economies. This paper provides the first quantitative analysis of optimal patent policy in trading economies. We develop a new model of trade, growth and patenting in which patent protection affects both innovation and market power. The model is estimated using data on patent applications to calibrate patent protection by country and the geography of innovation. Counterfactual analysis yields three main results. First, the potential gains from international cooperation over patent policies are large. However, achieving these gains requires more innovative economies to offer stronger protection. Second, only a small share of these gains has been realized so far. And third, by pushing towards policy harmonization, the TRIPS agreement hurts developing countries without generating global welfare gains. Overall, there is substantial scope for policy reforms to increase efficiency.

*We thank Pol Antràs, Giovanni Maggi, Phillip McCalman, Marc Melitz, Monika Mrázová, Fabrizio Zilibotti and many seminar and conference participants for helpful comments and discussions. Hémous is grateful for support from the University Research Priority Program “URPP Equality of Opportunity” of the University of Zürich.

1 Introduction

How should intellectual property be protected in the global economy? This is perhaps the most contentious question in modern trade policy, leading to recurring frictions between the Global North and the Global South that have undermined progress in trade negotiations. Rich countries argue that strong intellectual property rights (IPR) are needed to stimulate innovation and are willing to grant innovators substantial monopoly rights. Poor countries counter that strong IPR inflate consumer prices and amount to a transfer from poor-country households to rich-country firms.

The tensions surrounding the Trade-Related Aspects of Intellectual Property Rights (TRIPS) agreement are an important case in point. TRIPS was the most controversial part of the Uruguay Round negotiations that led to the creation of the World Trade Organization (WTO). It sought to strengthen IPR in developing economies by requiring countries to adopt policies similar to those already implemented in rich countries. TRIPS was pushed through by the United States (US), Europe, and Japan, against strong opposition from Brazil and India. A common assessment is that these tensions drove a lasting wedge between WTO members, contributing to the failure of the Doha Round and, ultimately, the stalemate at the WTO. IPR have continued to be a source of disputes since TRIPS, including as a motivating factor in the current US-China trade war.

Yet, so far, the literature lacks the tools to undertake quantitative analysis of patent policy in economies open to international trade. The goal of our paper is to provide these tools. We develop a new model of trade and patenting that integrates Helpman-Krugman trade and Eaton-Kortum trade with expanding variety growth. The model can be calibrated to patent and trade data, and we use it to conduct counterfactual experiments that shed new light on optimal patent policy. Our analysis implies that the potential gains from international cooperation are large, but to obtain these gains developed countries should provide stronger IPR than developing economies. Moreover, existing patent policies only realize a small share of the cooperative gains and, by pushing towards policy harmonization, TRIPS reduces welfare in the Global South.

The classic trade-off faced by patent policy is that stronger protection brings dynamic benefits from faster innovation, but also generates static costs through higher market power (Nordhaus 1969). Grossman and Lai (2004) build an elegant theory of how this trade-off shapes optimal IPR in a North-South setting. But their model misses important elements needed to connect to the data and developing a framework that embeds the innovation versus market power trade-off in a quantitative model of trade and patenting has proved challenging. Quantitative studies of IPR in open economies such as Eaton and Kortum (1999) and Santacreu (2024) allow intellectual property protection to affect innovation and technology licensing, but not market power.

Our theoretical contribution in Section 2 is to develop a quantitatively tractable model of the innovation-market power trade-off. Newly invented varieties are produced and sold around the

world by monopolists as in Helpman and Krugman (1987). But once technology diffuses, products that are not under patent protection are produced and traded competitively as in Eaton and Kortum (2002). In equilibrium, the split between Helpman-Krugman trade and Eaton-Kortum trade determines the degree of market power and this split depends upon patent policy in all countries. Stronger patent protection increases the expected duration of an innovator's monopoly, which encourages innovation, but also generates market power, leading to higher prices.

The model implies that in open economies stronger patent protection has global benefits, but local costs. While all countries reap dynamic benefits from the increase in product variety brought about by innovation, only households in the country issuing the patent pay the static costs. This is because a patent gives an innovator the exclusive right to sell in a particular market, thus establishing local monopoly power and raising local prices. Since patent policies have cross-border spillovers, there is scope for international policy coordination.

We show that the global growth rate is more sensitive to the level of patent protection in markets that are more profitable for innovators. This means that a country that strengthens IPR generates greater dynamic benefits when it buys a larger share of highly innovative economies' exports. Consequently, optimal patent protection differs across countries and depends on the geographies of innovation and trade integration. Optimal protection tends to be weaker in less innovative countries (due to home bias in trade) and in more isolated countries.

In Section 3 we calibrate the model to the world economy in 2015 divided into twelve countries. Our theory implies that patenting decisions depend upon both market access and IPR. Therefore, we infer the strength of patent protection in each country using variation in cross-border patent applications conditional on observed bilateral trade. The model matches well both targeted and untargeted moments. Our estimates imply that a US patent has an expected duration of 11 years, which captures a combination of the statutory patent length and the probability of patent enforcement. We find that Europe, Japan, Canada and Korea offer similar patent protection to the US, while developing countries, such as China, Brazil and India, provide weaker protection.

We begin our counterfactual analysis in Section 4 by studying countries' incentives to unilaterally change patent protection starting from the 2015 status quo. We calculate welfare effects accounting for transition dynamics between steady states. The results imply that all countries have an incentive to weaken their own patent protection, since the local static costs of protection exceed the dynamic benefits. But when any country reduces its patent protection, welfare falls in all other countries due to lower innovation and growth. Moreover, the magnitude of these international spillovers varies greatly since patent protection in bigger or more innovative countries has a larger effect on growth. The US, Europe and Japan together account for 78 percent of global innovation, implying that growth is especially dependent on patent policy in these economies.

Given these unilateral incentives, we simulate a full breakdown of international cooperation

over patent policies by solving for the Nash equilibrium. We find that no country offers any patent protection in the non-cooperative equilibrium, leading to reductions in both growth and market power. However, the dynamic costs dominate the static gains and welfare declines in all countries. World welfare is 2.1 percent lower in the Nash equilibrium compared to the calibrated steady state. This finding highlights the need for international cooperation.

To study the gains from cooperation, we solve for the patent policies that maximize world welfare. In the baseline case with equal welfare weights for all individuals, the US, Europe, Japan, Canada, Korea and Mexico provide complete patent protection, while other countries do not offer protection. Policy divergence is optimal because it is globally efficient to delegate the task of incentivizing innovation to those countries that are either highly innovative themselves, or that are closely integrated with innovative economies (as Mexico is with the US). By contrast, countries such as China that are large, but not (yet) very innovative, do not provide protection in the cooperative equilibrium. Although stronger protection in these economies raises growth, the static costs of this protection are greater because they impact more people. World welfare is 8.7 percent higher in the cooperative equilibrium, which is around half as large as the total gains from trade relative to autarky in our model. However, these large gains mask substantial distributional effects. Developing countries such as China, Brazil and India that free ride on protection provided elsewhere are the big winners and experience welfare gains of around 10 percent. But the gains are much smaller in countries that bear the static costs of protection. Welfare in the US, Europe, Japan and Mexico increases by only around 2 percent.

Finally, we use our framework to evaluate the welfare effects of TRIPS. Re-calibrating the model using 1992 data, we find that IPR in China, India and Russia were considerably weaker before TRIPS. And counterfactual analysis implies that strengthening patent protection in developing countries, as TRIPS sought to do through IPR harmonization, reduced welfare in those countries. However, this result comes with an important caveat. WTO membership is a single undertaking that includes market access commitments in addition to TRIPS. When we combine TRIPS with observed tariff changes since 1992, we find that the welfare gains from lower tariffs outweigh the costs of TRIPS suggesting that, on net, developing countries gain from WTO membership.

The paper's main insights are robust to moderate variation in the moments used for calibration and to introducing an additional benefit of IPR by allowing patent protection to affect export market entry on the extensive margin. However, this does not mean that optimal patent policies are insensitive to changes in the global economy, at least not if such changes are sufficiently large. For example, we find that lower trade costs exacerbate the incentive to free ride when choosing patent protection, which magnifies the need for international cooperation.

This paper contributes to a number of literatures. The theoretical framework builds upon extensive work on trade, growth and IPR. Helpman (1993) provides the first general equilibrium

analysis of the welfare effects of IPR in a two country endogenous growth model. Grossman and Lai (2004) model countries' incentives to provide patent protection with international trade and, like us, conclude that harmonization of patent regimes is not optimal. Lai and Yan (2013) use a calibrated extension of Grossman and Lai's model to argue that harmonization of patent protection at US levels raises global welfare, but their quantitative application remains illustrative in nature – for instance, they do not calibrate the model to match trade and patent flows. Our contribution to this literature lies in developing a quantitative framework for studying IPR that can be calibrated using trade and patent data.¹ The theory is also related to recent research on international technology diffusion by Lind and Ramondo (2022) and Hsieh et al. (2022) and on optimal innovation policy in open economies by Akcigit et al. (2021) and Borota Milicevic et al. (2023). However, these papers do not study IPR.

Existing quantitative research on IPR in open economies uses models of technology diffusion that build on Eaton and Kortum (1999). In Eaton and Kortum's (1999) model, innovators make profits by licensing abroad once their technologies have diffused internationally. However, patenting does not affect market power and patented products are non-tradable. McCalman (2001) applies the Eaton-Kortum model to quantify how TRIPS reallocates producer surplus across countries holding innovation and market power fixed. McCalman (2005) relaxes these constraints and finds that TRIPS may benefit all countries in the long-run through higher innovation. In contrast to these papers – and motivated by the fact that TRIPS formed part of the WTO's single undertaking – our framework places the interaction between trade and patent policy at centre stage. In addition, we are able to characterize the optimal set of cooperative and non-cooperative patent policies with many countries, which the existing quantitative literature does not do.²

More recently, Santacreu (2024) studies the joint determination of tariffs and IPR in bilateral trade agreements when intellectual property protection affects revenue from technology licensing. And LaBelle et al. (2024) analyze the relationship between globalization, cross-border patenting and development. However, neither of these papers studies the effect of IPR on market power. By quantifying the classic innovation versus market power trade-off, our analysis generates new insights into cross-country heterogeneity in incentives to protect patents, the links between trade integration and patent policy, and the gains from cooperative and non-cooperative policy setting.

Consistent with our model, recent evidence establishes that innovation responds positively to patent protection (Williams 2017). Moscona (2021) finds that the introduction of patent protection for plants in the US in 1985 led to the development of new varieties of affected crops. TRIPS itself

¹To be specific, the main features of our model not found in Grossman and Lai (2004) are: (i) endogenous patenting; (ii) a quantitative trade model featuring trade in both monopolistic and competitive products; (iii) stochastic technology diffusion, and; (iv) endogenous growth with international knowledge spillovers. Each of these features facilitates the model to the data and plays an important role in the counterfactual analysis.

²There is an extensive literature on optimal patent policy in closed economies, e.g. Acemoglu and Akcigit (2012).

generated exogenous variation in the duration of patent protection. Kyle and McGahan (2012) exploit cross-country variation in disease prevalence and patent laws during the implementation of TRIPS. They find that stronger patent protection in developed countries induced more innovation in pharmaceuticals, but there is no significant effect for developing countries. This result is in line with our model, since we find that the impact of patent rights on innovation depends upon countries’ size and innovativeness. Turning to the price effects of patents, Duggan et al. (2016) study the effect of India’s TRIPS-induced patent reform on pharmaceutical prices. They find moderate price increases for molecules that receive a patent.³ Also in line with our model, De Rassenfosse et al. (2022) document a positive association between exports and patenting within firm-product pairs and show that exports decline when firms lose patent protection in a market.

Our paper also fits into an emerging literature on “deep” integration agreements. While “shallow” agreements focus on reducing conventional trade barriers such as tariffs, deep agreements seek to achieve additional economic integration in areas such as investment, regulation, or intellectual property. Recent theoretical contributions on deep integration include Antràs and Staiger (2012), Grossman et al. (2020), Maggi and Ossa (2023), and Ossa et al. (2023). However, none of these papers study IPR.

2 A Theory of Trade and Patents

We develop a dynamic model of trade and patenting with endogenous innovation and market power. Following Nordhaus (1969) and Grossman and Lai (2004) patent protection incentivizes innovation, but creates a static distortion due to monopoly pricing. In addition, greater patent protection generates a sourcing distortion because monopoly control over production restricts buyers’ ability to source from the lowest cost supplier.

2.1 Economic Environment

The economy has N countries and $S + 1$ sectors. Sectors $s \neq 0$ feature endogenous innovation and patenting. In sector zero innovation is exogenous and there is no patenting. Each country n has a fixed labor endowment L_n and labor is the only factor of production. Time t is continuous.

Demand. In each country and sector, non-tradable sectoral output is produced competitively as a constant elasticity of substitution aggregate of tradable intermediate product varieties indexed by ω . Products differ in their quality $\psi(\omega)$, except in sector zero where $\psi(\omega) = 1$ for all varieties. Let

³Relatedly, Chaudhuri et al. (2006) use a structural model to estimate price and expenditure elasticities and quantify the static welfare costs of product patent enforcement under TRIPS in the fluoroquinolones subsegment of the Indian anti-bacterials market.

M_t^s denote the mass of products available in sector s at time t . Since the model does not feature an extensive margin of trade, M_t^s is the same in all countries. Sectoral output Y_{nt}^s then satisfies:

$$Y_{nt}^s = \left(\int_0^{M_t^s} \psi(\omega)^{\frac{1}{\sigma^s}} c_{nt}^s(\omega)^{\frac{\sigma^s-1}{\sigma^s}} d\omega \right)^{\frac{\sigma^s}{\sigma^s-1}}, \quad (1)$$

where $c_{nt}^s(\omega)$ denotes demand for product ω in country n at time t and σ^s is the elasticity of substitution between varieties. Optimization yields a constant elasticity demand function for each variety ω . The output of each sector is combined using a Cobb-Douglas aggregator to produce a non-tradable final good. Let β^s denote the share of sector s in final good production costs.

The final good is used for consumption and as an intermediate input in variety production. Consumption demand comes from each country's representative agent whose intertemporal preferences are given by:

$$U_{nt} = \int_t^\infty e^{-\rho(\tilde{t}-t)} \frac{C_{nt}^{1-1/\gamma}}{1-1/\gamma} d\tilde{t}. \quad (2)$$

In this equation, ρ is the discount rate, γ the elasticity of intertemporal substitution and C_{nt} is aggregate final good consumption in country n . Agents earn income from wages w_{nt} and by investing in a diversified portfolio of domestic firms to obtain risk free return r_{nt} . The representative agent also receives a transfer B_{nt} from the government, where B_{nt} equals tariff revenues. We assume that asset markets are closed and clear at the national level. However, we allow for exogenous trade imbalances. Let TB_{nt} be the trade surplus of country n .

Variety Production. Variety production combines labor with intermediate inputs that are produced one-to-one from the final good. Producers of variety ω in country i have productivity $z_i^s(\omega)$ and produce output $y_i^s(\omega)$ given by:

$$y_i^s(\omega) = \frac{z_i^s(\omega)}{(\alpha^s)^{\alpha^s} (1-\alpha^s)^{1-\alpha^s}} l_i^s(\omega)^{\alpha^s} q_i^s(\omega)^{1-\alpha^s}, \quad (3)$$

where $l_i^s(\omega)$ and $q_i^s(\omega)$ denote the quantities of labor and intermediate inputs, respectively, used to produce variety ω in country i . Labor and intermediate inputs are purchased in competitive markets and the parameter $\alpha^s \in (0, 1)$ equals labor's share of production costs.

Following Eaton and Kortum (2002), productivity is drawn from a country-sector specific Fréchet distribution $F_i^s(z) = \exp(-T_i^s z^{-\theta^s})$. The scale parameter T_i^s captures variables, such as institutions and infrastructure, that affect productivity conditional on which varieties can be produced. The shape parameter $\theta^s > \sigma^s - 1$ is an inverse measure of productivity dispersion across varieties. Productivity draws are independent across varieties and across countries within a variety.

Product varieties are tradable subject to iceberg trade costs and tariffs. To sell one unit of output

in country n , a producer in country i must ship $\tilde{\tau}_{ni}^s$ units. In addition, country n imposes a tariff b_{ni}^s on imports from i in sector s . The tariff is charged on the import price inclusive of iceberg trade costs. Let $\tau_{ni}^s = (1 + b_{ni}^s) \tilde{\tau}_{ni}^s$ denote total trade costs. We assume $\tilde{\tau}_{ii}^s = 1$ and $b_{ii}^s = 0$ for all i .

Innovation and Patenting. In all sectors $s \neq 0$, new products are created by risk neutral innovators using labor. Each worker employed in innovation in country i and sector s successfully innovates at Poisson rate $\eta_i^s (L_{Rit}^s)^{-\kappa}$. The parameter η_i^s determines the efficiency of R&D, while L_{Rit}^s denotes total employment of R&D workers in innovation in sector s and country i . We assume $\kappa \in (0, 1)$ implying that innovation is subject to a stepping-on-the-toes externality whereby the marginal productivity of R&D labor declines as the innovation sector expands. Imposing $\kappa > 0$ also ensures that all countries innovate in all sectors in equilibrium.

Let $\Psi_t^s = \int_0^{M_t^s} \psi(\omega) d\omega$ be the aggregate quality of all varieties produced in sector s . Assume that, when innovation occurs, each invention creates Ψ_t^s new product varieties. This assumption introduces knowledge spillovers into the innovation technology and is sufficient to ensure there is balanced growth in the steady state equilibrium.⁴

There is free entry into innovation. Let V_{it}^s be the expected value of inventing a new variety in country i and sector s at time t . The free entry condition requires that the wage rate equals the product of the probability of innovation, the number of products an invention creates and the expected value of each product. That is:

$$w_{it} = \eta_i^s (L_{Rit}^s)^{-\kappa} \Psi_t^s V_{it}^s. \quad (4)$$

Prior to invention, both product quality $\psi(\omega)$ and the productivity $z_n^s(\omega)$ with which varieties can be produced in each country are unknown. When innovation occurs, the innovator immediately learns the quality of their invention, which is drawn from a Pareto distribution $H(\psi) = 1 - \psi^{-k}$ with shape parameter $k > 1$ and scale parameter 1. All Ψ_t^s products that compose an invention have the same quality. Before commencing production, inventors in each country i also learn the domestic productivity $z_i^s(\omega)$ with which their products can be produced.

Initially, only the innovator knows how to produce its new varieties giving them a *technological* monopoly. However, as technologies are non-rival and imperfectly excludable we assume that technology diffusion occurs at Poisson rate $\nu^s > 0$. Before an invention diffuses, only domestic production in the innovator's home country is possible. After diffusion, any firm in any country can produce the diffused products. Moreover, all firms in a country that produce variety ω have the same country-specific productivity $z_n^s(\omega)$ and the variety has quality $\psi(\omega)$ regardless of where it is produced. Thus, diffusion strips the innovator of its technological monopoly. Innovators may also lose their monopoly due to product obsolescence, which occurs at Poisson rate ζ^s .

⁴Equivalently, one could assume that each invention creates one new product variety and that innovation and patenting costs are inversely proportional to Ψ_t^s .

Anticipating the possibility of technology diffusion, the innovator may also secure a *legal* monopoly over their invention by purchasing a patent. Patents are country-specific and cover all Ψ_t^s product varieties created by an invention. We assume that an inventor who holds a country n patent has the monopoly right to sell varieties covered by the patent to country n .⁵ Once purchased, patent protection is lost at Poisson rate δ_n^s , where δ_n^s is an inverse measure of the strength of patent protection in country n and sector s . The patent protection parameter δ_n^s is a summary measure that captures both the length of protection available and how effectively patent rights are enforced.⁶ An increase in patent protection reduces δ_n^s .

This set-up implies that there are two types of products sold in each destination: Helpman-Krugman products and Eaton-Kortum products. Varieties for which either the technology has not diffused or the inventor holds a patent are Helpman-Krugman products. These varieties are sold by a monopolist inventor who faces constant elasticity demand under monopolistic competition (Helpman and Krugman 1987). Varieties that are not under patent protection and for which technology diffusion has occurred are Eaton-Kortum products. These varieties are produced and sold competitively as in Eaton and Kortum (2002). Because patents are country-specific, whether a variety is a Helpman-Krugman product or an Eaton-Kortum product may differ across destinations.

A successful innovator must choose whether to patent their invention in each country after learning the quality of their invention, but before learning their productivity and commencing production. This restriction reflects the fact that patent law requires patenting to take place before, or very shortly after, a product is commercialized. Consequently, innovators have a strong incentive to file a patent application as soon as possible in order to assert priority over an invention (Dechezleprêtre et al. 2017).⁷ The assumption that innovators know their quality, but not their productivity when patenting captures the idea that inventors are well-informed about the potential of their inventions, but, prior to commercialization, know less about whether an invention is commercially viable or how much it will cost to produce – perhaps because innovators may choose to sell their blueprints rather than producing themselves.

Patenting benefits an innovator by extending the expected duration of their monopoly over an invention. An innovator who purchases a country n patent loses their monopoly in country n only

⁵We assume that patents give monopoly rights over sales, but not over production. This assumption is a useful simplification that makes the quantitative analysis feasible because it implies that patenting decisions are independent across markets. In practice, the value of patents in open economies comes primarily from obtaining a sales monopoly not a production monopoly because an innovator who holds a production monopoly in one country still faces competition from producers in other countries. We show this quantitatively in Section 3.2, see footnote 20.

⁶For example, suppose that a patent issued by country n expires at Poisson rate $\delta_{n,1}^s$, where $1/\delta_{n,1}^s$ is the expected lifetime of the patent. And suppose that the probability a patent issued t periods ago is enforced is $e^{-\delta_{n,2}^s t}$, where $\delta_{n,2}^s$ is an inverse measure of the strength of patent enforcement and the probability of enforcement declines with patent age. Then patent protection is lost at rate $\delta_n^s = \delta_{n,1}^s + \delta_{n,2}^s$.

⁷In general, inventors must file applications at different patent offices within 12 months (or 30 months for applications under the Patent Cooperation Treaty) to have a valid claim to priority.

when both the technology has diffused and patent protection is lost. When choosing whether to patent, the innovator compares this benefit to the costs of patenting. To patent, an innovator from country i must first hire $f_i^{s,o}$ units of domestic labor to prepare their patent application by codifying their invention in terms comprehensible to patent offices. We refer to $f_i^{s,o}$ as the patent preparation cost. This cost need only be paid once, even for inventions that are patented in many countries. Let $L_{it}^{s,o}$ denote total labor employed in the preparation of patent applications.⁸

After paying the patent preparation cost, an inventor can purchase a patent in country n by hiring $f_n^{s,e}$ units of country n labor. This country-specific patenting cost captures the fees a firm pays to submit an application (for example, application fees, translation fees and maintenance fees) and any other costs the firm incurs to make the application (for example, agent payments and internal costs of managing the application process). We refer to $f_n^{s,e}$ as the patent application cost. Let $L_{int}^{s,e}$ denote total labor employed in country n by innovators from country i to cover patent application costs. Payments to these workers are an export of patenting services from n to i .

To complete the specification of the model we return to sector zero. Sector zero is an Eaton and Kortum (2002) sector with no endogenous innovation or patenting. Instead, all varieties are produced competitively and the aggregate quality Ψ_t^0 , which equals the mass of varieties produced M_t^0 , grows exogenously at rate g^0 .

The model incorporates several familiar components: expanding variety growth, Helpman-Krugman trade, Eaton-Kortum trade, product-level heterogeneity and fixed costs of patenting. Our theoretical contribution is to integrate these elements in a quantitative model that captures the trade-off between the static costs and dynamic benefits of patent protection. The key step in achieving this goal is combining Helpman-Krugman and Eaton-Kortum trade while maintaining tractability. The assumption that makes this step possible is that productivity $z_n^s(\omega)$ is independent of quality $\psi(\omega)$ and is unknown at the moment innovators patent, which implies that optimal sourcing choices for Eaton-Kortum varieties are uncorrelated with patenting decisions. This separability reduces the number of state variables and, as shown below, guarantees that trade in goods that are produced competitively follows the Eaton-Kortum model.

The model does not allow for foreign direct investment (FDI). Ruling out FDI implies that firms' market access decisions are independent across countries, which facilitates the quantitative analysis. Previous studies building on the Eaton and Kortum (1999) framework have analyzed the interplay between FDI and patenting in economies where patented products are non-tradable. By contrast, in our model trade flows follow a gravity equation and patenting decisions depend upon trade costs. Consequently, the model is well-suited to studying whether optimal IPR should be – as the name TRIPS suggests – trade-related. Developing a quantitative model with trade, patenting

⁸We include the patent preparation cost in the theory to allow the calibrated model to better fit observed levels of domestic relative to international patenting (see Section 3.3).

and FDI would be an interesting avenue for future research.

2.2 Equilibrium

To solve the model, we start by decomposing the aggregate quality Ψ_t^s produced in sector s by product type. Let Ψ_{Mnit}^s denote the aggregate quality of all Helpman-Krugman products sold monopolistically from country i to destination n and let Ψ_{Cnt}^s be the aggregate quality of all Eaton-Kortum products sold competitively in country n . Then we have that for any destination n :

$$\Psi_t^s = \Psi_{Cnt}^s + \sum_{i=1}^N \Psi_{Mnit}^s. \quad (5)$$

2.2.1 Static Equilibrium

A convenient feature of the model is that the equilibrium conditions can be split into a static equilibrium and a dynamic equilibrium. The static equilibrium solves for wages, output levels, prices and trade flows conditional on knowing for all i , n and s the aggregate quality of products sold competitively Ψ_{Cnt}^s and monopolistically Ψ_{Mnit}^s , total labor employed in output production L_{Yit} and total labor employed in patenting $L_{int}^{s,e}$. The dynamic equilibrium solves for optimal innovation and patenting decisions.

In this section we sketch the static equilibrium. A formal definition and the full set of static equilibrium conditions are in Appendix A.1. Although all variables are time dependent, to simplify notation we henceforth drop the time subscript t except where needed to avoid confusion.

Solving the static equilibrium requires decomposing production and trade into Helpman-Krugman products sold monopolistically and Eaton-Kortum products sold competitively. Start by considering Helpman-Krugman varieties. Monopoly producers face constant elasticity demand with demand elasticity σ^s . Consequently, they charge a mark-up $\sigma^s / (\sigma^s - 1)$ above their marginal cost of serving a market. Using the production function (3) and recalling that exports are subject to iceberg trade costs $\tilde{\tau}_{ni}^s$ and tariffs b_{ni}^s , it follows that the consumer price $p_{ni}^s(\omega)$ of a Helpman-Krugman variety ω produced in i and sold in n satisfies:

$$p_{ni}^s(\omega) = \frac{\sigma^s}{\sigma^s - 1} \frac{\tau_{ni}^s w_i^{\alpha^s} P_i^{1-\alpha^s}}{z_i^s(\omega)}, \quad (6)$$

where P_i denotes the final good price index in country i . Since $z_i^s(\omega)$ is drawn after the patenting decision, the distribution of productivity z (but not quality ψ) is independent of whether varieties are patented. This allows us to aggregate prices and sales across varieties to obtain the tariff-inclusive value X_{Mni}^s of exports of Helpman-Krugman products from i to n :

$$X_{Mni}^s = \frac{\Psi_{Mni}^s (\Phi_{ni}^s)^{\frac{\sigma^s-1}{\theta^s}}}{\sum_{j=1}^N \Psi_{Mnj}^s (\Phi_{nj}^s)^{\frac{\sigma^s-1}{\theta^s}}} \left(\frac{P_{Mn}^s}{P_n^s} \right)^{1-\sigma^s} P_n^s Y_n^s, \quad (7)$$

where P_{Mn}^s is the consumer price index for Helpman-Krugman products sold in country n , P_n^s is the sector s price index in country n and Φ_{ni}^s is the supply potential of country i in country n . The supply potential is inversely related to the average cost of supplying n from i and is given by:

$$\Phi_{ni}^s \equiv T_i^s (\tau_{ni}^s w_i^{\alpha^s} P_i^{1-\alpha^s})^{-\theta^s}. \quad (8)$$

Helpman-Krugman trade is increasing in both the aggregate quality Ψ_{Mni}^s of products sold monopolistically from i to n and the supply potential Φ_{ni}^s of i in n . Substituting equation (8) into equation (7) implies that the elasticity of Helpman-Krugman trade to trade costs τ_{ni}^s equals $\sigma^s - 1$.

Now, consider Eaton-Kortum varieties. As in Eaton and Kortum (2002), each variety is sourced from the lowest cost supplier. Consequently, tariff-inclusive exports X_{Cni}^s of Eaton-Kortum products from i to n satisfy:

$$X_{Cni}^s = \frac{\Phi_{ni}^s}{\sum_{j=1}^N \Phi_{nj}^s} \left(\frac{P_{Cn}^s}{P_n^s} \right)^{1-\sigma^s} P_n^s Y_n^s, \quad (9)$$

where P_{Cn}^s is the consumer price index for Eaton-Kortum products sold in n . It follows that the elasticity of Eaton-Kortum trade to τ_{ni}^s is given by the Fréchet dispersion parameter θ^s .

Using equations (6)–(9), the remaining static equilibrium conditions can be obtained by aggregating Helpman-Krugman with Eaton-Kortum products in each sector and then imposing output market clearing and trade balance conditions (see Appendix A.1 for details).

2.2.2 Dynamic Equilibrium

The dynamic equilibrium solves for R&D investment levels, patenting decisions and how the aggregate quality of each type of product changes over time.

Value of firms and patenting decisions. Consider an innovator in country i and sector s that creates an invention with quality ψ at time t_0 . Let $V_{nit_0}^{s,NP}(\psi)$ denote the expected present discounted value of profits per variety that the innovator's firm makes in destination n if it chooses not to patent in n . A firm that does not patent loses its monopoly when its technology either diffuses (at rate ν^s) or becomes obsolete (at rate ζ^s). Let $\mathbb{E}_z \pi_{ni}^s(\psi, z)$ denote expected profits per variety computed over the distribution of productivity z . Since firms make patenting decisions before learning their productivity, we have:

$$V_{nit_0}^{s,NP}(\psi) = \int_{t_0}^{\infty} \mathbb{E}_z \pi_{ni}^s(\psi, z) \exp\left(-\int_{t_0}^t (r_{it} + \zeta^s + \nu^s) d\tilde{t}\right) dt, \quad (10)$$

By contrast, a firm that patents in n loses its monopoly only when its technology becomes obsolete, or both its patent has expired and its technology has diffused. For a product invented at t_0 , the probability that both diffusion and patent expiration occur before t is $[1 - e^{-\nu^s(t-t_0)}] [1 - e^{-\delta_n^s(t-t_0)}]$. Therefore, the expected present discounted value of profits per variety conditional on patenting $V_{nit_0}^{s,P}(\psi)$ satisfies:

$$V_{nit_0}^{s,P}(\psi) = \int_{t_0}^{\infty} \mathbb{E}_z \pi_{ni}^s(\psi, z) \exp\left(-\int_{t_0}^t (r_{it} + \zeta^s) d\tilde{t}\right) [e^{-\nu^s(t-t_0)} - e^{-(\nu^s+\delta_n^s)(t-t_0)} + e^{-\delta_n^s(t-t_0)}] dt. \quad (11)$$

We can now solve for expected profits $\mathbb{E}_z \pi_{ni}^s(\psi, z)$ by noting that profits per variety are proportional to quality ψ and equal to a fraction $1/\sigma^s$ of sales. Therefore, aggregate profits made by monopolists in i from sales to n are given by $X_{Mni}^s / [\sigma^s (1 + b_{ni}^s)]$ and expected profits are:

$$\mathbb{E}_z \pi_{ni}^s(\psi, z) = \psi \mathbb{E}_z \pi_{ni}^s(1, z) = \frac{\psi}{1 + b_{ni}^s} \frac{X_{Mni}^s}{\sigma^s \Psi_{Mni}^s}. \quad (12)$$

Substituting this expression into equations (10) and (11) implies that the value functions are proportional to quality ψ , i.e. $V_{nit_0}^{s,J}(\psi) = \psi V_{nit_0}^{s,J}(1)$ for $J = NP, P$.

After paying the patent preparation cost $w_i f_i^{s,o}$, a firm patents in country n if the difference between $V_{nit_0}^{s,P}(\psi)$ and $V_{nit_0}^{s,NP}(\psi)$ exceeds the patent application cost per variety $w_n f_n^{s,e} / \Psi^s$. Because value functions are proportional to ψ , this condition defines a quality threshold $\psi_{ni}^{s,e*}$ such that only firms with quality above the threshold patent in n (conditional on having paid the application preparation cost). As ψ is drawn from a distribution with lower bound one, we have:

$$\psi_{ni}^{s,e*} = \max\left(\frac{w_n f_n^{s,e}}{\Psi^s [V_{nit_0}^{s,P}(1) - V_{nit_0}^{s,NP}(1)]}, 1\right). \quad (13)$$

Next, we need to determine which firms pay the patent preparation cost $w_i f_i^{s,o}$. Appendix A.2 shows that there exists a second quality threshold $\psi_i^{s,o*}$ such that only firms with quality above this threshold pay the preparation cost. It follows that firms from country i with quality below $\psi_i^{s,o*}$ do not patent anywhere and that firms patent in country n if and only if $\psi \geq \psi_{ni}^{s,*}$ where:

$$\psi_{ni}^{s,*} = \max(\psi_{ni}^{s,e*}, \psi_i^{s,o*}). \quad (14)$$

Patenting increases the expected duration of an innovator's monopoly, which is more valuable to

higher quality firms. Consequently, the benefits exceed the fixed costs of patenting only for firms with quality above the patenting threshold $\psi_{ni}^{s,*}$.

Using the optimal patenting thresholds, summing across destinations and taking expectations over the quality distribution, Appendix A.2 shows that the value V_{it}^s of inventing a new variety is:

$$V_{it}^s = \sum_{n=1}^N \left\{ \frac{k}{k-1} \left[V_{nit}^{s,NP} (1) \left(1 - (\psi_{ni}^{s,*})^{-k+1} \right) + V_{nit}^{s,P} (1) (\psi_{ni}^{s,*})^{-k+1} \right] - (\psi_{ni}^{s,*})^{-k} \frac{w_n f_n^{s,e}}{\Psi^s} \right\} - (\psi_i^{s,o*})^{-k} \frac{w_i f_i^{s,o}}{\Psi^s}, \quad (15)$$

Thus, V_{it}^s equals the sum over all markets n of the expected value from innovations that are not patented plus the expected value from patented innovations less expected patenting costs.

Laws of motion for aggregate qualities. We can decompose the aggregate quality Ψ_{Mni}^s of Helpman-Krugman products sold from i to n into the aggregate quality of products that are not patented $\Psi_{Mni}^{s,NP}$, the aggregate quality of products that are patented but whose technology has not diffused $\Psi_{Mni}^{s,P,ND}$ and the aggregate quality of products that are patented and whose technology has diffused $\Psi_{Mni}^{s,P,D}$. We have:

$$\Psi_{Mni}^s = \Psi_{Mni}^{s,NP} + \Psi_{Mni}^{s,P,ND} + \Psi_{Mni}^{s,P,D}. \quad (16)$$

Together with the aggregate quality of Eaton-Kortum products Ψ_{Cn}^s , these aggregate qualities compose the state variables of the economy. To solve for the dynamic equilibrium, we need to characterize how they evolve over time.

The law of motion for the aggregate quality of Helpman-Krugman products that are not patented $\Psi_{Mni}^{s,NP}$ is given by:

$$\dot{\Psi}_{Mni}^{s,NP} = \eta_i^s (L_{Ri}^s)^{1-\kappa} \Psi^s \frac{k}{k-1} \left[1 - (\psi_{ni}^{s,*})^{1-k} \right] + \delta_n^s \Psi_{Mni}^{s,P,ND} - (\nu^s + \zeta^s) \Psi_{Mni}^{s,NP}. \quad (17)$$

The first term on the right hand side of this expression gives the aggregate quality of new goods invented in country i and sector s that are not patented in country n . There are L_{Ri}^s R&D workers each of whom innovates at rate $\eta_i^s (L_{Ri}^s)^{-\kappa}$ and each innovation produces Ψ^s new varieties. Innovations with quality below $\psi_{ni}^{s,*}$ are not patented in country n , implying that a unit mass of innovations contributes aggregate quality $\int_1^{\psi_{ni}^{s,*}} \psi dH(\psi) = \frac{k}{k-1} \left[1 - (\psi_{ni}^{s,*})^{1-k} \right]$ to $\Psi_{Mni}^{s,P,ND}$. Combining these observations yields the first term. The second term gives the increase in $\Psi_{Mni}^{s,NP}$ due to patent expiration among Helpman-Krugman varieties whose technology has not diffused. And the third term captures the decline in $\Psi_{Mni}^{s,NP}$ due to technology diffusion and product obsolescence.

Analogous reasoning gives the laws of motion for the other state variables. Helpman-Krugman products that are patented, but whose technology has not diffused are generated by patenting and

lost due to patent expiration, technology diffusion and product obsolescence, which yields:

$$\dot{\Psi}_{Mni}^{s,P,ND} = \eta_i^s (L_{Ri}^s)^{1-\kappa} \Psi^s \frac{k}{k-1} (\psi_{ni}^{s*})^{1-k} - (\delta_n^s + \nu^s + \zeta^s) \Psi_{Mni}^{s,P,ND}. \quad (18)$$

Helpman-Krugman products that are patented and whose technology has diffused are generated by technology diffusion and lost due to patent expiration and product obsolescence, implying:

$$\dot{\Psi}_{Mni}^{s,P,D} = \nu^s \Psi_{Mni}^{s,P,ND} - (\delta_n^s + \zeta^s) \Psi_{Mni}^{s,P,D}. \quad (19)$$

Finally, Eaton-Kortum products are generated by either technology diffusion among not patented products or patent expiration among products whose technology has already diffused. And Eaton-Kortum products are destroyed by product obsolescence. Therefore:

$$\dot{\Psi}_{Cn}^s = \sum_{i=1}^N \left(\nu^s \Psi_{Mni}^{s,NP} + \delta_n^s \Psi_{Mni}^{s,P,D} \right) - \zeta^s \Psi_{Cn}^s. \quad (20)$$

Combining these laws of motion with the patenting thresholds and firm value functions above and imposing labor market clearing gives the dynamic equilibrium, see Appendix A.2 for details.

Let g^s denote the growth rate of aggregate quality Ψ^s in sector s . Using equations (5) and (16) to decompose the growth rate of Ψ^s in terms of the growth rates of Eaton-Kortum and Helpman-Krugman products and then combining equations (17)–(20) implies that in any equilibrium:

$$g^s = \sum_{i=1}^N \frac{k}{k-1} \eta_i^s (L_{Ri}^s)^{1-\kappa} - \zeta^s, \quad \text{for } s \neq 0. \quad (21)$$

This expression shows how growth in aggregate quality depends upon R&D employment in all N countries. The first term on the right hand side captures the contribution of innovation to growth. Innovations occur at rate $\eta_i^s (L_{Ri}^s)^{1-\kappa}$ in country i and $k/(k-1)$ is the average quality of an innovation. The second term captures the decline in aggregate quality due to product obsolescence.

2.3 Steady State

We define a steady state of the global economy as a balanced growth path equilibrium in which all aggregate and industry-level variables grow at constant rates. This section describes the main features of a steady state, while Appendix A.3 provides further details.

Knowledge spillovers in the innovation and patenting technologies are global in scope. Consequently, steady state growth rates do not vary by country. The aggregate qualities of all product types grow at rate g^s , which leads to growth in wages w_i , final good output Y_i , trade flows X_{ni}^s and final consumption C_i at rate g given by:

$$g = \frac{1}{\sum_{s=0}^S \beta^s \alpha^s} \sum_{s=0}^S \frac{\beta^s}{\sigma^s - 1} g^s, \quad (22)$$

Computing the integrals in equation (10) gives that the steady state value of a variety with quality one that is not patented satisfies:

$$V_{nit}^{s,NP}(1) = R^{s,NP} \mathbb{E}_z \pi_{nit}^s(1, z) \text{ where } R^{s,NP} \equiv \frac{1}{r + \zeta^s + \nu^s - g + g^s}. \quad (23)$$

The term $R^{s,NP}$ is the inverse of the firm's effective discount rate, which equals the interest rate r augmented for obsolescence ζ^s and technology diffusion ν^s and reduced by $g - g^s$ to account for future profit growth. Similarly, equation (11) implies that the value of a patented variety is:

$$V_{nit}^{s,P}(1) = (R^{s,NP} + \Delta R_n^s) \mathbb{E}_z \pi_{nit}^s(1, z), \quad (24)$$

where $\Delta R_n^s \equiv \frac{1}{r + \zeta^s + \delta_n^s - g + g^s} - \frac{1}{r + \zeta^s + \nu^s + \delta_n^s - g + g^s} > 0$. Patenting reduces the firm's effective discount rate by extending the expected duration of its monopoly. Consequently, its valuation of future profit flows increases by ΔR_n^s , which we will refer to as the benefit of patenting in n . Stronger patent protection increases ΔR_n^s by reducing δ_n^s . The benefit of patenting is also increasing in the rate of technology diffusion ν^s , implying that patenting is complementary to technology diffusion. The complementarity arises because the probability that patent protection is needed to maintain the firm's monopoly is greater when technology diffusion is faster. Indeed, if $\nu^s = 0$ then $\Delta R_n^s = 0$ and firms have no incentive to patent. Intuitively, ΔR_n^s is also decreasing in r and ζ^s , but increasing in the growth rate of profits $g - g^s$.

To characterize the steady state equilibrium we detrend all variables to obtain normalized variables that are constant in steady state, which we denote using tildes. We normalize variables that grow at rate g^s by writing them relative to Ψ^s and normalize variables that grow at rate g by writing them relative to $\Psi \equiv \left[\prod_{s=0}^S (\Psi^s)^{\frac{\beta^s}{\sigma^s - 1}} \right]^{\frac{1}{\sum_{s=0}^S \beta^s \alpha^s}}$. Thus, $\tilde{\Psi}_{Cn}^s = \Psi_{Cn}^s / \Psi^s$ and $\tilde{w}_i = w_i / \Psi$, for example. Likewise, we normalize variables that grow at rate $g - g^s$, such as profits and value functions, by writing them relative to Ψ / Ψ^s . In particular, we define normalized expected profits as:

$$\tilde{\pi}_{ni}^s \equiv \Psi^s \frac{\mathbb{E}_z \pi_{ni}^s(1, z)}{\Psi} = \frac{1}{1 + b_{ni}^s \sigma^s} \frac{\tilde{X}_{Mni}^s}{\tilde{\Psi}_{Mni}^s}, \quad (25)$$

where the equality uses equation (12). This allows us to write the patenting threshold $\psi_{ni}^{s,e*}$ as:

$$\psi_{ni}^{s,e*} = \max \left(\frac{\tilde{w}_n f_n^{s,e}}{\Delta R_n^s \tilde{\pi}_{ni}^s}, 1 \right). \quad (26)$$

A higher patenting cost $\tilde{w}_n f_n^{s,e}$ increases the patenting threshold. By contrast, an increase in either

the benefit of patenting ΔR_n^s or normalized profits $\tilde{\pi}_{ni}^s$ reduces the patenting threshold. Appendix A.3 shows that the patent preparation threshold $\psi_i^{s,o*}$ satisfies a similar expression.

Using equation (15), the normalized expected value of inventing a new variety \tilde{V}_i^s is:

$$\tilde{V}_i^s = \sum_{n=1}^N \left[\frac{k}{k-1} \tilde{\pi}_{ni}^s \left(R_n^{s,NP} + \Delta R_n^s (\psi_{ni}^{s*})^{1-k} \right) - \tilde{w}_n f_n^{s,e} (\psi_{ni}^{s*})^{-k} \right] - \tilde{w}_i f_i^{s,o} (\psi_i^{s,o*})^{-k}. \quad (27)$$

The expected value of invention comprises four terms. The first term gives the expected value if there is no patenting. The second term captures the additional value that arises because firms have the opportunity to patent their inventions. The value that patenting creates is increasing in profitability $\tilde{\pi}_{ni}^s$ and in the benefit of patenting ΔR_n^s , but decreasing in the patenting threshold ψ_{ni}^{s*} . The final two terms in equation (27) give the expected patenting costs a firm pays.

Finally, equation (2) implies that steady state welfare is given by:

$$U_{nt} = \frac{\Psi_t^{1-1/\gamma}}{1-1/\gamma} \frac{\tilde{C}_n^{1-1/\gamma}}{\rho - g \left(1 - \frac{1}{\gamma} \right)}. \quad (28)$$

Conditional on the initial value Ψ_t , an increase in either normalized consumption \tilde{C}_n or growth g raises steady state welfare in country n . The trade-off between static costs and dynamic benefits of patent protection arises when stronger protection raises growth g , but reduces consumption \tilde{C}_n .

2.4 Patent Protection and Welfare

To develop more intuition about the channels through which patent protection affects welfare, we characterize the direct effect of patent policy δ_n^s in country n on steady state outcomes in all countries, without allowing for general equilibrium adjustments.

We start by characterizing the static effect on normalized consumption levels \tilde{C}_n . The laws of motion for aggregate qualities imply that the share of aggregate quality $\tilde{\Psi}_{Mni}^s$ sold in country n and sector s that is supplied monopolistically by producers from country i satisfies:

$$\tilde{\Psi}_{Mni}^s = \frac{k}{k-1} \eta_i^s (L_{Ri}^s)^{1-\kappa} \left[\frac{1}{g^s + \nu^s + \zeta^s} + \frac{(\psi_{ni}^{s*})^{1-k}}{g^s + \delta_n^s + \zeta^s} \frac{\nu^s}{g^s + \delta_n^s + \nu^s + \zeta^s} \right], \quad (29)$$

implying that a reduction in δ_n^s increases $\tilde{\Psi}_{Mni}^s$, all else constant. Since this relationship holds for any i , it follows that a fall in δ_n^s directly reduces the share of aggregate quality sold competitively in country n , $\tilde{\Psi}_{Cn}^s = 1 - \sum_{i=1}^N \tilde{\Psi}_{Mni}^s$. Thus, stronger patent protection directly increases the market

power of suppliers to country n .⁹ And this change in market power affects consumption levels through its impact on prices, profits and real wages. We next consider each of these effects in turn.

Market power creates two distortions that raise prices: a mark-up distortion and a sourcing distortion. The mark-up distortion arises because monopolists set a mark-up $\sigma^s / (\sigma^s - 1)$ above marginal cost. The sourcing distortion arises because country n sources Eaton-Kortum products from its lowest cost supplier, whereas Helpman-Krugman products can only be sourced from the monopolist's home country, which is not necessarily the lowest cost supplier.¹⁰ Differentiating the sectoral price index P_n^s with respect to $\tilde{\Psi}_{Mni}^s$ and accounting for the decline in $\tilde{\Psi}_{Cn}^s$ yields:

$$\frac{\partial P_n^s}{\partial \tilde{\Psi}_{Mni}^s} \propto 1 - \underbrace{\left(\frac{\sigma^s - 1}{\sigma^s} \right)^{\sigma^s - 1}}_{\text{Mark-up distortion}} \underbrace{\frac{\left(\tilde{\Phi}_{ni}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}}{\left(\sum_{j=1}^N \tilde{\Phi}_{nj}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}}}_{\text{Sourcing distortion}} > 0. \quad (30)$$

Thus, the increase in $\tilde{\Psi}_{Mni}^s$ due to stronger patent protection in n raises the domestic price level P_n^s . Moreover, the price increase is greater when the mark-up is higher (i.e. σ^s is lower) and when the supply potential of i in n $\tilde{\Phi}_{ni}^s$ is low relative to the supply potential of other countries.

In addition to raising prices, market power generates profits for innovators. Aggregate normalized profits made by innovators from i in country n and sector s are given by:

$$\frac{1}{1 + b_{ni}^s} \frac{\tilde{X}_{Mni}^s}{\sigma^s} = \frac{1}{1 + b_{ni}^s} \Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \frac{(\sigma^s - 1)^{\sigma^s - 1}}{(\sigma^s)^{\sigma^s}} \beta^s \tilde{\Psi}_{Mni}^s \left(\tilde{\Phi}_{ni}^s \right)^{\frac{\sigma^s - 1}{\theta^s}} \left(\tilde{P}_n^s \right)^{\sigma^s - 1} P_n \tilde{Y}_n,$$

showing that the direct effect of stronger patent protection in n is to increase the profits that all countries make in n by raising both $\tilde{\Psi}_{Mni}^s$ and \tilde{P}_n^s . The increase in profits is greater when country n has higher final good expenditure $P_n \tilde{Y}_n$ and when trade costs for exports from i to n are lower.

The level of normalized consumption \tilde{C}_i in each country i is affected by both price distortions and profit levels. Using the static equilibrium trade balance and market clearing conditions yields:

$$\tilde{C}_i = \frac{\tilde{w}_i}{P_i} \left(L_{Yi} + \sum_{s=1}^S \sum_{n=1}^N L_{ni}^{s,e} - \sum_{s=1}^S \sum_{n=1}^N \frac{\tilde{w}_n}{\tilde{w}_i} L_{in}^{s,e} \right) + \frac{\tilde{\Pi}_i}{P_i} - \frac{\tilde{T}B_i}{P_i} + \frac{\tilde{B}_i}{P_i}$$

where $\tilde{\Pi}_i \equiv \sum_{s=1}^S \sum_{n=1}^N \tilde{X}_{Mni}^s / [\sigma^s (1 + b_{ni}^s)]$ denotes aggregate normalized profits made by country i . This expression decomposes consumption into terms that depend upon the real wage \tilde{w}_i / P_i ,

⁹For given profits, wages, growth rates and patenting thresholds, a decline in δ_n^s also increases ΔR_n^s and, consequently, R&D employment L_{Ri}^s , which raises $\tilde{\Psi}_{Mni}^s$ by equation (29) leading to a further increase in market power.

¹⁰Helpman (1993) refers to the sourcing distortion as the ‘‘production composition’’ effect since it arises from an inefficient allocation of production across countries.

real profits $\tilde{\Pi}/P_i$, trade imbalances and tariff revenues. In turn, the real wage can be written as:

$$\frac{\tilde{w}_i}{P_i} = \left\{ \prod_{s=0}^S \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \right]^{\frac{\beta^s}{\sigma^s - 1}} \left(\frac{T_i}{\lambda_{Cii}^s} \right)^{\frac{\beta^s}{\theta^s}} \left(\frac{\tilde{\Psi}_{Ci}^s}{\mu_{Ci}^s} \right)^{\frac{\beta^s}{\sigma^s - 1}} \right\}^{\frac{1}{\sum_{s=0}^S \alpha^s \beta^s}}, \quad (31)$$

where $\lambda_{Cii}^s \equiv \tilde{X}_{Cii}^s / \sum_{j=1}^N \tilde{X}_{Cij}^s = \tilde{\Phi}_{Cii}^s / \sum_{j=1}^N \tilde{\Phi}_{Cij}^s$ denotes the share of expenditure on Eaton-Kortum products in country i and sector s that is allocated to domestic production, while μ_{Ci}^s denotes the expenditure share of Eaton-Kortum products in sector s and country i . When all products are sold competitively $\tilde{\Psi}_{Ci}^s = \mu_{Ci}^s = 1$ and equation (31) reduces to a multi-sector version of the Arkolakis et al. (2012) formulation of the gains from trade.

But in our model, real wages depend not only upon the domestic trade share for competitive products λ_{Cii}^s , but also upon the ratio of the share of aggregate quality supplied competitively $\tilde{\Psi}_{Ci}^s$ to the expenditure share of competitive products μ_{Ci}^s . Because of the pricing distortions for monopolistic products this ratio is less than one, meaning that expenditure on Eaton-Kortum products exceeds their share of aggregate quality and that real wages are lower than when all products are supplied competitively. In fact, we have:

$$\frac{\tilde{\Psi}_{Ci}^s}{\mu_{Ci}^s} = 1 - \sum_{j=1}^N \tilde{\Psi}_{Mij}^s \left[1 - \left(\frac{\sigma^s - 1}{\sigma^s} \right)^{\sigma^s - 1} \frac{\left(\tilde{\Phi}_{ij}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}}{\left(\sum_{k=1}^N \tilde{\Phi}_{ik}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}} \right], \quad (32)$$

implying that real wages are decreasing in the share of aggregate quality supplied monopolistically by each exporter j , $\tilde{\Psi}_{Mij}^s$. Moreover, comparing this expression with equation (30) shows that the impact of an increase in $\tilde{\Psi}_{Mij}^s$ on real wages is greater when the pricing distortions are larger. Equations (31) and (32) formalize how static inefficiencies reduce real wages in this economy. By increasing market power, stronger patent protection exacerbates these inefficiencies.

Putting everything together, we can now characterize the direct effect of a reduction in δ_n^s on steady state normalized consumption \tilde{C}_i in each country i . For countries $i \neq n$, the only direct effect is a rise in profits due to an increase in the share of aggregate quality $\tilde{\Psi}_{Mni}^s$ supplied to country n monopolistically. This means that stronger domestic patent protection generates positive direct spillovers to foreign countries by giving their suppliers greater market power.

However, in country n itself, increased market power raises not only profits, but also prices. And these higher prices reduce real wages. Thus, the direct cost of the price distortions generated by increased monopoly power is borne by country n itself. Section 4 quantifies these channels allowing for general equilibrium adjustments and transition dynamics between steady states in addition to the direct effects characterized in this section.

Next, we analyze the dynamic effect of stronger patent protection on growth. For simplicity,

we use a version of the model where $f_i^{s,o} = 0$, meaning that there are no patent preparation costs. Taking the partial derivative of equations (21), (26), (27) and the innovation free entry condition with respect to δ_n^s , while holding \tilde{w}_i , $\tilde{\pi}_{mi}^s$ and $R^{s,NP}$ constant for all countries i and m , implies that the direct effect of stronger patent protection on the sectoral growth rate g^s is given by:

$$\frac{\partial \ln g^s}{\partial \ln \delta_n^s} = \frac{\partial \ln \Delta R_n^s}{\partial \ln \delta_n^s} k \frac{1-\kappa}{\kappa} \underbrace{\sum_{i=1}^N \frac{\frac{k}{k-1} \eta_i^s (L_{Ri}^s)^{1-\kappa}}{g^s}}_{\text{Contribution of } i \text{ to global innovation}} \underbrace{\frac{\tilde{V}_{ni}^s}{\tilde{V}_i^s}}_{\text{Contribution of } n \text{ to value of innovation in } i} \underbrace{\frac{\Delta R_n^s (\psi_{ni}^{s*})^{1-k}}{k R^{s,NP} + \Delta R_n^s (\psi_{ni}^{s*})^{1-k}}}_{\text{Contribution of patent protection to value of supplying } n}. \quad (33)$$

The growth elasticity is negative because ΔR_n^s is decreasing in δ_n^s . Therefore, stronger patent protection has a positive direct effect on growth. Intuitively, this occurs because stronger protection raises the returns to innovation by extending the expected duration of an innovator's monopoly, which leads to increased R&D.

The decomposition on the right hand side of equation (33) shows that the effect of δ_n^s on growth is greater when n accounts for a larger share of the value of innovating. That is, when $\tilde{V}_{ni}^s/\tilde{V}_i^s$ is higher. Consequently, growth is more sensitive to patent protection provided by larger economies, since larger economies generate higher profits for innovators.

In addition, the effect of δ_n^s on growth is greater when $\tilde{V}_{ni}^s/\tilde{V}_i^s$ is higher for countries i that innovate more, i.e. countries where $\eta_i^s (L_{Ri}^s)^{1-\kappa}$ is larger. It follows that growth is more sensitive to patent protection in destinations that are closely integrated with more innovative countries. For example, home bias in trade implies that each country contributes more to the value of innovation for domestic innovators than for foreign innovators, i.e. $\tilde{V}_{nn}^s/\tilde{V}_n^s$ exceeds $\tilde{V}_{ni}^s/\tilde{V}_i^s$ for $i \neq n$. Therefore, all else equal, growth is more sensitive to δ_n^s when country n contributes a greater share of global innovation. This observation will play an important role in the quantitative analysis.

Finally, we note that, whereas the static costs of patent protection are domestic in scope, the dynamic benefits are global because knowledge spillovers are international, meaning that the steady state growth rate is the same in all countries. This contrast generates an incentive for countries acting unilaterally to choose weaker patent protection than is globally optimal.

3 Model Calibration

We calibrate the model's steady state to fit the world economy in 2015. The calibration uses patent data to estimate how patent protection varies across countries. In the model, patenting

thresholds depend upon both patent policy and market access (see equation 26). To separate these two channels, the calibration chooses trade costs to exactly match trade flows and then uses data on cross-border patent applications to calibrate the strength of patent protection δ_n^s in each country. This approach ensures the calibrated model fits bilateral data on both trade and patenting, which we argue is a pre-requisite for quantitative analysis of patent policy in open economies.

The calibrated model has two sectors, i.e. $S = 1$, one sector with patenting and one without. We map the Manufacturing, Information, and Professional, scientific and technical services industries to the patenting sector. These industries accounted for 93 percent of US patent applications in 2008 (NSF 2013) while producing 31 percent of gross output (BEA 2022). All other industries are mapped to the no patenting sector.

Countries are aggregated into $N = 12$ economies: US, Europe, Japan, China, Brazil, India, Canada, Korea, Russia, Mexico, South Africa and the rest of the world. This sample of countries includes the ten largest economies (by real GDP in 2015) for which the patent data we use is available.¹¹ We also add South Africa, so that the sample includes all the BRICS economies. Our Europe aggregate comprises the 32 countries that are members of the European Patent Office and are also included in the OECD’s Input-Output Tables (OECD 2021). Collectively, the eleven economies not included in the rest of the world aggregate account for 78 percent of global real GDP and 89 percent of global R&D expenditure in 2015 (World Bank 2023).

3.1 Data

We obtain data on patent applications from PATSTAT (2022) and WIPO (2023). We use PATSTAT to group applications into patent families that cover the same invention. We also observe the origin country for each patent family and the destination of each application. This allows us to measure the flow of applications at the patent family level between each pair of countries in our sample. Appendix B.1 provides further details about how we measure bilateral patent flows.

Our data shows that international patent flows in 2015 mostly originate in richer economies. US, Europe and Japan together are the origins of 75 percent of cross-border flows, whereas China is the origin of only 6 percent of flows and less than one percent of flows come from India or Brazil. By contrast, the destination of international flows is more dispersed and less correlated with income levels. The US is the most popular destination accounting for 32 percent of flows, followed by China with 23 percent. Europe, Japan, India, Canada and Korea are each the destination for between 5 and 10 percent of flows.

Data on trade, output, expenditure and intermediate input costs are from the OECD’s Input-Output Tables 2021 (OECD 2021). Tariff data is from Caliendo et al. (2023). Country-level GDP,

¹¹Indonesia is one of the ten largest economies, but is omitted from our sample due to poor quality patent data.

working age population, R&D expenditure, price level data and GDP deflator data are from the World Development Indicators (World Bank 2023). And we obtain sectoral price index and gross output data for the US from the Bureau of Economic Analysis (BEA 2022).

3.2 Calibration

We start by setting some parameters to values from the prior literature and choosing others to exactly match selected data moments. We then jointly estimate the remaining parameters using simulated method of moments. We take particular care to target moments that are closely related to the parameters we estimate. This section describes the moments we use and how we implement the calibration. Appendix B.2 provides further details on how we measure moments in the data. Appendix B.3 explains how we calculate moments in the model.

A first set of parameters ρ , γ , κ , σ^1 and θ^0 are standard and chosen based on previous work. Drawing on Acemoglu et al. (2018), we let the discount rate $\rho = 0.02$, the intertemporal elasticity of substitution $\gamma = 0.5$, the concavity of the innovation technology $\kappa = 0.5$ and the demand elasticity $\sigma^1 = 2.9$. Setting $\kappa = 0.5$ is consistent with evidence that the elasticity of R&D expenditure to R&D costs $(1 - \kappa) / \kappa$ is around one (Bloom et al. 2019). The calibrated level of σ^1 implies that the mark-up ratio $\sigma^1 / (\sigma^1 - 1)$ for monopoly producers in the patenting sector equals 1.53. This value is similar to De Loecker et al.’s (2020) estimate of the median mark-up in US manufacturing in 2012. We set the trade elasticity in the no patenting sector θ^0 to five (Head and Mayer 2014).

We choose TB_i , β^s , α^s , $\tilde{\tau}_{ni}^s$, b_{ni}^s and L_i to exactly match: trade imbalances relative to world output; sectoral expenditure shares; sectoral value-added to output ratios; bilateral trade shares X_{ni}^s / X_{nn}^s , and; the working age population, respectively. We also parameterize patenting costs as $f_n^o = f^o h_n$ and $f_n^e = f^e h_n$. In the model, patenting costs are denominated in labor, which implies large cost differences between high and low wage countries. However, measures of patent application costs are not strongly correlated with income levels (Park 2010, De Rassenfosse and Van Pottelsberghe 2013), likely because patenting costs reflect wages for skilled workers, which vary less across countries than average wages (Hjort et al. 2022). To capture this feature of the data, we set $h_n = (\text{Real GDP per capita in US} / \text{Real GDP per capita in } n)^{1-\chi}$, implying that the elasticity of patenting costs to real GDP per capita equals χ . We set $\chi = 0.16$ based on the estimated elasticity of real middle management costs to real GDP per capita in Hjort et al. (2022).¹²

This leaves $2NS + N(S + 1) + 3S + 4 = 55$ parameters: δ_n^1 , $f^{1,e}$, $f^{1,o}$, η_i^1 , T_i^s , ν^1 , ζ^1 , θ^1 , k and

¹²Given the challenges of obtaining comprehensive measures of patenting costs that cover both application fees and the labor costs firms incur during the application process, and the fact that available measures differ considerably by source, we choose not to use international variation in estimates of patent application costs to calibrate f_n^e . In practice, this means that our calibration will load unmodeled cross-country differences in patenting costs onto the patent protection parameters δ_n . We adjust f_{Europe}^e upwards to account for the fact that applicants must pay patent fees in multiple countries to obtain patent protection in Europe (see Appendix B.2 for details).

g^0 . We calibrate these parameters by simulated method of moments estimation using 169 moments that capture information on patent applications, the value and costs of patenting, R&D expenditure, growth, prices, production and the trade elasticity. To simplify notation, we henceforth drop the one superscript from parameters that only apply to the patenting sector, e.g. δ_n^1 becomes δ_n .

A series of moments allow us to identify patenting and innovation parameters. From the patent data we calculate the patent flow PAT_{ni} from i to n defined as the number of applications belonging to distinct patent families in destination n by applicants from country i in 2015. For each pair of economies with $n \neq i$, we target international patent shares, defined as the ratio of PAT_{ni} to total international patents $\sum_{n=1, n \neq i}^N \sum_{i=1}^N PAT_{ni}$. The $N(N-1) = 132$ international patent shares are the most important moments in our estimation. Countries with stronger patent protection receive more inward flows, all else equal. Consequently, conditional on patenting costs and trade flows, variation in applications by destination allow us to infer how the strength of patent protection δ_n differs across countries. At the same time, differences in applications by origin are informative about relative values of the R&D efficiency parameter η_i , since the flow of applications depends on relative levels of innovation in each origin.

In addition to cross-border flows, we include four other patenting moments. Kogan et al. (2017) use stock market responses to news about patent grants to estimate the private value of patenting. Averaging over 1995-2007 they estimate that the aggregate increase in the value of innovations from patenting amounts to 9.3 percent of R&D expenditure in the US. Targeting this moment allows us to infer the level of patent protection in the US δ_{US} . We target total US expenditure on domestic patent applications, which we compute by multiplying observed PAT_{nn} for the US by the estimated cost of a US patent application from Park (2010). This moment allows us to calibrate dispersion in the quality of innovations, k . And we target two moments that are useful in pinning down the patenting costs f^o and f^e : the share of US innovations that are patented, and; the share of domestic patents in total inward patents for the US, $PAT_{nn} / \sum_{i=1}^N PAT_{ni}$ with $n = US$.¹³

Data on R&D expenditures disciplines the allocation of resources to innovation, which depends on the level of the R&D efficiency parameters η_i . We target the ratios of R&D expenditure to GDP in the developed countries (US, Europe, Japan, Canada and Korea). But we do not use R&D data from developing economies, since measured R&D in countries far from the technological frontier is frequently used for knowledge absorption, rather than innovation (Griffith et al. 2004).¹⁴

To identify the remaining parameters we use several additional moments. We target turnover in the origin of US imports, which depends on the technology diffusion rate ν as in Hsieh et al.

¹³In order to reduce measurement error resulting from differences across patent offices in the average scope of a patent, we do not use information on domestic patenting in any countries other than the US. Dechezleprêtre et al. (2017) show that international patents are more comparable across patent offices than domestic patents.

¹⁴Consistent with this observation, the calibrated model under-predicts R&D expenditure relative to GDP in China, Brazil, India, Russia, Mexico, South Africa and the rest of the world.

(2022). We target aggregate growth g and the difference between price growth in the patenting and no patenting sectors. Conditional on innovation and patenting, these moments pin down the obsolescence rate ζ and the exogenous growth rate in sector zero g^0 . We target an average trade elasticity in the patenting sector equal to five (Head and Mayer 2014), which depends on θ^1 and σ^1 . Finally, we target world gross output and each economy's price index relative to the US and share of world real GDP. These moments are informative about the Fréchet scale parameters T_i^s .

Let Ω denote the set of parameters that we estimate using the simulated method of moments and K the set of targeted moments. Using m^k to denote moment k with dimension $\dim(m^k)$ and elements m_i^k that have target values $m_i^{k,target}$, the objective function we minimize is:

$$F(\Omega) = \sum_{k \in K} \sum_{i=1}^{\dim(m^k)} \left[\frac{v_k}{\sqrt{\dim(m^k)} \sum_{j \in K} v_j} \mathcal{L}^k \left(m_i^k(\Omega), m_i^{k,target} \right) \right]^2,$$


where v_k is the weight given to moment k , $m_i^k(\Omega)$ denotes the simulated value of element i of moment k and $\mathcal{L}^k(\cdot)$ is a loss function. Appendix B.4 describes the algorithm we use to solve for the model's steady state conditional on knowing Ω . Appendix B.5 provides further details about the calibration procedure we use to estimate the parameters in Ω . And Appendix D analyzes the sensitivity of the calibration to varying the targeted moments (see especially Figure 8).

3.3 Model Fit and Calibrated Parameters

Figure 1 and Table 1 report how the calibrated model matches targeted moments. Figure 1, panel (a) plots international patent shares implied by the model against their observed values. The model performs well in matching both cross-country and within-country variation in patenting shares. Regressing the log of the observed shares against their model-implied counterparts yields an elasticity of 0.95 with a standard error of 0.02 and an R-squared of 0.95.

Notably, the calibrated model matches observed variation in bilateral patenting flows without assuming any country-pair specific differences in patenting costs or patent protection. And although lower bilateral trade costs increase both trade and patenting, the fit evident in Figure 1 is not simply the result of trade and patent flows following the same gravity equation. Figure 1, panel (b) plots international trade shares (which the model matches exactly) and model-implied international patent shares against observed patent shares. We see that the match between model-implied and observed patent shares is much better than the match between trade and patent shares, implying that the model successfully captures how patent flows deviate from trade flows.

Table 1 shows that the model does a good job in matching the remaining targeted moments. The most notable discrepancy between the targeted and model-implied moments is that the model over-predicts the share of domestic patents in inward patents in the US by eleven percentage points.



Notes: Panel (a) shows model-implied versus observed international patent shares. Points labelled: (destination, origin). Panel (b) shows international trade shares and model-implied international patent shares plotted against observed international patent shares. Dashed line is 45 degree line.

Figure 1: International patent shares

There is a tension in the estimation between this moment, the share of innovations patented in the US and expenditure on domestic patent applications in the US, since all three moments are increasing in domestic patenting. The calibration resolves this tension by matching the share of innovations patented and expenditure on domestic applications more closely than the domestic patenting share. The model closely matches almost all other moments including the private value of patenting, turnover in US imports, and R&D expenditures relative to GDP in developed countries. It slightly over-predicts China's share of world real GDP and price level, but does well in matching cross-country variation in price levels and real GDP.

The calibrated parameters are reported in Table 2. We calibrate $\delta_{US} = 0.088$, which implies an expected patent duration of $1/\delta_{US} = 11.3$ years. The legal term of US patents is 20 years. In practice, the expected duration of protection can be less than 20 years because of imperfect patent enforcement and because not all inventions are patentable. Our calibration implies reasonably high levels of patent enforcement in the US.¹⁵

Other developed economies have similar levels of patent protection as the US, with slightly stronger protection in Canada and Korea, but slightly weaker protection in Europe and Japan. However, we find that patent protection is substantially weaker in developing countries. The expected duration of patent protection is 6.8 years in China, 5.9 years in Brazil and between 4 and 5 years in India, Russia and Mexico.¹⁶ It follows that, even after the implementation of TRIPS, there

¹⁵Consider a US patent that has not expired and has been enforced prior to this year. With a legal term of 20 years, our estimate that $\delta_{US} = 0.088$ implies that there is a 3.8 percent probability the patent is not enforced if challenged this year.

¹⁶Calibrated patent protection is substantially weaker in the rest of the world than in any other country. This gap largely reflects the need to patent in multiple jurisdictions to obtain a "rest of world" patent. As we do not adjust

Table 1: Model fit

Moment	Target	Model
International patent shares	See Figure 1	
Share of innovations patented in US	0.40	0.40
Share of domestic patents in inward patents in US	0.58	0.69
Value of patents relative to R&D expenditure in US	0.093	0.091
Turnover in US imports	0.0713	0.0715
Expenditure on domestic patent applications in US (trillion \$)	0.0047	0.0048
Aggregate growth rate	0.017	0.017
Price growth difference (non-patenting minus patenting)	0.0088	0.0088
Trade elasticity in patenting sector	5.0	5.0
World output (trillion \$)	145	145
R&D expenditure relative to GDP in US	0.028	0.028
R&D expenditure relative to GDP in Europe	0.020	0.020
R&D expenditure relative to GDP in Japan	0.032	0.033
R&D expenditure relative to GDP in Canada	0.017	0.017
R&D expenditure relative to GDP in Korea	0.040	0.039

Moment	World real GDP shares		Price indices relative to US	
Country	Target	Model	Target	Model
US	0.17	0.15		
Europe	0.20	0.20	0.87	0.87
Japan	0.05	0.05	0.85	0.85
China	0.17	0.21	0.62	0.78
Brazil	0.03	0.03	0.60	0.57
India	0.07	0.07	0.29	0.30
Canada	0.02	0.01	0.98	0.89
Korea	0.02	0.02	0.76	0.86
Russia	0.03	0.03	0.39	0.36
Mexico	0.02	0.02	0.53	0.51
South Africa	0.01	0.01	0.46	0.46
Rest of world	0.22	0.21	0.57	0.55

Notes: Targets and model-implied values for moments used in simulated method of moments.

exists substantial cross-country variation in patent rights. This finding is consistent with evidence from the Ginarte-Park patent rights index, which suggests that TRIPS narrowed, but did not close, the gap in patent rights between developed and developing economies (Park 2008).¹⁷

The calibration implies technology diffusion is moderately slow. We estimate $\nu = 0.059$, so that it takes 17.1 years on average for an innovation to diffuse. For comparison, Eaton and Kortum (1999) estimate that international technology diffusion takes 21 years on average, while Comin and Mestieri (2014, Table 2.3) report an average adoption lag of 13 years for seven technologies

patenting costs to reflect this need, the model infers that the rest of world has weaker patent protection.

¹⁷The correlation between the Ginarte-Park patent rights index in 2005 (the latest year available) and our calibrated δ_n parameters is -0.84 , which is very high given the different methodologies used to obtain the two measures. Note, however, that calibrating δ_n directly from the Ginarte-Park index is not feasible because the two measures are expressed in incommensurable units.

Table 2: Calibrated parameter values

Parameter	Value
Technology diffusion rate, ν	0.059
Product obsolescence rate, ζ	0.010
Shape parameter of Pareto quality distribution, k	1.25
Shape parameter of Fréchet productivity distribution in patenting sector, θ^1	6.2
Growth rate of no patenting sector, g^0	0.011
Patent preparation cost, f^o	0.048
Patent application cost, f^e	0.079

	Patent protection		R&D efficiency, $\eta_i \times 100$	Patenting sector productivity, $(T_i^1)^{\wedge}(1/\theta^1)$	No patenting sector productivity, $(T_i^0)^{\wedge}(1/\theta^0)$
	δ_i	$1 / \delta_i$			
US	0.088	11.3	2.8	5.5	9.2
Europe	0.095	10.5	1.6	5.1	7.3
Japan	0.107	9.4	2.8	5.1	8.7
China	0.148	6.8	0.4	4.6	4.0
Brazil	0.170	5.9	0.1	3.3	4.5
India	0.241	4.2	0.1	2.2	2.7
Canada	0.073	13.8	1.2	2.5	12.3
Korea	0.066	15.1	1.6	5.1	7.6
Russia	0.228	4.4	0.1	3.7	6.1
Mexico	0.205	4.9	0.1	4.0	4.8
South Africa	0.170	5.9	0.2	3.2	4.5
Rest of world	1.510	0.7	0.2	2.6	3.3

Notes: Table reports parameters calibrated using simulated method of moments.

invented since 1950. We also estimate that $f^o/f^e = 0.6$ implying that the patent preparation cost and patent application cost are of similar magnitudes. A relatively large patent preparation cost is required to match the observed prevalence of domestic patents. Setting $f^o = 0$ would lead to an equilibrium with too much domestic, relative to international, patenting.

R&D efficiency η_i is highest in the US and Japan, followed by Korea, Europe and Canada. And it is an order of magnitude greater in these economies than in the developing countries. Consequently, innovation is highly concentrated in developed economies. In the calibrated steady state, the US accounts for 34 percent of world innovations, Japan and Europe 22 percent each and Korea a further 10 percent. By contrast, China accounts for only 3 percent of innovations and Brazil, India, Russia, Mexico and South Africa each account for fewer than one percent. Calibrated productivity T_i^s is also higher in developed than in developing economies in both sectors. However, the implied productivity gaps are substantially smaller than the variation in R&D efficiency.

In the calibrated economy monopolistic products account for 9 percent of world output and 29 percent of world trade in the patenting sector. However, they account for larger shares of output and trade in more innovative countries. Helpman-Krugman products account for 60 percent of

US exports in the patenting sector, 75 percent of Japanese exports and 39 percent of European exports. By contrast, Eaton-Kortum products account for more than 90 percent of exports in the patenting sector in each of the developing countries. Indeed, while Europe and the US are the biggest exporters of Helpman-Krugman products, China is the largest exporter of Eaton-Kortum products in the patenting sector. These differences illustrate how variation in R&D efficiency leads to stark within-sector specialization in exports across product types.

To assess whether the calibrated model delivers reasonable counterfactuals, we consider two untargeted moments. First, in 1995 the US altered its patent length from 17 years post-grant to 20 years post-application in order to implement TRIPS. This change had differential impacts across technology fields depending upon the average lag between applications and grants. Using this variation, Bertolotti (2024) estimates that, without including anticipation effects, a one month increase in US patent protection raised patenting in the US by around 1.8 percent. Correspondingly, we increase the expected duration of US patent protection by one month starting from our calibrated steady state. We find a similar value: domestic patenting by US innovators $PAT_{US,US}$ increases by 1.0 percent at time zero.¹⁸

Second, Coelli et al. (2022) use firm-level data to estimate the effect of market access on patenting. They estimate that a one percentage point decrease in the tariffs faced by exporters leads to a 2.2 percent increase in patenting. We carry out the analogous tariff reduction exercise in our calibrated model and find that patenting increases by 2.5 percent on average across countries.¹⁹ While Coelli et al. (2022) estimate a partial equilibrium effect, we calculate the change including general equilibrium adjustments. Nevertheless, the similarity of our results is, once again, reassuring.²⁰

¹⁸Bertolotti (2024) estimates the change in patenting relative to a pre-TRIPS baseline. For comparison, using the pre-TRIPS calibration of our model from Section 4.4 below, we find that increasing the expected duration of US patent protection by one month raises domestic patenting by US innovators at time zero by 1.4 percent.

¹⁹Coelli et al. (2022) calculate tariff exposure as the weighted average of country-level tariffs where the weights are firm-specific and measure firms' exposure to different countries based on their patenting history. The set of countries they use includes the firm's own country, where it faces no tariffs. Therefore, a one percentage point decline in tariff exposure implies a larger tariff reduction in export markets when firms patent domestically. To simulate an equivalent change in the export costs faced by country i in our model, we reduce the tariff b_{ni} in each foreign destination n by $\left(\sum_{n=1}^N PAT_{ni} / \sum_{n=1, n \neq i}^N PAT_{ni}\right)$ percentage points. We report the change in patenting eight years after the trade cost shock to match the window used by Coelli et al. and weight changes in patenting by each country's share of global patented innovations.

²⁰As noted in footnote 5, we also use the calibrated model to quantify the impact of assuming patents provide monopoly rights over production in addition to sales. Consider an innovator who holds a production monopoly in their domestic market, but not in any foreign markets. At the calibrated steady state, when the innovator loses their sales monopoly in a foreign market their expected profits in that market drop by 97 percent on average (where the expectation is taken over productivity z and the drop is averaged across all country pairs). It follows that assuming a domestic patent protects production makes little difference to the value of patenting since innovators still faces competition from producers in other countries. Specifically, the expected value of domestic patent protection for firms that choose to patent increases by only 2.7 percent in the US, 5.1 percent in Europe, 7.3 percent in China and 7.8 percent on average across countries.

4 Patent Policy

Using the calibrated model, we study optimal patent policy by simulating counterfactual changes in the strength of patent protection δ_n . We start by characterizing the effect of unilateral changes in patent protection in a single country. Next, we solve for the non-cooperative Nash equilibrium in patent protection levels, followed by the cooperative equilibrium where countries jointly choose patent protection to maximize global welfare. Then, we analyze the welfare effects of the TRIPS agreement. Finally, we study how shocks to the global economy, such as the rise of China, affect optimal patent policy and we discuss the robustness of our results.

For each counterfactual, we study the impact of unanticipated, one-off, permanent changes in patent protection δ_n at time zero and assume that the global economy is in the calibrated steady state initially. We also assume that from time zero onwards the new protection levels apply to all patents that had not expired prior to the change in policy.

We measure welfare changes using the equivalent variation in consumption. The equivalent variation EV_i for country i is defined as the percentage increase in consumption in the initial steady state that delivers the same welfare as the new equilibrium. To calculate EV_i we need to solve for the entire dynamic equilibrium. Appendix C.1 explains how we solve for the transition dynamics between steady states. To obtain a fast and stable solution algorithm, we compute numerical derivatives using a projection method that decomposes functions on the space of Chebyshev polynomials (Judd 1992), rather than using finite differences.²¹ For comparison, we also compute the equivalent variation in steady state welfare EV_i^{SS} . The difference $EV_i - EV_i^{SS}$ captures how the transition dynamics affect welfare changes.

When aggregating welfare changes to the global level, we consider two alternative measures. First, an equal weights measure that sums the welfare of each individual: $U_{W,Equal} = \sum_{i=1}^N L_i u_i = \sum_{i=1}^N L_i^{\frac{1}{\gamma}} U_i$, where u_i denotes the welfare of an individual with normalized consumption per capita \tilde{C}_i/L_i and U_i denotes welfare as defined in equation (2). Second, a measure that uses Negishi (1960) weights based on individuals' initial inverse marginal utilities of consumption $U_{W,Negishi} = \sum_{i=1}^N \left(\tilde{C}_i^{SS,O} / L_i \right)^{\frac{1}{\gamma}} L_i u_i = \sum_{i=1}^N \left(\tilde{C}_i^{SS,O} \right)^{\frac{1}{\gamma}} U_i$ where SS, O denotes the initial steady state. By putting greater weight on the welfare of richer economies, the Negishi measure ensures that a social planner has no incentive to redistribute income across countries in the initial steady state.

²¹To the best of our knowledge, this projection method has not previously been used in the trade and growth literature. We strongly recommend it. Our counterfactual analysis would not be feasible with standard finite differences methods. See Schesch (2024) for more on the advantages of projection methods using Chebyshev polynomials.

4.1 Unilateral Patent Policy


To better understand the channels through which patent protection affects welfare, we start by analyzing the impact of unilateral policy changes that vary δ_n in one country at a time. As an example, suppose we reduce δ_{US} from 0.088 to 0.05. Figure 2 plots how key variables adjust over 60 years relative to their initial steady state values. Stronger patent protection in the US leads to a reallocation of labor into R&D in both the US and, to a lesser extent, other countries (panel a). The growth in R&D employment is greater in countries, such as Mexico and Canada, that trade relatively more with the US. Increased R&D raises the global growth rate g creating dynamic benefits for all countries (panel b).

Stronger protection also increases market power, generating static distortions that primarily affect the US. The share of aggregate quality sold competitively in the patenting sector $\tilde{\Psi}_{Ci}^1$ declines by around 13 percent in the US compared to a fall of around 2 percent in other countries (panel c). Increased market power raises real profits $\tilde{\Pi}_i/P_i$ in all countries, with Mexico and Canada experiencing the biggest increases since they are most dependent on US sales (panel d).²² Higher real profits are offset by lower real wages, but the wage decline is largest in the US where market power effects are strongest. It follows that normalized consumption \tilde{C}_i falls further in the US than in other countries (panel e). Combining the dynamic benefits with the static costs, we find that welfare falls by 0.5 percent in the US, but rises by around 1 percent elsewhere.

Panels (b) and (e) show that the growth rate adjusts more quickly than normalized consumption. Intuitively, changing protection affects growth immediately, whereas the static distortions adjust more slowly because the share of aggregate quality sold competitively in each market is a state variable. Panel (f) highlights this difference by plotting the dynamics of g and \tilde{C}_{US} with the changes in both variables between steady states normalized to positive one. After 20 years, the growth rate has completed 88 percent of its adjustment, whereas normalized consumption has completed 73 percent. This ranking of adjustment speeds holds across all the counterfactuals we consider below. Consequently, the change in welfare accounting for transition dynamics is systematically higher than the steady state change in countries that strengthen patent protection (i.e. $EV_i > EV_i^{SS}$) because the costs from increased market power materialize more slowly than the change in the growth rate. Likewise, transition dynamics reduce welfare relative to the steady state change EV_i^{SS} in countries that weaken patent protection.

Next, we analyze how the effects of unilateral patent policy depend upon which country changes its patent protection. For these and subsequent counterfactuals we focus on reporting changes in growth and welfare. Varying δ_n in one country n at a time, Figure 3 plots the steady state world

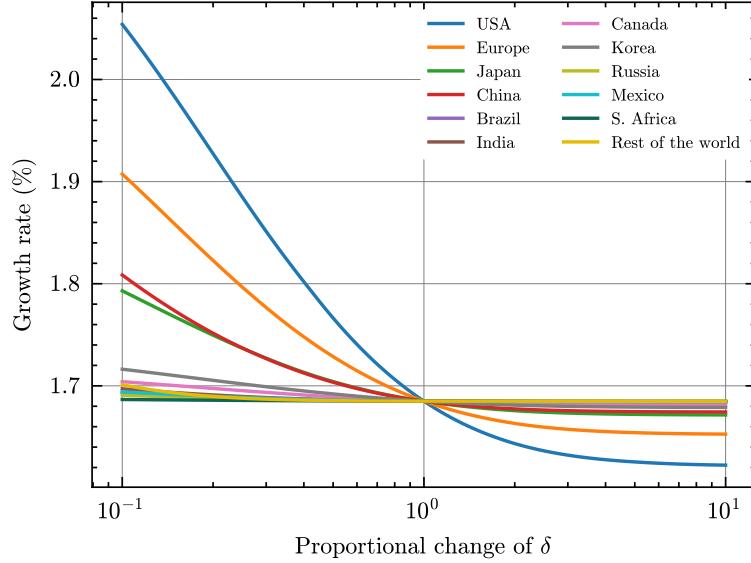
²²The US experiences a large rise in nominal profits (due to home bias in trade), but this is offset by an increase in the US price level relative to other countries. As a result, real profits increase less in the US than for all countries except Europe, Japan and Korea.



Notes: This figure plots the effects of reducing δ_{US} from 0.088 to 0.05 at time zero on R&D employment (panel a), growth (panel b), the share of aggregate quality sold competitively in the patenting sector (panel c), aggregate real profits (panel d) and normalized consumption (panel e). Panels (a)–(e) show the values of each variable relative to the calibrated steady state. The solid dots at year zero and year 60 show the values of the initial and new steady states, respectively. Panel (f) plots changes in growth and normalized US consumption relative to the changes in these variables between the initial and new steady states.

Figure 2: Transition dynamics with stronger US patent protection

growth rate g as a function of the proportional change in δ_n . The figure shows that stronger patent protection (i.e. lower δ) increases growth, but that the magnitude of the effect varies greatly across countries. The effect is largest for the US followed by Europe, and then Japan and China together. By contrast, patent protection in Korea and Canada has only a small effect on growth and the growth rate is effectively inelastic to changes in protection in the remaining countries. These results are consistent with our observation in Section 2.4 that growth is more sensitive to the level of patent protection in larger and more innovative countries.

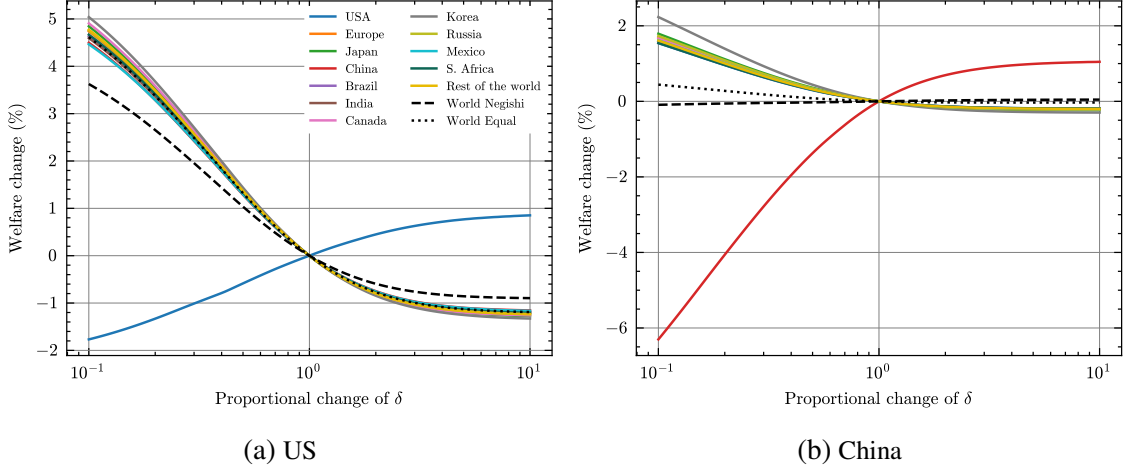


Notes: This figure plots the effect of a proportional change in calibrated patent protection δ_n of country n on the steady state growth rate g .

Figure 3: Unilateral patent policy changes and growth

Turning to welfare, Figure 4 plots welfare changes EV_i by country, and for the world, against proportional changes in δ_n . Panel (a) shows the impact of variation in US patent protection δ_{US} . Since stronger protection raises growth, it increases welfare in all countries (other than the US) and for the world. However, in the US, the dynamic benefits of stronger protection are offset by the static costs of higher prices from increased market power. As in the $\delta_{US} = 0.05$ counterfactual above, we find that these static costs are slightly larger than the dynamic benefits for the US itself. Consequently, reducing δ_{US} decreases US welfare. For comparison, panel (b) shows welfare effects from varying Chinese policy δ_{CHN} . Countries other than China gain from stronger Chinese protection. However, the benefits are smaller than in panel (a) and the costs to China itself are greater. World welfare is almost flat in δ_{CHN} and falls using Negishi weights.

Varying patent protection in other countries gives similar results to those shown in Figure 4, but with different magnitudes (see Appendix C.2). Strengthening protection in country n raises welfare in countries $i \neq n$, but reduces welfare in country n itself. Thus, all countries have a uni-



Notes: This figure plots the effect of proportional changes in calibrated patent protection on welfare in all countries relative to the calibrated steady state. Panel (a) changes δ_{US} and panel (b) changes δ_{CHN} . The legend in panel (a) applies to both panels. Welfare changes are expressed as the equivalent variation in consumption and account for transition dynamics.

Figure 4: Unilateral patent policy changes and welfare

lateral incentive to weaken patent protection starting from the calibrated steady state. Moreover, the gains from stronger protection in country n are larger when the growth rate is more sensitive to δ_n . Compared to the US case in panel (a), stronger protection always generates smaller foreign benefits and greater domestic costs. For policy changes in Europe, Japan, Canada, Korea and Mexico, the benefits of stronger protection to foreign countries exceed the domestic costs and world welfare rises. However, for the remaining countries, stronger protection reduces world welfare using Negishi weights.

4.2 Nash Equilibrium

Do existing international agreements over patent policy increase welfare? To answer this question, we compare the calibrated steady state to an equilibrium where each country independently chooses its patent protection δ_n to maximize its own welfare. Specifically, we solve for a Nash equilibrium in which each country chooses a one-off, permanent change in its patent protection to maximize its welfare including transition dynamics and starting from the calibrated steady state.²³ We bound the expected duration of each country's protection between one month and 100 years. This range corresponds to values of δ_n between 12 and 0.01, which we refer to as no protection and complete protection, respectively. Numerically, we find that the Nash equilibrium is unique.²⁴

²³Alternatively, we could assume countries seek to maximize steady state welfare. However, we find that this alternative makes no difference to optimal patent protection in either the Nash equilibrium or the cooperative equilibria.

²⁴Appendix C.3 provides further details about how we solve for the Nash equilibrium and for the cooperative equilibria.

Table 3, column (a) reports patent protection levels in the Nash equilibrium together with changes in growth and welfare relative to the calibrated steady state. From the analysis above, we know that all countries have a unilateral incentive to weaken patent protection. Therefore, it is not surprising that in the Nash equilibrium there is no patent protection in any country.²⁵

Table 3: Nash and cooperative equilibria

	(a) Nash			(b) Cooperative: equal			(c) Cooperative: Negishi		
	Patent protection, δ_i	Welfare change (percent)		Patent protection, δ_i	Welfare change (percent)		Patent protection, δ_i	Welfare change (percent)	
		Total, EV_i	Steady state, EV_i^{ss}		Total, EV_i	Steady state, EV_i^{ss}		Total, EV_i	Steady state, EV_i^{ss}
US	None	-0.5	-0.4	Complete	1.6	-2.9	Complete	1.4	-3.2
Europe	None	-1.0	-0.8	Complete	2.4	-2.0	Complete	2.2	-2.2
Japan	None	-1.3	-1.2	Complete	2.4	-2.0	Complete	2.2	-2.2
China	None	-1.3	-1.0	None	10.0	13.0	None	9.8	12.8
Brazil	None	-1.9	-1.6	None	9.3	12.2	None	9.2	12.0
India	None	-2.1	-1.8	None	9.0	11.9	None	8.9	11.8
Canada	None	-0.6	-0.4	Complete	4.2	2.4	Complete	4.1	2.2
Korea	None	-0.4	-0.2	Complete	4.2	2.8	Complete	4.0	2.6
Russia	None	-2.2	-1.9	None	9.2	12.3	None	9.1	12.1
Mexico	None	-2.1	-1.8	Complete	1.9	-5.6	None	8.7	11.0
South Africa	None	-2.0	-1.8	None	9.2	12.2	None	9.0	12.1
Rest of world	None	-2.5	-2.3	None	8.7	11.9	None	8.6	11.7
World Equal		-2.1	-1.9		8.7	11.2		8.6	11.3
World Negishi		-1.5	-1.2		6.0	5.3		6.0	5.5
Growth rate change (percentage)		-0.13		0.68		0.67			

Notes: This table reports levels of patent protection and welfare changes relative to the calibrated steady state in the Nash equilibrium (column a), the cooperative equilibrium when the social planners uses equal weights for all individuals (column b), and the cooperative equilibrium when the social planner uses Negishi weights (column c). Welfare changes are expressed as the equivalent variation in consumption. The table reports welfare changes computed accounting for transition dynamics, EV_i , and between steady states, EV_i^{ss} . It also reports the change in the steady state global growth rate, g . Complete protection corresponds to $\delta_i = 0.01$ and no protection corresponds to $\delta_i = 12$, which are the bounds we impose on our solutions.

With all countries offering weaker patent protection, innovation falls, and steady state growth g is 0.13 percentage points lower than in the calibrated economy. The reduction in growth is small because the slow pace of technology diffusion implies that innovators expect a relatively long-lasting technological monopoly even when there is no patent protection. Lower growth is partially offset by a decline in market power. The share of monopolistic products in world output of the patenting sector falls from 8.6 percent to 6.7 percent. However, in welfare terms, the fall in growth

²⁵Note that this is a quantitative results that depends on the calibration of the model. With sufficiently high trade costs some countries do provide protection in the Nash equilibrium (see Section 4.5).

dominates and world welfare is 2.1 percent lower in the Nash equilibrium using equal weights and 1.5 percent lower using Negishi weights. It follows that, by providing some patent protection, the observed equilibrium in 2015 generates higher welfare than the Nash equilibrium. In other words, existing patent rights agreements have been somewhat successful.

All countries are worse off in the Nash equilibrium. Losses range from 0.4 percent in Korea to 2.5 percent in the rest of the world. Countries with stronger patent protection in the calibrated steady state experience smaller losses on average because they implement bigger reductions in patent protection and, therefore, experience greater falls in domestic market power. Consistent with the analysis in Section 4.1, Table 3 also shows that accounting for transition dynamics increases the welfare costs of moving to the Nash equilibrium.

4.3 Cooperative Equilibrium

What could be achieved through global cooperation over patent policy? To assess the potential gains, we next solve for the cooperative equilibrium where countries choose all δ_n jointly to maximize world welfare starting from the calibrated steady state. Again, we consider one-off, permanent changes in patent protection and allow δ_n to vary between 0.01 and 12 for each country. The results are shown in Table 3 for equal weights in column (b) and Negishi weights in column (c). In each case we find that there is a unique cooperative equilibrium.

With equal weights, the cooperative equilibrium is for the US, Europe, Japan, Canada, Korea and Mexico to provide complete patent protection, while all other countries offer no protection. Why is this set of policies optimal? We showed in Section 2.4 and Figure 3 that the growth benefits of stronger protection are greater when protection is provided by countries that are larger or are more closely integrated with innovative economies. At the same time, the static costs of increased market power are greater when borne by a larger population, but are not directly affected by how innovative a country is. Consequently, home bias in trade means it is optimal for more innovative countries to provide stronger protection in the cooperative equilibrium, whereas country size does not determine optimal policy.

US, Europe, Japan, Canada and Korea are the five most innovative economies in the calibrated steady state and together account for 92 percent of global innovations. In our calibration, it is efficient for the world to delegate the job of incentivizing innovation to these countries. By contrast, although Figure 3 shows that protection in China and Japan has similar effects on the growth rate, it is not optimal for China to provide protection as it has a much larger population than Japan. The exception to this logic is Mexico, which accounts for only 0.1 percent of global innovations in the calibrated steady state. Nevertheless, Mexico provides complete protection due to its high levels of trade with the US. The benefits that stronger protection in Mexico generates through increased

US innovation outweigh the static costs from higher market power in Mexico.

Steady state growth is 0.68 percentage points greater in the cooperative equilibrium than in the calibrated steady state. And faster growth leads to higher welfare. World equal weights welfare is 8.7 percent higher in the cooperative equilibrium. This increase is approximately half as large as the total gains from trade (relative to autarky) in the patenting sector in our model, demonstrating that the gains from cooperation over patent policy are substantial relative to the benefits of trade integration.²⁶ However, the overall gains mask strong distributional effects. While welfare in China, Brazil, India, Russia, South Africa and the rest of the world rises by nearly 10 percent, the gains for countries that provide complete protection are much smaller, ranging from 1.6 percent in the US to 4.2 percent in Canada and Korea. This cross-country variation in the gains from cooperation arises because the static costs of encouraging innovation are borne by the economies that offer protection, while other countries free ride on the dynamic benefits of higher growth.

The cooperative equilibrium with Negishi weights is similar to the equal weights equilibrium except that Mexico switches from complete protection to no protection. This change slightly dampens the increase in growth, leading to lower gains for all countries except Mexico. However, in both the Negishi weights and equal weights equilibria all countries gain from cooperation.²⁷ To sum up, there is scope for reforms to the international patenting system to bring large welfare gains.

Table 3 shows that, in the Nash and cooperative equilibria, countries choose boundary, rather than interior, patent policies. This behavior occurs because the model features heterogeneous, open economies. By contrast, when patent policy is harmonized across countries (or in closed economies), we find that there is typically an interior optimal level of protection at which the static costs and dynamic benefits of stronger protection are balanced. For example, under the constraint $\delta_n = \delta$ for all n , steady state world welfare is hump-shaped as a function of δ and the maximum occurs when $\delta = 0.11$ using equal weights or when $\delta = 0.07$ using Negishi weights. However, as the cooperative equilibria demonstrate, harmonizing patent policy is globally inefficient. Even when the harmonized δ is chosen optimally, world welfare is lower than in either the cooperative equilibria or the calibrated steady state, though it is higher than in the Nash equilibrium.²⁸

²⁶Specifically, we find that shutting down trade in the patenting sector reduces world equal weights welfare by 18.4 percent with country-level losses ranging from 15.8 percent in China to 21.8 percent in Mexico. The gains from trade are larger in our model than in static trade models because trade integration generates dynamic gains through higher growth. Holding constant the allocation of labor to R&D and patenting, we find that the static cost of shutting down trade in the patenting sector is a 5.9 percent fall in world equal weights welfare.

²⁷This result holds only when welfare changes include transition dynamics. In the US, Europe and Japan, steady state welfare is lower in the cooperative equilibria than in the calibrated economy, which illustrates the importance of accounting for transition dynamics when calculating welfare effects.

²⁸At the equal weights optimum $\delta = 0.11$, steady state world welfare using equal weights is 1.4 percent lower than in the calibrated steady state. For the Negishi weights optimum $\delta = 0.07$, steady state world welfare using Negishi weights is 0.2 percent lower than in the calibrated steady state.

4.4 TRIPS

The TRIPS Agreement came into effect in 1995, but allowed developing countries to phase-in implementation over ten years, while giving least developed countries even longer to adjust. TRIPS sought to narrow the gap between the strength of IPR in developed and developing economies by: introducing minimum standards of protection and enforcement for all WTO members; applying the principles of national treatment and most-favoured nation treatment to intellectual property, and; placing IPR under the remit of the WTO's dispute settlement mechanism. For patents, TRIPS mandated that countries make patents available for inventions in all fields of technology and that patents should be enforceable for at least 20 years. In practice, implementation of TRIPS required developing countries to strengthen IPR, but had little or no impact on policies in developed countries (Saggi 2016).

Between 1990 and 2015 there was a rapid increase in international patent applications filed in China, India, Russia and, to a lesser extent, Brazil. In the model, the bilateral patenting threshold $\psi_{ni}^{s,e*}$ is decreasing in both the profitability of destination n for innovators from i and the strength of patent protection in n (see equation 26). Consequently, although the observed rise in cross-border patenting in developing countries could be caused by stronger patent rights, it may also result from increased market size due to rapid economic growth. We use the model to disentangle these alternative explanations and quantify changes in patent rights by calibrating patent protection prior to TRIPS while accounting for how changes in market size affect the patenting threshold.

Since countries may have initiated patent reforms in anticipation of TRIPS, we use 1992 data for the pre-TRIPS calibration. We recalibrate parameters that are country-specific or that are inferred by exact moment matching, while holding other parameters fixed at their values from the 2015 calibration. In particular, we recalibrate patent protection δ_i , R&D efficiency η_i and productivity T_i^s using simulated method of moments estimation with 1992 data (see Appendix B.6 for details). Column (a) of Table 4 reports the estimated patent protection levels in 2015 and in 1992. We find that China, India and Russia had considerably weaker protection in 1992 than 2015, but that protection in Brazil, Mexico and the developed economies is similar in both periods.²⁹

Using these estimates, we study the welfare effects of TRIPS by simulating a counterfactual return to pre-TRIPS patent protection starting from the calibrated steady state in 2015. The assumption that TRIPS was the primary cause of changes in protection between 1992 and 2015 is more plausible for developing than for developed countries.³⁰ Therefore, our main pre-TRIPS

²⁹Interestingly, we estimate that South Africa had slightly stronger protection in 1992 than in 2015. This finding is consistent with the fact that South Africa passed legislation in 1997 to weaken patent protection for pharmaceuticals by allowing for compulsory licensing and parallel imports of patented drugs on public health grounds (Marc 2001).

³⁰China and Russia did not join the WTO until after 1995 (China in 2001 and Russia in 2012), but both countries reformed their intellectual property policies to comply with the TRIPS agreement as part of the accession process. However, in addition to TRIPS, the transition towards capitalism may have contributed to increases in the strength of

Table 4: Pre-TRIPS counterfactuals

	(a)		(b)	(c)	(d)	(e)
	Patent protection, δ_i		Reverse TRIPS Developing countries	All countries	Reverse Uruguay Round	Harmonize with US
	Baseline, 2015	Pre-TRIPS, 1992	Welfare change, EV_i (percent)			
US	0.088	0.115	-0.11	-0.08	-2.06	1.07
Europe	0.095	0.115	-0.12	-0.34	-2.27	0.97
Japan	0.107	0.105	-0.13	-0.69	-2.55	0.83
China	0.148	0.223	0.47	-0.06	-1.65	-0.19
Brazil	0.170	0.194	0.01	-0.52	-2.00	0.13
India	0.241	0.468	0.19	-0.34	-3.47	-0.42
Canada	0.073	0.044	-0.12	-0.67	-2.08	1.10
Korea	0.066	0.080	-0.16	-0.31	-3.44	1.32
Russia	0.228	0.422	0.14	-0.41	-2.39	-0.09
Mexico	0.205	0.215	-0.09	-0.61	-2.01	0.08
South Africa	0.170	0.128	-0.34	-0.64	-2.75	0.39
Rest of world	1.510	1.287	-0.11	-0.65	-2.11	-0.96
World Equal			0.06	-0.46	-2.49	-0.56
World Negishi			0.03	-0.34	-2.19	0.16

Notes: This table reports calibrated levels of patent protection in 1992 (pre-TRIPS) and in 2015 (column a) and welfare changes for different patent policy and trade cost counterfactuals (columns b to e). In column (b), China, Brazil, India, Russia, Mexico and South Africa revert to pre-TRIPS patent protection levels. In column (c), all countries with stronger patent protection in 2015 than in 1992 revert to pre-TRIPS patent protection. In column (d), China, Brazil, India, Russia, Mexico and South Africa revert to pre-TRIPS patent protection and tariffs in the patenting sector for all countries are set equal to their values in 1992. In column (e), all countries with weaker patent protection than the US in 2015 set $\delta_i = \delta_{US}$. Welfare changes are expressed as the equivalent variation in consumption relative to the calibrated steady state in 2015 and account for transition dynamics.

counterfactual in column (b) of Table 4 analyzes the case where patent protection in developing countries (China, Brazil, India, Russia, Mexico and South Africa) reverts to pre-TRIPS levels, while protection elsewhere remains unchanged. For completeness, column (c) sets protection to pre-TRIPS levels in all countries where estimated protection is stronger in 2015 than in 1992.

The counterfactual welfare changes in column (b) show that returning developing countries to pre-TRIPS patent policies benefits China, India, Russia, Brazil and the world as a whole, while having small negative effects on other countries due to a slight decline in the global growth rate. Welfare increases by 0.47 percent in China, 0.19 percent in India, 0.14 percent in Russia and 0.06 percent for the world using equal weights. We conclude that, to the extent TRIPS increased protection in developing economies, it reduced welfare in those countries that strengthened protection. These countries faced local static costs from increased market power, but did not realize offsetting dynamic benefits because growth is largely inelastic to patent policy in developing countries.³¹

patent protection in China and Russia during our sample period.

³¹Table 4 reports pre-TRIPS counterfactuals relative to the 2015 baseline. Simulating TRIPS relative to the 1992

However, there are two important caveats to the conclusion that TRIPS harmed developing countries. First, the welfare changes in column (b) are small relative to those implied by other counterfactuals, such as the cooperative equilibrium. Consequently, small increases in protection in developed countries are sufficient to offset the costs of TRIPS. The results in column (c) demonstrate this point. When all countries with stronger protection in 2015 revert to pre-TRIPS protection levels welfare declines in all countries, with world equal weights welfare falling by 0.46 percent. Thus, the overall welfare effects of changes in patent protection between 1992 and 2015 are dominated by the impact of small increases in patent protection in developed economies (whether due to TRIPS or other causes), not by TRIPS-induced patent reform in developing countries.

Second, TRIPS formed part of a single undertaking that countries agreed as part of the Uruguay Round or in order to join the WTO. The single undertaking led to reduced trade costs and changes in subsidies and other domestic policies, as well as IPR reforms. There has not been a comprehensive ex-post quantification of the welfare effects of the Uruguay Round, likely due to the challenge of quantifying changes in non-tariff policies.³²

To quantify the trade-off between TRIPS and market access in our model, we change patent protection and tariffs simultaneously. Column (d) reports welfare changes for a counterfactual where patent protection in developing countries reverts to pre-TRIPS levels (as in column b) and tariffs in the patenting sector revert to their 1992 values for all country pairs. The average tariff increase in this counterfactual is 9.0 percentage points. We find that the tariff rise dominates the impact of weaker patent protection and welfare declines in all countries, with world equal weights welfare falling by 2.5 percent.³³ This counterfactual implies that relatively small reductions in market access are sufficient to offset the benefits to developing countries of reverting to pre-TRIPS patent protection levels. We conclude that the net welfare effect of WTO membership is positive even for countries that suffer losses due to TRIPS.

Finally, although our baseline calibration implies that TRIPS did not achieve its goal of harmonizing patent protection across countries, we can study what would happen under international harmonization. Starting from the 2015 calibration, column (e) reports welfare changes from setting patent protection in all countries with weaker protection than the US equal to δ_{US} , i.e. letting $\delta_i = \min\{\delta_i, \delta_{US}\}$. Harmonization increases welfare in developed economies, but reduces welfare in most developing countries. Welfare falls by 0.19 percent in China, 0.42 percent in India

calibration leads to the same conclusions. For example, if we start in the calibrated steady state in 1992 and increase patent protection in developing countries to 2015 levels, welfare falls by 0.34 percent in China, 0.18 percent in India, 0.21 percent in Russia and 0.18 percent for the world using equal weights, but increases slightly in developed countries.

³²Ex-ante forecasts suggest that the benefits of greater trade integration may have been sufficient to offset the costs of TRIPS. For example, reviewing thirteen ex-ante studies of the Uruguay Round, Martin and Winters (1995) conclude that cuts to merchandise trade protection will increase real incomes in developing countries by 1.2–2.0 percent.

³³For comparison, Caliendo et al. (2023) estimate that most-favored nation tariff cuts between 1990 and 2010 increased welfare by 2–3 percent.

and 0.56 percent for the world (using equal weights). Thus, harmonizing policy to US levels is globally inefficient, benefits richer countries at the expense of poorer nations and has larger welfare consequences than those due to changes in patent protection between 1992 and 2015. This finding supports the argument that it would be costly for developing countries to further strengthen protection and adopt US-style patent rights.

4.5 The Changing World Economy

To conclude the counterfactual analysis we briefly discuss how changes in the global economy, relative to the calibrated steady state in 2015, affect optimal patent policy. We consider the consequences of trade integration, faster technology diffusion and the rise of China. Appendix C.4 provides further details on the impacts of trade integration and faster technology diffusion.

Trade integration. Lower trade costs lead to increased divergence between unilateral incentives and global efficiency, strengthening the case for patent policy cooperation. Reduced home bias implies stronger international spillovers from patent policy, raising the incentive for countries to free ride by unilaterally weakening protection. In particular, the derivative of each country's welfare with respect to its own δ becomes more positive as trade costs fall, implying the benefits from reducing protection are greater.³⁴ In addition, as trade becomes less home biased, R&D investment becomes more sensitive to foreign patent protection because exports account for a greater share of innovators' profits. This channel increases the dynamic benefits obtained when less innovative countries strengthen patent protection, which expands the set of countries that offer protection in the cooperative equilibrium. Under free trade all countries other than China, India and the rest of the world provide complete patent protection in the equal weights cooperative equilibrium. This example illustrates how optimal patent policy depends on levels of trade integration.

Faster technology diffusion. Faster technology diffusion increases the value of patenting because it reduces the expected duration of an innovator's technological monopoly. Consequently, in contrast to lower trade costs, an increase in the rate of technology diffusion ν raises the incentive for countries to provide patent protection. If we double the rate of technology diffusion by setting $\nu = 0.12$, then US welfare is almost inelastic to unilateral changes in US patent protection (see Appendix C.4) and in the Nash equilibrium the US provides complete patent protection (although other countries continue to offer no protection). Since patenting is more valuable with faster technology diffusion, the welfare gains from optimal patent policy also rise. When $\nu = 0.12$ world equal weights welfare is 14.3 percent higher in the cooperative equilibrium with equal weights, compared to 8.7 percent higher in the baseline calibration with lower ν .

³⁴Conversely, higher trade costs reduce the benefits of unilaterally weakening protection and for sufficiently high trade costs some countries offer protection in the Nash equilibrium. For example, if we double iceberg trade costs in the patenting sector, then the US provides complete patent protection in the Nash equilibrium.

Rise of China. As China moves closer to the technology frontier, Chinese firms have become more innovative. In our data China is the origin of 0.1 percent of international patent applications in 1992, but 6.3 percent in 2015. In our calibrations, China’s R&D efficiency increased from 1.7 percent of US R&D efficiency in 1992 to 13 percent in 2015. These changes imply that global growth is becoming more sensitive to China’s patent policy. All else equal, China will join the club of nations that provide patent protection in the cooperative equilibrium with equal weights when its R&D efficiency relative to the US reaches 28 percent. Therefore, if current trends continue, China will join the patent protection club in the coming decades. By contrast, India’s calibrated R&D efficiency relative to the US in 2015 is 3.1 percent, implying that India is much less innovative than China and much further away from joining the patent protection club.

4.6 Sensitivity and Robustness

We have undertaken a battery of checks to systematically assess the sensitivity of our calibration and the robustness of our counterfactual results to variation in the parameters and moments that we use. These checks are summarized in Appendix D. They show that the main findings of the unilateral, Nash, cooperative and TRIPS counterfactuals in Sections 4.1–4.4 are robust to moderate changes in the externally calibrated parameters and the targeted moments. We do find that the magnitudes of counterfactual welfare changes differ from the baseline results in some robustness checks. For example, welfare gains in the cooperative equilibrium are much greater when the turnover moment is larger (implying faster technology diffusion), which increases the sensitivity of growth to patent protection. However, even in such cases, the qualitative variation in welfare effects across countries is similar to the baseline calibration and optimal patent policies are unchanged.

4.7 Extensive Margin of Trade

We have also analyzed the robustness of our findings to extending the model by including an extensive margin of exporting. Suppose innovators must pay a sunk, destination-specific entry cost to start exporting. In this environment, greater patent protection encourages export entry by raising profits, which creates an additional incentive for countries to offer stronger protection. Ivas (2015) and Cockburn et al. (2016) provide evidence consistent with the existence of this extensive margin channel. However, we find that allowing for sunk export costs makes little difference to our quantitative results (see Appendix E for details). The intuition for this finding is simple. The marginal innovator that is induced to start exporting by stronger protection is a relatively small exporter. By contrast, the increase in market power due to stronger protection is infra-marginal as the expected duration of protection increases for all innovators. Consequently, the extensive margin benefits of stronger protection are quantitatively small compared to the static costs.

5 Conclusion

Whether and how patent rights should vary across countries has long been controversial. But the debate has been hampered by a lack of evidence on the welfare effects of patent policy. We address this gap by developing a new quantitative framework that models the trade-off between the static costs and dynamic benefits of patent protection. The framework integrates Helpman-Krugman and Eaton-Kortum trade in a growth model with endogenous patenting. We calibrate the model using data on patent applications, trade flows and R&D, and use the calibrated economy to study optimal patent policy and the impact of TRIPS.

The counterfactual results imply that there are large potential gains from global IPR cooperation. However, realizing these gains requires not that policies are harmonized across countries, but that more innovative economies offer greater protection. Two mechanisms drive these findings. First, the dynamic benefits from increased innovation spillover across borders, whereas the static costs of higher prices are borne domestically. Consequently, non-cooperative patent rights are weaker than is globally efficient. Second, stronger patent protection generates greater benefits when provided by more innovative economies. Consequently, developed countries provide stronger patent rights than developing economies in the cooperative equilibrium. And, by pushing towards policy harmonization, TRIPS harms developing countries.

We caution that our analysis represents only a first attempt at quantifying the effects of patent policy in open economies. In the tradition of Nordhaus (1969) and Grossman and Lai (2004), we have quantified the trade-off between innovation and market power, which is at the heart of policy debates over IPR. But there exist other channels, such as FDI, through which patent policy may affect economic outcomes. We hope that future research can build upon our framework to better understand this important topic.

References

- Acemoglu, Daron, and Ufuk Akcigit. 2012. “Intellectual Property Rights Policy, Competition and Innovation.” *Journal of the European Economic Association* 10(1): 1-42.
- Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. “Innovation, Reallocation, and Growth.” *American Economic Review* 108(11): 3450–3491.
- Akcigit, Ufuk, Sina Ates, and Giammario Impullitti. 2021. “Innovation and Trade Policy in a Globalized World.” NBER Working Paper 24543.
- Antràs, Pol, and Robert Staiger. 2012. “Offshoring and the Role of Trade Agreements.” *American Economic Review* 102(7): 3140–83.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. “New Trade Models,

- Same Old Gains?” *American Economic Review* 102(1): 94-130.
- Bertolotti, Fabio. 2024. “Patent Term, Innovation, and the Role of Technology Disclosure Externalities.” Mimeo.
- Bloom, Nicholas, John Van Reenen, and Heidi Williams. 2019. “A Toolkit of Policies to Promote Innovation.” *Journal of Economic Perspectives* 33(3): 163–184.
- Borota Milicevic, Teodora, Fabrice Defever, Giammario Impullitti, and Adam Hal Spencer. 2023. “Innovation Union: Costs and Benefits of Innovation Policy Coordination.” Mimeo.
- Bureau of Economic Analysis (BEA). 2022. “Gross Output by Industry.” (accessed April 22, 2022 and February 18, 2023).
- Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor. 2023. “Tariff Reductions, Heterogeneous Firms, and Welfare: Theory and Evidence for 1990–2010.” *IMF Economic Review*: 1-35.
- Chaudhuri, Shubham, Pinelopi K. Goldberg, and Panle Jia. 2006. “Estimating the Effects of Global Patent Protection in Pharmaceuticals: A Case Study of Quinolones in India.” *American Economic Review* 96(5): 1477-1514.
- Cockburn, Iain M., Jean O. Lanjouw, and Mark Schankerman. 2016. “Patents and the Global Diffusion of New Drugs.” *American Economic Review* 106(1): 136-164.
- Coelli, Federica, Andreas Moxnes, and Karen Helene Ulltveit-Moe. 2022. “Better, Faster, Stronger: Global Innovation and Trade Liberalization.” *The Review of Economics and Statistics* 104(2): 205–216.
- Comin, Diego, and Marti Mestieri. 2014. “Technology Diffusion: Measurement, Causes, and Consequences.” In *Handbook of Economic Growth*, Vol.2, 565-622.
- Dechezleprêtre, Antoine, Yann Ménière, and Myra Mohnen. 2017. “International Patent Families: From Application Strategies to Statistical Indicators.” *Scientometrics* 111: 793-828.
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. “The Rise of Market Power and the Macroeconomic Implications.” *Quarterly Journal of Economics* 135(2): 561-644.
- De Rassenfosse, Gaetan, Marco Grazzi, Daniele Moschella, and Gabriele Pellegrino. 2022. “International Patent Protection and Trade: Transaction-level Evidence.” *European Economic Review*, 147: 104-160.
- De Rassenfosse, Gaetan, and Bruno van Pottelsberghe de la Potterie. 2013. “The Role of Fees in Patent Systems: Theory and Evidence.” *Journal of Economic Surveys* 27(4): 696-716.
- Duggan, Mark, Craig Garthwaite, and Aparajita Goyal. 2016. “The Market Impacts of Pharmaceutical Product Patents in Developing Countries: Evidence from India.” *American Economic Review* 106(1): 99-135.
- Eaton, Jonathan, and Samuel Kortum. 1999. “International Technology Diffusion: Theory and Measurement.” *International Economic Review* 40(3): 537-570.

- Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741-1779.
- Griffith, Rachel, Stephen Redding, and John Van Reenen. 2004. "Mapping the Two Faces of R&D: Productivity Growth in a Panel of OECD Industries." *Review of Economics and Statistics* 86(4): 883-895.
- Grossman, Gene M., and Edwin Lai. 2004. "International Protection of Intellectual Property." *American Economic Review* 94(5): 1635-1653.
- Grossman, Gene M., Phillip McCalman, and Robert W. Staiger. 2021. "The "New" Economics of Trade Agreements: From Trade Liberalization to Regulatory Convergence?" *Econometrica* 89(1): 215-49.
- Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." In *Handbook of International Economics*, Vol.4, 131-195.
- Helpman, Elhanan. 1993. "Innovation, Imitation, and Intellectual Property Rights." *Econometrica* 61(6): 1247-80.
- Helpman, Elhanan, and Paul Krugman. 1987. *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. MIT Press.
- Hjort, Jonas, Hannes Malmberg, and Todd Schoellman. 2022. "The Missing Middle Managers: Labor Costs, Firm Structure, and Development." NBER Working Paper 30592.
- Hsieh, Chang-Tai, Pete Klenow, and Kazuatsu Shimizu. 2022. "Romer or Ricardo?" Mimeo.
- Ivus, Olena. 2015. "Does Stronger Patent Protection Increase Export Variety? Evidence from US Product-level Data." *Journal of International Business Studies* 46: 724-731.
- Judd, Kenneth L. 1992. "Projection Methods for Solving Aggregate Growth Models." *Journal of Economic Theory* 58(2): 410-452.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics* 132(2): 665-712.
- Kyle, Margaret, and Anita McGahan. 2012. "Investments in Pharmaceuticals Before and After TRIPS." *Review of Economics and Statistics*, 94(4): 1157-1172.
- LaBelle, Jesse, Inmaculada Martinez-Zarzoso, Ana Maria Santacreu, and Yoto V. Yotov. 2024. "Cross-border Patenting, Globalization, and Development."
- Lai, Edwin, and Isabel Yan. 2013. "Would Global Patent Protection be too Weak without International Coordination?" *Journal of International Economics* 89(1): 42-54.
- Lind, Nelson, and Natalia Ramondo. 2022. "Global Knowledge and Trade Flows: Theory and Measurement." NBER Working Paper 30590.
- Maggi, Giovanni, and Ralph Ossa. 2023. "The Political Economy of International Regulatory Cooperation." *American Economic Review* 113(8): 2168-2200.

- Marc, Patrick. 2001. "Compulsory Licensing and the South African Medicine Act of 1997: Violation or Compliance of the Trade Related Aspects of Intellectual Property Rights Agreement." *NYL Sch. J. Int'l & Comp. L.* 21: 109.
- Martin, Will, and L. Alan Winters, eds. 1995. "The Uruguay Round and the Developing Economies." Vol. 307. World Bank Publications.
- McCalman, Phillip. 2001. "Reaping What you Sow: An Empirical Analysis of International Patent Harmonization." *Journal of International Economics* 55(1): 161-186.
- McCalman, Phillip. 2005. "Who Enjoys 'TRIPS' Abroad? An Empirical Analysis of Intellectual Property Rights in the Uruguay Round." *Canadian Journal of Economics* 38(2): 574-603.
- Moscona, Jacob. 2021. "Flowers of Invention: Patent Protection and Productivity Growth in US Agriculture." Mimeo.
- National Science Foundation (NSF). 2013. "Business Research and Development and Innovation 2008-10." (accessed February 18, 2023).
- Negishi, Takashi. 1960. "Welfare Economics and Existence of an Equilibrium for a Competitive Economy." *Metroeconomica* 12(2-3): 92-97.
- Nordhaus, William D. 1969. "Invention, Growth and Welfare: A Theoretical Treatment of Technological Change." MIT Press.
- Organization for Economic Co-operation and Development (OECD). 2021. "Input-Output Tables." Ossa, Ralph, Robert Staiger, and Alan Sykes. 2023. "Standing in International Investment and Trade Disputes." *Journal of International Economics* 145: 103791.
- Park, Walter. 2008. "International Patent Protection: 1960–2005." *Research Policy* 37(4): 761-766.
- Park, Walter. 2010. "On Patenting Costs." *The WIPO Journal* 2(1): 38-48.
- PATSTAT. 2022. "Worldwide Patent Statistical Database." European Patent Office, Autumn 2022 Global Edition.
- Saggi, Kamal. 2016. "Trade, Intellectual Property Rights, and the World Trade Organization." *Handbook of Commercial Policy* 1: 433-512.
- Santacreu, Ana Maria. 2024. "Dynamic Gains from Trade Agreements with Intellectual Property Provisions." Mimeo.
- Schesch, Constantin. 2024. "Pseudospectral Methods for Continuous-Time Heterogeneous-Agent Models." *Journal of Economic Dynamics and Control* 163, 104856.
- Williams, Heidi. 2017. "How do Patents Affect Research Investments?" *Annual Review of Economics* 9: 441-69.
- World Intellectual Property Organization (WIPO). 2023. "WIPO Statistics Database." Last updated: February 2023.
- World Bank. 2023. "World Development Indicators." (accessed March 10, 2023).

Appendices

A Theory

A.1 Static Equilibrium

Final good output is given by $Y_{nt} = \prod_{s=0}^S \left(\frac{Y_{nt}^s}{\beta^s} \right)^{\beta^s}$. Therefore, cost minimization implies that sectoral output Y_n^s and the final good price index P_n satisfy:

$$P_n^s Y_n^s = \beta^s P_n Y_n, \quad (34)$$

$$P_n = \prod_{s=0}^S (P_n^s)^{\beta^s}. \quad (35)$$

Aggregating prices across source countries and varieties yields:

$$P_{Mn}^s = \left(\frac{\sigma^s}{\sigma^s - 1} \right) \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \right]^{\frac{1}{1-\sigma^s}} \left(\sum_{j=1}^N \Psi_{Mnj}^s (\Phi_{nj}^s)^{\frac{\sigma^s-1}{\theta^s}} \right)^{\frac{1}{1-\sigma^s}}, \quad (36)$$

$$P_{Cn}^s = \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \Psi_{Cn}^s \right]^{\frac{1}{1-\sigma^s}} \left(\sum_{j=1}^N \Phi_{nj}^s \right)^{-\frac{1}{\theta^s}}, \quad (37)$$

where $\Gamma(\cdot)$ is the Gamma function. Cost minimization using equation (1) then gives:

$$P_n^s = \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \right]^{\frac{1}{1-\sigma^s}} \left[\Psi_{Cn}^s \left(\sum_{j=1}^N \Phi_{nj}^s \right)^{\frac{\sigma^s-1}{\theta^s}} + \left(\frac{\sigma^s}{\sigma^s - 1} \right)^{1-\sigma^s} \sum_{j=1}^N \Psi_{Mnj}^s (\Phi_{nj}^s)^{\frac{\sigma^s-1}{\theta^s}} \right]^{\frac{1}{1-\sigma^s}}. \quad (38)$$

We now impose market clearing conditions. Income from producing sector s output in country i is divided between wages paid to production workers, expenditure on intermediate inputs and profits. Let L_{Yi}^s denote production labor and Q_i^s denote total intermediate input usage in country i and sector s . Each monopolist's profits equal a fraction $1/\sigma^s$ of sales and sales equal exports net of tariffs. Consequently, aggregate profits made by monopolists in i from sales to n are given by $X_{Mni}^s / [\sigma^s (1 + b_{ni}^s)]$. Setting production income equal to total sales in each sector then yields:

$$w_i L_{Yi}^s + P_i Q_i^s + \frac{1}{\sigma} \sum_{n=1}^N \frac{X_{Mni}^s}{1 + b_{ni}^s} = \sum_{n=1}^N \frac{X_{ni}^s}{1 + b_{ni}^s}, \quad (39)$$

where total exports X_{ni}^s from i to n equals the sum of Eaton-Kortum and Helpman-Krugman trade:

$$X_{ni}^s = X_{Cni}^s + X_{Mni}^s. \quad (40)$$

The variety production technology (3) implies that intermediate input expenditure equals a fraction $1 - \alpha^s$ of production costs meaning: $\alpha^s P_i Q_i^s = (1 - \alpha^s) w_i L_{Yi}^s$. Substituting this expression into equation (39) and rearranging yields:

$$L_{Yi}^s = \frac{\alpha^s}{w_i} \left(\sum_{n=1}^N \frac{X_{ni}^s}{1 + b_{ni}^s} - \frac{1}{\sigma} \sum_{n=1}^N \frac{X_{Mni}^s}{1 + b_{ni}^s} \right). \quad (41)$$

Since $X_{Mni}^0 = 0$ for all n, i , the equation above also holds in sector zero. Summing over sectors gives:

$$w_i L_{Yi} = \sum_{s=0}^S \alpha^s \sum_{n=1}^N \frac{X_{ni}^s}{1 + b_{ni}^s} - \frac{1}{\sigma^s} \sum_{s=1}^S \alpha^s \sum_{n=1}^N \frac{X_{Mni}^s}{1 + b_{ni}^s}, \quad (42)$$

Final good market clearing in each country requires that:

$$Y_i = C_i + \sum_{s=0}^S Q_i^s = C_i + \sum_{s=0}^S \frac{1 - \alpha^s}{\alpha^s} \frac{w_i}{P_i} L_{Yi}^s. \quad (43)$$

Accounting for trade in both varieties and patenting services and setting the trade balance plus imports equal to exports gives:

$$TB_i + P_i Y_i - B_i + \sum_{s=1}^S \sum_{n=1}^N w_n L_{in}^{s,e} = \sum_{s=0}^S \sum_{n=1}^N \frac{X_{ni}^s}{1 + b_{ni}^s} + \sum_{s=1}^S \sum_{n=1}^N w_i L_{ni}^{s,e}, \quad (44)$$

where aggregate tariff revenues B_i are given by:

$$B_i = \sum_{s=0}^S \sum_{n=1}^N \frac{b_{in}^s}{1 + b_{in}^s} X_{in}^s. \quad (45)$$

Letting $P_1 = 1$ be the numeraire, we can now define the static equilibrium.

Definition 1. Static equilibrium. Assume that the aggregate quality of products sold competitively Ψ_{Cn}^s and monopolistically Ψ_{Mni}^s , labor allocated to output production L_{Yi} and labor allocated to patent purchases $L_{in}^{s,e}$ are known for all countries i and n and sectors s . Then a static equilibrium is defined as a set of N wage rates w_n , N final good output levels Y_n and N final good price indices P_n that solve:

- N final good price index equations (35) subject to the normalization $P_1 = 1$; N income equals sales equations (42), and; N trade balance equations (44), where:
- P_n^s are defined in (38); Φ_{ni}^s are defined in (8); X_{ni}^s are defined in (40); B_i are defined in (45); X_{Mni}^s are defined in (7); P_{Mn}^s are defined in (36); X_{Cni}^s are defined in (9); P_{Cn}^s are defined in (37), and; Y_n^s are defined in (34). The labor allocation L_{Yi}^s is then given by (41) and aggregate consumption C_i by (43).

A.2 Dynamic Equilibrium

Intertemporal demand. Solving the representative agent's intertemporal optimization problem yields the Euler equation:

$$r_{nt} = \rho + \frac{1}{\gamma} \left(\frac{\dot{C}_{nt}}{C_{nt}} + \frac{\dot{P}_{nt}}{P_{nt}} \right), \quad (46)$$

and, letting W_{nt} denote the representative agent's assets, the transversality condition is:

$$\lim_{t \rightarrow \infty} \exp \left(- \int_{t_0}^t r_{nt} \tilde{dt} \right) W_{nt} = 0. \quad (47)$$

Patenting thresholds. A firm that innovates with quality ψ at time t_0 pays the patent preparation cost $w_i f_i^{s,o}$ if and only if:

$$\sum_{n | \psi \geq \psi_{ni}^{s,e*}} \left[\Psi^s \left(V_{nit_0}^{s,P}(\psi) - V_{nit_0}^{s,NP}(\psi) \right) - w_n f_n^{s,e} \right] \geq w_i f_i^{s,o}.$$

Let $n_i^* \equiv \arg \min_n \psi_{ni}^{s,e*}$ denote the country n with the lowest patenting threshold $\psi_{ni}^{s,e*}$ in equation (13). The left hand side of the expression above is strictly increasing in ψ whenever $\psi \geq \psi_{n_i^*}^{s,e*}$. Consequently, there exists a unique threshold $\psi_i^{s,o*}$ such that only firms with quality $\psi \geq \psi_i^{s,o*}$ pay the patent preparation cost. The threshold is given by:

$$\begin{aligned} \psi_i^{s,o*} &= \psi_{n_i^*}^{s,e*} \text{ if } \sum_n \max \left[\psi_{n_i^*}^{s,e*} \left(V_{nit_0}^{s,P}(1) - V_{nit_0}^{s,NP}(1) \right) - \frac{w_n f_n^{s,e}}{\Psi^s}, 0 \right] \geq \frac{w_i f_i^{s,o}}{\Psi^s}, \\ \sum_n \max \left[\psi_i^{s,o*} \left(V_{nit_0}^{s,P}(1) - V_{nit_0}^{s,NP}(1) \right) - \frac{w_n f_n^{s,e}}{\Psi^s}, 0 \right] &= \frac{w_i f_i^{s,o}}{\Psi^s} \quad \text{otherwise,} \end{aligned} \quad (48)$$

Value of invention. Quality is drawn from a Pareto distribution with scale parameter one and shape parameter k . Therefore, an innovator's expected total patenting costs per variety equal:

$$\sum_{n=1}^N (\psi_{ni}^{s*})^{-k} \frac{w_n f_n^{s,e}}{\Psi^s} + (\psi_i^{s,o*})^{-k} \frac{w_i f_i^{s,o}}{\Psi^s}.$$

The expected value of profits per variety that a time t innovator makes from sales to destination n is:

$$\int_1^{\psi_{ni}^{s*}} V_{nit}^{s,NP}(\psi) k \psi^{-k-1} d\psi + \int_{\psi_{ni}^{s*}}^{\infty} V_{nit}^{s,P}(\psi) k \psi^{-k-1} d\psi = \frac{k}{k-1} \left[V_{nit}^{s,NP}(1) \left(1 - (\psi_{ni}^{s*})^{-k+1} \right) + V_{nit}^{s,P}(1) (\psi_{ni}^{s*})^{-k+1} \right].$$

Summing this expression over n and subtracting expected patenting costs per variety yields that the expected value V_{it}^s of inventing a new variety at time t satisfies equation (15).

Labor market clearing. Innovation in country i and sector s occurs at rate $\eta_i^s (L_{Ri}^s)^{1-\kappa}$ and a fraction

$(\psi_{ni}^{s*})^{-k}$ of innovations are patented in country n . Therefore, total labor employed by firms in country i to purchase patents in country n satisfies:

$$L_{in}^{s,e} = \eta_i^s (L_{Ri}^s)^{1-\kappa} (\psi_{ni}^{s*})^{-k} f_n^{s,e}. \quad (49)$$

Likewise total labor employed in country i for the preparation of patent applications is:

$$L_i^{s,o} = \eta_i^s (L_{Ri}^s)^{1-\kappa} (\psi_i^{s,o*})^{-k} f_i^{s,o}. \quad (50)$$

The labor market clearing condition is then given by:

$$L_i = L_{Yi} + \sum_{s=1}^S L_{Ri}^s + \sum_{s=1}^S L_i^{s,o} + \sum_{s=1}^S \sum_{n=1}^N L_{ni}^{s,e}. \quad (51)$$

We can now define a dynamic equilibrium.

Definition 2. Dynamic equilibrium. A dynamic equilibrium is defined as a set of labor allocations to R&D, patenting and production L_{Ri}^s , $L_{in}^{s,e}$, $L_i^{s,o}$ and L_{Yi} ; aggregate qualities of Helpman-Krugman and Eaton-Kortum products Ψ_{Mni}^s , $\Psi_{Mni}^{s,NP}$, $\Psi_{Mni}^{s,P,ND}$, $\Psi_{Mni}^{s,P,D}$ and Ψ_{Cn}^s ; patenting thresholds ψ_{ni}^{s*} and $\psi_i^{s,o*}$; value functions $V_{nit}^{s,NP}(1)$, $V_{nit}^{s,P}(1)$ and V_{it}^s ; interest rates r_i ; wage rates w_i ; final good output levels Y_i , and; final good price indices P_i , such that in all time periods:

- w_i , Y_i and P_i obey a static equilibrium according to Definition 1;
- labor market clearing (51) holds with patenting employment given by (49) and (50) and L_{Ri}^s satisfies (4);
- $\Psi_{Mni}^{s,NP}$, $\Psi_{Mni}^{s,P,ND}$, $\Psi_{Mni}^{s,P,D}$ and Ψ_{Cn}^s satisfy the laws of motion in (17) – (20) and Ψ_{Mni}^s is given by (16);
- V_{it}^s is given by (15) and $V_{nit}^{s,NP}(1)$ and $V_{nit}^{s,P}(1)$ are defined by (10) and (11) with $\psi = 1$ and expected profits obeying (12);
- ψ_{ni}^{s*} and $\psi_i^{s,o*}$ are defined by (13), (14) and (48), and;
- r_i satisfies the Euler equation (46) and the transversality condition (47) holds.

A.3 Steady State

Growth rates. We solve for a steady state equilibrium. Labor market clearing (51) implies that the allocation of labor to R&D, patenting and production in each country is constant in steady state. Equations (49) and (50) then imply that the patenting thresholds ψ_{ni}^{s*} and $\psi_i^{s,o*}$ are constant.

Equations (5) and (16) imply that the aggregate qualities Ψ_{Mni}^s , $\Psi_{Mni}^{s,NP}$, $\Psi_{Mni}^{s,P,ND}$, $\Psi_{Mni}^{s,P,D}$ and Ψ_{Cn}^s all grow at rate g^s in steady state. Using equations (8), (36), (43) and (44) it then follows that the growth rates of wages w_i , final good output Y_i , consumption C_i , final good price indices P_i and trade flows X_{ni}^s are all constant across countries. Since the final good in country one is the numeraire, we must have that steady state final good prices are constant in all countries.

Let g denote the growth rate of consumption C_i . The final good clearing condition (43) implies that wages w_i and final good output Y_i also grow at rate g , while the trade balance condition (44) implies that

trade flows X_{ni}^s grow at rate g . Equation (8) then gives that Φ_{ni}^s grows at rate $-\alpha^s \theta^s g$ and combining this result with equation (38) yields that the sectoral price index P_n^s grows at rate:

$$g_{P^s} = \alpha^s g - \frac{g^s}{\sigma^s - 1}. \quad (52)$$

From equation (34) we have that sectoral output Y_n^s growth $g_{Y^s} = g - g_{P^s}$ and the final good production technology implies $g = \sum_{s=0}^S \beta^s g_{Y^s}$. Combining these equations implies that g satisfies equation (22).

The Euler equation (46) then implies that the steady state interest rate is given by:

$$r = \rho + \frac{g}{\gamma}, \quad (53)$$

and since total assets W_n grow at rate g the transversality condition is satisfied if and only if $r > g$, which requires $\rho > g(1 - 1/\gamma)$. We also note from (10), (11) and (15) that the value functions $V_{nit}^{s,NP}(1)$, $V_{nit}^{s,P}(1)$ and V_{it}^s grow at the same rate as expected profits $\mathbb{E}_z \pi_{ni}^s(1, z)$, which equals $g - g^s$ by (12).

Laws of motion for normalized aggregate qualities. Normalizing each of the aggregate quality variables by Ψ^s , the laws of motion in equations (17)–(20) can be rewritten as:

$$\begin{aligned} (g^s + \nu^s + \zeta^s) \tilde{\Psi}_{Mni}^{s,NP} &= \eta_i^s (L_{Ri}^s)^{1-\kappa} \frac{k}{k-1} \left[1 - (\psi_{ni}^{s*})^{1-k} \right] + \delta_n^s \tilde{\Psi}_{Mni}^{s,P,ND}, \\ (g^s + \delta_n^s + \nu^s + \zeta^s) \tilde{\Psi}_{Mni}^{s,P,ND} &= \eta_i^s (L_{Ri}^s)^{1-\kappa} \frac{k}{k-1} (\psi_{ni}^{s*})^{1-k}, \\ (g^s + \delta_n^s + \zeta^s) \tilde{\Psi}_{Mni}^{s,P,D} &= \nu^s \tilde{\Psi}_{Mni}^{s,P,ND}, \\ (g^s + \zeta^s) \tilde{\Psi}_{Cn}^s &= \sum_{i=1}^N \left(\nu^s \tilde{\Psi}_{Mni}^{s,NP} + \delta_n^s \tilde{\Psi}_{Mni}^{s,P,D} \right). \end{aligned} \quad (54)$$

Patenting thresholds. Substituting equations (23) and (24) into (48) and using the definition of normalized profits in (25), the threshold for paying the application preparation cost satisfies:

$$\begin{aligned} \psi_i^{s,o*} &= \psi_{n_i^* i}^{s,e*} \text{ if } \sum_n \max \left(\psi_{n_i^* i}^{s,e*} \Delta R_n^s \tilde{\pi}_{ni}^s - \tilde{w}_n f_n^{s,e}, 0 \right) \geq \tilde{w}_i f_i^{s,o}, \\ \sum_n \max \left(\psi_i^{s,o*} \Delta R_n^s \tilde{\pi}_{ni}^s - \tilde{w}_n f_n^{s,e}, 0 \right) &= \tilde{w}_i f_i^{s,o} \quad \text{otherwise.} \end{aligned} \quad (55)$$

Innovation. Free entry into innovation (4) requires:

$$(L_{Ri}^s)^\kappa = \eta_i^s \frac{\tilde{V}_i^s}{\tilde{w}_i}, \quad (56)$$

which determines L_{Ri}^s . We can now define a steady state equilibrium.

Definition 3. Steady state. A steady state equilibrium is defined as a set of labor allocations to R&D, patenting and production L_{Ri}^s , $L_{in}^{s,e}$, $L_i^{s,o}$ and L_{Yi} ; normalized aggregate qualities of Helpman-Krugman and Eaton-Kortum products $\tilde{\Psi}_{Mni}^s$, $\tilde{\Psi}_{Mni}^{s,NP}$, $\tilde{\Psi}_{Mni}^{s,P,ND}$, $\tilde{\Psi}_{Mni}^{s,P,D}$ and $\tilde{\Psi}_{Cn}^s$; patenting thresholds ψ_{ni}^{s*} and $\psi_i^{s,o*}$; normalized value functions $\tilde{V}_{ni}^{s,NP}(1)$, $\tilde{V}_{ni}^{s,P}(1)$ and \tilde{V}_i^s ; normalized wage rates \tilde{w}_i ; normalized final good

- output levels \tilde{Y}_i ; final good price indices P_i ; growth rates g^s and g , and; interest rate r such that:
- \tilde{w}_i , \tilde{Y}_i and P_i obey a static equilibrium according to Definition 1 (with all variables normalized);
 - labor market clearing (51) holds with patenting employment given by (49) and (50) and L_{Ri}^s given by (56);
 - $\tilde{\Psi}_{Mni}^{s,NP}$, $\tilde{\Psi}_{Mni}^{s,P,ND}$, $\tilde{\Psi}_{Mni}^{s,P,D}$ and $\tilde{\Psi}_{Cn}^s$ satisfy the laws of motion in (54) and $\tilde{\Psi}_{Mni}^s$ is given by the normalized version of (16);
 - \tilde{V}_{it}^s satisfies (27) and $\tilde{V}_{ni}^{s,NP}$ (1) and $\tilde{V}_{ni}^{s,P}$ (1) are given by the normalized versions of (23) and (24) with normalized profits obeying (25);
 - ψ_{ni}^{s*} and $\psi_i^{s,o*}$ are defined by (14), (26) and (55), and;
 - g^s and g are given by (21) and (22), r satisfies (53) and the transversality condition holds.

B Calibration

B.1 Patent Flows

We use PATSTAT (2022) to obtain data on applications for “Patent of Inventions” filed at patent offices around the world. We use applications rather than patent grants for two reasons. First, in general, data coverage is better for applications. Second, in some countries (e.g. Japan and Korea) patent applications are not granted unless they are challenged.

Patent applications covering the same invention can be grouped into families. Since we are interested in unique innovations, we aggregate patent applications to the level of DOCDB simple patent families. A DOCDB family is a collection of patent documents that are considered to cover a single invention and have the same priorities. Each application belongs to exactly one DOCDB family. We date each patent family to the year of the earliest filing date of the root priority application. We then use the steps below to compute bilateral patent flows by year at the family level from 1990 onwards.

Using the probability mappings from Lybbert and Zolas (2012) we map the CPC/IPC technology classes associated with each patent family to ISIC sectors. We then drop patent families for which all CPC/IPC codes map to our patenting sector with probability less than one half. This leads to us dropping around 5 percent of patent families. We keep patent families for which CPC/IPC codes are not recorded.

We determine the origin country for each patent family based on the location of applicants. When different applicants within a patent family have different origins, we assign the patent fractionally across origins based on the share of applicants from each origin that are listed on any application belonging to the family. When applicant information is not available, we use the location of inventors. When data on both applicants and inventors is missing, but all applications in the family are filed at the same patent office, we assign the origin of the patent family using the location of the patent office. Otherwise, we drop the patent family. This leads to us dropping around 1 percent of patent families.

We assign a patent family to a destination country if any of the applications belonging to the family are filed in the destination (including national phase entries for applications filed under the Patent Cooperation Treaty). For patents granted by the European Patent Office (EPO), we use data on PGFP (Post Grant Fees

Paid) events to determine which EPO countries the application is transferred to. For non-granted EPO applications, which account for around two-thirds of EPO applications, we use a machine learning algorithm to predict which countries each application would have been transferred to if granted. We train and test a multi-label classifier on granted EPO patents using the following family-level features: year, number of applicants, number of inventors, number of other patent offices applied to, number of citations, number of applications in the family, and share of other offices that have granted applications in the family. For applications filed at the Eurasian Patent Office we use the designated events ‘MM4A’ to allocate applications across destinations. For patents filed at the Organisation Africaine de la Propriété Intellectuelle, the African Regional Intellectual Property Organization and the GCC Patent Office we assume the application covers all member countries. Finally, since Europe and the rest of the world comprise many individual countries, we weight counts by GDP shares when aggregating patent flows into Europe and into the rest of the world.

PATSTAT has poor coverage of applications filed at the Indian Patent Office. Consequently, to compute patent flows into India we use data from WIPO (2023) on patent applications (direct and Patent Cooperation Treaty national phase entries) filed in India by applicant’s origin. The WIPO data is at the application (not family) level, includes patents in all sectors and assigns origin using the first named applicant on the root priority application. We adjust for these differences by using PATSTAT to construct origin-year specific deflators based on applications filed in other large developing countries (China, Brazil, Russia and Mexico). In 2015 the cross-origin averages of the deflators are: 1.02 applications per family; 1.04 ratio of all patent families to families that map to our patenting sector, and; 0.99 adjustment for assigning origin using first named applicant. For 1990-93 and 2002-04, country of origin is missing for more than 10 percent of applications filed in India. For these years, we impute the origin of applications with missing origins using the origin of applications in the closest year that has fewer than 10 percent of unknown origins. For example, country of origin is missing for around two-thirds of applications filed in India in 1992. We impute origin countries for these applications using the origin of applications filed in India in 1994.

In the 1990s, the US did not report information on non-granted patents for inclusion in PATSTAT. PATSTAT uses applications filed in other countries to infer the existence of non-granted US applications where possible, but its coverage is incomplete. We deduce from the time series of patent applications in the US in PATSTAT that information on non-granted patents became available between 1997 and 2001. Therefore, to correct for under-reporting, we inflate patent flows into the US before 2001. To calculate the inflation factor, we start by restricting the PATSTAT sample to patents granted in the US and computing bilateral patent flows by year for the restricted sample using the procedure described above. We then define the inflation factor for origin i in year t by computing the ratio of the restricted sample flow into the US from origin i in year t_0 to the full sample flow and then dividing this ratio by the same ratio computed in 2001. We set $t_0 = 1997$ when $t < 1998$ and $t_0 = t$ when $1998 \leq t \leq 2000$. The cross-origin average inflation factor for years before 1998 equals 1.48.

B.2 Calibration Moments: Data

Exact moment matching. We take bilateral tariffs at the SITC Revision 2 4-digit level from Caliendo et al. (2023). We use tariffs for 2010, which is the latest year available in their data. We aggregate tariffs to our country sample and to the Agriculture and Fishing, Mining and Quarrying, and Manufacturing industries using 2010 import-share weights provided by Caliendo et al. (2023). We then use these industry-level tariffs to compute tariffs for the patenting and no patenting sectors by assuming that there are no tariffs in services industries and using 2010 import-share weights calculated from the OECD’s Input-Output Tables 2021 (OECD 2021) to aggregate across industries within the patenting and no patenting sectors.

Using the input-output tables, expenditure shares are set to each sector’s share of world output, which gives $\beta^0 = 0.61$. We also calibrate labor’s share of production costs to equal the ratio of value-added to output by sector, which implies $\alpha^0 = 0.64$ and $\alpha^1 = 0.39$. And trade costs $\tau_{ni}^s = (1 + b_{ni}^s) \tilde{\tau}_{ni}^s$ are chosen such that the equilibrium trade flows exactly match observed trade shares X_{ni}^s / X_{nn}^s . Using observed tariffs b_{ni}^s , we then infer iceberg trade costs $\tilde{\tau}_{ni}^s$. Population is the working age population aged 15-64 from the World Development Indicators (World Bank 2023).

Share of innovations patented. Cohen, Nelson and Walsh (2000) survey R&D labs in US manufacturing in 1994. Weighting responses by R&D expenditure, they find that respondents apply for patents on 49 percent of their product innovations and 31 percent of their process innovations. To obtain our target that 40 percent of US innovations are patented domestically, we take the simple average across product and process innovations, which is consistent with data from Bena and Simintzi (2022) on the share of patents that correspond to process innovations.

Patenting costs. Park (2010) estimates the cost of a US patent application in 2010 to be \$17,078 measured in real 2005 dollars. We inflate this cost to 2015 using the growth rate of the US GDP deflator.

Prior to the introduction of the unitary European patent in June 2023, patents granted by the European Patent Office (EPO) were only protected in countries where the patent was validated, which required the payment of national fees. Inventors could also seek protection in individual European countries by filing applications with national patent offices directly. The absence of a unitary European patent increased the cost of patenting in Europe. Based on survey data from Berger (2004), Park (2010) reports that 32 percent of the cost of patenting through the EPO is due to national renewal fees and a further 23 to 27 percent of the cost is due to national validation fees. Moreover, the average EPO patent is only validated in seven member countries. Inflating Park’s estimate of total EPO fees by $1 / (1 - 0.32 - 0.25)$ and adjusting for the share of European GDP covered by the seven largest European economies in 2015 implies that a European patent is 3.59 times more expensive than a US patent (using Park’s estimate of US patent application costs for comparison). Therefore, for $j = e, o$ we set $f_{Europe}^j = 3.87 f_{US}^j h_{Europe}$, which implies $w_{Europe} f_{Europe}^j = 3.59 w_{US} f_{US}^j$ when wages are proportional to observed real GDP per capita.

Turnover in US imports. We compute the rate at which the origin of US imports switches across countries. In the model such a switch occurs when a Helpman-Krugman variety produced abroad becomes an Eaton-Kortum variety that is sourced from a different foreign country. We use US trade data at the 10-digit level from Schott (2008) and aggregate countries to the regions used in our calibration. For any base year t , we

restrict the sample to 10-digit products for which the US was a net importer in both t and $t + 5$ and for which the leading country (in terms of US imports) has an import share at least 10 percentage points higher than any other country in year t and in year $t + 5$ (noting that the leading country may differ across years). These conditions are chosen to identify products for which there is a clear market leader in US imports. We then define the turnover rate as the import-weighted share of products for which the leading country changes between t and $t + 5$. For consistency with the model, we use US imports inclusive of tariff duties to measure import values. The average turnover rate calculated over all base years from 2006-2016 is 7.13 percent.

Other moments. The target aggregate growth rate is computed by regressing the ratio of US real GDP to working age population on a time trend using data for 1980-2019 from the World Development Indicators (World Bank 2023). We construct sectoral price indices for the patenting and no patenting sectors using Bureau of Economic Analysis gross output price indices (BEA 2022). Industries are weighted using gross output shares in 2000 and we compute the trend growth in each sector from 1997-2019.

Country characteristics in 2015 from the World Development Indicators (World Bank 2023) are defined as follows. GDP and R&D expenditure are measured in current US dollars. The price level is the ratio of the PPP conversion factor of GDP to market exchange rates. For Europe and the rest of the world we take the GDP weighted average of the price level in all countries with available data. We compute real GDP as the ratio of GDP in current US dollars to the price level.

B.3 Calibration Moments: Model

This section derives expressions for the moments used in the simulated method of moments calibration. All moments are computed in the model's steady state equilibrium.

The patent flow from origin i to destination n is $PAT_{ni}^s = \eta_i^s (L_{Ri}^s)^{1-\kappa} (\psi_{ni}^{s*})^{-k}$. International patent shares are then given by $PAT_{ni}^s / \sum_{n=1, n \neq i}^N \sum_{i=1}^N PAT_{ni}^s$, while the share of domestic patents in inward patents for country n equals $PAT_{nn}^s / \sum_{i=1}^N PAT_{ni}^s$. In addition, the share of innovations patented in the US is given by $(\psi_{ii}^{s*})^{-k}$ for $i = US$. And we calculate total expenditure on domestic patent applications in the US as $\sum_{s=1}^S PAT_{ii}^s \tilde{w}_i f_i^{s,e}$ for $i = US$.

The private value of holding a patent in destination n for an invention of quality ψ invented at time t in origin i equals $\Psi^s \psi \left[V_{nit}^{s,P}(1) - V_{nit}^{s,NP}(1) \right]$. Therefore, the aggregate value of patents purchased by US innovators in the US at time t is:

$$\sum_{s=1}^S \eta_i^s (L_{Ri}^s)^{1-\kappa} \Psi^s \left[V_{nit}^{s,P}(1) - V_{nit}^{s,NP}(1) \right] \int_{\psi_{ni}^{s*}}^{\infty} \psi dH(\psi),$$

where $i = n = US$. We compute the value of patents relative to R&D expenditure in the US by taking the ratio of this expression to R&D expenditure $RD_i = \sum_{s=1}^S \left(w_i L_{Ri}^s + w_i L_i^{s,o} + \sum_{n=1}^N w_n L_{in}^{s,e} \right)$ for $i = US$. Taking the ratio allows us to write all variables in their normalized forms.

We match the turnover moment TO_n^s with $n = US$ to the model-implied value of US imports that switch origin between t and $t + 5$ relative to total US imports of products imported in both t and $t + 5$. In

the model, the US sources each variety from a single country and the origin of imports only changes when varieties switch from Helpman-Krugman to Eaton-Kortum products (due to either technology diffusion or patent expiration) and the previous monopolist's country is not the lowest cost Eaton-Kortum supplier. Let $\epsilon(x) \equiv 1 - e^{-x}$ and $\Delta t = 5$. Then a little calculation yields:

$$TO_n^s = \frac{\left(\sum_k \tilde{\Phi}_{nk}^s \right)^{\frac{\sigma^s-1}{\theta^s}-1} \sum_{i \neq n} \sum_{j \neq i, n} \left[\tilde{\Psi}_{Mni}^{s,NP} \epsilon(\nu \Delta t) + \tilde{\Psi}_{Mni}^{s,P,ND} \epsilon(\nu \Delta t) \epsilon(\delta_n^s \Delta t) + \tilde{\Psi}_{Mni}^{s,P,D} \epsilon(\delta_n^s \Delta t) \right] \tilde{\Phi}_{nj}^s}{\left[\begin{aligned} & \left(\sum_k \tilde{\Phi}_{nk}^s \right)^{\frac{\sigma^s-1}{\theta^s}-1} \tilde{\Psi}_{Cn}^s \tilde{\Phi}_{ni}^s \\ & + \left(\frac{\sigma^s}{\sigma^s-1} \right)^{1-\sigma^s} \left[\tilde{\Psi}_{Mni}^{s,NP} e^{-\nu \Delta t} + \tilde{\Psi}_{Mni}^{s,P,ND} \left(e^{-\nu \Delta t} + \epsilon(\nu \Delta t) e^{-\delta_n^s \Delta t} \right) + \tilde{\Psi}_{Mni}^{s,P,D} e^{-\delta_n^s \Delta t} \right] \left(\tilde{\Phi}_{ni}^s \right)^{\frac{\sigma^s-1}{\theta^s}} \\ & + \left(\sum_k \tilde{\Phi}_{nk}^s \right)^{\frac{\sigma^s-1}{\theta^s}-1} \sum_{j \neq n} \left[\tilde{\Psi}_{Mni}^{s,NP} \epsilon(\nu \Delta t) + \tilde{\Psi}_{Mni}^{s,P,ND} \epsilon(\nu \Delta t) \epsilon(\delta_n^s \Delta t) + \tilde{\Psi}_{Mni}^{s,P,D} \epsilon(\delta_n^s \Delta t) \right] \tilde{\Phi}_{nj}^s \end{aligned} \right]},$$

The aggregate growth rate equals g . Using equation (52), the difference between price growth in the non-patenting and patenting sectors is $g_{P0} - g_{P1} = (\alpha^0 - \alpha^1) g + g^1 / (\sigma^1 - 1) - g^0 / (\sigma^0 - 1)$.

The trade elasticity TE_{ni}^s for exports from country i to country n is defined as the negative of the elasticity of trade value X_{ni}^s to trade costs τ_{ni}^s . In our model, the trade elasticity is the trade-share weighted average of the trade elasticity for Helpman-Krugman products $\sigma^s - 1$ and the trade elasticity for Eaton-Kortum products θ^s , which gives:

$$TE_{ni}^s = \frac{\tilde{X}_{Mni}^s}{\tilde{X}_{ni}^s} (\sigma^s - 1) + \frac{\tilde{X}_{Cni}^s}{\tilde{X}_{ni}^s} \theta^s.$$

We target the average trade elasticity across all country pairs defined as $\frac{1}{N(N-1)} \sum_{n=1, n \neq i}^N \sum_{i=1}^N TE_{ni}^s$.

We define the nominal GDP of country i as $GDP_i = P_i C_i + TB_i + RD_i$. This allows us to compute R&D expenditure relative to GDP as RD_i / GDP_i and world real GDP shares as GDP_i / P_i divided by $\sum_{n=1}^N GDP_n / P_n$. Price levels relative to the US are given by P_i / P_{US} . Finally, we calculate world gross output X^W as:

$$X^W = \sum_{s=0}^S \sum_{n=1}^N \sum_{i=1}^N \frac{X_{ni}^s}{1 + b_{ni}^s}. \quad (57)$$

B.4 Steady State Solution Algorithm

Let $\tilde{Z}_i \equiv P_i \tilde{Y}_i$, $\tilde{J}_{ni}^s \equiv \tilde{\pi}_{ni}^s / \tilde{w}_i$ and $\tilde{\phi}_{ni}^s \equiv \left(\tilde{\Phi}_{ni}^s \right)^{\frac{1}{\theta^s}}$. We solve for the steady state equilibrium using a fixed point approach in the vector of fundamental variables $VF = \left(\tilde{w}_i, \tilde{Z}_i, L_{Ri}^s, \tilde{J}_{ni}^s, \tilde{\phi}_{ni}^s \right)$.

Given an initial guess for VF , we compute the auxiliary variables as follows. Profits: $\tilde{\pi}_{ni}^s = \tilde{w}_i \tilde{J}_{ni}^s$. Growth rates: equation (21) gives g^s , equation (22) gives g and equation (53) gives r . Patenting thresholds: equation (26) gives $\psi_{ni}^{s,e*}$, equation (55) gives $\psi_i^{s,o*}$ and equation (14) gives ψ_{ni}^{s*} . Aggregate qualities: equation (29) gives $\tilde{\Psi}_{Mni}^s$ for $s \neq 0$, $\tilde{\Psi}_{Mni}^0 = 0$ and $\tilde{\Psi}_{Cn}^s = 1 - \sum_{i=1}^N \tilde{\Psi}_{Mni}^s$. Sectoral relative prices:

$$\frac{P_{Mn}^s}{P_n^s} = \left[\frac{\left(\frac{\sigma^s}{\sigma^s-1}\right)^{1-\sigma^s} \sum_{j=1}^N \tilde{\Psi}_{Mnj}^s \left(\tilde{\phi}_{nj}^s\right)^{\sigma^s-1}}{\left(\frac{\sigma^s}{\sigma^s-1}\right)^{1-\sigma^s} \sum_{j=1}^N \tilde{\Psi}_{Mnj}^s \left(\tilde{\phi}_{nj}^s\right)^{\sigma^s-1} + \tilde{\Psi}_{Cn}^s \left(\sum_{j=1}^N \left(\tilde{\phi}_{nj}^s\right)^{\theta^s}\right)^{\frac{\sigma^s-1}{\theta^s}}} \right]^{\frac{1}{1-\sigma^s}}, \quad (58)$$

$$\frac{P_{Cn}^s}{P_n^s} = \left[\frac{\tilde{\Psi}_{Cn}^s \left(\sum_{j=1}^N \left(\tilde{\phi}_{nj}^s\right)^{\theta^s}\right)^{\frac{\sigma^s-1}{\theta^s}}}{\left(\frac{\sigma^s}{\sigma^s-1}\right)^{1-\sigma^s} \sum_{j=1}^N \tilde{\Psi}_{Mnj}^s \left(\tilde{\phi}_{nj}^s\right)^{\sigma^s-1} + \tilde{\Psi}_{Cn}^s \left(\sum_{j=1}^N \left(\tilde{\phi}_{nj}^s\right)^{\theta^s}\right)^{\frac{\sigma^s-1}{\theta^s}}} \right]^{\frac{1}{1-\sigma^s}}. \quad (59)$$

Labour allocations: equation (49) gives $L_{in}^{s,e}$, equation (50) gives $L_i^{s,o}$ and equation (51) gives L_{Yi} . Trade flows:

$$\tilde{X}_{Mni}^s = \frac{\tilde{\Psi}_{Mni}^s \left(\tilde{\phi}_{ni}^s\right)^{\sigma^s-1}}{\sum_{j=1}^N \tilde{\Psi}_{Mnj}^s \left(\tilde{\phi}_{nj}^s\right)^{\sigma^s-1}} \left(\frac{P_{Mn}^s}{P_n^s}\right)^{1-\sigma^s} \beta^s \tilde{Z}_n, \quad (60)$$

$$\tilde{X}_{Cni}^s = \frac{\left(\tilde{\phi}_{ni}^s\right)^{\theta^s}}{\sum_{j=1}^N \left(\tilde{\phi}_{nj}^s\right)^{\theta^s}} \left(\frac{P_{Cn}^s}{P_n^s}\right)^{1-\sigma^s} \beta^s \tilde{Z}_n, \quad (61)$$

$$\tilde{X}_{ni}^s = \tilde{X}_{Cni}^s + \tilde{X}_{Mni}^s. \quad (62)$$

Tariff revenues: normalized version of equation (45) gives \tilde{B}_i . World gross output: normalized version of equation (57) gives \tilde{X}^W . Final good price indices:

$$P_n = \prod_{s=0}^S \left\{ \Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \left[\left(\frac{\sigma^s}{\sigma^s-1} \right)^{1-\sigma^s} \sum_{j=1}^N \tilde{\Psi}_{Mnj}^s \left(\tilde{\phi}_{nj}^s\right)^{\sigma^s-1} + \tilde{\Psi}_{Cn}^s \left(\sum_{j=1}^N \left(\tilde{\phi}_{nj}^s\right)^{\theta^s}\right)^{\frac{\sigma^s-1}{\theta^s}} \right] \right\}^{\frac{\beta^s}{1-\sigma^s}}. \quad (63)$$

We then update the fundamental variables using:

$$\frac{(L_{Ri}^s)^\kappa}{\eta_i^s} = \sum_{n=1}^N \left[\frac{k}{k-1} \tilde{J}_{ni}^s R^{s,NP} + (\psi_{ni}^{s*})^{-k} \frac{\tilde{w}_n h_n f^{s,e}}{\tilde{w}_i} \left(\frac{k}{k-1} \frac{\psi_{ni}^{s*} \tilde{w}_i \tilde{J}_{ni}^s \Delta R_n^s}{\tilde{w}_n h_n f^{s,e}} - 1 \right) \right] - (\psi_i^{s,o*})^{-k} h_i f^{s,o},$$

$$\tilde{J}_{ni}^s = \frac{1}{1 + b_{ni}^s} \frac{\tilde{X}_{Mni}^s}{\sigma^s \tilde{\Psi}_{Mni}^s \tilde{w}_i},$$

$$\tilde{\phi}_{ni}^s = \frac{(T_n^s)^{\frac{1}{\theta^s}}}{\tilde{w}_n^{\alpha^s} P_n^{1-\alpha^s}} \left[\left(\frac{X_{ni}^s}{X_{nn}^s} \right)^{\text{Data}} (1 + b_{ni}^s) \frac{\frac{\tilde{\Psi}_{Mnn}^s (\tilde{\phi}_{nn}^s)^{\sigma^s - 1 - \theta^s}}{\sum_{j=1}^N \tilde{\Psi}_{Mnj}^s (\tilde{\phi}_{nj}^s)^{\sigma^s - 1}} \left(\frac{P_{Mn}^s}{P_n^s} \right)^{1-\sigma^s} + \frac{1}{\sum_{j=1}^N (\tilde{\phi}_{nj}^s)^{\theta^s}} \left(\frac{P_{Cn}^s}{P_n^s} \right)^{1-\sigma^s}}{\frac{\tilde{\Psi}_{Mni}^s (\tilde{\phi}_{ni}^s)^{\sigma^s - 1 - \theta^s}}{\sum_{j=1}^N \tilde{\Psi}_{Mnj}^s (\tilde{\phi}_{nj}^s)^{\sigma^s - 1}} \left(\frac{P_{Mn}^s}{P_n^s} \right)^{1-\sigma^s} + \frac{1}{\sum_{j=1}^N (\tilde{\phi}_{nj}^s)^{\theta^s}} \left(\frac{P_{Cn}^s}{P_n^s} \right)^{1-\sigma^s}} \right]^{\frac{1}{\theta^s}},$$

$$\tilde{w}_i = \frac{1}{L_{Yi}} \sum_{s=0}^S \alpha^s \sum_{n=1}^N \frac{1}{1 + b_{ni}^s} \left(\tilde{X}_{ni}^s - \frac{1}{\sigma^s} \tilde{X}_{Mni}^s \right), \quad (64)$$

$$\tilde{Z}_i = \sum_{s=0}^S \sum_{n=1}^N \frac{\tilde{X}_{ni}^s}{1 + b_{ni}^s} + \sum_{s=1}^S \sum_{n=1}^N \tilde{w}_i L_{ni}^{s,e} + \tilde{B}_i - \left(\frac{TB_i}{X^W} \right)^{\text{Data}} \tilde{X}^W - \sum_{s=1}^S \sum_{n=1}^N \tilde{w}_n L_{in}^{s,e}, \quad (65)$$

and iterate to stabilize VF using a fixed-point iterative algorithm. Note that the algorithm takes the trade shares $(X_{ni}^s/X_{nn}^s)^{\text{Data}}$ directly from the data since the trade costs τ_{ni}^s are chosen to match these shares exactly. We also take trade surpluses relative to world output $(TB_i/X^W)^{\text{Data}}$ from the data and hold these ratios fixed in the counterfactual analysis.

We implement the solution algorithm using a type-I stationary Anderson accelerated process following the implementation of Zhang et al. (2020). Convergence is ensured by enforcing a damping hyperparameter, which is weakly optimized following Chen and Vuik (2022). The advantage of this approach is that we do not need to compute any Jacobian or Hessian matrices, either analytically or numerically. We measure a time complexity of $O(N^2 S \ln S)$ for the solver up to 100 countries and 100 sectors.

After the iteration stabilizes, we apply the transformation $(\tilde{w}_i, \tilde{Z}_i, L_{Ri}^s, \tilde{J}_{ni}^s, \tilde{\phi}_{ni}^s) \rightarrow (\tilde{w}_i/P_1, \tilde{Z}_i/P_1, L_{Ri}^s, \tilde{J}_{ni}^s, \tilde{\phi}_{ni}^s P_1)$, which yields a solution that respects our numeraire condition $P_1 = 1$.

To solve for the steady state when undertaking counterfactual analysis, we follow the same procedure described above, except that when updating $\tilde{\phi}_{ni}^s$ we use:

$$\tilde{\phi}_{ni}^s = \frac{(T_i^s)^{\frac{1}{\theta^s}}}{\tau_{ni}^s \tilde{w}_i^{\alpha^s} P_i^{1-\alpha^s}}. \quad (66)$$

B.5 Calibration Algorithm

For the international patent shares moment, the loss function is:

$$\mathcal{L}^k(m_i^k(\Omega), m_i^{k,target}) = \log \left[\frac{m_i^k(\Omega)}{m_i^{k,target}} \right] \left(1 + \frac{\xi}{\left| \log(m_i^{k,target}) \right|} \right),$$

with $\xi = 6$. We introduce the skewering factor ξ to give more weight to larger patent flows. For the aggregate growth rate, price growth difference, trade elasticity and world output moments, the loss function is the absolute value of the difference of the ratio of the simulated and targeted moments from one. For all other moments, the loss function is the log difference of the simulated and targeted moments.

The weights are chosen to optimize the model’s match to the targeted moments using an informal application of the epsilon constraint method of multi-objective optimization. The R&D expenditure relative to GDP moment has weight 10. The turnover, aggregate growth rate, trade elasticity and world output moments have weight 5. All other moments have weight one.

The calibration uses a trust-region algorithm, which we find to be more robust and less prone to finding a local minimum than variations of Newton or gradient-descent methods. We use the formulation of trust-region sub-problems of Branch et al. (1999), and the solving of the sub-problems in the trust regions uses the Levenberg-Marquardt algorithm from Moré (2006). We use the reflective characterization of the trust-region algorithm in Coleman and Li (1996) to avoid stepping directly into bounds.

B.6 Pre-TRIPS Calibration

For the pre-TRIPS calibration, we divide parameters into two groups: parameters that we allow to vary over time, which we calibrate in 1992, and; other parameters, which we hold fixed at their values from the baseline 2015 calibration. The time-varying parameters are: TB_i , β^s , α^s , $\tilde{\tau}_{ni}^s$, b_{ni}^s , L_i , h_i , δ_i , η_i and T_i^s . As before, we use exact moment matching to infer TB_i , β^s , α^s , $\tilde{\tau}_{ni}^s$, b_{ni}^s and L_i and differences in real GDP per capita to determine the patenting cost adjustment h_i .

We calibrate δ_i , η_i and T_i^s using simulated method of moments estimation as in the baseline calibration, except that we only target a subset of the baseline moments. We target those moments that we observe pre-TRIPS and that are informative about patent protection, innovation and productivity. In particular, we use: international patent shares; the share of domestic patents in inward patents in the US; expenditure on domestic patent applications in the US; world output; R&D expenditure relative to GDP in the US, Europe, Japan, Canada and Korea; world real GDP shares, and; price indices relative to the US. Implementing this partial calibration approach using our 2015 data yields patent protection estimates δ_i that are indistinguishable from the baseline estimates in Table 2 to at least two significant figures.

The data moments for the 1992 calibration are calculated following the same procedures used for the 2015 calibration with the following exceptions. To compute expenditure on domestic patent applications in the US in 1992, we deflate Park’s (2010) estimate of the cost of a US patent application from 2005 dollars to 1992 dollars using the US GDP deflator. The OECD data (OECD 2021) used to compute trade, output, expenditure and intermediate input costs is from 1995, the earliest year available. Likewise, we measure R&D expenditure relative to GDP in US, Europe, Japan, Canada and Korea in 1996, the earliest year for which it is available in the World Development Indicators (World Bank 2023). Finally, when aggregating the tariffs from Caliendo et al. (2023) to our patenting and non-patenting sectors we use the same 2010 import weights used for the baseline calibration to ensure that tariff differences before and after TRIPS are not caused by changes in trade patterns.

C Counterfactual Analysis

C.1 Transition Dynamics

Suppose there is an unanticipated change in one or more parameters at time zero and that the economy was in steady state before time zero. To characterize the transition dynamics between steady states we need to derive expressions for the time derivatives of the value functions $\tilde{V}_{nit}^{s,NP}(1)$, $\tilde{V}_{nit}^{s,P}(1)$ and $\tilde{V}_{nit}^{s,P,D}(1)$, where $\tilde{V}_{nit}^{s,P,D}(1)$ denotes the expected present discounted value of profits per variety that a firm from country i makes in destination n if at time t it owns a non-expired patent over an invention with quality one for which the technology has already diffused. We have:

$$\tilde{V}_{nit}^{s,P,D}(1) = \frac{\Psi_t^s}{\Psi_t} \int_t^\infty \mathbb{E}_z \pi_{nit}^s(1, z) \exp\left(-\int_t^{\hat{t}} (r_{it} + \zeta^s + \delta_n^s) d\hat{t}\right) d\hat{t},$$

and differentiating this expression with respect to t yields:

$$\dot{\tilde{V}}_{nit}^{s,P,D}(1) = (r_{it} + \zeta^s + \delta_n^s + g_t^s - g_t) \tilde{V}_{nit}^{s,P,D}(1) - \tilde{\pi}_{nit}^s. \quad (67)$$

Likewise, differentiating (10) and (11) gives:

$$\dot{\tilde{V}}_{nit}^{s,NP}(1) = (r_{it} + \zeta^s + \nu^s + g_t^s - g_t) \tilde{V}_{nit}^{s,NP}(1) - \tilde{\pi}_{nit}^s, \quad (68)$$

$$\dot{\tilde{V}}_{nit}^{s,P}(1) = (r_{it} + \zeta^s + \nu^s + \delta_n^s + g_t^s - g_t) \tilde{V}_{nit}^{s,P}(1) - \nu^s \tilde{V}_{nit}^{s,P,D}(1) - \delta_n^s \tilde{V}_{nit}^{s,NP}(1) - \tilde{\pi}_{nit}^s. \quad (69)$$

To solve for the transition dynamics, we use a fixed point algorithm with fundamental variables $\tilde{\Psi}_{Cnt}^s$, $\tilde{\Psi}_{Mnit}^{s,NP}$, $\tilde{\Psi}_{Mnit}^{s,P,ND}$, $\tilde{\Psi}_{Mnit}^{s,P,D}$, $\tilde{V}_{nit}^{s,NP}(1)$, $\tilde{V}_{nit}^{s,P}(1)$, $\tilde{V}_{nit}^{s,P,D}(1)$, P_{it} , \tilde{w}_{it} and $\tilde{Z}_{it} \equiv P_{it} \tilde{Y}_{it}$. We start by guessing time paths for the fundamental variables on the time interval $[0, T]$ under the assumption that the economy is in the new steady state from time $T/2$ onwards and that at time zero the state variables $\tilde{\Psi}_{Cnt}^s$, $\tilde{\Psi}_{Mnit}^{s,NP}$, $\tilde{\Psi}_{Mnit}^{s,P,ND}$, $\tilde{\Psi}_{Mnit}^{s,P,D}$ equal their values in the old steady state. In practice, we set $T = 500$ and our results show that the economy is always extremely close to the new steady state after 100 years.

Given our initial guess for the fundamental variables, we compute the auxiliary variables as follows. Equation (66) gives $\tilde{\phi}_{nit}^s$. Normalized version of equation (16) gives $\tilde{\Psi}_{Mnit}^s$. Equation (58) gives P_{Mnt}^s/P_{nt}^s and equation (59) gives P_{Cnt}^s/P_{nt}^s . Equation (13) gives $\psi_{nit}^{s,e*}$, equation (48) gives $\psi_{it}^{s,o*}$ and equation (14) gives ψ_{nit}^{s*} . Normalized version of equation (15) gives \tilde{V}_{it}^s . Equation (56) gives L_{Rit}^s . Equation (21) gives g_t^s and equation (22) gives g_t . Equation (49) gives $L_{int}^{s,e}$, equation (50) gives $L_{it}^{s,o}$ and equation (51) gives L_{Yit} . Equation (60) gives \tilde{X}_{Mnit}^s , equation (61) gives \tilde{X}_{Cnt}^s and equation (62) gives \tilde{X}_{nit}^s . Normalized version of equation (45) gives \tilde{B}_i . Normalized version of equation (57) gives \tilde{X}^W . Equation (25) gives $\tilde{\pi}_{nit}^s$. Normalized version of equation (43) gives $P_{it} \tilde{C}_{it}$ and equation (46) gives:

$$r_{it} = \rho + \frac{1}{\gamma} \left[\frac{\frac{\partial}{\partial t} (P_{it} \tilde{C}_{it})}{P_{it} \tilde{C}_{it}} + g_t \right].$$

Next, we compute numerical derivatives and update the fundamental variables. We calculate numerical derivatives using a pseudo-spectral (or projection) method that projects functions on the space of Chebyshev polynomials. We use Chebyshev polynomials for the projection because their differentiation matrix allows us to efficiently impose boundary conditions for state variables, ensuring continuity with the steady state at the initial time and a zero time derivative at infinity. We work with 25 Chebyshev collocation nodes.

Using equations (17)-(20) we set:

$$\begin{aligned} \tilde{\Psi}_{Cnt}^s &= \frac{1}{\zeta^s + g_t^s} \left[\sum_{i=1}^N \left(\nu^s \tilde{\Psi}_{Mnit}^{s,NP} + \delta_n^s \tilde{\Psi}_{Mnit}^{s,P,D} \right) - \dot{\psi}_{Cnt}^s \right], \\ \tilde{\Psi}_{Mnit}^{s,NP} &= \frac{1}{\zeta^s + \nu^s + g_t^s} \left(\eta_i^s (L_{Rit}^s)^{1-\kappa} \frac{k}{k-1} \left[1 - (\psi_{nit}^{s*})^{1-k} \right] + \delta_n^s \tilde{\Psi}_{Mnit}^{s,P,ND} - \dot{\psi}_{Mnit}^{s,NP} \right), \\ \tilde{\Psi}_{Mnit}^{s,P,ND} &= \frac{1}{\zeta^s + \nu^s + \delta_n^s + g_t^s} \left(\eta_i^s (L_{Rit}^s)^{1-\kappa} \frac{k}{k-1} \left[1 - (\psi_{nit}^{s*})^{1-k} \right] - \dot{\psi}_{Mnit}^{s,P,ND} \right), \\ \tilde{\Psi}_{Mnit}^{s,P,D} &= \frac{1}{\zeta^s + \delta_n^s + g_t^s} \left(\nu^s \tilde{\Psi}_{Mnit}^{s,P,ND} - \dot{\psi}_{Mnit}^{s,P,D} \right). \end{aligned}$$

And from equations (67)-(69) we set:

$$\begin{aligned} \tilde{V}_{nit}^{s,P,D}(1) &= \frac{1}{r_{it} + \zeta^s + \delta_n^s + g_t^s - g_t} \left[\tilde{\pi}_{nit}^s + \dot{\tilde{V}}_{nit}^{s,P,D}(1) \right], \\ \tilde{V}_{nit}^{s,NP}(1) &= \frac{1}{r_{it} + \zeta^s + \nu^s + g_t^s - g_t} \left[\tilde{\pi}_{nit}^s + \dot{\tilde{V}}_{nit}^{s,NP}(1) \right], \\ \tilde{V}_{nit}^{s,P}(1) &= \frac{1}{r_{it} + \zeta^s + \nu^s + \delta_n^s + g_t^s - g_t} \left[\tilde{\pi}_{nit}^s + \nu^s \tilde{V}_{nit}^{s,P,D} + \delta_n^s \tilde{V}_{nit}^{s,NP}(1) + \dot{\tilde{V}}_{nit}^{s,P}(1) \right]. \end{aligned}$$

Finally, we update P_{it} using equation (63), \tilde{w}_{it} using equation (64) and \tilde{Z}_{it} using equation (65). We iterate this procedure until the fundamental variables stabilize using an Anderson-accelerated fixed point routine as used in the steady state solver.

C.2 Unilateral Patent Policy

Figure 5 plots welfare changes EV_i for each country, and for the world as a whole, against proportional changes in patent protection δ_n provided by Europe, Japan, Brazil and India. Figure 5 is equivalent to Figure 4, which shows welfare changes when the US and China vary patent protection.

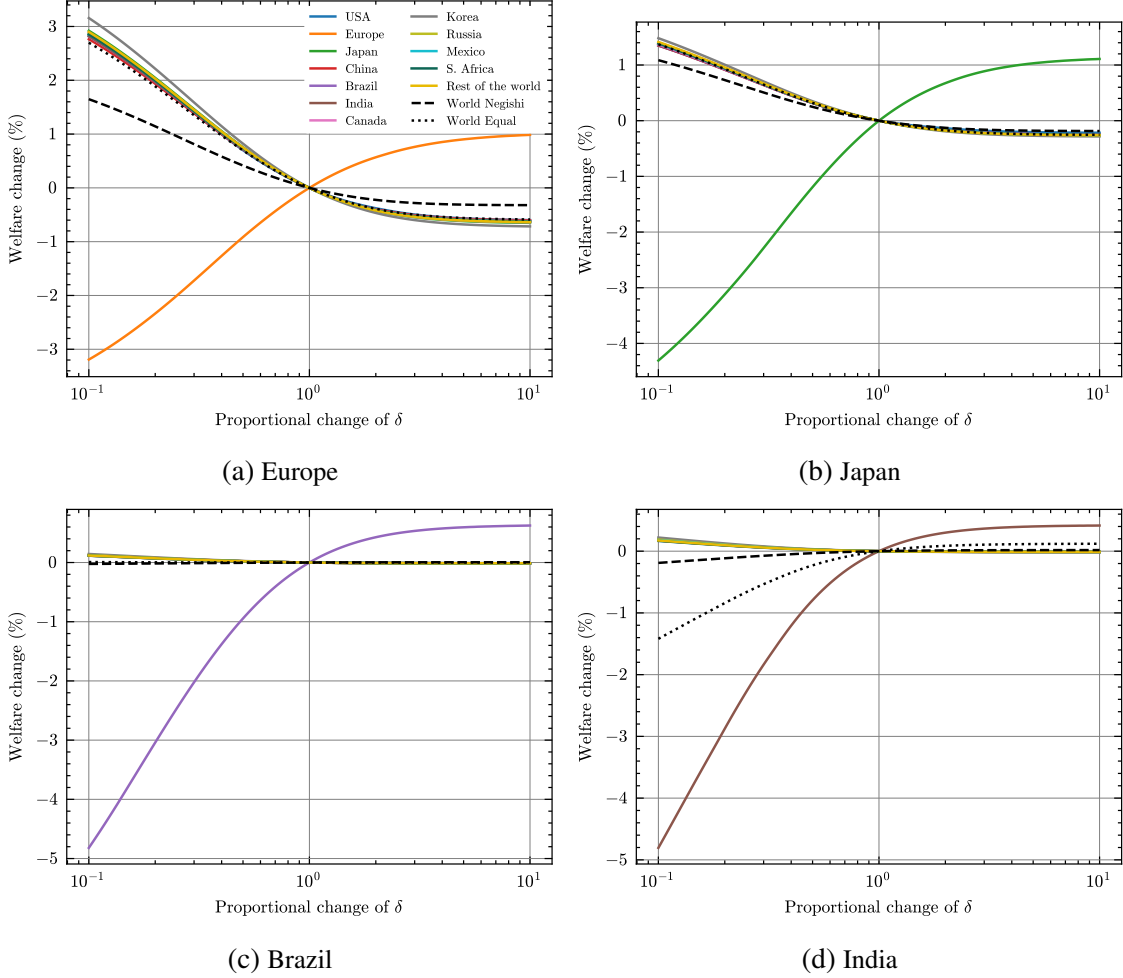


Figure 5: More welfare effects of unilateral patent policy changes

C.3 Nash and Cooperative Equilibria


We solve for the Nash equilibrium using an iterative algorithm. Given any set of δ_n for all n , for each country i we find δ'_i that maximizes welfare (including transition dynamics) in country i when country i changes its patent protection from δ_i to δ'_i and all other countries keep δ_n unchanged. Then we update δ_n to δ'_n for all n and iterate until convergence. The main challenges in implementing this algorithm are that welfare may be non-monotonic in δ'_i and there may be local maxima. Consequently, we need a global optimizer. For the case without transition dynamics, we use a Brent (2013) algorithm. For the case with transition dynamics, we use a simplicial homology global optimization algorithm, which is well-suited for blackbox optimization of low dimensional problems to global optimality (Endres et al. 2018).

Solving for the cooperative equilibrium is straightforward because the search space is bounded and low

dimensional. We use the quasi-Newton method of Broyden, Fletcher, Goldfarb and Shanno as described in Nocedal and Wright (1999).

C.4 Trade Costs and Technology Diffusion Rate Counterfactuals

Figure 6 shows how the welfare effects of patent protection depend on trade costs. It plots the derivative of each country's welfare $EV_i - 1$ with respect to its own protection δ_i as a function of iceberg trade costs in the patenting sector. We see that the welfare gains from unilaterally weakening patent protection (by increasing δ_i) tend to be greater when trade costs are lower.




Notes: This figure plots the derivative of welfare in country i with respect to patent protection in country i as a function of the proportional change (relative to the calibrated steady state) in international iceberg trade costs in the patenting sector. Welfare changes are expressed as the equivalent variation in consumption and account for transition dynamics.

Figure 6: Welfare effects and trade costs

While Figure 6 shows the welfare effects of marginal changes in patent protection, Figure 7 reports larger changes. Panel (a) plots EV_{US} against δ_{US} for the baseline calibration and two variants: a free trade economy with no iceberg trade costs or tariffs, and; an economy where we double the rate of technology diffusion ν . The figure confirms that removing trade costs increases the welfare losses to the US from strengthening domestic protection. But in contrast to the baseline calibration, US welfare is approximately flat in δ_{US} when we double ν . However, for all countries other than the US, stronger patent protection continues to reduce domestic welfare even when ν is doubled. Panel (b) illustrates this observation by plotting how European welfare depends upon the strength of European patent protection in each case.

Table 5 summarizes the Nash equilibria and the equal weights cooperative equilibria for the economies with free trade or doubled technology diffusion rate. The table shows whether any countries change their optimal policy compared to the baseline calibration (as reported in Table 3) and how moving to the Nash



Notes: This figure plots the effect of proportional changes in calibrated patent protection δ_i on welfare in country i for the baseline 2015 calibration and two variants of the baseline: a free trade economy with no trade costs or tariffs, and; an economy where the rate of technology diffusion ν is doubled. Plots vary δ_{US} in panel (a) and δ_{Europe} in panel (b). The legend in panel (a) applies to both panels. Welfare changes are expressed as the equivalent variation in consumption relative to each calibration's steady state and account for transition dynamics.

Figure 7: Unilateral patent policy and welfare: free trade and doubling technology diffusion rate

and cooperative equilibria affects US and equal weights world welfare. The only country that ever changes its baseline policy in the Nash equilibrium is the US, which provides complete patent protection when the technology diffusion rate is doubled. And since the US provides complete patent protection, growth increases, which raises world welfare. The cooperative equilibrium with faster technology diffusion is similar to the baseline case, except that the gains from cooperation are larger because faster technology diffusion raises the value of patenting. Under free trade, Brazil, Russia and South Africa join the group that provides complete patent protection in the cooperative equilibrium, since trade integration increases the sensitivity of innovation to foreign patent policy.

D Sensitivity and Robustness Checks

Figure 8 plots the sensitivity of the baseline parameter estimates in Table 2 to varying the targeted moments and the externally calibrated parameters (θ^0 , γ , κ , ρ and σ^1). The columns show the estimated parameters and the rows correspond to targeted moments and externally calibrated parameters. Each row reports the elasticities of the estimated parameters to varying one of the targeted moments or externally calibrated parameters (relative to the baseline 2015 calibration). For example, the first row shows sensitivity to varying θ^0 . Darker reds indicate more positive elasticities and darker blues more negative elasticities. To simplify the presentation, we vary all cross-border patent flows into, or out of, a given country at the same time. Thus, the row labelled “Patents in USA” reports elasticities when we increase all patent flows into the US by the same proportion and then recompute the international patent shares.

Table 5: Nash and cooperative equilibria: free trade and doubling technology diffusion rate


	(a) Nash			(b) Cooperative: equal		
	Changes in patent protection from baseline	Welfare change (percent)		Changes in patent protection from baseline	Welfare change (percent)	
Calibration		US	World		US	World
Baseline	n/a	-0.5	-2.1	n/a	1.6	8.7
No trade costs or tariffs	No changes	-0.5	-1.8	Brazil, Russia, South Africa complete	0.2	5.2
Double technology diffusion rate, v	US complete	-1.9	6.1	No changes	4.4	14.3

Notes: This table reports counterfactual results for two alternatives to the baseline calibration: an economy with no trade costs or tariffs, and; an economy where the rate of technology diffusion v is doubled. Column (a) reports results for the Nash equilibrium. Column (b) reports results for the cooperative equilibrium when the social planner uses equal weights for all individuals. Changes in patent protection are relative to the baseline patent policies in the Nash and cooperative equilibria reported in Table 3. Welfare changes are expressed as the equivalent variation in consumption relative to each calibration's steady state in 2015 and account for transition dynamics. Welfare changes are aggregated across countries using equal weights. Complete protection corresponds to $\delta_i = 0.01$, which is the lower bound we impose on our solutions.

We highlight two patterns in Figure 8. First, as expected, the international patent shares and R&D expenditure ratios play an important role in calibrating patent protection δ_i and R&D efficiency η_i . Except for the US, δ_i is strongly decreasing in the share of international patent flows into country i , i.e. more cross-border patenting implies stronger calibrated protection. As noted in Section 3.2, the level of δ_{US} is not pinned down by patent flows because international patent shares depend upon the relative strength of patent protection in different countries. Instead, Figure 8 shows that δ_{US} is sensitive to the moments capturing the value of patents relative to R&D expenditure in the US, turnover in US imports and the trade elasticity in the patenting sector, as well as to the parameter σ^1 . We also see that for developing countries η_i is strongly increasing in the share of international patent flows originating in country i , while for developed economies η_i is more sensitive to the R&D expenditure ratios.

Second, Figure 8 shows that our estimates are relatively more sensitive to three of the externally calibrated parameters (κ , ρ and σ^1) and five of the targeted moments (the value of patents relative to R&D expenditure in the US, the trade elasticity in the patenting sector, turnover in US imports, expenditure on domestic patent applications in the US, and the aggregate growth rate). Therefore, to study the robustness of our counterfactual results, we analyze the consequences of increasing or reducing each of these eight parameters/moments by 20 percent, while keeping all other calibration inputs unchanged.³⁵ For each robustness check, we recalibrate the model following the same steps used for the baseline calibration in Section 3.2

³⁵When varying σ^1 we change the mark-up ratio $1/(\sigma^1 - 1)$ by 20 percent and when varying ρ we change the steady state interest rate $r = \rho + g/\gamma$ by 20 percent.

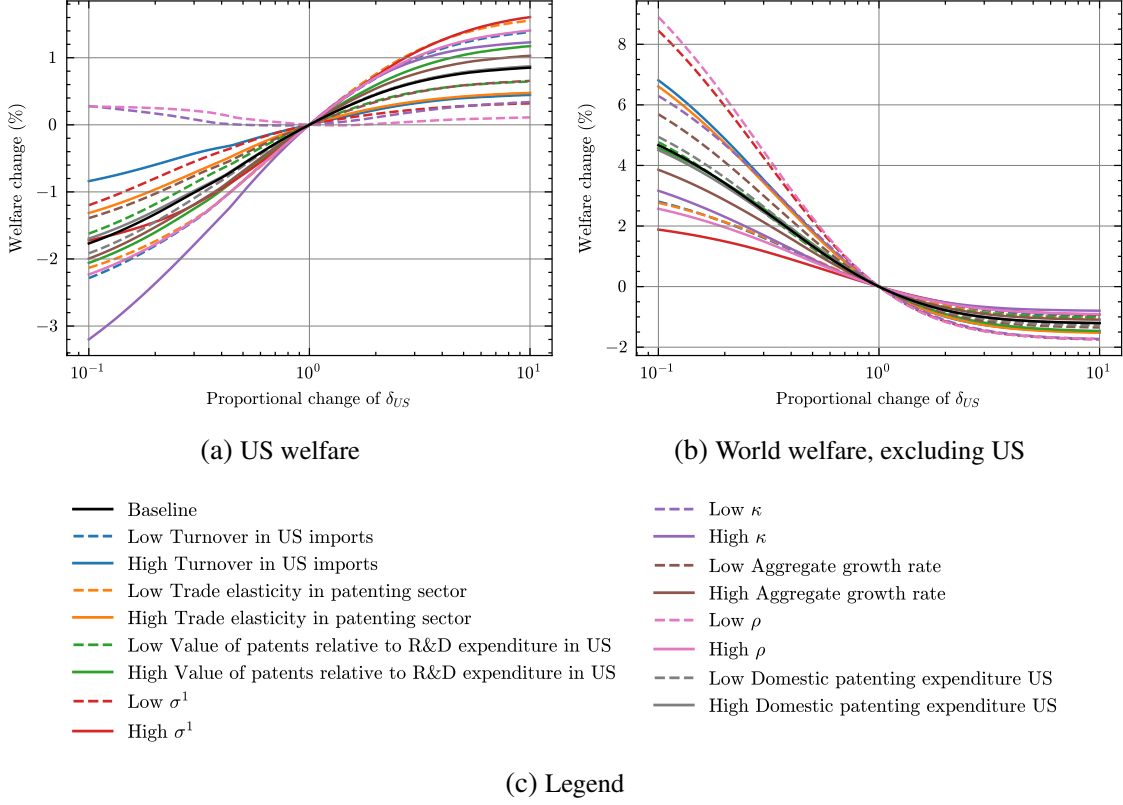


Notes: This figure shows elasticities of simulated method of moments estimates (in columns) to externally calibrated parameters and targeted moments (in rows). Colormap range is bounded at the 2nd and 98th percentiles of the elasticity distribution. For the international patenting moments, we vary all patent flows into (or out of) a given country by the same proportion. Patenting sector productivity in country i defined as $(T_i^1)^{\frac{1}{\theta^1}}$. No patenting sector productivity is defined as $(T_i^0)^{\frac{1}{\theta^0}}$.

Figure 8: Sensitivity analysis

and then compute the unilateral, Nash, cooperative and TRIPS counterfactuals as in Sections 4.1–4.4. The results of the robustness checks are summarized in Figure 9 and Table 6. When discussing these robustness checks we always use equal weights to aggregate welfare changes across countries

Figure 9 plots counterfactual results for unilateral changes in US patent policy like in Section 4.1. Panel (a) plots changes in US welfare EV_{US} against δ_{US} for the baseline calibration and each of the sixteen robustness checks (the legend is in panel c). Note that each line in panel (a) reports welfare changes starting from a different calibration of the model. Stronger patent protection reduces US welfare in all cases, except in the low κ and low ρ scenarios where welfare is U-shaped as a function of δ_{US} . Panel (b) plots the welfare effects of varying δ_{US} on the world excluding the US. As in the baseline calibration, we find that stronger



Notes: This figure plots the effect of proportional changes in calibrated patent protection in the US δ_{US} on welfare for alternative calibrations of the model. Welfare changes are expressed as the equivalent variation in consumption relative to each calibration's steady state and account for transition dynamics. Welfare changes aggregated across countries using equal weights.

Figure 9: Robustness checks for unilateral patent policy counterfactuals

protection in the US raises welfare abroad by increasing the growth rate. For countries other than the US, increasing the strength of domestic patent protection reduces domestic welfare and raises foreign welfare in the baseline calibration and in all of the robustness checks.

Table 6 summarizes the results of the robustness checks for the Nash, cooperative and TRIPS counterfactuals. For each robustness check, we report whether any countries change their patent protection in the Nash (column a) and equal weights cooperative equilibria (column b) compared to their choices in the baseline calibration (the baseline Nash and cooperative choices are shown in Table 3). We also report how moving to the Nash and cooperative equilibria affects world welfare. For TRIPS, we report the welfare effects of reverting to pre-TRIPS patent protection in developing countries (column c).

For the Nash equilibrium, the only country that ever changes its baseline policy is the US, which deviates in the low κ and low ρ scenarios. Consistent with the unilateral counterfactuals in Figure 9 the US chooses strong protection in both these cases; although with low κ optimal US protection is interior ($\delta_{US} = 0.016$). Because the US provides strong protection growth increases. Therefore, in these cases welfare for countries outside the US, and for the world as a whole, is higher in the Nash equilibrium than in the calibrated steady

Table 6: Robustness checks

Robustness check	(a) Nash		(b) Cooperative: equal		(c) Pre-TRIPS		
	Changes in patent protection from baseline	World welfare change (percent)	Changes in patent protection from baseline	World welfare change (percent)	Welfare change (percent)		
					Developing	Developed	World
Baseline	n/a	-2.1	n/a	8.7	0.25	-0.12	0.06
Low Turnover in US imports	No changes	-1.5	No changes	4.8	0.29	-0.09	0.10
High Turnover in US imports	No changes	-2.8	No changes	13.4	0.21	-0.17	0.03
Low Trade elasticity in patenting sector	No changes	-1.6	No changes	4.7	0.22	-0.05	0.09
High Trade elasticity in patenting sector	No changes	-2.8	No changes	13.2	0.22	-0.18	0.02
Low Value patents to R&D expenditure in US	No changes	-1.8	No changes	9.1	0.19	-0.09	0.05
High Value patents to R&D expenditure in US	No changes	-2.6	No changes	8.4	0.34	-0.18	0.08
Low σ^1	No changes	-3.2	No changes	17.4	0.19	-0.21	-0.01
High σ^1	No changes	-1.4	No changes	2.9	0.32	-0.08	0.12
Low κ	$\delta_{US} = 0.016$	3.5	No changes	11.4	0.21	-0.18	0.02
High κ	No changes	-1.3	No changes	5.9	0.33	-0.11	0.11
Low Aggregate growth rate	No changes	-2.4	No changes	10.6	0.23	-0.14	0.04
High Aggregate growth rate	No changes	-1.9	No changes	7.1	0.26	-0.11	0.08
Low ρ	US complete	7.6	No changes	16.6	0.17	-0.19	-0.01
High ρ	No changes	-1.4	No changes	4.7	0.28	-0.08	0.10
Low Domestic patenting expenditure in US	No changes	-2.3	No changes	9.7	0.24	-0.12	0.06
High Domestic patenting expenditure in US	No changes	-2.1	No changes	8.0	0.25	-0.12	0.06

Notes: This table reports counterfactual results for different calibrations of the model. Column (a) reports results for the Nash equilibrium. Column (b) reports results for the cooperative equilibrium when the social planner uses equal weights for all individuals. Column (c) reports welfare effects when developing countries revert to their pre-TRIPS patent protections level. Changes in patent protection are relative to the baseline patent policies in the Nash and cooperative equilibria reported in Table 3. Developing countries are: China, Brazil, India, Russia, Mexico and South Africa. Developed countries are: US, Europe, Japan, Canada and Korea. Welfare changes are expressed as the equivalent variation in consumption relative to each calibration's steady state in 2015 and account for transition dynamics. Welfare changes are aggregated across countries using equal weights. Complete protection corresponds to $\delta_i = 0.01$, which is the lower bound we impose on our solutions.

state. However, moving to Nash still reduces US welfare because the US bears the static costs of stronger protection. In all other robustness checks, moving to the Nash equilibrium reduces world welfare.

For the equal weights cooperative equilibrium, no country deviates from its baseline optimal patent protection in any of the robustness checks. World welfare gains in the cooperative equilibrium are always positive, but the magnitude of the gains differs across calibrations. Gains are largest when the discount rate ρ is low (implying future growth is more valuable), the demand elasticity σ^1 is low (implying higher mark-ups), or turnover in US imports is high (implying faster technology diffusion).

For TRIPS, we recalibrate the model for each robustness check using 1992 data as in Section 4.4. Then

we simulate a return to pre-TRIPS protection levels in developing countries as in column (b) of Table 4. Column (c) of Table 6 reports welfare changes in this counterfactual aggregated across developing countries, developed countries and the entire world. In all of the robustness checks, the pre-TRIPS calibration implies that patent protection in China, India and Russia was weaker in 1992 than 2015, whereas protection in other developing countries differs less between periods. In all cases, we find that a return to pre-TRIPS protection in developing countries increases the welfare of developing countries and reduces welfare in developed economies. Reverting to pre-TRIPS protection also increases world welfare in most of the robustness checks. These findings demonstrate the robustness of our baseline findings for the TRIPS counterfactual.

E Sunk Export Cost Model

Suppose that, in order to export to country n , an innovator from country i must pay a sunk export cost of a_{ni}^s/ϵ units of country i labor per innovation. $a_{ni}^s \geq 0$ governs bilateral sunk export costs, while ϵ is a producer-specific inverse cost shock drawn from a Pareto distribution with scale parameter one and shape parameter d . Let $a_{ii}^s = 0$ for all i , implying that there is no sunk cost of domestic market entry. We assume that innovators learn ϵ after deciding whether to pay the patent preparation cost. Knowing ϵ and ψ producers then make simultaneous entry and patenting decision for each destination before learning their productivity. Apart from the addition of sunk export costs for innovators, the model is unchanged from Section 2.

This extension can be solved using the same steps as in the baseline model, with the complication that entry and patenting decisions in each market are interdependent and the set of available Helpman-Krugman varieties differs across destinations. There exist threshold values of the inverse cost shock ϵ above which innovators choose to export. These entry thresholds vary by origin-destination-sector, depend upon whether an innovation is patented and are decreasing in innovation quality ψ , implying that higher quality innovations are sold in more countries on average. To conserve space, we omit the details of the solution.

To calibrate the extended model, we start by fixing d , which determines the dispersion of sunk cost draws. Using data on export participation and lifecycle dynamics for US firms, Alessandria and Choi (2014) calibrate a dynamic model of firm-level exporting with heterogeneity in sunk export costs. We set d such that the coefficient of variation of sunk export costs in our model equals the value implied by Alessandria and Choi's calibration, which yields $d = 0.19$. Next, we assume $a_{ni} = a\tilde{\tau}_{ni}$ for $n \neq i$, implying that sunk export costs are proportional to iceberg trade costs. Finally, we calibrate the model as in Section 3.2, but including one additional moment. We target that 70 percent of US innovators enter at least one export market based on Foster et al. (2016) who report that 70 percent of US firms that perform R&D in 2005 participate in international trade in future years. This moment allows us to calibrate a and we obtain $a = 0.48$. Otherwise, the calibrated parameters and model fit are similar to the baseline calibration. The calibrated patent protection levels are shown in column (a) of Table 7.

Using the calibrated sunk cost model, we perform the same set of counterfactual experiments that we analyzed for the baseline model in Section 4. However, to reduce the computational burden, we assume that countries maximize steady state welfare when computing the Nash and cooperative equilibria and we only

Table 7: Sunk export cost model counterfactuals

	(a) Patent protection, δ_i		(b) Nash		(c) Cooperative: equal		(d) Reverse TRIPS	
	Baseline	Sunk costs	Baseline	Sunk costs	Baseline	Sunk costs	Baseline	Sunk costs
	Steady state welfare change, EV_i^{SS} (percent)							
US	0.088	0.091	-0.4	-0.5	-2.9	-1.5	-0.10	-0.12
Europe	0.095	0.101	-0.8	-1.0	-2.0	-1.3	-0.12	-0.14
Japan	0.107	0.114	-1.2	-1.4	-2.0	-0.3	-0.13	-0.16
China	0.148	0.157	-1.0	-1.1	13.0	13.7	0.58	0.62
Brazil	0.170	0.185	-1.6	-1.7	12.2	12.8	0.04	0.05
India	0.241	0.258	-1.8	-2.0	11.9	12.6	0.23	0.20
Canada	0.073	0.073	-0.4	-0.5	2.4	3.8	-0.11	-0.13
Korea	0.066	0.070	-0.2	-0.4	2.8	4.0	-0.18	-0.21
Russia	0.228	0.249	-1.9	-2.1	12.3	12.8	0.17	0.14
Mexico	0.205	0.221	-1.8	-1.9	-5.6	-4.8	-0.07	-0.04
South Africa	0.170	0.187	-1.8	-1.9	12.2	12.7	-0.41	-0.34
Rest of world	1.510	1.592	-2.3	-2.4	11.9	12.6	-0.11	-0.13
World Equal			-1.9	-2.0	11.2	12.1	0.09	0.07
World Negishi			-1.2	-1.4	5.3	6.0	0.06	0.06

Notes: This table reports calibrated levels of patent protection in the baseline and sunk export cost models (column a) and welfare changes in these two models for different patent policy counterfactuals (columns b to d). Column (b) reports results for the Nash equilibrium. Column (c) reports results for the cooperative equilibrium when the social planner uses equal weights for all individuals. In column (d), China, Brazil, India, Russia, Mexico and South Africa revert to pre-TRIPS patent protection levels. Pre-TRIPS patent protection is calibrated separately for each model. The Nash and cooperative equilibria in columns (b) and (c) are computed assuming countries maximize steady state welfare. Welfare changes are expressed as the equivalent variation in steady state consumption relative to the calibrated steady state in 2015.

report changes in steady state welfare. Subject to this restriction, we find that the counterfactual results with sunk export costs are very similar to those obtained using the baseline model.

For both the Nash and equal weights cooperative equilibria, we find that each country's optimal patent policy with sunk export costs is the same as in the baseline model.³⁶ In addition, counterfactual steady state welfare changes differ little between the sunk export cost model and the baseline model as shown in Table 7 with the Nash equilibrium in column (b) and the equal weights cooperative equilibrium in column (c). Welfare losses from moving to Nash are slightly larger and gains from cooperation slightly bigger, highlighting the small extra benefit generated by patents with sunk export costs. Finally, using the extended model to calibrate pre-TRIPS patent protection and simulating a return to pre-TRIPS protection levels in developing countries also leads to similar welfare changes to those obtained in the pre-TRIPS counterfactual for the baseline model (column d).

³⁶Recall from footnote 23 that, in the baseline model, optimal patent policies in the Nash and cooperative equilibria are independent of whether or not countries maximize steady state welfare or welfare including transition dynamics. Thus, optimal patent policies in the sunk export cost model are the same as those reported in Table 3.

References

- Alessandria, George, and Horag Choi. 2014. "Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization." *Journal of International Economics* 94(2): 207-223.
- Bena, Jan, and Elena Simintzi. 2022. "Machines Could Not Compete with Chinese Labor: Evidence from US Firms' Innovation." Available at SSRN: <https://ssrn.com/abstract=2613248>.
- Berger, Roland. 2004. "A Study on the Cost of Patenting."
- Branch, Mary Ann, Thomas F. Coleman, and Yuying Li. 1999. "A Subspace, Interior and Conjugate Gradient Method for Large-Scale Bound-Constrained Minimization Problems." *SIAM Journal on Scientific Computing* 21(1): 1-23.
- Brent, Richard. 2013. "Algorithms for Minimization Without Derivatives." Courier Corporation.
- Bureau of Economic Analysis (BEA). 2022. "Gross Output by Industry." (accessed April 22, 2022 and February 18, 2023).
- Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor. 2023. "Tariff Reductions, Heterogeneous Firms, and Welfare: Theory and Evidence for 1990–2010." *IMF Economic Review*: 1-35.
- Chen, Kewang, and Cornelis Vuik. 2022. "Non-stationary Anderson Acceleration with Optimized Damping." arXiv preprint arXiv:2202.05295.
- Cohen, Wesley, Richard Nelson, and John Walsh. 2000. "Protecting their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (Or Not)." NBER Working Paper 7552.
- Coleman, Thomas F., and Yuying Li. 1996. "An Interior Trust Region Approach for Nonlinear Minimization Subject to Bounds." *SIAM Journal on Optimization* 6(2): 418-445.
- Endres, Stefan C., Carl Sandrock, and Walter W. Focke. 2018. "A Simplicial Homology Algorithm for Lipschitz Optimisation." *Journal of Global Optimization* 72: 181-217.
- Foster, Lucia, Cheryl Grim, and Nikolas Jason Zolas. 2016. "A Portrait of Firms that Invest in R&D." US Census Bureau Center for Economic Studies Paper No. CES-WP-16-41.
- Lybbert, Travis J., and Nikolas J. Zolas. 2012. "Getting Patents and Economic Data to Speak to Each Other: An "Algorithm Links with Probabilities" Approach for Joint Analyses of Patenting and Economic Activity." WIPO Working Paper No. 5.
- Moré, Jorge J. 2006. "The Levenberg-Marquardt Algorithm: Implementation and Theory." In *Numerical Analysis: Proceedings of the Biennial Conference Held at Dundee, June 28–July 1, 1977*, pp. 105-116.
- Nocedal, Jorge, and Stephen J. Wright, eds. 1999. "Numerical Optimization." Springer New York.
- Organization for Economic Co-operation and Development (OECD). 2021. "Input-Output Tables."
- Park, Walter. 2010. "On Patenting Costs." *The WIPO Journal* 2(1): 38-48.
- PATSTAT. 2022. "Worldwide Patent Statistical Database." European Patent Office, Autumn 2022.
- Schott, Peter. 2008. "The Relative Sophistication of Chinese Exports." *Economic Policy* 23(53): 6-49.
- World Intellectual Property Organization (WIPO). 2023. "WIPO Statistics Database."
- World Bank. 2023. "World Development Indicators." (accessed March 10, 2023).
- Zhang, Junzi, Brendan O'Donoghue, and Stephen Boyd. 2020. "Globally Convergent Type-I Anderson Acceleration for Nonsmooth Fixed-point Iterations." *SIAM Journal on Optimization* 30(4): 3170-3197.