

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

Working Paper

Trade, Innovation and Optimal Patent Protection

CESifo Working Paper, No. 10777

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Hémous, David; Lepot, Simon; Sampson, Thomas; Schärer, Julian (2023) : Trade, Innovation and Optimal Patent Protection, CESifo Working Paper, No. 10777, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/282465>

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.

Trade, Innovation and Optimal Patent Protection

David Hémous, Simon Lepot, Thomas Sampson, Julian Schärer

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Trade, Innovation and Optimal Patent Protection

Abstract

This paper provides a first comprehensive quantitative analysis of optimal patent policy in the global economy. We introduce a new framework, which combines trade and growth theory into a tractable tool for quantitative research. Our application delivers three main results. First, the potential gains from international cooperation over patent policies are large. Second, only a small share of these gains has been realized so far. And third, the WTO's TRIPS agreement has been counterproductive, slightly reducing welfare in the Global South and for the world. Overall, there is substantial scope for policy reform.

JEL-Codes: F130, D430, O340.

Keywords: trade policy, innovation, growth, patents, TRIPS.

David Hémous
University of Zurich / Switzerland
david.hemous@econ.uzh.ch

Simon Lepot
University of Zurich / Switzerland
simon.lepot@econ.uzh.ch

Thomas Sampson
London School of Economics and Political
Science, London / United Kingdom
t.a.sampson@lse.ac.uk

Julian Schärer
University of Geneva / Switzerland
julian.schaerer@unige.ch

November 2023

We thank Pol Antràs, Giovanni Maggi, Phillip McCalman, Marc Melitz, Monika Mrázová and Fabrizio Zilibotti for helpful comments and discussions. We also thank seminar and conference participants at ETH Zürich, Geneva, Glasgow, Harvard-MIT, KU Leuven, Manchester, Milano Bicocca, Oslo, Paris School of Economics, St Gallen, Tilburg, Vienna, Zürich, the Barcelona Macro Workshop, the Princeton IES Summer Workshop, the Villars Workshop on International Trade and Multilateralism, and the IX Workshop on Structural Transformation and Macroeconomic Dynamics in Cagliari.

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. Rich countries argue that strong intellectual property rights are needed to stimulate innovation and are willing to grant innovators substantial monopoly rights. Poor countries counter that monopoly rights inflate consumer prices and argue that strong intellectual property rights amount to a transfer from poor-country households to rich-country firms. These frictions have been amplified by the rapidly rising importance of the Global South in the world economy, which has brought questions about intellectual property rights in these countries to the forefront of the policy debate.

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 intellectual property rights in developing economies by requiring countries to adopt intellectual property policies similar to those already implemented in rich countries. As Saggi (2016) reports, 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, thereby contributing to the failure of the Doha Round and, ultimately, the stalemate at the WTO.

This paper provides a first comprehensive quantitative analysis of optimal patent policy in the global economy. We make two main contributions. First, we develop a new framework of trade, growth, and patenting, which combines elements from trade and growth theory into a tractable tool for quantitative work. Second, we use this framework to conduct a series of counterfactual experiments that shed new light on international patent policy. Our analysis yields three main results. First, the potential gains from international cooperation in setting patent policies are large. Second, only a small share of these gains has been realized so far. And third, the WTO's TRIPS agreement has been counterproductive, slightly reducing welfare in the Global South and for the world as a whole. The main reason is that the best-case scenario of cooperative patent policy requires much weaker patent protection in the Global South than in the Global North according to our model, while the TRIPS agreement seeks to harmonize patent protection across the board. Overall, we conclude that there is substantial scope for policy reform.

The classic trade-off patent policy faces is that stronger protection brings dynamic benefits from faster innovation, but also generates static costs through higher prices (Nordhaus 1969). We formalize this trade-off in a model of trade, growth and patents that builds upon Grossman and Lai (2004) and allows innovation, patenting and market power to respond endogenously to changes

in patent policy. The model implies that in open economies stronger patent protection has global benefits, but local costs. While all countries reap the dynamic benefits of 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 a firm 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.

A key property of the model is that the global growth rate is more sensitive to the level of patent protection in markets that are more profitable for innovators. Such markets account for a higher fraction of the value of innovation and, therefore, have a disproportionate effect on research and development (R&D) investment. It follows that not only larger countries, but also (because of home bias in trade) more innovative countries generate greater dynamic benefits by increasing patent rights. Consequently, optimal patent protection differs across economies and tends to be weaker in less innovative and more isolated countries.

The paper's theoretical contribution is to develop a patenting model that is quantitatively tractable in an open economy setting. We achieve this by integrating Helpman-Krugman trade theory and Eaton-Kortum trade theory within an expanding variety growth model. Innovators create varieties and choose whether to patent their inventions in each country. 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 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.

We calibrate the model to the world economy in 2015 divided into the US, Europe, Japan, China, Brazil, India, Canada, Korea, Russia, Mexico and a residual rest of the world. The calibration exactly matches bilateral trade flows and uses data on R&D expenditure and cross-border patent applications to estimate, for each country, the efficiency of R&D and the strength of patent protection. In the calibrated economy, innovation is highly concentrated in the richer economies. The US alone accounts for over one-third of global innovation and the US, Europe, and Japan together account for 79 percent of innovation. Since trade is home biased, this pattern of specialization in innovation implies that growth is much more dependent upon patent policy in Europe, Japan and especially the US, than upon patent policy in other countries.

The key policy parameter that captures the strength of patent protection is the Poisson rate at which a patent expires. This parameter can be interpreted as capturing a combination of the statutory patent length and the probability of patent enforcement. We infer how this parameter differs across countries using variation in levels of international patent applications relative to

bilateral trade flows. Intuitively, stronger protection increases the value of patenting, leading to more patent applications conditional on trade flows. Our estimates imply that a patent in the US has an expected duration of 12 years. Since the legal term of US patents is 20 years, this value is consistent with relatively high levels of patent enforcement in the US. We find that Europe, Japan, Canada and Korea have similar patent protection to the US, while protection is weaker elsewhere. Expected patent duration is around 7 years in China, 5 years in Brazil and 4 years in India.

We begin our counterfactual analysis by studying countries' incentives to unilaterally deviate from the 2015 status quo. We simulate how changes to patent protection in one country affect the global economy and compute welfare effects accounting for transition dynamics between steady states. We find that all countries have an incentive to weaken their patent protection, because 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 larger or more innovative countries has a much stronger effect on growth than patent protection in developing economies such as Brazil or India.

Next, we turn to the noncooperative scenario and 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 Nash 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 2.4 percent lower in the Nash equilibrium compared to the calibrated steady state (measured as the equivalent variation in consumption and weighing all individuals equally). Losses are lower for developed countries which have larger static gains because their initial levels of protection are higher. For example, welfare declines by only 0.5 percent in the US and 0.3 percent in Korea.

Then, we consider the cooperative scenario that maximizes world welfare. In the baseline case with equal welfare weights for all individuals, efficiency requires that 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 encouraging innovation only to countries that are either highly innovative themselves, or that are closely integrated with innovative economies (such as Mexico with the US). By contrast, countries such as China that are large, but not very innovative, do not provide protection in the cooperative equilibrium because, although stronger protection in these economies raises growth, the static costs of this protection impact more people.

We find that the gains from patent policy cooperation are considerable. World welfare is 8.8 percent higher in the cooperative equilibrium, which is over 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.¹

Finally, we study the effects of TRIPS. We begin by re-calibrating the model using 1992 data to estimate pre-TRIPS patent protection levels. This exercise shows that protection in China, India and Russia was considerably lower before the implementation of TRIPS. Starting from the calibrated equilibrium in 2015, we then consider two counterfactuals. A pre-TRIPS counterfactual in which we return patent protection in developing countries to pre-TRIPS levels. And a harmonization counterfactual in which we set patent protection in all countries with weaker protection than the US equal to the US level. Both counterfactuals lead to the same conclusion. Increasing patent protection in developing countries, as TRIPS sought to do, reduces welfare in those countries and for the world as a whole, while slightly benefiting developed economies that do not strengthen their protection.

Taken together, our results imply that there is significant scope for TRIPS reform. The main reason is that TRIPS pushes policy towards common patent rights across countries, which is suboptimal according to our analysis. Consistent with concerns about the equity of the TRIPS agreement, we also find that strengthening patent rights in developing economies raises global inequality by redistributing consumption from poorer to richer countries. However, these results come with an important caveat. WTO membership is a single undertaking that includes market access commitments in addition to TRIPS. When we combine our pre-TRIPS counterfactual with changes in trade costs, we find that relatively small trade cost increases are sufficient to offset any benefits developing countries reap from reverting to pre-TRIPS patent protection. This suggests that, on net, developing countries still gain from WTO membership due to the benefits of improved market access.

Our paper contributes to a number of literatures. The theory builds upon extensive work on trade and growth.² It is most closely related to the theories of innovation and patenting in Eaton and Kortum (1999) and Grossman and Lai (2004). However, we go beyond the existing literature by introducing a quantitative model of trade and patenting that incorporates trade in both monopolistic and competitive products in an endogenous growth framework. In our model, the trade elasticity depends upon the diffusion of existing technologies, a feature shared by Lind and Ramondo (2022)

¹The incentive to free-ride on other countries' patent protection is reminiscent of Acemoglu et al. (2017), who show that countries can have an incentive to adopt greater levels of social protection while free-riding on countries that are more egalitarian, but also more innovative.

²Helpman (1993) provides the first general equilibrium analysis of the welfare effects of intellectual property rights in an endogenous growth model. Saggi (2016) reviews the literature on trade and intellectual property rights. Akcigit and Melitz (2021) and Melitz and Redding (2021) review the broader trade and growth literature.

who study how innovation and diffusion shape trade patterns. The model is also related to Hsieh et al. (2022) who quantify the sources of growth in a global economy where new varieties start out being produced in a single country, before diffusion leads to international competition. But neither of these papers endogenizes innovation or considers the effects of patent policy on growth and welfare. Borota Milicevic et al. (2023) analyze optimal innovation policy in a two country world, but focus on subsidies rather than intellectual property rights. Akcigit et al. (2021) study the interplay between trade policy and innovation policy and use patent data to calibrate their model. But, like other quantitative dynamic trade models that use patent data for calibration (e.g. Cai et al. 2022, Sampson 2023), they do not model patenting decisions or institutions, nor do they analyze strategic interactions between countries.

There is a small quantitative literature on TRIPS and intellectual property policy in open economies. Eaton and Kortum (1996, 1999) use a model of innovation and patenting to shed light on which countries drive growth and the extent of international technology diffusion. However, in their model mark-ups are unaffected by patent protection. Building on Eaton and Kortum's theory, McCalman (2001) quantifies how patent policy harmonization under TRIPS reallocates producer surplus across countries, while McCalman (2005) shows that TRIPS benefits all countries through higher innovation. However, McCalman's analysis does not allow TRIPS to affect market power. Chaudhuri et al. (2006) use estimated price and expenditure elasticities to quantify the welfare effects of product patent enforcement under TRIPS in the fluoroquinolones subsegment of the Indian anti-bacterials market, but do not study changes in innovation.

More recently, Lai and Yan (2013) calibrate a version of the Grossman and Lai (2004) model and argue that harmonization of patent protection at US levels is globally welfare increasing. However, their main contribution is to extend Grossman and Lai's theory and their quantitative application is mainly illustrative in nature – for instance, they do not match trade or patent flows. In related work, Jakobsson and Segerstrom (2016) find that TRIPS raises welfare in a two-region product cycle model of foreign direct investment. And Santacreu (2023) analyzes the joint determination of tariffs and intellectual property rights in bilateral trade agreements when intellectual property protection affects revenue from technology licensing. But none of these studies models patenting choices. Relative to existing quantitative research, our main contribution is to study optimal patent policy when patenting is endogenous and there exists a trade-off between the dynamic benefits and static costs of stronger protection.

Consistent with our model, recent evidence establishes that innovation responds 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 for the affected crops. TRIPS itself 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

to assess whether stronger patent protection induces more innovation in pharmaceuticals. They find a positive effect for protection in developed economies, but no 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 market size and levels of innovation. Looking at 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. 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 rights. 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 intellectual property rights.

The remainder of the paper is organized as follows. Section 2 introduces the model and characterizes equilibrium behavior. Section 3 explains how we calibrate the model and discusses model fit and the calibrated parameters. Section 4 presents the counterfactual analysis. Finally, Section 5 concludes.

2 A Theory of Trade and Patents

We develop a dynamic model of trade and patenting with endogenous innovation. In the model, patent protection determines the expected duration of innovators’ monopoly control over their ideas. Consequently, as in Grossman and Lai (2004), stronger patent rights incentivize innovation, but generate a static distortion due to monopoly pricing. In our model, greater patent protection also generates a sourcing distortion because monopoly control over production restricts buyers’ ability to source from the lowest cost supplier. The theory embeds the trade-off between static costs and dynamic benefits into a quantitative open economy model that is suitable for counterfactual analysis of changes in patent protection.

2.1 Economic Environment

We consider an economy with N countries and $S + 1$ sectors. Sectors $s \neq 0$ feature endogenous innovation and patenting, while 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 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 given by:

$$c_{nt}^s(\omega) = \psi(\omega) \left(\frac{p_{nt}^s(\omega)}{P_{nt}^s} \right)^{-\sigma^s} Y_{nt}^s,$$

where $p_{nt}^s(\omega)$ is the price of variety ω in country n and P_{nt}^s is the sectoral price index.

Output from each sector is combined using a Cobb-Douglas aggregator to produce a non-tradable final good according to:

$$Y_{nt} = \prod_{s=0}^S \left(\frac{Y_{nt}^s}{\beta^s} \right)^{\beta^s}. \quad (2)$$

The final good is used for consumption and as an intermediate input in variety production.

Final consumption demand comes from each country's representative agent whose intertemporal preferences are given by:

$$U_{nt} = \int_t^\infty e^{-\rho(\tilde{t}-t)} \frac{C_{nt}^{1-1/\gamma}}{1-1/\gamma} d\tilde{t}. \quad (3)$$

In this equation ρ is the discount rate, γ is the elasticity of intertemporal substitution and C_{nt} denotes aggregate final good consumption in country n . Agents earn income from wages w_{nt} and by investing in a risk free asset with interest rate r_{nt} . Consequently, the representative agent's intertemporal budget constraint is:

$$\dot{W}_{nt} = w_{nt}L_n + r_{nt}W_{nt} - P_{nt}C_{nt},$$

where W_{nt} denotes total assets owned by the representative agent and P_{nt} is the price index of the final good in country n . We assume there is no international borrowing or lending, implying that asset markets clear at the national level.

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:

$$y_i^s(\omega) = \frac{z_i^s(\omega)}{(\alpha^s)^{\alpha^s} (1 - \alpha^s)^{1 - \alpha^s}} l_i^s(\omega)^{\alpha^s} q_i^s(\omega)^{1 - \alpha^s}, \quad (4)$$

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 levels are 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 innovation and diffusion outcomes. The shape parameter $\theta^s > \sigma^s - 1$ is an inverse measure of productivity dispersion across varieties. Productivity draws are independent across varieties and also across countries within a variety.

Product varieties are tradable subject to iceberg trade costs. In order to sell one unit of output in country n , a producer in country i must ship τ_{ni}^s units. We assume $\tau_{ii}^s = 1$ for all countries i .

Innovation and Patenting. In all sectors $s \neq 0$, new products are created by risk neutral innovators. Innovation uses labor and 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, which is country-sector specific, 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 (Jones 1995). Imposing $\kappa > 0$ also ensures that all countries innovate in all sectors in equilibrium.

Let Ψ_t^s be the aggregate quality of all varieties produced in sector s , defined by:

$$\Psi_t^s = \int_0^{M_t^s} \psi(\omega) d\omega.$$

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.³

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:

³Equivalently, one could assume that each invention creates one new product variety and that innovation and patenting costs are inversely proportional to Ψ_t^s .

$$w_{it} = \eta_i^s (L_{Rit}^s)^{-\kappa} \Psi_t^s V_{it}^s. \quad (5)$$

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. Allowing variety-level quality and country-level productivity to vary independently ensures the tractability of the model because it implies that patenting and international sourcing decisions are separable. This separability reduces the number of state variables and, as shown below, guarantees that trade in goods that are not produced monopolistically follows the Eaton-Kortum model.

Initially, only the innovator knows how to produce its new varieties giving them a technological monopoly over their invention. 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 produce products with the same quality and all firms in a given country have the same country-specific productivity $z_n^s(\omega)$. Consequently, diffusion strips the innovator of its technological monopoly. Innovators may also lose their monopoly due to product obsolescence. We assume each variety become obsolete at Poisson rate ζ^s .

Anticipating the possibility of technology diffusion, the innovator may also secure a legal monopoly over their invention by obtaining 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 .⁴ Once purchased, a patent expires 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 captures both the length of protection available and the effective enforcement of patent rights. An increase in patent protection

⁴We 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, an innovator who holds a production monopoly in one country will still face competition from producers in other countries. Therefore, the value of patents in open economies comes primarily from obtaining a sales monopoly. Consider an innovator who holds a production monopoly in their domestic market, but not in any foreign markets. At the steady state equilibrium calibrated in Section 3.2, when the innovator loses their sales monopoly in a foreign market their expected profits drop by 97.4 percent on average (where the expectation is taken over productivity z and the drop is averaged across all country pairs). Consequently, assuming that a domestic patent protects production, as well as sales, would make little difference to the value of holding a domestic patent. At our calibrated steady state, the expected value of domestic patent protection for firms that choose to patent would increase by 2.0 percent in the US, 3.1 percent in Europe, 5.8 percent in China and 5.7 percent on average across countries.

reduces δ_n^s .

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).⁵ 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.

The benefit of patenting for an innovator is that it extends the expected duration of their monopoly over an invention. An innovator who purchases a country n patent loses their monopoly in country n only when both the technology has diffused and the patent has expired. When choosing whether to patent, the innovator compares this benefit to the costs of patenting. In order 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 will 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.

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 when making the application (for example, agent payments and internal costs of managing the application process). We will 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 $w_{nt}L_{int}^{s,e}$ represent 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 .

⁵Using data on patenting by French firms, De Rassenfosse et al. (2022) note that the application year and the priority year of the invention coincide in around 85 percent of cases.

2.2 Equilibrium

The model has 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 non-expired 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.

Before solving the model, it is useful to decompose aggregate quality Ψ_t^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 at time t and let Ψ_{Cnt}^s be the aggregate quality of all Eaton-Kortum products sold competitively in country n at time t . Then we have:

$$\Psi_t^s = \Psi_{Cnt}^s + \sum_{i=1}^N \Psi_{Mnit}^s. \quad (6)$$

Note that since all products are sold to all countries, this equation holds for any destination n .

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 to purchase patents $L_{int}^{s,e}$. The dynamic equilibrium solves for optimal innovation and patenting decisions.

In this section we sketch the main features of the static equilibrium. A formal definition together with the full set of static equilibrium conditions can be found 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 (4) to solve for marginal cost and recalling that exports are subject to iceberg trade costs τ_{ni}^s , it follows that the price $p_{ni}^s(\omega)$ of a Helpman-Krugman

variety ω produced in i and sold in n satisfies:

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

Therefore, the monopolists profits per variety $\pi_{ni}^s(\omega)$ are given by:

$$\pi_{ni}^s(\omega) = \psi(\omega) \frac{(\sigma^s - 1)^{\sigma^s - 1}}{(\sigma^s)^{\sigma^s}} \left(\frac{\tau_{ni}^s w_i^{\alpha^s} P_i^{1-\alpha^s}}{z_i^s(\omega)} \right)^{1-\sigma^s} (P_n^s)^{\sigma^s} Y_n^s. \quad (7)$$

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 across source countries and varieties to derive a subprice index P_{Mn}^s for Helpman-Krugman products sold in country n :

$$P_{Mn}^s = \left(\frac{\sigma^s}{\sigma^s - 1} \right) \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \right]^{\frac{1}{1-\sigma^s}} \left(\sum_{j=1}^N \Psi_{Mnj}^s (\Phi_{nj}^s)^{\frac{\sigma^s - 1}{\theta^s}} \right)^{\frac{1}{1-\sigma^s}}, \quad (8)$$

where $\Gamma(\cdot)$ is the Gamma function and:

$$\Phi_{ni}^s \equiv T_i (\tau_{ni}^s w_i^{\alpha^s} P_i^{1-\alpha^s})^{-\theta^s}, \quad (9)$$

gives the supply potential of country i in country n , which is an inverse measure of the average cost of producing for country n in country i .

Aggregation also yields that the value of exports X_{Mni}^s of Helpman-Krugman products from i to n is given by:

$$X_{Mni}^s = \frac{\Psi_{Mni}^s (\Phi_{ni}^s)^{\frac{\sigma^s - 1}{\theta^s}}}{\sum_{j=1}^N \Psi_{Mnj}^s (\Phi_{nj}^s)^{\frac{\sigma^s - 1}{\theta^s}}} \left(\frac{P_{Mn}^s}{P_n^s} \right)^{1-\sigma^s} P_n^s Y_n^s. \quad (10)$$

We see that 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 (9) into equation (10) also 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, the subprice index P_{Cn}^s for Eaton-Kortum products sold in country n is:

$$P_{Cn}^s = \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \Psi_{Cn}^s \right]^{\frac{1}{1-\sigma^s}} \left(\sum_{j=1}^N \Phi_{nj}^s \right)^{-\frac{1}{\theta^s}}, \quad (11)$$

and exports X_{Cni}^s of Eaton-Kortum products from i to n satisfy:

$$X_{Cni}^s = \frac{\Phi_{ni}^s}{\sum_{j=1}^N \Phi_{nj}^s} \left(\frac{P_{Cn}^s}{P_n^s} \right)^{1-\sigma^s} P_n^s Y_n^s. \quad (12)$$

It follows that the elasticity of Eaton-Kortum trade to τ_{ni}^s is given by the Fréchet dispersion parameter θ^s .

Using equations (8)–(12), the remaining static equilibrium conditions can be obtained by aggregating Helpman-Krugman with Eaton-Kortum products in each sector to obtain price indices and trade flows 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 product type changes over time.

Value of firms and patenting decisions. Consider a firm 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 firm makes from sales 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). Recalling that firms make patenting decisions before learning their productivity, we have:

$$V_{nit_0}^{s,NP}(\psi) = \int_{t_0}^{\infty} \mathbb{E}_z \pi_{ni}^s(\psi, z) \exp \left(- \int_{t_0}^t (r_i + \zeta^s + \nu^s) d\tilde{t} \right) dt, \quad (13)$$

where $\mathbb{E}_z \pi_{ni}^s(\psi, z)$ denotes expected profits computed over the distribution of productivity z .

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 in destination n conditional on patenting $V_{nit_0}^{s,P}(\psi)$ satisfies:

$$V_{nit_0}^{s,P}(\psi) = \int_{t_0}^{\infty} \mathbb{E}_z \pi_{ni}^s(\psi, z) \exp \left(- \int_{t_0}^t (r_i + \zeta^s + \nu^s) d\tilde{t} \right) \times \{1 - \exp[-\delta_n^s(t-t_0)] + \exp[(\nu^s - \delta_n^s)(t-t_0)]\} dt. \quad (14)$$

We can now solve for expected profits by noting from equation (7) that profits per variety are

proportional to the monopolist's quality ψ and equal to a fraction $1/\sigma^s$ of revenue. Therefore, aggregate profits made by monopolists in i from sales to n are given by X_{Mni}^s/σ^s and expected profits are:

$$\mathbb{E}_z \pi_{ni}^s(\psi, z) = \psi \mathbb{E}_z \pi_{ni}^s(1, z) = \psi \frac{X_{Mni}^s}{\sigma^s \Psi_{Mni}^s}. \quad (15)$$

Substituting this expression into equations (13) and (14) 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. Since each invention comprises Ψ^s varieties, it follows that the firm patents in n if at the time of application t_0 :

$$\Psi^s \left[V_{nit_0}^{s,P}(\psi) - V_{nit_0}^{s,NP}(\psi) \right] \geq w_n f_n^{s,e}.$$

Because the value functions are proportional to ψ , this inequality defines a quality threshold $\psi_{ni}^{s,e*}$ such that only firms with quality ψ above the threshold opt to patent in n (conditional on having paid the application preparation cost). Rearranging the expression above and remembering that ψ is drawn from a distribution with lower bound one, we have:

$$\psi_{ni}^{s,e*} = \max \left(\frac{w_n f_n^{s,e}}{\Psi^s \left[V_{nit_0}^{s,P}(1) - V_{nit_0}^{s,NP}(1) \right]}, 1 \right). \quad (16)$$

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 the patenting threshold $\psi_{ni}^{s,*}$ is defined by:

$$\psi_{ni}^{s,*} = \max(\psi_{ni}^{s,e*}, \psi_i^{s,o*}). \quad (17)$$

Patenting increases the expected duration of an innovator's monopoly over their varieties. A longer monopoly is more valuable to higher quality firms since expected profits are proportional to quality by equation (15). Consequently, the benefits exceed the fixed costs of patenting only for firms with quality above the patenting thresholds defined in equation (17).

The expected value V_{it}^s of inventing a new variety at t equals the expected present discounted value of profits in all markets less expected patenting costs. Using the optimal patenting thresholds, summing across destinations and taking expectations over the quality distribution, Appendix A.2 shows that V_{it}^s is given by:

$$V_{it}^s = \sum_{n=1}^N \left\{ \frac{k}{k-1} \left[V_{nit}^{s,NP} (1) \left(1 - (\psi_{ni}^{s*})^{-k+1} \right) + V_{nit}^{s,P} (1) (\psi_{ni}^{s*})^{-k+1} \right] - (\psi_{ni}^{s*})^{-k} \frac{w_n f_n^{s,e}}{\Psi^s} \right\} - (\psi_i^{s,o*})^{-k} \frac{w_i f_i^{s,o}}{\Psi^s}. \quad (18)$$

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:

$$\Psi_{Mni}^s = \Psi_{Mni}^{s,NP} + \Psi_{Mni}^{s,P,ND} + \Psi_{Mni}^{s,P,D}. \quad (19)$$

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:

$$\dot{\Psi}_{Mni}^{s,NP} = \eta_i^s (L_{Ri}^s)^{1-\kappa} \Psi^s \frac{k}{k-1} \left[1 - (\psi_{ni}^{s*})^{1-k} \right] + \delta_n^s \Psi_{Mni}^{s,P,ND} - (\nu^s + \zeta^s) \Psi_{Mni}^{s,NP}. \quad (20)$$

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 ψ_{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:

$$\dot{\Psi}_{Mni}^{s,P,ND} = \eta_i^s (L_{Ri}^s)^{1-\kappa} \Psi^s \frac{k}{k-1} (\psi_{ni}^{s*})^{1-k} - (\delta_n^s + \nu^s + \zeta^s) \Psi_{Mni}^{s,P,ND}. \quad (21)$$

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:

$$\dot{\Psi}_{Mni}^{s,P,D} = \nu^s \Psi_{Mni}^{s,P,ND} - (\delta_n^s + \zeta^s) \Psi_{Mni}^{s,P,D}. \quad (22)$$

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:

$$\dot{\Psi}_{Cn}^s = \sum_{i=1}^N \left(\nu^s \Psi_{Mni}^{s,NP} + \delta_n^s \Psi_{Mni}^{s,P,D} \right) - \zeta^s \Psi_{Cn}^s. \quad (23)$$

Combining these laws of motion with the patenting thresholds and firm value functions derived above and imposing labor market clearing gives the dynamic equilibrium, which is formally defined in Appendix A.2.

Let g^s denote the growth rate of aggregate quality Ψ^s in sector s . Using equations (6) and (19) 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 (20)–(23) implies that in any dynamic equilibrium:

$$g^s = \sum_{i=1}^N \frac{k}{k-1} \eta_i^s (L_{Ri}^s)^{1-\kappa} - \zeta^s, \quad \text{for } s \neq 0. \quad (24)$$

This expression shows how sector-level 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. Steady state also requires that the aggregate qualities of Eaton-Kortum products and of each type of Helpman-Krugman products in sector s grow at rate g^s . Let g be the growth rate of final consumption C_i in any country i . Wages w_i , final good output Y_i , trade flows X_{ni}^s and aggregate profits also grow at rate g , which is given by:

$$g = \frac{1}{\sum_{s=0}^S \beta^s \alpha^s} \sum_{s=0}^S \frac{\beta^s}{\sigma^s - 1} g^s, \quad (25)$$

Thus, in steady state, growth results from increases in the aggregate quality of varieties produced in each sector. In turn, growth in aggregate quality results from innovation as shown by equation (24).

Computing the integrals in equation (13) gives that the steady state value of a variety with quality one that is not patented satisfies:

$$V_{nit}^{s,NP}(1) = R^{s,NP} \mathbb{E}_z \pi_{nit}^s(1, z) \text{ where } R^{s,NP} \equiv \frac{1}{r + \zeta^s + \nu^s - g + g^s}. \quad (26)$$

The term $R^{s,NP}$ in this expression captures the expected value that a firm that does not patent obtains from future profit flows. It is the inverse of the firm's effective discount rate and is decreasing in the interest rate r , the product obsolescence rate ζ^s and the technology diffusion rate ν^s , but increasing in the growth rate of profits per variety $g - g^s$.

Similarly, equation (14) implies that the value of a patented variety is:

$$V_{nit}^{s,P}(1) = R_n^{s,P} \mathbb{E}_z \pi_{nit}^s(1, z), \quad (27)$$

where $R_n^{s,P} = R^{s,NP} + \Delta R_n^s$ and:

$$\Delta R_n^s \equiv \frac{1}{r + \zeta^s + \delta_n^s - g + g^s} - \frac{1}{r + \zeta^s + \nu^s + \delta_n^s - g + g^s} > 0.$$

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 . We will refer to ΔR_n^s 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 $\nu^s = 0$, meaning that there is no technology diffusion, 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 it is convenient to detrend all variables. Detrending yields 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 Ψ defined by:

$$\Psi \equiv \left[\prod_{s=0}^S (\Psi^s)^{\frac{\beta^s}{\sigma^s-1}} \right]^{\frac{1}{\sum_{s=0}^S \beta^s \alpha^s}}.$$

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:

$$\tilde{\pi}_{ni}^s \equiv \Psi^s \frac{\mathbb{E}_z \pi_{ni}^s(1, z)}{\Psi} = \frac{\tilde{X}_{Mni}^s}{\sigma^s \tilde{\Psi}_{Mni}^s}, \quad (28)$$

where the equality uses equation (15). Using equations (26) and (27) to substitute for $V_{nit}^{s,NP}$ (1) and $V_{nit}^{s,P}$ (1) in equation (16), we now obtain that the patenting threshold $\psi_{ni}^{s,e*}$ satisfies:

$$\psi_{ni}^{s,e*} = \max \left(\frac{\tilde{w}_n f_n^{s,e}}{\Delta R_n^s \tilde{\pi}_{ni}^s}, 1 \right). \quad (29)$$

The (interior) patenting threshold depends upon the cost of patenting in n , the benefit of patenting in n and the profitability of market n . 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, but accounting for the option value of patenting in all destinations.

Using equation (18), we can also write the normalized expected value of inventing a new variety \tilde{V}_i^s as:

$$\tilde{V}_i^s = \sum_{n=1}^N \left[\frac{k}{k-1} \tilde{\pi}_{ni}^s \left(R^{s,NP} + \Delta R_n^s (\psi_{ni}^{s*})^{1-k} \right) - \tilde{w}_n f_n^{s,e} (\psi_{ni}^{s*})^{-k} \right] - \tilde{w}_i f_i^{s,o} (\psi_i^{s,o*})^{-k}. \quad (30)$$

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 (30) give the expected patenting costs a firm pays.

Free entry into innovation (5) implies:

$$(L_{Ri}^s)^\kappa = \eta_i^s \frac{\tilde{V}_i^s}{\tilde{w}_i}. \quad (31)$$

Together with equation (30), this expression determines the allocation of labor to R&D and, therefore, the sectoral growth rate g^s by equation (24).

Finally, equation (3) implies that steady state welfare is given by:

$$U_{nt} = \frac{\Psi_t^{1-1/\gamma}}{1-1/\gamma} \frac{\tilde{C}_n^{1-1/\gamma}}{\rho - g \left(1 - \frac{1}{\gamma}\right)}. \quad (32)$$

Conditional on the initial value Ψ_t , an increase in either the normalized consumption level \tilde{C}_n or the growth rate 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 Understanding the Model

Before calibrating the model, it is useful to develop more intuition about how patent protection affects the steady state equilibrium. Therefore, in this section we characterize the direct effect of changes in the strength of patent protection δ_n^s in country n on steady state outcomes in all countries, without allowing for general equilibrium adjustments.

Suppose country n increases the strength of its patent protection by reducing δ_n^s . The resulting change in steady state welfare can be decomposed into a static effect on normalized consumption and a dynamic effect on growth, as shown in equation (32). We start by characterizing the direct static effect on consumption levels. Using the laws of motion for aggregate qualities we have:

$$\tilde{\Psi}_{Mni}^s = \frac{k}{k-1} \eta_i^s (L_{Ri}^s)^{1-\kappa} \left[\frac{1}{g^s + \nu^s + \zeta^s} + \frac{(\psi_{ni}^{s*})^{1-k}}{g^s + \delta_n^s + \zeta^s} \frac{\nu^s}{g^s + \delta_n^s + \nu^s + \zeta^s} \right], \quad (33)$$

implying that a reduction in δ_n^s increases, all else constant, the share of aggregate quality $\tilde{\Psi}_{Mni}^s$ sold in country n and sector s that is supplied monopolistically by country i . And since the equation above holds for any i , it follows that stronger patent protection 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 .⁶ In turn, changes in market power affect consumption levels through their impact on prices, profits and real wages.

Market power creates two distortions that raise prices in country n : 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 their marginal costs. The sourcing distortion arises because country n sources Eaton-Kortum products from its lowest cost supplier, whereas Helpman-Krugman products can only be sourced

⁶For given profits, wages, growth rates and patenting thresholds, a decline in δ_n^s also increases ΔR_n^s and, consequently, L_{Rn}^s by equations (30) and (31). Higher R&D employment raises $\tilde{\Psi}_{Mnn}^s$ by equation (33) leading to a further increase in market power.

from the monopolist's home country, which is not necessarily the lowest cost supplier.⁷ 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:

$$\frac{\partial P_n^s}{\partial \tilde{\Psi}_{Mni}^s} \propto 1 - \underbrace{\left(\frac{\sigma^s - 1}{\sigma^s} \right)^{\sigma^s - 1}}_{\text{Mark-up distortion}} \underbrace{\frac{\left(\tilde{\Phi}_{ni}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}}{\left(\sum_{j=1}^N \tilde{\Phi}_{nj}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}}}_{\text{Sourcing distortion}}. \quad (34)$$

The right hand side of this expression is positive. It follows that the increase in $\tilde{\Psi}_{Mni}^s$ due to stronger patent protection in country 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 country i in country n given by $\tilde{\Phi}_{ni}^s = T_i (\tau_{ni}^s \tilde{w}_i^{\alpha^s} P_i^{1-\alpha^s})^{-\theta^s}$ 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 satisfy:

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

This equation shows 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 $P_n \tilde{Y}_n$.

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

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

where $\tilde{\Pi}_i \equiv \sum_{s=1}^S \sum_{n=1}^N \tilde{X}_{Mni}^s / \sigma^s$ 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 , real profits $\tilde{\Pi} / P_i$ and trade imbalances. In turn, the real wage can be written as:

$$\frac{\tilde{w}_i}{P_i} = \left\{ \prod_{s=0}^S \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \right]^{\frac{\beta^s}{\sigma^s - 1}} \left(\frac{T_i}{\lambda_{Ci}^s} \right)^{\frac{\beta^s}{\theta^s}} \left(\frac{\tilde{\Psi}_{Ci}^s}{\mu_{Ci}^s} \right)^{\frac{\beta^s}{\sigma^s - 1}} \right\}^{\frac{1}{\sum_{s=0}^S \alpha^s \beta^s}}, \quad (35)$$

⁷ Helpman (1993) refers to the sourcing distortion as the “production composition” effect since it arises from an inefficient allocation of production across countries.

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 domestic share of expenditure on Eaton-Kortum products in sector s and country i , while μ_{Ci}^s denotes the expenditure share of Eaton-Kortum products in sector s and country i . When all products are sold competitively $\mu_{Ci}^s = 1$ and equation (35) 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:

$$\frac{\tilde{\Psi}_{Ci}^s}{\mu_{Ci}^s} = 1 - \sum_{j=1}^N \tilde{\Psi}_{Mij}^s \left[1 - \left(\frac{\sigma^s - 1}{\sigma^s} \right)^{\sigma^s - 1} \frac{\left(\tilde{\Phi}_{ij}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}}{\left(\sum_{k=1}^N \tilde{\Phi}_{ik}^s \right)^{\frac{\sigma^s - 1}{\theta^s}}} \right], \quad (36)$$

implying that real wages are decreasing in the share of aggregate quality supplied monopolistically by each exporter j , $\tilde{\Psi}_{Mij}^s$. And comparing this expression with equation (34) shows that the impact of an increase in $\tilde{\Psi}_{Mij}^s$ on real wages is greater when the pricing distortions are larger. Equations (35) and (36) formalize how the static inefficiencies due to market power reduce real wages in this economy. By increasing market power, stronger patent protection exacerbates these inefficiencies.

Putting everything above 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 that occurs because of an increase in the share of aggregate quality $\tilde{\Psi}_{Mni}^s$ supplied to country n monopolistically by country i . 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 higher profits are offset by a decline in real wages caused by an increase in the share of aggregate quality supplied to country n monopolistically. Thus, the direct cost of the static pricing distortions generated by increased monopoly power is borne by country n itself. Section 4 quantifies the importance of 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 turn to the dynamic effect of stronger patent protection on growth. To understand this effect, we use a version of the model where $f_i^{s,o} = 0$, meaning that there are no patent preparation costs and that $\psi_{ni}^{s*} = \psi_{ni}^{s,e*}$ for all n and i . With this simplification, we can write the normalized value of inventing a variety as $\tilde{V}_i^s = \sum_{n=1}^N \tilde{V}_{ni}^s$ where \tilde{V}_{ni}^s denotes the value that comes from supplying destination n . Assuming that all patenting thresholds are interior and substituting equation

(29) into equation (30) yields:

$$\tilde{V}_{ni}^s = \frac{k}{k-1} \tilde{\pi}_{ni}^s \left(R^{s,NP} + \frac{\Delta R_n^s (\psi_{ni}^{s*})^{1-k}}{k} \right),$$

showing that destination n is more valuable when it generates higher expected profits $\tilde{\pi}_{ni}^s$, when the benefits of patenting ΔR_n^s in n are greater and when the threshold ψ_{ni}^{s*} for patenting in n is lower.

Taking the partial derivative of equations (24), (29), (30) and (31) with respect to δ_n^s while holding \tilde{w}_i , $\tilde{\pi}_{ni}^s$ and $R^{s,NP}$ constant implies that the direct effect of stronger patent protection on the sectoral growth rate g^s is given by:

$$\frac{\partial \ln g^s}{\partial \ln \delta_n^s} = \frac{\partial \ln \Delta R_n^s}{\partial \ln \delta_n^s} \frac{1-\kappa}{k} \sum_{i=1}^N \underbrace{\frac{\frac{k}{k-1} \eta_i^s (L_{Ri}^s)^{1-\kappa}}{g^s}}_{\text{Contribution of } i \text{ to global innovation}} \underbrace{\frac{\tilde{V}_{ni}^s}{\tilde{V}_i^s}}_{\text{Contribution of } n \text{ to value of innovation in } i} \underbrace{\frac{\Delta R_n^s (\psi_{ni}^{s*})^{1-k}}{k R^{s,NP} + \Delta R_n^s (\psi_{ni}^{s*})^{1-k}}}_{\text{Contribution of patent protection to value of supplying } n}. \quad (37)$$

The growth elasticity is negative because ΔR_n^s is decreasing in δ_n^s as discussed above. 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.

The decomposition of the growth elasticity in equation (37) shows that the effect of patent protection in country n on growth is greater when destination n accounts for a larger share of the value of innovating \tilde{V}_i^s in more innovative countries. That is, 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 larger countries, in more profitable markets and in destinations that are closely integrated with countries that undertake more innovation. In particular, home bias in trade implies that each country contributes more to the value of domestic innovation than to the value of innovating in foreign countries, i.e. $\tilde{V}_{nn}^s/\tilde{V}_n^s$ exceeds $\tilde{V}_{ni}^s/\tilde{V}_i^s$ for i not equal to n . Therefore, all else equal, growth is more sensitive to δ_n^s when country n contributes a greater share of global innovation.

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

(37) will play an important role in the quantitative analysis.

Our theory captures the trade-off between static costs and dynamic benefits of patent protection when cross-border spillovers result from trade and international knowledge flows. The model does not allow firms to produce abroad through foreign direct investment (FDI). Ruling out FDI implies that firms' market access decisions are independent across countries, which facilitates the quantitative analysis. Incorporating FDI would increase openness by enabling firms to sell abroad through exports or overseas production. It might also mitigate the sourcing distortion, although some distortion would remain in any model with frictions to FDI investment. Developing a quantitative model of FDI and patenting would be an interesting avenue for future research.

3 Model Calibration

We calibrate the model's steady state to fit the world economy in 2015. 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 $N = 11$ economies: US, Europe, Japan, China, Brazil, India, Canada, Korea, Russia, Mexico and the rest of the world. Europe 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). The economies not included in the rest of the world aggregate are the ten largest economies (by real GDP in 2015) for which the patent data we use to calibrate the model is available.⁸ They account for 77 percent of global real GDP and 88 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 obtain 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 patent flows.

Our data shows that international patent flows in 2015 mostly originate in larger, richer economies. US, Europe and Japan together are the origins of 75 percent of cross-border flows, whereas China

⁸Indonesia is omitted due to poor quality patent data.

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). Country-level GDP, 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

To calibrate the model we set some parameters equal to values from the prior literature, choose others to exactly match selected data moments and then jointly calibrate the remaining parameters using simulated method of moments estimation. This section describes the moments that 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 simulate moments in the model.

The parameters ρ , γ , κ , σ^1 and θ^0 are 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, which in our model equates to $(1 - \kappa) / \kappa$, 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 also set the trade elasticity in the no patenting sector θ^0 equal to 5 (Head and Mayer 2014).

We use exact moment matching to infer TB_i , β^s , α^s , τ_{ni}^s and L_i . Trade imbalances TB_i are measured relative to world output. Expenditure shares are set to each sector’s share of world output, which gives $\beta^0 = 0.61$. We 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 $\alpha^1 = 0.39$. Trade costs τ_{ni}^s are chosen such that the equilibrium trade flows exactly match observed trade shares X_{ni}^s / X_{nn}^s in the input-output tables. Finally, the population of each country L_i is set equal to its working age population.

This leaves $4NS + N(S+1) + 3S + 2 = 71$ parameters: δ_n^1 , $f_n^{1,e}$, $f_i^{1,o}$, η_i^1 , T_i^s , ν^1 , ζ^1 , θ^1 , k and g^0 . We calibrate these parameters using 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 .

From the patent data we observe 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 $N(N-1) = 110$ international patent shares are the most important moments in our estimation. Differences in flows by origin provide information about relative values of η_i because they depend upon the relative levels of innovation in each country. At the same time, variation in flows by destination reveal how the strength of patent protection δ_n differs across countries. Countries with stronger patent protection receive more inward flows, all else equal.

Firm-level surveys find that patent applications are made for around 40 percent of US innovations (Cohen, Nelson and Walsh 2000). We match this value to the share of innovations that are patented domestically by US innovators. We also target the share of domestic patents in total inward patents for the US, which is given by the ratio of PAT_{nn} to $\sum_{i=1}^N PAT_{ni}$ with $n = US$.⁹ These moments are informative about the size of the patent preparation cost f_n^o and the patent application cost f_n^e .

We target two moments that discipline the level of patent protection in the US δ_{US} and the technology diffusion rate ν . Kogan et al. (2017) use stock market responses to news about patents to estimate the private value of holding a patent. Averaging their estimates for 1995-2007, we target an aggregate value of patents relative to R&D expenditure of 9.3 percent for the US. We also use trade data from Schott (2008) to compute a measure of turnover in US imports. As in Hsieh et al. (2022), turnover depends upon the rate at which technology diffusion leads to changes in where products are produced. We calculate the rate at which US imports switch origins at the 10-digit product level as described in Appendix B.2.

In the model, patenting costs are denominated in labor units, 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). This may be 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 parameterize patenting costs as:¹⁰

⁹In 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 other countries. Dechezleprêtre et al. (2017) show that international patents are more comparable across patent offices than domestic patents.

¹⁰Given 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

$$f_n^o = f^o h_n, \quad f_n^e = f^e h_n, \quad \text{where } h_n \equiv \left(\frac{\text{Real GDP per capita in US}}{\text{Real GDP per capita in } n} \right)^{1-\chi},$$

where f^o and f^e are common across countries. Imposing this parameterization reduces the number of parameters to calibrate to 51. The cost adjustment h_n shrinks cross-country variation in patenting costs. We compute h_n using observed data on real GDP per capita and setting $\chi = 0.16$ based on the estimated elasticity of real middle management costs to real GDP per capita in Hjort et al. (2022). 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). We also 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 .

Data on R&D expenditures allow us to discipline the allocation of resources to innovation and, consequently, infer 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 do not use R&D data from developing economies. There are two reasons for this choice. The share of R&D investment targeted at knowledge absorption, rather than innovation, is greater in countries further from the technological frontier (Griffith et al. 2004). And, as with patenting costs, R&D workers in developing economies are likely to be relatively more expensive, compared to the average wage, than in high income countries.¹¹

We target two growth rates: aggregate growth g and the difference between price growth in the patenting and no patenting sectors. We measure g as trend growth in US real GDP per capita from 1980-2019 and sector-level price growth using US gross output price indices from 1997-2019. Conditional on innovation and patenting, these moments pin down the product obsolescence rate in the patenting sector ζ and the exogenous growth rate in sector zero g^0 .

The trade elasticity in the patenting sector is a weighted average of $\sigma^1 - 1$ and θ^1 , where the weights depend upon the share of trade in Helpman-Krugman versus Eaton-Kortum products, which varies by country pair. We calibrate θ^1 by targeting an average trade elasticity across all international trade flows equal to five (Head and Mayer 2014). Finally, we target world gross output, each economy's share of world real GDP and each economy's price index relative to the US. These moments are informative about the Fréchet scale parameters T_i^s , which capture any

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 .

¹¹Consistent with these observations, the calibrated model under-predicts R&D expenditure relative to GDP in China, Brazil, India, Russia, Mexico and the rest of the world.

productivity variation across countries that does not stem from differences in innovative capacity.

Let Ω denote the set of parameters that we calibrate 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 that the calibration seeks to 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 .

3.3 Model Fit and Calibrated Parameters

Figure 1 and Table 1 report how the calibrated model matches the targeted moments. Figure 1 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 0.96 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. Moreover, the fit is not simply due to trade and patent flows both following a gravity equation. Figure 2 plots international trade shares (which the model matches exactly) and model-implied international patent shares against observed patent shares. Evidently, the calibrated patent shares match the observed patent data better than the trade shares, implying that the model captures how patent flows deviate from trade flows.

Table 1 shows that the model mostly 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. 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 coun-




Figure 1: International patent shares




Figure 2: Trade and patenting

tries. 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.

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	0.095
Turnover in US imports	0.0175	0.0176
Expenditure on domestic patent applications in US (trillion \$)	0.0047	0.0046
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 \$)	134	134
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	0.032
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	0.16		
Europe	0.20	0.21	0.87	0.88
Japan	0.05	0.05	0.85	0.89
China	0.17	0.22	0.62	0.80
Brazil	0.03	0.03	0.60	0.59
India	0.07	0.07	0.29	0.30
Canada	0.02	0.01	0.98	0.92
Korea	0.02	0.02	0.76	0.89
Russia	0.03	0.03	0.39	0.37
Mexico	0.02	0.02	0.53	0.52
Rest of world	0.23	0.17	0.57	0.43

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

The calibrated parameters are reported in Table 2. We calibrate $\delta_{US} = 0.083$, which implies an expected patent duration of $1/\delta_{US} = 12.1$ years. The legal term of US patents is 20 years. In practice, the expected duration of protection may 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.

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 7.0 years in China, 5.3 years in Brazil and around 4 years in

India, Russia and Mexico.¹² It follows that, even after the implementation of TRIPS, there 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).

Table 2: Calibrated parameter values

Parameter	Value				
Technology diffusion rate, ν	0.055				
Product obsolescence rate, ζ	0.011				
Shape parameter of Pareto quality distribution, k	1.20				
Shape parameter of Fréchet productivity distribution in patenting sector, θ^1	6.6				
Growth rate of no patenting sector, g^0	0.012				
Patent preparation cost, f^o	0.049				
Patent application cost, f^e	0.079				

	Patent protection		R&D efficiency, $\eta_i \times 100$	Patenting sector productivity, $(\tau_i^1)^{\wedge}(1/\theta^1)$	No patenting sector productivity, $(\tau_i^0)^{\wedge}(1/\theta^0)$
	δ_i	$1 / \delta_i$			
US	0.083	12.1	2.5	6.4	12.0
Europe	0.101	9.9	1.4	8.4	7.4
Japan	0.107	9.3	2.7	3.9	14.6
China	0.143	7.0	0.3	3.3	6.5
Brazil	0.187	5.3	0.1	3.1	6.0
India	0.258	3.9	0.1	1.6	4.0
Canada	0.079	12.7	1.0	5.9	10.0
Korea	0.065	15.3	1.5	5.2	10.7
Russia	0.248	4.0	0.1	3.8	7.9
Mexico	0.229	4.4	0.1	3.0	7.6
Rest of world	1.626	0.6	0.2	2.5	3.4

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

The calibration implies technology diffusion is moderately slow. We estimate $\nu = 0.055$, implying it takes 18.2 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 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.

We find that R&D efficiency η_i is highest in the US and Japan, followed by Korea, Europe

¹²Calibrated 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 patenting costs to reflect this need, the model infers that the rest of world has weaker patent protection.

and Canada. Moreover, it is an order of magnitude greater in these economies than in the developing countries. Combining these differences in R&D efficiency with variation in the allocation of labor to R&D implies that innovation is highly concentrated in developed economies. In the calibrated steady state, the US accounts for 34 percent of world innovations, Japan 23 percent, Europe 22 percent and Korea a further 10 percent. By contrast, China accounts for only 3 percent of innovations and Brazil, India, Russia and Mexico each account for fewer than one percent of innovations. Calibrated productivity T_i^s is also higher in developed than in developing economies in both the patenting and no patenting sectors. However, the implied productivity gaps are substantially smaller than the estimated variation in R&D efficiency.

Although we calibrate trade costs to exactly match trade shares in both the patenting and no patenting sectors, the division of trade between monopolistic Helpman-Krugman and competitive Eaton-Kortum products within the patenting sector is endogenous. In the calibrated economy monopolistic products account for 10 percent of world output of the patenting sector. However, since monopolistic products are more traded than competitive products, they account for 36 percent of international trade in the patenting sector. Nevertheless, there is still substantial home bias in sales of monopolistic products. Domestic sales account for 48 percent of world output of monopolistic products.

Unsurprisingly, monopolistic products account for larger shares of output and trade in more innovative countries. Helpman-Krugman products account for 66 percent of US exports in the patenting sector, 83 percent of Japanese exports and 50 percent of European exports. By contrast, Eaton-Kortum products account for more than 90 percent of patenting sector exports in each of China, Brazil, India, Russia and Mexico. 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 international variation in innovation levels leads to stark within-sector specialization in exports across product types.

To assess whether our calibrated model delivers reasonable counterfactuals, we carry out two exercises that have a counterpart in the literature. First, in 1995 the US altered its patent length from 17 years after grant to 20 years after application in order to implement TRIPS. This change had differential impacts on patent protection across technology fields depending upon the average lag between applications and grants. Using this variation, Bertolotti (2023) estimates that, without including anticipation effects, a one month increase in US patent protection raised patenting in the US by around 2 percent. Correspondingly, we compute the effect on patenting of increasing the expected duration of US patent protection by one month (starting from our calibrated steady state at time zero). We find a similar value: domestic patenting by US innovators $PAT_{US,US}$ increases by 0.9 percent at time zero.¹³

¹³Bertolotti (2023) estimates the change in patenting relative to a pre-TRIPS baseline. For comparison, using the

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 a similar trade cost reduction exercise in our calibrated model and find that patenting increases by 1.9 percent on average across countries.¹⁴ While Coelli et al. (2022) estimate a partial equilibrium effect of trade integration on patenting, we calculate the change including general equilibrium effects. Nevertheless, the similarity of our results is, again, reassuring.

4 Patent Policy

We use the calibrated model to study optimal patent policy by analyzing 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. 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 between the initial and new steady states affect welfare changes.

Obtaining steady state welfare from equation (32) and using SS, O to denote the initial steady state and SS, N to denote the new steady state gives a simple expression for the change in steady

pre-TRIPS calibration of our model from Section 4.4 below, we compute that increasing the expected duration of US patent protection by one month raises domestic patenting by US innovators at time zero by 1.2 percent.

¹⁴Coelli 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, particularly when domestic patenting is more prevalent. To simulate an equivalent change in the export costs faced by country i in our model, we reduce the trade costs τ_{ni} to each foreign destination n by $(1/1.032) \left(\sum_{n=1}^N PAT_{ni} / \sum_{n=1, n \neq i}^N PAT_{ni} \right)$ percent, where 1.032 is the average tariff in the Coelli et al. sample in 1992. We also report the change in patenting eight years after the trade cost shock in line with the timing of their exercise.

state welfare:

$$EV_i^{SS} = \frac{\tilde{C}_i^{SS,N}}{\tilde{C}_i^{SS,O}} \left[\frac{\rho - g^{SS,O} \left(1 - \frac{1}{\gamma}\right)}{\rho - g^{SS,N} \left(1 - \frac{1}{\gamma}\right)} \right]^{\frac{\gamma}{\gamma-1}}.$$

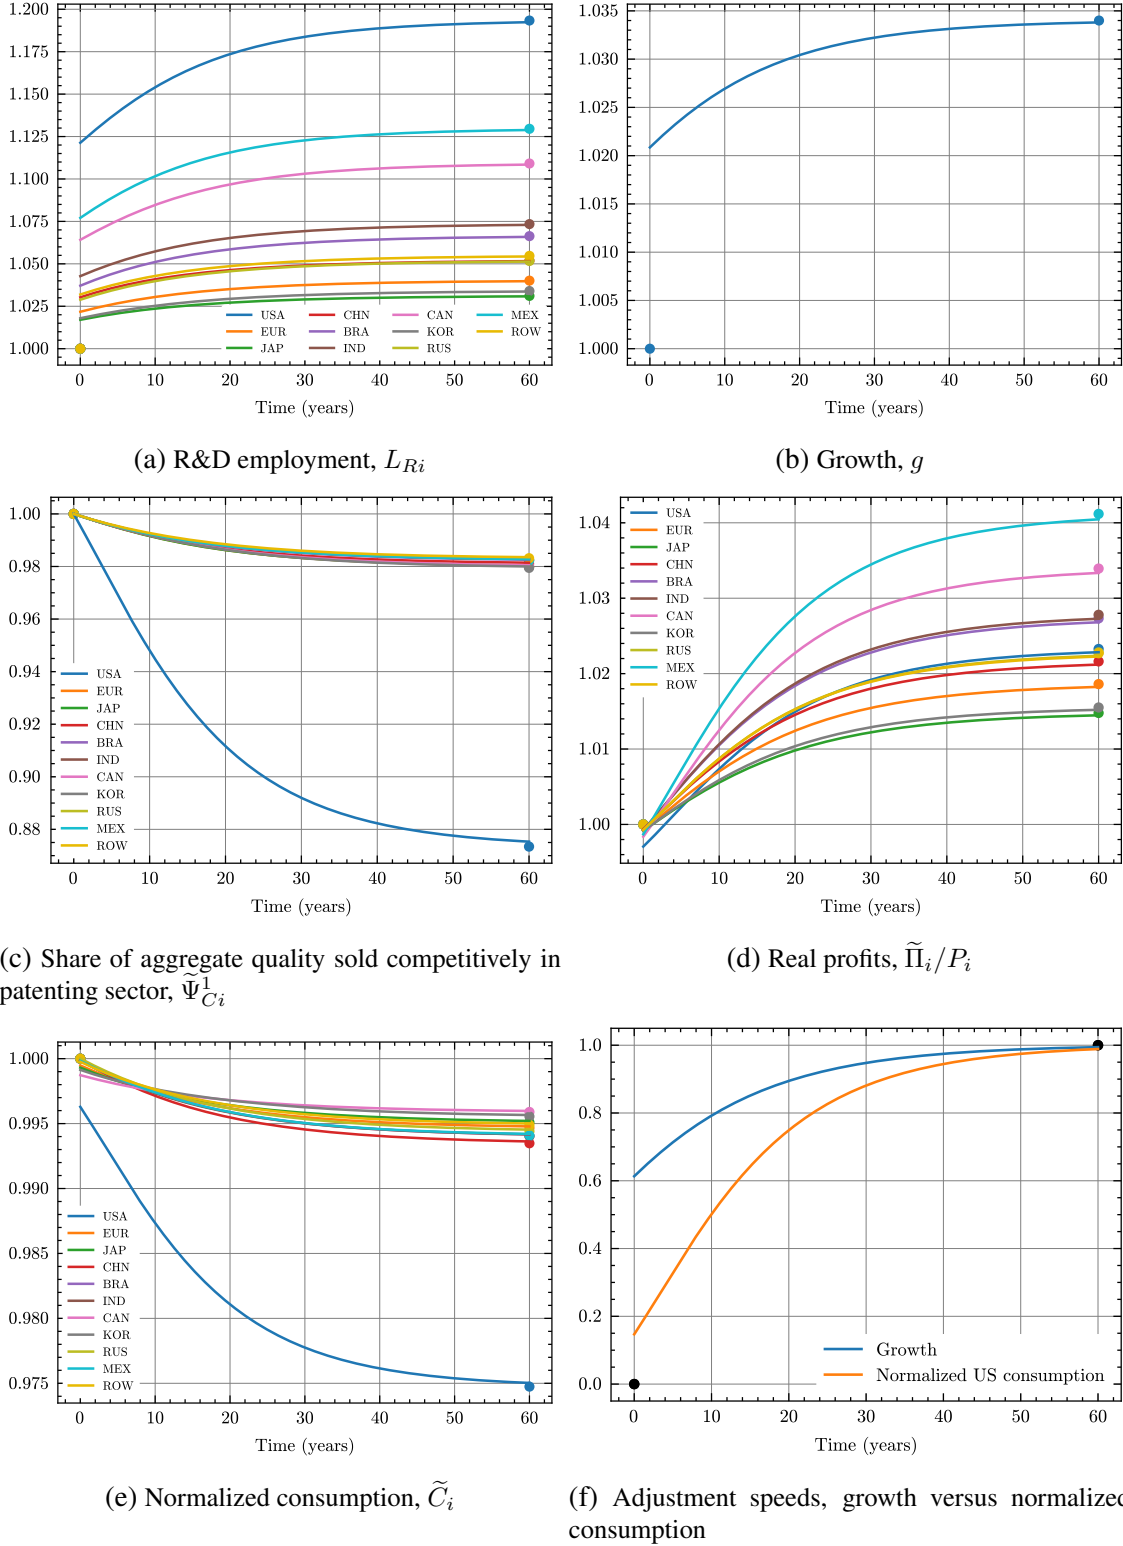
By contrast, to calculate EV_i we need to know the entire dynamic equilibrium starting from the initial steady state. Appendix C.1 explains how we solve for the transition dynamics between steady states. To obtain a fast and stable solution algorithm, we compute numerical derivatives using a projection method that decomposes functions on the space of Chebyshev polynomials (Judd 1992, Trefethen 2000), rather than using finite differences.

When calculating the equivalent variation for the world as a whole, we consider two alternative measures of global welfare. 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 aggregate welfare in country i as defined by equation (3). Second, a measure that uses Negishi (1960) weights based on individuals' inverse marginal utility of consumption in the initial steady state $U_{W,Negishi} = \sum_{i=1}^N \left(\tilde{C}_i^{SS,O}/L_i\right)^{\frac{1}{\gamma}} L_i^{\frac{1}{\gamma}} U_i = \sum_{i=1}^N \left(\tilde{C}_i^{SS,O}\right)^{\frac{1}{\gamma}} U_i$. The Negishi measure puts greater weight on the welfare of richer economies than the equal weights measure since it implies a social planner has no incentive to redistribute income across countries in the initial steady state.

4.1 Unilateral Patent Policy

We explore the effect of unilateral changes in patent policy by varying δ_n in one country at a time, while holding patent protection in all other countries constant. To illustrate how the global economy adjusts to changes in patent policy, we start by analyzing a reduction in δ_{US} from 0.083 to 0.05. Figure 3 shows how key variables change 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 leads to increased 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 12 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 because they are most dependent on US sales (panel



Notes: This figure plots the effects of reducing δ_{US} from 0.083 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). The vertical axis in 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 the changes in growth and normalized US consumption relative to the changes in these variables between the initial and new steady states.

34
Figure 3: Transition dynamics with stronger US patent protection


d).¹⁵ 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 0.4 percent in the US, but rises by around 1 percent elsewhere.

The transition dynamics in panels (b) and (e) show that the growth rate adjusts more quickly than normalized consumption to changes in patent protection. 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 around 90 percent of its adjustment to the new steady state, whereas normalized consumption has completed 75 percent of its adjustment. This ranking of adjustment speeds holds across all the counterfactuals we consider below. Consequently, accounting for transition dynamics increases the welfare change in countries that strengthen patent protection, i.e. $EV_i - EV_i^{SS} > 0$, because the benefits from higher growth materialize more quickly than the costs from increased market power. Likewise, transition dynamics reduce the welfare change in countries that weaken protection because the growth rate adjusts to its new lower steady state level more quickly than normalized consumption increases.

Next, we consider 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 4 plots the steady state world 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 Brazil, India, Russia, Mexico and the rest of the world. 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.

Turning to welfare, Figure 5 plots welfare effects EV_i by country, and for the world as a whole, against proportional changes in δ_n . In each panel we change patent protection in a different country. 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 also for the world as a

¹⁵The 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 in Mexico, Canada, India and Brazil.



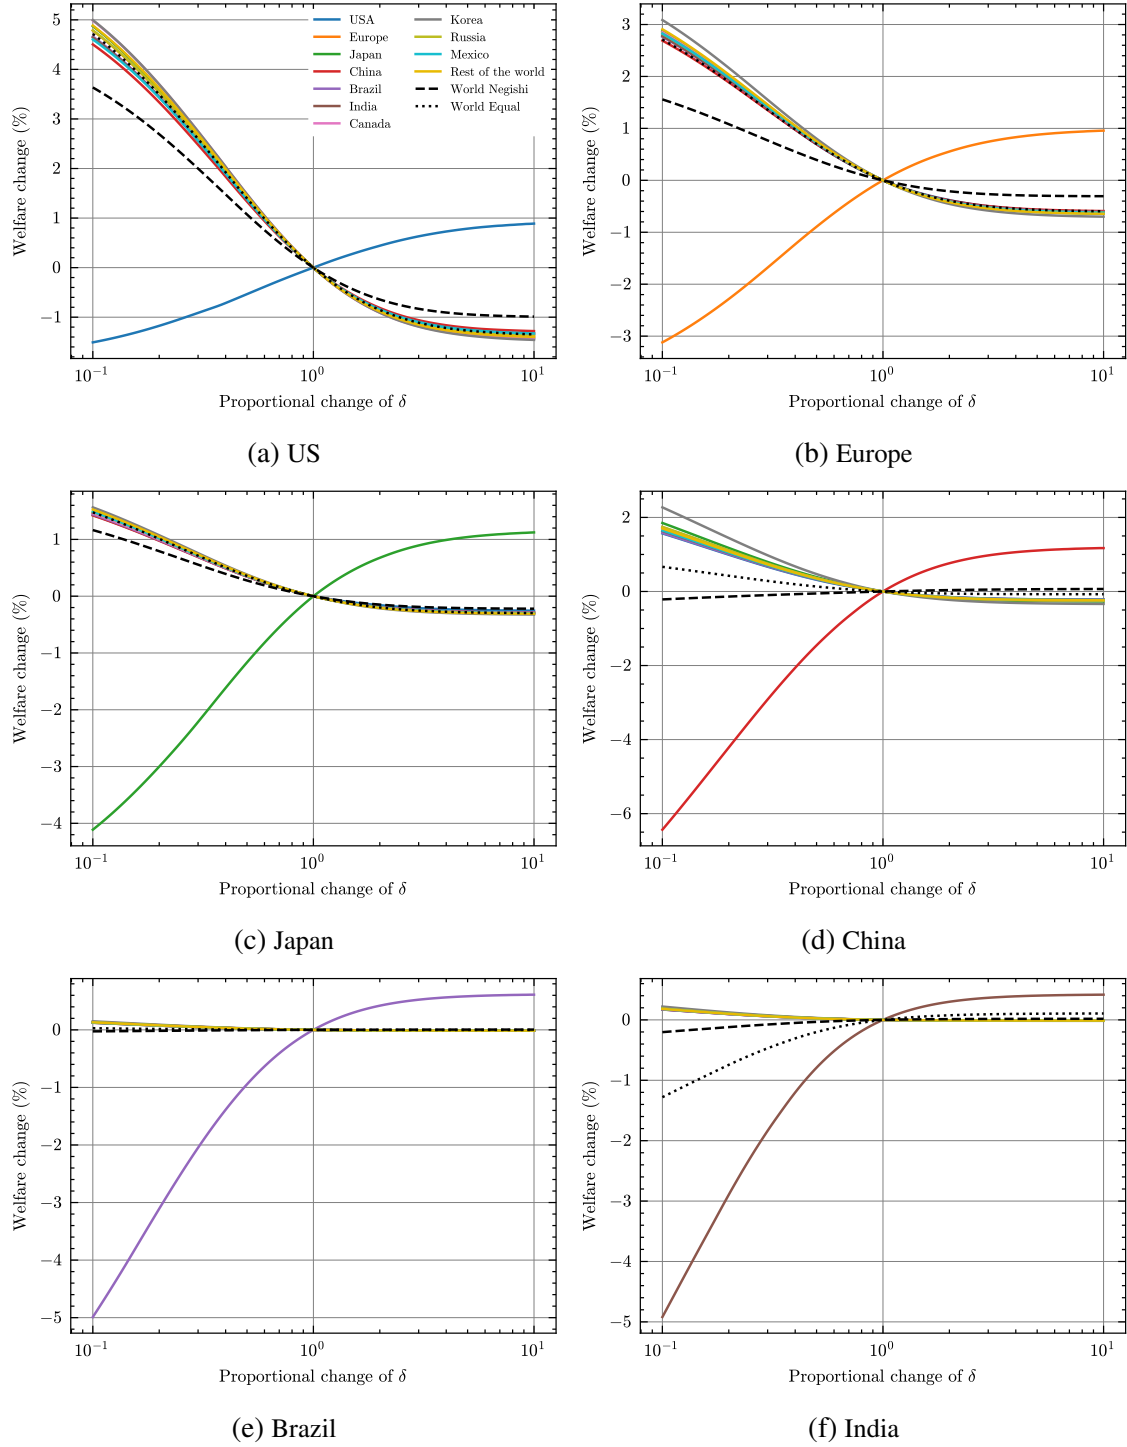
Notes: This figure plots the effect of a proportional change in calibrated patent protection δ_n of country n on the steady state growth rate g .

Figure 4: Unilateral patent policy changes and growth

whole. However, in the US, the dynamic benefits of stronger protection are offset by the static costs of higher prices due to increased market power in the patenting sector. As in the $\delta_{US} = 0.05$ counterfactual considered above, we find that these static costs are slightly larger than the dynamic benefits for the US. Consequently, reducing δ_{US} leads to small falls in US welfare.

Panels (b)-(f) show welfare effects from varying δ_n in Europe, Japan, China, India and Brazil, respectively. Each figure shows similar patterns, but with different magnitudes (note that the scale of the vertical axes differs across panels). Strengthening patent protection in country n raises welfare in countries $i \neq n$, but reduces welfare in country n itself. Thus, all countries have a unilateral incentive to weaken patent protection starting from the calibrated steady state. Compared to the US case in panel (a), stronger protection generates smaller spillover benefits for foreign countries because it has a weaker effect on growth (as shown in Figure 4). For Europe and Japan, the benefits of stronger protection to foreign countries exceed the domestic costs and world welfare rises. However, for China, Brazil and India, domestic costs and foreign benefits are more evenly balanced and stronger protection reduces world welfare using Negishi weights.¹⁶

¹⁶To conserve space we do not show the plots for all countries. The welfare effects of changing δ_n in Canada, Korea, Russia or Mexico are similar to those shown for Brazil in panel (e), except that stronger protection in Canada, Korea or Mexico generates slightly larger welfare increases for other countries than stronger protection in Brazil. Consequently, world welfare rises as protection becomes stronger in Canada, Korea or Mexico.



Notes: This figure plots the effect of proportional changes in calibrated patent protection δ_n of country n on welfare in all countries relative to the calibrated steady state. The legend in panel (a) applies to all panels. Welfare changes are expressed as the equivalent variation in consumption and account for transition dynamics.

Figure 5: Unilateral patent policy changes and welfare

4.2 Nash Equilibrium

How does the calibrated steady state compare to a Nash equilibrium where each country chooses patent protection δ_n to maximize its own welfare taking the response functions of other countries as given? To address this question, we solve for a Nash equilibrium in which each country makes a one-off, permanent change in its patent protection in order to maximize its welfare including transition dynamics and starting from the calibrated steady state.¹⁷ We bound the expected duration of each country's patent 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 of this game is unique.¹⁸

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.

With all countries offering weaker patent protection, there is less innovation and the steady state growth rate g is 0.14 percentage points lower than in the calibrated economy. The reduction in growth is small because of the slow pace of technology diffusion, which means that innovators expect to have a relatively long-lasting technological monopoly even when there is no patent protection. Lower growth is partially offset by a reduction in market power. The share of monopolistic products in world output of the patenting sector falls from 9.6 percent to 7.3 percent. However, in welfare terms, the fall in growth dominates and world welfare is 2.4 percent lower in the Nash equilibrium using equal weights and 1.6 percent lower using Negishi weights. It follows that, by providing some patent protection, the observed global equilibrium in 2015 generates moderately higher welfare than the Nash equilibrium.

All countries are worse off in the Nash equilibrium. Losses range from 0.3 percent in Korea to 2.8 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 changes in patent protection and, therefore, experience greater reductions in domestic market power. Table 3 also shows that accounting for transition dynamics increases the welfare costs of moving to the Nash equilibrium. Consistent with the analysis in Section 4.1, this occurs because the costs of lower growth are realized more quickly than the benefits of reduced market power.

¹⁷Alternatively, we could assume countries seek to maximize steady state welfare. However, we find that this alternative approach makes no difference to patent protection levels either in the Nash equilibrium or in the cooperative equilibria analyzed in Section 4.3.

¹⁸Appendix C.2 provides further details about how we solve for the Nash equilibrium and for the cooperative equilibria.

Table 3: Nash and cooperative equilibria

	(a) Nash			(b) Cooperative: equal weights			(c) Cooperative: Negishi weights		
	Patent protection, δ_i	Welfare change (percent) Total, EV_i	Steady state, EV_i^{ss}	Patent protection, δ_i	Welfare change (percent) Total, EV_i	Steady state, EV_i^{ss}	Patent protection, δ_i	Welfare change (percent) Total, EV_i	Steady state, EV_i^{ss}
US	None	-0.5	-0.4	Complete	2.1	-1.4	Complete	1.9	-1.7
Europe	None	-1.3	-1.1	Complete	2.3	-1.8	Complete	2.2	-2.0
Japan	None	-1.5	-1.3	Complete	2.5	-1.4	Complete	2.3	-1.6
China	None	-1.3	-1.0	None	10.2	13.1	None	10.0	12.9
Brazil	None	-2.1	-1.8	None	9.3	12.1	None	9.2	12.0
India	None	-2.3	-2.0	None	9.1	11.9	None	9.0	11.8
Canada	None	-0.9	-0.7	Complete	3.9	1.8	Complete	3.8	1.6
Korea	None	-0.3	-0.1	Complete	3.8	1.9	Complete	3.6	1.6
Russia	None	-2.4	-2.1	None	9.2	12.2	None	9.1	12.0
Mexico	None	-2.3	-2.1	Complete	1.8	-5.9	None	8.9	11.3
Rest of world	None	-2.8	-2.5	None	8.9	12.4	None	8.8	12.2
World Equal		-2.4	-2.1		8.8	11.6		8.8	11.6
World Negishi		-1.6	-1.3		6.0	5.4		6.0	5.6

Growth rate change (percentage points) -0.14 0.69 0.68

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.

4.3 Cooperative Equilibrium

We assess the potential gains from global cooperation over patent policy by solving 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. Column (b) reports the case where the objective function uses equal weights and column (c) the case with Negishi weights. In each case we find that there is a unique cooperative equilibrium.

Starting with equal weights, we see that 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. This pattern occurs because growth is more sensitive to patent protection in more innovative economies, as shown in Section 2.4. 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. Consequently, it is efficient for the world to delegate the job of incentivizing innovation to these economies. The exception to this logic is Mexico, which accounts for only 0.1 percent of global innovations in the calibrated steady state. Nevertheless, it is optimal for Mexico to provide complete patent protection in the cooperative equilibrium due to its relatively high levels of trade with the US. The dynamic benefits that stronger protection in Mexico generates through increased innovation in the US outweigh the static costs from higher market power in Mexico.¹⁹

With stronger incentives to innovate, steady state growth is 0.69 percentage points greater than in the calibrated equilibrium. And faster growth leads to higher welfare. World equal weights welfare is 8.8 percent higher in the cooperative equilibrium. This increase is over 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.²⁰ However, the overall gains mask strong distributional effects. While welfare in China, Brazil, India, Russia and the rest of the world rises by nearly 10 percent, the gains for countries that provide complete patent protection are much smaller, ranging from 1.8 percent in Mexico to 3.9 percent in Canada. This cross-country variation occurs because the static costs of encouraging innovation are borne by the economies that offer protection, while other countries free ride on the dynamic benefits from higher growth.

The cooperative equilibrium with Negishi weights is similar to the equal weights equilibrium

¹⁹Section 2.4 and Figure 4 also show that growth is more sensitive to patent protection in larger economies. However, the static costs of increased market power are greater when borne by a larger population. Consequently, it is a country's R&D efficiency not its size that primarily determines its patent policy in the cooperative equilibrium.

²⁰Specifically, we find that shutting down trade in the patenting sector reduces world equal weights welfare by 16.6 percent with country-level losses ranging from 14.3 percent in China to 19.6 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 4.6 percent fall in world equal weights welfare.

except that Mexico switches from complete protection to no protection. This change slightly dampens the increase in growth, leading to lower welfare gains for all countries except Mexico. However, abandoning patent protection brings large gains for Mexico whose welfare increases by 8.9 percent. On a Negishi weights basis, world welfare rises by 6.0 percent. In both the Negishi weights and equal weights equilibria, all countries experience welfare gains. However, this conclusion holds only when welfare changes are calculated including transition dynamics. Steady state welfare in the US, Europe and Japan is lower in the cooperative equilibria.

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.²¹


4.4 TRIPS

The TRIPS Agreement came into effect at the start of 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 intellectual property rights 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 intellectual property rights 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 intellectual property rights, but had little or no impact on policies in developed countries (Saggi 2016).

Figure 6 shows the evolution of the number of international patent family applications by destination in our data. Between 1990 and 2015 there was a rapid increase in applications filed in China, India, Russia and, to a lesser extent, Brazil. This observation is consistent with patent rights be-

²¹At the equal weights optimum $\delta = 0.11$, steady state world welfare using equal weights is 1.8 percent lower than in the calibrated steady state. For the Negishi weights optimum $\delta = 0.07$, steady state world welfare using Negishi weights is 0.4 percent lower than in the calibrated steady state.

coming stronger in these countries during this period, but it could also result from increased market size due to rapid economic growth. We use the model to disentangle these alternatives and quantify changes in patent rights by calibrating patent protection prior to TRIPS.



Notes: Count of cross-border patent applications by destination at the family level.

Figure 6: International patent flows by destination

Since countries may have initiated patent reforms in anticipation of TRIPS, we use 1992 data for the pre-TRIPS calibration. We recalibrate those 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.

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 patent protection between 1992 and 2015 is more plausible for developing than for developed countries.²² Therefore, our main pre-TRIPS counterfactual in column (b) of Table 4 analyzes the case where patent protection in developing countries (China, Brazil, India, Russia and Mexico) 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

²²China 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 patent protection in China and Russia during our sample.

Table 4: Pre-TRIPS counterfactuals

	(a)		(b)	(c)	(d)	(e)
	Patent protection, δ_i		Developing countries	All countries	Developing countries and trade costs	Harmonize patent protection with US
	Baseline, 2015	Pre-TRIPS, 1992	Welfare change, EV_i (percent)			
US	0.083	0.104	-0.16	-0.11	-1.01	1.12
Europe	0.101	0.112	-0.18	-0.46	-0.95	0.77
Japan	0.107	0.103	-0.19	-0.67	-1.07	0.73
China	0.143	0.269	0.70	0.24	-0.06	-0.31
Brazil	0.187	0.217	-0.03	-0.49	-0.87	-0.14
India	0.258	0.594	0.17	-0.29	-0.76	-0.68
Canada	0.079	0.051	-0.17	-0.66	-1.51	1.16
Korea	0.065	0.080	-0.23	-0.24	-1.37	1.37
Russia	0.248	0.440	0.07	-0.40	-0.98	-0.32
Mexico	0.229	0.212	-0.22	-0.62	-1.64	-0.12
Rest of world	1.626	1.392	-0.17	-0.65	-1.24	-1.08
World Equal			0.03	-0.43	-0.96	-0.79
World Negishi			0.05	-0.28	-0.84	0.09

Notes: This table reports calibrated levels of patent protection in 1992 (pre-TRIPS) and 2015 (post-TRIPS) (column a) and welfare changes for different patent policy and trade cost counterfactuals (columns b to e). Welfare changes are expressed as the equivalent variation in consumption relative to the calibrated steady state in 2015 and account for transition dynamics. In column (b), China, Brazil, India, Russia and Mexico 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 and Mexico revert to pre-TRIPS patent protection and all trade costs in the patenting sector increase by 3 percent. In column (e), all countries with weaker patent protection than the US in 2015 set $\delta_i = \delta_{US}$.

pre-TRIPS patent policies benefits China, India, Russia 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 0.70 percent in China, 0.17 percent in India, 0.07 percent in Russia and 0.03 percent for the world using equal weights. We conclude that, to the extent TRIPS increased patent protection in developing countries, it reduced welfare in those countries that substantially strengthened protection. These countries faced local static costs from increased market power, but did not realize offsetting dynamic benefits because growth is inelastic to patent policy in developing countries.²³

However, there are two important caveats to this conclusion. First, the welfare changes in column (b) are small relative to those implied by other patent policy counterfactuals, such as the

²³Table 4 reports pre-TRIPS counterfactuals relative to the 2015 baseline, but simulating TRIPS using the 1992 calibration leads to the same conclusions. For example, if we start in the calibrated steady state in 1992 and then increase patent protection in developing countries to 2015 levels, welfare falls by 0.56 percent in China, 0.23 percent in India, 0.19 percent in Russia and 0.25 percent for the world using equal weights, but increases slightly in developed countries.

cooperative equilibrium in Section 4.3. Consequently, small increases in patent protection in developed countries are sufficient to offset the costs of TRIPS. The findings in column (c) demonstrate this point. When all countries with stronger protection in 2015 revert to pre-TRIPS protection levels, Chinese welfare increases by only 0.24 percent and welfare declines in all other countries with world equal weights welfare falling by 0.43 percent. These results imply that the welfare effects of changes in calibrated patent protection between 1992 and 2015 are primarily driven by 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 reforms in intellectual property laws. 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. However, existing research suggests that the benefits from greater trade integration may have offset our estimates of the welfare costs of TRIPS. For example, Harrison et al. (1997) forecast that reductions in trade barriers and agricultural subsidies due to the Uruguay Round will raise GDP in developing countries by 0.4–1.4 percent. And reviewing thirteen ex-ante assessments of the Uruguay Round, Martin and Winters (1995) conclude that cuts to merchandise trade protection will increase real incomes in developing countries by between 1.2 and 2.0 percent. Similarly, Caliendo et al. (2023) find that most-favored nation tariff cuts between 1990 and 2010 increased welfare by 2–3 percent, with larger gains in emerging and developing countries.

To quantify the trade-off between TRIPS and market access in our model, we study counterfactuals that change both patent protection and trade costs simultaneously. Column (d) reports welfare changes when we combine the pre-TRIPS counterfactual in column (b) with a 3 percent increase in all trade costs in the patenting sector. The increase in trade costs dominates and welfare declines in all countries. We can also use the model to solve for the trade cost increase that offsets the benefits of reverting to pre-TRIPS patent protection in China, India and Russia. For each of these countries in turn, we simulate the pre-TRIPS counterfactual in column (b) while also increasing the country’s trade costs in the patenting sector (holding trade costs fixed for all other country pairs), and solve for the uniform increase that leaves the country’s welfare unchanged. We find that a 6.0 percent increase makes China indifferent, while for India and Russia the numbers are 1.2 percent and 0.4 percent, respectively. These results show that relatively small increases in trade costs are sufficient to offset the benefits to developing countries of reverting to pre-TRIPS patent protection. It follows that the net welfare effect of joining the WTO may have been positive even for countries that suffered losses due to TRIPS.

Our baseline calibration implies that TRIPS did not lead to harmonization of patent protec-

tion across countries. Developing countries continue to offer weaker protection than developed economies in 2015. Nevertheless, we can also use our model to 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 all developing countries. Welfare falls by 0.3 percent in China, 0.7 percent in India and 0.8 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 Additional Counterfactuals

To conclude the counterfactual analysis we briefly discuss how changes in the global economy (compared to the calibrated steady state in 2015) would affect optimal patent policy. We consider the consequences of trade integration, faster technology diffusion and the rise of China. Appendix C.3 includes more detailed results for the trade integration and faster technology diffusion cases.

Trade integration. Lower trade costs reduce home bias in consumption leading to stronger international spillovers from patent policy. Counterfactual simulations with different trade costs show that this effect increases the incentive for countries to free ride by unilaterally weakening patent protection. In particular, the derivative of each country's welfare with respect to its own δ becomes more positive as trade costs fall, implying that there are greater benefits from reducing patent protection. Conversely, higher trade costs reduce the benefits of unilaterally weakening protection and, for sufficiently high trade costs, some countries provide patent protection in the Nash equilibrium. For example, if we double all trade costs in the patenting sector and then solve for the Nash equilibrium, we find that the US chooses to provide complete patent protection, while all other countries offer no protection. However, patent protection levels in the cooperative equilibria are insensitive to changes in trade costs since countries jointly internalize spillovers when they cooperate. Doubling trade costs in the patenting sector does not affect the patent protection any country offers in the cooperative equilibrium with either equal weights or Negishi weights.

Faster technology diffusion. Faster technology diffusion increases the value of patenting because it reduces the expected duration of an innovator's technological monopoly making R&D investment more sensitive to patent protection. 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 $\nu = 0.11$, then US welfare is

almost inelastic to unilateral changes in US patent protection, i.e. the US welfare line in Figure 5 (a) becomes flat (see Appendix C.3). In the Nash equilibrium with $\nu = 0.11$ the US provides complete patent protection, although other countries continue to offer no protection. And in the cooperative equilibrium with equal weights Russia joins the club of countries with complete patent protection. Since patenting is more valuable with faster technology diffusion, the welfare gains from optimal patent policy also rise. World equal weights welfare is 14.3 percent higher in the cooperative equilibrium with equal weights, compared to 8.8 percent higher in the baseline calibration.

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 6.4 percent in 2015. In our calibrations, China's R&D efficiency increased from 1.5 percent of US R&D efficiency in 1992 to 12 percent in 2015. With China's economy becoming larger and innovating more, global growth is becoming more sensitive to China's patent policy. We calculate that, all else constant, China should 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 22 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 2.7 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 assess the sensitivity of our calibration and the robustness of our counterfactual results to plausible variation in the parameters and moments used in the calibration. These checks are summarized in Appendix D. They show that the main findings of the unilateral, Nash, cooperative and TRIPS counterfactuals analyzed in Sections 4.1–4.4 are robust to moderate changes in the externally calibrated parameters and targeted moments.

The magnitude of counterfactual changes in welfare differ from the baseline results for some robustness checks. For example, welfare gains in the cooperative equilibrium are much larger when the demand elasticity σ^1 is smaller (implying innovators have greater market power), or the turnover moment is larger (implying faster technology diffusion). Both these alternatives increase the sensitivity of growth to patent protection, which raises the welfare gains from cooperation. However, even in these cases, qualitative variation in welfare effects across countries is similar to the baseline.

5 Conclusions

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 optimal patent policy in the global economy. To address this gap, we introduce a new quantitative model of trade and patents. By allowing innovation, patenting, and market power to respond endogenously to domestic and foreign patent policy, the model captures the trade-off between the static costs and dynamic benefits of stronger patent protection.

We calibrate the steady state of the model to match the world economy in 2015 and use the calibrated economy to analyze the impact of changes in patent policy. We study unilateral changes in patent protection, cooperative and non-cooperative equilibria, and the TRIPS agreement. The counterfactual results imply that there are large potential gains from global cooperation over patent policy. However, realizing these gains requires not that policies are harmonized across countries, but that larger and more innovative economies offer stronger protection. Moreover, the gains and losses from optimal patent policy are not equally shared between countries.

Two mechanisms drive these findings. First, international spillovers. The dynamic benefits from the higher growth caused by stronger patent rights are global in scope due to international knowledge spillovers, whereas the static costs due to higher prices are primarily borne domestically. Consequently, non-cooperative patent policies tend to be weaker than is globally efficient because countries do not fully internalize the dynamic benefits. Starting from the calibrated economy, all countries have an incentive to weaken their patent protection and in the Nash equilibrium no countries provide patent protection.

Second, heterogeneity in growth effects. The global growth rate is more sensitive to patent protection provided by countries that account for higher shares of innovators' profits. This means that stronger protection has greater benefits if provided by a large, innovative country such as the US, than if provided by a smaller, less innovative, and more isolated economy such as India. Consequently, optimal patent policy varies greatly across countries. We find that in the cooperative equilibrium, developed economies such as the US, Europe and Japan offer complete patent protection, whereas developing countries such as China, Brazil and India provide no protection. This pattern of protection increases growth and raises global welfare, but not all economies benefit equally. Countries that do not provide protection gain more than those that shoulder the burden of encouraging innovation.

The TRIPS agreement required developing countries to increase patent protection towards levels provided in developed economies. Our results imply that TRIPS reforms reduced welfare in developing countries and that further policy harmonization would only exacerbate these effects. This finding suggests that some of the opposition to TRIPS may be well-founded and that there is

scope for patent policy reform that both increases efficiency and reduces international inequality. However, our results do not imply that WTO membership has hurt developing countries, since the gains from improved market access likely dominate any costs due to TRIPS.

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. And although this trade-off is at the heart of policy debates over intellectual property rights, there are other channels through which patent policy may affect economic outcomes. Future research could extend our analysis by allowing patent protection to impact export market entry (Ivus 2015, Cockburn et al. 2016), foreign investment choices (Branstetter et al. 2011, Bilir 2014), knowledge spillovers (Moser and Voena 2012, Hegde et al. 2023), technology transfer within multinational firms (Branstetter et al. 2006) and/or investment in imitation and technology adoption (Santacreu 2023). It would also be interesting to shift the focus from multilateral to bilateral policies and study the TRIPS-plus provisions in many recent free trade agreements that further strengthen intellectual property rights (Mercurio 2006). The theory and calibration method that we have developed provide a framework that future research can build upon to better understand this important topic.

References

- Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. “Innovation, Reallocation, and Growth.” *American Economic Review* 108(11): 3450–3491.
- Acemoglu, Daron, James Robinson, and Thierry Verdier. 2017. “Asymmetric Growth and Institutions in an Interdependent World.” *Journal of Political Economy* 125(5): 1245–1305.
- Akcigit, Ufuk, Sina Ates, and Giammario Impullitti. 2021. “Innovation and Trade Policy in a Globalized World.” NBER Working Paper 24543.
- Akcigit, Ufuk, and Marc Melitz. 2021. “International Trade and Innovation.” NBER Working Paper 29611.
- 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. 2023. “Patent Term, Innovation, and the Role of Technology Disclosure Externalities.” Mimeo.
- Bilir, L. Kamran. “Patent Laws, Product Life-cycle Lengths, and Multinational Activity.” *American Economic Review* 104(7): 1979–2013.

- 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.
- Branstetter, Lee G., Raymond Fisman, and C. Fritz Foley. 2006. "Do Stronger Intellectual Property Rights Increase International Technology Transfer? Empirical Evidence from US Firm-level Panel Data." *Quarterly Journal of Economics* 121(1): 321-349.
- Branstetter, Lee, Ray Fisman, C. Fritz Foley, and Kamal Saggi. 2011. "Does Intellectual Property Rights Reform Spur Industrial Development?" *Journal of International Economics* 83(1): 27-36.
- Bureau of Economic Analysis (BEA). 2022. "Gross Output by Industry." (accessed April 22, 2022 and February 18, 2023).
- Cai, Jie, Nan Li, and Ana Maria Santacreu. 2022. "Knowledge Diffusion, Trade, and Innovation Across Countries and Sectors." *American Economic Journal: Macroeconomics* 14(1): 104-45.
- 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.
- 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.
- 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. 1996. "Trade in Ideas. Patenting and Productivity in the OECD." *Journal of International Economics* 40(3-4): 251-278.
- 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.
- Harrison, Glenn W., Thomas F. Rutherford, and David G. Tarr. 1997. "Quantifying the Uruguay Round." *Economic Journal* 107(444): 1405-1430.
- Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." In *Handbook of International Economics*, Vol.4, 131-195.
- Hegde, Deepak, Kyle Herkenhoff, and Chenqi Zhu. 2023. "Patent Publication and Innovation." *Journal of Political Economy* 131(7): 1845-1903.
- 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.
- Jakobsson, Amanda, and Paul Segerstrom. 2016. "The TRIPS Agreement: Imitation and Foreign

Direct Investment.”

- Jones, Charles I. 1995. “R&D-based Models of Economic Growth.” *Journal of Political Economy* 103(4): 759-784.
- 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.
- 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.
- 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.
- Melitz, Marc, and Stephen Redding. 2021. “Trade and Innovation.” NBER Working Paper 28945.
- Mercurio, Bryan Christopher. 2006. “Trips-Plus Provisions in FTAs: Recent Trends.” in *Regional Trade Agreements and the WTO Legal System*, Lorand Bartels, Federico Ortino, eds., 215-237, Oxford University Press.
- Moscona, Jacob. 2021. “Flowers of Invention: Patent Protection and Productivity Growth in US Agriculture.” Mimeo.
- Moser, Petra, and Alessandra Voena. 2012. “Compulsory Licensing: Evidence from the Trading with the Enemy Act.” *American Economic Review* 102(1): 396-427.
- 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.
- Sampson, Thomas. 2023. “Technology Gaps, Trade, and Income.” *American Economic Review* 113(2): 472-513.
- Santacreu, Ana Maria. 2023. “Dynamic Gains from Trade Agreements with Intellectual Property Provisions.” Mimeo.
- Schott, Peter. 2008. “The Relative Sophistication of Chinese Exports.” *Economic Policy* 23(53): 6-49.
- Trefethen, Lloyd N. 2000. “Spectral Methods in MATLAB.” Society for Industrial and Applied Mathematics.
- 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

Cost minimization using the sectoral production function (1) implies that the sectoral price index satisfies:

$$\begin{aligned}
 P_n^s &= \left(\int_0^{M^s} \psi(\omega) p_n^s(\omega)^{1-\sigma^s} d\omega \right)^{\frac{1}{1-\sigma^s}}, \\
 &= \left[(P_{Cn}^s)^{1-\sigma^s} + (P_{Mn}^s)^{1-\sigma^s} \right]^{\frac{1}{1-\sigma^s}}, \\
 &= \left[\Gamma \left(\frac{\theta^s + 1 - \sigma^s}{\theta^s} \right) \right]^{\frac{1}{1-\sigma^s}} \left[\Psi_{Cn}^s \left(\sum_{j=1}^N \Phi_{nj}^s \right)^{\frac{\sigma^s-1}{\theta^s}} + \left(\frac{\sigma^s}{\sigma^s-1} \right)^{1-\sigma^s} \sum_{j=1}^N \Psi_{Mnj}^s (\Phi_{nj}^s)^{\frac{\sigma^s-1}{\theta^s}} \right]^{\frac{1}{1-\sigma^s}},
 \end{aligned} \tag{38}$$

where the final equality uses equations (8) and (11) to substitute for P_{Mn}^s and P_{Cn}^s , respectively.

Likewise, cost minimization using the final good production function (2) implies that sectoral output satisfies:

$$P_n^s Y_n^s = \beta^s P_n Y_n, \tag{39}$$

where the final good price index is given by:

$$P_n = \prod_{s=0}^S (P_n^s)^{\beta^s}. \tag{40}$$

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 revenue. Consequently, aggregate profits made by monopolists in i from sales to n are given by X_{Mni}^s/σ^s . Setting production income equal to total sales in each sector then yields:

$$w_i L_{Yi}^s + P_i Q_i^s + \frac{1}{\sigma} \sum_{n=1}^N X_{Mni}^s = \sum_{n=1}^N X_{ni}^s, \tag{41}$$

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

$$X_{ni}^s = X_{Cni}^s + X_{Mni}^s. \quad (42)$$

The variety production technology (4) implies that intermediate input expenditure equals a fraction $1 - \alpha^s$ of production costs. Therefore, at the sectoral level we have:

$$P_i Q_i^s = \frac{1 - \alpha^s}{\alpha^s} w_i L_{Yi}^s. \quad (43)$$

Substituting this expression into equation (41) and rearranging yields:

$$L_{Yi}^s = \frac{\alpha^s}{w_i} \left(\sum_{n=1}^N X_{ni}^s - \frac{1}{\sigma} \sum_{n=1}^N X_{Mni}^s \right). \quad (44)$$

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

$$w_i L_{Yi} = \sum_{s=0}^S \alpha^s \sum_{n=1}^N X_{ni}^s - \frac{1}{\sigma} \sum_{s=1}^S \alpha^s \sum_{n=1}^N X_{Mni}^s, \quad (45)$$

where L_{Yi} is total labor employed in production in country i , which we take as given in the static equilibrium.

Final good market clearing in each country requires that:

$$Y_i = C_i + \sum_{s=0}^S Q_i^s,$$

and substituting for Q_i^s using equation (43) yields:

$$Y_i = C_i + \sum_{s=0}^S \frac{1 - \alpha^s}{\alpha^s} \frac{w_i}{P_i} L_{Yi}^s. \quad (46)$$

We allow for the possibility of exogenous trade imbalances. Let TB_i be the trade surplus of country i relative to the value of global final output. Accounting for trade in both varieties and patenting services and setting the trade balance plus imports equal to exports gives:

$$TB_i \sum_{n=1}^N P_n Y_n + \sum_{s=0}^S \sum_{n=1}^N X_{in}^s + \sum_{s=1}^S \sum_{n=1}^N w_n L_{in}^{s,e} = \sum_{s=0}^S \sum_{n=1}^N X_{ni}^s + \sum_{s=1}^S \sum_{n=1}^N w_i L_{ni}^{s,e}, \quad (47)$$

where we note that $\sum_{n=1}^N X_{in}^s = P_i^s Y_i^s$, which implies $\sum_{s=0}^S \sum_{n=1}^N X_{in}^s = P_i Y_i$.

Letting the final good in country one be the numeraire, meaning $P_1 = 1$, we can now define the static equilibrium as follows.

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 (40) subject to the normalization $P_1 = 1$;
- N income equals sales equations (45), and;
- N trade balance equations (47), where:
- P_n^s are defined in (38); Φ_{ni}^s are defined in (9); X_{ni}^s are defined in (42); X_{Mni}^s are defined in (10); P_{Mn}^s are defined in (8); X_{Cni}^s are defined in (12); P_{Cn}^s are defined in (11), and; Y_n^s are defined in (39). The allocation of production labor across sectors L_{Yi}^s is then given by (44) and aggregate consumption C_i by (46).

A.2 Dynamic Equilibrium

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

$$r_{nt} = \rho + \frac{1}{\gamma} \left(\frac{\dot{C}_{nt}}{C_{nt}} + \frac{\dot{P}_{nt}}{P_{nt}} \right), \quad (48)$$

and the transversality condition:

$$\lim_{t \rightarrow \infty} \exp \left(- \int_{t_0}^t r_{nt} \tilde{dt} \right) W_{nt} = 0. \quad (49)$$

Patenting thresholds. Equation (16) gives the quality threshold above which firms in country i choose to patent in country n if they have paid the patent preparation cost $w_i f_i^{s,o}$. Paying this cost gives firms the option of applying for a patent in each destination. Therefore, a firm that creates an invention with quality ψ at time t_0 opts to pay the preparation cost if and only if:

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

The left hand side of this expression gives the value of patenting net of application costs, while the right hand side is the patent preparation cost. We can rewrite the inequality as:

$$\sum_n \max \left[\psi \left(V_{nit_0}^{s,P} (1) - V_{nit_0}^{s,NP} (1) \right) - \frac{w_n f_n^{s,e}}{\Psi^s}, 0 \right] \geq \frac{w_i f_i^{s,o}}{\Psi^s}.$$

Let $n_i^* \equiv \arg \min_n \psi_{ni}^{s,e*}$ denote the country n with the lowest threshold in equation (16). 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 $\psi_i^{s,o*}$ defined by:

$$\begin{aligned} \psi_i^{s,o*} &= \psi_{n_i^*}^{s,e*} \text{ if } \sum_n \max \left[\psi_{n_i^*}^{s,e*} \left(V_{nit_0}^{s,P} (1) - V_{nit_0}^{s,NP} (1) \right) - \frac{w_n f_n^{s,e}}{\Psi^s}, 0 \right] \geq \frac{w_i f_i^{s,o}}{\Psi^s}, \\ \sum_n \max \left[\psi_i^{s,o*} \left(V_{nit_0}^{s,P} (1) - V_{nit_0}^{s,NP} (1) \right) - \frac{w_n f_n^{s,e}}{\Psi^s}, 0 \right] &= \frac{w_i f_i^{s,o}}{\Psi^s} \text{ otherwise,} \end{aligned} \quad (50)$$

such that only firms with quality $\psi \geq \psi_i^{s,o*}$ pay the patent preparation cost.

Value of invention. Quality is drawn from a Pareto distribution with scale parameter one and shape parameter k . An innovator from country i and sector s pays the patent preparation cost if quality exceeds $\psi_i^{s,o*}$ and patents in country n if quality exceeds ψ_{ni}^{s*} . Each innovation creates Ψ^s varieties. 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 present discounted value of profits per variety that a time t innovator makes from sales to destination n is:

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

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 (18).

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:

$$L_{in}^{s,e} = \eta_i^s (L_{Ri}^s)^{1-\kappa} (\psi_{ni}^{s*})^{-k} f_n^{s,e}. \quad (51)$$

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

$$L_i^{s,o} = \eta_i^s (L_{Ri}^s)^{1-\kappa} (\psi_i^{s,o*})^{-k} f_i^{s,o}. \quad (52)$$

The labor market clearing condition is then given by:

$$L_i = L_{Yi} + \sum_{s=1}^S L_{Ri}^s + \sum_{s=1}^S L_i^{s,o} + \sum_{s=1}^S \sum_{n=1}^N L_{ni}^{s,e}. \quad (53)$$

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 (53) holds with employment in patenting given by (51) and (52);
- $\Psi_{Mni}^{s,NP}$, $\Psi_{Mni}^{s,P,ND}$, $\Psi_{Mni}^{s,P,D}$ and Ψ_{Cn}^s satisfy the laws of motion in (20) – (23) and Ψ_{Mni}^s is given by (19);
- $V_{nit}^{s,NP}(1)$ and $V_{nit}^{s,P}(1)$ are defined by (13) and (14) with $\psi = 1$ and expected profits obeying (15);
- V_{it}^s is given by (18);
- ψ_{ni}^{s*} and $\psi_i^{s,o*}$ are defined by (16), (17) and (50);
- L_{Ri}^s satisfies the innovation free entry condition (5), and;
- r_i satisfies the Euler equation (48) and the transversality condition (49) holds.

A.3 Steady State

Growth rates. We solve for a steady state equilibrium. Labor market clearing (53) implies that the allocation of labor to R&D, patenting and production in each country is constant in steady state. Equations (51) and (52) then imply that the patenting thresholds ψ_{ni}^{s*} and $\psi_i^{s,o*}$ are constant.

Let g^s denote the growth rate of Ψ^s . Equations (6) and (19) imply that the aggregate qualities of Helpman-Krugman and Eaton-Kortum products in sector s denoted by 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. Note that the growth rate of aggregate quality is the same in all countries. Using equations (8), (9), (46) and (47) 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 (46) implies that wages w_i and final good output Y_i also grow at rate g , while the trade balance condition (47) implies that trade flows X_{ni}^s grow at rate g . Equation (9) 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:

$$g_{P^s} = \alpha^s g - \frac{g^s}{\sigma^s - 1}. \quad (54)$$

From equation (39) we have that sectoral output Y_n^s grows at rate:

$$g_{Y^s} = g - g_{P^s},$$

and the production technology (2) implies:

$$g = \sum_{s=0}^S \beta^s g_{Y^s}.$$

Combining the three equations above implies that g satisfies equation (25) in the main text.

The Euler equation (48) then implies that the steady state interest rate is constant across countries and given by:

$$r = \rho + \frac{g}{\gamma}, \quad (55)$$

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 (13), (14) and (18) 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 (15).

Laws of motion for normalized aggregate qualities. Normalizing each of the aggregate quality variables by Ψ^s , the laws of motion in equations (20)–(23) can be rewritten as:

$$\begin{aligned} (g^s + \nu^s + \zeta^s) \tilde{\Psi}_{Mni}^{s,NP} &= \eta_i^s (L_{Ri}^s)^{1-\kappa} \frac{k}{k-1} \left[1 - (\psi_{ni}^{s*})^{1-k} \right] + \delta_n^s \tilde{\Psi}_{Mni}^{s,P,ND}, \\ (g^s + \delta_n^s + \nu^s + \zeta^s) \tilde{\Psi}_{Mni}^{s,P,ND} &= \eta_i^s (L_{Ri}^s)^{1-\kappa} \frac{k}{k-1} (\psi_{ni}^{s*})^{1-k}, \\ (g^s + \delta_n^s + \zeta^s) \tilde{\Psi}_{Mni}^{s,P,D} &= \nu^s \tilde{\Psi}_{Mni}^{s,P,ND}, \\ (g^s + \zeta^s) \tilde{\Psi}_{Cn}^s &= \sum_{i=1}^N \left(\nu^s \tilde{\Psi}_{Mni}^{s,NP} + \delta_n^s \tilde{\Psi}_{Mni}^{s,P,D} \right). \end{aligned} \quad (56)$$

Patenting thresholds. Substituting equations (26) and (27) into (50) and using the definition of normalized profits in (28), the threshold for paying the application preparation cost satisfies:

$$\begin{aligned} \psi_i^{s,o*} &= \psi_{n_i^*}^{s,e*} \text{ if } \sum_n \max \left(\psi_{n_i^*}^{s,e*} \Delta R_n^s \tilde{\pi}_{ni}^s - \tilde{w}_n f_n^{s,e}, 0 \right) \geq \tilde{w}_i f_i^{s,o}, \\ &\sum_n \max \left(\psi_i^{s,o*} \Delta R_n^s \tilde{\pi}_{ni}^s - \tilde{w}_n f_n^{s,e}, 0 \right) = \tilde{w}_i f_i^{s,o} \text{ otherwise.} \end{aligned} \quad (57)$$

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 (53) holds with employment in patenting given by (51) and (52);
- $\tilde{\Psi}_{Mni}^{s,NP}$, $\tilde{\Psi}_{Mni}^{s,P,ND}$, $\tilde{\Psi}_{Mni}^{s,P,D}$ and $\tilde{\Psi}_{Cn}^s$ satisfy the laws of motion in (56) and $\tilde{\Psi}_{Mni}^s$ is given by the normalized version of (19);
- $\tilde{V}_{ni}^{s,NP}(1)$ and $\tilde{V}_{ni}^{s,P}(1)$ are given by the normalized versions of (26) and (27) with normalized profits obeying (28);
- \tilde{V}_{it}^s satisfies (30);
- ψ_{ni}^{s*} and $\psi_i^{s,o*}$ are defined by (17), (29) and (57);
- L_{Ri}^s is given by (31);
- g^s and g are given by (24) and (25), and;
- r satisfies the Euler equation (55) and the transversality condition $r > g$ 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. Patent applications covering the same invention are 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 with 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 as the ratio of the restricted sample flow into the US from origin i in year t_0 relative to the same ratio 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

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 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.

Turnover in US imports. The turnover measure captures 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, for any base year t , restrict the sample to 10-digit products for which the US was a net importer in both t and $t + 5$. We then aggregate countries to the regions used in our calibration and define the turnover rate as the import-weighted share of products for which there is a significant change in the origin of US imports between t and $t + 5$. We classify a product as experiencing a significant change in origin

if three conditions are met: (i) the leading country (in terms of US imports) changes; (ii) the initial leader has an import share at least 25 percentage points higher than any other country in year t , and; (iii) the new leader has an import share at least 25 percentage points higher than any other country in year $t + 5$. These conditions are chosen to identify products that experience a clear switch in the origin of imports. The average turnover rate calculated over all base years from 2000-2015 is 1.75 percent.

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 $f_{Europe}^j = 3.87f_{US}^j$, which implies $w_{Europe}f_{Europe}^j = 3.59w_{US}f_{US}^j$ when wages are proportional to observed real GDP per capita.

Other moments. The target 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. Population is the working age population aged 15-64. 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 PAT_{ni}^s 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 given by:

$$RD_i = \sum_{s=1}^S \left(w_i L_{Ri}^s + w_i L_i^{s,o} + \sum_{n=1}^N w_n L_{in}^{s,e} \right).$$

for $i = US$. Taking the ratio allows us to write all variables in their normalized forms.

We match the turnover moment to the model-implied ratio of the value of US imports that switch origin between t and $t + 5$ to total US imports of products imported in both t and $t + 5$, which we denote TO_n^s with $n = US$. 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 $\epsilon(x) \equiv 1 - e^{-x}$. Then a little calculation yields:

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

where $\Delta t = 5$.

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:

$$\frac{1}{N(N-1)} \sum_{n=1, n \neq i}^N \sum_{i=1}^N TE_{ni}^s.$$

The aggregate growth rate equals g . Using equation (54), the difference between price growth in the non-patenting and patenting sectors is:

$$g_{P^0} - g_{P^1} = (\alpha^0 - \alpha^1) g + \frac{g^1}{\sigma^1 - 1} - \frac{g^0}{\sigma^0 - 1}.$$

We calculate world gross output as $\sum_{i=1}^N P_i \tilde{Y}_i$. Price levels relative to the US are given by P_i/P_{US} . We define the nominal GDP of country i as:

$$GDP_i = P_i C_i + TB_i \sum_{n=1}^N P_n Y_n + w_i (L_i - L_{Yi}).$$

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 $\sum_{n=1}^N GDP_n/P_n$.

B.4 Steady State Solution Algorithm

Let $\tilde{Z}_i \equiv P_i \tilde{Y}_i$, $\tilde{B}_{ni}^s \equiv \tilde{\pi}_{ni}^s / \tilde{w}_i$ 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 $VF = \left(\tilde{w}_i, \tilde{Z}_i, L_{Ri}^s, \tilde{B}_{ni}^s, \tilde{\phi}_{ni}^s \right)$.

Given an initial guess for VF , we compute the auxiliary variables as follows. Profits: $\tilde{\pi}_{ni}^s = \tilde{w}_i \tilde{B}_{ni}^s$. Growth rates: equation (24) gives g^s , equation (25) gives g and equation (55) gives r . Patenting thresholds: equation (29) gives $\psi_{ni}^{s,e*}$, equation (57) gives $\psi_i^{s,o*}$ and equation (17) gives ψ_{ni}^{s*} . Aggregate qualities: equation (33) 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 (51) gives $L_{in}^{s,e}$, equation (52) gives $L_i^{s,o}$ and equation (53) gives L_{Yi} . Trade flows:

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

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

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

Final good price indices:

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

We then update the fundamental variables using:

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

$$\tilde{B}_{ni}^s = \frac{\tilde{X}_{Mni}^s}{\sigma^s \tilde{\Psi}_{Mni}^s \tilde{w}_i},$$

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

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

$$\tilde{Z}_i = \sum_{s=0}^S \sum_{n=1}^N \tilde{X}_{ni}^s + \sum_{s=1}^S \sum_{n=1}^N \tilde{w}_i L_{ni}^{s,e} - \left(T B_i \sum_{n=1}^N \tilde{Z}_n + \sum_{s=1}^S \sum_{n=1}^N \tilde{w}_n L_{in}^{s,e} \right), \quad (65)$$

and iterate to stabilize VF using a fixed-point iterative algorithm. Note that the algorithm takes the trade shares $\left(\tilde{X}_{ni}^s / \tilde{X}_{nn}^s \right)^{\text{Data}}$ directly from the data since the trade costs τ_{ni}^s are chosen to match these shares exactly. We use a type-I stationary Anderson accelerated process (Anderson 1965), following the implementation of Zhang et al. (2020). Convergence is ensured by enforcing a damping hyper-parameter (Evans et al. 2020). The damping hyper-parameter is weakly optimized, following a modified implementation of Chen and Vuik (2022). The key 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 $(\tilde{w}_i, \tilde{Z}_i, L_{Ri}^s, \tilde{B}_{ni}^s, \tilde{\phi}_{ni}^s) \rightarrow (\tilde{w}_i / P_1, \tilde{Z}_i / P_1, L_{Ri}^s, \tilde{B}_{ni}^s, \tilde{\phi}_{ni}^s P_1)$, 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:

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

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 1.

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 follows an implementation of 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 , τ_{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 , τ_{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 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).

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 (13) implies:

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

while differentiating (14) gives:

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

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 (19) 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 (16) gives $\psi_{nit}^{s,e*}$, equation (50) gives $\psi_{it}^{s,o*}$ and equation (17) gives ψ_{nit}^{s*} . Normalized version of equation (18) gives \tilde{V}_{it}^s . Equation (31) gives L_{Rit}^s . Equation (24) gives g_t^s and equation (25) gives g_t . Equation (51) gives $L_{int}^{s,e}$, equation (52) gives $L_{it}^{s,o}$ and equation (53) 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 . Equation (28) gives $\tilde{\pi}_{nit}^s$. Normalized version of equation (46) gives $P_{it} \tilde{C}_{it}$ and equation (48) 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 work with 25 Chebyshev collocation nodes. Using equations (20)-(23) we set:

$$\tilde{\Psi}_{Cnt}^s = \frac{1}{\zeta^s + g_t^s} \left[\sum_{i=1}^N \left(\nu^s \tilde{\Psi}_{Mnit}^{s,NP} + \delta_n^s \tilde{\Psi}_{Mnit}^{s,P,D} \right) - \dot{\psi}_{Cnt}^s \right],$$

$$\tilde{\Psi}_{Mnit}^{s,NP} = \frac{1}{\zeta^s + \nu^s + g_t^s} \left(\eta_i^s (L_{Rit}^s)^{1-\kappa} \frac{k}{k-1} \left[1 - (\psi_{nit}^{s*})^{1-k} \right] + \delta_n^s \tilde{\Psi}_{Mnit}^{s,P,ND} - \dot{\psi}_{Mnit}^{s,NP} \right),$$

$$\tilde{\Psi}_{Mnit}^{s,P,ND} = \frac{1}{\zeta^s + \nu^s + \delta_n^s + g_t^s} \left(\eta_i^s (L_{Rit}^s)^{1-\kappa} \frac{k}{k-1} \left[1 - (\psi_{nit}^{s*})^{1-k} \right] - \dot{\psi}_{Mnit}^{s,P,ND} \right),$$

$$\tilde{\Psi}_{Mnit}^{s,P,D} = \frac{1}{\zeta^s + \delta_n^s + g_t^s} \left(\nu^s \tilde{\Psi}_{Mnit}^{s,P,ND} - \dot{\psi}_{Mnit}^{s,P,D} \right).$$

From equations (67)-(69) we set:

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

$$\tilde{V}_{nit}^{s,NP}(1) = \frac{1}{r_{it} + \zeta^s + \nu^s + g_t^s - g_t} \left[\tilde{\pi}_{nit}^s + \dot{V}_{nit}^{s,NP}(1) \right],$$

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

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 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 is 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.3 Additional Counterfactual Results


Figure 7 shows how variation in trade costs affects the derivative of each country's welfare with respect to its own patent protection, where the derivative is computed based on the welfare change $EV_i - 1$. The figure plots the derivative for each country i as a function of trade costs in the patenting sector. It shows that the welfare gains from unilaterally weakening patent protection (by increasing δ_i) tend to be smaller when trade costs are greater.



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 trade costs in the patenting sector. Welfare changes are expressed as the equivalent variation in consumption and account for transition dynamics.

Figure 7: Welfare effects and trade costs

While Figure 7 focuses on the welfare effects of small changes in patent protection, Figure 8 reports the impact of larger changes similar to those analyzed in Figure 5. Panel (a) plots changes in US welfare EV_{US} against δ_{US} for the baseline calibration and two variants: an economy where we double the rate of technology diffusion ν , and; an economy where we double international trade costs in the patenting sector. The figure shows that, in contrast to the baseline calibration, US welfare is approximately flat in δ_{US} when either the trade costs or the technology diffusion rate are doubled. However, for all countries other than the US, stronger patent protection reduces domestic welfare in all three cases. Panel (b) illustrates this observation by plotting how European welfare depends upon the strength of European patent protection in each of the calibrations. In line with Figure 7, Europe gains less from reducing patent protection when trade costs are higher (and the same pattern holds for all other countries outside the US).



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: an economy where the rate of technology diffusion ν is doubled, and; an economy where all international trade costs in the patenting sector are 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 8: Unilateral patent policy and welfare: doubling technology diffusion rate and trade costs

Table 5 summarizes the Nash equilibria and the equal weights cooperative equilibria for the economies with doubled trade costs and doubled technology diffusion rate. We report whether any countries change their optimal patent protection compared to their choices in the baseline calibration and how moving to the Nash and cooperative equilibria affects US and equal weights world welfare. The only country that changes its baseline policy in the Nash equilibria is the US, which provides complete patent protection in both cases. And since the US provides complete patent protection, growth increases, which raises world welfare. The cooperative equilibria are similar to the baseline cooperative equilibrium, except that doubling the technology diffusion rate increases the welfare gains from cooperation and leads to Russia providing complete patent protection.

D Sensitivity and Robustness Checks

Figure 9 plots the sensitivity of the baseline simulated method of moments parameter estimates shown in Table 2 to variation in the targeted moments and in the externally calibrated parameters (θ^0 , γ , κ , ρ and σ^1). The columns are the estimated parameters and the rows are the targeted moments and the externally calibrated parameters. Each row reports the elasticities of the estimated parameters to variation in one of the targeted moments or externally calibrated parameters (starting from the baseline 2015 calibration). For example, the first row shows the sensitivity of the estimated parameters to varying θ^0 . Darker reds indicate more positive elasticities and darker

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

Calibration	(a) Nash			(b) Cooperative: equal weights		
	Changes in patent protection from baseline	Welfare change (percent)		Changes in patent protection from baseline	Welfare change (percent)	
		US	World		US	World
Baseline	n/a	-0.5	-2.4	n/a	2.1	8.8
Double technology diffusion rate, ν	US complete	-1.9	5.7	Russia complete	5.0	14.3
Double international trade costs	US complete	-0.5	4.0	No changes	3.1	8.7


Notes: This table reports results for two variants of the baseline 2015 calibration: an economy where the rate of technology diffusion ν is doubled, and; an economy where all international trade costs in the patenting sector are doubled. The 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. 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. Complete protection corresponds to $\delta_i = 0.01$, which is the lower bound we impose on our solutions.

blues more negative elasticities. To simplify the presentation, we vary all 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.

We highlight two patterns in Figure 9. First, as expected, the R&D expenditure ratios and international patent shares play an important role in calibrating patent protection δ_i and R&D efficiency η_i . To better visualize these relationships, Figure 10 zooms in on Figure 9 and shows only the columns for the δ_i and η_i parameters and the rows for the R&D expenditure and international patent share moments.

From Figure 10 we see that, except for the US, δ_i is strongly decreasing in the share of international patent flows into country i . 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 9 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 from Figure 10 that η_i is strongly increasing in the ratio of R&D expenditure to GDP for developed countries and in the share of international patent flows originating in country i for developing countries.

Second, Figure 9 shows that our estimates are relatively sensitive to three of the externally calibrated parameters (κ , ρ and σ^1) and five of the targeted moments (the value of patents relative




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 defined as $(T_i^0)^{\frac{1}{\theta^0}}$.

Figure 9: Sensitivity analysis: full

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.²⁴ In each case, we recalibrate the model following the same procedure as in Section 3.2 and then repeat the unilateral, Nash, cooperative and TRIPS counterfactuals from

²⁴When 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.



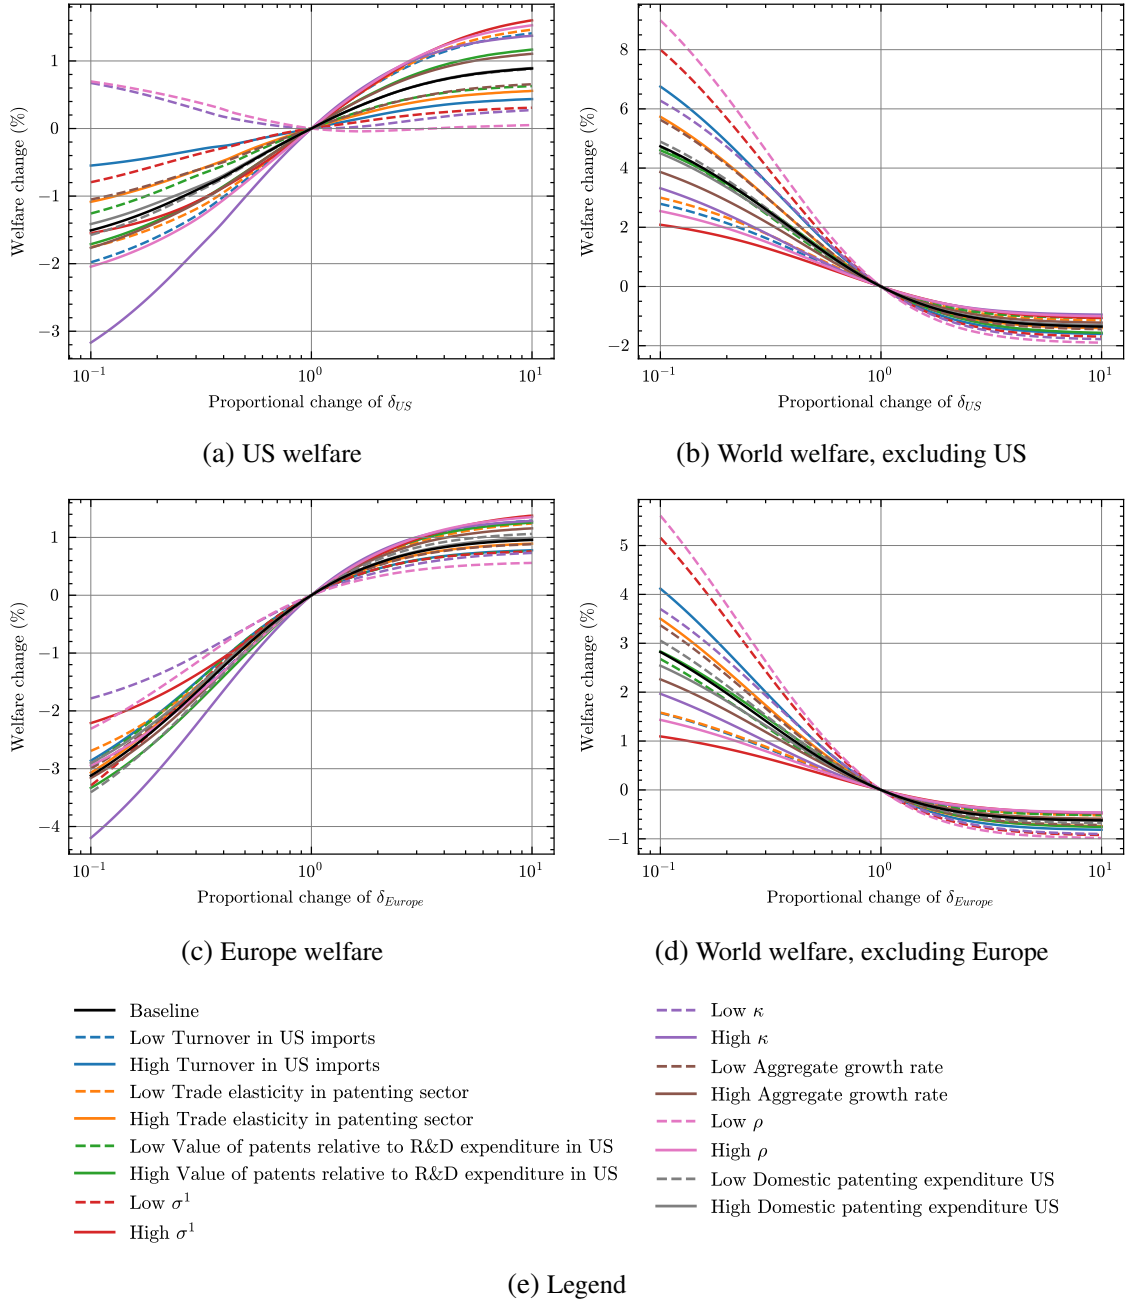
Notes: This figure shows elasticities of simulated method of moments estimates (in columns) to targeted moments (in rows). Colormap range is bounded at the 2nd and 98th percentiles of the elasticity distribution. For the international patenting moments, we vary all patent flows into (or out of) a given country by the same proportion.

Figure 10: Sensitivity analysis: patenting and innovation

Sections 4.1–4.4. The results of these robustness checks are summarized in Figure 11 and Table 6. In the discussion of the robustness checks below, we always use equal weights to aggregate welfare changes across countries

Figure 11 shows selected results for unilateral patent policy counterfactuals similar to those analyzed in Section 4.1. Panel (a) plots changes in US welfare EV_{US} against US patent protection δ_{US} for the baseline calibration and each of the sixteen robustness checks (the legend is in panel e). 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 protection in the US raises welfare abroad by increasing the growth rate. Increasing the strength of patent protection in countries other than the US reduces domestic welfare and raises foreign welfare in the baseline and in all of the robustness checks. Panels (c) and (d) illustrate this observation by showing charts analogous to panels (a) and (b), but for variation in δ_{Europe} . In all scenarios, stronger patent protection in Europe reduces European welfare, but benefits the world excluding Europe.

Table 6 summarizes the results of the robustness checks for the Nash, cooperative and TRIPS



Notes: This figure plots the effect of proportional changes in calibrated patent protection on welfare for different calibrations of the model. Plots vary δ_{US} in panels (a) and (b), and δ_{Europe} in panels (c) and (d). 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 11: Robustness checks for unilateral patent policy counterfactuals

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 11, the US chooses complete patent protection in both these cases. Because the US provides complete protection, growth increases and welfare for countries other than the US, and for the world as a whole, is higher in the Nash equilibrium than in the calibrated steady state. However, moving to Nash still reduces US welfare in these cases because the US bears the static costs of stronger protection. In all other scenarios moving to the Nash equilibrium reduces world welfare.

For the equal weights cooperative equilibrium, China and Russia switch to complete patent protection in some of the robustness checks, but no other countries deviate. 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, the demand elasticity σ^1 is low (implying higher mark-ups), or turnover in US imports is high (implying faster technology diffusion). For the Negishi weights cooperative equilibrium, the only country that ever deviates from its baseline policy is Mexico, which provides complete protection in the low κ and low ρ scenarios. However, variation in welfare gains across scenarios is similar for the Negishi weights and equal weights equilibria.

For TRIPS, we start by recalibrating the model for each robustness check using 1992 data as in Section 4.4. Then we simulate a return to pre-TRIPS patent 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 world as a whole. 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 countries differs less between periods. Consequently, in all cases, simulating a return to pre-TRIPS patent 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 results demonstrate the robustness of our baseline findings for the TRIPS counterfactual.

Table 6: Robustness checks

Robustness check	(a) Nash		(b) Cooperative: equal weights		(c) Pre-TRIPS		
	Changes in patent protection from baseline	World welfare change (percent)	Changes in patent protection from baseline	World welfare change (percent)	Developing	Developed	World
Baseline	n/a	-2.4	n/a	8.8	0.30	-0.18	0.03
Low Turnover in US imports	No changes	-1.7	No changes	4.7	0.32	-0.11	0.08
High Turnover in US imports	No changes	-3.0	No changes	13.2	0.23	-0.21	-0.02
Low Trade elasticity in patenting sector	No changes	-1.9	No changes	5.0	0.29	-0.09	0.07
High Trade elasticity in patenting sector	No changes	-2.7	No changes	10.9	0.25	-0.20	0.00
Low Value patents to R&D expenditure in US	No changes	-1.9	No changes	8.9	0.20	-0.11	0.03
High Value patents to R&D expenditure in US	No changes	-2.9	No changes	8.4	0.35	-0.20	0.03
Low σ^1	No changes	-3.4	No changes	16.0	0.24	-0.34	-0.09
High σ^1	No changes	-1.8	No changes	3.3	0.30	-0.08	0.09
Low κ	US complete	4.4	No changes	11.2	0.22	-0.22	-0.03
High κ	No changes	-1.6	No changes	6.3	0.34	-0.12	0.08
Low Aggregate growth rate	No changes	-2.7	China, Russia complete	10.6	0.25	-0.18	0.01
High Aggregate growth rate	No changes	-2.2	Russia complete	7.1	0.31	-0.15	0.05
Low ρ	US complete	7.3	China, Russia complete	17.6	0.18	-0.26	-0.07
High ρ	No changes	-1.7	Russia complete	4.7	0.33	-0.11	0.08
Low Domestic patenting expenditure in US	No changes	-2.5	China, Russia complete	9.5	0.27	-0.15	0.03
High Domestic patenting expenditure in US	No changes	-2.3	China, Russia complete	8.1	0.28	-0.16	0.03

Notes: This table reports counterfactual results for different calibrations of the model. The 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 for the world, for developing countries (China, Brazil, India, Russia and Mexico), and for developed countries (US, Europe, Japan, Canada and Korea). 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 protection level. Complete protection corresponds to $\delta_i = 0.01$, which is the lower bound we impose on our solutions.

References

- Anderson, Donald. 1965. "Iterative Procedures for Nonlinear Integral Equations." *Journal of the ACM (JACM)* 12(4): 547-560.
- 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).
- 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.
- Evans, Claire, Sara Pollock, Leo G. Rebholz, and Mengying Xiao. 2020. "A Proof that Anderson Acceleration Improves the Convergence Rate in Linearly Converging Fixed-point Methods (But Not in those Converging Quadratically)." *SIAM Journal on Numerical Analysis* 58(1): 788-810.
- 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. Springer Berlin Heidelberg.
- 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

Global Edition.

Schott, Peter. 2008. “The Relative Sophistication of Chinese Exports.” *Economic Policy* 23(53): 6-49.

World Intellectual Property Organization (WIPO). 2023. “WIPO Statistics Database.” Last updated: February 2023.

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.