Taxation and the Global Allocation of Intangibles

Jesse LaBelle, Fernando M. Martin and Ana María Santacreu*

May 21, 2026

Abstract

We study how international tax systems shape the global allocation of intangible assets. Firm-level data show that cross-border patent ownership transfers respond strongly to tax differentials, especially within multinational firms, while arm's-length transfers are more closely associated with intellectual property rights (IPR) protection. We develop a model in which patent owners choose whether to license, sell, or transfer patents to a foreign affiliate, linking tax rates, IPR protection and transfer-pricing wedges to patent location. Counterfactuals suggest that tax harmonization and stronger enforcement reduce profit shifting and expand the US royalty tax base, whereas a global minimum tax has limited effects.

JEL Codes: F12, O33, O41, O47.

Keywords: Intangibles, cross-border patent sales, licensing, profit shifting, intellectual property rights, taxation, tax harmonization, global minimum tax.

*LaBelle, Northwestern University, email: jesselabelle2029@u.northwestern.edu; Martin, Federal Reserve Bank of St. Louis, email: fernando.m.martin@stls.frb.org; Santacreu, Federal Reserve Bank of St. Louis, email: am.santacreu@gmail.com. The views expressed in this paper do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

1 Introduction

Intangible assets have become a central input in production, innovation and international competitiveness. Unlike physical capital, intellectual property is non-rival, scalable and can be separated from the location of invention, production or final sales. As a result, the global allocation of intangibles is sensitive to economic fundamentals, intellectual property rights (IPR) enforcement and international tax differences. A growing literature shows that multinational firms respond to cross-country tax differentials by relocating intangible assets to low-tax jurisdictions.[1] Yet the mechanisms through which multinational firms reallocate intangible assets across jurisdictions remain less understood.

We study how international tax systems shape the global allocation of intangibles by focusing on patents, whose ownership can be tracked across countries and firms. Our approach distinguishes arm's-length transactions between unrelated firms from transfers within multinational groups, allowing us to separate market-based reallocations from intra-firm movements that may be more closely tied to profit shifting, namely the relocation of intangible ownership to lower-tax jurisdictions in order to reduce global tax liabilities.

The analysis proceeds in three steps. First, we construct a firm-level dataset of cross-border patent ownership changes from 1995 to 2020, distinguishing arm's-length transactions between unrelated firms from intra-firm reallocations within multinational groups. Second, we develop a model in which patent owners choose whether to license a technology, sell it to an unrelated buyer, or transfer it to a foreign affiliate. The model links these organizational choices to country fundamentals, tax differentials and transfer-pricing opportunities, thereby determining the location of patent ownership and royalty income. Third, we calibrate the model to evaluate how alternative tax policies—including tax harmonization, tighter transfer-pricing enforcement, higher relocation costs and global minimum taxes—affect profit shifting and the location of intangible income.

The empirical analysis documents four facts. First, patent transfers from high-tax to low-tax jurisdictions increased sharply during the 2000s, with tax havens absorbing a disproportionate share of inflows despite their limited infrastructure and production capability. Second, patent transfers respond strongly to tax differentials, but this response is concentrated within multinational firms: intra-firm transfers are substantially more sensitive to tax differences than

[1] For example, see Santacreu (2026), Tørsløv, Wier, and Zucman (2023) and Guvenen, Mataloni, Rassier, and Ruhl (2022).

arm's-length transfers. Third, stronger intellectual property rights attract inter-firm inflows, while tax havens receive substantial intra-firm inflows despite weaker protection. Fourth, the geography of patent transfers closely mirrors the global distribution of royalty income, suggesting that patent ownership relocation is an important channel through which multinational firms shift intangible income across jurisdictions.

These facts reveal trade-offs in the decision of where to locate patent ownership. If patent location were driven primarily by fundamentals such as IPR protection and local market conditions, ownership would be concentrated in countries where patents generate the highest licensing value. Instead, a substantial share of patents is located in jurisdictions with weak fundamentals but favorable tax treatment. At the same time, the strong response of intra-firm transfers to tax differentials suggests that an important adjustment margin operates within multinational firms, where ownership can be rearranged and transfer prices may deviate from arm's-length values.

To organize these findings, we develop a model that connects country-specific fundamentals and tax incentives to patent location and organizational choice. In the model, a patent owner can license the technology globally, sell the patent to an unrelated foreign owner, or establish a foreign affiliate and transfer the patent within the firm. When non-tax fundamentals, such as productivity and IPR protection, are the main consideration, patents locate where they generate the highest licensing value, consistent with the positive relationship between IPR strength and inter-firm patent inflows documented in the data. Tax regimes create additional motives for relocation. A low domestic capital-gains tax rate makes arm's-length sales more attractive, particularly when sellers capture a larger share of the transaction surplus. For intra-firm transfers, low destination income tax rates help multinationals overcome setup costs and other frictions, while deviations from arm's-length transfer prices allow firms to further exploit tax differentials across jurisdictions. These mechanisms can rationalize why tax havens attract substantial patent inflows despite weak fundamentals and why tax sensitivity is concentrated among intra-firm transfers.

We calibrate the model to match US-origin patent transfers and royalty receipts across low-tax countries, tax havens and other destinations. The calibrated framework reproduces both the selective relocation of patents and the concentration of royalty income in low-tax jurisdictions, which allows us to evaluate policy counterfactuals. Policies that target the underlying tax wedge or transfer-pricing margin have large effects. Harmonizing foreign royalty tax rates to the US level removes tax differentials, substantially reduces profit shifting and increases the US royalty tax base by more than 20 percent. In the baseline calibration, this corresponds to roughly $3 billion in additional US royalty tax revenues per year, or about $30 billion over a decade. Policies that increase the cost of relocating intangible assets or tighten arm's-length transfer-pricing enforcement have similar effects. By contrast, a global minimum tax has limited effects because it leaves substantial tax differentials in place and therefore does little to alter firms' location incentives. In this environment, reducing the incentive to relocate intangible assets matters more than imposing a floor on where profits can be booked.

Related literature This paper relates to a broad literature on the international allocation of intangible assets and multinational profit shifting. A growing body of work shows that multinational firms relocate intangible assets toward low-tax jurisdictions and that patent ownership responds to international tax incentives, particularly within multinational groups and under preferential tax regimes (Crouzet, Eberly, Eisfeldt, and Papanikolaou, 2022; Karkinsky and Riedel, 2012; Becker and Riedel, 2012; Santacreu, 2024; Ciaramella, 2023). Our empirical analysis contributes by distinguishing arm's-length from intra-firm patent transfers, showing that tax sensitivity is concentrated within multinational firms, whereas inter-firm transfers are more closely associated with intellectual property rights protection. We further connect these ownership reallocations to the global allocation of royalty income.

Our analysis is also connected to the broader literature on profit shifting and international tax policy. Existing empirical work shows that multinational firms reallocate substantial profits toward low-tax jurisdictions, distorting both taxable income and measured macroeconomic aggregates (Clausing, 2009, 2016; Tørsløv, Wier, and Zucman, 2023; Guvenen, Mataloni, Rassier, and Ruhl, 2022). In parallel, recent quantitative frameworks study how intangible capital and profit shifting shape the effects of international tax policy in general equilibrium (Dyrda, Hong, and Steinberg, 2024a,b; Dyrda et al., 2025). We complement these approaches by linking profit shifting to observable reallocations of patent ownership and royalty income across jurisdictions.

The paper also relates to work on patent boxes, licensing and the institutional organization of intellectual property. Preferential tax regimes have been shown to attract patent relocations, especially when firms can benefit from lower tax rates without relocating the underlying innovative activity (Gaessler, Hall, and Harhoff, 2021), while international licensing responds jointly to tax incentives and intellectual property rights protection (Santacreu, 2026). Consistent with this evidence, we find that inter-firm transfers are more closely associated with IPR protection, whereas intra-firm transfers are primarily tax-driven. Related evidence further suggests that recent international tax reforms shifted firms toward less observable avoidance margins (Santacreu and Stewart, 2025). Finally, Fosfuri, Helmers, and Roux (2012) show that patent ownership is typically concentrated within corporate groups rather than shared across unrelated firms, a feature that facilitates the intra-firm reallocation of intangible assets across jurisdictions.

Our counterfactual results also support the broader view that policies reducing effective tax differentials and limiting avoidance margins are more effective than policies that simply impose minimum tax rates, consistent with Devereux (2022).

2 Empirical Analysis

We begin by constructing a dataset that tracks global patent ownership and transfers. We then document key empirical patterns in international patent transactions and their relationship with cross-country tax differences. Finally, we estimate a gravity model to quantify how tax incentives shape firms' decisions about where to locate intellectual property.

2.1 A Dataset on the Location of Intangibles

We construct a firm-level cross-country dataset of patent ownership and transactions over 1995–2020 by combining PATSTAT Global with ktMINE transaction data. This dataset allows us to track the ownership history of individual patents and to distinguish between inter-firm and intra-firm transfers, capturing both market transactions and relocations within multinational firms.

We focus on international patent transfers, which occur either between unrelated firms (inter-firm) or within multinational groups (intra-firm). This distinction is central: intra-firm transfers are more likely to reflect tax-motivated reallocations, whereas inter-firm transfers more closely reflect market-based transactions.

We augment the data with information on statutory corporate income tax rates, patent box regimes, and measures of IPR protection in both origin and destination countries, as well as standard gravity controls. These variables allow us to study how tax and institutional differences shape the location of patent ownership.

The resulting dataset is an unbalanced panel covering 54 countries, more than 50,000 firms, and a total of 4,835,765 observations, capturing the connections between country and firm, at a given time. Patent transfers are rare events, accounting for 0.4% of observations. We classify destinations into non-haven, low-tax, and tax haven countries following Tørsløv, Wier, and Zucman (2023). While most transactions occur between non-haven countries, intra-firm transfers are disproportionately concentrated in tax havens.

Additional details on data construction, ownership tracking, and the identification of intra-firm relationships are provided in the Appendix.

2.2 Stylized Facts

We begin by documenting a set of empirical patterns that help guide the analysis. While these patterns are descriptive, they suggest a systematic relationship between tax incentives, institutional characteristics, and the location of patent ownership.

Dynamics of Patent Transfers We first examine how the destination of patent transfers has evolved over time. Figure 1 shows the evolution of patent sales from non-tax havens to different destination groups between 1995 and 2012, relative to 1995. Transfers to tax havens increase markedly over this period, particularly in the 2000s, while transfers to low-tax countries also rise, albeit more gradually. In contrast, transfers to other non-haven destinations remain comparatively stable.

These patterns are consistent with a gradual reallocation of patent ownership toward lower-tax jurisdictions. While the underlying innovation activity may remain concentrated in certain countries, the location of ownership appears to respond to additional factors beyond the place of invention.

Tax Differentials and IPR We next examine the cross-sectional relationship between patent transfers, tax differences, and the quality of IPR. Figure 2 presents two simple correlations.

Panel (a) shows a positive association between bilateral tax differences and patent transfers: larger tax gaps between seller and buyer countries are correlated with greater transfers toward lower-tax jurisdictions. Panel (b) shows that countries with stronger IPR protection tend to attract more patent inflows, consistent with the role of institutions in supporting the value of intellectual property.

These patterns suggest that both tax incentives and institutional quality are associated with cross-border patent transfers. At the same time, the presence of tax havens—which attract

Figure 1: Evolution of Patent Sales from Non-Tax Havens

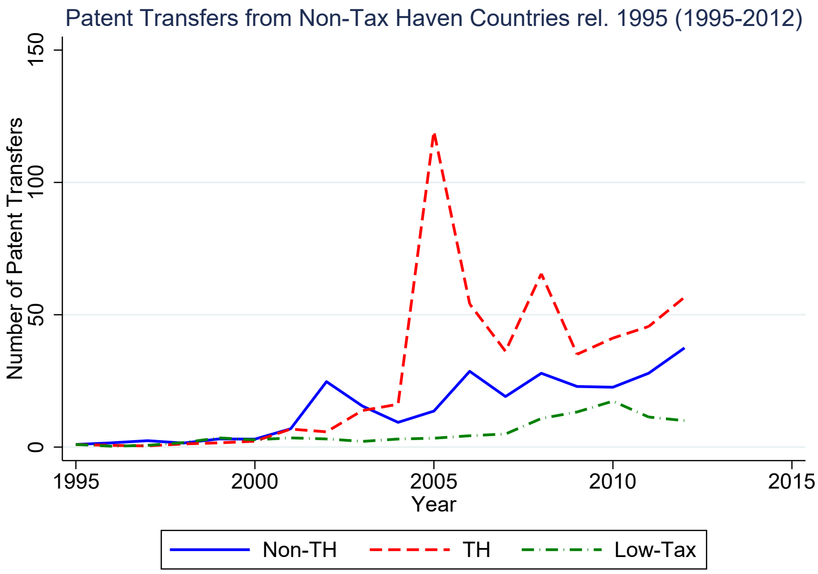

Note: The figure shows the evolution of patent sales from non-tax havens to non-tax havens (solid line), low-tax countries (dash-dotted line), and tax havens (dashed line) between 1995 and 2012 (relative to 1995).

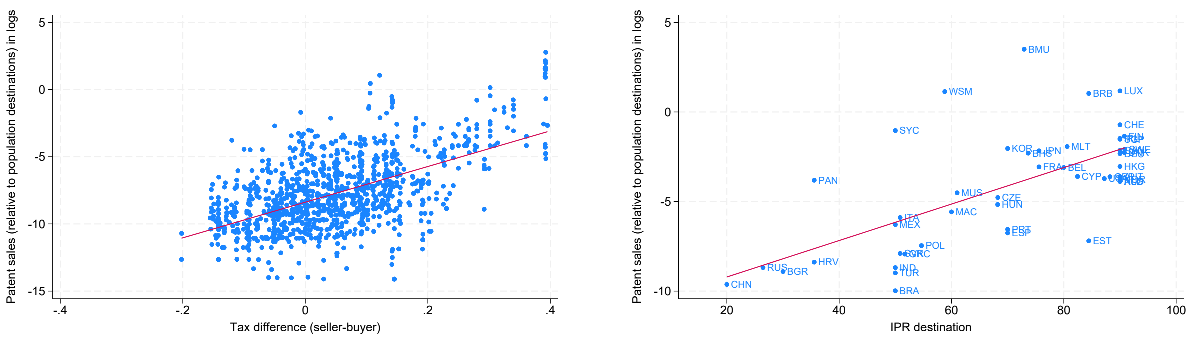

(a) Patent transfers and tax differences

(b) Patent purchases and IPR

Figure 2: Determinants of cross-border patent transfers

Note: Panel (a) plots patent sales against bilateral tax differences. Panel (b) plots patent purchases against IPR. Correlations are approximately 0.54 and 0.42, respectively. Tax havens appear as outliers with relatively low IPR but large inflows.

substantial inflows despite weaker IPR protection—highlights that tax considerations can, in some cases, outweigh institutional differences.

While these correlations are informative, they do not quantify the strength of these relationships or disentangle the underlying mechanisms. In particular, the descriptive evidence does not allow us to separate the role of tax incentives from other confounding factors, nor to distinguish between transactions occurring within firms and those taking place across firms.

2.3 Empirical Analysis

To address these issues, we now turn to a formal econometric analysis that quantifies the relationship between tax differentials and patent transfers, while controlling for a rich set of fixed effects and separating intra- and inter-firm transactions. Specifically, we estimate a Poisson Pseudo–Maximum Likelihood (PPML) gravity model with rich fixed effects (Yotov, Piermartini, Monteiro, and Larch, 2016; Weidner and Zylkin, 2021). Throughout, we interpret these estimates as conditional relationships reflecting equilibrium sorting across locations, rather than causal effects of specific tax policies.

We specify the following estimating equation:

$$P_{inj,t} = \exp\left[\beta_0 + \beta_1(\tau_n - \tau_i) + \mu_{nt} + \mu_{it} + \mu_j + \mu_t + \mu_n + \mu_i + \mu_{in}\right] \varepsilon_{inj,t}, \qquad (1)$$

where $P_{inj,t}$ represents the number of patents sold from country n to country i by firm j at time t. The term $(\tau_n - \tau_i)$ captures the difference in tax rates between origin and destination countries. The terms μ_{nt} and μ_{it} include controls that capture time-varying country characteristics (such as GDP and GDP per capita), while μ_j and μ_t control for firm-specific and aggregate time shocks. Finally, μ_n, μ_i, and μ_{in} capture origin, destination, and bilateral pair effects, respectively. The error term $\varepsilon_{inj,t}$ captures residual variation in patent transfers.

Given the count nature of patent transfers, we estimate the model using PPML, which accommodates zero flows and is robust to heteroskedasticity in gravity-type settings.[2] This specification is not designed to identify the causal effect of a particular tax reform. Rather, it provides a reduced-form characterization of how bilateral patent transfers co-vary with tax differences, conditional on a rich set of fixed effects that absorb country characteristics, bilateral relationships, and common shocks. The estimated coefficients should therefore be interpreted

[2] See Santos Silva and Tenreyro (2006) and Santos Silva and Tenreyro (2021).

as conditional elasticities that reflect equilibrium responses across firms and margins.

2.3.1 Baseline: Tax Differentials and Patent Transfers

Table 1 reports estimates of how statutory tax differentials predict patent transfers across three periods: 1995–2006, 2007–2012, and 2013–2020. Across specifications, the response is substantially stronger for intra-firm transfers than for inter-firm transfers, consistent with multinational firms using internal reallocations to shift intangible ownership across jurisdictions.

During 1995–2006, the relationship between tax differentials and patent transfers is positive but relatively modest. The coefficient for intra-firm transfers is 7.7 and statistically significant, implying that a 10 percentage-point increase in the tax advantage of the destination country is associated with roughly a 116 percent increase in intra-firm patent transfers. In contrast, the coefficient for inter-firm transfers is much smaller and statistically insignificant, suggesting that arm's-length transactions were less responsive to tax considerations during this period.

The response becomes substantially stronger during 2007–2012. The coefficient on intra-firm transfers rises to 14.8, implying that a 10 percentage-point tax differential is associated with approximately a 338 percent increase in intra-firm patent reallocations. Even the coefficient for total transfers nearly doubles relative to the earlier period. These magnitudes indicate that patent ownership became very responsive to international tax differences during the peak years of multinational profit shifting.

After 2013, the estimates become substantially less precise and are no longer statistically significant for total or inter-firm transfers, although the point estimate for intra-firm transfers remains large. This pattern is consistent with changes in the global tax environment—including the OECD's Base Erosion and Profit Shifting (BEPS) initiative, a set of reforms aimed at limiting profit shifting by multinational firms—which may have altered the channels through which firms respond to tax incentives. Rather than eliminating tax-motivated reallocations altogether, the evidence suggests that firms increasingly shifted toward less observable or more complex margins of adjustment.

2.3.2 Destination Attractiveness and IPR

While the previous results emphasize the role of tax differentials, countries also differ in their institutional environments. We examine how IPR protection relates to the location of patent ownership.

Table 1: Patent Transfers and Statutory Tax Rate Differentials: Evidence Across Three Periods

	$Patents_{injt}$								
	1995-2006			2007-2012			2013-2020		
	Total	Intra	Inter	Total	Intra	Inter	Total	Intra	Inter
$\tau_o - \tau_d$	4.587	7.703*	1.910	8.479**	14.77**	3.252	0.621	14.12	1.579
	(2.817)	(3.686)	(2.614)	(3.277)	(4.811)	(3.999)	(7.522)	(10.52)	(9.157)
N	279,931	150,182	186,999	343,777	180,366	249,128	211,661	93,535	149,241
pseudo R^2	0.569	0.582	0.550	0.505	0.508	0.503	0.561	0.613	0.556
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Origin FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Destination FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Pair FE	YES	YES	YES	YES	YES	YES	YES	YES	YES

Standard errors in parentheses clustered by firm
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Table reports PPML estimates. The dependent variable is the count of patent transfers from origin country i to destination country j in year t. $\tau_o - \tau_d$ represents the difference between statutory corporate tax rates in origin and destination countries. All specifications include a full set of fixed effects and controls for origin and destination GDP and GDP per capita. Standard errors are clustered at the firm level.

Figure 3 plots destination fixed effects against IPR strength. A clear pattern emerges. Among non-tax havens, stronger IPR protection is associated with higher destination attractiveness, consistent with the idea that secure property rights support market-based transactions. In contrast, tax havens attract substantial patent inflows despite weaker or less relevant variation in IPR, suggesting that other factors—most notably tax considerations—play a larger role in these jurisdictions.

Figure 3: Relationship between Destination Fixed Effects and IPR

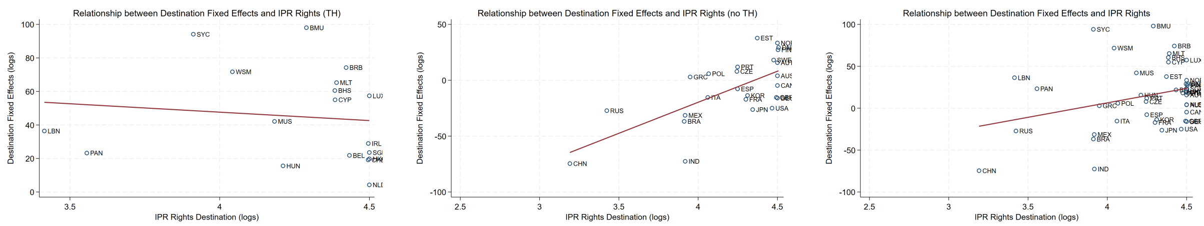

(a) Tax Havens (b) Non-Tax Havens (c) All Countries

Notes: Charts plot destination fixed effects against the IPR index (in logs). Both measures are averaged over the sample period.

These results suggest that the determinants of patent location differ across margins: tax differentials primarily shape intra-firm reallocations, whereas IPR protection is more closely associated with inter-firm (market-based) transactions.

We analyze patent box regimes—preferential tax schemes for intellectual property income that can influence firms' location of patent ownership—in the Appendix, where we find similar patterns to our baseline results.

2.4 From Patent Transfers to Royalty Income in Tax Havens

The previous analysis shows that patent ownership responds systematically to international tax differences. We next ask whether these reallocations of ownership translate into corresponding reallocations of income.

This step is important because the location of ownership of an intangible asset determines where the associated royalty income is recorded and taxed. When a multinational firm transfers ownership of a patent to an affiliate in a lower-tax jurisdiction, future royalty payments generated by that patent are also more likely to accrue in that jurisdiction. As a result, cross-border patent transfers provide a mechanism through which multinational firms can shift taxable income across countries without relocating real production or innovation activity.

Using bilateral royalty data, Figure 4 shows that tax havens and low-tax countries account for a disproportionate share of global royalty receipts. Despite their small economic size, tax havens capture a non-negligible fraction of worldwide royalties, while low-tax countries account for roughly a quarter to a third of the total.

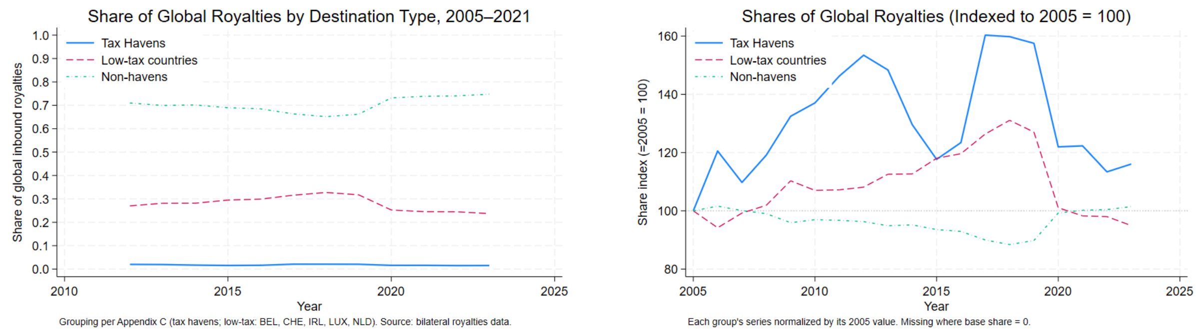

(a) Share of global royalties by destination type, 2012–2023.

(b) Shares of global royalties, indexed to 2005 = 100.

Figure 4: Global distribution of royalty receipts by destination type.

Over time, the importance of tax havens rises sharply during the 2000s and declines somewhat after 2015, consistent with changes in international tax enforcement (e.g., OECD BEPS and the US TCJA). Low-tax countries follow a similar, though more moderate, pattern.

Figure 5 further highlights the magnitude of these flows: in several small jurisdictions, royalty income represents a large share of GDP, in stark contrast to major economies. Barbados and Antigua and Barbuda report royalty receipts close to 40 and 30 percent of GDP, respectively, while Luxembourg, Cyprus, Ireland, and Malta also exhibit unusually large royalty-to-GDP ratios.

These magnitudes are difficult to reconcile with the local scale of innovative activity, domestic market size, or production capacity in these economies. Instead, they suggest that these jurisdictions function as locations for the legal ownership of intellectual property and the booking of royalty income generated elsewhere. In other words, the geographic location of patent ownership and the location of underlying innovation activity are often disconnected.

The presence of advanced economies such as the Netherlands, Ireland, and Switzerland in the top rankings also highlights that the concentration of royalty income is not limited to traditional offshore tax havens. More broadly, the figure illustrates how a relatively small set of low-tax jurisdictions captures a disproportionate share of global intangible income.

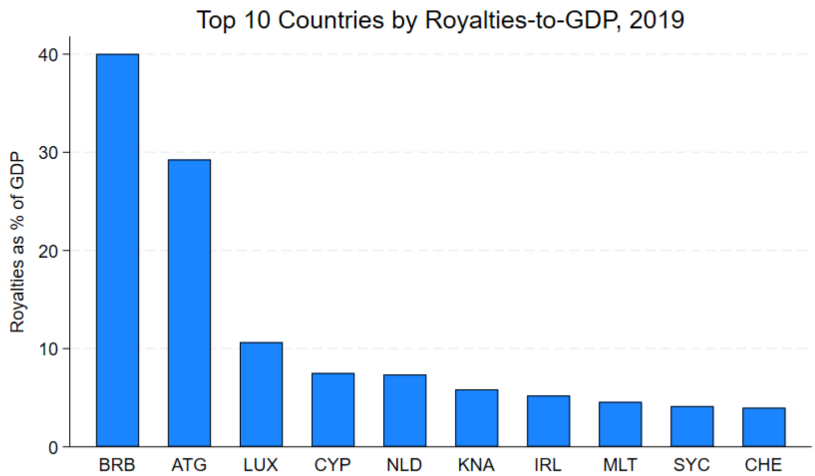

Figure 5: Top 10 countries by royalties-to-GDP, 2019.

These patterns are consistent with the idea that tax-motivated patent transfers are reflected in the global allocation of royalty income. Details on data construction are provided in the Appendix.

3 Theory

This section develops a framework to study how taxes and country characteristics shape the global allocation of patent ownership. The model links firms' organizational choices—licensing, arm's-length sales, and intra-firm transfers—to the location of royalty income and profit shifting incentives.

Patents represent valuable technologies that generate profits through licensing.[3] When a firm owns a patent, it can license this technology to producers across multiple countries simultaneously, reflecting the non-rival nature of intellectual property. Each license generates a cash flow that depends on the productivity of the technology and local market conditions. The patent owner must then decide whether to retain ownership and collect licensing revenues across countries, sell the patent to an unrelated foreign entity, or transfer ownership within the multinational firm.

This decision depends on the present value of licensing revenues under different tax regimes, interest rates, and risks of patent expiration or imitation. A central consideration is profit shifting, through which a multinational firm transfers patent ownership to a subsidiary in another jurisdiction at a potentially manipulated transfer price. This strategy allows firms to maximize global after-tax profits by exploiting differences between capital-gains taxes in the origin country and patent-income taxes in the destination country. As a result, the transfer price relative to arm's-length valuations becomes a key decision variable.

The framework formalizes how country-specific characteristics—particularly tax differences and transfer-pricing opportunities—shape both the location of patent ownership and the allocation of royalty income across jurisdictions.

3.1 Environment

There are $i = 1, \ldots, M$ countries, where $M > 1$. Each country i is characterized by a set of parameters $\{\tau_i, \tau_i^c, \tau_i^k, r_i, \phi_i\}$, where $\tau_i \in [0,1]$ is the patent income tax rate, $\tau_i^c \in [0,1]$ is the corporate income tax rate, $\tau_i^k \in [0,1]$ is the capital-gains tax rate, r_i is the interest rate and $\phi_i \in [0,1]$ is the probability that a patent expires or is imitated.

A patent owned in country i and licensed in country j generates a pre-tax cash flow of $y_{ij} \geq 0$ for the licensee-producer. Licenses last for one period and cost a negotiated fee ℓ_{ij} per period. Thus, a patent owned in country i and licensed in country j yields the patent owner after-tax income $(1-\tau_i)\ell_{ij}$ and the licensee-producer, $(1-\tau_j^c)(y_{ij} - \ell_{ij})$.

The licensing fee ℓ_{ij} is negotiated every period between patent owner and producer. Assuming the patent owner has bargaining power $\rho \in [0,1]$, the fee solves the following bargaining

[3] Within our analytical framework, we focus primarily on the licensing channel. In principle, patents can generate value through alternative mechanisms: direct commercialization in products manufactured by the patent owner, strategic blocking of competitors, improving the firm's bargaining position in cross-licensing negotiations, signaling technological competence to investors, or serving as collateral for financing.

problem

$$\max_{\ell_{ij}} \; [(1-\tau_i)\ell_{ij}]^{\rho} \left[(1-\tau_j^c)(y_{ij}-\ell_{ij})\right]^{1-\rho}$$

The solution is $\ell_{ij} = \rho y_{ij}$. Hence, the after-tax income for the owner in country i and licensee in country j are $(1-\tau_i)\rho y_{ij}$ and $(1-\tau_j^c)(1-\rho)y_{ij}$, respectively.

A patent licensed in country j may expire (or be imitated) with probability ϕ_j at the end of the period. If the patent does not expire, then the patent owner may license it again next period. If the patent expires, then the owner receives a pre-tax scrap value $\varepsilon_{ij} \geq 0$ next period and loses the opportunity to license it ever again in country j. Expiration or imitation of a patent is country-specific and independent across countries. That is, if a patent expires or is imitated in country j the owner in country i can no longer license it in country j but may still license it in other countries.

3.2 Licensing

Let V_{ij} be the pre-tax value of a patent owned in country i and licensed in country j. We obtain the following recursion

$$V_{ij} = \rho y_{ij} + \frac{1}{1+r_i}\left[(1-\phi_j)V_{ij} + \phi_j \varepsilon_{ij}\right] \tag{2}$$

Solving for V_{ij} we obtain

$$V_{ij} = \frac{(1+r_i)\rho y_{ij} + \phi_j \varepsilon_{ij}}{r_i + \phi_j} \tag{3}$$

Licensing of patents is non-rival. The following result follows.

Lemma 1. *Patents are optimally licensed in all countries.*

Proof. Follows from $V_{ij} \geq 0$. □

The pre-tax value of a patent owned in country i is the sum of licenses to all countries,

$$W_i = \sum_{j=1}^{M} V_{ij} \tag{4}$$

Using (3) we obtain,

$$W_i = \sum_{j=1}^{M} \frac{(1+r_i)\rho y_{ij} + \phi_j \varepsilon_{ij}}{r_i + \phi_j} \tag{5}$$

3.3 Selling

A patent owner may prefer to sell the patent instead of licensing it. In this case, the seller of the patent pays a capital-gains tax τ_i^k from the proceeds of the sale. Let P_{ij} be the sale price of a patent transferred from country i to country j. A seller chooses the country that offers the highest price. The pre-tax value associated with selling a patent is

$$S_i = \max_j \; P_{ij} \tag{6}$$

The sale price P_{ij} is the outcome of a Nash-bargaining game and thus, maximizes the joint surplus of buyer and seller. The surplus for the seller is the after-tax sale price, $(1-\tau_i^k)P_{ij}$, minus the outside option of not selling, i.e., the after-tax value of the patent, $(1-\tau_i)W_i$. The surplus for the buyer is the after-tax profit, which consists of the licensing value minus the sale price, after taxes, $(1-\tau_j)(W_j - P_{ij})$.

Let $\theta \in [0,1]$ be the seller's bargaining power. The sale price P_{ij} is the solution to the following bargaining problem

$$\max_{P_{ij}} \; \left[(1-\tau_i^k)P_{ij} - (1-\tau_i)W_i)\right]^{\theta}\left[(1-\tau_j)(W_j - P_{ij})\right]^{1-\theta}$$

which yields

$$P_{ij} = (1-\theta)\left(\frac{1-\tau_i}{1-\tau_i^k}\right)W_i + \theta W_j \tag{7}$$

We need to verify that surpluses are positive for buyers and sellers. That is, $(1-\tau_i^k)P_{ij} - (1-\tau_i)W_i \geq 0$ and $(1-\tau_j)(W_j - P_{ij}) \geq 0$. Using expression (7) it turns out that both these requirements are satisfied by the same condition. Specifically, a sale from i to j is profitable for both buyers and sellers when

$$W_j \geq \left(\frac{1-\tau_i}{1-\tau_i^k}\right)W_i \tag{8}$$

Combining (7) and (8), we get that $dP_{ij}/d\theta \geq 0$, i.e., the market price of a patent is increasing in the seller's bargaining power. Thus, if there are gains from trade from a patent sale, i.e., (8) is satisfied, then (7) implies $P_{ij} \in [W_i(1-\tau_i)/(1-\tau_i^k), W_j]$.

Using (6) and (7) the pre-tax value from selling the patent to the country with the highest bidder is

$$S_i = (1-\theta)\left(\frac{1-\tau_i}{1-\tau_i^k}\right)W_i + \theta \max_j \; W_j \tag{9}$$

The decision of which country to sell to is independent of the tax rates paid by potential buyers, as reflected by the second term in (9). The reason is that such taxes are internalized in the sale price, P_{ij}, as shown in the bargaining problem. However, the pre-tax value of selling, S_i, does depend on local (seller's) taxes, through the difference between the patent income tax rate, τ_i, and the capital-gains tax rate, τ_i^k.

The result derived above means that if a patent owner decides to sell, then the buyer's country is determined by characteristics other than tax rates in the destination country. Therefore, if countries only differed in their tax rates, then W_j would be the same for all j and the buyer's country would be indeterminate. This result highlights an important distinction between arm's-length transactions and intra-firm reallocations. Under arm's-length sales, destination-country tax rates do not directly affect the location choice, as they are fully capitalized into the sale price through bargaining. By contrast, in the data we observe a strong relationship between tax differentials and patent transfers, particularly within multinational firms. This pattern is therefore consistent with tax considerations operating primarily through intra-firm transfers—where pricing and organizational structure allow firms to exploit cross-country tax differences—rather than through arm's-length sales.

3.4 Licensing vs. selling

The decision to license or sell comes down to the maximum between after-tax licensing and selling values. The owner of a patent prefers licensing to selling when $(1-\tau_i)W_i \geq (1-\tau_i^k)S_i$. The following result states when licensing is preferred to selling.

Proposition 1. *When $\theta = 0$, a patent owner is indifferent between licensing and selling. When $\theta > 0$, a patent owner in country i prefers licensing over selling if and only if there is no country $j \neq i$ for which (8) holds with strict inequality.*

Proof. Take condition (9). When $\theta = 0$, $(1-\tau_i^k)S_i = (1-\tau_i)W_i$ and so, the values of licensing and selling are the same. When $\theta > 0$, licensing is preferred to selling when

$$\left(\frac{1-\tau_i}{1-\tau_i^k}\right) W_i \geq \max_j \; W_j \tag{10}$$

Suppose there exists a $j \neq i$ for which (8) holds with strict inequality, i.e., there is a strictly positive surplus for buyers and sellers when the patent is sold from country i to country j. Then, (10) cannot be satisfied, so selling is preferred to licensing. Now suppose that there is no $j \neq i$ for which (8) holds with strict inequality. Then, (10) is satisfied and so, licensing is preferred to selling. □

When buyers have all the bargaining power, $\theta = 0$, they extract all the surplus from a sale and thus, the owner of a patent is indifferent between licensing and selling. When $\theta > 0$, licensing is preferred over selling unless there exists at least one country in which a sale would generate a strictly positive surplus for both buyer and seller. The intuition behind Proposition 1 is simple: if relocating a patent generates a positive surplus then it is optimal to do so (i.e., no money is left on the table) and the only issue is how to divide the surplus.

3.5 Profit shifting

Now we extend our analysis to consider profit-shifting strategies. Consider a patent owner in country i that establishes a subsidiary in country j and transfers the patent rights to this subsidiary at a potentially manipulated transfer price. This strategy allows the multinational firm to optimize its global tax burden.

The patent owner's objective is to maximize total global profits across both the headquarters and the subsidiary. For simplicity, we assume that once the patent is transferred to the subsidiary, that subsidiary can only license the patent (not sell it further). Let T_{ij} be the transfer price between a patent owner residing in country i and a subsidiary in country j. Let κ_{ij} represent the setup cost of the subsidiary. This subsidiary earns pre-tax profits consisting of the licensing value, W_j, net of the transfer price, T_{ij}. After-tax profits derived from profit shifting are

$$\Pi_i = \max_j \ \underbrace{(1-\tau_i^k)(T_{ij}-\kappa_{ij})}_{\text{Headquarter profits}} + \underbrace{(1-\tau_j)(W_j - T_{ij})}_{\text{Affiliate profits}}$$

which can be rearranged as

$$\Pi_i = \max_j \ (1-\tau_j)W_j - (1-\tau_i^k)\kappa_{ij} + (\tau_j - \tau_i^k)T_{ij} \tag{11}$$

If the domestic capital gains and the foreign patent income tax rates are the same, i.e., $\tau_i^k = \tau_j$, then the value of profit shifting to country j is independent of the transfer price T_{ij}. In such a case, the gains from profit shifting depend only on licensing value in the foreign country, net of subsidiary setup costs. More generally, how the transfer price impacts global profits depends on tax differences.

The market price for a patent sale, P_{ij}, serves as a reference for the transfer price, T_{ij}. Suppose that $T_{ij} = \sigma_{ij} P_{ij}$, where $\sigma_{ij} \geq 0$ and implies a surcharge ($\sigma_{ij} > 1$) or discount ($\sigma_{ij} < 1$) relative to market prices. When $\sigma_{ij} = 1$ we obtain what is known as "arm's-length" pricing. In practice, multinational companies often set $\sigma_{ij} < 1$ for strategic reasons, including tax optimization. In contrast, $\sigma_{ij} > 1$ might occur when patents are transferred to subsidiaries in high-tax jurisdictions. The deviation from arm's-length pricing ($\sigma_{ij} \neq 1$) represents a key focus for tax authorities concerned with base erosion and profit shifting (BEPS) activities. In the Appendix, we develop an extension in which multinational firms can additionally manipulate intra-firm royalty payments to foreign affiliates through transfer pricing, providing an additional channel for profit shifting beyond the relocation of patent ownership.

Given $T_{ij} = \sigma_{ij} P_{ij}$ and using (7) to replace the expression for P_{ij}, we can rewrite the profit-shifting problem (11) as:

$$\Pi_i = \max_j \; (1-\tau_j)W_j - (1-\tau_i^k)\kappa_{ij} + (\tau_j - \tau_i^k)\sigma_{ij} \Big[\underbrace{(1-\theta)\Big(\tfrac{1-\tau_i}{1-\tau_i^k}\Big) W_i + \theta W_j}_{P_{ij}} \Big]. \qquad (12)$$

There are three relevant terms determining the value of profit shifting. The first reflects the after-tax revenue derived from licensing the patent in the destination country. The second term is the affiliate setup cost, which is net of the domestic tax deduction. The third term reflects the tax impact of transfer-pricing. When the foreign income tax rate is larger than the domestic capital-gains tax rate, $\tau_j > \tau_i^k$, a higher transfer price (larger σ_{ij}) increases profits as the foreign tax deduction outweighs the domestic tax obligation. The converse occurs when the foreign income tax rate is lower than the domestic capital-gains tax rate. As mentioned above, the impact disappears when these two tax rates are the same.

4 The Mechanism: The Allocation of Patents

In the previous section we presented three alternatives for a patent owner: licensing, selling and profit shifting. The decision is based on their after tax value. Specifically, the owner of a patent solves the following decision problem:

$$\max \Big\{ \underbrace{(1-\tau_i)W_i}_{\text{Licensing}} \; , \; \underbrace{(1-\tau_i^k)S_i}_{\text{Selling}} \; , \; \underbrace{\Pi_i}_{\text{Profit shifting}} \Big\} \qquad (13)$$

where W_i, S_i and Π_i are defined in expressions (5), (9) and (12), respectively.

Patent allocation depends on two forces: cross-country differences in licensing value and cross-country tax differences. The first determines where patents generate the highest economic surplus, while the second creates incentives to relocate ownership across jurisdictions. Differences in licensing value may make it optimal to sell a patent or transfer it to a foreign affiliate, whereas tax differences create incentives for profit shifting, particularly when multinational firms can manipulate transfer prices across jurisdictions.

Proposition 1 established that selling is preferred to licensing whenever selling generates a positive surplus. Adding the possibility of profit shifting complicates characterizing the decision for the patent owner, due to the added elements of price manipulation and setup costs. In what follows, we analyze the decision focusing on differences across countries in licensing value and tax regimes.

4.1 The role of licensing values

Consider the case when there are no differences in tax rates across countries and when capital income and patent income tax rates are the same. That is, suppose $\tau_i = \tau_i^k = \tau$ for all $i = 1, \ldots, M$. Note that we still allow corporate income tax rates, τ_i^c, to differ across countries as they do not affect a patent owner's decision problem. Under these assumptions, problem (13) simplifies to

$$\max\left\{W_i, (1-\theta)W_i + \theta \max_j \; W_j, \max_j \; W_j - \kappa_{ij}\right\} \tag{14}$$

Note that tax rates drop from every term, as they are the same across alternatives. We can subtract W_i from every term in (14) to further highlight the relevant margin:

$$\max\left\{0, \theta(\max_j \; W_j - W_i), \max_j \; W_j - W_i - \kappa_{ij}\right\} \tag{15}$$

The decision of licensing vs. selling boils down to licensing value. The patent should be located in the country that can generate the highest licensing income—though recall from Proposition 1 that the patent owner is indifferent when buyers have all the bargaining power, i.e., $\theta = 0$. If the current location i generates the maximum licensing income, $W_i = \max_j W_j$, then licensing dominates both selling and profit shifting.

In contrast, when $W_i < \max_j W_j$, the relevant alternatives depend on the seller's bargaining power, θ. When $\theta = 0$, the patent owner is indifferent between licensing and selling. If $W_j - \kappa_{ij} >$

W_i, profit shifting is preferred to licensing or selling, as relocating the patent to a foreign subsidiary generates enough additional income to cover setup costs. When $\theta > 0$, licensing is dominated by both selling and profit shifting. In this case, the decision is a race between the seller's bargaining power, θ, and the setup costs, κ_{ij}. Selling becomes more attractive as we increase the seller's bargaining power, while profit shifting becomes more desirable as we lower the subsidiary's setup costs. In the extreme case of $\kappa_{ij} = 0$, the patent owner is indifferent between selling and profit shifting when $\theta = 1$, and strictly prefers profit shifting when $\theta < 1$.

4.2 The role of tax rates

We now switch to the second channel that impacts the location decision: tax differences across countries. To focus the analysis, we assume that countries only differ in tax rates. That is, assume $y_{ij} = y$, $r_i = r$, $\phi_j = \phi$, $\varepsilon_{ij} = \varepsilon$, $\kappa_{ij} = \kappa$ and $\sigma_{ij} = \sigma$. By (5) these assumptions imply $W_i = W$ for all $i = 1, \dots, M$. Now, the patent location decision does not depend on licensing value, but on tax differences (including factors that interact with tax differences).

The patent owner's problem (13) simplifies to

$$\max\left\{1-\tau_i, (1-\theta)(1-\tau_i)+\theta(1-\tau_i^k), \max_j\ (1-\tau_j)-(1-\tau_i^k)\left(\tfrac{\kappa}{W}\right)+(\tau_j-\tau_i^k)\sigma\left[(1-\theta)\left(\tfrac{1-\tau_i}{1-\tau_i^k}\right)+\theta\right]\right\} \tag{16}$$

where each term was divided by the common term, W. Note that W now scales the subsidiary's setup costs, κ, in the profit-shifting term.

We can further rearrange the patent owner's problem by subtracting $1 - \tau_i$ from every term:

$$\max\left\{0, \theta(\tau_i - \tau_i^k), \max_j\ \tau_i - \tau_j - (1 - \tau_i^k)\left(\tfrac{\kappa}{W}\right) + (\tau_j - \tau_i^k)\sigma\left[(1 - \theta)\left(\tfrac{1-\tau_i}{1-\tau_i^k}\right) + \theta\right]\right\} \tag{17}$$

As discussed above, when $\theta = 0$ the patent owner is indifferent between licensing and selling. When $\theta > 0$, the decision between licensing and selling depends on the difference between patent income and capital income tax rates, $\tau_i - \tau_i^k$. The lower tax rate determines the preferred option: if the capital income tax rate is lower (higher) then selling (licensing) is preferred.

The value of profit shifting now depends critically on foreign patent income tax rates through two channels. The first channel is direct: a higher foreign patent income tax rate, τ_j, makes shifting profit to a subsidiary less attractive. The second channel works through transfer pricing and the tax savings that arise from tax rate differences across the origin and destination countries. In this case, a higher foreign tax rate raises profits as it increases the foreign tax deduction on the patent payment by the subsidiary.

As a starting point, consider the case $\kappa = \sigma = 0$. That is, profit shifting has no setup costs and the transfer is done for free. In this case, (17) simplifies to $\max\{0, \theta(\tau_i - \tau_i^k), \tau_i - \min_j \tau_j\}$. All alternatives yield the same value when there are no differences in tax rates. Licensing is preferred when the domestic patent income tax, τ_i, is lower than domestic capital income taxes, τ_k and the minimum foreign patent income tax, τ_j. Selling is preferred to profit shifting when the minimum foreign income tax is high relative to domestic taxes; specifically, when $\min_j \tau_j \geq (1-\theta)\tau_i + \theta\tau_i^k$. In general, the decision on where to allocate the patent depends both on tax differences and market structure.

Now suppose, $\kappa > 0$. This has no effect on the values of licensing and selling, but lowers the value of profit shifting. This impact is mitigated by the deduction on domestic capital income taxes.

Finally, consider $\sigma > 0$. Again, this only impacts the value of profit shifting, through the difference between foreign income and domestic capital income tax rates, $\tau_j - \tau_i^k$. We refer to term multiplying σ in the profit-shifting expression as the *manipulation term*, as it reflects the gains from transfer-price manipulation. The sign of this term depends on the difference between the foreign patent-income tax rate, τ_j, and the domestic capital-gains tax rate, τ_i^k. When $\tau_i^k < \tau_j$, the manipulation term is positive. As a result, increasing σ, which brings the transfer price closer to arm's-length pricing, raises the profit-shifting payoff and expands the region of lower values of τ_i^k, where selling was previously preferred. In contrast, when $\tau_i^k > \tau_j$, the manipulation term is negative, so the profitability of shifting profits declines. In the Appendix we further expand this analysis with some comparative statics.

In the next section, we extend the model to incorporate heterogeneity in patent values, which allows us to quantify how these margins shape the distribution of global royalty income and perform royalty counterfactuals.

5 Extended Model and Policy Counterfactuals

To evaluate tax policy counterfactuals, we extend the model to incorporate heterogeneity in patent quality and discipline it using observed patent transfers and royalty flows. This extension allows organizational choices over patent ownership to vary systematically across patents and generates a mapping between patent reallocations, royalty income, and tax revenues across jurisdictions.

To keep the analysis tractable and focus on the main economic mechanisms, we make several simplifying assumptions. First, the interest rate and the imitation probability are identical across countries, i.e., $r_i = r$ and $\phi_i = \phi$. Second, we set the pre-tax scrap value ε_{ij} to zero. Third, we equate capital-gains and corporate income tax rates, $\tau_i^k = \tau_i$, which is a reasonable approximation for most countries. In addition, as shown in the previous section, when $\tau_i^k = \tau_i$ the owner of a patent is indifferent between licensing and selling. We therefore assume that licensing is chosen in this case, implying that arm's-length sales play a limited role in equilibrium relative to licensing and intra-firm transfers. Empirically, outright cross-border patent sales are rare and, when they occur, tend to reflect strategic acquisitions or mergers rather than systematic responses to tax differentials. Consistent with this, our gravity regressions show that tax differentials strongly predict intra-firm patent transfers but have little explanatory power for arm's-length sales. Finally, the United States is the largest source of cross-border patent transfers in our dataset, so we treat it as the origin country and focus on the reallocation of patents from the US to foreign destinations. We therefore suppress the origin-country subscript i in what follows, so unsubscripted variables refer to the United States.

5.1 The global allocation of heterogeneous patents

In the data, most patents remain in their country of origin, while the patents that relocate account for a disproportionate share of royalty income. We capture this selection using a Pareto distribution for patent cash flows,

$$F(y) = 1 - \left(\frac{y_{\min}}{y}\right)^{\alpha}, \qquad y \geq y_{\min},$$

where $\alpha > 1$ governs the thickness of the upper tail.[4]

Figure 6 motivates this selection mechanism. Royalty income and royalties per transferred patent are highly concentrated in a small set of destinations, particularly low-tax jurisdictions and tax havens. These patterns suggest that the upper tail of the patent-value distribution plays a central role in shaping observed transfers and royalty allocations.

[4] In the Appendix, we develop a version of the model in which patent quality is endogenous.

Figure 6: Royalties

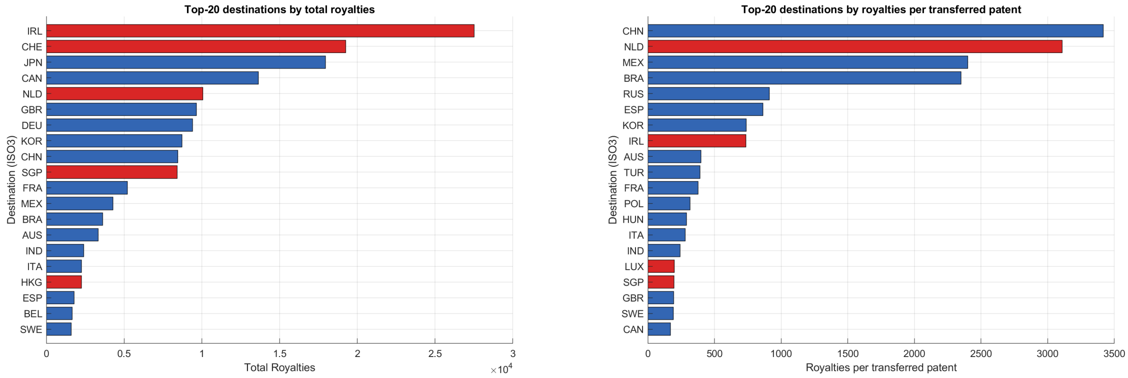

(a) Top-20 destinations by total royalties.

(b) Top-20 destinations by royalties per transferred patent.

Note. Royalty *levels* (left) and *royalties per transferred patent* (right) are both highly concentrated in a small set of destinations, with low-tax jurisdictions and tax havens over-represented (red). These facts motivate a selection framework with a thick upper tail for patent values.

For a patent with cash flow y, the after-tax payoff from licensing is

$$\mathcal{W}(y) = (1-\tau)M\eta y, \tag{18}$$

where

$$\eta \equiv \frac{(1+r)\rho}{r+\phi},$$

and M denotes the number of licensing destinations.

If instead the patent is shifted to destination j, the after-tax payoff is

$$\Pi_j(y) = [1-\sigma_j\tau-(1-\sigma_j)\tau_j]\, M\eta y - (1-\tau)\kappa_j, \tag{19}$$

where τ_j is the destination tax rate, $\sigma_j \leq 1$ captures deviations from arm's-length transfer pricing, and $\kappa_j > 0$ is the fixed cost of establishing a foreign affiliate.

Because both payoffs are linear in patent value, there exists a cutoff level $y_j^\star$ above which profit shifting dominates licensing:

$$y_j^\star = \frac{(1-\tau)}{(1-\sigma_j)(\tau-\tau_j)}\frac{\kappa_j}{M\eta}. \tag{20}$$

Only sufficiently valuable patents are relocated abroad. Lower foreign tax rates, lower relocation costs, and larger transfer-pricing discounts reduce the cutoff and increase the fraction of patents shifted. Additional derivations and characterization results are provided in the Appendix.

The fraction of patents shifted to destination j is therefore

$$s_j = \left(\frac{y_{\min}}{y_j^\star} \right)^{\alpha}, \tag{21}$$

while the corresponding share of royalty income is

$$\tilde{s}_j = \left(\frac{y_{\min}}{y_j^\star} \right)^{\alpha-1}. \tag{22}$$

Because royalty income is proportional to patent value, the relocation of a relatively small number of high-value patents can generate large shifts in the global allocation of royalty income.

The heavy-tailed nature of the Pareto distribution implies that small changes in relocation incentives generate disproportionately large changes in the allocation of royalty income.

We next extend the framework to multiple foreign destinations. Each destination is characterized by a destination-specific slope parameter,

$$\beta_j \equiv [1 - \sigma_j \tau - (1 - \sigma_j)\tau_j]\, M\eta, \tag{23}$$

so that the profit-shifting payoff can be written as

$$\Pi_j(y) = \beta_j y - (1 - \tau_j)\kappa_j.$$

This structure generates sorting across destinations. Destinations with higher slopes attract higher-value patents because payoff differences scale with patent value, while destinations with lower fixed costs may attract intermediate-value patents. As a result, the allocation of patents across jurisdictions depends jointly on tax rates, transfer-pricing opportunities, and relocation costs.

A deterministic cutoff rule, however, implies that only the destination with the highest payoff receives patent inflows for a given patent value. To allow for smoother substitution patterns across destinations, we introduce destination-specific amenities δ_j and idiosyncratic

Type-I extreme-value shocks ε_j. The resulting payoff from shifting a patent to destination j is

$$\Pi_j(y) + \delta_j + \varepsilon_j.$$

Under the extreme-value assumption, the probability that a patent of quality y is shifted to destination j is

$$\pi_j(y) = \frac{\exp\left[(\Pi_j(y) + \delta_j)/\mu\right]}{\exp\left[\mathcal{W}(y)/\mu\right] + \sum_{m=1}^{M} \exp\left[(\Pi_m(y) + \delta_m)/\mu\right]},$$

where $\mu > 0$ governs the dispersion of unobserved heterogeneity. This formulation preserves the underlying economic structure of the model while allowing all destinations to receive positive patent inflows.

US tax revenues arise from two sources: royalty income booked domestically when patents are licensed, and capital-gains taxation on patents transferred to foreign affiliates. Aggregating across patents and destinations, total US tax revenue is

$$\mathcal{T}_{\text{US}} = \tau M \eta \int_{y_{\min}}^{\infty} \left[y\pi^L(y) + \sum_{j=1}^{M} \sigma_j y \pi_j(y) \right] f_Y(y)\, dy, \tag{24}$$

where $\pi^L(y)$ denotes the probability that a patent remains licensed from the United States. We use this framework in the next section to evaluate how alternative tax policies affect patent allocation, royalty income, and US tax revenues.

5.2 Calibration Strategy

This section describes how we discipline the quantitative model developed in Section 5.1. All payoff expressions, cutoff rules, and notation are defined there and are not repeated here. Our calibration maps observed US-origin patent transfers and royalty receipts into a small set of structural parameters governing patent heterogeneity, profit-shifting costs, transfer-pricing frictions, and destination choice.

Destinations are partitioned into four mutually exclusive regions: the United States (HOME), low-tax advanced economies (LOWTAX), tax havens (HAVEN), and other countries (OTHER). This partition is based on statutory corporate tax rates, grouping destinations with similar tax environments into a small number of regions used consistently throughout the quantitative analysis (Tørsløv, Wier, and Zucman, 2023; Santacreu, 2026).[5]

[5]LOWTAX: Switzerland (CHE), Ireland (IRL), Luxembourg (LUX), the Netherlands (NLD), and Singapore

The calibration targets four empirical objects: the concentration of royalty income across destinations, the fraction of US-origin patents shifted abroad, the allocation of shifted patents across regions, and statutory tax rates. Table 2 reports these moments. We normalize the lower bound of the patent-quality distribution to $y_{\min} = 1$, which fixes the units of patent quality without loss of generality. The statutory tax rate for the US is set to 35%.

Table 2: Calibration targets

Region	Fraction shifted	Number of transferred patents	Royalty income	Statutory tax rate
LOWTAX	2.30%	1,803	74,942	12%
HAVEN	0.20%	208	3,110	0%
OTHER	5.53%	5,343	88,501	30%

Notes: The table reports empirical objects used to discipline the quantitative analysis. The fraction shifted, number of transferred patents, and royalty income are computed using US-origin patents. Royalty income is measured in thousands of USD. Statutory tax rates are region averages taken directly from the data and are held fixed throughout the analysis.

These moments summarize the key empirical patterns used to discipline the model. The fraction shifted measures the share of US-origin patents that are relocated to each destination group, capturing the extensive margin of profit shifting. The number of transferred patents reflects the relative importance of each region in terms of transaction counts, while royalty income captures the value associated with relocated patents and highlights the concentration of income across destinations. Taken together, these moments reveal that patent transfers are highly skewed: a small number of destinations account for a disproportionate share of both transactions and value. This joint distribution of flows and value is central to identifying the model's parameters, as it informs both the selection of high-value patents into relocation and the relative attractiveness of different destinations.

Table 3 reports the calibrated parameters. Patent value heterogeneity is disciplined using the selection mechanism described in Section 5.1. The Pareto tail parameter α governs the thickness of the upper tail of the patent-value distribution. A lower α implies that a small fraction of high-value patents accounts for a large share of total income generated by patents. We therefore identify α from the observed concentration of royalty income, ensuring that the model reproduces the empirical fact that a small fraction of patents accounts for a disproportionate share of total returns.

Region-specific fixed costs of profit shifting are identified from the extensive margin of re-

(SGP). HAVEN: Barbados (BRB), Hong Kong (HKG), St. Kitts and Nevis (KNA), Malta (MLT), Panama (PAN), and Seychelles (SYC).

location decisions. In the model, κ_j shifts the cutoff above which patents are transferred to destination j. Higher fixed costs reduce the mass of patents that are profitably relocated. We therefore use observed shifting frequencies across regions to discipline κ_j, mapping differences in the fraction of patents transferred into differences in the implied cutoff levels.

Transfer-pricing discounts are identified from discrepancies between the allocation of R&D service receipts and the distribution of patent transfer flows across regions. In the data, we do not observe separate transaction values for affiliated and unaffiliated transfers; instead, we allocate total BEA R&D receipts using observed affiliated shares of patent transfers. In the model, a lower σ_j reduces the effective valuation at which patents are transferred to affiliates, generating a gap between the share of affiliated transactions in total value and their share in observed flows. We therefore discipline σ_j using cross-region differences between the allocation of R&D receipts and the distribution of transfer counts. Values of $\sigma_j < 1$ are interpreted as reflecting lower implied valuations of intra-firm transactions relative to arm's-length benchmarks.

Destination amenities capture systematic differences in location attractiveness that are not explained by tax rates or transfer-pricing incentives. In the discrete-choice structure of the model, δ_j shifts the level of the payoff associated with each destination without affecting its slope in patent value. As a result, amenities primarily affect the relative allocation of patents across destinations conditional on relocation. We use them to match residual differences in destination shares that are not accounted for by taxes, transfer-pricing frictions, or fixed costs.

In the quantitative analysis, all patents not transferred abroad are assumed to be licensed at home, and the model is scaled so that baseline US licensing revenues match observed royalty receipts. This normalization anchors the level of the tax base and allows us to focus on how tax policy reallocates patent ownership across locations.

Overall, the calibration exploits distinct moments to separately identify each parameter: the upper tail of the patent-value distribution (via royalty concentration), the extensive margin of relocation (via shifting frequencies), the relative valuation of affiliated transactions (via value-flow discrepancies), and residual destination attractiveness (via location shares).

5.3 Evaluating Tax Policies and US Tax Revenue

We use the calibrated model to evaluate how different tax policies affect the location of patent ownership and, in turn, US tax revenue. The objective is to quantify how changes in tax differentials and profit-shifting frictions alter the US royalty base and the associated tax revenues.

Table 3: Calibrated parameters

Parameter	Value	Description
Pareto tail (α)	2.26	Patent value heterogeneity
Licensing share (ρ)	0.50	Bargaining weight in licensing
Discount rate (r)	0.02	Annual discount rate
Imitation rate (ϕ)	0.30	Patent expiration/obsolescence
Logit scale (μ)	0.058	Dispersion in destination choice
Transfer-pricing discount (HAVEN)	0.22	Relative valuation (proxy)
Transfer-pricing discount (LOWTAX)	0.42	Relative valuation (proxy)
Transfer-pricing discount (OTHER)	0.66	Relative valuation (proxy)
Profit-shifting fixed cost (LOWTAX)	3.61	Extensive-margin discipline
Profit-shifting fixed cost (HAVEN)	11.79	Extensive-margin discipline
Profit-shifting fixed cost (OTHER)	0.49	Extensive-margin discipline
Destination amenity (LOWTAX)	0.42	Location-specific attractiveness
Destination amenity (HAVEN)	3.44	Location-specific attractiveness
Destination amenity (OTHER)	0.00	Location-specific attractiveness

Notes: Transfer-pricing discounts capture relative valuation differences across regions. Fixed costs and destination amenities are calibrated to match observed shifting frequencies and destination shares. All calibration details are provided in the Appendix.

In each counterfactual, we hold fixed the distribution of patent quality and the underlying economic environment, and vary only the policy parameters governing profit shifting. For each scenario, we compute the share of US patents shifted abroad, the US royalty base, and total US tax revenue. We report tax revenues over a decade, while royalty flows are expressed in annual terms.

Table 4 summarizes the results. In the baseline, about 8% of US-origin patents are shifted abroad. This implies US tax revenues of roughly \$115 billion over a decade (about \$11.5 billion per year), consistent with annual US royalty receipts of about \$32.6 billion.

We consider four policy experiments: harmonizing foreign royalty tax rates to the US level, strengthening transfer-pricing enforcement, increasing the fixed costs of relocating patents abroad, and introducing a global minimum tax. The first three policies substantially reduce or eliminate the incentive to shift profits abroad, while the fourth—a global minimum tax—only partially reduces it.

Policies that eliminate shifting incentives deliver very similar outcomes. Strengthening transfer-pricing enforcement (high σ), increasing the fixed costs of shifting (high κ), or harmonizing foreign tax rates to the US level all reduce the share of shifted patents to essentially zero. As a result, the US royalty base expands substantially, reaching about \$41 billion annually, and tax revenues rise to about \$143–144 billion over a decade. This corresponds to an increase of roughly \$28 billion, or about 25% relative to the baseline.

In contrast, imposing a 20% minimum tax on foreign income has almost no effect. The share of shifted patents remains close to its baseline level, and US tax revenue is essentially unchanged. The reason is that, even with a minimum tax, foreign tax rates remain below the US level, so the incentive to shift high-value patents persists.

Table 4: US tax revenue from royalties: baseline and counterfactuals

Scenario	US tax revenue (Decade, \$bn)	Change (\$bn)	US royalties (Annual, \$bn)	Patents shifted (%)
Baseline	115	–	32.6	7.8
High enforcement (High σ)	143	+28	41.0	0.1
High costs (High κ)	144	+29	41.0	0.0
Tax harmonization ($\tau_j = 0.35$)	143	+28	41.0	0.2
Minimum tax ($\tau_j \geq 0.20$)	115	0	32.6	7.8

Notes: "Patents shifted" is the model-implied share of US-origin patents located abroad. "US royalties" are annual receipts from US-origin patents booked domestically. "US tax revenue" includes both royalty taxation and capital-gains taxation from transfers. All values are rounded for presentation.

These results underscore that the incentive to shift profits depends primarily on cross-country tax differentials rather than tax levels per se. Policies that eliminate these differences—either directly through tax harmonization or indirectly through enforcement or higher relocation costs—substantially increase the US tax base. Policies that leave tax gaps in place have much more limited effects.

More broadly, the results highlight that the location of intangible income is highly sensitive to relatively small differences in international tax treatment. Because a small number of high-value patents account for a disproportionate share of royalty income, persistent tax wedges generate large reallocations of the tax base across jurisdictions. In this environment, policies that target the profitability of relocating intangible assets are considerably more effective than policies that simply impose a minimum floor on foreign tax rates.

6 Final Remarks

International tax regimes play a central role in shaping where multinational firms locate and transfer their intangible assets. Using a new dataset on cross-border patent ownership, we documented that intra-group transfers are highly sensitive to statutory tax differentials, especially during the peak profit-shifting period of 2007–2012, and that tax havens continue to absorb a disproportionate share of patent relocations and royalty income even after BEPS and related reforms. We rationalize these facts with a model in which firms choose between licensing, selling, or shifting patents in response to corporate and capital-gains tax differences and transfer-pricing frictions. Our counterfactual analysis shows that narrowing cross-country tax differentials through rate harmonization or higher shifting costs produces sizable reductions in profit shifting and large revenue gains for high-tax countries, whereas modest minimum taxes have more limited effects on the allocation of high-value patents.

Looking ahead, the implementation of Pillar Two and further tightening of transfer-pricing rules will provide an opportunity to test these mechanisms in real time. More broadly, an important next step is to understand how international tax systems affect not only the location of intangible assets, but also the incentives to innovate in the first place. Extending the framework to incorporate endogenous innovation and patent quality would allow future work to study how profit-shifting opportunities shape the global allocation of innovative activity itself.

References

Becker, Johannes and Nadine Riedel. 2012. "Cross-border tax effects on affiliate investment—evidence from European multinationals." *European Economic Review* 56 (3):436–450.

Ciaramella, Laurie. 2023. "Taxation and the transfer of patents: Evidence from Europe." *European Economic Review* 151:104312.

Clausing, Kimberly A. 2009. "Multinational firm tax avoidance and tax policy." *National Tax Journal* 62 (4):703–725.

———. 2016. "The effect of profit shifting on the corporate tax base in the United States and beyond." *National Tax Journal* 69 (4):905–934.

Crouzet, Nicolas, Janice C. Eberly, Andrea L. Eisfeldt, and Dimitris Papanikolaou. 2022. "The Economics of Intangible Capital." *Journal of Economic Perspectives* 36 (3):29–52.

De Rassenfosse, Gaetan, Jan Kozak, and Florian Seliger. 2021. "Geocoding of Worldwide Patent Data." *Available at SSRN 3425764* .

Devereux, Michael P. 2022. "The OECD Global Minimum Tax: Incentives and Interaction with Tax Competition." Working Paper WP22/25, Oxford University Centre for Business Taxation.

Dyrda, Sebastian, Guangbin Hong, Muhammad Ali Sajid, and Joseph B. Steinberg. 2025. "Global Ripple Effects of Corporate Tax Reforms." Working paper.

Dyrda, Sebastian, Guangbin Hong, and Joseph B Steinberg. 2024a. "A macroeconomic perspective on taxing multinational enterprises." *Journal of international economics* 152:104022.

———. 2024b. "Optimal taxation of multinational enterprises: A Ramsey approach." *Journal of Monetary Economics* 141:74–97.

Fosfuri, Andrea, Christian Helmers, and Catherine Roux. 2012. "Are Joint Patents Collusive? Evidence from the US and Europe." Working Paper, SSRN 2164222.

Gaessler, Fabian, Bronwyn H Hall, and Dietmar Harhoff. 2021. "Should there be lower taxes on patent income?" *Research Policy* 50 (1):104129.

Guvenen, Fatih, Raymond J. Mataloni, Dylan G. Rassier, and Kim J. Ruhl. 2022. "Offshore Profit Shifting and Aggregate Measurement: Balance of Payments, Foreign Investment, Productivity, and the Labor Share." *American Economic Review* 112 (6):1848–1884.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Alexander Hassan, Josh Lerner, and Ahmed Tahoun. 2021. "The Diffusion of New Technologies." Working Paper 28999, National Bureau of Economic Research. URL http://www.nber.org/papers/w28999.

Karkinsky, Tom and Nadine Riedel. 2012. "Corporate Taxation and the Choice of Patent Location within Multinational Firms." *Journal of International Economics* 88 (1):176–185.

LaBelle, Jesse, Inmaculada Martinez-Zarzoso, Ana Maria Santacreu, and Yoto V Yotov. 2023. "Cross-border Patenting, Globalization, and Development." .

Santacreu, Ana Maria. 2024. "Comment on: "Optimal Taxation of Multinational Enterprises: A Ramsey Approach"." *Journal of Monetary Economics* 141:98–100.

———. 2026. "International technology licensing, intellectual property rights, and tax havens." *Review of Economics and Statistics* :1–45.

Santacreu, Ana Maria and Ashley Stewart. 2025. "The Impact of the 2017 Tax Cuts and Jobs Act on US Multinationals' Intangible Assets." *Federal Reserve Bank of St. Louis Review* 107 (4):1–16.

Santos Silva, J.M.C. and Silvana Tenreyro. 2006. "The Log of Gravity." *Review of Economics and Statistics* 88 (4):641–658.

Santos Silva, J.M.C. and Silvana Tenreyro. 2021. "The Log of Gravity At 15." School of economics discussion papers, School of Economics, University of Surrey.

Tørsløv, Thomas, Ludvig Wier, and Gabriel Zucman. 2023. "The missing profits of nations." *The Review of Economic Studies* 90 (3):1499–1534.

Weidner, Martin and Thomas Zylkin. 2021. "Bias and consistency in three-way gravity models." *Journal of International Economics* 132:103513.

Yotov, Yoto V, Roberta Piermartini, José-Antonio Monteiro, and Mario Larch. 2016. *An advanced guide to trade policy analysis: The structural gravity model.* World Trade Organization Geneva.

Appendix

A Data Construction

This appendix provides details on the construction of the patent ownership dataset. We combine PATSTAT Global (Spring 2024) with ktMINE transaction data to track the ownership history of individual patents over time. Following De Rassenfosse, Kozak, and Seliger (2021), we identify primary patents within each family and assign inventors based on the lead applicant.

To construct ownership chains, we link patent-level records across datasets and track changes in assignees over time. We distinguish between intra-firm and inter-firm transactions using name-matching algorithms based on Kalyani, Bloom, Carvalho, Hassan, Lerner, and Tahoun (2021), supplemented with additional cleaning procedures to reduce false matches.

We further incorporate country-level variables, including statutory corporate income tax rates, patent box regimes, and measures of IPR protection. Additional details on variable construction, matching procedures, and robustness checks are provided below.

We create a comprehensive dataset of patent agreements. We define patent agreements broadly as instances where one entity transfers its interest—either partially or entirely—in a patent or a patent application to another entity. These transfers may occur before or after the patent has been granted by the respective patent authority as they reflect only the ability to use the courts to enforce the patent protections if applicable. The primary data source used for constructing this dataset is ktMINE, which collects data on patent assignments from the United States Patent and Trademark Office (USPTO) and the International Patent Documentation (INPADOC) maintained by the European Patent Office (EPO). Given the structural differences in these data sources, different methods of data cleaning were required to achieve a uniform dataset. In addition, variables from the EPO's PATSTAT database were incorporated to supplement missing or incomplete information.

We begin by processing the data from the USPTO. The ktMINE dataset provides details on assignments recorded by the USPTO, covering patents where at least one member of the patent family was filed in the United States. Since the US is one of the most popular destinations for patents filed internationally, this dataset encompasses a wide array of patent assignments dating back to the mid-20th century.

Patent datasets often suffer from inconsistencies due to errors and variations in the way forms are filled out. For instance, the name of a company on one patent application may be listed as “Patent Holdings Incorporated,” while on another it might appear as “Patent Holdings Inc.” Despite both names referring to the same entity, such discrepancies result in multiple listings for the same company. Moreover, mistakes like writing “Patent Holding Inc.” instead of “Patent Holdings Inc.” can introduce further noise. While ktMINE provides a field for “ultimate company name” to link subsidiaries to their parent companies, these inconsistencies still lead to incomplete corporate trees. To address this issue and ensure a robust comparison between intra-firm and arm’s-length patent transfers, we employ multiple methods, including a string matching algorithm based on Kalyani, Bloom, Carvalho, Hassan, Lerner, and Tahoun (2021) and a system that utilizes inventor information to infer intra-firm transfers when company data is incomplete.

A.1 String Matching

Our first step in constructing the dataset was to clean and standardize the company names provided by ktMINE. Each company name in the dataset, along with its associated “ultimate company name” (where applicable), is linked to location information. However, as mentioned earlier, this data is prone to inconsistencies, and companies that should logically belong to the same corporate tree might not appear as such in the raw data.

To address this, we use a string matching algorithm to identify and group company names that are likely to refer to the same entity. While this is a useful approach, it must be applied carefully to avoid over-linking unrelated companies. For instance, while “Microsoft Computing Technologies LLC” and “Microsoft Software LLC” should be linked, “Toyota Motor Company” and “Ford Motor Company” should not, despite sharing a large portion of their string characters. The methodology developed in Kalyani, Bloom, Carvalho, Hassan, Lerner, and Tahoun (2021) is designed to handle such complexities.

The string matching process begins by breaking each company name into four-character segments called “grams.” Common terms such as “the,” “llc,” “inc” are excluded to prevent spurious matches. A vector of these grams is created for each company, and the similarity between companies is measured using Term Frequency-Inverse Document Frequency (TF-IDF) scores. This approach allows us to quantify the degree of similarity between company names, with higher TF-IDF scores indicating greater likelihood of the names referring to the same corporate entity. We then group companies with TF-IDF scores above a certain threshold (set at 75% in our main analysis) under the same corporate tree. Additionally, if Company A is linked to Company B, and Company B is linked to Company C, all three are grouped together.

A.2 Data Construction

After cleaning the company names, we merge the USPTO transfer data with variables from PATSTAT. This merging process is performed in multiple stages due to the absence of clean, unique identifiers across the datasets. The primary key for the merge is the DOCDB Document Number, a unique identifier consisting of the patent's publication authority code, publication number, and publication kind. For patents where the DOCDB Document Number is unavailable, we use the application number as the secondary key.

To prevent double-counting and ensure consistency in tracking ownership chains, we analyze patent assignments at the family level. Patent families consist of multiple applications for the same invention filed in different jurisdictions, and assignments generally involve entire families rather than individual applications. Counting at the application level would lead to inflated assignment counts and inaccurate ownership chains. To retain information about the size of the patent family, we include a variable from PATSTAT that records the number of family members. We removed approximately 0.02% of the sample (63 patent families) that exhibited inconsistencies in family member movements during the assignment process, ensuring that the remaining data is reliable.

We organize the data so that each family has a single observation for each time it is assigned, whether or not the assignment involves a change in ownership. Each observation includes details such as whether the assignment constitutes an ownership transfer, the assignor and assignee companies, the countries of both assignor and assignee, and the year of the assignment. Each assignment is sequentially numbered, with separate counts for total assignments and ownership transfers only. Additionally, we fill missing country information based on previous assignments and the country of the first applicant in the family.

Finally, we merge the patent data with gravity variables from the CEPII database to include control variables in our regression analysis. Our final dataset, based on USPTO records, consists of 194,151 patent family assignments across 54 countries between 1995 and 2020, with 39% identified as intra-firm assignments.

Table A1: List of Countries

Country	**Country**
Antigua and Barbuda	Luxembourg
Australia	Latvia
Austria	Macau
Belgium	Mexico
Bulgaria	Malta
Bahamas	Mauritius
Belize	Netherlands
Bermuda	Norway
Brazil	Panama
Barbados	Poland
Canada	Portugal
Switzerland	Russia
China	Singapore
Cayman Islands	Sweden
Cyprus	Seychelles
Czech Republic	Turkey
Germany	United States
Denmark	Vietnam
Spain	Vanuatu
Estonia	Samoa
Finland	
France	
United Kingdom	
Greece	
Hong Kong	
Croatia	
Hungary	
India	

A.3 Intra-firm Assignments

To accurately capture intra-firm assignments, we take several steps. The first is the most straightforward: if the parent company of the assignor matches that of the assignee, the assignment is classified as intra-firm. However, this method alone does not account for all intra-firm transfers, as the company name data may still be incomplete or inaccurate, even after cleaning.

To supplement this, we use ktMINE's "is_ownership_transfer" and "is_inventor" variables. The former indicates whether an assignment involves a change in the official owner of the patent, while the latter identifies whether the assignor is the original inventor of the patent. By examining sequential assignments within a patent family, we can infer intra-firm transfers even when the company data is missing or unclear. For instance, if an ownership change occurs but the assignor is still listed as the inventor, we can reasonably infer that the previous transfer was intra-firm. We adjust our data accordingly, marking such transfers as intra-firm.

A.4 Summary Statistics

Our constructed dataset reveals several important patterns in patent assignment behavior, both across countries and over time. Tables A2 and A3 present the top 10 countries by total patent assignments for the periods before 2010 and from 2010 to present, respectively.

Table A2: Top 10 Countries by Total Patent Assignments (Pre-2010)

Country	**Arm's Length**	**Intra-Firm**	**Total**	**% Intra-Firm**
SGP	1,489	3,753	5,242	71.6%
USA	16,339	39,815	56,154	70.9%
CAN	2,142	3,690	5,832	63.3%
ISR	544	919	1,463	62.8%
CHL	1,726	2,105	3,831	54.9%
KOR	2,399	2,896	5,295	54.7%
GBR	2,168	2,603	4,771	54.6%
AUS	660	680	1,340	50.7%
FRA	1,561	1,468	3,029	48.5%
JPN	16,438	10,836	27,274	39.7%

Examining the pre-2010 period (Table A2), we find that Singapore and the United States exhibit the highest proportions of intra-firm transfers, with 71.6% and 70.9% of all assignments, respectively, occurring within corporate structures. This suggests that firms in these countries may be engaging in strategic management of intellectual property, possibly motivated by tax optimization. In contrast, Japan, despite having the second-highest total number of patent assignments, exhibits a much lower intra-firm transfer rate of 39.7%, indicating a more diverse market with a greater prevalence of arm's-length transactions.

Table A3: Top 10 Countries by Total Patent Assignments (2010 to Present)

Country	Arm's Length	Intra-Firm	Total	% Intra-Firm
BRB	367	1,295	1,662	77.9%
SGP	1,670	5,168	6,838	75.6%
LUX	1,472	3,333	4,805	69.4%
CAN	2,412	5,273	7,685	68.6%
KOR	3,238	6,041	9,279	65.1%
BMU	1,237	2,135	3,372	63.3%
USA	22,779	18,908	41,687	45.4%
CHL	5,274	4,038	9,312	43.4%
CHE	7,014	5,039	12,053	41.8%
GBR	2,005	1,425	3,430	41.5%

In the post-2010 period (Table A3), the landscape shifts notably. Barbados and Singapore emerge as leaders in intra-firm transfers, with rates of 77.9% and 75.6%, respectively. This increase may be linked to favorable tax regimes in these jurisdictions. Notably, the United States sees a significant decline in intra-firm transfers, dropping from 70.9% to 45.4%, potentially due to changes in US tax laws or corporate strategies. Meanwhile, countries such as Luxembourg and Bermuda enter the top ranks, suggesting their growing importance as hubs for intellectual property management.

Tables A4 and A5 present intra-firm transfers per million people, offering insights into the relative intensity of intra-firm patent activity across countries.

Table A4: Intra-firm Transfers per Million People (Pre-2010)

Country	Population (millions)	Intra-firm	Intra-firm per million people
SGP	4.372	3,753	858.406
CHE	7.439	1,911	256.890
USA	296.751	39,815	134.170
CHL	16.431	2,105	128.108
CAN	32.235	3,690	114.470
JPN	127.622	10,836	84.907
KOR	47.997	2,896	60.338
GBR	60.500	2,603	43.025
DEU	77.950	3,085	39.577
FRA	63.575	1,468	23.091

When we analyze intra-firm transfers on a per capita basis, smaller jurisdictions emerge as major players. Singapore leads in both periods, with particularly high per capita intra-firm transfers. In the post-2010 period, countries like the Cayman Islands and Bermuda exhibit extraordinarily high per capita rates, further indicating their role as patent havens. The Cayman Islands, for instance, shows an astounding 37,389 intra-firm transfers per million people, followed closely by Bermuda with 33,071. These figures are orders of magnitude higher than

Table A5: Intra-firm Transfers per Million People (2010 to Present)

Country	**Population (millions)**	**Intra-firm**	**Intra-firm per million people**
CYM	0.060	2,261	37,389.021
BMU	0.065	2,135	33,071.018
LUX	0.540	3,333	6,169.934
BRB	0.284	1,295	4,566.971
SGP	5.446	5,168	948.890
CHE	8.122	5,039	620.443
IRL	4.605	2,076	450.774
CHL	17.714	4,038	227.951
CAN	34.963	5,273	150.819
KOR	49.937	6,041	120.971

those of larger economies, highlighting the outsized role these small jurisdictions play in global intellectual property strategies.

In contrast, larger economies such as the United States and Japan do not feature prominently in the per capita rankings, despite their high absolute number of assignments.

B List of countries and classification

Tax Havens include Anguilla (AIA), The Bahamas (BHS), Belize (BLZ), Bermuda (BMU), Barbados (BRB), Cayman Islands (CYM), Cyprus (CYP), Gibraltar (GIB), Hong Kong (HKG), Saint Kitts and Nevis (KNA), Lebanon (LBN), Macao (MAC), Malta (MLT), Mauritius (MUS), Panama (PAN), Seychelles (SYC), Turks and Caicos Islands (TCA), British Virgin Islands (VGB), and Samoa (WSM).

Low Tax Countries comprise Belgium, Switzerland, Ireland, Luxembourg, the Netherlands, and Singapore (SGP).

Other encompass Australia, Austria, Bulgaria, Brazil, Canada, China, Czech Republic, Germany, Denmark, Spain, Estonia, Finland, France, United Kingdom, Greece, Croatia, Hungary, India, Italy, Japan, South Korea, Latvia, Mexico, Norway, Poland, Portugal, Romania, Russia, Sweden, Turkey, and Vietnam.

The United States (USA) is treated separately as the **HOME** country in the quantitative analysis.

The countries with missing values for tax effectiveness are Bermuda (BMU), Cayman Islands (CYM), Gibraltar (GIB), Saint Kitts and Nevis (KNA), Macao (MAC), Turks and Caicos Islands (TCA), and British Virgin Islands (VGB).

C Robustness Analysis: Patent Box Regimes

In the main text, we show that patent ownership responds strongly to international tax differentials, particularly within multinational firms. We also document that low-tax jurisdictions attract substantial patent inflows and royalty income despite often having limited local innovation activity. These findings suggest that tax incentives play an important role in shaping the global allocation of intangible assets.

In this appendix, we examine whether preferential tax regimes specifically targeting intellectual property income—so-called patent box regimes—affect patent location decisions differently from broader statutory corporate tax rates. Patent boxes are tax incentives designed to encourage local R&D and innovation by offering preferential tax treatment on income generated from intellectual property. While their stated objective is to stimulate domestic innovative activity, they may also create incentives for multinational firms to relocate patent ownership across jurisdictions for tax purposes.[6]

Studying patent boxes separately is useful because these regimes target the taxation of intangible income directly, rather than corporate income more broadly. As a result, they provide a particularly relevant setting for understanding how multinational firms respond to tax incentives tied specifically to intellectual property ownership.

Patent boxes, which offer preferential tax treatment for intellectual property income, appear to influence patent transfer decisions differently across our three study periods.

During 1995-2006, the relationship between patent box tax differentials and transfers mirrors the pattern we observed with statutory rates, but with some important differences. The coefficient for intra-firm transfers (7.975, $p < 0.05$) is statistically significant, while the total transfers coefficient (4.526) remains positive but imprecisely estimated. This suggests that even before patent boxes became widespread, firms were responsive to specialized intellectual property tax incentives when available.

[6] The introduction of a patent box may alter the dynamics of where firms choose to locate intellectual property. Countries that have introduced patent box regimes during our sample period include Belgium, Cyprus, France, Greece, Hungary, Italy, Liechtenstein, Luxembourg, Malta, the Netherlands, Portugal, Ireland, South Korea, Spain, Switzerland, Turkey, and the United Kingdom.

Table A1: Patent sales and patent box taxation by period

	$Patents_{injt}$								
	1995-2006			2007-2012			2013-2020		
	Total	Intra	Inter	Total	Intra	Inter	Total	Intra	Inter
$\tau_o - \tau_d^{PB}$	4.526	7.975*	1.758	2.223	4.808*	0.371	1.730	5.722	1.040
	(2.981)	(3.934)	(2.679)	(1.186)	(2.006)	(1.769)	(3.291)	(6.048)	(4.939)
N	279,931	150,182	186,999	343,777	180,366	249,128	211,661	93,535	149,241
pseudo R^2	0.569	0.582	0.550	0.504	0.507	0.503	0.561	0.612	0.556
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Origin FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Destination FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Pair FE	YES	YES	YES	YES	YES	YES	YES	YES	YES

Standard errors in parentheses clustered by firm

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes. Regression done with PPML methods. The dependent variable is the count of patent transfers. $\tau_o - \tau_d^{PB}$ represents the difference between origin country's corporate tax rate and destination country's patent box rate. All specifications include origin, destination, time, firm, and origin-destination pair fixed effects. Controls include origin and destination GDP and GDP per capita.

The 2007-2012 period shows an interesting departure from our statutory rate results. While still positive and significant for intra-firm transfers (4.808, $p < 0.05$), the magnitude is notably smaller than for statutory rate differentials during the same period. The coefficient for total transfers (2.223) is also considerably reduced. This pattern may reflect the limited adoption of patent box regimes during this period— while some countries like Netherlands and Belgium had introduced such schemes, they were not yet widespread.

The post-2013 period shows coefficients becoming negative but statistically insignificant across all specifications (total: -2.716, intra: -1.879, inter: -3.294). This differs from the statutory rate results in magnitude but maintains the pattern of reduced predictive power. Several factors may explain this finding.

First, patent box regimes proliferated significantly during this period, with many countries adopting increasingly similar preferential rates. This convergence in effective tax rates on patent income may have reduced the tax motivation for transfers. Second, post-BEPS reforms often required companies to demonstrate real business operations and R&D activities to qualify for patent box benefits, making simple transfer strategies less effective. Third, firms may have adjusted their intellectual property planning to focus on development location rather than post-development transfers, as many patent box regimes began requiring substantial research activity in the benefiting jurisdiction.

The consistently smaller coefficients for patent box differentials compared to statutory rate differentials suggest that firms may view permanent statutory rate differences as more reliable for long-term planning than potentially reversible preferential regimes. This interpretation aligns with the observation that patent transfers, being relatively costly to reverse, require stable tax incentives to justify their transaction costs.

D Royalty Data Construction

We construct bilateral royalty flows using data from the IMF Balance of Payments Statistics (BOPS) and the OECD International Trade in Services databases over 2005–2021. These data report cross-border payments for the use of intellectual property, including patents, trademarks, and licensing fees.

For each destination country i and year t, total royalty inflows are computed as

$$R_{it} = \sum_{n} R_{nit},$$

where R_{nit} denotes royalty payments from origin country n to destination country i. To facilitate comparisons across countries of different size, we normalize royalty inflows by GDP:

$$\frac{R_{it}}{GDP_{it}}.$$

We classify destination countries into three groups: tax havens, low-tax countries, and other countries, following Tørsløv, Wier, and Zucman (2023). Tax havens include jurisdictions such as Bermuda, the Cayman Islands, and Barbados, while low-tax countries include Ireland, Luxembourg, the Netherlands, Switzerland, Belgium, and Singapore.

For each group g, we compute its share of global royalty inflows:

$$\text{Share}_{gt} = \frac{\sum_{i \in g} R_{it}}{\sum_{i} R_{it}}.$$

To highlight changes over time, we also construct an index relative to 2005:

$$\text{Index}_{gt} = \frac{\text{Share}_{gt}}{\text{Share}_{g,2005}} \times 100.$$

Because royalty income is recorded where intellectual property is owned, these measures provide a proxy for the global location of intangible income and complement the analysis of patent ownership transfers.

E The Mechanism: Two-country comparative statics

This section uses a two-country version of the model to illustrate the allocation of patents across organizational forms. We denote the domestic (origin) country as i and the foreign (destination) country as j, and examine how variation in tax parameters maps into patent location decisions.

We focus on the role of domestic versus foreign taxation, as well as the factors that are intrinsic to profit shifting. Countries only differ in their tax rates: We fix the foreign income tax, τ_j, and compute optimal strategies for varying domestic tax rates, $\{\tau_i, \tau_i^k\}$. Then, we recompute the decisions for different values of the transfer-price discount, σ, and the foreign subsidiary setup cost, κ to study the interaction between tax regimes and profit-shifting incentives. The remaining parameters are held at their baseline values, as shown on Table A1.[7] Below, we also show how our results change when we vary the seller's bargaining power, θ.

Table A1: Baseline Parameter Values

Parameter	Description	Value/Range
τ_i	Tax on patent income, domestic	[0%,35%]
τ_i^k	Tax on capital gains, domestic	[0%,35%]
τ_j	Tax on patent income, foreign	17.5%
σ	Transfer-price discount	[0, 1]
κ	Subsidiary setup cost	[0,100]
r	Interest rate	5%
ϕ	Probability of patent imitation	10%
ε	Patent scrap value	0.02
y	Cash flow	100
ρ	Patent owner's bargaining power when licensing	0.5
θ	Patent owner's bargaining power when selling	0.5

Given our assumptions, the pre-tax value of a patent is the same across countries, $W_i = W = \frac{2(1+r)\rho y+\phi\varepsilon}{r+\phi}$, while the pre-tax selling price of a patent depends only on the tax rates in the country of origin, $P_{ij} = P_i = \frac{W}{1-\tau_i^k}\left[1 - \frac{\tau_i+\tau_i^k}{2}\right]$.

The regions displayed in Figure A1 correspond to the optimal decision for given tax rates and profit-shifting parameters. Blue regions represent **Licensing**, where the patent remains in the home jurisdiction and generates after-tax revenues $(1-\tau_i)W$ from global licensing. Green regions represent **Selling**, in which the patent is transferred to another entity at the negotiated

[7] Note that the corporate income tax rate, tau_j^c, is irrelevant for our calculations as it only affects the income of licensees.

price P, with the seller paying capital-gains tax τ_i^k and retaining $(1-\tau_i^k)P_i$. Red regions represent **Profit Shifting**, where the patent is transferred to a foreign affiliate at a transfer price $T_i = \sigma P_i$. In this case, the headquarters records $(1-\tau_i^k)(T_i-\kappa)$, while the affiliate obtains $(1-\tau_j)(W-T_i)$. Total after-tax income generated by profit shifting is the sum of the headquarters' and affiliate's after-tax incomes.

Figure A1: Strategy Heatmaps for Selected σ and κ

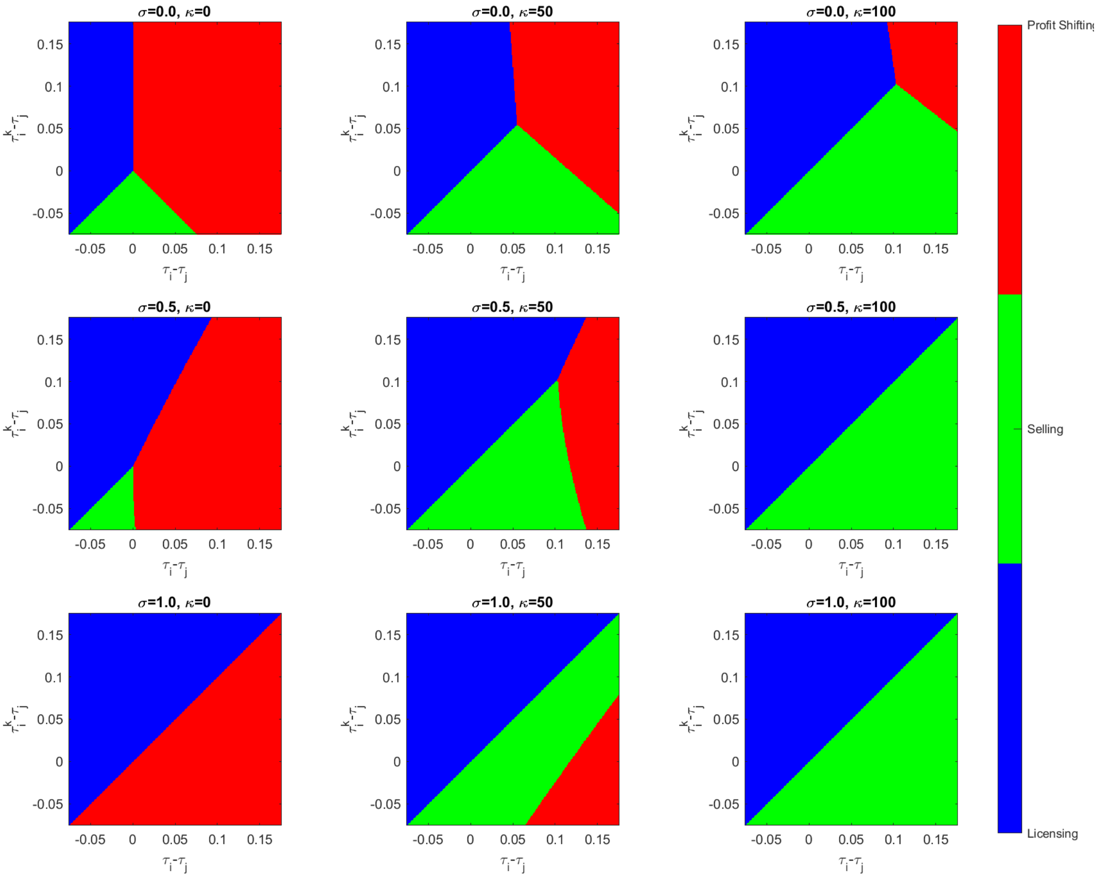

Note. Horizontal axis: τ_i (domestic patent-income tax rate) minus τ_j (foreign patent-income tax rate). Vertical axis: τ_i^k (domestic capital-gains tax rate) minus τ_j (foreign patent-income tax rate). Rows vary the transfer-price discount, $\sigma \in [0, 0.5, 1]$, and columns vary the foreign subsidiary setup cost, $\kappa \in [0, 50, 100]$. Blue = licensing, green = selling, red = profit shifting.

We fix the destination royalty tax at $\tau_j = 0.175$ and vary the home-country royalty tax τ_i and capital gains tax τ_i^k. On the horizontal axis of each chart we have the difference between the domestic patent-income tax, τ_i and the foreign income tax rate, τ_j, while on the vertical axis we have the difference between the capital-gains tax, τ_i^k, and the foreign income tax rate, τ_j. The diagonal line on each chart marks the boundary at which the capital-gains tax equals the tax on licensing income. Charts differ in the values of the transfer-price discount, σ, and the fixed affiliate setup cost, κ. The three rows of charts correspond to different values of the transfer-price parameter σ, and the columns to different values of the affiliate setup cost κ.

Licensing vs. selling. The choice between licensing and selling is given by the maximum of $(1-\tau_i)W$ and $(1-\tau_i^k)P_i$. Given the common pre-tax components in both alternatives (and the normalization $P_i = W/2$), the decision boils down to comparing domestic tax rates, specifically τ_i and $\frac{\tau_i+\tau_i^k}{2}$. Therefore, licensing is preferred to selling when $\tau_i < \tau_i^k$, while selling is preferred to licensing when $\tau_i > \tau_i^k$. In the absence of profit-shifting, licensing would correspond to the area above the 45-degree line on the charts in Figure A1, while selling would appear below the 45-degree line. As we vary σ and κ, we change the value of profit shifting and identify regions where this option dominates licensing and selling, as described below.

Starting point: $\sigma = 0$, $\kappa = 0$ (top-left). With $\sigma = \kappa = 0$, equation (12) simplifies to $\Pi_i = (1-\tau_j)W$, so the profit-shifting payoff is simply "book royalties at τ_j." The value of profit-shifting is independent of the domestic tax rates and so, is constant along the vertical and horizontal axes.

The choice of licensing, selling and profit shifting is given by the maximum of $(1-\tau_i)W$, $(1-\tau_i^k)P_i$ and $(1-\tau_j)W$. Given the common terms in all three alternatives, the decision boils down to comparing tax rates, i.e.,

$$\min\left\{\tau_i, \frac{\tau_i+\tau_i^k}{2}, \tau_j\right\}$$

When the are no differences in tax rates, $\tau_i = \tau_i^k = \tau_j$, all alternatives offer the same value. As derived above, licensing dominates selling when $\tau_i < \tau^k$ and vice versa. Profit shifting is the best option when the foreign tax rate, τ_j is lower than both the domestic patent-income tax rate, τ_i, and the average of the domestic patent-income and capital-gains tax rates, $\frac{\tau_i+\tau_i^k}{2}$.

Raising σ with $\kappa = 0$. When $\sigma > 0$, equation (12) acquires an additional term,

$$\Pi_i = \left[1-\tau_j+\sigma\left(\frac{\tau_j-\tau_i^k}{1-\tau_i^k}\right)\left(1-\frac{\tau_i+\tau_i^k}{2}\right)\right]W$$

We refer to this component as the *manipulation term*, as it reflects the gains from transfer-price manipulation. The sign of this term depends on the difference between the foreign patent-income tax rate, τ_j, and the domestic capital-gains tax rate, τ_i^k. When $\tau_i^k < \tau_j$ (negative values on the vertical axis), the manipulation term is positive. As a result, increasing σ, which brings the transfer price closer to arm's-length pricing, raises the profit-shifting payoff and expands the red region into lower values of τ_i^k, where selling (the green area) was previously preferred.

When $\tau_i^k > \tau_j$ (positive values on the vertical axis), the manipulation term is negative, so the profitability of shifting profits declines, the red region recedes and licensing (the blue area) becomes more prominent.

As enforcement becomes stricter, i.e., σ increases, profit-shifting gains ground over selling when $\tau_i^k < \tau_j$ while licensing gains ground over profit shifting when $\tau_i^k > \tau_j$. When $\sigma = 1$, the value equations imply that profit shifting is preferred to selling if $\tau_i > \tau_i^k$. That is, in the region where selling dominates licensing (below the 45-degree line in the chart), profit shifting dominates selling. As a result, selling is never preferred to the alternative strategies. Overall, when $\sigma = 1$ and $\kappa = 0$, licensing is preferred above the 45-degree line, while profit shifting is preferred below the 45-degree line.

Effect of κ. Increasing the setup cost of establishing an affiliate uniformly lowers the profit-shifting payoff by the term $-(1-\tau_i^k)\kappa$ in equation (12). As κ rises, the red region contracts and the released space is absorbed by licensing (blue) when $\tau_i^k > \tau_i$ and by selling (green) when $\tau_i^k < \tau_i$. This reallocation of regions is evident when moving to the right along each row of Figure A1. For sufficiently large κ, licensing and selling become the only preferred alternatives.

Case $\theta = 0$. When the seller has no bargaining power, the selling margin effectively disappears. The outcome is a stark trade-off between licensing (blue) and profit shifting (red), with profit shifting dominating whenever the tax difference favors relocating income. As shown in Figure A2, the green region is absent across all parameter values.

Case $\theta = 1$. When the seller has full bargaining power, the selling margin becomes more attractive. In Figure A3, the green region expands relative to $\theta = 0.5$, particularly when $\tau_i^k < \tau_i$. Profit shifting is present when $\sigma < 1$, but vanishes completely at $\sigma = 1$, leaving a clean division between licensing and selling along the $\tau_i = \tau_i^k$ diagonal.

The cases $\theta = 0$ and $\theta = 1$ highlight the role of bargaining power. With $\theta = 0$, the seller has no bargaining power and the sale margin is essentially inactive, so the choice is between licensing and profit shifting, with red regions dominating whenever the tax difference favors shifting. With $\theta = 1$, the seller captures the buyer's full value and the selling region (green) is more prominent, while profit shifting disappears entirely under arm's-length enforcement

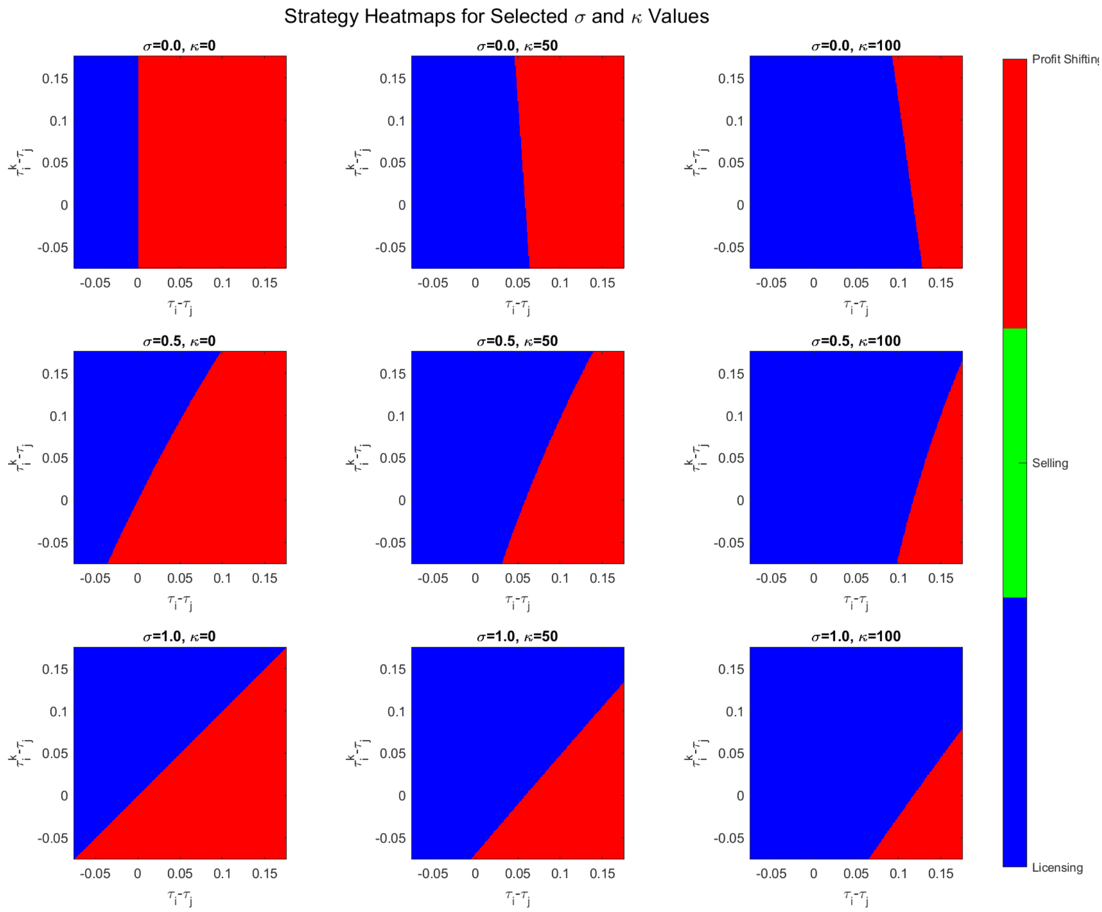

Figure A2: Strategy heatmaps for $\theta = 0$, fixing $\tau_j = 0.175$. Blue = licensing, red = profit shifting. The selling region is absent.

($\sigma = 1$). The intermediate case $\theta = 0.5$, discussed above, falls between these extremes: the green region contracts relative to $\theta = 1$, while the red region is more pervasive than under $\theta = 1$ but less than under $\theta = 0$.

F Counterfactual Analysis: Calibration

This appendix describes the calibration strategy used in the counterfactual analysis. We calibrate the model to match key features of the global allocation of patent ownership and royalty income, with particular emphasis on the concentration of intangible income in low-tax jurisdictions. The calibration targets observed cross-country differences in tax rates, royalty shares, and patent relocations, and links the empirical elasticities estimated in the data to the structural parameters governing firms' organizational choices.

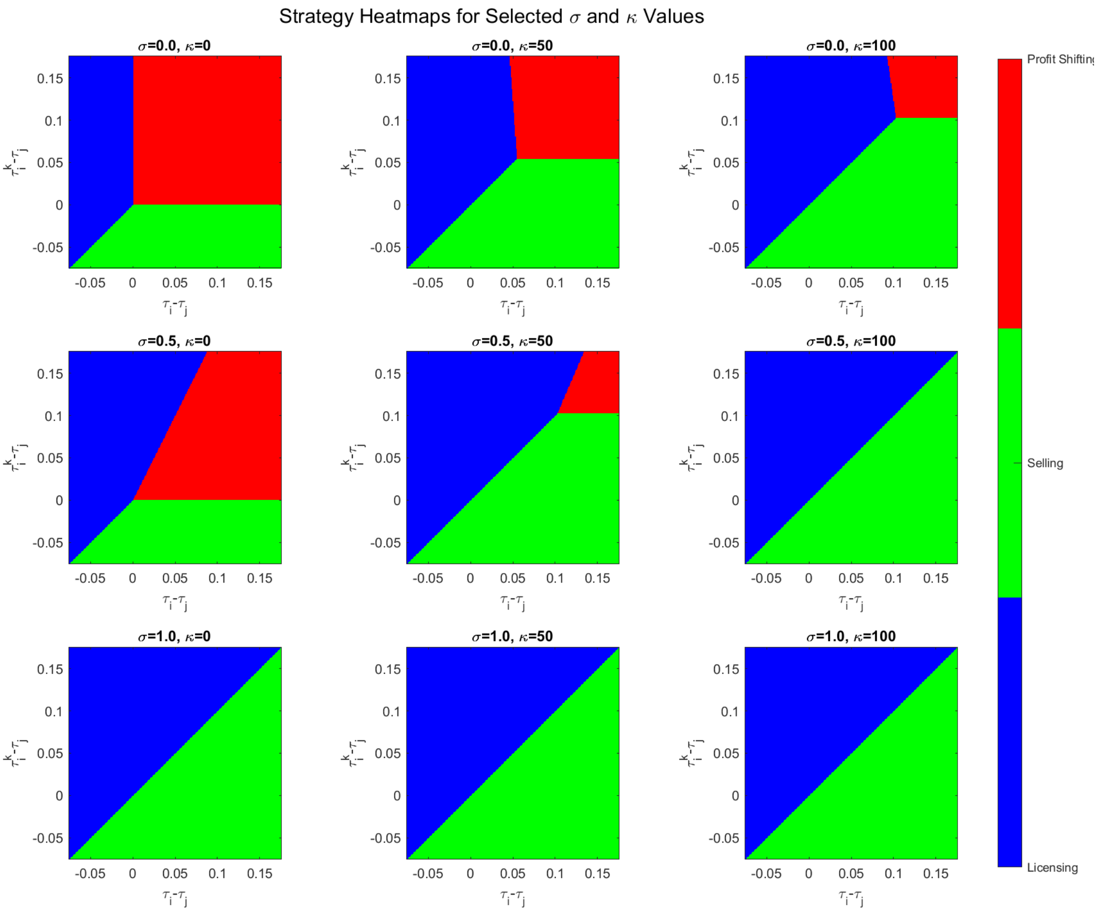

Figure A3: Strategy heatmaps for $\theta = 1$, fixing $\tau_j = 0.175$. Blue = licensing, green = selling, red = profit shifting. Profit shifting disappears completely when $\sigma = 1$.

F.1 Calibration of the Transfer-Pricing Parameter σ

We calibrate the transfer-pricing parameter σ using data on international patent transfers and R&D receipts. The empirical exercise combines two main sources:

(i) **Patent transfer flows:** Bilateral counts of patent reassignments are taken from the *ktMINE Transactions Database*, which provides annual flows of patent transfers by country and type of relationship (affiliated vs. unaffiliated) for 2007–2017.[8]

(ii) **R&D service receipts:** Data on receipts for R&D and intellectual property (IP) services are obtained from the *US Bureau of Economic Analysis (BEA)*. We use the series *"Research and Development Services"* from *Table 2.2, US Trade in Services by Type of Service and by Country or Affiliation*, which includes contract R&D as well as *"outright*

[8]Source: ktMINE, *Intellectual Property Transactions Data*, accessed via ktMINE Analytics, 2024.

sales of the outcomes of R&D, such as patents, copyrights, and industrial processes".[9]

For each country i and year t, we observe the number of affiliated (intra-firm) and unaffiliated (arm's-length) patent transfers, denoted by $direct\_intra_{it}$ and $direct\_inter_{it}$, respectively. The total number of transfers is $direct\_total_{it} = direct\_intra_{it} + direct\_inter_{it}$, and the affiliated share of transfers is:

$$share^{A}_{it} = \frac{direct\_intra_{it}}{direct\_total_{it}}.$$

To combine patent flows with BEA receipts, we proceed in two steps. First, we compute affiliated shares at the country level and then average them across countries within tax-status groups and years. Let $share^{A,g}_{it}$ denote this group-level affiliated share. We use this smoothed share to allocate total R&D receipts into affiliated and unaffiliated components:

$$RD^{A}_{it} = RD_{it} \times share^{A,g}_{it}, \qquad RD^{U}_{it} = RD_{it} \times (1 - share^{A,g}_{it}),$$

where RD_{it} denotes total R&D receipts.

We then compute the implied value per patent transferred for each category:

$$ratio^{A}_{it} = \frac{RD^{A}_{it}}{direct\_intra_{it}}, \qquad ratio^{U}_{it} = \frac{RD^{U}_{it}}{direct\_inter_{it}}.$$

The transfer-pricing parameter is defined as the ratio:

$$\sigma_{it} = \frac{ratio^{A}_{it}}{ratio^{U}_{it}}.$$

Substituting the definitions above, RD_{it} cancels out of the ratio. As a result, σ_{it} depends on the comparison between the group-level allocation of R&D receipts and the country-level composition of patent transfers. Intuitively, σ_{it} captures whether affiliated transactions are underrepresented in value relative to their frequency in observed transfer flows.

We therefore interpret σ_{it} as a reduced-form proxy for the relative discount at which headquarters transact with affiliates. Values of $\sigma_{it} < 1$ indicate that affiliated transfers are associated with lower effective valuations relative to arm's-length transactions, consistent with transfer-pricing behavior.

Finally, we restrict attention to the period 2012–2016 and aggregate σ_{it} in two steps. First,

[9] Source: US Bureau of Economic Analysis (BEA), *International Transactions Accounts*, Table 2.2 (latest release, 2025).

for each country and tax-status group, we take the minimum value of σ_{it} across years, which places greater weight on episodes with relatively low implied valuations of affiliated transfers. Second, we average these values across countries within each group to obtain the calibration targets.

We compute the average value of σ across tax regimes, following the BEA classification of countries into high-tax, low-tax, and tax-haven groups.[10] The resulting averages are summarized in Table A1.

Table A1: Estimated transfer-pricing discount factor σ by tax group

Tax group	$\sigma = ratio^A/ratio^U$	**Implied discount** $(1-\sigma)$
Havens	0.22	0.78
Low-tax	0.42	0.58
High-tax	0.66	0.34

Notes: Computed using flow data for 2012–2017. Ratios are constructed as the average value per patent transferred from the United States, distinguishing affiliated (intra-firm) and unaffiliated (inter-firm) transactions. The discount $(1-\sigma)$ represents the percentage undervaluation of affiliated transfers relative to market-based (unaffiliated) valuations. Sources: ktMINE Transactions Database (patent transfers) and US BEA International Transactions Accounts, Table 2.2 "R&D Services" by country and affiliation.

F.2 Calibration of the Pareto Distribution of Patent Values

In the model, each patent draws a value v from a common upper–tail distribution. A patent is *transferred internationally* if and only if its value exceeds a cutoff v^*_{it} that summarizes the destination i's environment (taxes, IPR/enforcement, contracting frictions). Royalties received in i are the cash flows accruing to the entity that holds the transferred IP. We map these flows to patent values via the bargaining split,

$$v = \frac{\text{royalties}}{\rho},$$

where $\rho \in (0,1)$ is the owner's bargaining weight. Hence v is the ex–post payoff of patents that *arrive* in i.

F.2.1 Data and empirical moments

We construct two empirical moments that the model pins down:

[10] Tax status follows BEA and OECD country classifications used in international service trade and balance-of-payments data. Low-tax countries include, among others, Ireland, Singapore, and the Netherlands; tax havens include Bermuda, the Cayman Islands, and other offshore financial centers.

(i) **Patent transfers (N_{it}^T).** From the *ktMINE Transactions Database* we count annual patent reassignments from the United States to destination i in year t (affiliated and unaffiliated).

(ii) **Patent applications (N_{it}).** From *INPACT-S* (LaBelle et al., 2023) we take total patent applications by destination i and year t to measure the pool of innovations.

(iii) **Royalties (royalties$_{it}$).** From OECD EBOPS we take receipts for the use of IP received in destination i from the rest of the world. These are the cash flows paid to the firm holding the transferred IP in i (typically a local affiliate), so they proxy the ex–post payoff to patents that *arrive* in i. We deflate and aggregate to an annual destination–year series consistent with N_{it}^T.

The *extensive* margin is the fraction of (potential) innovations that move to i,

$$f_{it} = \frac{N_{it}^T}{N_{it}},$$

and the *intensive* margin is the mean value among those that move,

$$\bar{v}_{it} = \frac{\text{royalties}_{it}/\rho}{N_{it}^T}.$$

F.2.2 Model mapping and identification

Assume a Pareto upper tail with shape $\alpha > 1$ and scale $v_{\min}$,

$$F(v) = 1 - \left(\tfrac{v_{\min}}{v}\right)^{\alpha}, \qquad v \geq v_{\min}.$$

The selection rule $v \geq v_{it}^*$ implies

$$f_{it} = \Pr(v \geq v_{it}^*) = \left(\tfrac{v_{\min}}{v_{it}^*}\right)^{\alpha}, \qquad E[v \mid v \geq v_{it}^*] = \frac{\alpha}{\alpha - 1}\, v_{it}^* = \bar{v}_{it}.$$

Eliminating v_{it}^* yields the estimable log–log relation

$$\ln f_{it} = \alpha\big[\ln\alpha - \ln(\alpha - 1) - \ln\bar{v}_{it}\big],$$

whose slope is $-\alpha$. Intuitively, a higher cutoff (fewer transfers) raises the average value among transferees, generating a negative trade-off between f_{it} and $\bar{v}_{it}$.

We treat α as a *global* technology parameter common to all (i,t) and fit it to the cross-section of moments by minimizing the NT–weighted squared gap

$$\min_{\alpha>1} \sum_{i,t} N_{it}^{T} \Big\{ \ln f_{it} - \alpha\big[\ln\alpha - \ln(\alpha-1) - \ln\bar{v}_{it}\big]\Big\}^2,$$

re-parameterizing $\alpha = 1+\exp(\theta)$ to enforce $\alpha > 1$. We set $v_{\min} = 1$ and benchmark $\rho = 0.5$.

The baseline (NT–weighted) estimate is

$$\widehat{\alpha} = 2.263.$$

A trimmed robustness that keeps $N_{it}^{T} \geq 5$ and winsorizes $\bar{v}_{it}$ at the 1–99th percentiles yields

$$\widehat{\alpha}^{\text{trim}} = 2.511,$$

implying a slightly thinner tail but overall stability of the calibration.

Economic interpretation and implied cutoff Given $(\widehat{\alpha}, f_{it})$ and $v_{\min} = 1$, the model implies the destination–year transfer threshold

$$v_{it}^{*} = f_{it}^{-1/\widehat{\alpha}}.$$

For reference, if $f = 10^{-3}$ then $v^{*} \approx 21$; if $f = 10^{-4}$ then $v^{*} \approx 59$. Thus only upper–tail patents are profitable to relocate internationally. In elasticity form,

$$\frac{d\ln f_{it}}{d\ln\bar{v}_{it}} = -\alpha,$$

so with $\widehat{\alpha} \approx 2.26$, selecting a thinner, higher–value slice (a 1% rise in $\bar{v}_{it}$) is associated with a $\approx 2.3\%$ fall in f_{it}.

Figure A1 shows the scatter of $\ln f_{it}$ against $\ln\bar{v}_{it}$. The cloud exhibits a slope close to $-\widehat{\alpha}$, matching the model-implied selection mapping. Figure A2 ranks destinations by mean f_{it} and reveals a highly skewed distribution: a handful of innovation hubs and selected low–tax jurisdictions account for a disproportionate share of transfers. Figure A3 contrasts havens and non–havens: the mean (median) f is **0.00111** (**0.000203**) in havens versus **0.00086** (**0.000005**) in non–havens, indicating a higher mass of transferees into low–tax destinations, consistent with a lower effective cutoff v^*.

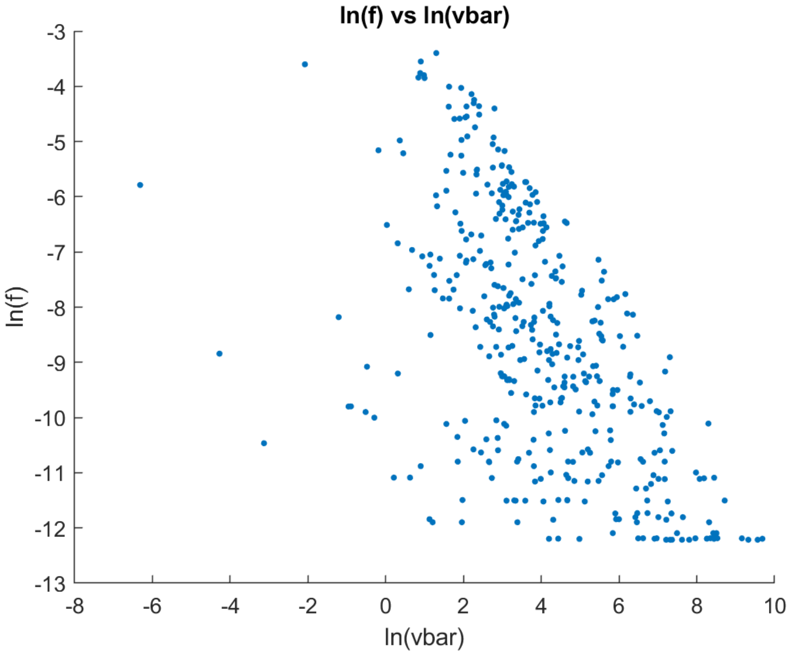

Figure A1: Diagnostic: $\ln(f)$ versus $\ln(\bar{v})$ across destination–years. The downward envelope is consistent with $\ln f = \alpha[\ln\alpha - \ln(\alpha-1) - \ln\bar{v}]$.

F.2.3 Calibrating the Fixed Cost of Profit Shifting κ_P

We recover the constant affiliate cost κ_P from the model's cutoff condition that separates "license at home" from "profit-shifting abroad," combined with the Pareto selection rule that we have already calibrated. We take the following steps:

Step 1: Cutoff from selection. With a Pareto upper tail of patent values and global tail parameter $\widehat{\alpha}$, the observed fraction of movers $f_{it} = N^T_{it}/N_{it}$ identifies the value cutoff v^*_{ijt} for destination j in year t (normalizing $v_{\min} = 1$):

$$v^*_{ijt} \;=\; f_{it}^{-1/\widehat{\alpha}}. \tag{F.1}$$

Step 2: Cutoff from the model's payoff comparison. In the heterogeneous–patent version of the model, the cutoff equates the payoff from keeping the patent domestically with the payoff from profit shifting to j. This yields

$$v^*_{ijt} \;=\; \frac{(1-\tau_{jt})\,\kappa_P}{\sigma^P_{ijt} - \sigma^L_{it}}, \tag{F.2}$$

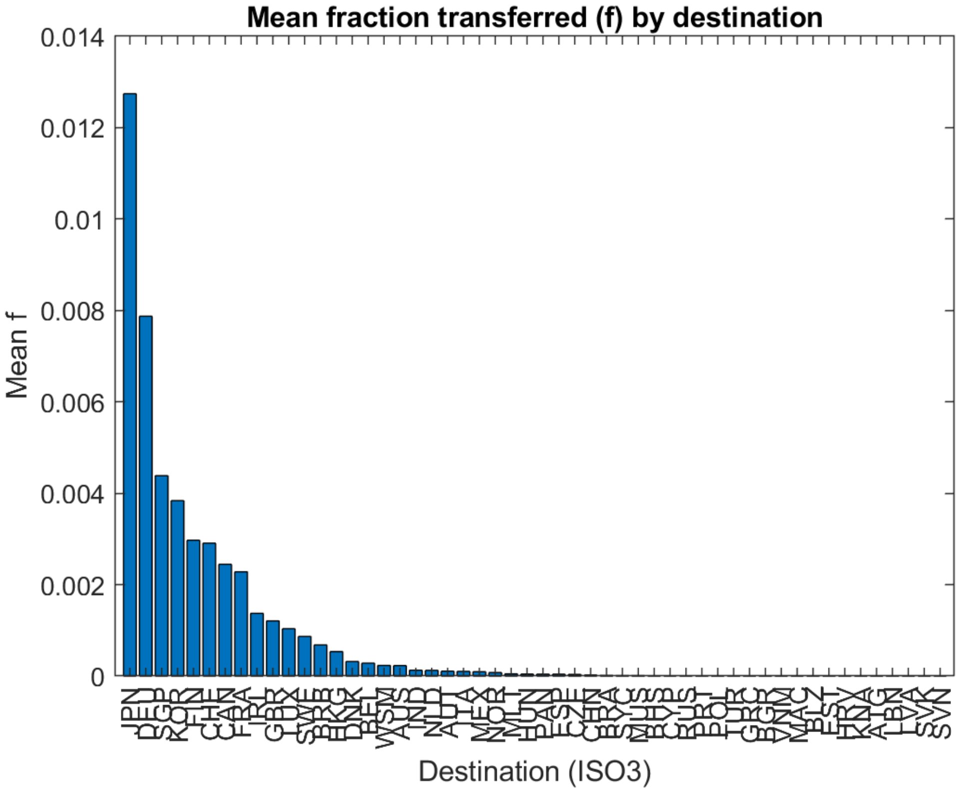

Figure A2: Mean fraction transferred f by destination (ranked). A small set of innovation hubs and selected havens account for a disproportionate share of transfers.

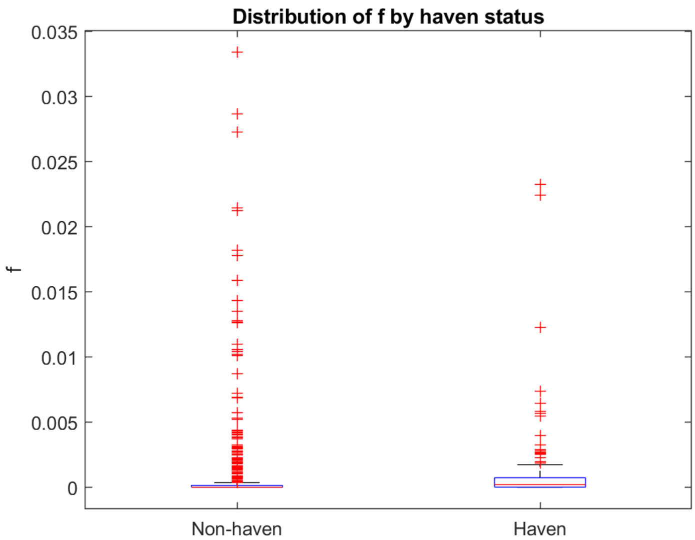

Figure A3: Distribution of f by haven status (boxplot). Havens display higher central tendency and a fatter upper tail, consistent with a lower transfer cutoff v^*.

where

$$\sigma^L_{it} \;=\; (1-\tau_{it})\,\eta_t, \tag{F.3}$$

$$\sigma^P_{ijt} \;=\; \sigma\,(1-\theta)\,\frac{\tau_{jt}-\tau^k_{it}}{1-\tau^k_{it}}\,(1-\tau_{it})\,\eta_{it} \;+\; \big(1-\tau_{jt}+\sigma\,\theta(\tau_{jt}-\tau^k_{it})\big)\,\eta_t, \tag{F.4}$$

$$\eta_{jt} \;=\; \frac{(1+r)\,\rho}{r+\phi_{jt}}, \qquad \eta_{it} \;=\; \sum_m \eta_{mt}, \qquad \eta_t \;=\; \eta_{it}. \tag{F.5}$$

Here τ_{it} and τ_{jt} are origin and destination effective tax rates, τ^k_{it} the origin tax rate on the relevant IP base, $\sigma \in (0,1]$ the transfer-pricing discount used in the policy payoff, $\theta \in [0,1]$ the weight on technological leadership, r the discount rate, ρ the bargaining share in licensing, and ϕ_{jt} the destination-specific depreciation/obsolescence parameter entering the valuation kernel η_{jt}.

Step 3: κ_P by cell and constant $\widehat{\kappa}_P$. Equating (F.1) and (F.2) gives the per–destination–year implied cost:

$$\kappa_{P,ijt} \;=\; \frac{\sigma^P_{ijt}-\sigma^L_{it}}{1-\tau_{jt}}\,f_{it}^{-1/\widehat{\alpha}}. \tag{F.6}$$

We treat κ_P as a constant across destinations and years and set it to the average.

F.3 Discrete Location Choice with Extreme-Value Shocks

Here, we provide a step-by-step derivation of the choice probabilities, cutoffs, and aggregate shares in the discrete-choice version of the model to show how the Gumbel (extreme-value) shocks and the Pareto distribution of patent values interact.

F.3.1 Setup: Patent Value and Location Options

Consider a patent of value (or "quality") $y \in [y_{\min}, \infty)$ owned by a firm in origin country i. The firm must choose where to locate the associated intangible asset. We distinguish between:

- a *home licensing option*, denoted L, and
- a set of *foreign profit-shifting destinations*, indexed by $j \in J$ (e.g. LOWTAX, HAVEN, OTHER).

For each option, we define a *deterministic payoff* that depends on y and on destination characteristics (tax rates, relocation costs, etc.).

Home licensing. We write the deterministic payoff from keeping the patent at home as

$$V_i^L(y). \tag{F.7}$$

In many applications this payoff is linear in y; for example,

$$V_i^L(y) = \sigma_L y, \tag{F.8}$$

where $\sigma_L > 0$ is the effective profit per unit of patent value when the IP is licensed at home.

Foreign profit-shifting. For a foreign destination j, we denote the deterministic payoff as

$$V_{ij}^P(y), \tag{F.9}$$

which again is often taken to be linear in y. A convenient parametric specification is

$$V_{ij}^P(y) = -(1 - \tau_j)\,\kappa_j + \sigma_{Pj}\,y, \tag{F.10}$$

where

- τ_j is the corporate tax rate in destination j,
- κ_j is an up-front relocation or setup cost (scaled appropriately),
- σ_{Pj} is the effective profit per unit of patent value in j.

F.3.2 Extreme-Value Shocks and Systematic Utility

We assume that the firm's actual utility from each option includes a random idiosyncratic component. Let ε_k denote the shock associated with option k (home L or foreign $j \in J$). The utility of choosing option k for a patent of value y is

$$U_k(y) = \bar{U}_k(y) + \varepsilon_k, \tag{F.11}$$

where $\bar{U}_k(y)$ is the *systematic* component of utility.

Destination fixed effects. For foreign destinations $j \in J$ we introduce a destination-specific shifter δ_j, so that

$$\bar{U}_{ij}(y) \equiv V_{ij}^{P}(y) + \delta_j, \qquad j \in J. \tag{F.12}$$

For home licensing we normalize

$$\bar{U}_{iL}(y) \equiv V_i^{L}(y), \qquad \delta_L = 0. \tag{F.13}$$

The parameter δ_j captures all systematic features that make destination j more or less attractive, over and above the payoff $V_{ij}^{P}(y)$ captured by the model (e.g. legal environment, IP services, unmodeled institutional features). In estimation or calibration, δ_j behaves like a destination fixed effect.

Distributional assumption. We assume that the shocks ε_k are i.i.d. across options and patents, and follow a Type-I extreme value (Gumbel) distribution with scale parameter $\mu > 0$. The cumulative distribution function (CDF) and density (pdf) are

$$F(\varepsilon) = \Pr(\varepsilon_k \leq \varepsilon) = \exp\Big(-e^{-\varepsilon/\mu}\Big), \tag{F.14}$$

$$f(\varepsilon) = \frac{1}{\mu} e^{-\varepsilon/\mu} \exp\Big(-e^{-\varepsilon/\mu}\Big). \tag{F.15}$$

The parameter μ governs the amount of idiosyncratic noise in the choice. As $\mu \to 0$, choices become nearly deterministic; larger μ implies more randomness.

F.3.3 Choice Rule and Definition of Choice Probabilities

For a given patent with value y, the firm chooses the location option with the highest utility:

$$k^*(y) = \arg \max_{k \in \{L\} \cup J} \{U_k(y)\}. \tag{F.16}$$

We focus on the *choice probability* that a patent of value y selects destination $j \in J$:

$$P_{ij}(y) \equiv \Pr\big(k^*(y) = j \,\big|\, y\big). \tag{F.17}$$

To derive $P_{ij}(y)$ analytically, we proceed in several steps.

Event "Destination j Is Chosen" The event "j is chosen" means that option j yields no lower utility than any other alternative:

$$\{k^*(y) = j\} \quad \Longleftrightarrow \quad U_{ij}(y) \geq U_{ik}(y) \quad \text{for all } k \in \{L\} \cup J. \tag{F.18}$$

In terms of systematic utility and shocks, this is

$$\bar{U}_{ij}(y) + \varepsilon_j \;\geq\; \bar{U}_{ik}(y) + \varepsilon_k \quad \forall\, k. \tag{F.19}$$

Rearranging, we can write this event as

$$\varepsilon_k \;\leq\; \varepsilon_j + \bar{U}_{ij}(y) - \bar{U}_{ik}(y) \quad \forall\, k. \tag{F.20}$$

Define the deterministic differences

$$\Delta_{jk}(y) \equiv \bar{U}_{ij}(y) - \bar{U}_{ik}(y). \tag{F.21}$$

Then "j is chosen" is equivalent to

$$\varepsilon_k \leq \varepsilon_j + \Delta_{jk}(y) \quad \forall\, k. \tag{F.22}$$

Conditioning on the Shock ε_j We can now apply the law of iterated expectations, conditioned on the realization of ε_j. The choice probability is

$$P_{ij}(y) = \Pr\big(\varepsilon_k \leq \varepsilon_j + \Delta_{jk}(y),\ \forall k\big) \tag{F.23}$$

$$= \int_{-\infty}^{\infty} \Pr\big(\varepsilon_k \leq t + \Delta_{jk}(y),\ \forall k \,\big|\, \varepsilon_j = t\big)\, f_\varepsilon(t)\, dt, \tag{F.24}$$

where f_ε is the Gumbel pdf.

Conditional on $\varepsilon_j = t$, the shocks ε_k for $k \neq j$ are independent, so

$$\Pr\big(\varepsilon_k \leq t + \Delta_{jk}(y),\ \forall k \,\big|\, \varepsilon_j = t\big) = \prod_{k \neq j} \Pr\big(\varepsilon_k \leq t + \Delta_{jk}(y)\big) \tag{F.25}$$

$$= \prod_{k \neq j} F\big(t + \Delta_{jk}(y)\big). \tag{F.26}$$

Thus,

$$P_{ij}(y) = \int_{-\infty}^{\infty} \left[\prod_{k \neq j} F\big(t + \Delta_{jk}(y)\big) \right] f_\varepsilon(t)\, dt. \tag{F.27}$$

F.3.4 Using the Gumbel Distribution to Obtain the Logit Form

We now exploit the specific functional form of the Gumbel CDF and pdf to simplify this expression.

Product of CDF Terms Recall that

$$F(x) = \exp\Big(-e^{-x/\mu}\Big). \tag{F.28}$$

For a given $k \neq j$ and value of t, we have

$$F\big(t + \Delta_{jk}(y)\big) = \exp\Big(-e^{-(t+\Delta_{jk}(y))/\mu}\Big) \tag{F.29}$$

$$= \exp\Big(-e^{-t/\mu} e^{-\Delta_{jk}(y)/\mu}\Big). \tag{F.30}$$

Therefore, the product over all $k \neq j$ is

$$\prod_{k \neq j} F\big(t + \Delta_{jk}(y)\big) = \prod_{k \neq j} \exp\Big(-e^{-t/\mu} e^{-\Delta_{jk}(y)/\mu}\Big) \tag{F.31}$$

$$= \exp\left(-e^{-t/\mu} \sum_{k \neq j} e^{-\Delta_{jk}(y)/\mu}\right). \tag{F.32}$$

Multiplying by the Gumbel pdf The Gumbel pdf is

$$f_\varepsilon(t) = \frac{1}{\mu} e^{-t/\mu} \exp\Big(-e^{-t/\mu}\Big). \tag{F.33}$$

Multiplying the product term and the pdf yields

$$\prod_{k \neq j} F\big(t + \Delta_{jk}(y)\big) f_\varepsilon(t) = \exp\left(-e^{-t/\mu} \sum_{k \neq j} e^{-\Delta_{jk}(y)/\mu}\right) \cdot \frac{1}{\mu} e^{-t/\mu} \exp\Big(-e^{-t/\mu}\Big) \tag{F.34}$$

$$= \frac{1}{\mu} e^{-t/\mu} \exp\left(-e^{-t/\mu} \Big[1 + \sum_{k \neq j} e^{-\Delta_{jk}(y)/\mu}\Big]\right). \tag{F.35}$$

Define

$$A_j(y) \equiv 1 + \sum_{k \neq j} e^{-\Delta_{jk}(y)/\mu}. \tag{F.36}$$

Then the integrand simplifies to

$$\frac{1}{\mu} e^{-t/\mu} \exp\big(- A_j(y)\, e^{-t/\mu} \big). \tag{F.37}$$

Hence

$$P_{ij}(y) = \int_{-\infty}^{\infty} \frac{1}{\mu} e^{-t/\mu} \exp\big(- A_j(y)\, e^{-t/\mu} \big)\, dt. \tag{F.38}$$

Change of Variables Let

$$z = A_j(y)\, e^{-t/\mu}. \tag{F.39}$$

Then

$$e^{-t/\mu} = \frac{z}{A_j(y)}, \qquad dt = -\mu \frac{dz}{z}. \tag{F.40}$$

As t goes from $-\infty$ to $+\infty$, $e^{-t/\mu}$ goes from $+\infty$ to 0, so z goes from $+\infty$ to 0. Substituting these expressions into the integral gives

$$\begin{aligned}
P_{ij}(y) &= \int_{+\infty}^{0} \frac{1}{\mu} \left(\frac{z}{A_j(y)} \right) \exp(-z) \left(-\mu \frac{dz}{z} \right) && \text{(F.41)} \\
&= \int_{0}^{\infty} \frac{1}{A_j(y)} e^{-z}\, dz && \text{(F.42)} \\
&= \frac{1}{A_j(y)} \int_{0}^{\infty} e^{-z}\, dz && \text{(F.43)} \\
&= \frac{1}{A_j(y)}. && \text{(F.44)}
\end{aligned}$$

Expressing $A_j(y)$ in Terms of Systematic Utilities Recall that

$$\Delta_{jk}(y) = \bar{U}_{ij}(y) - \bar{U}_{ik}(y), \tag{F.45}$$

so we can write

$$\begin{aligned}
e^{-\Delta_{jk}(y)/\mu} &= \exp\left(- \frac{\bar{U}_{ij}(y) - \bar{U}_{ik}(y)}{\mu} \right) && \text{(F.46)} \\
&= \exp\left(\frac{\bar{U}_{ik}(y)}{\mu} \right) \exp\left(- \frac{\bar{U}_{ij}(y)}{\mu} \right). && \text{(F.47)}
\end{aligned}$$

Hence

$$A_j(y) = 1 + \sum_{k \neq j} e^{-\Delta_{jk}(y)/\mu} \tag{F.48}$$

$$= 1 + \sum_{k \neq j} \exp\left(\frac{\bar{U}_{ik}(y)}{\mu}\right) \exp\left(-\frac{\bar{U}_{ij}(y)}{\mu}\right). \tag{F.49}$$

Factor out $\exp(-\bar{U}_{ij}(y)/\mu)$:

$$A_j(y) = \exp\left(-\frac{\bar{U}_{ij}(y)}{\mu}\right) \left[\exp\left(\frac{\bar{U}_{ij}(y)}{\mu}\right) + \sum_{k \neq j} \exp\left(\frac{\bar{U}_{ik}(y)}{\mu}\right)\right] \tag{F.50}$$

$$= \exp\left(-\frac{\bar{U}_{ij}(y)}{\mu}\right) \sum_{k} \exp\left(\frac{\bar{U}_{ik}(y)}{\mu}\right). \tag{F.51}$$

Therefore,

$$P_{ij}(y) = \frac{1}{A_j(y)} = \frac{\exp\left(\frac{\bar{U}_{ij}(y)}{\mu}\right)}{\sum_k \exp\left(\frac{\bar{U}_{ik}(y)}{\mu}\right)}. \tag{F.52}$$

Substituting back $\bar{U}_{ij}(y) = V_{ij}^P(y) + \delta_j$ and $\bar{U}_{iL}(y) = V_i^L(y)$ (with $\delta_L = 0$) yields the familiar logit formula:

$$\boxed{P_{ij}(y) = \frac{\exp\left(\dfrac{V_{ij}^P(y) + \delta_j}{\mu}\right)}{\exp\left(V_i^L(y)/\mu\right) + \sum_{k \in J} \exp\left(\dfrac{V_{ik}^P(y) + \delta_k}{\mu}\right)}.} \tag{F.53}$$

F.3.5 Deterministic Cutoffs and Binary Choice

To clarify the role of *cutoffs*, it is helpful to first consider a binary choice between home licensing L and a single foreign destination j.

Deterministic Payoffs and the Cutoff y_j^* Suppose the deterministic payoffs are linear in y:

$$V_i^L(y) = a_L + b_L y, \tag{F.54}$$

$$V_{ij}^P(y) = a_j + b_j y. \tag{F.55}$$

Ignoring random shocks (i.e. taking the limit $\mu \to 0$), the firm chooses the option with the higher deterministic payoff. The cutoff y_j^* at which the firm is indifferent between L and j

solves

$$a_L + b_L y_j^* = a_j + b_j y_j^*. \tag{F.56}$$

Solving for y_j^* gives

$$y_j^* = \frac{a_j - a_L}{b_L - b_j}. \tag{F.57}$$

Under the empirically relevant assumption that $b_j > b_L$ (so shifting is more attractive for high-value patents), we have

- for $y < y_j^*$, home licensing is optimal;
- for $y > y_j^*$, shifting to j is optimal.

Logit Probability as a Smooth Cutoff With extreme-value shocks ($\mu > 0$), the choice is no longer deterministic. Using the logit formula for the binary case (home vs j), we obtain

$$P_{ij}(y) = \frac{\exp\big((a_j + b_j y)/\mu\big)}{\exp\big((a_L + b_L y)/\mu\big) + \exp\big((a_j + b_j y)/\mu\big)} \tag{F.58}$$

$$= \frac{\exp\big(\{(a_j - a_L) + (b_j - b_L)y\}/\mu\big)}{1 + \exp\big(\{(a_j - a_L) + (b_j - b_L)y\}/\mu\big)}. \tag{F.59}$$

Define

$$\Delta a \equiv a_j - a_L, \qquad \Delta b \equiv b_j - b_L. \tag{F.60}$$

Then

$$P_{ij}(y) = \frac{\exp\big((\Delta a + \Delta b\, y)/\mu\big)}{1 + \exp\big((\Delta a + \Delta b\, y)/\mu\big)}. \tag{F.61}$$

Using the definition of y_j^*, one can show

$$\Delta a + \Delta b\, y = \Delta b\,(y - y_j^*), \tag{F.62}$$

so the choice probability can be written as a logistic function centered at the cutoff y_j^*:

$$\boxed{P_{ij}(y) = \frac{1}{1 + \exp\big(-\frac{\Delta b}{\mu}(y - y_j^*)\big)}.} \tag{F.63}$$

This expression makes clear that:

- $P_{ij}(y_j^*) = \frac{1}{2}$;
- as $\mu \to 0$, $P_{ij}(y)$ converges to a step function $\mathbf{1}\{y > y_j^*\}$;

- for $\mu > 0$, the extreme-value shocks *smooth* the cutoff, but do not change its location y_j^*.

In the parametrization used in the main text,

$$V_i^L(y) = \sigma_L y, \qquad V_{ij}^P(y) = -(1-\tau_j)\,\kappa_j + \sigma_{Pj} y + \delta_j, \tag{F.64}$$

the cutoff between L and j is given by

$$y_j^* = \frac{(1-\tau_j)\,\kappa_j - \delta_j}{\sigma_{Pj} - \sigma_L}, \tag{F.65}$$

and the choice probability can be written as

$$P_{ij}(y) = \frac{1}{1 + \exp\big(-\frac{\sigma_{Pj}-\sigma_L}{\mu}(y - y_j^*)\big)}. \tag{F.66}$$

F.4 Pareto Distribution of Patent Values and Aggregate Shares

The derivations above characterize the choice probability for a given patent value y. To connect the model to data, we must aggregate across patents. We assume that patent values follow a Pareto distribution.

Pareto Distribution of Patent Quality Let y be distributed according to a Pareto law with shape (tail) parameter $\alpha > 0$ and lower bound $y_{\min} > 0$:

$$f(y) = \alpha y_{\min}^{\alpha} y^{-(\alpha+1)}, \qquad y \ge y_{\min}, \tag{F.67}$$

$$F(y) = 1 - \left(\frac{y_{\min}}{y}\right)^{\alpha}, \qquad y \ge y_{\min}. \tag{F.68}$$

The tail probability is

$$\Pr(y \ge z) = \Big(\frac{y_{\min}}{z}\Big)^{\alpha} \quad \text{for } z \ge y_{\min}. \tag{F.69}$$

This heavy-tailed structure captures the idea that a small set of very valuable patents accounts for a large share of total profits.

F.4.1 Deterministic Cutoff and Aggregate Share

In the deterministic case ($\mu \to 0$), the firm shifts to destination j if and only if $y \geq y_j^*$. The fraction of patents that are shifted to j is then

$$s_j^{\text{det}} = \Pr(y \geq y_j^*) = \left(\frac{y_{\min}}{y_j^*}\right)^{\alpha}, \quad \text{for } y_j^* \geq y_{\min}. \tag{F.70}$$

This expression makes clear how the Pareto tail parameter α and the location of the cutoff y_j^* jointly determine the share of patents allocated to destination j. A small change in y_j^*, especially if it lies far in the tail, can induce a substantial change in s_j^{det}.

Logit Choice and Aggregate Share With extreme-value shocks and logit probabilities, the aggregate share of patents assigned to destination j is

$$s_j = \int_{y_{\min}}^{\infty} P_{ij}(y)\, f(y)\, dy. \tag{F.71}$$

In the binary case (home vs j), using the logistic representation above,

$$s_j = \int_{y_{\min}}^{\infty} \frac{1}{1 + \exp\big(-\frac{\Delta b}{\mu}(y - y_j^*)\big)}\, \alpha y_{\min}^{\alpha} y^{-(\alpha+1)}\, dy, \tag{F.72}$$

where $\Delta b = b_j - b_L$ and y_j^* is the deterministic cutoff.

In general, this integral does not admit a closed-form expression and is evaluated numerically. However, the structure is intuitive:

- For $y \ll y_j^*$, $P_{ij}(y)$ is close to 0, but there is substantial mass of patents.
- For $y \gg y_j^*$, $P_{ij}(y)$ is close to 1, but the mass of patents is thin due to the Pareto tail.

The shape of the Pareto density (through α) and the steepness of the logit transition (through μ and Δb) together determine the aggregate share s_j.

Small Approximation When the noise parameter μ is small, the logit probability transitions sharply around y_j^*. In the limit $\mu \to 0$, $P_{ij}(y)$ approaches the indicator $\mathbf{1}\{y \geq y_j^*\}$, and the aggregate share s_j converges to the deterministic expression

$$s_j \approx \Pr(y \geq y_j^*) = \left(\frac{y_{\min}}{y_j^*}\right)^{\alpha}. \tag{F.73}$$

Thus, even when we allow for extreme-value shocks to obtain a smooth and tractable logit structure at the micro level, the Pareto distribution of y and the deterministic cutoffs y_j^* remain the key drivers of selection and of the aggregate fractions of patents (and patent value) allocated to each destination.

In summary, this appendix shows how (i) the Type-I extreme-value assumption on location-specific shocks leads to the logit formula for choice probabilities, (ii) deterministic cutoffs in patent value emerge from the underlying payoff comparisons, and (iii) a Pareto distribution of patent values maps those cutoffs into aggregate location shares in a transparent and tractable way.

G Endogenous Patent Quality

In our baseline model, innovators do not make an explicit effort to improve the quality of their patents, even though firms maximizing global profits would have incentives to do so. Allowing for endogenous quality choice would introduce an additional margin through which tax policy could affect outcomes, by changing not only the allocation of patents across licensing and profit shifting, but also the scale of royalty income. We abstracted from this margin to keep the analysis focused and the quantitative exercise tractable. Here, we sketch a simple extension in which firms choose a costly quality-improving effort that scales patent value.

Quality Effort Upon innovating, a firm chooses a non-negative quality effort $e \geq 0$ at cost $C(e)$, with $C(0) = 0$, $C'(e) > 0$ and $C''(e) > 0$. Effort scales the realized quality of the patent multiplicatively:

$$y = \psi(e)\,\tilde{y},$$

where $\tilde{y}$ is an exogenous draw from the baseline Pareto distribution and $\psi(e)$ is increasing and weakly concave, normalized so that $\psi(0) = 1$. Hence, no effort on the part of innovators corresponds to our baseline model which exogenous patents.

Alternatively, one could specify $\psi(e)$ as the probability of success of a patent project. In this case, natural assumptions would be $\psi(0) = 0$ and $\lim_{e\to\infty} \psi(e) = 1$.

Allocation and Expected Payoffs As in the baseline model, these probabilities arise from a logit allocation rule with idiosyncratic destination-specific shocks. Conditional on realized quality $y = \psi(e)\tilde{y}$, the firm chooses between licensing at home and profit shifting to a foreign destination. As in the baseline model, this allocation is summarized by the choice probabilities $\pi_L(y)$ for licensing and $\pi_j(y)$ for profit shifting to destination j, which satisfy

$$\pi_L(y) + \sum_j \pi_j(y) = 1.$$

Quality-improving effort affects the allocation decision only through its effect on realized patent value, and therefore shifts these probabilities by scaling all payoffs proportionally.

Using these probabilities, the firm's expected payoff from the allocation stage can be written as the probability-weighted average of underlying payoffs:

$$\mathbb{E}[\text{payoff} \mid \tilde{y}, e] = \pi_L\big(\psi(e)\tilde{y}\big)\, \mathcal{W}\big(\psi(e)\tilde{y}\big) + \sum_j \pi_j\big(\psi(e)\tilde{y}\big)\, \Pi_j\big(\psi(e)\tilde{y}\big).$$

This representation parallels the construction of tax bases in the quantitative analysis and makes explicit that policy affects outcomes both by changing payoffs and by reallocating probability mass across licensing and profit-shifting destinations, while quality effort operates by scaling the value of the underlying tax base.

Firm Problem The firm's static problem is therefore

$$\max_{e \geq 0} \; \mathbb{E}_{\tilde{y}} \left[\pi_L\big(\psi(e)\tilde{y}\big)\, \mathcal{W}\big(\psi(e)\tilde{y}\big) + \sum_j \pi_j\big(\psi(e)\tilde{y}\big)\, \Pi_j\big(\psi(e)\tilde{y}\big) \right] - C(e).$$

The optimal effort balances the marginal cost of quality improvement against the expected increase in profits, taking into account how higher quality affects both the level of payoffs and the allocation probabilities.

Implications for the Tax Base An increase in quality effort raises the scale of patent value and therefore increases the tax base associated with both licensing income and profit-shifted royalties. Policies that affect the returns to high-quality patents—such as changes in foreign tax rates or enforcement—would therefore operate not only through reallocation across destinations, but also through incentives to invest in patent quality.

Relation to the Quantitative Analysis In the quantitative analysis, we abstract from endogenous quality effort by setting $e = 0$, which implies $\psi(0) = 1$ and $C(0) = 0$. This isolates the allocation channel and allows us to discipline the model using observed patent shares and royalty flows. The extension outlined here preserves the basic structure of our model and shows that endogenous quality would primarily operate by expanding or contracting the relevant tax bases, without altering the allocation mechanisms emphasized in the paper.

H Profit shifting through intra-firm royalty pricing

Our baseline analysis focuses on profit shifting through the relocation of patent ownership. An alternative margin operates through the manipulation of intra-firm royalty payments. In this case, a firm transfers the patent to an affiliate located in a low-tax jurisdiction and then licenses the technology back to its domestic operations at a distorted royalty rate. By setting royalty payments above (or below) their arm's-length level, the firm can reallocate taxable profits across jurisdictions. In particular, a higher royalty payment shifts profits from a high-tax country to a low-tax affiliate.

In the benchmark model, we separate patent ownership from production, treating the patent owner and the producing entity as distinct agents. To analyze intra-firm royalty manipulation, we adapt the framework so that the patent originator is also the domestic producer. We impose two simplifying assumptions. First, we assume $\tau_i = \tau_i^c$, so that the tax rate on patent income equals the corporate income tax rate. Second, we assume $\rho = 1$, implying that the patent owner has full bargaining power in licensing and therefore extracts the entire surplus, i.e., $\ell_{ij} = y_{ij}$. Under these assumptions, the outcome of licensing to a domestic affiliate is equivalent to in-house production, and the model can be interpreted as one in which the patent owner also produces domestically.

We now extend the profit-shifting problem to allow for manipulation of royalty payments. In the absence of manipulation, the baseline problem remains unchanged: all domestic surplus is attributed to the patent owner, and foreign profits are captured by the term $(1 - \tau_j)W_j$. To introduce pricing distortions, let Δ_{ij} denote the deviation in the net present value of royalty payments from their arm's-length level when the patent is transferred to an affiliate in country j. A positive Δ_{ij} corresponds to excess royalty payments (i.e., inflated royalty fees), which shift profits toward the affiliate, while a negative value captures the opposite case.

The firm's profit-shifting problem becomes:

$$\Pi_i = \max_j \left\{ \underbrace{(1-\tau_i^k)(T_{ij}-\kappa_{ij}) - (1-\tau_i)\Delta_{ij}}_{\text{Headquarter profits}} + \underbrace{(1-\tau_j)\big(W_j + \Delta_{ij} - T_{ij}\big)}_{\text{Affiliate profits}} \right\}.$$

In the benchmark analysis, we impose $\Delta_{ij} = 0$, ruling out intra-firm pricing manipulation. Allowing $\Delta_{ij} \neq 0$ introduces an additional margin through which firms can reallocate profits across jurisdictions via internal pricing, rather than through the location of ownership alone.

Recall that we set $T_{ij} = \sigma P_{ij}$, where $\sigma \geq 0$ and P_{ij} is given by (7). Similarly, we can pose that Δ_{ij} is proportional to the market value of a license, as derived in (3). In particular, suppose $\Delta_{ij} = \zeta V_{ji}$, which implies

$$W_j + \Delta_{ij} = (1+\zeta)V_{ji} + \sum_{m \neq i} V_{jm}$$

After replacing T_{ij} and Δ_{ij}, and rearranging, we can write the profit-shifting problem as follows

$$\Pi_i = \max_j \; (1-\tau_j)W_j - (1-\tau_i^k)\kappa_{ij} + (\tau_j - \tau_i^k)\sigma[(1-\theta)\left(\tfrac{1-\tau_i}{1-\tau_i^k}\right)W_i + \theta W_j] - (\tau_j - \tau_i)\zeta V_{ji}.$$

If we assume $\tau_i^k = \tau_i$, as we did in Section 5, then the problem simplifies to

$$\Pi_i = \max_j \; (1-\tau_j)W_j - (1-\tau_i)\kappa_{ij} + (\tau_j - \tau_i)\{\sigma[(1-\theta)W_i + \theta W_j] - \zeta V_{ji}\}.$$

A positive value for ζ, meaning excess royalty payments to the affiliate ($\Delta_{ij} > 0$), increases profits when the affiliate resides in a low-tax country, i.e., when $\tau_j < \tau_i$. In this case, profit-shifting leads to increased global profits when lowering σ (discount transfer price) or increasing ζ (high royalty fees). As with transfer pricing, the benefits of price/fee manipulation disappear when there are no differences in tax rates.