Government Reputation, FDI, and Profit-Shifting *

Yeonggyu Yun†

University of Wisconsin-Madison

July 2026

Abstract

Countries with high expropriation risk receive less FDI and lose more profit to shifting, yet their statutory corporate tax rates are no lower than those of safe countries. I explain this with a model in which the government may renege on the rate it announces, and firms can shift profits abroad once the rate is set. The central force is that profit-shifting disciplines a government that cannot commit. It bounds how much the government can extract after investment is sunk, since taking too much drives the tax base abroad, so the burden firms bear stays close to the announced rate, and closer the worse the government's reputation, as in the data. Profit-shifting also makes credibility a reason to tax less: when firms place more weight on the announcement, investment responds more strongly to it, so a government with a good reputation competes for capital by announcing a low rate, while one with a poor reputation gains little from cutting and announces a high rate. This inverts the usual hold-up logic, in which only trusted governments attract capital and therefore tax it most—the benchmark my model nests without profit-shifting, and one whose predicted pattern the data do not display. A two-period version endogenizes reputation, with profit-shifting lowering the revenue a government sacrifices to stay credible; and investment treaties, modeled as a penalty on reneging, raise investment and lower statutory rates where credibility is scarce.

Keywords: Commitment, Corporate income tax, Foreign direct investment, Government reputation, Profit-shifting

JEL Classification: E61, F21, F23, F55, H21, H26

---

*I am deeply indebted to advice and guidance by Charles Engel, Rishabh Kirpalani, and Kim Ruhl. I appreciate helpful comments from Carter Braxton, Louphou Coulibaly, Lydia Cox, Agustín Gutiérrez, Nora Strecker, Heejin Yoon, and seminar participants at the University of Wisconsin-Madison, Midwest Macro Meeting in Fall 2023, International Institute of Public Finance in 2024, and Asian Meeting of the Econometric Society in Summer 2024. Financial support from the University of Wisconsin-Madison and Korea Foundation for Advanced Studies are gratefully acknowledged.

†Email: yyun23@wisc.edu.

1 Introduction

A statutory corporate tax rate is announced before investment and applied after it. The government legislates the rate at which profits will be taxed before a multinational sinks capital into the host country; once the capital is sunk, some governments apply the announced rate and others do not. India amended its tax code retroactively in 2012 to reach Vodafone, Hungary imposed windfall taxes on selected sectors in 2010, and Argentina nationalized private pension assets in 2008 and the oil company YPF in 2012. A government with a reputation for such reneging attracts less foreign capital, and it has an instrument with which to compensate investors for the risk, namely a lower tax rate. Countries with the worst reputations do not use it. They receive less FDI and lose more profit to shifting, yet their statutory corporate tax rates are no lower than those of safe countries, and in the raw cross-section they run about three percentage points higher per point of expropriation risk. Why do governments with bad reputations tax at high rates?

The answer is that the statutory rate and the burden firms expect to pay diverge, by an amount that depends on the government's reputation. What firms price is not the posted rate but the rate they anticipate once capital is sunk. I document five facts about corporate taxation across countries, proxying reputation by investment-risk scores from the International Country Risk Guide. Riskier countries receive less FDI and lose more profit to shifting, yet set statutory rates that are no lower—and if anything higher—than those of safe countries. The resolution is the fourth fact: the gap between the effective and the statutory tax burden, which I call the effective–statutory gap, is smaller where reputation is worse, rising by three to five percentage points per risk point, so effective burdens there lie closer to, and more often above, the posted rate. Crucially, the gap is at most a few percentage points wide, never the tens of points that outright confiscation would imply. Finally, reputation moves slowly: less than a third of a risk deterioration is undone within four years. The statutory rate and the gap are the two objects the model must explain—why risky governments refuse to undercut on the posted rate, and why the burden they actually impose departs from it only modestly.

I explain these facts with a model of corporate taxation that rests on two ingredients: firms do not know whether the government will honor its announced rate, and they can shift profits abroad once the rate is set. I build on the finite-horizon capital-taxation model of Chari, Kehoe, and Prescott (1988), add cross-border profit-shifting, and model reputation as the prior probability that the government is a commitment type, following Dovis and Kirpalani (2021); the empirical risk scores map to its complement. Reputation has two opposing effects on the rate a committing government announces. A better reputation expands the tax base by attracting investment, which raises the optimal rate; it also makes firms rely more heavily on the announcement, so a given increase deters more investment, which lowers it. Neither effect harms firms: better reputation always lowers the tax they expect to pay and raises investment. The two effects therefore contend only over the level of the announced rate, and profit-shifting determines which one dominates.

The static model isolates the mechanism. A small open economy faces a government that is either a commitment type, which honors its announcement, or an opportunistic type, which re-optimizes once capital is sunk; firms do not observe the type, invest against the reputation-weighted rate, and shift profits after the rate is realized. Two results follow. First, profit-shifting caps the deviation. A reneging government stops strictly short of full confiscation, because a confiscatory rate would drive the entire base abroad—I prove the deviation is interior by an explicit margin—so realized gaps stay within the few percentage points seen in the data. Second, profit-shifting reverses the comparative static of the announced rate. In the calibrated model the committing government lowers its announcement as its reputation improves, matching the sign in the data, because when firms place more weight on the announcement their investment responds more strongly to it, so the government competes for capital by announcing less. The benchmark without shifting gives the opposite comparative static: I show as a theorem that when profits cannot move, the announced rate has a closed form that rises with reputation, so risky countries would post the lowest rates and reneging would be confiscatory—neither of which appears in the data.

The static model treats the deviation as mechanical: the opportunistic type reneges whenever it moves, so honesty is an assumption about types, and under that type the realized rate counterfactually rises with reputation. I therefore extend the model to two periods, which makes honoring the announcement a choice by giving the opportunistic type a reputation to protect. It trades off the one-period gain from reneging against the future investment lost once its type is revealed, and for reputation and patience not too low it honors the announcement and pools with the commitment type. Two implications follow. First, because both types announce the same rate, statutory rates trace announcements while the gap traces deviations, so the model's two margins map onto the two the data measure. Second, the cost of sustaining honesty depends on profit-shifting: without it, keeping a low-reputation government honest requires announcements so close to confiscation that revenue vanishes, whereas with it an honest, positive-revenue equilibrium exists at every reputation level, though the revenue-maximizing government may still tolerate occasional deviation at very low reputation. The extension thus endogenizes reputation and gives both types a reason to behave alike.

A final extension applies the mechanism to investment treaties. I model a bilateral treaty with investor-state dispute settlement as a penalty on reneging and show that treaties compress the gap, raise investment, and, contrary to the intuition that a safer base should be taxed more heavily, lower statutory rates, with all three effects concentrated in the risky countries where the credibility constraint binds. In the data, statutory rates fall after treaty ratification in risky countries but not safe ones, and decline with treaty coverage among risky hosts in the cross-section, patterns consistent with these predictions, though anticipation and a fifty-year accumulated stock make this a configuration consistent with the mechanism rather than an identified test.

Literature Review This paper builds on three literatures. The first studies policy under limited commitment and the reputational forces that substitute for it. The time-consistency problem in capital taxation originates with Chari et al. (1988), and a recent literature models reputation as learning about a hidden government type: Dovis and Kirpalani (2021) study fiscal rules, Dovis and Kirpalani (2020) study federal bailouts, and Amador and Phelan (2021) study sovereign default. I adopt this approach and bring it to open-economy corporate taxation, where the instrument that disciplines the government is neither default nor a fiscal rule but firms' ability to move the tax base after the rate is set.

The paper is also related to the literature on how a government's inability to commit shapes the FDI it receives (Aguiar & Amador, 2011; Schnitzer, 1999; Thomas & Worrall, 1994). Thomas and Worrall (1994) and Aguiar and Amador (2011) model expropriation risk as an enforcement constraint on self-enforcing investment and show that it slows capital accumulation, and Azzimonti (2018) traces the risk to political instability. Relative to these papers, I model reputation over the corporate tax, so that deviation occurs on the equilibrium path and leaves a measurable trace in the gap, rather than only tightening a constraint that never binds observably. The disciplining force also differs from Janeba (2000), where firms install excess capacity ex ante to deter ex-post taxation: in my model it operates ex post through mobile profits rather than production, and so survives even when capital is fully sunk. Langenmayr and Simmler (2021) show that the mobility of a jurisdiction's tax base shapes its rate; I embed that ex-post mobility margin in a reputation game.

Finally, the paper contributes to the literature on international corporate taxation and profit-shifting. Peralta, Wauthy, and van Ypersele (2006), Devereux, Lockwood, and Redoano (2008), and Johannesen (2010) study how governments compete for firms' investment and reported profits. Relative to these papers, I cast profit-shifting as a force that disciplines a time-inconsistent government rather than only a channel of tax competition, while sharing their strategic-complementarity implication that the optimal domestic rate rises with the foreign rate. This role rationalizes the finding of Johannesen, Tørsløv, and Wier (2019) that profit-shifting is more aggressive in low-income countries: worse reputation raises the realized burden, and with it the incentive to shift.

Outline Section 2 documents the facts. Section 3 presents the static model and isolates the role of profit-shifting. Section 4 develops the two-period model with endogenous reputation and the treaty extension. Section 5 concludes, and proofs are collected in Appendix A.

2 Data and Empirical Facts

2.1 Data Sources

My main sample for analysis is an annual panel of countries from 2000 through 2021 that consists of information on government reputation, statutory corporate income tax rates, and net FDI inflows. I use investment risk scores provided in the International Country Risk Guide (ICRG) written by the PRS Group to proxy government reputation. The ICRG investment profile score aggregates three sources of risk associated with business activities in a country, namely risk of capital expropriation, risk to profit repatriation, and risk of payment delays. All of these risks arise when the government is likely to arbitrarily seize firms' assets, profits, or revenues, so the scores capture the credibility of the government in precisely the sense of the model. Because it corresponds most closely to the model's notion of a deviation, I use the expropriation sub-score as my baseline measure. The sub-scores are available monthly for about 140 countries from 2001, and I recode them so that a higher value means a higher risk of investing in the country. Because the scores begin in 2001 and enter the regressions with a one-year lag, the risk specifications effectively cover 2002 through 2021, and the year 2000 appears only in descriptive objects that do not involve risk. The aggregate investment profile score, which combines all three risk dimensions, is the regressor for the profit-shifting analysis and a robustness measure elsewhere. I combine the risk scores with statutory corporate income tax rates collected by Enache (2022), which include sub-national taxes, and with FDI net inflows from the World Bank Development Indicators database.

For profit-shifting and effective tax burdens I use two sources. The first is a set of cross-sectional profit-shifting estimates by country for 2016, estimated by Garcia-Bernardo and Jansky (2021) from OECD Country-by-Country Reporting (CbCR) of corporate profits. I use estimates from their misalignment model, which adjusts reported profits according to real economic activity, together with the accompanying mean and median effective tax rates on accrued corporate income. The second extends the effective-rate evidence to a panel. From the OECD CbCR aggregates for 2016–2022, I compute each host country's accrued effective tax rate as corporate taxes accrued over positive profits of foreign-controlled sub-groups. The 2016 wave covers only a small set of early-reporting parents and the 2022 wave was incomplete at the vintage I downloaded, so panel estimates use 2017 through 2021. Because the set of reporting parent jurisdictions expands over time even within these years, panel specifications include year fixed effects, and I verify robustness on the balanced set of parents reporting throughout.

I do not exclude tax havens classified by Hines (2010) and Tørsløv, Wier, and Zucman (2022) from my main analysis.[1] Within the scope of this paper, tax havens are countries that set a very low tax rate to attract FDI and profit-shifting with a minimal risk of deviation. I view these countries as maintaining a good reputation to keep their positions as tax havens. Excluding them from the sample does not qualitatively change the empirical facts I show henceforth, and I report haven-excluded estimates alongside the baseline throughout.

[1] Table A1 provides the list of tax havens.

2.2 Empirical Facts

I organize the evidence around the puzzle from the introduction: riskier countries receive less FDI and lose more profit to shifting, yet their statutory rates are no lower. After documenting these three facts, I show that effective rates lie closer to statutory rates where reputation is worse, which reconciles them, and finally that the risk scores themselves move slowly.

2.2.1 FDI net inflows are higher in countries with better reputation.

Countries that are riskier on average receive less FDI. Figure 1 plots each country's average FDI net inflows relative to GDP from 2000 to 2021 against its average risk score, and the relationship is negative.[2]

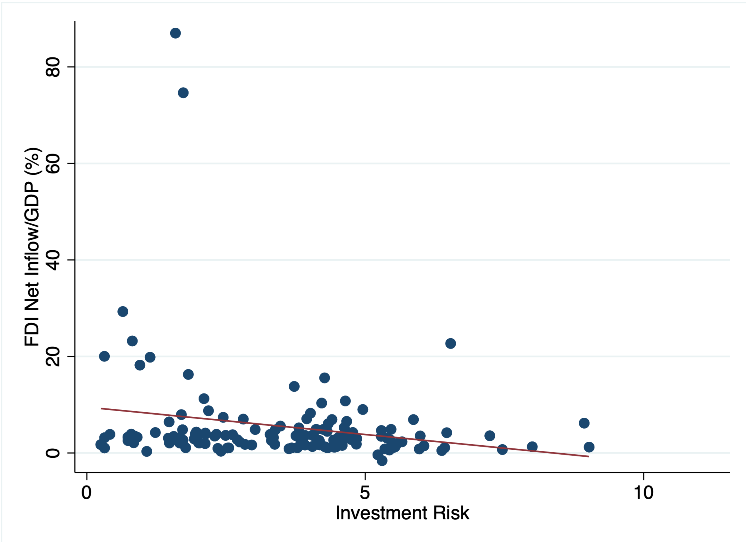

FIGURE 1. Government Reputation and FDI Net Inflows (% of GDP)

Table 1 confirms this relation in the panel with year fixed effects and standard errors clustered by country. A one-point increase in lagged expropriation risk lowers log real FDI net inflows by 14.8 percent, and given the persistence of FDI the implied long-run response is substantially larger. Scaling inflows by GDP and adding the lagged dependent variable, log GDP, GDP growth, investment growth, and the capital control index of Fernandéz, Klein, Rebucci, Schindler, and Uribe (2016) yields the same sign with less precision. Previous work by Klein and Rosengren (1994) and Desai, Foley, and Hines (2006) emphasizes capital-account frictions as determinants of FDI, and the association with reputation keeps its sign, though not its precision, when these are held

[2]Normalizing by GDP addresses the concern that a levels-based association merely reflects country size.

fixed, so the log-inflow specification carries the clearest evidence. Estimates with the standardized aggregate risk score are less precise and are collected in Appendix D. The within-country variation in these slow-moving scores is limited, so the event studies of Section 2.2.5 provide a sharper view of timing.

TABLE 1. Regression of FDI Net Inflows

	log FDI net inflows (1)	FDI net inflows / GDP (2)
Expropriation Risk$_{t-1}$	-0.1475***	-0.2897
	(0.0337)	(0.2494)
Lagged dep. var.	0.8979***	0.7347***
	(0.0146)	(0.0510)
Capital Control$_t$		-2.0512*
		(1.1655)
$\log \text{GDP}_t$		0.2595
		(0.1910)
ΔGDP_t		-0.0317
		(0.0952)
ΔI_t		0.0289
		(0.0179)
Year FE	Yes	Yes
N	2311	1951
Countries	139	122

Note: Annual panel, 2002–2021. Standard errors clustered by country. * and *** denote significance at 10% and 1% levels, respectively. Estimates with the standardized aggregate risk score are in Table A6.

2.2.2 Less profit is shifted out of countries with better reputation.

Multinational firms shift more profits out of countries with higher investment risk. Figure 2 plots the Garcia-Bernardo and Jansky (2021) profit-shifting estimates against investment risk, where a negative value means that firms shift profits out of the country. The fact might appear puzzling, because profit-shifting responds to tax differentials rather than to reputation directly. I attribute it to the higher realized tax burdens in riskier countries documented in Section 2.2.4, consistent with what the model implies.

The negative relationship holds when I control for the capital control index, exchange rate risk, real GDP growth, the manufacturing share of GDP, and the growth of FDI relative to GDP, as summarized in Table 2. This table uses the aggregate investment profile score rather than the

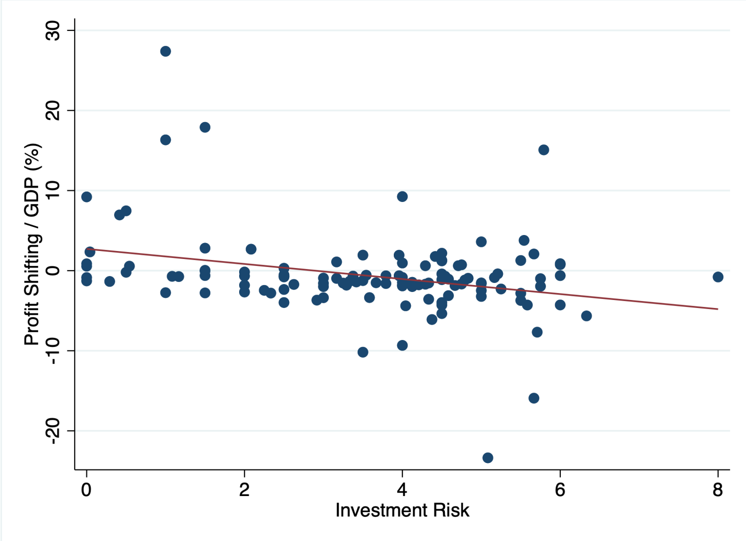

FIGURE 2. Government Reputation and Profit-Shifting / GDP (%)

expropriation sub-score used elsewhere, because profit-shifting responds to the whole realized burden—expropriation, repatriation restrictions, and payment delays alike—so the composite is the natural regressor for this margin. The expropriation sub-score alone preserves the negative sign in all four specifications, with point estimates from -1.39 to -0.94 and standard errors of 0.70 and 0.60, significant at the five-to-ten percent level in all but the specification with the full set of controls.

TABLE 2. Regression of Profit-shifting/GDP

	(1)	(2)	(3)	(4)
Risk_t	-0.9844***	-0.9288**	-0.8110***	-0.7402**
	(0.3453)	(0.4048)	(0.2856)	(0.3481)
Capital Control_t	-2.2289	-2.1167	-1.3246	-1.1759
	(1.4297)	(1.5010)	(1.2701)	(1.3273)
$\text{Exchange Rate Risk}_t$		-0.1261		-0.1532
		(0.3889)		(0.4011)
ΔGDP_t	0.0806	0.0664	0.0512	0.0347
	(0.1531)	(0.1835)	(0.1537)	(0.1870)
$\text{Manufacturing Share of GDP}_t$	0.0511	0.0432	0.0843	0.0748
	(0.1076)	(0.1169)	(0.0933)	(0.1037)
$\Delta\text{FDI}_t/\text{GDP}_t$			-0.1376***	-0.1384***
			(0.0514)	(0.0519)
R^2	0.1248	0.1198	0.2398	0.2367
N	116	115	116	115

Note: Risk is the aggregate ICRG investment profile score recoded so that a higher value is riskier. Standard errors are robust to heteroskedasticity. ** and *** denote significance at 5% and 1% levels, respectively.

2.2.3 Statutory tax rates are higher in riskier countries, and never lower.

The previous two facts show that a bad reputation costs a country investment and tax base. A government could compensate firms for this risk with a low statutory rate. Riskier countries do not do this. Figure 3 plots each country's average statutory rate from 2001 to 2021 against its average expropriation risk. The relationship is positive. A country that is riskier by one point has a statutory rate that is higher by 3.3 percentage points, with a standard error of 1.0.

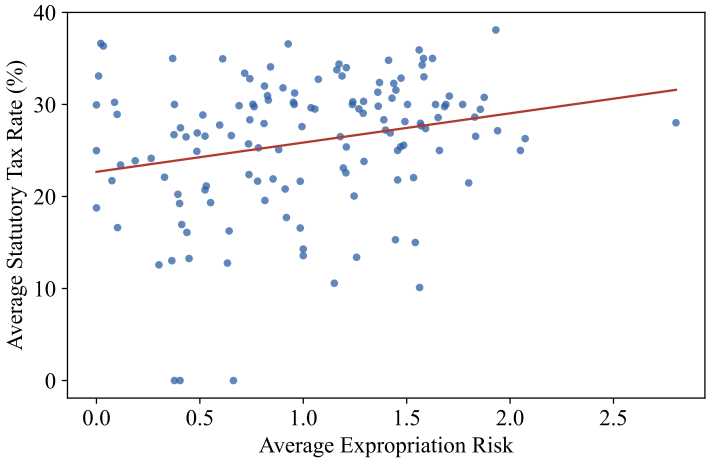

FIGURE 3. Government Reputation and Statutory Corporate Income Tax Rate

Table 3 examines this relationship. Column (1) reports the cross-country regression underlying the figure. Columns (2) and (3) estimate the same relationship in the annual panel with year fixed effects and lagged risk. With log GDP as the income control the coefficient stays positive and loses precision. With log GDP per capita it is zero. Column (4) reports the specification of the first version of this paper, which uses the aggregate investment profile score with country and year fixed effects. There a within-country deterioration of the score raises the statutory rate in the following year at the ten percent level.[3] There are two ways to read the table. Statutory rates rise with risk, or the raw correlation reflects differences in development. No specification yields a significantly negative coefficient. The conservative conclusion, and all the theory requires, is that risky countries do not set lower rates. The raw data suggest they set higher ones.

This is the puzzle the model must explain. The benchmark without profit-shifting predicts, as a theorem rather than a calibration, that the announced rate rises with reputation, so risky countries

[3]The aggregate score bundles expropriation risk with impediments to profit repatriation and payment delays. A within-country deterioration of the aggregate score may coincide with fiscal distress that raises tax rates for budgetary reasons, so I read column (4) as an association rather than a causal response.

TABLE 3. Regression of Statutory Corporate Income Tax Rate

	(1) Between	(2) Panel	(3) Panel	(4) Panel
Expropriation Risk	3.2753***			
	(1.0293)			
Expropriation Risk$_{t-1}$		0.2754	−0.0669	
		(0.7463)	(0.7762)	
Aggregate Risk$_{t-1}$				0.2331*
				(0.1152)
Income control	—	log GDP	log GDP p.c.	log GDP
Fixed effects	—	Year	Year	Country, Year
N	138	2430	2251	2791
Countries	138	136	125	135

Note: Dependent variable is the statutory corporate income tax rate. Column (1): cross-country averages over 2001–2021, heteroskedasticity-robust standard errors. Columns (2)–(3): annual panel 2002–2021, expropriation sub-score, GDP growth controlled, standard errors clustered by country. Column (4): aggregate investment profile score recoded so that a higher value is riskier, GDP growth and log GDP controlled, country and year fixed effects, standard errors clustered by country and year. * and *** denote significance at 10% and 1% levels.

should set the lowest statutory rates. The data do not display that pattern in any specification. The model with profit-shifting generates a declining schedule at the calibration, in line with the raw gradient. The next fact turns to what firms actually pay. The statutory rate is not the burden firms expect to bear, and the gap between the two narrows as reputation worsens.

2.2.4 Effective tax burdens are closer to statutory rates where reputation is worse.

The most direct trace of the government's deviation is the difference between the effective tax rate that firms actually bear and the statutory rate. I call this difference the effective–statutory gap. Effective rates are generally below statutory rates because of deductions and profit-shifting itself. The question is how the gap varies with risk. It closes sharply as risk rises. Moving from the safest to the riskiest tercile, the mean effective–statutory gap rises from −8.2 to −3.6 percentage points, and the share of countries whose effective rate exceeds their statutory rate doubles. Table 4 quantifies this margin. In the 2016 cross-section, the gap rises by 3.37 percentage points per point of expropriation risk, and the estimate is essentially unchanged when tax havens are excluded. The full CbCR panel is diluted by very small host economies with noisy effective rates. Among hosts with more than one billion dollars of foreign-affiliate profits, which account for the bulk of the tax base, the gap rises by 4.92 points per risk point, and by 3.97 when havens are excluded.

Equivalently, effective rates are higher in riskier countries even after controlling for the statutory rate, which implies tax-related costs to business beyond the statute. This pattern is consistent with the retroactive amendments and repatriation impediments that the model calls deviations, though base and enforcement differences that the controls do not absorb may contribute as well.

TABLE 4. The Effective–Statutory Gap and Expropriation Risk

	(1) CS 2016	(2) CS 2016 no havens	(3) Panel large hosts	(4) Panel large, no havens
Expropriation Risk$_t$	3.365**	3.408**	4.918**	3.974*
	(1.699)	(1.720)	(2.387)	(2.126)
Income control	GDP p.c.	GDP p.c.	GDP	GDP
Year FE	—	—	Yes	Yes
N	115	97	310	251
Countries	115	97	76	63

Note: Dependent variable is the effective–statutory gap. Columns (1)–(2): 2016 cross-section (median accrued effective rates, Garcia-Bernardo and Jansky (2021) vintage), heteroskedasticity-robust standard errors, controls are log GDP per capita and GDP growth. Columns (3)–(4): OECD CbCR panel 2017–2021, standard errors clustered by country, log GDP control. "Large hosts" are country-years with over $1bn of positive foreign-affiliate profits. Robustness (balanced parent set, cash-based effective rates, alternative specifications) in Appendix D. * and ** denote significance at 10% and 5% levels.

The magnitude of the gap is as informative as its sign. Gaps of a few percentage points are what an interior deviation produces. If deviations meant full confiscation, the conditional gap would be the entire untaxed profit, upward of sixty percentage points at observed statutory rates, and expected gaps would be an order of magnitude larger than anything in the data. The evidence therefore calls for a force that stops a reneging government short of full confiscation. The model in the next section shows that profit-shifting provides one. Two measurement caveats qualify the comparison. Deductions and loss treatment push effective rates below statutory rates everywhere, which is why the average gap is negative, so the regressions read the slope of the gap in risk rather than its level. And the most severe expropriations can remove an affiliate from the reporting population altogether instead of raising its recorded rate, so the gap understates the deviation margin where risk is highest.

Appendix D.1 makes the mapping between this gap and the model's deviation margin exact and confronts it with each measurement objection. Because the same taxable base enters the numerator and denominator of the accrued rate, the effective rate constructed from Country-by-Country Reporting recovers the realized statutory rate, so the gap measures the rate a government actually imposes in excess of the one it announced, and its gradient in risk recovers the model's deviation margin. Deductions, intra-group dividend double-counting, and affiliate selection shift the level of the gap but, resource rents aside, not its slope in risk, and the slope survives the corresponding checks (Table A4). Excluding conduit jurisdictions, conditioning on the intra-group share of revenue, and an upper-bound dividend correction (Blouin & Robinson, 2025) leave it between three and six points—the dividend double-counting is concentrated in low-risk conduits and if anything works against the slope—and it holds under region and income-group fixed effects and a per-capita income control. Two limits are worth stating plainly. Country fixed effects cannot identify a slope in a country attribute that barely moves within countries over these years, so the fact is a cross-country one; and a resource-rent control halves the large-host panel slope, because resource-rich states are both riskier and subject to special high-rate resource taxation, though the cross-section is unaffected and resource nationalism is itself a form of the deviation the model describes.

2.2.5 Reputational damage is persistent.

Finally, I use the monthly sub-scores to examine the dynamics of reputation itself. For this exercise I keep the score in its original ICRG orientation, in which a higher value means a safer environment, and I date a risk-deterioration event as a twelve-month decline of at least one point in the sub-score on its 0–4 scale, with a 36-month re-event exclusion window. The 162 resulting country-events include the familiar episodes. Argentina in 2002, 2008, and 2020, Bolivia in 2004, Venezuela in 2005, Hungary in 2009, and India in 2012 all appear. Figure 4 traces the average score around events relative to its level thirteen months prior. There is no pre-trend. The score falls by one point at impact and remains 0.7 to 0.8 points below baseline for the following four years, so less than a third of the damage is recovered within 48 months. In local projections with year fixed effects, FDI net inflows decline by 2.9 percent of GDP two years after an event, with a standard error of 1.9. Statutory rates do not respond at these horizons, consistent with their role as legislated announcements rather than instruments of high-frequency adjustment.

Since only a cross-section of profit-shifting estimates is available and risk scores are discrete, I check the validity of these observations with an alternative proxy for the government's time inconsistency, the World Bank Government Effectiveness index, in Appendix C. Further robustness

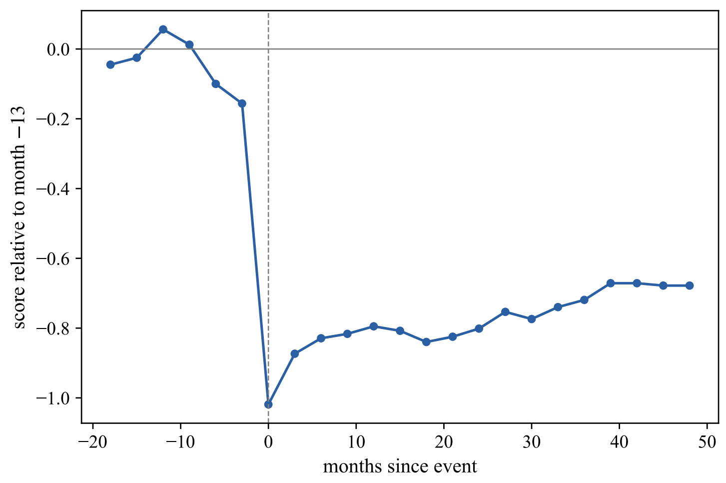

FIGURE 4. Average monthly expropriation sub-score, in its original orientation with higher values safer, around risk-deterioration events, relative to the level 13 months before the event.

checks, including the balanced CbCR parent set and cash-based effective rates, are collected in Appendix D. The patterns are unchanged.

3 Simple Model of Optimal Taxation on Multinational Firms

I rationalize these facts with a model of optimal corporate taxation, relating government reputation to tax rates, FDI, and profit-shifting. Government reputation is modeled as the prior probability that the government is a commitment type, one that honors the corporate tax rate announced at the beginning of a period, so the empirical risk scores map to its complement, and I interpret reneging broadly to include regulation of firms' business activities through fees and non-tax costs.[4] The corporate tax rate chosen at the beginning of the period can be regarded as the statutory tax rate explicitly stated in law. The mechanism concerns the part of the government's target that remains mobile after investment is sunk, namely current profits that can be relocated on paper before the tax is collected. Claims on physical assets or on past liabilities cannot be relocated in this way, and the model without profit-shifting studied below corresponds to that case, since no part of the base can be moved after a deviation. The empirical risk measures mix both kinds of exposure, so the two versions of the model cover the two polar cases, and the comparison between them indicates

[4]There have been court disputes between multinational firms and governments over the legality of retroactive tax amendments in many developing countries such as Kenya and Nigeria, but this is also relevant for some developed countries, including the U.S., where retroactive tax amendment is allowed for the federal government in principle, as Ortigueira and Pereira (2021) point out.

which component disciplines tax outcomes.

I set up a static Bayesian game with a government that either honors or reneges on its announced corporate income tax rate and a continuum of homogeneous multinational firms that choose investment and determine how much profit to shift out of the host country. There are two economies, the host country and the foreign economy. The host country is a small open economy and takes the foreign tax rate as given. In the host country, the government is either a commitment type or an opportunistic type, and firms do not know which.[5] The commitment type chooses a tax rate at the beginning of the period to maximize revenue and commits to it. The opportunistic type lacks commitment and re-optimizes after firms invest.[6] Both types announce the same rate at the beginning of the period, because the opportunistic type does not want to reveal its type. I focus on this pooling equilibrium and denote the announced rate by τ^R.

Firms have a common prior that the government is the commitment type, which I call reputation and denote $p \in [0, 1]$. Since firms do not know the type, the tax rate to be imposed at the end of the period is uncertain, and firms invest to maximize expected profit. Profit-shifting, however, is determined after the government's type is revealed and the tax rate is finalized.[7]

Timeline Timing within the period is illustrated in Figure 5. At the beginning of the period, the government announces the corporate tax rate to which it commits if it is the commitment type. Firms choose investment k to maximize expected profit. If the government is opportunistic, it re-optimizes the tax rate to τ^O taking aggregate investment K as given. Otherwise the announced rate stands. After the tax is finalized, firms choose profit-shifting $\theta(k, \tau)$. I solve the game backward, starting from the last stage.

[5]Schnitzer (1999) analyzes a dynamic contract in which an enforcement constraint disciplines the government when multinational firms decide how much to invest, characterizing when the constraint binds but ruling out deviation on the equilibrium path; Azzimonti (2018), by contrast, studies a Markov-perfect equilibrium in which the government over-expropriates along the path. I intend to capture the incidence of government deviation directly, so I build a model of reputation over commitment to the corporate tax.

[6]Revenue maximization is the natural objective when taxing foreign-owned capital. Firms are owned abroad and capital earns the world return, so host welfare coincides with revenue up to the surplus of domestic factors, which taxation of the multinational sector leaves untouched. A revenue-maximizing government whose base can flee is also the object of the Leviathan tradition, in which base mobility disciplines an over-extracting government (Brennan & Buchanan, 1980; Edwards & Keen, 1996; Keen & Konrad, 2013); here the discipline runs through reputation rather than contemporaneous competition. With domestically owned factors the deviation calculus would acquire a term penalizing over-taxation, and Section 3.9 verifies that introducing such a term leaves the comparative statics intact.

[7]This reflects that profit-shifting techniques like transfer pricing or reallocation of debt can be implemented quickly once the government's tax decision is finalized.

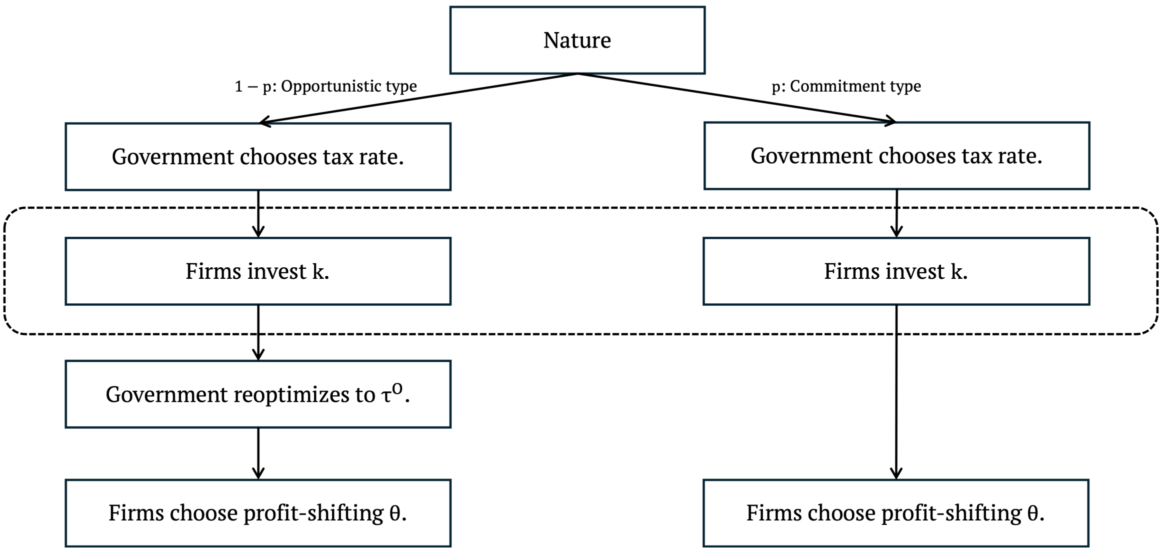

FIGURE 5. Timeline within a Period

3.1 Stage 4: Firms' Profit-shifting Choice

An atomistic firm chooses how much profit to shift to its foreign affiliate after observing the tax rate imposed at the end of the period. This timing matters because end-of-period shifting is what punishes the government's deviation. If the government commits to the announced rate, the firm's profit-shifting plan does not change from the plan it would have established at the beginning of the period.

A representative firm maximizes the sum of after-tax profits in the host country and the foreign country given tax rates τ and τ^* and before-shifting profits $\rho(k)$ at home and $\rho(\bar{k} - k)$ abroad. I assume ρ is strictly concave and strictly increasing[8] and that profit-shifting incurs a convex cost following the literature.[9] To be specific, I impose a real quadratic cost of profit-shifting as in Hines and Rice (1994), with parameter γ governing its costliness, and I interpret γ as reflecting potential legal costs associated with income- and profit-shifting.[10]

The risk of the government's deviation is distinct from the cost of profit-shifting γ. Reputation concerns whether the government seizes profits net of all costs, so bad reputation is not the same as a high shifting cost. I focus on the implications of reputation p throughout.

The firm chooses the amount of profit-shifting θ, where a negative value means shifting profits

[8] I impose $\rho(k) = zk^\alpha - r^*k$ with $\alpha < 1$ in the quantitative analysis.

[9] Kallen (2023) shows that a convex cost of profit-shifting can arise from a monitoring game between a firm and an auditor.

[10] Kallen (2023) remarks that γ can be heterogeneous across firms with different intangible intensity. I set γ common across firms and symmetric across countries for simplicity. The sensitivity analysis in Section 3.9 relaxes the symmetry and considers non-deductible costs, without qualitative changes.

from the host country to the foreign country:

$$\max_{\theta}\ (1-\tau)\left[\rho(k)+\theta-\frac{\gamma}{2}\frac{\theta^2}{\rho(k)}\right]+(1-\tau^*)\left[\rho(\bar{k}-k)-\theta-\frac{\gamma}{2}\frac{\theta^2}{\rho(\bar{k}-k)}\right] \tag{1}$$
$$\text{s.t. } \rho(k)+\theta-\frac{\gamma}{2}\frac{\theta^2}{\rho(k)}\geq 0, \qquad \rho(\bar{k}-k)-\theta-\frac{\gamma}{2}\frac{\theta^2}{\rho(\bar{k}-k)}\geq 0.$$

The objective is strictly concave in θ, so whenever the nonnegativity constraints are slack the unique solution is the first-order rule

$$\theta=\frac{1}{\gamma}\frac{(\tau^*-\tau)\rho(k)\rho(\bar{k}-k)}{(1-\tau)\rho(\bar{k}-k)+(1-\tau^*)\rho(k)}\equiv\theta(k,\tau). \tag{2}$$

At tax rates high enough that the implied host base would turn negative, the host constraint binds instead, a case A.2 treats explicitly. How shifting responds to the tax rate determines the rates the government can impose later in the game, so I record its properties first.

Lemma 1. *$\theta(k,\tau)$ defined in (2) is strictly decreasing and strictly concave in τ.*

Proof. See A.1. ∎

Lemma 1 states that firms shift more profits out of the country when tax rates are higher, and it supplies the curvature on which the government's problems below rely. Profit-shifting is also decreasing in γ, because a high γ makes shifting costly. Note that $\theta(k,\tau)$ is neither globally increasing nor decreasing in k.

3.2 Stage 3: Opportunistic Type's Re-optimization

The opportunistic government re-optimizes its tax rate after firms' investments are realized. It takes aggregate investment K and aggregate before-shifting profits $\Lambda(K)$ as given, and internalizes firms' aggregate profit-shifting Θ.[11] Its problem is

$$\max_{\tau^O\in[0,1]}\ \tau^O\left[\Lambda(K)+\Theta\left(K,\tau^O\right)-\frac{\gamma}{2}\frac{\Theta\left(K,\tau^O\right)^2}{\Lambda(K)}\right], \tag{3}$$

[11]With a unit measure of homogeneous atomistic firms, $k=K$, $\rho(k)=\Lambda(K)$, and $\theta(k,\tau)=\Theta(K,\tau)$ by aggregation.

where the uniform rate across firms is without loss of generality, since the re-optimization problem over a homogeneous continuum is symmetric.[12] The first-order condition is

$$\Lambda + \Theta^O - \frac{\gamma}{2}\frac{\Theta^{O2}}{\Lambda} + \tau^O\left[1 - \frac{\gamma\Theta^O}{\Lambda}\right]\frac{\partial\Theta}{\partial\tau^O} = 0, \tag{4}$$

where Θ^O denotes profit-shifting under τ^O. The opportunistic type chooses the best-response deviation rate to a given announced rate τ^R through the aggregate investment it induces, $\tau^O(K(\tau^R))$. The key difference from the commitment type, discussed below, is that the opportunistic type does not internalize firms' investment decisions.

Optimization alone rules out confiscation. By Lemma 1, the tax base in brackets is strictly decreasing in the deviation rate, and when shifting is effective enough, a rate close to one drives so much profit abroad that the base is exhausted strictly before the rate reaches one. Revenue is positive at interior rates and non-positive beyond the exhaustion point, so the best-response deviation is interior. The condition for shifting to be effective enough is an upper bound on its cost,

$$\gamma \;<\; \tilde{\gamma}(K) \;\equiv\; \frac{\rho(\bar{k}-K)}{\rho(K)} + \frac{\rho(\bar{k}-K)^2}{2\rho(K)^2}, \tag{5}$$

which I maintain over the range of investment levels arising in the analysis, since the bound depends on K. The following proposition states the result and the size of the margin.

Proposition 1. *Fix aggregate investment K and write $A = \rho(K)$, $B = \rho(\bar{k}-K)$ and $s = \sqrt{1+2\gamma}-1$. If $\gamma < \tilde{\gamma}(K)$ as in (5), the host tax base is exhausted at a unique rate $\bar{\tau}(K) \in (0,1)$ satisfying*

$$1 - \bar{\tau}(K) \;=\; \frac{(1-\tau^*)\,(1-sA/B)}{1+s} \;>\; 0, \tag{6}$$

and the opportunistic type's best-response deviation rate satisfies $\tau^O < \bar{\tau}(K) < 1$. Without profit-shifting the base is predetermined, revenue is strictly increasing in the rate, and $\tau^O = 1$.

Proof. See A.2. ■

Because the bound is endogenous, I verify that this window and the lower bound Γ_{lo} below hold with slack uniformly over the equilibrium investment range at the calibration: the admissible

[12] With heterogeneous firms, for example in shifting ability as in Kallen (2023), a government freed from its announcement might tax firms differentially. Statutory rates are uniform by law, but deviations need not be. I leave this margin to future work.

interval for γ is $[0.32, 2.69]$, which contains $\gamma = 1.25$, and it stays non-empty across the reputation range and the sensitivity exercises. Proposition 1 is what separates the two versions of the model at the deviation margin. The distance from confiscation, $1 - \bar{\tau}(K)$, determines the size of the effective–statutory gap, so it is this contrast that holds the model's gaps to the few percentage points documented in Section 2.2.4 while the benchmark without profit-shifting produces confiscatory gaps.

The best-response deviation rate is also decreasing in τ^R, and not always higher than the announced rate, as shown in Figure 6 for $\gamma = 1.25$ and $p = 0.5$. When the announced rate is low, investment and profits are high, and the opportunistic type exploits the large base by deviating to a higher rate. As the announced rate rises, investment shrinks, reducing the incentive to deviate upward. If the announced rate is so high that investment is severely deterred, the opportunistic type deviates downward to attract profit-shifting into the country and secure the base.

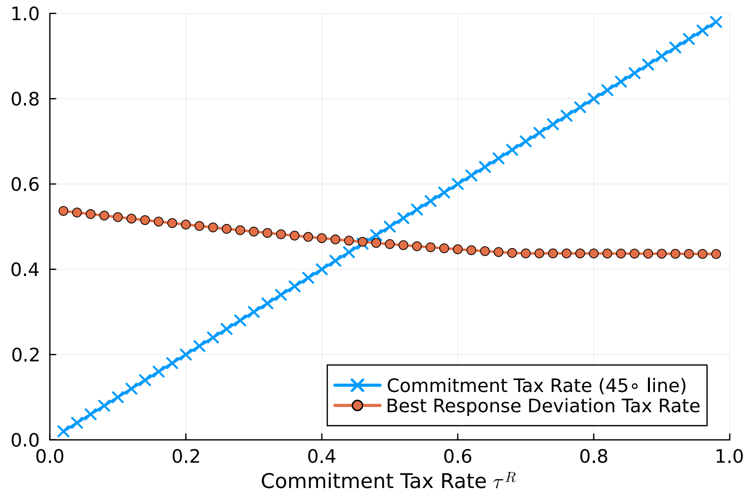

FIGURE 6. Best Response Deviation Tax Rate to Commitment Tax Rate when $p = 0.5$

3.3 Stage 2: Firms' Investment Choice

Individual firms set investment k to maximize the expected sum of after-tax profits at home and abroad. Each firm has an exogenous capital endowment $\bar{k}$ satisfying $\bar{k} < \left(\frac{\alpha z}{r^*}\right)^{\frac{1}{1-\alpha}}$, so that the marginal revenue of capital exceeds its cost r^* in both countries. I thereby abstract from the extensive margin of FDI and confine attention to the intensive margin.

The government's time inconsistency adds uncertainty to the firm's choice as in Ortigueira and Pereira (2021). The host tax rate is a binary random variable equal to τ^R with probability p and τ^O

with probability $1-p$. The firm's problem is

$$\max_{k\in[0,\bar{k}]} \mathbb{E}_\tau \left[(1-\tau)\left[\rho(k)+\theta(k,\tau)-\frac{\gamma}{2}\frac{\theta(k,\tau)^2}{\rho(k)}\right]+(1-\tau^*)\left[\rho(k^*)-\theta(k,\tau)-\frac{\gamma}{2}\frac{\theta(k,\tau)^2}{\rho(k^*)}\right]\right] \tag{7}$$

$$\text{s.t. } k+k^*=\bar{k},$$

with first-order condition

$$\mathbb{E}_\tau\left[(1-\tau)\left(1+\frac{\gamma\theta(k,\tau)^2}{2\rho(k)^2}\right)\rho'(k)-(1-\tau^*)\left(1+\frac{\gamma\theta(k,\tau)^2}{2\rho(k^*)^2}\right)\rho'(k^*)\right]=0. \tag{8}$$

Optimal investment $k(\tau^R,\tau^O)$ is a function of the announced rate and the anticipated deviation rate. Without profit-shifting, the firm equates after-tax marginal profits across countries, so a higher home tax lowers home investment. With profit-shifting, the terms $\gamma\theta^2/2\rho^2$ in (8) show that shifting mitigates the tax distortion on investment by preserving after-tax returns, as discussed by Kallen (2023). The extent of the mitigation depends on γ.[13]

3.4 Stage 1: Commitment Type's Optimal Taxation

At the beginning of the period, the commitment type chooses τ^R to maximize revenue while internalizing firms' investment and contingent profit-shifting plans, taking into account that the opportunistic type announces the same rate but deviates at Stage 3:

$$\max_{\tau^R\in[0,1]} \tau^R\left[\Lambda(K(\tau^R,\tau^O))+\Theta(K(\tau^R,\tau^O),\tau^R)-\frac{\gamma}{2}\frac{\Theta(K(\tau^R,\tau^O),\tau^R)^2}{\Lambda(K(\tau^R,\tau^O))}\right], \tag{9}$$

with first-order condition

$$\Lambda+\Theta^R-\frac{\gamma}{2}\frac{\Theta^{R2}}{\Lambda}+\tau^R\left[\left[1-\frac{\gamma\Theta^R}{\Lambda}\right]\frac{\partial\Theta}{\partial\tau^R}+\frac{\partial}{\partial K}\left[\Lambda+\Theta^R-\frac{\gamma}{2}\frac{\Theta^{R2}}{\Lambda}\right]\frac{\partial K}{\partial\tau^R}\right]=0. \tag{10}$$

Here $\partial K/\partial\tau^R$ is the total derivative along the fixed point in which firms anticipate the deviation best response $\tau^O(K)$, so the channel through the induced deviation rate is included. The bracketed terms capture the distortionary effect of a higher announced rate on the base. The first term is the loss from firms' profit-shifting. The second term, absent from (4), is the loss from deterred investment, which the commitment type internalizes and the opportunistic type does not.[14] Throughout, I impose that $\gamma\geq\Gamma_{\text{lo}}$ so that the tax base increases in investment, and A.2 states the

[13] In the extreme case $\gamma\to 0$, a higher τ could induce higher investment, as firms produce at home and shift most profits abroad. I rule this out by the lower bound Γ_{lo} on γ imposed below and stated in A.2.

[14] The second-order condition of the opportunistic type's problem can be verified analytically from Lemma 1. That of the commitment type's problem involves the curvature of $K(\tau^R)$ and is verified numerically at the computed optimum. See A.2.

closed form.

The deviation rate exceeds the announced rate in equilibrium. The logic is visible in the two optimality conditions. Conditions (4) and (10) are identical except for the internalization term $\tau^R \, \partial(\text{base})/\partial K \cdot \partial K/\partial \tau^R < 0$, so at any common rate the commitment type's marginal revenue is strictly below the opportunistic type's, and its optimum must lie at a lower rate.

Proposition 2. *The deviation tax rate τ^O is higher than the optimal commitment tax rate τ^R in equilibrium as long as investment K is decreasing in τ^R and firms' profit net of profit-shifting is increasing in capital K.*

Proof. See A.3. ■

In reality, the deviation rate can be interpreted as $\tau^O = \tau^R + \Delta$, where τ^R is the statutory rate stated in law and $\Delta > 0$ reflects impediments or costs associated with retrieving profits, such as retroactive taxation or non-tax costs of repatriation. Proposition 2 states that the opportunistic type always makes such a deviation under the commitment type's optimal announcement. Because the deviation rate is interior by Proposition 1, this pins down the observable. With the effective rate measured as accrued taxes over the taxable base, the model's effective rate equals the realized rate, so the expected effective–statutory gap equals the probability of a deviation times the conditional gap. Conditional on a deviation, the effective–statutory gap is the entire untaxed profit $1 - \tau^R$ without profit-shifting, while with profit-shifting it is $\tau^O - \tau^R$ with an interior τ^O. In the calibrated equilibrium below, this conditional gap amounts to a few percentage points, an order of magnitude below confiscation, which is the magnitude documented in Section 2.2.4.

3.5 Equilibrium

With all four stages characterized, I now define the equilibrium. Equilibrium consists of firms' investment and profit-shifting choices $\{k, \theta\}$ and the government's tax rates $\{\tau^R, \tau^O\}$ jointly solving (2), (4), (8), and (10). I focus on the pooling equilibrium in which both types announce τ^R at Stage 1. This is a weak perfect Bayesian equilibrium, and restricting attention to it is without loss, because no equilibrium in which the announcement reveals the government's type exists. Any belief advantage an announcement confers raises investment, and a larger sunk base benefits the opportunistic type at least as much as the commitment type, since it retains the option to deviate on that base. No announcement is therefore attractive to the commitment type yet unprofitable for the opportunistic type to imitate; a formal statement and proof, including off-path beliefs and standard refinements, are in A.6, where I restrict the signal space to the policy announcement. This pooling property is what the two-period extension turns into a statement about what statutory rates measure.

3.6 Comparative Statics

I now examine how the equilibrium varies with reputation. I solve the model with the parameter values in Table 5. The key parameter is γ, and I impose $\gamma = 1.25$, the mean value estimated by Kallen (2023) from the Orbis database. The remaining parameters are set independently of the facts: $\tau^* = 0.15$ is a low-tax shifting destination, $\alpha = 0.66$ fixes the curvature of the before-shifting profit function and is varied over $[0.55, 0.75]$ in Section 3.9, $r^* = 0.04$ is the world real return, and z and $\bar{k}$ are scale normalizations that satisfy the interior-FDI restriction $\bar{k} < (\alpha z/r^*)^{1/(1-\alpha)}$ with ample slack. At these parameter values γ satisfies the two maintained bounds along the computed equilibria. The calibration targets no moment other than γ, yet the model reproduces two it does not target. The conditional deviation gap $\tau^O - \tau^R$ stays between about four and thirteen percentage points across the reputation range, an order of magnitude below the fifty-plus points of confiscation that the no-shifting benchmark implies and in line with the gap magnitudes of Section 2.2.4; and the outward-shifted fraction of home profits is on the order of fifteen to thirty percent, the same order as the estimates of Garcia-Bernardo and Jansky (2021) and Tørsløv et al. (2022). Given the parameters, I vary reputation p and trace aggregate investment K, profit-shifting Θ, and the tax rates.

TABLE 5. Parameter Values

Parameter	z	α	γ	τ^*	r^*	$\bar{k}$
Value	1	0.66	1.25	0.15	0.04	10

Comparative statics are summarized in Figures 7 and 8. Consistent with the data, investment increases with reputation. Because the announced rate lies below the deviation rate and falls with reputation, a better reputation lowers the firms' expected tax rate and raises the expected return on investment. The deviation rate τ^O increases with reputation, as firms invest more when reputation is high and the government's gain from deviating rises with the base.

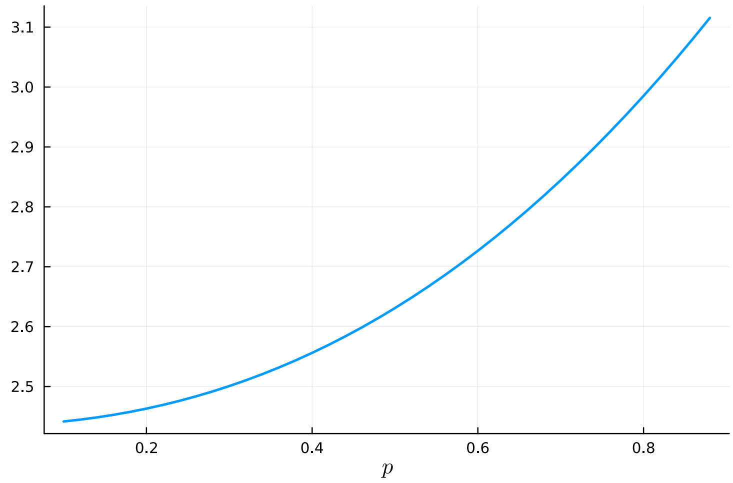

FIGURE 7. Government Reputation p and Aggregate Investment K

The announced rate τ^R is lower than the deviation rate, and it is decreasing in reputation. The driving force is the response of investment to the announcement. When reputation is high, firms place more weight on the announced rate, so a given increase in it deters more investment: the tax base becomes more elastic and the revenue-maximizing rate falls, moving the peak of the Laffer curve to the left. A government with a good reputation therefore competes for capital by announcing less, while one with a poor reputation, whose firms discount the announcement, gains little from cutting and posts a higher rate. Profit-shifting is what makes this investment channel decisive: by capping the deviation it keeps the announcement informative about the rate firms will face, even though the shifting margin itself works in the opposite direction, as the next subsection formalizes. Without profit-shifting the announced rate instead rises with reputation, the case I turn to next.

In comparing the model to the data, one challenge is that the government's type is not observed. Under the commitment type, the optimal rate is lower when reputation is higher, generating investment and profit-shifting patterns consistent with the data. Under the opportunistic type, the realized rate rises with reputation, which is counterfactual. The two-period extension in Section 4 resolves this by inducing the opportunistic type to imitate the commitment type in equilibrium, aligning the realized behavior of both types with the data.

(A) Tax Rate τ (B) Profit-Shifting Portion $\frac{\theta}{\rho(k)}$

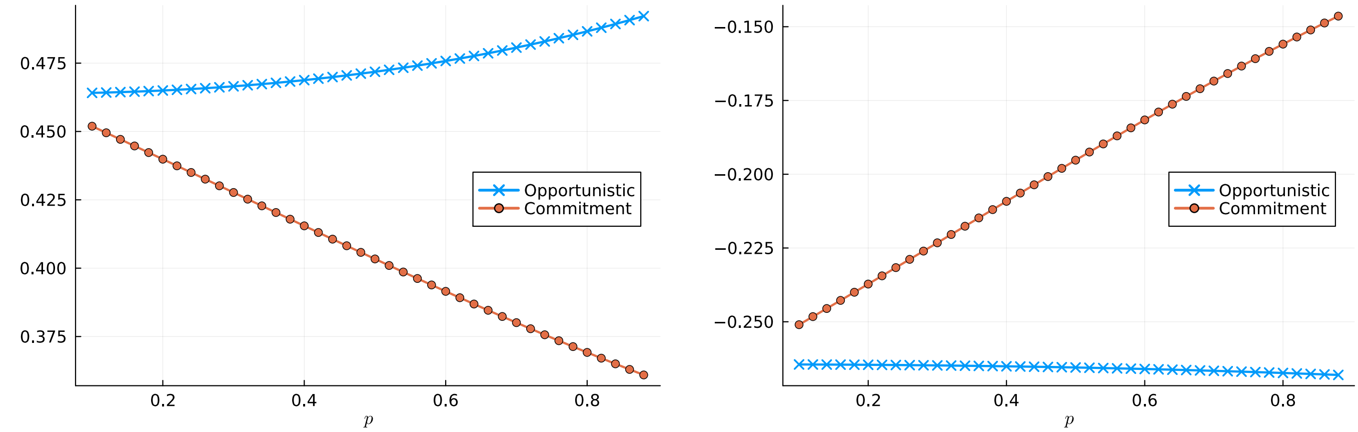

FIGURE 8. Government Reputation and Equilibrium Outcome

3.7 Role of Profit-Shifting

I now isolate what profit-shifting contributes by shutting it down. While profit-shifting is insufficient to deter the opportunistic type's deviation in the static game, it disciplines the commitment type's announcement. Introducing profit-shifting makes the optimal announced rate decreasing in reputation. This is not true in the model without profit-shifting, where the commitment type raises its rate with better reputation, as seen in Figure 9.

(A) Commitment Tax Rate τ^R (B) Investment k

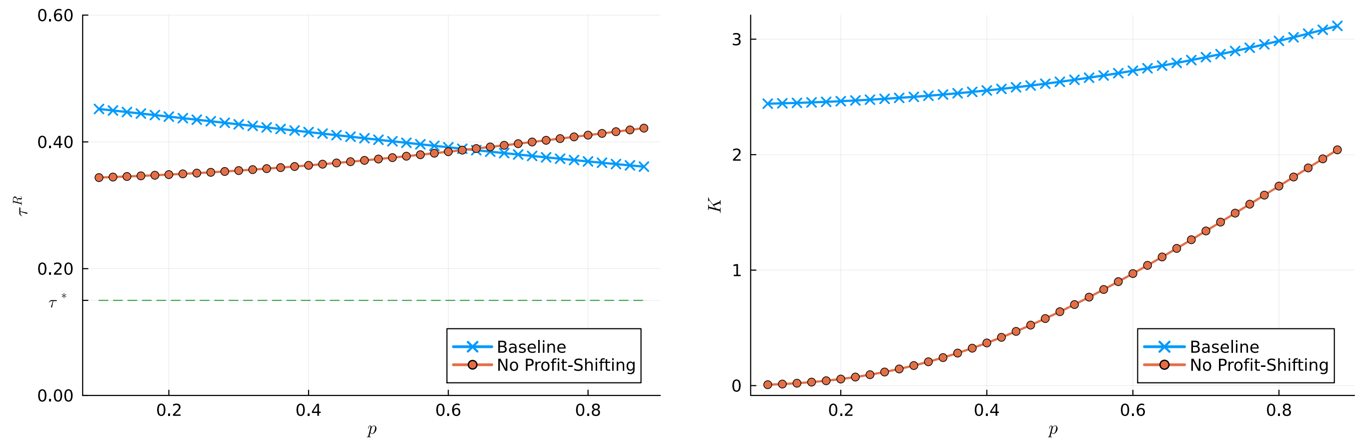

FIGURE 9. Model Results with and without Profit-Shifting

Consider the model without profit-shifting.[15] When firms cannot shift profits, the opportunistic type's re-optimization reduces to maximizing $\tau^O \Lambda(K)$ over a predetermined base, so $\tau^O = 1$ trivially, and nothing prevents full confiscation. Equilibrium investment is accordingly much

[15]Equivalent to $\gamma \to \infty$.

lower (Figure 9), since the downside risk is the loss of all profits. The effect of reputation on the announced rate can be characterized exactly in this case. The government's optimality condition, combined with the firm's, collapses to a locus that does not involve reputation, and the comparative static follows.

Proposition 3. *In the model without profit-shifting, the announced rate satisfies $\tau^R = T(K)$, where T does not depend on p. With $\rho(k) = zk^\alpha$, this is exactly*

$$\tau^R = \frac{1-\alpha}{1-\alpha x}, \qquad x = K/\bar{k}, \tag{11}$$

$\tau^R \in (1-\alpha, 1)$, and $d\tau^R/dp > 0$ for every $p \in (0,1)$, with no further conditions. With $r^* > 0$, the same conclusions hold under the strengthening of the capital-endowment restriction to $\alpha z \bar{k}^{\alpha-1} > (2-\alpha)r^*$.

Proof. See A.4. ■

The strengthened restriction for $r^* > 0$ holds with ample slack at the calibration. The economics of the proposition is the hold-up problem operating through the announcement. Sunk capital is exposed to ex-post extraction, so without profit-shifting only a government that firms expect to honor its announcement attracts capital. Better reputation raises the expected return to investing because the probability of full confiscation falls, the base expands steeply, and the government then exploits the enlarged base with a higher announced rate.

Introducing profit-shifting changes the sign of this comparative static, and it is worth isolating why. Write the commitment government's revenue along the equilibrium as $R = \tau^R\big[\Lambda(K) + \Theta - \frac{\gamma}{2}\Theta^2/\Lambda(K)\big]$, the Stage-1 objective, with K solving the firm's condition and the anticipated deviation rate $\tau^O(K)$ substituted, so that derivatives in τ^R are total and carry the deviation-rate feedback. At an interior optimum $\partial R/\partial \tau^R = 0$, and differentiating this identity in p gives

$$\frac{\mathrm{d}\tau^R}{\mathrm{d}p} = -\frac{\partial^2 R/\partial \tau^R \partial p}{\partial^2 R/\partial (\tau^R)^2}, \tag{12}$$

so under the second-order condition $\partial^2 R/\partial(\tau^R)^2 < 0$ the announced rate falls with reputation exactly when the cross-partial $\partial^2 R/\partial \tau^R \partial p$ is negative. Equivalently, it falls with reputation exactly when the tax base $\Lambda(K) + \Theta - \frac{\gamma}{2}\Theta^2/\Lambda(K)$ becomes more responsive to the rate as reputation rises, moving the peak of the Laffer curve in Figure 10 to the left. That responsiveness has two parts: an investment response, through how much the announced rate deters investment, and a profit-shifting response, through how much profit it drives abroad at a fixed level of investment. The next proposition signs the second of these.

Proposition 4. *Whenever $\tau^R > \tau^*$, so that profits are shifted out of the host, the profit-shifting contribution to how responsive the base is to the rate,*

$$-\frac{\tau^R}{\Lambda(K) + \Theta - \frac{\gamma}{2}\Theta^2/\Lambda(K)} \frac{\partial \Theta}{\partial \tau^R}\left(1 - \frac{\gamma\Theta}{\Lambda(K)}\right),$$

is strictly smaller the larger is investment K. Since investment rises with reputation, this profit-shifting response raises, rather than lowers, the announced rate as reputation improves, opposing the reversal.

Proof. See A.5. ■

Remark 1. *As $\gamma \to \infty$, $\Theta \to 0$ and $\tau^O \to 1$, so $R \to \tau^R \Lambda(K)$ and the model reduces to the no-profit-shifting benchmark of Proposition 3, in which $\mathrm{d}\tau^R/\mathrm{d}p > 0$. This is the one limit in which the sign of the comparative static is unambiguous.*

The investment response is what reverses the schedule. When reputation is high, firms weight the announced rate more, and the deviation rate they anticipate rises with the base they install, so the announced rate deters more investment and the base grows more responsive to it. The profit-shifting response works the other way: as reputation improves and home investment grows, home profit rises relative to the foreign base, so a given increase in the rate drives out a smaller share of profit and the base responds less through shifting, not more, which is Proposition 4. The reversal therefore requires the investment response to overcome the profit-shifting one. This is why it is a property of the calibrated model rather than an unconditional theorem: the two responses carry opposite signs, and no condition on the primitives alone settles which dominates. Remark 1 fixes the one unambiguous limit, in which shifting is shut down and the schedule provably rises with reputation. The discrimination between the two models does not rest on resolving this contest: Proposition 3 fixes the no-profit-shifting slope as a theorem, and Proposition 1 fixes the deviation magnitudes.

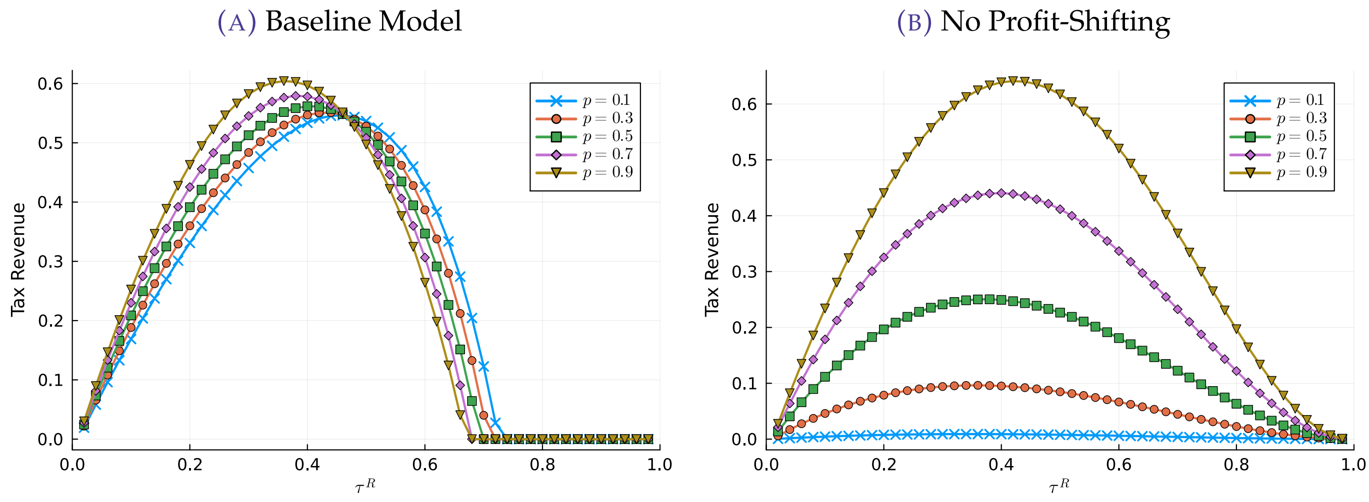

FIGURE 10. Laffer Curves under Different Values of Reputation p

3.8 Implication on Tax Competition

Although the model abstracts from strategic interaction between governments, it has an implication for tax competition. The optimal domestic rate increases in the foreign rate. Figure 11 displays the optimal rates of both types as the foreign rate varies, at $p = 0.5$ and $\gamma = 1.25$. When the foreign tax rate rises, firms invest more in the host country and shift less profit out, so the base expands and the government optimally raises its rate, consistent with the tax-competition literature (Devereux et al., 2008).

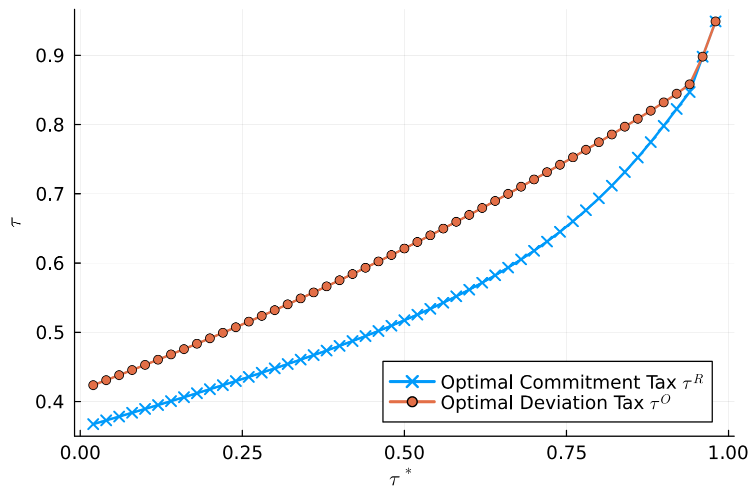

FIGURE 11. Optimal Tax Rate as a Best Response to Foreign Tax Rate

3.9 Sensitivity Analysis

Before extending the model, I verify that the comparative statics do not hinge on the specification of the shifting cost.

Convexity of the shifting cost. I check whether the results are sensitive to the convexity of the shifting cost by considering $\frac{\gamma}{2}\left(\frac{\theta}{\rho(k)}\right)^{\phi}\theta$ for $\phi \in \{1, 3, 5\}$, where $\phi = 1$ is the baseline. Figure 12 shows that the optimal announced rate decreases and the deviation rate increases in reputation under all three, as in the baseline.

(A) Commitment Tax Rate τ^R (B) Deviation Tax Rate τ^O

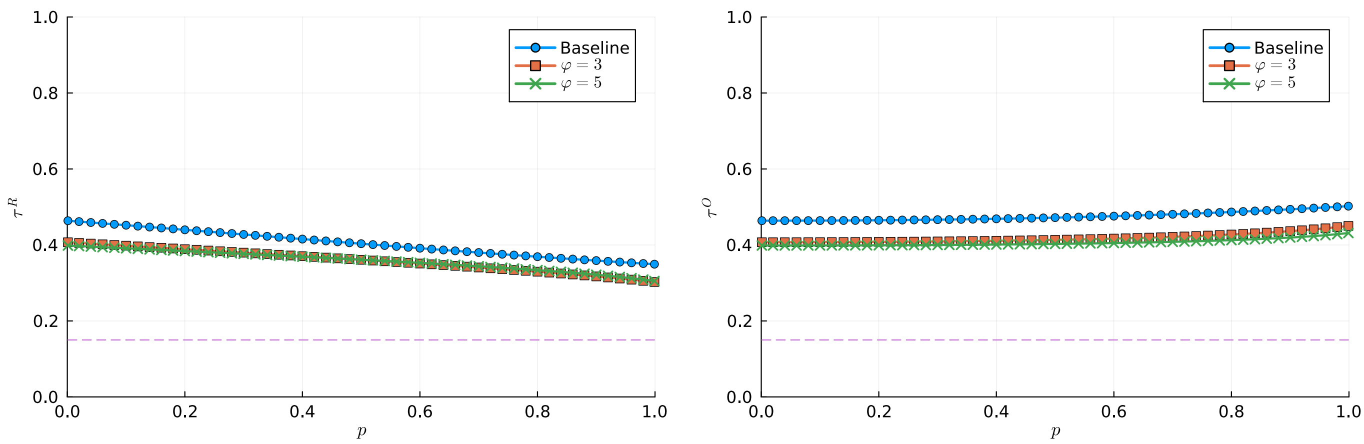

FIGURE 12. Optimal Tax Rate under Different Levels of Convexity

Costliness of shifting. I also impose $\gamma \in \{0.5, 1.25, 1.5\}$. Both types impose higher rates under higher γ, because firms have less recourse to shifting when it is expensive. Investment is correspondingly lower, while the shifted fraction of profits is non-monotone in γ (Figure 13).

Asymmetric shifting costs. The baseline assumes the same γ in both countries. Allowing the foreign cost γ^* to differ changes only the Stage-4 solution, which becomes $\theta = (\tau^* - \tau)\rho(k)\rho(\bar{k} - k)/[\gamma(1-\tau)\rho(\bar{k} - k) + \gamma^*(1-\tau^*)\rho(k)]$. A higher foreign cost dampens shifting and weakens its disciplining role, but the equilibrium remains qualitatively similar to the baseline.

Non-deductible shifting costs. When shifting costs are not deductible from taxable income, the shifting rule is linear in the tax rate and the deviation rate has the exact solution $\tau^O = \tau^*/2 + \gamma(A+B)/(2B)$, where $A = \rho(K)$ and $B = \rho(\bar{k} - K)$, so the distance from confiscation is

$[(2-\tau^*)B - \gamma(A+B)]/(2B)$. All disciplining results carry over with simpler algebra. A variant with costs deductible only in the origin country instead makes the deviation rate independent of investment and eliminates the margin, so I do not adopt it.

Beyond revenue maximization. The baseline government maximizes revenue, appropriate when the multinational sector is foreign-owned. Suppose instead that a fraction ϕ of after-tax host profits accrues to domestic residents, so the government maximizes $[\phi + (1-\phi)\tau]$ times the tax base $\Lambda(K) + \Theta - \frac{\gamma}{2}\Theta^2/\Lambda$, a weighting that penalizes over-taxation. Raising ϕ from zero to one-half lowers the whole announced-rate schedule but overturns neither mechanism: the announced rate stays strictly decreasing in reputation and the deviation rate stays interior, in fact further from confiscation, at every ϕ. The revenue-maximizing case is thus the conservative one for the disciplining role of profit-shifting, and the mechanism does not depend on the government valuing only revenue.

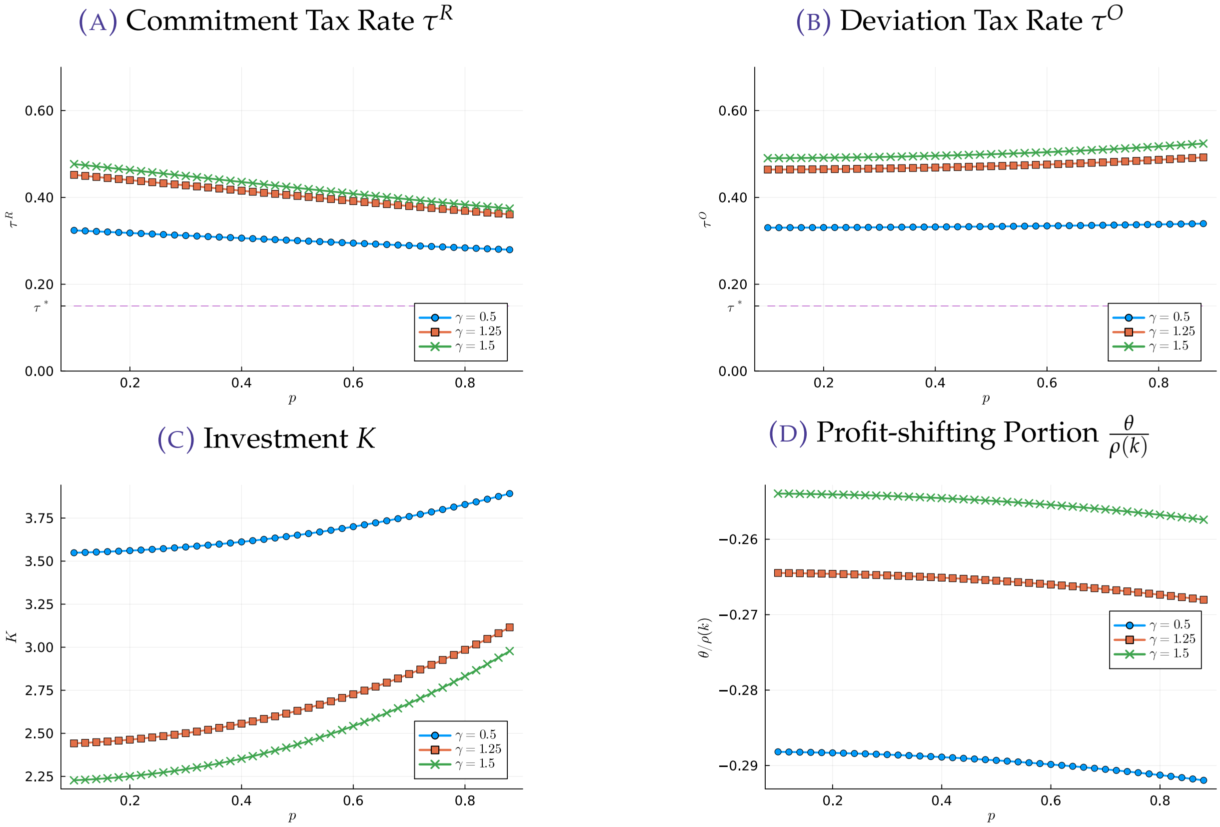

FIGURE 13. Equilibrium Results under Different γ

4 Extensions

The static model conveys the mechanism, but it treats the deviation as mechanical. The opportunistic type reneges whenever it moves, so honesty is an assumption about types rather than an outcome, and nothing in the model says whether reputation can discipline behavior. The two-period extension makes honoring the announcement a choice by giving the opportunistic type a reputation to protect, and asks what sustains honesty. The answer is the substantive result of this section, and it turns on profit-shifting. Profit-shifting lowers the cost of keeping the opportunistic type honest, because it caps the gain from reneging, while the benchmark without profit-shifting keeps a low-reputation government honest only at announcements approaching confiscation. A second extension uses the same machinery to study investment treaties.

4.1 Two-Period Model with Endogenous Reputation

I extend the baseline model to two periods, where the government's deviation in the first period is punished by a fall in investment in the second as its reputation deteriorates. The stage game repeats the static game. For simplicity, capital fully depreciates each period, so the periods are linked only through beliefs.[16] After the first period, firms update beliefs about the government's type from the observed tax rate, as in Dovis and Kirpalani (2020, 2021).

4.1.1 Government Reputation Dynamics

I first describe how reputation evolves. Let $\sigma \in [0,1]$ denote the probability that the opportunistic government does not deviate and mimics the commitment type in period 1. The probability that the announced rate is imposed is

$$\pi(p) = \Pr(\tau_1 = \tau_1^R) = p + \sigma(1-p), \tag{13}$$

and reputation in the second period follows Bayes' rule,

$$\pi'(p) = \frac{p}{p+\sigma(1-p)}, \tag{14}$$

[16] Allowing capital to accumulate would make firms even more reluctant to invest under bad reputation without changing the qualitative implications.

after an observed commitment. If the opportunistic type deviates in period 1, its type is revealed and reputation falls to zero. The prospect of preserving reputation disciplines the opportunistic type's first-period behavior.

4.1.2 The Opportunistic Government's Problem

The opportunistic government chooses its deviation rates and the probability σ of honoring the announcement to maximize the discounted sum of revenues, taking the announced rate τ_1^R and firms' beliefs as given. Write $R(\pi, \tau)$ for period revenue at belief π under realized rate τ. The opportunistic type's expected revenue decomposes as

$$\sigma\Bigg[\underbrace{\beta\left(R\left(\pi', \tau_2^{RO}\right) - R\left(0, \tau_2^{OO}\right)\right)}_{\text{Reputation Gain}} - \underbrace{\left(R\left(\pi, \tau_1^{O}\right) - R\left(\pi, \tau_1^{R}\right)\right)}_{\text{Deviation Gain}}\Bigg] + R\left(\pi, \tau_1^{O}\right) + \beta R\left(0, \tau_2^{OO}\right), \tag{15}$$

where τ_2^{RO} and τ_2^{OO} are its second-period deviation rates after commitment and after revelation, respectively, and β is the discount factor. If the reputation gain exceeds the deviation gain, the opportunistic type commits. If it falls short, the type deviates. If the two are equal, it mixes. Patience raises the reputation gain, so commitment is easier to sustain under high β. The announced rate affects the deviation gain non-monotonically, since the best-response deviation is itself non-monotone in τ_1^R (Section 3.2). Two remarks on reading (15) are in order. Both terms of the deviation gain are evaluated at the same sunk first-period investment, the one firms chose at belief π under the announcement, while the continuation revenues are evaluated at the investment induced by the corresponding second-period belief. And the opportunistic type takes the beliefs π and π' as given when choosing σ. Consistency then requires the beliefs to be generated by the chosen σ, so mixing arises from indifference rather than from the type internalizing the effect of its randomization on the posterior.

4.1.3 Firms and the Commitment Government

Firms invest given belief $\pi = p + (1-p)\tilde{\sigma}$ over the first-period rate, where consistency requires their conjecture $\tilde{\sigma}$ to equal the opportunistic type's optimal σ. The commitment type chooses τ_1^R and τ_2^R subject to these responses as implementability constraints, in the spirit of Dovis and

Kirpalani (2021). It effectively selects, among the consistent continuation equilibria indexed by τ_1^R, the one with the highest discounted revenue. The equilibrium concept is perfect Bayesian equilibrium. Beliefs follow Bayes' rule on the path, off-path announcements are attributed to the opportunistic type, and the commitment type's choice among the continuation equilibria is the selection device. Depending on τ_1^R, the opportunistic type deviates, commits, or mixes. The commitment type thus chooses between signalling, meaning a rate the opportunistic type declines to honor so that types are revealed at the cost of first-period revenue, and hiding, meaning a rate the opportunistic type prefers to honor so that both types behave alike. The full taxonomy of deviation, mixing, and commitment zones of τ_1^R, together with the Laffer-curve analysis underlying it, is in Appendix B.

4.1.4 Equilibrium

For reputation and patience not too low, the commitment type hides its type and the opportunistic type honors the announced rate in period 1. Figure 14 displays the commitment type's announced rate and the probability σ that the opportunistic type honors the announcement, across reputation p and for $\beta \in \{0.3, 0.6, 0.9\}$. The announced rate decreases in both p and β. Higher patience raises the reputation gain and lets the commitment type keep the opportunistic type honest with a lower and less distortionary announcement. Higher reputation loosens the same constraint.

Two distinct forces make the announced rate fall with reputation here, and it is worth separating them, since only the first is present in the static model. The first is the elasticity channel of Section 3.7: at a higher reputation firms place more weight on the announcement, the tax base is more responsive to it, and the revenue-maximizing rate falls. The second is the credibility constraint just described: sustaining honesty at a low reputation requires an announcement high enough that the opportunistic type prefers not to deviate, and both reputation and patience relax that requirement. The two work through different margins, the first through the position of the Laffer peak and the second through the lower edge of the set of announcements the opportunistic type will honor, but both lower the announced rate as credibility improves, and the schedule in Figure 14 combines them. Profit-shifting is necessary for each: it delivers the elasticity channel by capping the deviation, and Proposition 5 below shows that without it the credibility constraint can be met at low reputation only at announcements that leave no revenue. The treaty extension of Section 4.2

isolates the second channel, since a penalty on deviation moves the constraint without changing reputation.

(A) Commitment Tax Rate τ_1^R (B) Probability to Commit σ

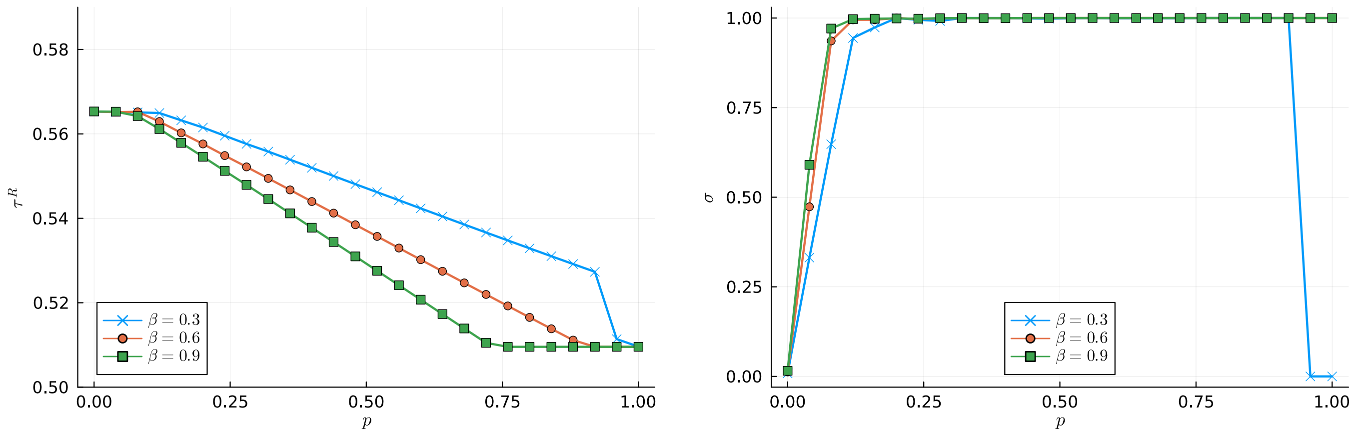

FIGURE 14. Government's Choice and Reputation in Two-period Model

Two properties of the equilibrium follow directly and hold with or without profit-shifting. Both types announce the same rate, by the same pooling argument as in the static game (A.6), so the statutory rate corresponds to the announcement, and the effective–statutory gap opens only when the opportunistic type deviates, with expected size $(1-\pi)(\tau_1^O - \tau_1^R)$. Because firms' belief π is at least p, the expected gap is at most $(1-p)$ times the conditional gap, so it is small when reputation is high and vanishes when reputation is perfect. Neither property separates the two models. What separates them is what it costs a government to make its announcement credible, and the following proposition shows that the answer depends on profit-shifting.

Proposition 5. *(i) Without profit-shifting, as $p \to 0$ every announcement that the opportunistic type honors with probability one either approaches confiscation or commands a tax base that vanishes, and first-period revenue under any such announcement converges to zero. Conditional on a deviation the effective–statutory gap is the entire untaxed profit $1 - \tau_1^R$.*

(ii) With profit-shifting, an equilibrium in which the opportunistic type always honors the announcement, with zero gap and strictly positive revenue, exists at every reputation level.

Proof. See A.7. ■

Proposition 5 states the clearest difference between the two models. Without profit-shifting, a government with low reputation can make its announcement credible only by setting rates so close to confiscation that revenue vanishes.[17] With profit-shifting, credibility need not be costly. There always exists an announcement at which deviating offers the opportunistic type no gain, because the deviation best response crosses the announced rate at an interior point. The gap in the profit-shifting model is therefore small, vanishing at high reputation, while the benchmark generates confiscatory gaps at low reputation. Part (ii) is an existence statement, not a claim about the selected equilibrium. The commitment type selects the revenue-maximizing continuation, and it may still prefer an announcement that tolerates occasional deviation; the computed equilibria below display small positive gaps at low reputation. Profit-shifting therefore lowers the cost of credibility rather than guaranteeing honesty.

Solving the two-period model without profit-shifting quantifies this contrast. Table 6 compares the two equilibria at $\beta = 0.6$. At low and intermediate reputation the benchmark cannot sustain honesty, deviation occurs on the path, and because a deviating government confiscates, the expected gap is enormous. It equals 31 percentage points at $p = 0.1$ and only approaches zero as reputation becomes high. The announced rate is also uniformly higher and displays no declining schedule. On both margins the benchmark's predictions run against the patterns in the data. It generates gaps an order of magnitude larger than those in Table 4 and a statutory-rate schedule unlike anything the data suggest. Profit-shifting lowers the cost of commitment, keeps deviations moderate, and produces gaps of the size the data display.

TABLE 6. Two-Period Equilibria with and without Profit-Shifting

Reputation p	0.1	0.3	0.5	0.7	0.9
With profit-shifting					
Announced rate τ_1^R	0.45	0.42	0.40	0.39	0.37
Expected gap (pp)	0.5	0.9	0.8	0.0	0.0
Without profit-shifting					
Announced rate τ_1^R	0.50	0.60	0.58	0.52	0.46
Expected gap (pp)	31.0	8.8	1.0	0.2	0.0

Note: Two-period equilibrium values at $\beta = 0.6$. The expected gap is $(1-\pi)(\tau_1^O - \tau_1^R)$ in percentage points, where $\tau_1^O = 1$ without profit-shifting.

A simple threshold describes when the gap disappears entirely. Honoring the announcement

[17] The two-period truncation understates the reputation stake relative to an infinite horizon, where the punishment is the capitalized loss of future investment. Lengthening the horizon scales the stake by a constant factor but leaves its vanishing rate as $p \to 0$ unchanged, so this conclusion is not an artifact of truncation.

at the unconstrained revenue peak is sustainable at high reputation exactly when the discount factor exceeds $\bar{\beta} = D(\tau_L)/[R_2^O(1) - R_2^O(0)]$, the ratio of the one-period gain from deviating at the peak to the difference in second-period revenue between perfect reputation and a revealed type.[18] At the calibration $\bar{\beta} = 0.55$, consistent with the computed equilibria, in which the gap vanishes at high reputation for $\beta \in \{0.6, 0.9\}$ but not for 0.3. For less patient governments the commitment type instead prefers a low announcement that the opportunistic type declines to honor, so that types separate. Tax holidays and special-zone rates offered by impatient governments in risky environments are consistent with this case, and the model predicts they are transitional, since reputations diverge in the following period.

Firms' choices follow the same logic. Because the opportunistic type commits and the announced rate decreases in reputation, investment increases and outward profit-shifting decreases in reputation under both types, unlike in the static model, where the opportunistic type's behavior was counterfactual. Figure 15 displays first-period aggregate investment and the expected shifted fraction of profits.

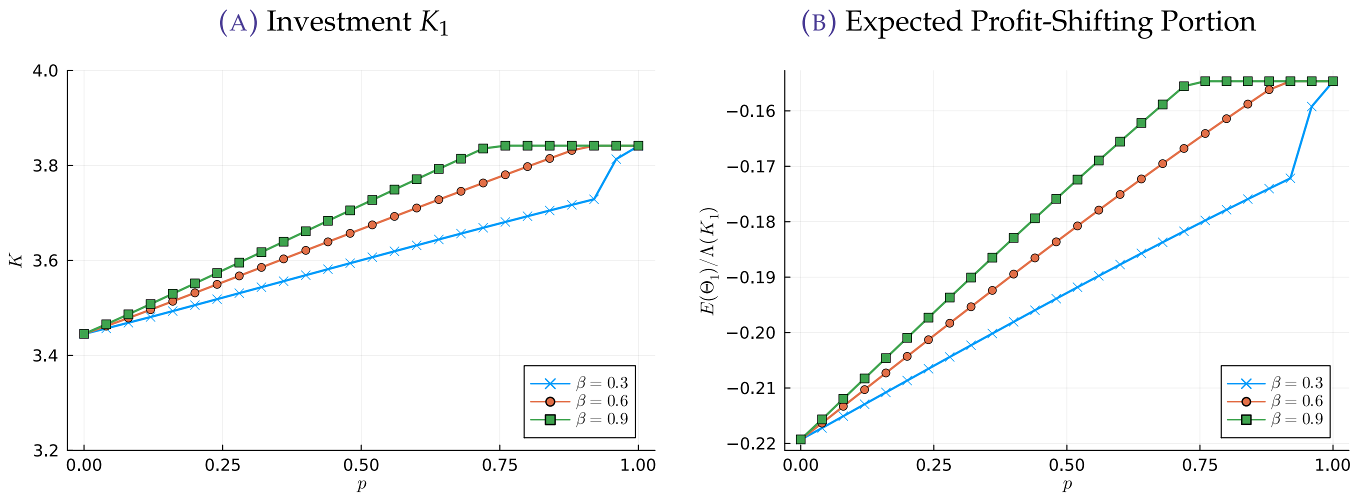

FIGURE 15. Firms' Choice and Reputation in Two-period Model

Taking stock, the two-period model is consistent with each of the facts of Section 2, and it is worth being precise about what that consistency rests on. The evidence in Section 2 is descriptive, so I do not read it as a formal test of either model. The useful observation is that the two models make very different predictions for the same descriptive patterns. On the statutory margin, the benchmark's announced rate rises with reputation by Proposition 3, so risky countries should set

[18] A.8 contains the derivation.

the lowest rates, while in Table 3 statutory rates are higher in riskier countries in the raw cross-section and never significantly lower in any specification. The declining schedule of Figure 14 is a property of the calibrated profit-shifting model, and it is the sign the raw gradient displays. On the gap margin, the profit-shifting model keeps the gap within a few percentage points, vanishing at high reputation, by the interior deviation rate and the bound on the deviation probability, while the benchmark puts the conditional gap at the entire untaxed profit, far beyond anything in Table 4. The FDI and profit-shifting facts are shared predictions. The calibrated equilibrium reproduces their signs, but so would a wider class of models, and I draw no comparison from them.

4.2 Investment Treaties

The facts of Section 2 contain no treaty variation, so this section is not about explaining them. Its purpose is to draw out the policy content of the mechanism. The analysis so far gives the deviation margin a central role. If reneging is what separates the tax outcomes of risky countries from those of safe ones, a policy that raises the cost of reneging should move each of the observables studied above. Bilateral investment treaties with investor-state dispute settlement are such a policy: they give firms legal recourse and monetary compensation when a host government reneges on the terms under which they invested, and Kohler and Stähler (2019) and Horn and Tangerås (2021) study these agreements as devices that compensate investors for harmful ex-post government action. Within the model such an agreement raises the cost of deviating without altering the distribution of types, so it operates through the credibility constraint of Section 4.1 and not through reputation itself. This distinction matters for the empirical design below, in which treaty coverage is interacted with risk rather than treated as a substitute for it. This section asks what the model predicts a treaty does, and then takes a first look at ratification data.

I model a treaty as a penalty $\chi \geq 0$ that a deviating government must pay. The penalty reduces the deviation gain in the trade-off (15) and changes nothing else, which makes its effects transparent. The penalty is paid to a third party rather than to the firms, which isolates the deterrence side of a treaty. An award that compensated investors would in addition insure firms against the deviation state and raise investment directly, a second channel that would reinforce the effects below and that I leave out for transparency.

Proposition 6. *An increase in the penalty χ has the following effects.*

(i) It enlarges the set of announcements the opportunistic type honors and raises the probability that it honors them.

(ii) It lowers the expected effective–statutory gap and raises investment.

(iii) It lowers the announced statutory rate wherever the credibility constraint binds.

(iv) It changes nothing where the announcement was already credible, which is the case at high reputation.

Proof. See A.9. ■

Parts (i) and (ii) are the direct effects one would expect from raising the cost of reneging. Part (iii) is the notable prediction, because it contradicts a natural intuition. One might expect a treaty to let the government tax a now-safer base more heavily. In the model the opposite happens. A low-reputation government must announce a higher rate than it would otherwise choose, because only sufficiently high announcements make deviating unattractive to its opportunistic self. The treaty relaxes this constraint, so the announced rate falls toward the unconstrained optimum. This is the same constraint that Proposition 5(ii) shows can always be satisfied at the crossing rate $\tau_{\times}$: with profit-shifting, credibility is attainable at every reputation level, but attaining it at the revenue-maximizing announcement is not, and the penalty is what closes that distance. Part (iv) gives the cross-sectional prediction. The treaty matters only where the credibility constraint binds, so all of its effects are concentrated in risky countries. At the calibration with $\beta = 0.6$, raising χ from zero to 0.05 raises the average probability of honoring the announcement from 0.92 to 1.00, eliminates the average gap, raises investment by eleven percent, and lowers the mean announced rate from 0.403 to 0.365, with an announced-rate response of 7.5 percentage points at $p = 0.02$ and zero at $p = 0.98$.

The model predicts that treaty ratification raises investment, compresses the effective–statutory gap, and lowers statutory rates, all concentrated in high-risk countries (Proposition 6). I take the statutory prediction to the data, where the accumulated treaty record is long enough to have power.

A bilateral investment treaty is an agreement between two countries granting each other's investors legal protections, the most consequential of which is access to investor-state dispute settlement: an international arbitration mechanism through which a firm can sue a host government for monetary compensation when it reneges on the terms under which the firm invested. I collect the universe of these treaties from the UNCTAD International Investment Agreements Navigator, which records 3,448 bilateral investment treaties with their signature and entry-into-force dates, of which 2,777 have entered into force. Because both parties gain access to dispute settlement once a treaty enters into force, I treat each entry into force as an event for both signatories. I date a ratification episode as a year in which a country brings at least one treaty into force after three years without any, the annual counterpart of the re-event exclusion of Section 2.2.5; this yields 487 episodes across 179 countries between 1962 and 2023. For the cross-section I measure a country's treaty coverage as the log of one plus the number of its treaties in force in 2016. Statutory rates come from the Tax Foundation panel for 1980 through 2023.

The statutory margin shows the predicted heterogeneity in the timing of ratification. Around each episode I estimate a distributed-lag specification with country and year fixed effects, regressing the statutory rate on leads and lags of the episode indicator and clustering by country, separately for the riskiest and safest quintiles of average expropriation risk. Figure 16 traces the estimated path. In the riskiest quintile the statutory rate falls by roughly two to three percentage points in the years after an episode; in the safest quintile it does not fall, and if anything rises by about a percentage point. This is the sign and the concentration in risky countries that parts (iii) and (iv) predict. The riskiest quintile also drifts down in the year before the episode, consistent with anticipation, since a treaty is signed and publicly known before it enters into force; the joint pre-trend test is marginal ($p = 0.08$), so I read the timing as consistent with the prediction rather than as an identified effect.

The same prediction shows in the accumulated stock, read through the cross-section, since part (iii) concerns the configuration that dense coverage should produce. In the 2016 cross-section of Section 2 I interact expropriation risk with treaty coverage. Table 7 reports the result. The interaction is negative: the statutory–risk gradient declines as coverage accumulates. Conditioning on risk sharpens it. Among countries in the riskier half of the sample, statutory rates decline with treaty coverage, by 1.8 percentage points per log point of coverage with a standard error of 0.7,

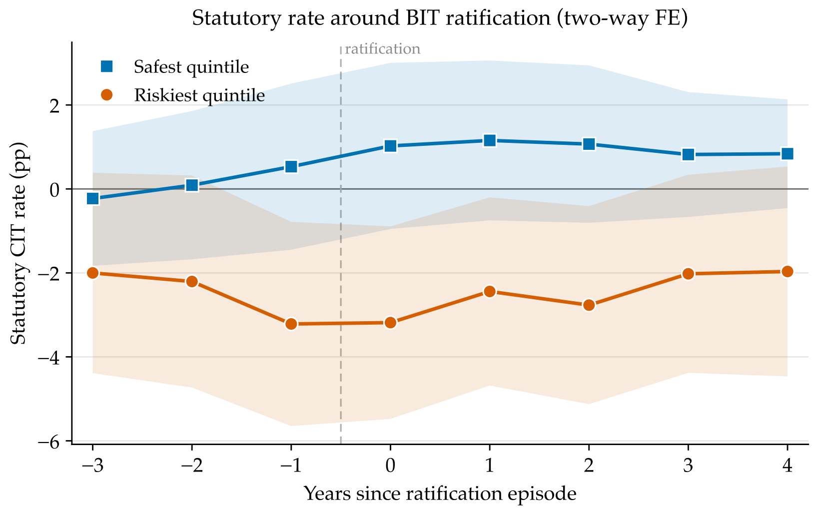

FIGURE 16. Statutory corporate income tax rate around BIT ratification episodes, from a distributed-lag specification with country and year fixed effects, estimated separately for the riskiest and safest quintiles of average expropriation risk. Coefficients are relative to years outside the event window, with 90 percent country-clustered confidence bands. An episode is a year in which a country brings at least one treaty into force after three quiet years.

while no such relation appears among safe countries—again the heterogeneity of parts (iii) and (iv). Two cautions apply: nearly every country holds some treaties by 2016, so the level of the risk gradient at low coverage is not separately identified, and the pattern weakens when tax havens are excluded, where the interaction is -0.58 with a standard error of 0.80. A treaty stock accumulated over five decades is not a ratification experiment, so I read the table as a configuration consistent with the prediction rather than as an estimate of a treaty effect.

TABLE 7. Statutory Rate, Expropriation Risk, and Treaty Coverage

	Statutory rate
Expropriation Risk	5.584**
	(2.694)
Treaty coverage	1.134
	(1.251)
Risk × Coverage	−2.040**
	(0.927)
N	123

Note: Dependent variable is the statutory corporate income tax rate, 2016 cross-section, heteroskedasticity-robust standard errors. Treaty coverage is the log of one plus the number of bilateral investment treaties in force in 2016, from the UNCTAD IIA Navigator. Controls are log GDP per capita and GDP growth. ** denotes significance at the 5% level.

5 Conclusion

This paper studies the relationship between government reputation and the optimal corporate tax rate through a Bayesian game between multinational firms and a government that is either a commitment type or an opportunistic type. The motivating fact is that the countries with the worst reputations, which receive the least FDI and lose the most profit to shifting, set statutory tax rates that are no lower than those of safe countries, and in the raw data higher. The model adds government reputation and firms' profit-shifting to the capital taxation framework of Chari et al. (1988). While reputation has counteracting effects on the optimal announced rate, profit-shifting makes firms' investment responsive enough to the announcement that the commitment type optimally lowers the statutory rate as its reputation improves, generating more investment and less outward profit-shifting. Profit-shifting also disciplines the deviation itself. I prove that a reneging government stops short of confiscation by an explicit margin, so tax rate gaps remain at the few percentage points observed in the data. The model without profit-shifting instead predicts a rising statutory schedule and confiscatory gaps, and neither pattern appears in the data. In the two-period extension, the opportunistic type honors its announcement to preserve its reputation, and profit-shifting lowers the cost of doing so, while in the benchmark commitment at low reputation requires giving up nearly all revenue. Because both types announce the same rate, statutory rates trace announcements while the effective–statutory gap traces deviations, and the data match both margins. Modeling investment treaties as penalties on deviation yields the prediction that treaties lower statutory rates precisely in risky countries, the opposite of the base-protection intuition. Statutory rates fall after treaty ratification in risky countries but not safe ones, and decline with treaty coverage among risky hosts in the cross-section, patterns in line with the prediction, though anticipation and a fifty-year accumulated stock make this a configuration consistent with the mechanism rather than an identified test.

References

Aguiar, M., & Amador, M. (2011). Growth in the shadow of expropriation. *The Quarterly Journal of Economics*, *126*(2), 651–697.

Amador, M., & Phelan, C. (2021). Reputation and sovereign default. *Econometrica*, *89*(4), 1979–2010.

Azzimonti, M. (2018). The politics of FDI expropriation. *International Economic Review*, *59*(2), 479-510.

Blouin, J., & Robinson, L. A. (2025). Double counting accounting: How much profit of multinational enterprises is really in tax havens? *Contemporary Accounting Research*, *42*(1), 5–40.

Brennan, G., & Buchanan, J. M. (1980). *The power to tax: Analytical foundations of a fiscal constitution*. Cambridge: Cambridge University Press.

Chari, V., Kehoe, P., & Prescott, E. (1988). *Time consistency and policy* (Staff Report No. 115). Federal Reserve Bank of Minneapolis.

Desai, M. A., Foley, C. F., & Hines, J., James R. (2006). Capital controls, liberalizations, and foreign direct investment. *The Review of Financial Studies*, *19*(4), 1433-1464.

Devereux, M. P., Lockwood, B., & Redoano, M. (2008). Do countries compete over corporate tax rates? *Journal of Public Economics*, *92*(5), 1210-1235.

Dovis, A., & Kirpalani, R. (2020). Fiscal rules, bailouts, and reputation in federal governments. *American Economic Review*, *110*(3), 860-88.

Dovis, A., & Kirpalani, R. (2021). Rules without commitment: Reputation and incentives. *The Review of Economic Studies*, *88*(6), 2833-2856.

Edwards, J., & Keen, M. (1996). Tax competition and Leviathan. *European Economic Review*, *40*(1), 113–134.

Enache, C. (2022). *Corporate tax rates around the world* (Tech. Rep.). Tax Foundation. Retrieved from `https://taxfoundation.org/data/all/global/corporate-tax-rates-by-country-2022/`

Fernandéz, A., Klein, M. W., Rebucci, A., Schindler, M., & Uribe, M. (2016). Capital control measures: A new dataset. *IMF Economic Review*, *64*(3), 548–574.

Garcia-Bernardo, J., & Jansky, P. (2021). *Profit shifting of multinational corporations worldwide* (Working Paper No. 14). Prague, Czechia: Institute of Economic Studies.

Hines, J., James R. (2010). Treasure islands. *Journal of Economic Perspectives*, *24*(4), 103-126.

Hines, J., James R., & Rice, E. M. (1994). Fiscal paradise: Foreign tax havens and American business. *Quarterly Journal of Economics*, *109*(1), 149–182.

Horn, H., & Tangerås, T. (2021). Economics of international investment agreements. *Journal of International Economics*, *131*, 103433.

Janeba, E. (2000). Tax competition when governments lack commitment: Excess capacity as a countervailing threat. *American Economic Review*, *90*(5), 1508–1519.

Johannesen, N. (2010). Imperfect tax competition for profits, asymmetric equilibrium and beneficial tax havens. *Journal of International Economics*, *81*(2), 253-264.

Johannesen, N., Tørsløv, T., & Wier, L. (2019). Are less developed countries more exposed to multinational tax avoidance? Method and evidence from micro-data. *The World Bank Economic Review*, *34*(3), 790-809.

Kallen, C. (2023). *Profit shifting and investment by multinationals.*

Keen, M., & Konrad, K. A. (2013). The theory of international tax competition and coordination. In A. J. Auerbach, R. Chetty, M. Feldstein, & E. Saez (Eds.), *Handbook of public economics* (Vol. 5, pp. 257–328). Elsevier.

Klein, M. W., & Rosengren, E. (1994). The real exchange rate and foreign direct investment in the united states: Relative wealth vs. relative wage effects. *Journal of International Economics*, *36*(3), 373-389.

Kohler, W., & Stähler, F. (2019). The economics of investor protection: ISDS versus national treatment. *Journal of International Economics*, *121*, 103254.

Langenmayr, D., & Simmler, M. (2021). Firm mobility and jurisdictions' tax rate choices: Evidence from immobile firm entry. *Journal of Public Economics*, *204*, 104530.

Ortigueira, S., & Pereira, J. (2021). Lack of commitment, retroactive taxation, and macroeconomic instability. *Journal of the European Economic Association*, *20*(1), 264-311.

Peralta, S., Wauthy, X., & van Ypersele, T. (2006). Should countries control international profit shifting? *Journal of International Economics*, *68*(1), 24-37.

Schnitzer, M. (1999). Expropriation and control rights: A dynamic model of foreign direct investment. *International Journal of Industrial Organization*, *17*(8), 1113-1137.

Thomas, J., & Worrall, T. (1994). Foreign direct investment and the risk of expropriation. *The*

Review of Economic Studies, *61*(1), 81–108.

Tørsløv, T., Wier, L., & Zucman, G. (2022). The missing profits of nations. *The Review of Economic Studies*, *90*(3), 1499-1534.

Appendix A Proofs and Derivations

A.1 Proof of Lemma 1

First and second derivatives of $\theta(k, \tau)$ with respect to τ are

$$\frac{\partial\theta}{\partial\tau} = -\frac{1}{\gamma}\frac{(1-\tau^*)\,\rho(k)\rho(\bar{k}-k)\,\left(\rho(k)+\rho(\bar{k}-k)\right)}{\left((1-\tau)\rho(\bar{k}-k)+(1-\tau^*)\,\rho(k)\right)^2}, \tag{A1}$$

$$\frac{\partial^2\theta}{\partial\tau^2} = -\frac{2}{\gamma}\frac{(1-\tau^*)\,\rho(k)(\rho(\bar{k}-k))^2\,\left(\rho(k)+\rho(\bar{k}-k)\right)}{\left((1-\tau)\rho(\bar{k}-k)+(1-\tau^*)\,\rho(k)\right)^3}, \tag{A2}$$

where profits $\rho(k)$ and $\rho(\bar{k}-k)$ are positive, and the denominators are positive on $\tau \in [0, 1]$. Hence both derivatives are strictly negative. ■

A.2 Proof of Proposition 1 and the Bounds on γ

Without profit-shifting the deviator maximizes $\tau\Lambda(K)$ over a predetermined base, which is strictly increasing in τ, so $\tau^O = 1$. With profit-shifting, let $A = \rho(K)$, $B = \rho(\bar{k}-K)$, and $s = \sqrt{1+2\gamma}-1$. Evaluating (2) at $\tau = 1$ gives $\Theta(K, 1) = -B/\gamma$, so the base at confiscation is $\widehat{B}(K, 1) = A - B/\gamma - B^2/(2\gamma A)$, which is negative if and only if $\gamma < \widetilde{\gamma}(K)$ in (5). By Lemma 1 and $\gamma\Theta/A < 1$, the base is strictly decreasing in τ. It is positive at $\tau = 0$ and negative at $\tau = 1$, so there is a unique base-exhaustion point $\bar{\tau}(K) \in (0, 1)$ solving $\widehat{B} = 0$, i.e. $\Theta = A(1-\sqrt{1+2\gamma})/\gamma$. Substituting into (2) and solving,

$$1-\bar{\tau} \;=\; \frac{(1-\tau^*)\,(1-sA/B)}{1+s} \;>\; 0, \tag{A3}$$

which is positive exactly under $\gamma < \widetilde{\gamma}(K)$. The rule (2) is the unconstrained maximizer of the firm's strictly concave shifting problem. For $\tau \geq \bar{\tau}$ the implied host base is negative, so the nonnegativity constraint in the firm's shifting problem binds and the constrained solution sets the host base to zero, making revenue zero rather than negative there. Either way revenue on $[\bar{\tau}, 1]$ is dominated by interior rates with strictly positive revenue, so $\tau^O < \bar{\tau}$. No concavity of the objective is used.

The lower bound on γ. The lower bound ensuring that the base increases in K obtains from the closed form $\gamma_{\mathrm{lo}}(K, \tau) = -\Theta_0^2/(2A^2) - (\Theta_{0,K}/\rho'(K))(1-\Theta_0/A)$, where $\Theta_0 \equiv \gamma\Theta$ does not involve γ. The second-order condition of the opportunistic problem follows from Lemma 1. ■

A.3 Proof of Proposition 2

We prove this by contradiction. First, $\tau^R = \tau^O$ is impossible. At a common rate the two first-order conditions (4) and (10) differ exactly by the investment-internalization term $\tau^R\,\partial\big[\Lambda + \Theta - \frac{\gamma}{2}\Theta^2/\Lambda\big]/\partial K \cdot \partial K/\partial\tau^R$, which is strictly negative because the base increases in K under $\gamma \geq \Gamma_{\text{lo}}$ while $\partial K/\partial\tau^R < 0$, so the two conditions cannot hold at the same rate. Assume then $\tau^O < \tau^R$. By Lemma 1, $\Theta^O > \Theta^R$; by assumption, $\partial K/\partial\tau^R < 0$ and the base is increasing in K. The optimality conditions (4) and (10) then yield

$$\underbrace{\Lambda + \Theta^R - \frac{\gamma}{2}\frac{\Theta^{R2}}{\Lambda}}_{A_1} + \underbrace{\tau^R\left[1 - \frac{\gamma\Theta^R}{\Lambda}\right]\frac{\partial\Theta}{\partial\tau^R}}_{B_1} > \underbrace{\Lambda + \Theta^O - \frac{\gamma}{2}\frac{\Theta^{O2}}{\Lambda}}_{A_2} + \underbrace{\tau^O\left[1 - \frac{\gamma\Theta^O}{\Lambda}\right]\frac{\partial\Theta}{\partial\tau^O}}_{B_2}. \tag{A4}$$

Firms' optimality (2) gives $\gamma\Theta/\Lambda < 1$ and $\partial(\Theta - \frac{\gamma}{2}\Theta^2/\Lambda)/\partial\tau < 0$, so $A_1 < A_2$; and since Θ is strictly concave in τ with $\tau^R > \tau^O$ under the contradiction hypothesis, $\partial\Theta/\partial\tau^R < \partial\Theta/\partial\tau^O < 0$ while $\tau^R(1 - \gamma\Theta^R/\Lambda) > \tau^O(1 - \gamma\Theta^O/\Lambda) > 0$, so $B_1 < B_2$. Hence $A_1 + B_1 < A_2 + B_2$, a contradiction. ■

A.4 Proof of Proposition 3

Without profit-shifting, the firm's condition is $F \equiv p(1-\tau)\rho'(K) - (1-\tau^*)\rho'(\bar{k} - K) = 0$ and the government's is $\rho(K) + \tau\rho'(K)K_\tau = 0$ with $K_\tau = p\rho'(K)/F_K$, $F_K < 0$. Substituting $(1 - \tau^*)\rho'(\bar{k} - K) = p(1-\tau)\rho'(K)$ to eliminate p collapses the government's condition to a locus free of reputation,

$$\rho'(\bar{k} - K)\left[(1-\tau)\rho(K)\rho''(K) + \tau\rho'(K)^2\right] + (1-\tau)\rho(K)\rho'(K)\rho''(\bar{k} - K) = 0, \tag{A5}$$

which solves to $\tau = T(K) \equiv |N|/(|N| + \rho'(K)^2\rho'(\bar{k} - K))$ with $|N| = \rho(K)(|\rho''(K)|\rho'(\bar{k} - K) + \rho'(K)|\rho''(\bar{k} - K)|) > 0$; in pure power this reduces to (11).

For the slope with $r^* = 0$: the firm's condition gives $p(1 - \tau^R) = (1 - \tau^*)(x/(1-x))^{1-\alpha}$, and the locus gives $1 - \tau^R = \alpha(1-x)/(1-\alpha x)$. Substituting and solving for p,

$$p = \frac{1 - \tau^*}{\alpha}\Phi(x), \qquad \Phi(x) = \frac{(1-\alpha x)\,x^{1-\alpha}}{(1-x)^{2-\alpha}}, \tag{A6}$$

whose logarithmic derivative, upon cross-multiplication by $x(1-x)(1-\alpha x) > 0$, has numerator $(1-\alpha)[1 + x(1-\alpha)] > 0$. Hence $p \leftrightarrow x$ is a strictly increasing bijection; since T is strictly increasing in x, τ^R is strictly increasing in p, unconditionally. For $r^* > 0$ it suffices that $T'(K) > 0$:

decomposing $g \equiv |N|/(\rho'(K)^2\rho'(\bar{k}-K)) = g_1 + g_2$ and substituting $u = zK^{\alpha-1}$, the derivative of $g_1 = \alpha(1-\alpha)u(u-r^*)/(\alpha u - r^*)^2$ in u has numerator $r^*[r^* - (2-\alpha)u] < 0$ under the endowment restriction; $\rho(K)/\rho'(K) = K\psi(u)$ with $\psi' < 0$ is increasing in K; and with $w = \bar{k} - K$, $d\ln(|\rho''(\bar{k} - K)|/\rho'(\bar{k}-K))/dw$ has numerator $(2-\alpha)r^* - \alpha z w^{\alpha-1}$, negative for all $w \leq \bar{k}$ exactly under $\alpha z \bar{k}^{\alpha-1} > (2-\alpha)r^*$. Hence $g' > 0$ and $T' > 0$.

It remains to sign dK/dp, which the argument above uses in both cases. Substituting $\tau = T(K)$ into the firm's condition gives

$$p = \frac{(1-\tau^*)\,\rho'(\bar{k}-K)}{(1-T(K))\rho'(K)}. \tag{A7}$$

The numerator is strictly increasing in K, since ρ' is strictly decreasing and $\bar{k} - K$ falls with K; the denominator is strictly positive, since $T \in (0,1)$ and $\rho' > 0$ under the endowment restriction, and strictly decreasing in K, since T is strictly increasing and ρ' strictly decreasing. Hence p is a strictly increasing function of K along the equilibrium locus, so $dK/dp > 0$ and $d\tau^R/dp = T'(K)\,dK/dp > 0$. ■

A.5 Proof of Proposition 4 and Remark 1

Proposition 4. The response of the base to the rate at fixed K is $\partial/\partial\tau\big[\Lambda(K) + \Theta - \frac{\gamma}{2}\Theta^2/\Lambda(K)\big] = (\partial\Theta/\partial\tau)\big(1 - \gamma\Theta/\Lambda(K)\big)$, with $\partial\Theta/\partial\tau < 0$ by Lemma 1 and $1 - \gamma\Theta/\Lambda(K) > 0$. Substituting Θ from (2) and $\partial\Theta/\partial\tau$ from Lemma 1 into the quantity displayed in the proposition, a direct computation shows it depends on K only through the ratio $\rho(\bar{k}-K)/\Lambda(K)$ and, when $\tau^R > \tau^*$, is strictly increasing in that ratio: with $\tau^R > \tau^*$ profit flows outward, and a larger foreign base relative to the home base—a larger $\rho(\bar{k}-K)/\Lambda(K)$—makes shifting more responsive to the rate. The ratio $\rho(\bar{k}-K)/\Lambda(K)$ is strictly decreasing in K, since the numerator falls and the denominator rises, so the quantity is strictly decreasing in K. At fixed rates, a higher p shifts weight from the deviation rate τ^O to the lower announced rate $\tau^R < \tau^O$ and so raises the expected after-tax return to investing; the total derivative $\mathrm{d}K/\mathrm{d}p$, which also carries the equilibrium responses of τ^R and τ^O, is positive throughout the computed equilibria (Figure 7). Hence the profit-shifting response decreases in p.

Remark 1. By (2), $\Theta = (\tau^* - \tau)\Lambda(K)\rho(\bar{k}-K)/\big(\gamma[(1-\tau)\rho(\bar{k}-K) + (1-\tau^*)\Lambda(K)]\big)$ is of order $1/\gamma$, so as $\gamma \to \infty$ both $\Theta \to 0$ and $\frac{\gamma}{2}\Theta^2/\Lambda(K) = \frac{\Lambda(K)}{2\gamma}\big(\gamma\Theta/\Lambda(K)\big)^2 \to 0$; the base tends to

$\Lambda(K)$ and $R \to \tau^R \Lambda(K)$. The deviator then maximizes $\tau\Lambda(K)$ over a predetermined base, giving $\tau^O \to 1$, and the firm's problem reduces to the no-profit-shifting problem. Proposition 3 then gives $\mathrm{d}\tau^R/\mathrm{d}p > 0$. ■

A.6 Pooling at the Announcement Stage

Claim. No equilibrium exists in which an announcement is made by the opportunistic type alone, so no announcement reveals the government's type; the pooled announcement below is an equilibrium and survives the standard refinements.

No revealing candidate survives. In any candidate equilibrium, an announcement made only by the opportunistic type carries belief zero by Bayes' rule, and at belief zero firms expect the deviation rate regardless of the announced rate, so investment is a constant K_0 and the announcer's payoff is its deviation revenue on K_0. Announcing instead a rate that the commitment type uses with positive probability yields a strictly positive belief and strictly higher investment, because the firm's after-tax home return is decreasing in the realized rate and the announced rate lies below the anticipated deviation rate by the argument of Proposition 2. Deviation revenue is increasing in the base under $\gamma \geq \Gamma_{\mathrm{lo}}$, so the opportunistic type strictly gains by mimicking, and no equilibrium assigns it an announcement the commitment type does not use. The opportunistic type's announcement therefore lies in the support of the commitment type's, and no announcement identifies the type; since its payoff is strictly increasing in the belief the announcement induces, it places all of its mass on the most favorable announcement in that support. I study the pooling equilibrium in which that announcement is the commitment type's optimal rate.

The pooled equilibrium and refinements. Assign off-path beliefs that any announcer of $\tau' \neq \tau_1^{R\star}$ is opportunistic, so the period runs at belief zero. The opportunistic type's Stage-3 option set is identical under any announcement, and announcements matter only through beliefs and hence investment. Announcing the pooled rate yields belief $\pi = p + \sigma(1-p) > 0$ and a larger base, with a weakly better continuation, so mimicry at the announcement stage is weakly dominant, and the commitment type strictly prefers its optimum. Sequential rationality under the stated beliefs is immediate. For the Intuitive Criterion, note that for any off-path τ' the most favorable belief raises investment to $K(1,\tau')$, which raises the opportunistic type's payoff at least as much as the committed type's, since it deviates on the larger base. Hence no announcement is equilibrium-dominated for the opportunistic type while attractive to the committed one, and the pessimistic off-path beliefs are admissible. ■

A.7 Proof of Proposition 5

The proof of part (i) rests on a scaling property of investment when the effective return on home investment is small.

Lemma 2. *Let $n = \alpha/(1-\alpha)$. Under the capital-endowment restriction there exist constants $0 < c' \leq C' < \infty$ and $\eta > 0$ such that, whenever $\pi(1-\tau) \leq \eta$,*

$$c' \left(\pi(1-\tau)\right)^n \;\leq\; \Lambda\big(K(\pi,\tau)\big) \;\leq\; C' \left(\pi(1-\tau)\right)^n. \tag{A8}$$

Proof. Without profit-shifting the firm's condition is $\pi(1-\tau)\rho'(K) = (1-\tau^*)\rho'(\bar{k}-K)$. As $K \to 0^+$, the right side converges to $(1-\tau^*)\rho'(\bar{k})$, which is finite and strictly positive under the endowment restriction, while $\rho'(K) = \alpha z K^{\alpha-1} - r^*$ is dominated by its first term. Solving for K gives two-sided bounds proportional to $(\pi(1-\tau))^{1/(1-\alpha)}$. Near zero, $\rho(K) = zK^\alpha - r^* K$ is dominated by zK^α, and substituting the bounds on K yields the claim with $n = \alpha/(1-\alpha)$. ■

Proof of Proposition 5(i). Without profit-shifting, $\tau^O \equiv 1$, so conditional on a deviation the realized gap is $1-\tau_1^R$, and $R_2^O(0) = 0$ because at prior zero firms invest nothing at home. Suppose the opportunistic type honors announcement τ with probability one. Beliefs then satisfy $\pi = 1$, so consistency requires the deviation gain not to exceed the reputation gain, $(1-\tau)\Lambda(K(1,\tau)) \leq \beta R_2^O(p)$, and Lemma 2 bounds the right side by $\beta C' p^n$. If $\tau \leq \frac{1}{2}$, the constraint forces $\Lambda(K(1,\tau)) \leq 2\beta C' p^n$, so the tax base vanishes and first-period revenue $\tau\Lambda$ is at most of order p^n. If $\tau > \frac{1}{2}$, consider first the announcements with $1-\tau > \eta$, on which Lemma 2 does not apply. There $\Lambda(K(1,\tau)) \geq \Lambda(K(1,1-\eta)) \geq c'\eta^n$, because $\Lambda(K(1,\cdot))$ is decreasing in the rate, so the constraint requires $\eta\, c'\eta^n \leq \beta C' p^n$, which fails for all p below a threshold; no such announcement is consistent with full mimicry once p is small. On the remaining range $1-\tau \leq \eta$, Lemma 2 gives $\Lambda(K(1,\tau)) \geq c'(1-\tau)^n$, so the constraint forces $(1-\tau)^{n+1} \leq (\beta C'/c')\, p^n$, hence $1-\tau \leq C''\, p^{n/(n+1)}$ and the announcement approaches confiscation, with revenue $\tau\Lambda \leq \Lambda \leq C'(1-\tau)^n \leq C'''\, p^{n^2/(n+1)}$. In either case first-period revenue converges to zero as $p \to 0$. ■

Proof of Proposition 5(ii). Continuity of $K(1,\cdot)$ follows from the implicit function theorem applied to the firm's first-order condition, whose derivative in k is strictly negative, and continuity of $\tau^O(\cdot)$ follows from the maximum theorem together with the strict concavity of the deviation objective on the relevant region, so the composition $\tau \mapsto \tau^O(K(1,\tau))$ is continuous. At $\tau = 0$ deviation revenue is zero at rate zero and strictly positive at interior rates, so $\tau^O(K(1,0)) > 0$. At the other end, the exhaustion point $\bar{\tau}(K)$ of A.2 is bounded away from one on the compact range of $K(1,\cdot)$, over which the maintained bound $\gamma < \tilde{\gamma}(K)$ holds, so for announcements close enough to one the best response satisfies $\tau^O(K(1,\tau)) \leq \bar{\tau} < \tau$. By the intermediate value theorem the map crosses the diagonal at an interior $\tau_\times$ with $K(1,\tau_\times) > 0$. Announcing $\tau_\times$ makes the deviation gain exactly zero, so honoring the announcement is consistent at every p and every β, and revenue $R(1,\tau_\times) > 0$ since the rate is interior and the base positive. ■

A.8 Derivation of the Patience Threshold

At $\sigma = 1$, beliefs are degenerate and the posterior after commitment is p, so honoring the announcement at the peak τ_L is consistent if and only if $D(\tau_L) \leq \beta[R_2^O(p) - R_2^O(0)]$. The left side is independent of p and β. The right side is continuous and nondecreasing in p with limit $\beta[R_2^O(1) - R_2^O(0)]$. If $\beta \geq \bar{\beta}$ the inequality holds at $p = 1$ and, by continuity, on an interval $(\bar{p}(\beta), 1]$, where the unconstrained peak is optimal and the gap is zero. If $\beta < \bar{\beta}$ it fails even at $p = 1$. ■

A.9 Proof of Proposition 6

Write $G(\sigma)$ for the deviation gain minus the penalty χ and minus the discounted reputation gain, evaluated at beliefs $\pi = p + \sigma(1-p)$, so that the opportunistic type honors the announcement when $G \leq 0$. The penalty enters additively, so $\partial G/\partial\chi = -1$. (i) The set of credible announcements $\{\tau_1^R : G(1) \leq 0\}$ is a sublevel set of a χ-shifted function, hence weakly expanding; for interior mixing, implicit differentiation gives $\partial\sigma^*/\partial\chi = 1/(\partial G/\partial\sigma) > 0$ wherever the generator is increasing in σ (deviation gain nondecreasing and reputation gain nonincreasing in σ), and the region statement holds regardless. (ii) Consider announcements below the crossing point $\tau_\times$, where the deviation is upward and $\tau^O > \tau^R$; the constrained optimum of (iii) and the computed equilibria lie in this region. The expected gap is $(1-\pi)(\tau^O - \tau^R)$. At a fixed announcement, a higher χ raises σ by (i) and hence π, which lowers $1-\pi$ and with it the expected gap, and lowers the expected tax rate, raising investment. Combining this with (iii), the induced fall in the announcement raises investment further, while its effect on the conditional gap $\tau^O - \tau^R$ is of ambiguous sign; the statement about the gap is therefore exact wherever the penalty raises σ to one, in which case the gap is zero, as in the computed equilibria reported below. (iii) Where the optimum sits at the binding lower edge $\tau_b(\chi)$ of the full-mimicry set, differentiating $G(1; \tau_b, \chi) = 0$ gives $d\tau_b/d\chi = 1/\partial_\tau D < 0$ below the crossing point $\tau_\times$, where the deviation gain is decreasing in the announcement. (iv) Since $\partial G/\partial\chi = -1$, $G(1)$ is decreasing in χ. Hence if $G(1) < 0$ at $\chi = 0$, so that the credibility constraint is slack and the optimum is the unconstrained revenue peak, the constraint remains slack and the optimum is unchanged for every $\chi \geq 0$, not merely for small penalties. ■

Appendix B Zones of the Two-Period Announcement

Given p and β, the announced rate τ_1^R partitions into zones according to the opportunistic type's best response: a deviation zone at low announcements (deviation gain exceeds reputation gain), a commitment zone at intermediate announcements, and mixing zones in between; at very high announcements the opportunistic type deviates downward to attract profit-shifting, revealing its type. The commitment type's Laffer curve across these zones, the associated investment schedule, and the deviation-gain/reputation-gain trade-off are displayed in Figures A1–A2 for $p = 0.5$, $\beta = 0.6$. The optimal hiding announcement lies in the commitment zone; expansion of the commitment zone (through patience, reputation, or treaty penalties) weakly lowers it.

(A) Commitment Government's Laffer Curve (B) Investment

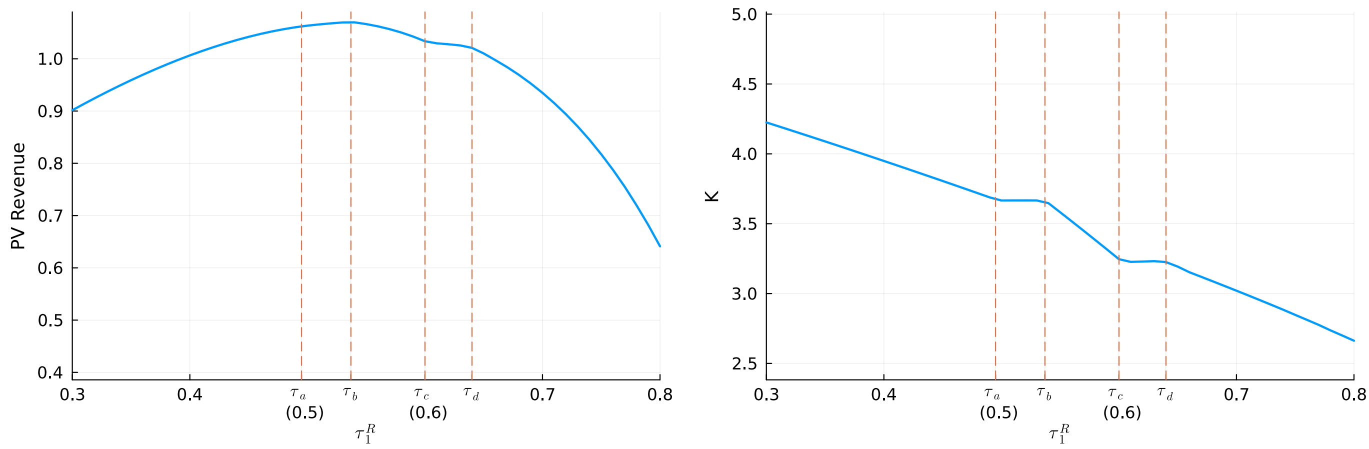

FIGURE A1. Laffer Curve and Investment

(A) Commitment Probability σ (B) Trade-off of Deviation Choice

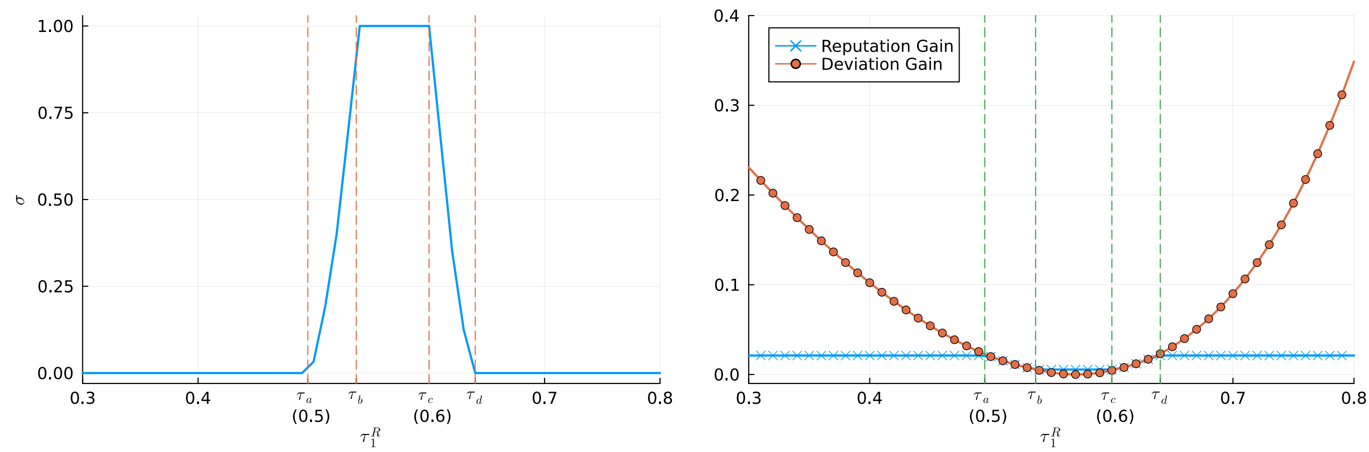

FIGURE A2. Opportunistic Type's Optimal Choice under τ_1^R

Appendix C Additional Tables and Figures

Table A1 lists the tax havens classified by Hines (2010) and Tørsløv et al. (2022).

Table A2 reports summary statistics of the baseline risk measures, and Table A3 lists the countries in the estimation sample of Table 3. Both sub-scores are persistent and their variation is predominantly between countries, which is why the panel estimates lean on between-country variation and event studies, and why I read the within-country column of Table 3 as supportive rather than primary.

TABLE A1. List of Tax Haven Countries

Tax Havens		
Antigua and Barbuda	Aruba	Bahamas
Barbados	Bahrain	Belgium
Belize	Bermuda	Costa Rica
Curacao	Cyprus	Djibouti
Dominica	Fiji	French Polynesia
Grenada	Guyana	Honduras
Hong Kong SAR	Ireland	Jordan
Lebanon	Liberia	Luxembourg
Macao SAR	Maldives	Malta
Marshall Islands	Mauritius	Micronesia
Netherlands	Palau	Panama
Qatar	Samoa	Seychelles
Singapore	Sint Maarten	Solomon Islands
St. Kitts and Nevis	St. Lucia	St. Vincent and the Grenadines
Switzerland	Timor-Leste	Tonga
Turks and Caicos Islands	United Arab Emirates	Vanuatu

TABLE A2. Summary Statistics of Risk Measures (annual, 2001–2021)

	Mean	SD	Between SD	Within SD	AR(1)	*N* [countries]
Expropriation risk	1.04	0.71	0.60	0.38	0.956	2742 [139]
Repatriation risk	1.10	0.79	0.68	0.38	0.980	2742 [139]

Note: Sub-scores recoded so that higher values mean higher risk (0–4 scale). AR(1) is the pooled first-order autocorrelation coefficient.

TABLE A3. Countries in the Estimation Sample

Albania	El Salvador	Liberia	Saudi Arabia
Algeria	Estonia	Libya	Senegal
Angola	Ethiopia	Lithuania	Serbia
Argentina	Finland	Luxembourg	Sierra Leone
Armenia	France	Madagascar	Singapore
Australia	Gabon	Malawi	Slovak Republic
Austria	Gambia, The	Malaysia	Slovenia
Azerbaijan	Germany	Mali	South Africa
Bahamas, The	Ghana	Malta	Spain
Bahrain	Greece	Mexico	Sri Lanka
Bangladesh	Guatemala	Moldova	Sudan
Belarus	Guinea	Mongolia	Suriname
Belgium	Guinea-Bissau	Montenegro	Sweden
Bolivia	Guyana	Morocco	Switzerland
Botswana	Haiti	Mozambique	Syrian Arab Republic
Brazil	Honduras	Myanmar	Tanzania
Brunei Darussalam	Hong Kong SAR, China	Namibia	Thailand
Bulgaria	Hungary	Netherlands	Togo
Burkina Faso	Iceland	New Zealand	Trinidad and Tobago
Cameroon	India	Nicaragua	Tunisia
Canada	Indonesia	Niger	Turkiye
Chile	Iran, Islamic Rep.	Nigeria	Uganda
China	Iraq	Norway	Ukraine
Colombia	Ireland	Oman	United Arab Emirates
Congo, Dem. Rep.	Israel	Pakistan	United Kingdom
Congo, Rep.	Italy	Panama	United States
Costa Rica	Jamaica	Papua New Guinea	Uruguay
Cote d'Ivoire	Japan	Paraguay	Uzbekistan
Croatia	Jordan	Peru	Vietnam
Cyprus	Kazakhstan	Philippines	Yemen, Rep.
Czechia	Kenya	Poland	Zambia
Denmark	Korea, Rep.	Portugal	Zimbabwe
Dominican Republic	Kuwait	Qatar	
Ecuador	Latvia	Romania	
Egypt, Arab Rep.	Lebanon	Russian Federation	

I also verify the empirical facts with the World Bank Government Effectiveness index, which captures the government's commitment to policy and is used as a reputation proxy by Dovis and Kirpalani (2020). Countries with better governance have more FDI inflows, lower statutory rates, less outward profit-shifting, and lower effective tax rates (Figures A3–A4). These figures display the two tax rates separately and therefore speak to their levels. The gap margin, which is their difference, is examined with the ICRG expropriation score in Section 2.2.4 and Appendix D.

(A) log FDI Net Inflows

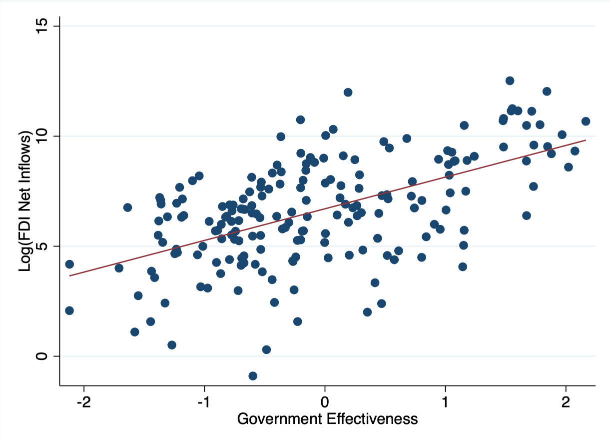

(B) Statutory Corporate Tax Rate

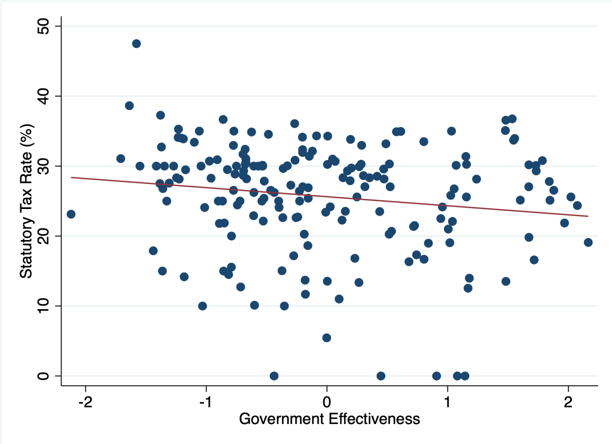

FIGURE A3. Government Effectiveness, FDI Net Inflows, and Statutory Corporate Tax Rate

(A) Profit-Shifting / GDP (%)

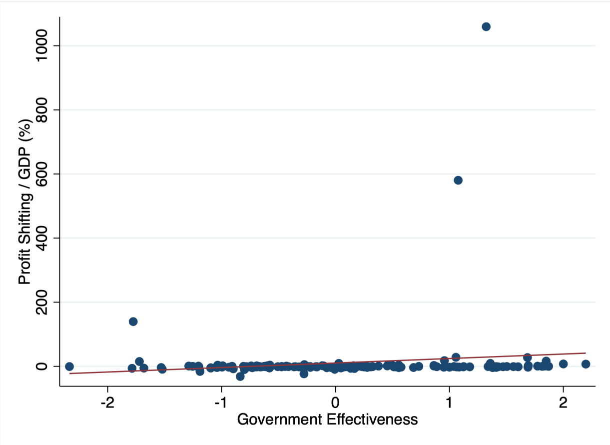

(B) Median Effective Tax Rate (%)

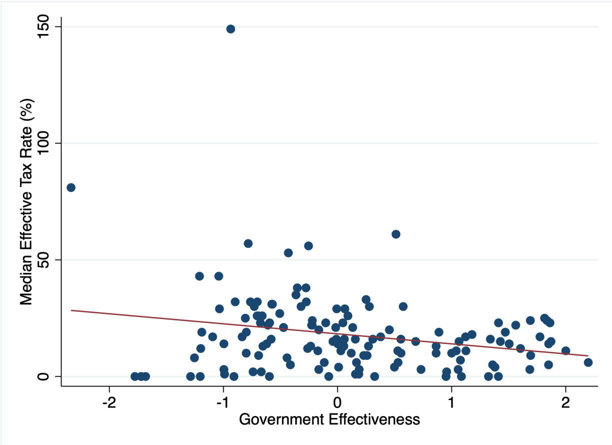

FIGURE A4. Government Effectiveness, Profit-Shifting, and Effective Tax Rate

Appendix D Empirical Robustness

D.1 What the CbCR accrued effective tax rate measures

The constructed gap subtracts the statutory rate from an accrued effective tax rate built from Country-by-Country Reporting (CbCR): for host country i in year t it is the ratio of accrued corporate tax to positive book profit of foreign-controlled affiliates. This subsection derives what that object equals inside the model and shows that it recovers the deviation margin, up to a level shift that the measurement objections govern.

Let $A = \rho(K)$ denote the host's pre-shifting operating profit. An affiliate shifts θ of book profit at cost $\frac{\gamma}{2}\theta^2/A$, so the host taxable base is $\widehat{B} = A + \theta - \frac{\gamma}{2}\theta^2/A$. The realized rate applied to that base is $\widetilde{\tau} = \tau^R$ with probability π and τ^O with probability $1-\pi$. Accrued host tax is the current charge $\widetilde{\tau}\,\widehat{B}$, and the reported CbCR denominator is post-shifting host book profit $\widehat{P}$. The mapping rests on two accounting identities that hold by construction of CbCR: reported book profit equals the tax base, $\widehat{P} = \widehat{B}$, because the shifting cost is a deductible real cost and the model has no book–tax differences; and reported tax accrued equals the current charge, $T = \widetilde{\tau}\,\widehat{B}$.

Measurement result. Under these identities the accrued effective rate equals the realized rate,

$$e \;\equiv\; \frac{T}{\widehat{P}} \;=\; \frac{\widetilde{\tau}\,\widehat{B}}{\widehat{B}} \;=\; \widetilde{\tau}, \tag{A9}$$

because the base cancels identically: e does not depend on A, θ, or γ. Hence the gap $w \equiv e - \tau^R = \widetilde{\tau} - \tau^R$ is a rate object, zero under commitment and $\tau^O - \tau^R$ under deviation, with expectation

$$\mathbb{E}[w] \;=\; \pi\tau^R + (1-\pi)\tau^O - \tau^R \;=\; (1-\pi)(\tau^O - \tau^R). \tag{A10}$$

Letting expropriation risk be x with $\pi = \pi(x)$, $\pi'(x) < 0$, the gap slope is $\mathrm{d}\mathbb{E}[w]/\mathrm{d}x = -(\tau^O - \tau^R)\pi'(x) > 0$: the regression coefficient on risk estimates the marginal deviation margin.

The CbCR identities do not hold exactly in the data, and the referee objections are precisely the gaps between the data and them. Let deductions, losses, withholding, and holidays narrow the taxed base by a factor $1-d_i$ and let intra-group dividend double-counting inflate the reported denominator by $m_i \geq 1$. Then $e_i^{\text{data}} = \widetilde{\tau}_i(1-d_i)/m_i = \widetilde{\tau}_i + L_i$ with $L_i \equiv \widetilde{\tau}_i[(1-d_i)/m_i - 1]$, so the measured gap decomposes as $w_i^{\text{data}} = (\widetilde{\tau}_i - \tau_i^R) + L_i$ and its slope as $-(\tau^O - \tau^R)\pi'(x) + \mathrm{d}\,\mathbb{E}[L \mid x]/\mathrm{d}x$. Identification of the deviation slope requires the measurement term to be a level, $\mathrm{d}\,\mathbb{E}[L \mid x]/\mathrm{d}x \approx 0$: the intercept absorbs the common shift L (which is why the gap is negative in levels, $L < 0$, though the deviation margin is nonnegative), and the coefficient on x recovers the deviation slope insofar as L is flat in risk. Each objection is a statement about $\mathrm{d}L/\mathrm{d}x$. Deductions, losses, and withholding are not systematically ordered by risk, so they move the intercept, not the slope. Intra-group dividend double-counting (Blouin & Robinson, 2025) is concentrated in conduit and holding jurisdictions, which sit at the low-risk end (the conduit set of Table A4), so the induced L_i is most negative where risk is low: it steepens the positive slope rather than generating it. Affiliate selection removes the most heavily assessed affiliates preferentially where risk is high, so $\mathrm{d}\,\mathbb{E}[L \mid x]/\mathrm{d}x > 0$ offsets the deviation slope and the estimate is a lower bound. Resource rents are the one component that plausibly orders in risk: special high-rate resource regimes make $d_i < 0$ exactly in resource-rich, high-risk states, so the resource-rent control is reported, though the cross-section is unaffected and resource nationalism is itself a form of the deviation the model describes.

Table A4 implements these checks. The gap–risk slope survives the dividend and conduit corrections (between three and six points), region and income-group fixed effects, and a per-capita control; it does not survive a country fixed effect (underpowered, since roughly ninety percent of the variance of risk is between countries and half the hosts have no within-country variation over 2017–2021) or a resource-rent control in the large-host panel (the cross-section is unaffected).

Table A5 examines whether the gap margin of Section 2.2.4 depends on particular measurement choices. Each row re-estimates the gap regression of Table 4 after changing one ingredient at a time. The first row computes effective rates from cash taxes paid rather than accrued taxes. The second restricts the panel to the balanced set of 36 parent jurisdictions that report in every year from 2017 to 2021, which removes composition drift in the reporting population. The third replaces the log GDP control with log GDP per capita, the control that absorbs the statutory gradient in Table 3, and the gap margin survives it. The fourth returns to the 2016 cross-section and regresses the weighted-mean effective rate on expropriation risk while controlling for the statutory rate, so that the gap is estimated rather than imposed as the dependent variable. Because the last specification conditions on the statutory rate rather than differencing it out, its magnitude is not directly comparable to the constructed-gap rows. The coefficient on risk is positive and significant in each case, so the association between the effective–statutory gap and expropriation risk is not an artifact of any one of these choices.

Table A6 reports the panel estimates of Tables 1 and 3 with the standardized aggregate risk

TABLE A4. Robustness of the effective–statutory gap to CbCR measurement objections

Specification	Slope on risk	s.e.	N
Baseline			
Cross-section 2016 (all)	$+3.37^{**}$	(1.70)	118
Cross-section 2016 (no havens)	$+3.41^{**}$	(1.72)	99
Panel, large hosts	$+4.92^{**}$	(2.39)	310
Panel, large hosts (no havens)	$+3.97^{*}$	(2.13)	251
(1) Intra-group dividend double-counting			
Cross-section, no conduits	$+3.26^{*}$	(1.74)	106
Panel, no conduits	$+3.90^{*}$	(2.20)	257
Panel, control related-party revenue share	$+4.65^{**}$	(2.35)	310
Panel, low related-party share (< 0.30)	$+6.23^{**}$	(2.52)	167
Panel, dividend-deflated ETR (upper bound)	$+6.04^{**}$	(3.01)	310
(2) Mechanical base differences			
Panel, region fixed effects	$+4.15^{*}$	(2.28)	310
Panel, income-group fixed effects	$+5.05^{**}$	(2.48)	310
Panel, country fixed effects (within)	-1.28	(2.24)	310
Panel, + log GDP per capita	$+4.94^{**}$	(2.38)	306
Cross-section, + resource rents	$+3.51^{**}$	(1.66)	118
Cross-section, + rents + manufacturing	$+3.86^{**}$	(1.66)	115
Panel, + resource rents	$+1.60$	(1.82)	310
Panel, low-resource hosts ($\leq 10\%$)	$+1.78$	(2.16)	270
(3) Selection out of the positive-profit denominator			
Diagnostic: ln(#parents) on risk	-0.16^{*}	(0.08)	332
Panel, stable/growing affiliate populations	$+5.07$	(3.12)	226
Balanced-parent panel	$+5.20^{**}$	(2.48)	309

Note: Dependent variable is the gap (effective minus statutory rate, percentage points); the coefficient is the slope on expropriation risk. Cross-section (Garcia-Bernardo and Jansky (2021) 2016) controls log GDP per capita and GDP growth, HC1 standard errors; the CbCR panel (2017–2021) controls log GDP and year fixed effects, country-clustered standard errors, with "large hosts" the country-years above \$1bn of positive foreign-affiliate profit. The country fixed-effects row is underpowered by construction: about ninety percent of the variance of risk is between countries and nearly half the hosts have no within-country variation over these years. The resource-rent control attenuates the large-host panel because resource-rich states are both riskier and subject to special high-rate resource taxation; the cross-section is unaffected. *, ** denote significance at 10%, 5%.

Table A5. Robustness of the Effective–Statutory Gap

	Coefficient	(s.e.)	N [countries]
Cash-tax (paid) ETR gap, large hosts	$+4.345^{*}$	(2.436)	309 [76]
Balanced parent set, large hosts	$+5.200^{**}$	(2.477)	309 [75]
Log GDP per capita control, large hosts	$+4.816^{**}$	(2.316)	310 [76]
Weighted-mean ETR on risk and statutory rate, 2016	$+6.537^{**}$	(2.712)	118

Note: First three rows use the OECD CbCR panel 2017–2021 with year fixed effects and standard errors clustered by country, with "large hosts" as in Table 4; the first two control for log GDP and the third for 2016 log GDP per capita. The last row uses the 2016 cross-section with heteroskedasticity-robust standard errors and controls for the statutory rate, log GDP per capita, and GDP growth. * and ** denote significance at 10% and 5% levels.

score, which standardizes the investment profile score within each year. The estimates are less precise than the sub-score baselines and do not alter any conclusion in the text.

Table A6. Panel Estimates with the Standardized Aggregate Risk Score

	Statutory rate	log FDI net inflows	FDI net inflows / GDP
	(1)	(2)	(3)
Standardized Risk_{t-1}	-0.0318	-0.0170	0.3095
	(0.1672)	(0.0157)	(0.3182)
Lagged dep. var.		0.9165***	0.7347***
		(0.0119)	(0.0511)
Controls	as in Table 3	—	as in Table 1
Year FE	Yes	Yes	Yes
N	2388	2268	1915
Countries	133	135	119

Note: Annual panel, 2001–2021. Standard errors clustered by country. Column (1) controls for log GDP and GDP growth as in column (2) of Table 3. Column (3) controls for the capital control index, log GDP, GDP growth, and investment growth as in column (2) of Table 1. *** denotes significance at the 1% level.