Coordination across Tax Havens: A Global Games Approach*

Joan Margalef†

June 9, 2025

Abstract

This paper develops a coordination game to study tax evasion across multiple tax havens. I extend the global games framework—a class of coordination games with incomplete information—by introducing multiple tax havens in which investors coordinate. This allows for strategic interaction across jurisdictions, capturing cross-haven coordination effects. I analyze how policies that raise the cost of being a tax haven affect overall evasion. Targeting a single haven can backfire by concentrating evaders elsewhere, increasing overall evasion. In contrast, uniform interventions across jurisdictions are more effective. This mechanism helps explain the limited success of past international efforts and highlights the value of a harmonized approach, as intended under the Global Minimum Tax.

*Margalef acknowledges the financial support from the FPI grant of the Spanish Ministry of Science, Innovation and Universities (PRE2020-093943), and from the Generalitat de Catalunya under the research grant 2021 SGR 00194. All errors are mine.

†Universitat Autònoma de Barcelona and Barcelona School of Economics (joan.margalef@bse.eu)

1 Introduction

Tax evasion has become a major concern in recent decades, as it undermines public revenues and exacerbates inequality (OECD 2024). A key channel, particularly among top-income earners, is using tax havens—jurisdictions that offer low tax rates and strong financial secrecy.[1] Estimates suggest that approximately 8% of global household financial wealth—equivalent to around 10% of world GDP—is held in tax havens (Zucman 2013, 2014; Alstadsæter et al. 2018).[2]

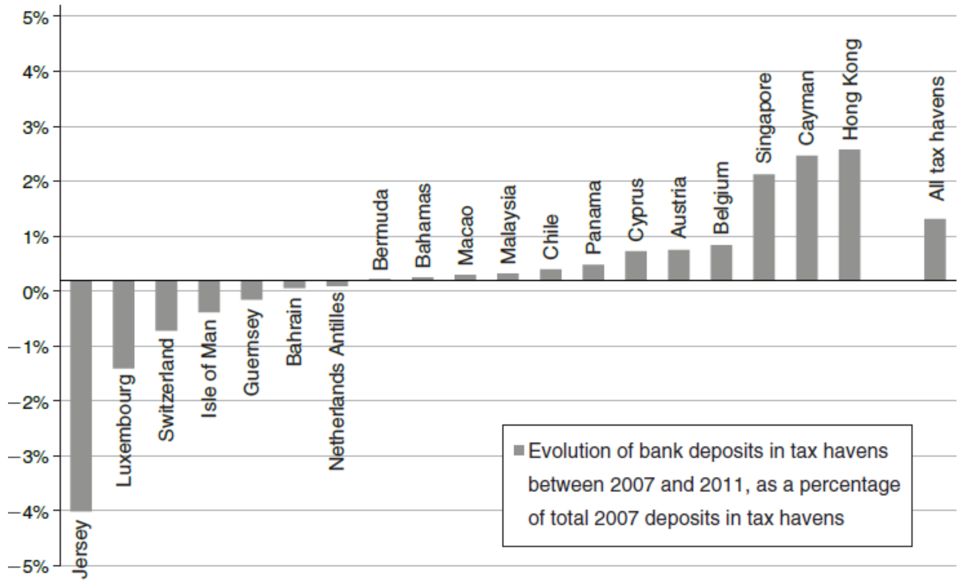

Figure 1: Evolution of foreign-owned deposits in tax havens.

Notes: The figure shows the evolution of foreign-owned deposits in each BIS-reporting tax haven. The comparison is based on the average deposits from the first semester of 2011 and 2007 (except for Cyprus, which began reporting in 2008:IV, and Malaysia, which started in 2007:IV). The difference is expressed as a fraction of the total deposits held in all tax havens in 2007 (2.6 trillion dollars).

Source: Bank for International Settlements (2002–2011), restricted bilateral locational banking statistics (Johannesen and Zucman 2014).

This has prompted supranational organizations such as the OECD and G20 to implement policies to curb these practices. Major initiatives were launched in 2008–2009, which pressured tax havens to sign a series of bilateral treaties of information exchange. However, overall evasion grew despite the international efforts (Johannesen and Zucman 2014). Figure 1 shows the evolution of foreign-owned before and after the implementation of the policies. While some jurisdictions experienced declines, others saw increases. On net, deposits held in tax havens continued to grow. The authors find that net deposits tend to decrease as more treaties are signed by each haven, suggesting that activity moved

1. A tax haven is a jurisdiction that attracts foreign wealth through a combination of low or zero tax rates and limited cooperation in sharing information about foreign asset holders.

2. Zucman (2014) estimates $7.6 trillion in offshore wealth in 2013, corresponding to roughly $190 billion in annual tax revenue losses. Other estimates are higher: Boston Consulting Group (2014) reports $8.9 trillion for the same year, while Henry (2012) suggests the total could reach as much as $32 trillion.

toward less compliant jurisdictions. This outcome reveals a tension between local improvements and global effectiveness, driven by a strategic relocation across jurisdictions.[3]

This paper develops a game-theoretical model to study tax evasion across multiple tax havens, capturing the strategic coordination effects that arise across them.[4] The model builds on the concept of global games, a class of coordination games with incomplete information.[5] The global game structure emphasizes how agents' beliefs, shaped by expectations about others' actions, determine equilibrium outcomes. I then analyze how different policies that raise the cost of being a tax haven affect evader behavior and overall levels of evasion.

The model is a two-stage sequential game with homogeneous investors who can choose to pay taxes in a high-tax country or attempt to evade through one of two tax havens. Agents face incomplete information about the number required to sustain each tax haven. The regime—whether a jurisdiction remains a tax haven or not—depends on the number of evaders who choose to operate through it. In the first stage, investors decide whether to specialize in one tax haven or the other. Specialization can be interpreted as completing all necessary steps to transfer funds into the chosen jurisdiction. It involves no direct cost, but choosing one haven implies losing the opportunity to move to the other. This decision is based on a common public signal. In the second stage, investors receive an additional private signal and choose whether to evade or comply. If enough investors choose to evade, the tax haven remains viable, and evaders benefit relative to compliance. If not enough investors evade, the haven falls, and evaders are reported to their home authorities, facing a punishment greater than if they had complied.

In a partial equilibrium setting, I analyze how different types of policies affect overall evasion. Interventions are modeled as changes to the public signal, resembling economic sanctions that raise the cost of sustaining a tax haven, thereby requiring more evaders. I then conduct a comparative statics exercise to study the effects of targeting one or both havens.

The model captures the strategic interaction between tax havens, making policy outcomes dependent not only on the absolute safety of a haven but also on its relative attractiveness. When one haven is relatively more attractive, investors tend to concentrate there, increasing the likelihood of sustaining it. However, when havens are similar, investors split more evenly, reducing the likelihood of sustaining them. As a result, a policy targeting a single tax haven reduces evasion only if it affects the most attractive one; otherwise, it may even increase it. In contrast, equal punishment to all havens always reduces evasion. This mechanism helps explain the limited effectiveness of the OECD initiatives and provides a rationale for more homogeneous policies, such as the Global Minimum Tax.

3. The coexistence of multiple tax havens with different levels of compliance has long been recognized as a challenge for policy effectiveness. Hines Jr (2005) already estimated that two-thirds of U.S. multinational investment was located in jurisdictions not classified as tax havens by the OECD's 1998 report.

4. The idea of tax evasion as a coordination game is well established in the literature (Bucovetsky 2014; Konrad and Stolper 2016). By lowering its tax rate, a small country may lose little revenue from its domestic base while gaining from attracting foreign taxable wealth. Therefore, the tax haven mechanism for evading is only available if a sufficient number of investors choose to evade. For foundational models of tax competition, see the ZMW model (Zodrow and Mieszkowski 1986; Wilson 1986) and the KK model (Kanbur and Keen 1991), as summarized in Keen and Konrad (2013).

5. Global games were introduced by Carlsson and Van Damme (1993) which often leads to a unique, iterative dominant equilibrium. See Morris and Shin (2001) for a complete explanation.

This paper relates to at least two strands of literature. The first concerns tax evasion, specifically the literature on tax havens. Most existing models do not treat investors' decisions as a coordination problem with endogenous beliefs. Elsayyad and Konrad (2012) analyze a sequential game in which the OECD offers compensation or punishment to tax havens in exchange for ceasing their activities. They find that offers should be made simultaneously rather than sequentially, as remaining tax havens become stronger and more costly to deter. A similar result appears in Slemrod and Wilson (2009), which models tax havens as juridical entrepreneurs selling protection from national taxation. The paper most closely related to this one is Konrad and Stolper (2016), which presents a global game in which investors do not know the fixed cost of a tax haven. While they extend the model to multiple havens, each is treated independently, thereby ruling out coordination effects across tax havens.

The second is the global games literature, which has been used to model phenomena such as speculative attacks on currency pegs (Morris and Shin 1998), bank runs (Goldstein and Pauzner 2005), and revolutions against governments (Angeletos et al. 2007). This model belongs to the class of regime change global games, where payoffs change at a threshold rather than continuously. However, these models have focused on a single entity—one currency, one bank, or one government. In my setting, multiple global games happen simultaneously, and agents, at an ex-ante stage, choose which one they participate in. The first stage captures the broader coordination problem across entities, while the second stage captures the classic coordination problem within the entity.

The contribution of this paper lies in extending global games to a multi-entity setting, enabling the analysis of coordination effects across entities and introducing a layer of strategic interaction that has not been previously studied. This framework proves useful for understanding tax evasion through tax havens. The model generates equilibrium outcomes that would not arise—or would be reversed—in single-haven settings. This helps explain recent evidence on the limited effectiveness of past international efforts and supports more harmonized approaches, such as the Global Minimum Tax.

This paper is structured as follows. Section 2 presents the theoretical model. In Section 3, I perform a comparative statics exercise to analyze how different policies affect overall evasion. Section 4 extends the analysis to a setting with heterogeneous agents. In Section 5, I discuss the policy implications. The paper concludes with a summary of the results.

2 Model

The model is a two-stage sequential game. In the first stage, investors choose which of the two tax havens to specialize in. In the second stage, they decide whether to evade through the haven they selected or comply and pay taxes in their home country.

There is a continuum of homogeneous investors, indexed by $i \in I$, with total mass normalized to one. Each investor owns one unit of mobile capital. The home country taxes this capital at a rate $t \in [0, 1]$, while both tax havens apply a rate $p \in [0, 1]$, with $p < t$. The two tax havens are indexed by

$j \in \{1, 2\}$ and referred to as TH_1 and TH_2.[6]

There is incomplete information about the number of investors required to sustain each tax haven, which is represented by θ_j. θ_j can be interpreted as the economic cost of being a tax haven due to, for example, the loss of domestic revenue and international sanctions. Agents do not know the true value of θ_j but have a common prior belief about it, which follows an independent normal distribution $N(\mu_j, \sigma_\mu)$. Tax havens differ in their μ_j, representing the public signal.

In the first stage, each investor chooses which tax haven to specialize in, enabling evasion through it at a later point. This represents an "access stage", capturing jurisdiction-specific requirements for evasion, such as identifying legal loopholes, obtaining residence or citizenship, or setting up shell companies. Once an investor specializes in a given haven, they cannot use the other. I assume that the cost of specialization is negligible compared to the potential gains from evasion, so it is always profitable to specialize in one of the tax havens before deciding whether to evade or not. Denote the specialization decision of each agent as $s_i = \{1, 2\}$. The proportion of agents specialized in each tax haven is denoted by S_j, s.t. $S_1 + S_2 = 1$.

By specializing in a specific tax haven, investors also gain additional information about the corresponding θ_j. Each receives a private signal $x_{ij} = \theta_j + \varepsilon_i$, where $\varepsilon_i \sim N(0, \sigma_x)$. They also observe the proportion of agents specialized in their chosen haven. Using this information, each agent decides whether to evade or comply and pay taxes in their home country. Let $a_i \in \{\text{Evade}, \text{Comply}\}$ denote the decision of agent i. The share of agents who choose to evade in each tax haven is denoted by A_j, with $A_j \leq S_j$ by construction.

If $A_j \geq \theta_j$, the tax haven country remains as a tax haven and taxes evaders according to p. However, if $A_j < \theta_j$, the tax haven's regime changes and reports evaders to their home country, which punishes evaders by expropriating all the capital. Thus, the payoff from compliance is $1 - t$, while the payoff from evasion is $1 - p$ if the haven survives, and 0 if it does not. The higher the value of θ_j, the more agents are required to sustain the haven. As a result, agents' incentives to evade decrease in both signals and increase in the number of evaders. Their actions are therefore strategic complements.

I denote the information set of each agent by $\mathcal{I}_i^\xi$, where $\xi \in \{1, 2\}$ indicates the stage of the game. In the first stage, all agents have the same public information, so the information set is $\mathcal{I}^1 = \{\mu_1, \mu_2, \sigma_\mu, \sigma_x, p, t\}$. In the second stage, they accumulate information on the number of specialized agents and receive private signals, making the information set differ across agents: $\mathcal{I}_i^2|_{s_i=k} = \{x_{ik}, S_k, \mathcal{I}^1\}$. Figure 2 shows the structure of the game with the two stages and their different information sets.

The policies implemented by international organizations, such as economic sanctions, increase the cost of being a tax haven, raising the number of agents required to sustain it. This is captured by an increase in θ_j. Although agents do not observe the exact value of θ_j, the impact of these policies is

6. The assumption of identical tax rates simplifies the analysis, allows the model to focus on the difference of public signals. In practice, tax haven rates are all close to zero.

reflected in the public signal μ_j. The following comparative statics exercise analyzes how changes in the public signals (μ_1, μ_2) affect the number of evaders in each haven, (A_1, A_2).

I solve the model by Backward induction.

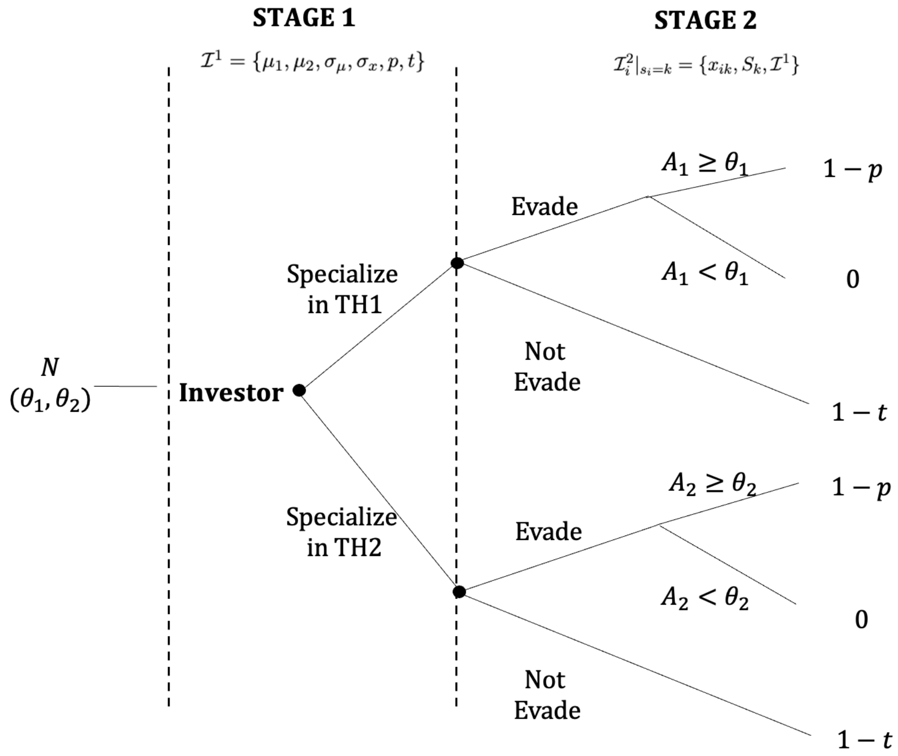

Figure 2: Structure of the game

Notes: The figure illustrates the sequential structure of the game. In the first stage, investors receive public information and choose which tax haven to specialize in. In the second stage, they receive additional private information about their selected haven and decide whether to evade taxes or comply. If the proportion of evaders A_j in a given haven is greater than or equal to θ_j, the haven remains active and taxes evaders at rate p. Otherwise, it collapses and reports evaders to their home country, where they are penalized. Investors who choose to comply are taxed at the standard rate t.
Source: Author's elaboration.

2.1 Stage 2: evasion decision

The decision problem in the second stage takes the form of a global game. Consider the case of an agent who has chosen TH_k.

If the agent chooses to comply, her payoff is $1 - t$. If she chooses to evade, the payoff depends on the aggregate behavior of others: it is $1 - p$ if the selected tax haven survives ($A_k \geq \theta_k$), and 0 if it collapses ($A_k < \theta_k$). Table 1 presents the corresponding payoff structure.

	TH_1 Remains	TH_1 Falls
	$(A_k \geq \theta_k)$	$(A_k < \theta_k)$
Evade	$1-p$	0
Comply	$1-t$	$1-t$

Table 1: Stage 2 payoff table $\forall i : s_i = k$.

Notes: This table presents the payoffs of an agent who has chosen to specialize in TH_k during the second stage.
Source: Author's elaboration.

An investor finds optimal to evade if

$$(1-p)\Pr(A_k \geq \theta_k) \geq 1-t. \tag{1}$$

Each agent forms a belief about the state of the world θ_k using the available signals. Since the risk of being reported increases with the signals, it is strictly dominant to evade when they are sufficiently low. According to this, suppose that they will adopt a *switching strategy* $a(x_{ik})$ based on a threshold value of the private signal $\hat{x}_k$, s.t.

$$a(x_{ik}) = \begin{cases} \text{Evade} & \text{if } x_{ik} \leq \hat{x}_k \\ \text{Comply} & \text{if } x_{ik} > \hat{x}_k. \end{cases} \tag{2}$$

Given a threshold value $\hat{x}_k$, the probability of evading corresponds to the proportion of investors specialized in TH_k who decide to evade, that is,

$$\Pr(x_{ik} \leq \hat{x}_k|\theta) = \Phi\left(\frac{\hat{x}_k - \theta_k}{\sigma_x}\right) = \frac{A_k(\theta_k)}{S_k} \tag{3}$$

where Φ is the cumulative distribution function (CDF) of the standard normal. Since $A_k(\theta_k)$ decreases with θ_k, there exist a unique state of θ_k, say $\hat{\theta}_k$, that is equal to $A_k(\hat{\theta}_k)$. Using this fact, TH_k will survive if

$$\theta_k \leq \hat{\theta}_k = S_k \cdot \Phi\left(\frac{\hat{x}_k - \hat{\theta}_k}{\sigma_x}\right) \tag{4}$$

which characterizes a fixed point.

Given the value of the signals, investors can update their beliefs about θ_k. By Bayesian updating, the posterior belief about θ_k conditional on the signals is normal with mean $(\sigma_x^2\mu_k + \sigma_\mu^2 x_{ik})/(\sigma_x^2 + \sigma_\mu^2)$ and variance $(\sigma_x^2\sigma_\mu^2)/(\sigma_x^2 + \sigma_\mu^2)$. Then, considering Equation 4, the posterior probability of a tax haven surviving is

$$\Pr(\theta_k \leq \hat{\theta}_k|\mathcal{I}_i^2|_{s_i=k}) = \Phi\left(\frac{\hat{\theta}_k - \frac{\sigma_x^2\mu_k + \sigma_\mu^2 x_{ik}}{\sigma_x^2+\sigma_\mu^2}}{\sqrt{\frac{\sigma_x^2\sigma_\mu^2}{\sigma_x^2+\sigma_\mu^2}}}\right). \tag{5}$$

Remember that the incentives to evade decrease as the private signal increases. When the private signal is exactly at the threshold $\hat{x}_k$, an agent, in equilibrium, should be indifferent between evading and complying. Therefore, $\hat{x}_k$ can be pinned down as the value of the private signal that satisfies

$$(1-p)\Phi\left(\frac{\hat{\theta}_k - \frac{\sigma_x^2\mu_k+\sigma_\mu^2\hat{x}_k}{\sigma_x^2+\sigma_\mu^2}}{\sqrt{\frac{\sigma_x^2\sigma_\mu^2}{\sigma_x^2+\sigma_\mu^2}}}\right) = 1-t. \tag{6}$$

Using some algebra, the equilibrium threshold can be expressed as

$$\hat{x}_k^* = \alpha\hat{\theta}_k - \beta\Phi^{-1}\left(\frac{1-t}{1-p}\right) - (\alpha-1)\mu_k \tag{7}$$

where $\alpha = \frac{\sigma_x^2+\sigma_\mu^2}{\sigma_\mu^2}$ and $\beta = \frac{\sigma_x}{\sigma_\mu}\sqrt{\sigma_x^2+\sigma_\mu^2}$.

A monotone equilibrium $\hat{x}_k^*$ is thus identified by solving the system of equations formed by Equation 4 and Equation 7.

Proposition 1 *A Bayesian NE for each tax haven regime change global game exists and is unique if and only if* $\sigma_\mu^2 > \sigma_x/\sqrt{2\pi}$.

The proof is provided in Appendix A.

Given that the equilibrium is defined by fixed points, the model needs to be solved computationally. However, we can state the relation between $\hat{x}_k^*$ and some parameters.

Corollary 1 *The equilibrium threshold $\hat{x}_k^*$ satisfies the following properties:*

- *$\hat{x}_k^*$ and $\hat{\theta}_k$ are complements.*
- *$\hat{x}_k^*$ is increasing in S_k.*
- *$\hat{x}_k^*$ is increasing in t but decreasing in p.*
- *$\hat{x}_k^*$ is decreasing in μ_k.*

The proof is provided in Appendix A. The equilibrium threshold $\hat{x}_k^*$ determines investors' incentives to evade; a higher threshold increases the likelihood of evasion, as it expands the set of private signals (x_{ik}) for which evasion is optimal. The first point reflects complementarity: a higher likelihood of evasion requires a higher state $\hat{\theta}_k$ for the tax haven to undergo a regime change, and conversely, a higher regime-change threshold encourages more evasion. The second point is intuitive: as more investors specialize in a tax haven, the mass of potential evaders increases, thus raising the likelihood of evasion. The third point captures comparative statics related to tax rates: a higher tax rate in the high-tax country (t) makes evasion more attractive, while higher tax haven rates (p) reduce incentives to evade. Finally, an increase in the public signal (μ_k) reduces the incentives to evade since it implies a higher expected value of θ_k.

2.2 Stage 1: specialization decision

In this stage, agents choose which tax haven to specialize in. The outcome is a specialization distribution (S_1^*, S_2^*) that is consistent with agents' optimal responses, based on their inference about the second stage. At the time of the decision, agents do not know the specialization proportions—which are determined endogenously—nor their private signals x_{ij}, which are revealed later. All agents share the same information set $\mathcal{I}^1 = \{\mu_1, \mu_2, \sigma_\mu, \sigma_x, p, t\}$. Then, the belief about θ_j conditional on the information set is normally distributed with mean μ_j and variance σ_μ^2, while for the signal x_{ij} is normally distributed with mean μ_j and variance $\sigma_\mu^2 + \sigma_x^2$.

Denote the expected payoff of specializing in TH_j as Π_j. An agent i will specialize in TH_1 if

$$\Pi_{i1}(\hat{x}_1^*, \hat{\theta}_1 | \mathcal{I}^1) \geq \Pi_{i2}(\hat{x}_2^*, \hat{\theta}_2 | \mathcal{I}^1). \tag{8}$$

with

$$\Pi_{ij}(\hat{x}_j^*, \hat{\theta}_j | \mathcal{I}^1) = (1-p)\Pr(x_{ij} \leq \hat{x}_j^* \ \cap \ \theta_j \leq \hat{\theta}_j | \mathcal{I}^1) + (1-t)\Pr(x_{ij} > \hat{x}_j^* | \mathcal{I}^1). \tag{9}$$

The specialization decision depends on the thresholds $(\hat{x}_j^*, \hat{\theta}_j)$, which determine the likelihood of successfully evading. These thresholds are functions of the tax haven parameters, as defined in Equations 4 and 7. The only factors differentiating expected payoffs are agents' beliefs about the specialization proportions (S_1, S_2) and the public signals (μ_1, μ_2).

Lemma 1 *The expected payoff of specializing in a TH_j (Π_{ij}) increases with the number of agents who specialize in it (S_j) and decreases in its public signal (μ_j) if*

$$(1-p)\frac{\sqrt{\sigma_x^2 + \sigma_\mu^2}}{\sqrt{2\pi}\sigma_\mu \sigma_x} e^{\left(-\frac{1}{2(1-\rho^2)}\right)} \int_{-\infty}^{\hat{\theta}_j(S_j,\mu_j)} e^{\left(-\frac{1}{2(1-\rho^2)}\left(-\frac{2\rho(\hat{x}_j(S_j,\mu_j)-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j - (1-t) > 0 \tag{10}$$

The proof is provided in Appendix A. The inequality formalizes a sufficient condition under which the strategic complementarity of agents' choices holds. As S_j increases, the thresholds $\hat{x}_j$ and $\hat{\theta}_j$ also rise making evasion more likely since there are more potential evaders (Corollary 1). This raises the expected payoff by increasing the likelihood of successful evasion but simultaneously lowers it by reducing the likelihood of compliance. Since these probabilities are not exact complements—one is based on the joint distribution and the other on the marginals—the decrease in the first probability does not necessarily match the decrease in the second. The overall effect of a change in the thresholds on the expected payoff depends on the relative magnitude of the changes in the probabilities weighted by the tax rates. The condition ensures that the evasion payoff dominates.

Conversely, an increase in μ_j lowers both thresholds, reducing the likelihood of evasion and thus the expected payoff.

With a continuum of agents, individual decisions have no impact on aggregate specialization proportions. Therefore, an agent's best response is only a function of the specialization distribution and the public signals. Let $\beta_i(S_1, S_2, \mu_1, \mu_2) = s_i$ denote the best response of agent i.

In equilibrium, the specialization proportions must be consistent with these best responses. Following the global games approach, assume that the strategy s_i is a switching rule that maps public signals into specialization choices, i.e., $s_i(\mu_1, \mu_2) \in \{1, 2\}$.

Definition 1 *A strategy profile $\{s_i^*, a_1^*, a_2^*\}_{i \in I}$ constitutes a **Perfect Bayesian Nash Equilibrium (PBNE)** of the two-stage game if:*

- ***(Stage 2 optimality)** For each agent i, given their information set $\mathcal{I}_i^2$, the evasion strategy $a_j(x_{ij})$ is a best response in the global game solved in Stage 2. That is, a_j^* solves the regime change problem with thresholds $(\hat{x}_j^*, \hat{\theta}_j)$.*

- ***(Stage 1 optimality)** For each agent i, given their information set $\mathcal{I}^1$ and the second-stage equilibrium thresholds $(\hat{x}_1^*, \hat{\theta}_1, \hat{x}_2^*, \hat{\theta}_2)$, the specialization strategy s_i is a best response, i.e.,:*

$$s_i^*(\mu_1, \mu_2) = \beta_i(S_1^*, S_2^*, \mu_1, \mu_2) = \arg\max_{j \in \{1,2\}} \Pi_{ij}(\hat{x}_j^*, \hat{\theta}_j | \mathcal{I}^1)$$

with the equilibrium proportions satisfying:

$$S_j^* = \int_I \mathbf{1}\{s_i^*(\mu_1, \mu_2) = j\}\, di \quad \textit{for } j \in \{1, 2\}$$

Before introducing equilibrium strategies, it is useful to understand how public signals influence best responses. Given Lemma 1, if μ_1 is sufficiently low compared to μ_2, the expected payoff from specializing in TH_1 dominates for all values of (S_1, S_2). This arises because changes in $\mu_j \in (-\infty, \infty)$ can shift the threshold $\hat{x}_j^*$ across the entire real line, while $S_j \in [0, 1]$ has a much more limited influence on expected payoffs. Conversely, if μ_1 is sufficiently high, specialization in TH_2 dominates.

Accordingly, the set of equilibria is determined by two threshold values, $\underline{\mu_1}(\mu_2)$ and $\overline{\mu_1}(\mu_2)$, such that:

- If $\mu_1 < \underline{\mu_1}$, specializing in TH_1 strictly dominates for all (S_1, S_2), yielding a unique equilibrium where all agents choose TH_1 ($S_1^* = 1$, $S_2^* = 0$).

- If $\mu_1 > \overline{\mu_1}$, specializing in TH_2 strictly dominates for all (S_1, S_2), yielding a unique equilibrium where all agents choose TH_2 ($S_1^* = 0$, $S_2^* = 1$).

- If $\mu_1 \in [\underline{\mu_1}, \overline{\mu_1}]$, the equilibrium outcome depends on (S_1, S_2), generating multiple equilibria. In this case, there exists a unique threshold $\hat{S}_1(\mu_1, \mu_2) \in [0, 1]$ that makes agents indifferent between the two tax havens s.t.:

– If $S_1 > \hat{S}_1$, then $\Pi_{i1} > \Pi_{i2}$, agents in TH_2 deviate, and the equilibrium converges to $(S_1^* = 1, S_2^* = 0)$.

– If $S_1 < \hat{S}_1$, then $\Pi_{i1} < \Pi_{i2}$, agents in TH_1 deviate, and the equilibrium becomes $(S_1^* = 0, S_2^* = 1)$.

– If $S_1 = \hat{S}_1$, then $\Pi_{i1} = \Pi_{i2}$, agents have no incentives to deviate, and the equilibrium is interior: $(S_1^* = \hat{S}_1, S_2^* = 1 - \hat{S}_1)$.

When the difference between the public signals is small, full coordination in the tax haven with the higher signal may still yield a higher expected payoff due to the strategic complementarity in specialization decisions. In other words, the mass of agents choosing the same haven can compensate for the unfavorable public signal. The threshold $\hat{S}_1(\mu_1, \mu_2)$ captures the precise share of agents required in TH_1 to make others indifferent between the two havens. This threshold is decreasing in $\mu_1 - \mu_2$: the more favorable the public signal of TH_1 relative to TH_2, the fewer agents are needed in TH_1 to sustain indifference.

Among all strategies that can support the equilibria described above, we focus on three symmetric switching strategies. The strategies are defined as follows:

- **Strategy 1 (Full coordination)** Agents specialize in the tax haven with the lower public signal. In the event of a tie, they coordinate on one tax haven. Without loss of generality, assume ties are broken in favor of TH_1:

$$s_i(\mu_1, \mu_2) = \begin{cases} \text{Specialize in } TH_1 & \text{if } \mu_1 \leq \mu_2 \\ \text{Specialize in } TH_2 & \text{if } \mu_1 > \mu_2. \end{cases} \tag{11}$$

- **Strategy 2 (Fair-Mixing)** Agents specialize in the tax haven with the lower public signal, and mix uniformly when the signals are equal:

$$s_i(\mu_1, \mu_2) = \begin{cases} \text{Specialize in } TH_1 & \text{if } \mu_1 < \mu_2 \\ \text{Specialize in } TH_2 & \text{if } \mu_1 > \mu_2 \\ \text{Mix (0.5, 0.5)} & \text{if } \mu_1 = \mu_2. \end{cases} \tag{12}$$

The full coordination strategy ensures coordination on the more favorable tax haven but introduces a discontinuous jump when public signals are equal. Finally, the hybrid strategy follows Strategy 1 when signals differ, but switches to Strategy 2 when they are equal, avoiding arbitrary tie-breaking.

Proposition 2 *If $\sigma_\mu^2 > \sigma_x/\sqrt{2\pi}$ and Lemma 1 holds, then Strategies 1 and 2 constitute Perfect Bayesian Nash Equilibria.*

The proof is provided in Appendix A. The first condition, established in Proposition 1, ensures that agents' second-stage strategies are optimal given the equilibrium thresholds and that the solution is unique. Lemma 1 ensures that expected payoffs are increasing in S_j and decreasing in μ_j, implying strategic complementarities and monotonicity in best responses. As a result, coordination emerges endogenously, and each strategy leads to equilibrium specialization consistent with the public signals.

Proposition 3 *Strategy 1 yields the equilibrium with the highest expected payoff for all agents; that is, it constitutes the payoff-dominant equilibrium.*

This result follows directly from Lemma 1 and the properties of Strategy 1. Since the expected payoff increases with the number of agents specializing in a given tax haven and decreases with the public signal, the highest payoff of the game is achieved when all agents coordinate on the tax haven with the lower public signal. In the case of a tie, full coordination on either haven still dominates any split. This is exactly what Strategy 1 prescribes. Hence, Strategy 1 attains the highest possible expected payoff, and, being an equilibrium strategy, it constitutes the payoff-dominant equilibrium.

3 Evasion analysis

For given switching strategies, the total number of agents who choose to evade in equilibrium is

$$A^* = A_1^* + A_2^* = S_1^* \Pr\left(x_{i1} \leq \hat{x}_1^* | \theta_1\right) + S_2^* \Pr\left(x_{i2} \leq \hat{x}_2^* | \theta_2\right). \tag{13}$$

Note that the realization of the private signals and the survival of the tax havens depend on the realizations of θ_1 and θ_2. As a result, A^* changes in each realization of the game.

We can construct an approximation of A^* by making a "first-stage" inference about the distribution of x_{ij}, i.e., using the public signals:

$$A^* \approx \tilde{A} = S_1^* \Phi\left(\frac{\hat{x}_1^* - \mu_1}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) + S_2^* \Phi\left(\frac{\hat{x}_2^* - \mu_2}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{14}$$

The higher σ_x and the lower σ_μ, the more accurate the approximation becomes, as the public signals provide a better estimate of θ_j.

Comparative statics on $\tilde{A}$ are conducted with respect to the public signals. An increase in only one of the signals, say μ_1, affects evasion through multiple channels: (i) it shifts the mean of the distribution of x_{i1} to the right, making it less likely that agents fall below the evasion threshold—thereby reducing evasion; (ii) it raises the evasion threshold $\hat{x}_1^*$ (as shown in Lemma 1), further decreasing the likelihood of evasion in TH_1; and (iii) depending on the specialization switching strategy, the increase in μ_1 may reduce S_1^* and increase S_2^*, shifting more agents toward TH_2. These changes feed back into the thresholds: $\hat{x}_1^*$ decreases due to the lower S_1^*, while $\hat{x}_2^*$ increases due to the higher S_2^*. The net effect of this third channel is ambiguous. While evasion in TH_1 becomes less attractive, the reallocation of agents toward TH_2 makes evasion there more appealing, potentially offsetting the initial reduction and even increasing total evasion.

The aggregate effect can be seen formally in this

$$\frac{\partial \tilde{A}}{\partial \mu_1} = \underbrace{S_1^* \frac{\phi_1(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}} \left(\frac{\partial \hat{x}_1^*}{\partial S_1^*} \frac{\partial S_1^*}{\partial \mu_1} - 1 \right)}_{\leq 0} + \underbrace{S_2^* \frac{\phi_2(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}} \left(\frac{\partial \hat{x}_2^*}{\partial S_2^*} \frac{\partial S_2^*}{\partial S_1^*} \frac{\partial S_1^*}{\partial \mu_1} \right)}_{\geq 0} + \underbrace{\frac{\partial S_1^*}{\partial \mu_1} \left[\Phi_1(\cdot) - \Phi_2(\cdot) \right]}_{\substack{<0 \text{ if } \mu_1 < \mu_2 \\ >0 \text{ if } \mu_1 > \mu_2 \\ =0 \text{ if } \mu_1 = \mu_2}} \tag{15}$$

Here, $\Phi_j(\cdot)$ and $\phi_j(\cdot)$ denote the CDF and PDF corresponding to TH_j. The first term captures the direct effect in TH_1: a higher μ_1 shifts the signal distribution (i) and raises the evasion threshold (ii), both reducing evasion. The second term reflects the reallocation effect (iii)—as agents switch from TH_1 to TH_2, S_2^* increases, potentially lowering $\hat{x}_2^*$ and raising evasion. The third term captures a composition effect: reallocating agents alters total evasion depending on which haven has the higher evasion rate.

When the change in μ_1 does not affect the specialization decision—that is, $\frac{\partial S_1^*}{\partial \mu_1} = 0$—the expression simplifies to:

$$\frac{\partial \tilde{A}}{\partial \mu_1} = -S_1^* \frac{\phi_1(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}} \leq 0 \tag{16}$$

This confirms that, holding specialization fixed, increasing μ_1 always reduces evasion through TH_1.

Proposition 4 *Given Strategy 1 and 2, when the public signal of one tax haven increases, the effect on evasion depends on the relative signal levels:*

(i) If the increase occurs in the more favorable haven (i.e., the one with the lower public signal), evasion decreases under both strategies.

(ii) If the increase brings the public signals to equality ($\mu_1 = \mu_2$), Strategy 2—which prescribes mixing—results in lower evasion than full coordination (Strategy 1).

(iii) If the signal increases beyond the point of equality, evasion rises under Strategy 2 due to a discontinuous shift in specialization, while it remains unchanged under Strategy 1.

(iv) If the increase occurs in the less favorable haven, it has no effect on evasion.

The proof is provided in Appendix A. When an increase in a public signal does not induce a shift in the specialization distribution, it can only reduce evasion. The interesting dynamics arise under Strategy 2 when the signals become equal. Relative to full coordination (Strategy 1), splitting the population across tax havens reduces evasion. However, due to the discontinuous nature of Strategy 2, a marginal increase beyond the equality point causes all agents to switch to the previously less favorable haven. This concentration effect reverses the earlier decline and leads to a sharp rise in evasion relative to the mixed case. By contrast, if the increase occurs in the haven with the higher public signal—one that no agents were choosing—then both strategies still prescribe full specialization in the other haven, and the increase has no effect on evasion.

Proposition 5 *An equal increase in both public signals leads to a monotonic decrease in evasion under both strategies.*

This result follows from the fact that when both public signals increase simultaneously while maintaining a constant difference (e.g., $\mu_2 = \mu_1 + c$), the specialization choices of agents remain unchanged. Consequently, the third channel—feedback from switching in specialization—is neutralized. Evasion decreases purely due to the direct effects: the shift in the distribution of private signals and the increase in the evasion thresholds, both of which reduce the probability of evading taxes.

Figure 3 illustrates how changes in public signals affect equilibrium evasion outcomes under the two specialization strategies, showcasing the results from Propositions 4 and 5. Panel (a) considers the case where μ_2 is fixed and μ_1 increases (Proposition 4). Panel (b) shows the case where both μ_1 and μ_2 increase equally (Proposition 5).

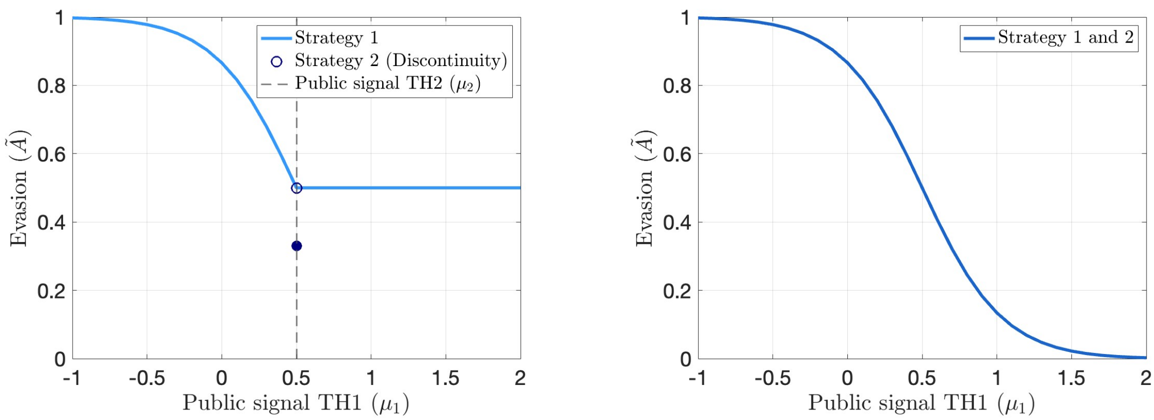

(a) Increasing μ_1 (b) Increasing μ_1 and μ_2

Figure 3: Comparative statics on evasion

Notes: The figure shows how equilibrium evasion levels vary under Strategies 1 and 2 for different values of the public signals. Panel (a) analyzes the effect of increasing μ_1 while holding μ_2 fixed at 0.5. Panel (b) considers the case where both public signals increase jointly such that $\mu_2 = \mu_1 + 1.5$. The results are computed under the following parameter values: $(p, t, \sigma_\mu, \sigma_x) = (0, 0.5, 1, 1)$.
Source: Author's elaboration.

4 Extension: investors with specialization bias

The main results rely on equilibrium switching strategies that generate sharp transitions in specialization, with discontinuous shifts around signal thresholds. In this section, I explore a more flexible framework that allows for smooth switching behavior by introducing individual biases toward one tax haven. While this extension does not yield a full equilibrium characterization, it provides useful insights into how gradual transitions affect evasion outcomes.

A simple example of a strategy that allows for smooth transitions is given by

$$s_i(\mu_1, \mu_2, \delta_i) = \begin{cases} \text{Specialize in } TH_1 & \text{if } \mu_1 - \delta_i < \mu_2, \\ \text{Mix } (0.5, 0.5) & \text{if } \mu_1 - \delta_i = \mu_2, \\ \text{Specialize in } TH_2 & \text{if } \mu_1 - \delta_i > \mu_2, \end{cases} \tag{17}$$

where $\delta_i \sim \mathcal{N}(0, \sigma_\delta^2)$ represents an idiosyncratic bias toward TH_1. By centering the distribution at zero, agents are on average indifferent between the two tax havens. The symmetry of the normal distribution ensures that half the population favors TH_1, while the other half prefers TH_2. The parameter σ_δ captures the extent of heterogeneity: higher values lead to greater dispersion in preferences.

Figure 4 presents the comparative statics of the base model and this extension, using alternative distributions of δ ($\sigma_\delta = 0.25$ and $\sigma_\delta = 0.5$). A key insight is that lower volatility in δ (i.e., smaller σ_δ) leads to behavior that more closely mirrors the base model, as specialization shifts more abruptly between tax havens.

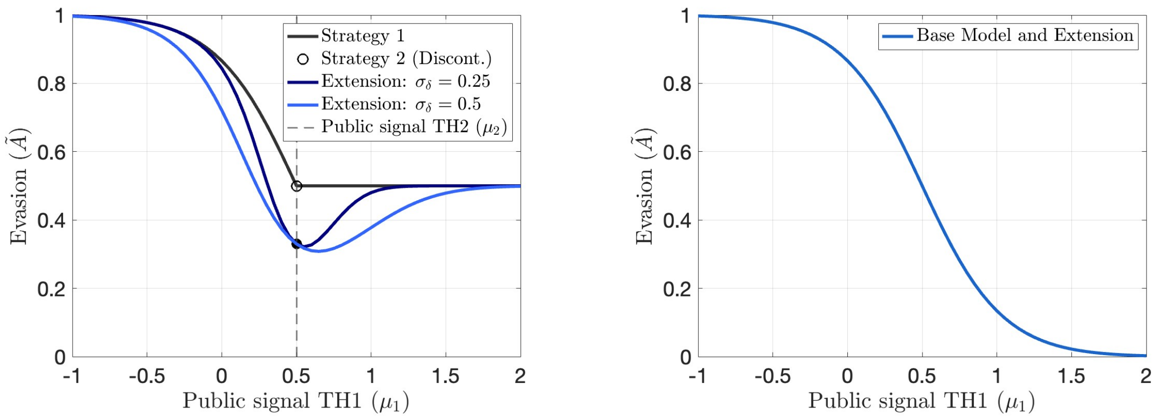

(a) Increasing μ_1 (b) Increasing μ_1 and μ_2

Figure 4: Comparative statics on evasion in extension

Notes: The figure illustrates how equilibrium evasion levels respond to changes in the public signals, comparing the base model from the previous section with an extension under alternative distributions ($\sigma_\delta = 0.25$ and $\sigma_\delta = 0.5$). Panel (a) examines the effect of increasing μ_1 while holding μ_2 fixed at 0.5. Panel (b) considers a joint increase in both signals such that $\mu_2 = \mu_1 + 1.5$. Results are computed using the following parameters: $(p, t, \sigma_\mu, \sigma_x) = (0, 0.5, 1, 1)$.
Source: Author's elaboration.

In the extended model, agents begin switching before the signals reach equality, which introduces a smoothing effect that reduces evasion relative to the base model. This reduction occurs gradually, reaching its lowest point precisely at equality—where, due to the symmetry of the distribution, half of the agents specialize in each tax haven, replicating Strategy 2 in the base model.

Interestingly, as the difference between μ_1 and μ_2 increases further, evasion continues to fall, reaching a minimum, and then begins to rise again, eventually converging to the base model values. This pattern suggests that, initially, the policy effect dominates—evasion lost in the taxed haven is not fully compensated by gains in the alternative haven. However, once the gap in public signals becomes large enough, the crowding-in effect takes over, and evasion begins to rise as more agents coordinate on the better tax haven.

5 Policy Implications

The model reveals a central insight: when coordination problems arise across multiple tax havens, agents respond not only to the absolute conditions of each jurisdiction but also to their relative attractiveness. If one jurisdiction stands out as more secure, agents concentrate there, making coordination easier and evasion more likely. When havens appear similar, agents are more uncertain about others' choices, coordination weakens, and aggregate evasion falls.

This has direct implications for policy design. A policy targeting a single tax haven reduces evasion only if it undermines the most attractive one; otherwise, it may simply redirect evaders to other jurisdictions, leaving overall evasion unchanged or even increased. By contrast, symmetric enforcementweakens coordination in all of them and reduces evasion.

This mechanism may help explain the limited effectiveness of past initiatives, such as those led by the OECD. Uneven enforcement created relocation rather than deterrence. In this light, the Global Minimum Tax appears more promising, as long as it is implemented uniformly and with broad international participation.

6 Conclusions

This paper studies tax evasion through the lens of a global game with multiple tax havens. By modeling evasion as a regime change problem occurring across several tax havens, it captures how investors coordinate both within and across jurisdictions under incomplete information. The key innovation is to allow for cross-haven strategic interactions, which are absent in single-entity models. I then analyze how different types of policies that raise the cost of being a tax haven affect evasion outcomes.

Coordination is easier when tax havens differ in perceived safety, as investors tend to concentrate on the most attractive one. This makes uneven enforcement potentially counterproductive, as it may increase overall evasion by shifting it across jurisdictions. In contrast, applying pressure evenly reduces coordination incentives in all havens and is more effective. This helps explain the limited success of the OECD's 2008–2009 initiatives and highlights the potential of the Global Minimum Tax, provided it is implemented uniformly and with broad coverage.

As a direction for future research, the model could be extended to include more than two tax havens with differing tax rates, or to a general equilibrium setting where both the OECD and tax havens are strategic players. Additionally, since individuals are assumed to be risk-neutral, the crowding-in effect arises from differences in expected value rather than risk reduction. Extending the model to risk-averse investors would offer new insights into how coordination effects operate under different preferences. Finally, the framework could be applied to other coordination settings where individuals choose among multiple entities—such as banks or currencies, as studied in other global game contexts. Ignoring cross-entity interactions in these environments may miss important strategic interactions.

References

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman. 2018. "Who owns the wealth in tax havens? Macro evidence and implications for global inequality." *Journal of Public Economics* 162:89–100.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2007. "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks." *Econometrica* 75 (3): 711–756.

Boston Consulting Group. 2014. *Global Wealth Report 2014: Riding a Wave of Growth.* Boston, MA: Boston Consulting Group.

Bucovetsky, Sam. 2014. "Honor among tax havens." *Journal of Public Economics* 110:74–81.

Carlsson, Hans, and Eric Van Damme. 1993. "Global games and equilibrium selection." *Econometrica: Journal of the Econometric Society,* 989–1018.

Elsayyad, May, and Kai A Konrad. 2012. "Fighting multiple tax havens." *Journal of international Economics* 86 (2): 295–305.

Goldstein, Itay, and Ady Pauzner. 2005. "Demand–deposit contracts and the probability of bank runs." *the Journal of Finance* 60 (3): 1293–1327.

Henry, James S. 2012. "The price of offshore revisited." *Tax Justice Network* 22:57–168.

Hines Jr, James R. 2005. "Do tax havens flourish?" *Tax policy and the economy* 19:65–99.

Johannesen, Niels, and Gabriel Zucman. 2014. "The end of bank secrecy? An evaluation of the G20 tax haven crackdown." *American Economic Journal: Economic Policy* 6 (1): 65–91.

Kanbur, SM Ravi, and Michael Keen. 1991. *Tax competition and tax coordination when countries differ in size.* Vol. 738. World Bank Publications.

Keen, Michael, and Kai A Konrad. 2013. "The theory of international tax competition and coordination." In *Handbook of public economics,* 5:257–328. Elsevier.

Konrad, Kai A, and Tim BM Stolper. 2016. "Coordination and the fight against tax havens." *Journal of International Economics* 103:96–107.

Morris, Stephen, and Hyun Song Shin. 1998. "Unique equilibrium in a model of self-fulfilling currency attacks." *American Economic Review,* 587–597.

———. 2001. "Global games: Theory and applications."

OECD. 2024. *Taxation and Inequality: OECD Report to G20 Finance Ministers and Central Bank Governors.* Paris. https://doi.org/10.1787/8dbf9a62-en. https://doi.org/10.1787/8dbf9a62-en.

Slemrod, Joel, and John D Wilson. 2009. "Tax competition with parasitic tax havens." *Journal of Public Economics* 93 (11-12): 1261–1270.

Wilson, John D. 1986. "A theory of interregional tax competition." *Journal of urban Economics* 19 (3): 296–315.

Zodrow, George R, and Peter Mieszkowski. 1986. "Pigou, Tiebout, property taxation, and the underprovision of local public goods." *Journal of urban economics* 19 (3): 356–370.

Zucman, Gabriel. 2013. "The missing wealth of nations: Are Europe and the US net debtors or net creditors?" *The Quarterly journal of economics* 128 (3): 1321–1364.

———. 2014. "Taxing across borders: Tracking personal wealth and corporate profits." *Journal of economic perspectives* 28 (4): 121–48.

Appendices

A Ommited proofs

A.1 Proof of Proposition 1

To proof existence and uniqueness I follow Angeletos et al. (2007) methodology.

Using the equilibrium Equation 4 and 7, we can create a function named $U(\hat{\theta}_k)$ s.t.:

$$U(\hat{\theta}_k) = \hat{\theta}_k - S_k \cdot \Phi\left(\frac{(\alpha-1)(\hat{\theta}_k - \mu_k) - \beta\Phi^{-1}\left(\frac{1-t}{1-p}\right)}{\sigma_x}\right). \tag{A1}$$

To proof the function is monotonic in $\hat{\theta}_k$ and hence, the FP exist and ! we need that

$$\frac{\partial U(\cdot)}{\partial \hat{\theta}_k} = 1 - S_k \frac{1}{\sigma_x}\phi(\cdot)(\alpha-1) > 0 \tag{A2}$$

To proof this derivative is positive, consider the maximum values of $\phi(\cdot) = 1/\sqrt{2\pi}$ and $S_k = 1$, then

$$\frac{\partial U(\cdot)}{\partial \hat{\theta}_1} > 0 \Rightarrow 1 - \frac{1}{\sigma_x}\frac{1}{\sqrt{2\pi}}\frac{\sigma_x^2}{\sigma_\mu^2} > 0 \Rightarrow \sigma_\mu^2 > \frac{\sigma_x}{\sqrt{2\pi}} \tag{A3}$$

Therefore the last condition is both necessary and sufficient for $U(\cdot)$ to be monotonic on $\hat{\theta}_k$, in which case the monotone equilibrium is unique. Finally, to prove that this equilibrium is the only one surviving iterated deletion of strictly dominated strategies, see Morris and Shin (2001). ■

A.2 Proof of Corollary 1

To proof the second point we need to use implicit differentiation in Equation A1 ,

$$\frac{\partial \hat{\theta}_k}{\partial \hat{x}_k} = -\frac{\partial U/\partial \hat{x}_k}{\partial U/\partial \hat{\theta}_k} = \frac{S_k \frac{\phi(\cdot)}{\sigma_x}}{1 - S_k\phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \tag{A4}$$

The latter inequality comes from the condition for existence and uniqueness. Using Equation 7,

$$\frac{\partial \hat{x}_k^*}{\partial \hat{\theta}_k} = \alpha > 0 \tag{A5}$$

Regarding the second point, using implicit differentiation in Equation A1 ,

$$\frac{\partial \hat{\theta}_k}{\partial S_k} = -\frac{\partial U/\partial S_k}{\partial U/\partial \hat{\theta}_k} = \frac{\Phi(\cdot)}{1 - S_k\phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \tag{A6}$$

The latter inequality comes from the condition for existence and uniqueness. Then, using Equation 7,

$$\frac{\partial \hat{x}_k^*}{\partial S_k} = \alpha \cdot \frac{\Phi(\cdot)}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \tag{A7}$$

Regarding the different tax rates,

$$\frac{\partial \hat{\theta}_k}{\partial t} = -\frac{\partial U/\partial t}{\partial U/\partial \hat{\theta}_k} = -\frac{S_k \frac{\phi(\cdot)}{\sigma_x}\beta\frac{\partial \Phi^{-1}(\cdot)}{\partial t}}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \quad \text{given that } \frac{\partial \Phi^{-1}(\cdot)}{\partial t} < 0, \tag{A8}$$

then

$$\frac{\partial \hat{x}_k}{\partial t} = \alpha \frac{\partial \hat{\theta}_k}{\partial t} - \beta \frac{\partial \Phi^{-1}(\cdot)}{\partial t} > 0. \tag{A9}$$

W.r.t. p

$$\frac{\partial \hat{\theta}_k}{\partial p} = -\frac{\partial U/\partial p}{\partial U/\partial \hat{\theta}_k} = -\frac{S_k \frac{\phi(\cdot)}{\sigma_x}\beta\frac{\partial \Phi^{-1}(\cdot)}{\partial p}}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} < 0 \quad \text{given that } \frac{\partial \Phi^{-1}(\cdot)}{\partial p} > 0, \tag{A10}$$

then

$$\frac{\partial \hat{x}_k}{\partial p} = \alpha \frac{\partial \hat{\theta}_k}{\partial p} - \beta \frac{\partial \Phi^{-1}(\cdot)}{\partial p} < 0. \tag{A11}$$

Finally, regarding μ_k,

$$\frac{\partial \hat{\theta}_k}{\partial \mu_k} = -\frac{\partial U/\partial \mu_k}{\partial U/\partial \hat{\theta}_k} = -\frac{S_k \phi(\cdot)\frac{\alpha-1}{\sigma_x}}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} < 0, \tag{A12}$$

then

$$\frac{\partial \hat{x}_k}{\partial \mu_k} = \alpha \frac{\partial \hat{\theta}_k}{\partial \mu_k} - (\alpha - 1) = -\alpha \frac{S_k \phi(\cdot)\frac{\alpha-1}{\sigma_x}}{1 - S_k \phi(\cdot)\frac{\alpha-1}{\sigma_x}} - (\alpha - 1) < 0. \tag{A13}$$

■

A.3 Proof of Lemma 1

We need to show that $\partial \Pi_j(S_j|\mathcal{I}^1)/\partial S_j > 0$, i.e.,

$$\frac{\partial}{\partial S_j}(1-p)\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1) + \frac{\partial}{\partial S_j}(1-t)\Pr(x_{ij} > \hat{x}_j^*(S_j)|\mathcal{I}^1) > 0 \tag{A14}$$

Rewriting the probability term,

$$\frac{\partial}{\partial S_j}(1-p)\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1) + \frac{\partial}{\partial S_j}(1-t)\left[1 - \Pr(x_{ij} \leq \hat{x}_j^*|\mathcal{I}^1)\right] > 0 \tag{A15}$$

which simplifies further to:

$$(1-p)\frac{\partial}{\partial S_j}\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1) - (1-t)\frac{\partial}{\partial S_j}\Pr(x_{ij} \leq \hat{x}_j^*(S_j)|\mathcal{I}^1) > 0 \tag{A16}$$

The joint distribution of x_{ij} and θ_j follows a bivariate normal distribution:

$$f_{x_{ij},\theta_j}(x_{ij},\theta_j) \sim \text{Bivariate Normal with } \boldsymbol{\mu} = \begin{pmatrix} \mu_j \\ \mu_j \end{pmatrix}, \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_x^2 + \sigma_\mu^2 & \sigma_\mu^2 \\ \sigma_\mu^2 & \sigma_\mu^2 \end{pmatrix}, \rho = \frac{\sigma_\mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}; \tag{A17}$$

The corresponding cumulative distribution function (CDF) w.r.t to the threshold $\hat{x}_j$ and $\hat{\theta}_j$ is

$$\begin{aligned} F_{x_{ij},\theta_j}(\hat{x}_j,\hat{\theta}_j) &= \int_{-\infty}^{\hat{x}_j}\int_{-\infty}^{\hat{\theta}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{-\frac{z}{2(1-\rho^2)}} d\theta_1 dx_{ij} \\ z &= \left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(x_{ij}-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2 \end{aligned} \tag{A18}$$

Applying Leibniz's rule, the derivative of $\Pr(x_{ij} \le \hat{x}_j^*(S_j)|\mathcal{I}^1)$ is:

$$\frac{1}{\sqrt{2\pi(\sigma_x^2+\sigma_\mu^2)}} e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} \frac{\partial \hat{x}_j^*}{\partial S_j} \tag{A19}$$

Similarly, the derivative of $\Pr(x_{ij} \le \hat{x}_j^*(S_j), \theta_j \le \hat{\theta}_j(S_j)|\mathcal{I}^1)$

$$\int_{-\infty}^{\hat{\theta}_j} f_{x_{ij},\theta_j}(\hat{x}_j,\theta_j) d\theta_j \frac{\partial \hat{x}_j^*}{\partial S_j} + \int_{-\infty}^{\hat{x}_j} f_{x_{ij},\theta_j}(x_{ij},\hat{\theta}_j) dx_{ij} \frac{\partial \hat{\theta}_j}{\partial S_j} \tag{A20}$$

Substituting the bivariate normal distribution, we obtain:

$$\begin{aligned} &\int_{-\infty}^{\hat{\theta}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j \frac{\partial \hat{x}_j^*}{\partial S_j} \\ &+ \int_{-\infty}^{\hat{x}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij} \frac{\partial \hat{\theta}_j}{\partial S_j} \end{aligned} \tag{A21}$$

Then, Equation A16 becomes

$$
\begin{aligned}
&(1-p)\Big[\int_{-\infty}^{\hat{\theta}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j \frac{\partial \hat{x}_j^*}{\partial S_j} \\
&+\int_{-\infty}^{\hat{x}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij} \frac{\partial \hat{\theta}_j}{\partial S_j}\Big] \\
&-(1-t)\frac{1}{\sqrt{2\pi(\sigma_x^2+\sigma_\mu^2)}} e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} \frac{\partial \hat{x}_j^*}{\partial S_j} > 0
\end{aligned}
\tag{A22}
$$

Considering $\hat{x}_j^* \propto \alpha\hat{\theta}_j$, and eliminating common terms ($\frac{1}{\sqrt{\sigma_x^2+\sigma^2\mu}}$), we obtain:

$$
\begin{aligned}
&(1-p)\frac{1}{2\pi\sigma_\mu\sqrt{1-\rho^2}}\Bigg[\alpha\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j \\
&+\int_{-\infty}^{\hat{x}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij}\Bigg] \\
&-(1-t)\frac{1}{\sqrt{2\pi}}\alpha e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} > 0
\end{aligned}
\tag{A23}
$$

Substituting ρ and multiplying both sides by $\sqrt{2\pi}$, we get:

$$
\begin{aligned}
&(1-p)\frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sqrt{2\pi}\sigma_\mu\sigma_x}\Bigg[\alpha\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j \\
&+\int_{-\infty}^{\hat{x}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij}\Bigg] \\
&-(1-t)\alpha e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} > 0
\end{aligned}
\tag{A24}
$$

Using $\sqrt{\alpha} = \frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sigma_\mu}$, we rewrite:

$$\begin{aligned}
&(1-p)\frac{1}{\sqrt{2\pi}\sqrt{\alpha}\sigma_x}\left[\alpha\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j\right.\\
&+\left.\int_{-\infty}^{\hat{x}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij}\right]\\
&-(1-t)e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} > 0
\end{aligned} \tag{A25}$$

Since all terms are positive, we simplify by disregarding the second integral, adopting a sufficient condition. The second integral "helps" in proving the statement. Extracting the leading exponent term further refines the condition to:

$$\begin{aligned}
&(1-p)\frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sqrt{2\pi}\sigma_\mu\sigma_x}e^{\left(-\frac{1}{2(1-\rho^2)}\right)}\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j\\
&-(1-t) > 0
\end{aligned} \tag{A26}$$

Since the coefficient of $(1-p)$ is positive, the condition holds when the difference between p and t is sufficiently large. For instance, setting $t=1$ ensures the condition is met for all $p>0$. ■

A.4 Proof of Proposition 2

To prove that these strategies constitute a Perfect Bayesian Nash Equilibrium (PBNE), we must verify two conditions: Stage 2 optimality and Stage 1 optimality. Stage 2 optimality has already been established in Proposition 1. We now focus on Stage 1.

Since all strategies considered are symmetric, we analyze the case where all agents follow the same strategy. Rather than expressing expected payoffs as functions of the thresholds—which themselves depend on specialization proportions and public signals—it is more convenient to define expected payoffs directly as a function of $(S_1, S_2, \mu_1, \mu_2) : \Pi_{ij}(S_j, \mu_j \mid \mathcal{I}^1) = \Pi_{ij}(\hat{x}_j(S_j, \mu_j), \hat{\theta}_j(S_j, \mu_j) \mid \mathcal{I}^1)$.

Assume all agents follow Strategy 1, which prescribes specializing in the tax haven with the lower public signal, with ties broken in favor of TH_1. Then:

- If $\mu_1 < \mu_2$, all agents specialize in TH_1, so $(S_1 = 1, S_2 = 0)$. By Lemma 1, $\Pi_{i1}(1, \mu_1 \mid \mathcal{I}^1) > \Pi_{i2}(0, \mu_2 \mid \mathcal{I}^1)$, $\forall i$. No agent has an incentive to deviate and $(S_1 = 1, S_2 = 0)$ is consistent.

- If $\mu_1 = \mu_2$, tie-breaking leads all agents to specialize in TH_1, again yielding $(S_1 = 1, S_2 = 0)$. By Lemma 1, $\Pi_{i1}(1, \mu_1 \mid \mathcal{I}^1) > \Pi_{i2}(0, \mu_1 \mid \mathcal{I}^1)$, $\forall i$. No agent has an incentive to deviate and $(S_1 = 1, S_2 = 0)$ is consistent.

- If $\mu_1 > \mu_2$, all agents specialize in TH_2, so $(S_1 = 0, S_2 = 1)$. By Lemma 1, $\Pi_{i1}(0, \mu_1 \mid \mathcal{I}^1) < \Pi_{i2}(1, \mu_2 \mid \mathcal{I}^1)$, $\forall i$. No agent has an incentive to deviate and $(S_1 = 0, S_2 = 1)$ is consistent.

Thus, Strategy 1 constitutes a Perfect Bayesian Nash Equilibrium.

Assume all agents follow Strategy 2, which prescribes specializing in the tax haven with the lower public signal, and mixing uniformly when the signals are equal. Let $\sigma_i(j)$ denote the probability of agent i specializing in tax haven j.

- If $\mu_1 < \mu_2$, all agents specialize in TH_1, so $(S_1 = 1, S_2 = 0)$. By Lemma 1, $\Pi_{i1}(1, \mu_1 \mid \mathcal{I}^1) > \Pi_{i2}(0, \mu_2 \mid \mathcal{I}^1)$, $\forall i$. No agent has an incentive to deviate and $(S_1 = 1, S_2 = 0)$ is consistent.
- If $\mu_1 > \mu_2$, all agents specialize in TH_2, so $(S_1 = 0, S_2 = 1)$. By Lemma 1, $\Pi_{i1}(0, \mu_1 \mid \mathcal{I}^1) < \Pi_{i2}(1, \mu_2 \mid \mathcal{I}^1)$, $\forall i$. No agent has an incentive to deviate and $(S_1 = 0, S_2 = 1)$ is consistent.
- If $\mu_1 = \mu_2$, all agents mix uniformly: $\sigma_i(1) = \sigma_i(2) = 0.5$, resulting in $(S_1 = 0.5, S_2 = 0.5)$. By symmetry,
$$\Pi_{i1}(0.5, \mu_1 \mid \mathcal{I}^1) = \Pi_{i2}(0.5, \mu_2 \mid \mathcal{I}^1), \quad \forall i.$$
Agents are indifferent and have no incentive to deviate and $(S_1 = 0.5, S_2 = 0.5)$ is consistent.

Thus, Strategy 2 also constitutes a Perfect Bayesian Nash Equilibrium.

■

A.5 Proof of Proposition 4

We analyze the effect of increasing μ_1 on equilibrium evasion $\tilde{A}$. By symmetry, the same logic applies to increases in μ_2 when μ_1 is held constant.

According to Equation 15, the effect of increasing μ_1 varies across three domains:

- **Case 1:** $\mu_1 < \mu_2$
In this region, all agents specialize in TH_1, so $S_1^* = 1$ and $S_2^* = 0$. Since specialization remains fixed as μ_1 increases reducing evasion. This proves (i).
- **Case 2:** $\mu_1 = \mu_2$
Under Strategy 1, agents coordinate fully on TH_1, resulting in:
$$\tilde{A}_{\text{Strat 1}} = \Phi\left(\frac{\hat{x}_1^*(1) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{A27}$$
Under Strategy 2, agents split evenly across the two havens:
$$\tilde{A}_{\text{Strat 2}} = 0.5 \cdot \Phi\left(\frac{\hat{x}_1^*(0.5) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) + 0.5 \cdot \Phi\left(\frac{\hat{x}_2^*(0.5) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{A28}$$

By symmetry, $\hat{x}_1^*(0.5) = \hat{x}_2^*(0.5)$, implying:

$$\tilde{A}_{\text{Strat 2}} = \Phi\left(\frac{\hat{x}_1^*(0.5) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{A29}$$

Since $\hat{x}1^(1) > \hat{x}_1^(0.5)$ (by Lemma 1), we conclude that $\tilde{A}$Strat 1 $> \tilde{A}_{\text{Strat 2}}$, proving part (ii). However, note that a marginal increase in μ_1 beyond the symmetry point immediately triggers a switch back to full specialization under Strategy 2. As a result, evasion jumps back to the level of Strategy 1, reversing the previous decline. This discontinuous increase confirms part (iii).

- **Case 3:** $\mu_1 > \mu_2$
  Once all agents are specialized in TH_2, further increases in μ_1 have no effect on specialization or evasion, as no one responds to TH_1's signal anymore. This proves part (iv).

■

A.6 Proof of Proposition 8

The first part of the proof is straightforwards, by substituting the value of the public signals of the extreme cases in Equation 8 and using $F_\delta(\mu_2 - \mu_1)$, we obtain the same expressions as in the Proof of Proposition 4.

■

A.7 Proof of Proposition 9

If we evaluate Equation 15 at $\mu_1 = \mu_2$, which implies $\hat{x}_1 = \hat{x}_2$, $\hat{\theta}_1 = \hat{\theta}_2$, $S_1 = S_2 = 0.5$, $\Phi_1 = \Phi_2$ and $\phi_1 = \phi_2$, the derivative becomes

$$0.5\frac{\phi(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\left(\frac{\partial \hat{x}_1^*}{\partial \mu_1} - 1\right) + 0.5\frac{\phi(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\left(\frac{\partial \hat{x}_2^*}{\partial \mu_1}\right) \tag{A30}$$

The fist term is negative whereas the second one it is positive. μ_1 decreases $\hat{x}_1$ through $S_1 = F_\delta$ and directly through Equation 7; whereas it increases $\hat{x}_2$ just through $S_2 = 1 - F_\delta$. Considering only the effect on $\hat{x}_1$ through S_1, the expression becomes the following one:

$$0.5\frac{\phi(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\left(\frac{\partial \hat{x}_1^*}{\partial S_1} - 1\right) + 0.5\frac{\phi(\cdot)}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\left(\frac{\partial \hat{x}_2^*}{\partial S_2}\right) \tag{A31}$$

The symmetry of the problem induces that

$$\frac{\partial \hat{x}_1^*}{\partial S_1} = \alpha \cdot \frac{\Phi(\cdot)}{1 - S_1\phi(\cdot)\frac{\alpha-1}{\sigma_x}} = \alpha \cdot \frac{\Phi(\cdot)}{1 - S_2\phi(\cdot)\frac{\alpha-1}{\sigma_x}} = \frac{\partial \hat{x}_2^*}{\partial S_2} \tag{A32}$$

and using the fact that $\partial S_2 = -\partial S_1$ the remaining term in (43) makes the equation negative, even when ignoring an effect that would reinforce this result. ■