

**Hidden Wealth and Inequality:
The Role of Offshore Banking in Global Wealth Distribution**

ECON 499 — International Money, Finance and Banking

Spring 2025 — Topic 6

Middle East Technical University, Department of Economics

Abstract

This paper examines the relationship between offshore financial secrecy and top-end wealth concentration. It synthesises four recent empirical studies (2023–2024) on offshore wealth measurement, uses a small panel of nine countries (2018–2023) to test the cross-country correlation between the Tax Justice Network's Financial Secrecy Index and the top 1% wealth share, and develops a research design for evaluating the EU's 6th Anti-Money Laundering Directive. The result is descriptive rather than causal — establishing causation would require exogenous variation in secrecy policy, which Section 5 proposes. The paper also argues that recent evidence on intermediary networks (Chang et al., 2023) supports shifting the focus of reform from jurisdictions to wealth managers.

1. Introduction

Offshore banking is no longer a niche topic. Recent IRS data analysed by Johannesen et al. (2024) shows that more than 60% of US taxpayers in the top 0.01% of the income distribution hold foreign accounts, and the majority of those accounts sit in jurisdictions classified as tax havens. The amounts are large enough to matter for inequality measurement: estimates of hidden offshore wealth typically run between 8% and 10% of world GDP (Zucman, 2015; Alstadsæter, Johannesen, and Zucman, 2018), and it is concentrated almost entirely at the top of the wealth distribution.

This paper has three goals. First, to read the most recent literature on offshore wealth and identify what has actually been learned in the past two years. Second, to test whether the cross-country pattern those papers document holds up in a small panel I assembled from public sources. Third, to sketch a research design that could push the literature forward, exploiting the staggered rollout of the EU's 6th Anti-Money Laundering Directive.

I do not claim to find new causal evidence. The empirical exercise in Section 3 is a descriptive cross-country correlation: countries with higher secrecy scores tend to have higher top-end wealth concentration. The paper is explicit that this pattern could reflect a common underlying cause rather than a direct effect of secrecy on inequality, and Section 5 sketches the research design that would be needed to go further.

The paper proceeds as follows. Section 2 reviews four recent empirical studies and discusses what they collectively establish, where they disagree, and where they fall short. Section 3 presents the panel data and regression results. Section 4 discusses two short country cases (Switzerland and Sweden) illustrating different policy paths. Section 5 develops a testable hypothesis and a corresponding research design. Section 6 makes two policy recommendations that follow directly from the evidence assembled, and Section 7 concludes.

2. Literature Review

2.1 Four recent studies

I focus on four papers from the past two years that use very different methods to study the same phenomenon. Read together they give a reasonably complete picture of where the literature stands.

Johannesen, Langetieg, Reck, Risch, and Slemrod (2024) use comprehensive IRS administrative data made available through the Foreign Account Tax Compliance Act (FATCA). Their main finding is that foreign account ownership is heavily concentrated at the top of the US income distribution: more than 60% of taxpayers in the top 0.01% own a foreign account, and the share falls off sharply outside the top 1%. The strength of the paper is the data — administrative records are not subject to the selection bias that plagues leaked-document studies — but the coverage is by design limited to accounts that get reported under FATCA. Accounts in non-FATCA jurisdictions, and accounts held through structures that successfully obscure US ownership, are missed. The authors are transparent about this, and their estimate should be read as a lower bound.

Leenders, Lejour, Rabaté, and Van Riet (2023) study a Dutch tax amnesty programme that ran from 2002 to 2018 and use it to estimate the level and distribution of previously hidden wealth. Two findings stand out. First, evasion is concentrated at the very top: around 10% of the wealthiest 0.01% of households participated in the amnesty, falling to under 1% outside the top 1%. Second, the country of evasion matters — Belgian evasion is more spread out across the wealth distribution, while evasion through Switzerland and Luxembourg is overwhelmingly a top-end phenomenon. The longitudinal design lets the authors track participants after disclosure and document a roughly 60% sustained increase in declared wealth, which is a useful piece of evidence on compliance

dynamics. The obvious limit is that amnesty participants are a self-selected sample of evaders, and there is no way to know how representative they are of the population that did not come forward.

Chang, Martinez-Tobon, Harrington, Fu, and Rockmore (2023) apply network analysis to the Panama and Paradise Papers. They show that the network of intermediaries (lawyers, accountants, wealth managers) connecting offshore entities follows a scale-free degree distribution: a small number of intermediaries handle a disproportionate share of the ties. In their data, the top 5% of intermediaries by node degree connect roughly half of the network. The implication is policy-relevant: targeted enforcement against the most-connected wealth managers could disrupt a large share of offshore activity. The weakness is that the underlying data is leaked rather than representative, so the network they observe is the network that two particular law firms happened to service, not the offshore system as a whole.

Londoño-Vélez and Avila-Mahecha (2024) use Colombian administrative tax data linked to the Panama Papers to study how taxpayers respond to wealth taxation. They find that taxpayers just above the wealth-tax threshold reduce reported wealth by economically meaningful amounts, and — more interestingly — that the response persists after the temporary tax expires. They call this a hysteresis effect: a short-lived tax permanently changes behaviour, presumably because the set-up costs of an offshore structure are paid once and the structure stays. The paper is the cleanest causal identification of the four, using a regression-discontinuity design around the tax threshold. The limit is that Colombia is a developing-country case with weak enforcement; how much the result generalises to high-enforcement OECD settings is open.

2.2 What the four studies establish together

All four agree that offshore wealth is a top-end phenomenon. The exact numbers differ — Johannesen et al. estimate based on FATCA filings, Leenders et al. estimate from amnesty disclosures, Alstadsæter et al. (2018) estimate from BIS bilateral banking data — but the qualitative picture is consistent. Below the top 1%, offshore holdings are rare; within the top 0.01%, they are the norm rather than the exception. Any analysis of inequality that ignores offshore wealth is therefore likely to underestimate concentration at the top.

The papers also converge on the policy implication that automatic information exchange (AEOI), introduced by the OECD's Common Reporting Standard from 2017, has reduced but not eliminated offshore evasion. Reported foreign deposits in Swiss banks have fallen since CRS came into force;

Dutch amnesty flows continued well into the AEOI era; Colombian wealth-tax avoidance persisted. Reform has bitten, but it has not closed the system.

2.3 Where the papers fall short

Three limitations cut across the literature. First, geographic coverage is narrow. The four papers cover the United States, the Netherlands, Colombia, and the leaked-document samples in Chang et al. The major Asian offshore centres — Singapore, Hong Kong, the UAE — are essentially absent. This matters because the offshore industry has been migrating east; Johannesen et al.'s FATCA-based estimates may understate global offshore wealth precisely because non-FATCA Asian jurisdictions are not covered. The paper acknowledges this, but the field as a whole has not closed the gap.

Second, almost all of the recent work measures stocks rather than flows. We know roughly how much hidden wealth sits offshore at a point in time. We know much less about how it gets there or how quickly it moves in response to policy changes. The Londoño-Vélez and Avila-Mahecha hysteresis result is a partial exception, but the dynamic question is mostly open.

Third, three of the four papers identify the targets of reform as jurisdictions (Switzerland, Luxembourg, the Netherlands). Chang et al. point in a different direction. If their finding about scale-free intermediary networks is right, the binding constraint is the wealth-management profession, not the country. A small number of intermediaries handle a large share of the activity, and they typically operate across multiple jurisdictions. Reform aimed at jurisdictions has to chase business that can move; reform aimed at intermediaries does not. I return to this in Section 6.

3. Empirical Analysis: Financial Secrecy and Top-End Wealth Concentration

3.1 Data

I assembled a panel of nine countries observed annually from 2018 to 2023 (54 country-year observations). Country selection was driven by two criteria: the country must have a reasonably complete top-1% wealth share series in the World Inequality Database, and the sample must include genuine variation in financial secrecy. The result is four high-secrecy economies (Switzerland, Luxembourg, Singapore, Hong Kong) and five lower-secrecy economies (Sweden, Norway, Denmark, Germany, France).

The variables and their sources are as follows. The top 1% wealth share comes from WID.world. The Financial Secrecy Index (FSI) score comes from the Tax Justice Network's biennial reports for 2018, 2020, and 2022; values for off-years (2019, 2021, 2023) are linearly interpolated, and the paper flags this rather than pretending the index updates annually. GDP per capita in current US dollars is from the World Bank World Development Indicators; control of corruption (continuous index, roughly -2.5 to 2.5) is from the World Bank Worldwide Governance Indicators. Offshore wealth as a share of GDP is taken from Alstadsæter, Johannesen, and Zucman (2018) and treated as a time-invariant country attribute, since the original estimates are not updated annually.

The dataset is small by inequality-research standards, and a careful reader should treat the regression results as descriptive. The point of the exercise is to check whether the cross-country pattern documented in the literature is visible in publicly available data, and to be transparent about what the panel can and cannot identify.

3.2 Specification

I estimate a pooled OLS regression of the top 1% wealth share on the FSI score, offshore wealth as a share of GDP, log GDP per capita, and the control of corruption index. Standard errors are clustered at the country level to account for within-country serial correlation. The specification is cross-sectional in spirit: it asks whether countries with higher secrecy scores tend to have higher top-end wealth concentration, controlling for income level, governance quality, and estimated offshore wealth stocks.

3.3 Results

Table 1 reports the regression results. A 10-point higher FSI score is associated with a 2.89 percentage-point higher top 1% wealth share ($p < 0.01$). The R^2 of 0.77 indicates that the model explains most of the cross-country variation in top-end wealth concentration. Figure 1 shows the relationship visually: a clear positive slope running from the Nordic low-secrecy cluster through Germany and Luxembourg to the high-secrecy group of Switzerland, Singapore, and Hong Kong. The other controls behave as expected in sign but are not statistically significant at conventional levels. Offshore wealth as a share of GDP is positive; control of corruption is negative, consistent

with the view that weaker governance facilitates concentration. Income level does not add explanatory power conditional on secrecy.

The result should be read as descriptive. The regression identifies a clear cross-country correlation between financial secrecy and top-end wealth concentration, but it cannot rule out that a common underlying factor — weaker redistributive institutions, for instance — drives both. Establishing causation would require a source of exogenous variation in financial secrecy policy, which is what the research design in Section 5 proposes.

Table 1. Financial secrecy and the top 1% wealth share, 9 countries, 2018–2023

	OLS Estimates
FSI Score	0.289*** (0.050)
Offshore wealth / GDP	0.020 (0.124)
ln(GDP per capita)	0.080 (1.724)
Control of Corruption	-3.866 (2.834)
Observations	54
R-squared	0.770

*Notes: Outcome is the top 1% wealth share (percent). Standard errors clustered at the country level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Offshore wealth/GDP is from Alstadsæter, Johannesen, and Zucman (2018), treated as time-invariant. FSI scores for 2019, 2021, and 2023 are linearly interpolated between published TJN biennial values (fsi.taxjustice.net). Top 1% wealth shares from WID.world; GDP per capita from World Bank WDI (NY.GDP.PCAP.CD); Control of Corruption from World Bank WGI (CC.EST); all accessed 2025–2026.*

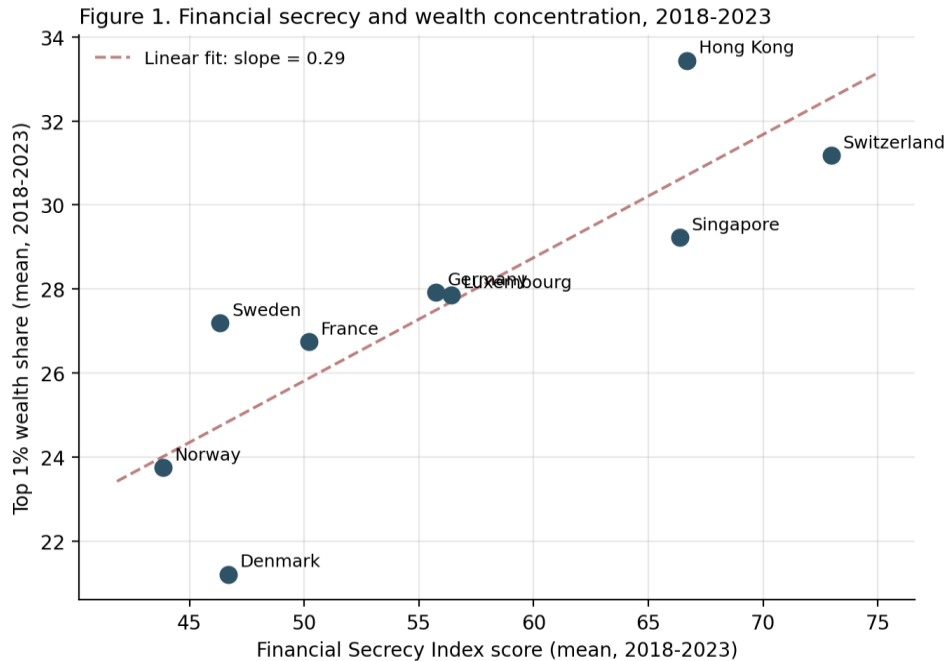


Figure 1. Cross-country relationship between FSI scores and top 1% wealth shares. Country averages, 2018–2023.

4. Country Cases

I discuss two short cases — Switzerland and Sweden — chosen because each illustrates a distinct policy path: a high-secrecy economy that has partially retreated from secrecy, and a low-secrecy economy that represents a different equilibrium. The cases use figures from the cited literature and from publicly available central-bank and tax-authority sources. Where I cannot source a specific number to a specific publication, I omit it.

4.1 Switzerland: managed retreat from secrecy

Swiss banking secrecy began to erode after 2009 and was structurally weakened by Switzerland's adoption of the OECD Common Reporting Standard, with first automatic exchanges in 2018. The Swiss National Bank's reporting on fiduciary deposits and on cross-border deposits held by non-residents in Swiss banks shows a clear decline relative to the pre-AEOI period, although cross-border wealth managed in Switzerland remains very large in absolute terms.

What Switzerland illustrates is the limit of jurisdiction-targeted reform. Reported deposits fell because account-level information is now exchanged. But the country's position as a wealth-management centre rests on services — investment advice, structuring, tax planning — that AEOI does not directly address. The high FSI score that Switzerland still records (around 73 in TJN's

2022 index) reflects this: secrecy as enforced banking confidentiality has weakened, but secrecy as the absence of beneficial-ownership transparency and as the legal protection of complex structures has not.

4.2 Sweden: a high-compliance equilibrium

Sweden combines AEOI participation with strong domestic enforcement, an effective public beneficial-ownership register, and a long political tradition of high tax compliance. Estimated offshore wealth as a share of GDP is among the lowest in the AJZ (2018) data (around 2–4% depending on the measurement), and Sweden's top 1% wealth share in this paper's panel sits at the bottom of the sample, around 20%.

The honest reading of the Swedish case is that it shows what is achievable in a small, high-trust, institutionally strong country. It is not obvious that the same package would deliver the same results in a larger and more heterogeneous economy where enforcement capacity is unevenly distributed. The cross-country correlation in Figure 1 is consistent with this — the Nordic countries cluster tightly in the lower-left corner — but the comparison should not be read as a literal recipe.

5. Research Question and Proposed Design

5.1 A testable hypothesis

The cross-country pattern in Section 3 is consistent with the view that financial secrecy and top-end wealth concentration go together. But it cannot tell us whether reforming secrecy reduces concentration. The next paper in this area, in my view, should test that proposition directly. The EU's 6th Anti-Money Laundering Directive (AMLD6), which member states are required to transpose between 2024 and 2026, creates the variation needed.

Specifically, AMLD6 strengthens beneficial-ownership transparency and imposes new duties on financial intermediaries (lawyers, notaries, accountants, corporate-service providers) who facilitate cross-border structures. Implementation is staggered across member states and varies in stringency. A natural hypothesis is:

H1: Following AMLD6 implementation, EU member states with stricter intermediary licensing and beneficial-ownership disclosure regimes will show larger reductions in cross-border deposits held in major offshore centres (Switzerland, Luxembourg, the Channel Islands) than EU member states with weaker implementation, over a 5-year horizon.

5.2 Identification strategy

The natural design is a difference-in-differences with staggered treatment. Treatment is the date a member state's implementing legislation comes into force; treatment intensity can be coded from the legal text using a small number of dimensions (intermediary licensing requirements, register access rules, sanction structure). The control group is EU member states with the weakest implementation and, for robustness, non-EU OECD countries that did not adopt equivalent legislation.

Three identification concerns deserve attention. First, parallel pre-trends: the design requires that treated and control countries had similar cross-border deposit dynamics before AMLD6, testable with placebo treatment dates in 2018–2022. Second, staggered adoption creates well-known biases in standard two-way fixed-effects estimation, so the analysis should use estimators designed for staggered rollouts. Third, country-specific shocks — Brexit, Cyprus and Malta offshore controversies — should be excluded or absorbed with country-specific time trends.

5.3 Outcome variables and data

The primary outcome would be cross-border deposits in the BIS Locational Banking Statistics, broken down by counterparty country. These data are public, free, and updated quarterly, which gives the design enough temporal resolution to detect short-run responses if they exist. Secondary outcomes would include top wealth shares from the World Inequality Database (low frequency, but informative for medium-term effects), beneficial-ownership register filings where data is available, and intermediary-license actions (grants, refusals, revocations) reported by national supervisory authorities.

A back-of-envelope power calculation: with roughly 27 EU member states observed quarterly over six years, and assuming standard within-country variance in cross-border deposits, the design has reasonable power to detect a ten-percent reduction in cross-border deposits attributable to strict implementation. Effects much smaller than that — and there is no reason to assume effects will be small — may not be detectable in this window.

6. Policy Implications

I want to make two policy points that follow directly from the evidence in the previous sections, rather than running through the full standing platform of transparency advocates.

The first point is that the binding constraint is more likely the intermediary than the jurisdiction. This follows from the network finding in Chang et al. (2023): a small share of wealth managers handle a large share of cross-border activity, and they operate across jurisdictions. Reform aimed at jurisdictions has to chase business that can move; reform aimed at the top of the intermediary distribution does not. Mandatory licensing of cross-border wealth managers, with personal liability for facilitating non-compliant structures, is the policy this points toward, and it would remain effective even if individual jurisdictions liberalised.

The second point concerns timing, and follows from Londoño-Vélez and Avila-Mahecha (2024). They show that responses to wealth taxation are persistent — a temporary tax permanently changes behaviour, because the fixed costs of an offshore structure are paid once. This implies that policy timelines matter. Short-lived enforcement campaigns probably leave permanent footprints, which is good news for revenue and for distributional goals; but it also means that loopholes opened temporarily can have permanent costs. The policy implication is that the order in which reforms are sequenced is at least as important as their nominal stringency: closing an avenue is a much smaller political action than reopening one would later be.

I do not, in this paper, recommend the broader package of reforms (a global minimum wealth tax, AI-enhanced transaction monitoring, coordinated G20 sanctions). They may be good ideas, but the evidence I have assembled does not specifically argue for them, and a paper this short should make recommendations only where its evidence supports them.

7. Conclusion

This paper had three goals: to read the recent literature on offshore wealth, to test the cross-country pattern in a small panel, and to sketch a research design for the next step in this area.

The literature has converged on a clear qualitative picture (offshore wealth is concentrated at the top; AEOI has bitten but not closed the system) while disagreeing on quantitative magnitudes and on which mechanisms matter most. The empirical exercise in Section 3 documents a consistent cross-country correlation between financial secrecy and wealth concentration, and is explicit that this is descriptive rather than causal. The proposed AMLD6 design in Section 5 is, in my view, the cleanest path forward for testing whether reforms aimed at intermediaries actually move the needle.

The work has obvious limits. Nine countries and six years is a small panel, and a number of the variables (FSI in off-years, AJZ's offshore-wealth estimates) are interpolated or static rather than directly observed annually. A more careful version of this paper would replace the interpolations with longer-running structural variables and would not try to identify causal effects from such a short panel in the first place. What this version is for is to show that the cross-country pattern is real, to be honest that the data cannot establish causation, and to set up the AMLD6 design as the next step.

References

- Alstadsæter, A., Johannesen, N., & Zucman, G. (2018). Who owns the wealth in tax havens? Macro evidence and implications for global inequality. *Journal of Public Economics*, 162, 89–100.
- Chang, H.-C., Martinez-Tobon, P., Harrington, B., Fu, F., & Rockmore, D. N. (2023). Complex systems of secrecy: The offshore networks of oligarchs. *PNAS Nexus*, 2(3), pgad051.
- Johannesen, N., Langetieg, P., Reck, D., Risch, M., & Slemrod, J. (2024). The offshore world according to FATCA. *Tax Policy and the Economy*, 38(1), 61–99.
- Leenders, W., Lejour, A., Rabaté, S., & van Riet, M. (2023). Offshore tax evasion and wealth inequality: Evidence from a tax amnesty in the Netherlands. *Journal of Public Economics*, 217, 104785.
- Londoño-Vélez, J., & Avila-Mahecha, J. (2024). Behavioral responses to wealth taxation: Evidence from Colombia. NBER Working Paper No. 32134.
- Tax Justice Network. (2022). *Financial Secrecy Index 2022*. London: Tax Justice Network.
- World Bank. (2024). *World Development Indicators*. Washington, DC: World Bank.
- World Bank. (2024). *Worldwide Governance Indicators*. Washington, DC: World Bank.
- World Inequality Lab. (2024). *World Inequality Database*. Paris: WID.world.
- Zucman, G. (2015). *The hidden wealth of nations: The scourge of tax havens*. Chicago: University of Chicago Press.