

**The Offshore World According to FATCA:
New Evidence on the Foreign Wealth of U.S. Households***

9 March 2023

**Niels Johannesen, University of Copenhagen
Daniel Reck, University of Maryland
Max Risch, Carnegie Mellon University
Joel Slemrod, University of Michigan
John Guyton, Internal Revenue Service
Patrick Langetieg, Internal Revenue Service**

Abstract: This paper uses account-level information, reported to the IRS by foreign financial institutions under the Foreign Account Tax Compliance Act (FATCA), to produce new evidence on the foreign financial wealth of U.S. households. We find that U.S. taxpayers hold around \$4 trillion in foreign accounts, almost half in jurisdictions usually considered tax havens. Combining the FATCA reports with other administrative tax data and tracing account ownership through partnerships, we document a steep income gradient in the propensity to hold assets in foreign financial institutions. Specifically, more than 60% of the individuals in the top 0.01% of the income distribution own foreign accounts, the vast majority in tax havens and more than half through a partnership. We discuss the likely implications of these findings for the overall impact of FATCA on tax compliance and government revenue.

* We thank Patricia Banta, Timothy Burke, Kathleen Doty, Aidia Greenlee, Michael Hayden, Anne Herlache, Barry Johnson, Camille Landais, Wouter Leenders, Barry Levine, Juliana Londoño-Vélez, Jim Omartian, Neviana Petkova, Ted Setzer, Steve Shay, Matt Smith, Bryan Stiernagle, Bill Strang, Johannes Spinnewijn, Owen Zidar, and numerous seminar and conference participants for helpful discussion, support, and comments on early versions of this work. Reck gratefully acknowledges financial support from the Economic and Social Research Council and the Laura and John Arnold Foundation. Jeanne Bomare, Baptiste Roux, and Francesco Armellei provided excellent research assistance. All findings, opinions, and errors are those of the authors alone and do not necessarily represent the opinions of the Internal Revenue Service or Treasury Department.

1. Motivation

Information flows from employers and financial institutions to tax authorities represent a key technology to limit tax evasion: when an employer has reported a salary payment or a bank has reported an interest payment, the scope for the taxpayer to evade taxes on those income items is generally limited (Kleven et al., 2011). Developed economies typically rely heavily on such third-party reporting and regularly seek to extend it to additional tax bases to improve tax compliance (Slemrod et al., 2017).

Considering the importance of this enforcement technology, income from foreign sources represents a major challenge. Foreign firms are typically not covered by the reporting obligations facing domestic firms. The resulting challenge is particularly salient in the context of capital taxation. Many taxpayers have long been able to evade capital income taxes, wealth taxes and inheritance taxes by holding wealth through banks in countries with a strong commitment not to share information, in the form of bank secrecy laws and other measures. Basic facts about the magnitude and composition of offshore financial wealth remain elusive, because of its characteristic opacity, which in turn creates difficulties for efforts to estimate the distribution of wealth across countries (Lane and Milesi-Ferretti, 2001, 2007) and within countries (Saez and Zucman, 2016). Despite all this, in the last 10 years, massive reforms to tax law and waves of international agreements have attempted to address the problems created by opacity of offshore financial wealth for tax enforcement by expanding cross-border information sharing.

In this paper, we analyze the information on offshore wealth disclosed by foreign banks under the Foreign Account Tax Compliance Act (FATCA), a major U.S. policy that for the first time extends comprehensive third-party reporting to foreign financial institutions. Since 2015, FATCA has essentially required all foreign banks, investment funds and other financial intermediaries, including those based in essentially every current or former tax haven, to provide information to the IRS about specified financial accounts controlled by U.S. taxpayers on the Form 8966. Crucially, the reporting obligations include accounts where the beneficial owner is a U.S. taxpayer even if, e.g., the immediate owner of the account is a closely held foreign company. About 45,000 foreign financial institutions in 190 countries provided financial information for 4.5 million U.S.-owned accounts in 2018.

By combining the account-level reports from foreign financial institutions with other administrative data, we construct new measures of the aggregate foreign financial wealth of U.S. households as well as the distribution of this wealth over income groups. Earlier papers used macro statistics to estimate aggregate foreign financial assets (Zucman, 2013; Alstadsæter et al., 2018) as well as different sources of information about foreign wealth at the individual-level to estimate the distribution across wealth groups and income groups (Alstadsæter et al., 2019; Londoño-Velez and Ávila-Mahecha, 2020; Leenders et al., 2020, Guyton et al, 2021). Our paper differs from these papers in two main respects. First, we use administrative data that in principle cover the full population of US owners of offshore wealth, and that looks through the complicated ownership structures often used by owners of offshore accounts. This means that

we can dispense with many of the assumptions necessary for earlier researchers to make inference from incomplete macro data or relatively small and potentially selected samples of micro data. Second, our data cover foreign financial assets after the massive shock to financial transparency associated with FATCA and related reforms. This means that the descriptive statistics we estimate are potentially shaped by responses to this shock, such as repatriation of wealth or investment in non-financial assets outside the scope of reporting.

According to our estimates, around 1.5 million U.S. taxpayers held foreign financial accounts with aggregate assets of around \$4 trillion in tax year 2018. By comparison, the total financial assets of U.S. households totaled roughly \$80 trillion according to official financial accounts (Federal Reserve, 2022). Around half of the assets in foreign accounts, just below \$2 trillion, were held in jurisdictions usually considered tax havens, such as Switzerland, Luxembourg and the Cayman Islands.¹ Just 14% of accounts are located in tax havens compared to nearly half the total wealth, which reflects that accounts in havens were on average larger.

While we have conducted a thorough cleaning of the FATCA records before computing these aggregates – e.g. removing duplicate records, accounting for jointly owned accounts, removing records with clear signs of being invalid – some uncertainty remains about the exact magnitudes. On the one hand, our estimates may understate the true foreign financial assets of U.S. households, perhaps because some foreign financial institutions are less than fully compliant, or because some financial assets are not required to be reported in our dataset. Some financial institutions may fail to identify all accounts with U.S. beneficial owners, whether strategically or not. On the other hand, our estimates may overstate the true foreign assets of U.S. households to the extent that foreign financial institutions report some assets to the IRS that in effect belong to foreign households. For most foreign assets in our sample, we can verify that foreign account owners are indeed U.S. taxpayers based on the Taxpayer Identification Numbers (TIN) in the FATCA records. However, in a substantial number of cases where the FATCA records do not include TINs, we cannot verify U.S. ownership in this way (see also TIGTA, 2018), and some U.S. taxpayers may be part of a foreign household (e.g. U.S. citizens residing abroad).

Our preferred aggregate estimates imply a ratio of tax haven assets to GDP of around 10% in 2018. Interestingly, this is somewhat higher than a comparable estimate for 2007 of around 7% based on macro statistics (Alstadsæter et al., 2018). There are two possible interpretations of this discrepancy. Financial assets in tax havens may have grown significantly faster than the overall U.S. economy since 2007 despite significant policy efforts to curb offshore tax evasion (e.g.

¹ We use the term havens as a shorthand descriptor of countries with offshore financial centers with low effective tax rates, and, at least before FATCA, sufficient commitment to financial secrecy to attract foreigners desiring to shelter income from home-country taxation. There is no single, universally accepted list of havens, and being so designated is often disputed by named countries. In our descriptive analysis, we use the same list of tax havens as Johannesen et al. (2020). This list does not have any official role in IRS enforcement efforts; the IRS does not have an officially accepted definition of a tax haven. We further note that insofar as FATCA/CRS reduced offshore non-compliance in these countries, some of them might no longer meet our definition of a tax haven. Our use of the term should not be misconstrued as taking a stance on whether individuals can still conceal income from tax authorities in these countries by exploiting banking secrecy protections. See also footnote 15.

Johannesen and Zucman, 2014; Johannesen et al., 2020) and numerous data leaks from tax havens penetrating offshore secrecy (Johannesen and Stolper, 2021). Alternatively, the discrepancy may result from biases in the different approaches to measurement, most likely that the methods adopted by previous studies in the absence of administrative data tended to underestimate offshore wealth.

We investigate how foreign assets are distributed across the individual income distribution by linking foreign accounts to individual U.S. owners. This is relatively straightforward in cases where the account owner is an individual and the FATCA report includes the individual's TIN, but is more challenging in other cases. First, the FATCA reports identify partnerships, the preferred organizational form of investment funds, and other pass-through entities as the owners of more than 30% of the foreign assets, including many of the largest accounts in the data. Pass-through entities are sometimes owned by other entities, which may themselves be owned by yet other entities. Following Cooper et al., (2016), we use the tax forms issued by partnerships and other pass-through entities to their shareholders to unwind ownership structures, which allows us to match the vast majority of foreign assets to the tax returns of the owners of partnerships that reportedly own the foreign assets. Second, for almost 40% of the foreign assets, the FATCA report does not contain information about the TIN or includes a TIN that does not yield an unambiguous match to an individual or an entity, in which case we cannot confidently link the assets to an individual.

We document a steep income gradient in the propensity to hold assets in foreign accounts even within the top of the income distribution. More than 60% of the individuals in the top 0.01% of the income distribution own foreign accounts, either directly or indirectly through a pass-through entity. By comparison, this fraction is less than 40% for the bottom half of the top 0.1%; less than 20% for the bottom half of the top 1%; and less than 5% for the bottom half of the top 10%. Interestingly, the income gradient is primarily driven by indirect ownership through partnerships and by accounts in tax havens: The fraction of individuals who hold foreign assets through directly-owned accounts in non-havens, and only hold foreign assets in this way, is roughly the same for all the income groups within the top 5%. These figures are all lower bounds: they would be unambiguously higher if foreign financial institutions reported all U.S.-owned accounts and all reports could be matched to individual beneficial owners.

Our finding of a steep income gradient of the propensity to own offshore wealth is qualitatively in line with a growing body of literature estimating how the propensity to own offshore wealth varies with income and wealth. However, because our estimates are based on comprehensive administrative data rather than much smaller subsamples from data leaks or voluntary disclosures, we can meaningfully interpret the estimates in absolute and not just in relative terms. For instance, Alstadsæter et al. (2019) show that the propensity of Nordic wealth owners to hold an account in a specific Swiss bank increases monotonically through the wealth distribution up to around 1% at the top. While this is suggestive of differences in the propensity to hold assets offshore more broadly, it is not directly informative about the overall fraction holding assets offshore in each group.

We also document how the value of the assets in foreign accounts is distributed across income groups. We find a strong concentration of offshore assets at the very top of the income distribution: Around 30% of all foreign assets belong to the top 0.01%, with a particularly high share for assets held through partnerships and assets held in tax havens. We note that our distributional metrics may be somewhat sensitive to potentially imperfect reporting by foreign financial institutions and unmatched FATCA reports: they overstate concentration if foreign accounts belonging to the top wealth group are more likely to be reported and matched than other foreign accounts, and understate concentration if selection goes the other way.

We attempt to use these data to shed light on the income accruing to offshore accounts. Our preferred estimates indicate that the average annual rate of return on offshore wealth lies between 3% and 6%, with a significantly higher return for wealth in tax havens. Our estimates are roughly in line with calibrations based on macro aggregates in prior work (Alstadsæter et al., 2019, Guyton et al. 2021). However, examining aggregate income statistics also highlights some limitations of the FATCA reporting where income reporting is concerned. Financial institutions provide no income information on the FATCA report for about 45% of accounts and 41% of total wealth. We conducted a detailed analysis of records with missing income, and conclude that the wealth for which income is not reported was likely generating significant income. Our attempts to estimate a rate of return focus on the subset of accounts where we do observe income information, and we must additionally confront uncertainty in the map from gross income reported by financial institutions on FATCA reports to owners' net taxable income.

Our paper contributes to the literature that uses various types of macroeconomic data to quantify the external wealth of nations (Lane and Milesi-Ferretti, 2007), with recent papers paying particular attention to the measurement problems created by the wealth held in tax havens (Zucman, 2013). We advance this literature by providing a new estimate of aggregate U.S. financial wealth abroad based on the account-level reports from foreign financial institutions. Our estimate is significantly higher than the most recent estimates based on macro-data (Alstadsæter et al., 2018).

We also contribute to the literature concerned with the distribution of offshore wealth in tax havens across income and wealth groups, and the implications for inequality. We document that foreign assets, notably those held in tax havens through complex legal structures, are overwhelmingly concentrated among taxpayers at the top of the income distribution. This is consistent with several recent studies (Alstadsæter et al., 2019; Londoño-Velez and Ávila-Mahecha, 2020, Guyton et al., 2021), but not all (Leenders et al., 2020). To the extent that the capital income accruing to foreign accounts is underreported, inequality in the U.S. may be even more pronounced than suggested by studies based on tax records (Saez and Zucman, 2016; Piketty et al., 2018, Guyton et al., 2021).

Finally, we also contribute to the literature on automatic exchange of information as enacted under FATCA as well as under the Common Reporting Standard (CRS), a policy similar to FATCA adopted by over 100 countries. These policies are controversial because of their allegedly high compliance costs for foreign financial institutions (Byrnes and Munro, 2017), considerable

uncertainty about the revenue gains (TIGTA, 2022), and the risk of unintended consequences (Dharmapala, 2016). Several existing studies find evidence of behavioral responses to recent reforms consistent with increased tax compliance: declines in offshore deposits (Menkhoff and Miethe, 2019; O'Reilly et al., 2019; Cusi et al., 2020); declines in “round-tripping” investment through tax havens (De Simone et al., 2020); and declines in incorporations of offshore shell companies (Omartian, 2017). Other studies find responses on other behavioral margins suggestive of continued evasion: increases in non-financial assets not covered by current automatic exchange of information agreements, such as real estate and art (De Simone et al., 2020; Bomarre and Le Guern Herry, 2022).

Rather than estimating taxpayer responses to FATCA, our paper measures the scope of the new third-party reporting designed to curtail tax evasion, overall and at different positions in the income distribution. We discuss how these data, in combination with assumptions about the increase in compliance rates might be used to gauge the tax compliance effects of FATCA. The large amount of offshore wealth we observe in the FATCA records suggests that the compliance response of those who maintained their offshore accounts after FATCA implementation will be an important determinant of the overall compliance effect of FATCA. However, the magnitude of this effect is highly uncertain without a better understanding of these owners’ voluntary compliance response to FATCA reporting, which we plan to address in future work.

2. Background on FATCA

In the decade prior to the introduction of FATCA, governments in most developed countries, including the United States, took some measures to reign in offshore tax evasion. The G20 spearheaded a coordinated action to establish information exchange with tax havens in tax evasion cases (Johannesen and Zucman, 2014); the European Union compelled a group of tax havens to collect withholding taxes on interest income accruing to undisclosed accounts and share the tax revenue with the home countries of the account owners (Johannesen, 2014; Omartian, 2017); and the United States took legal action against specific foreign banks suspected of facilitating tax evasion (Johannesen et al., 2020). While these policy efforts apparently produced some increases in tax compliance – for instance, one study estimates that the bundle of policies adopted by the U.S. in 2008-2009 increased revenue by around \$1 billion per year (Johannesen et al., 2020) – they were widely perceived to be insufficient, and calls for more ambitious policies gained traction among academics as well as policymakers.

In March 2010, the Hiring Incentives to Restore Employment Act of 2010 was passed into law. This act added Chapter 4, known as the Foreign Account Tax Compliance Act (FATCA), to the Internal Revenue Code. The FATCA extends the scope of the longstanding domestic information reporting regimes to include foreign financial institutions (FFIs), requiring FFIs to report information about their U.S. account holders to the IRS or be subject to substantial withholding. Prior to FATCA, due to a lack of automatic information reporting on foreign accounts, it would have been extremely difficult for the IRS to know if taxpayers were misreporting foreign income absent information provided by a whistleblower (Guyton et al., 2021).

Under FATCA, FFIs with U.S. account holders, as well as certain types of non-financial foreign entities (NFFEs) with U.S. owners, are required to report automatically to the IRS on U.S. taxpayers' offshore assets. If they do not comply, the U.S. imposes a large, complex and potentially quite costly penalty on financial institutions: any FFI that declines to report automatically on all U.S. accounts is subject to 30% withholding on their U.S. source income.² This leaves an FFI four effective choices, i) stop investing in U.S. source income, ii) stop serving U.S. clients, iii) comply with the FATCA reporting requirements, or iv) accept the 30% withholding.

Intergovernmental Agreements (IGAs) between the U.S. and foreign governments facilitate the administration of FATCA. Because of the variation in foreign national bank secrecy laws, the IRS and Treasury Department developed two models of IGA. Under a Model 1 IGA, a foreign government agrees to adopt rules that require FFIs located in that jurisdiction to report the information required under FATCA to the foreign government, which then reports the information to the IRS. Under a Model 2 IGA, the foreign government agrees to direct FFIs to comply with an "FFI Agreement," and to report information on U.S. accounts directly to the IRS. The majority of countries signed a Model 1 IGA - out of the 114 countries currently with an IGA, 100 operate under Model 1.³ The U.S. has signed IGAs with essentially all major tax havens.⁴ Among 22 tax havens with IGAs, only Switzerland, Hong Kong, and Bermuda operate under Model 2 IGAs; the rest operate under Model 1. IGAs signed in the early years of FATCA typically designate the years 2014 and 2015 as a transition period during which some information was not required to be reported. Relatedly, the IRS notified the public that it would regard these two years as a transition period, indicating that it would not pursue withholding on those attempting in "good faith" to comply (IRS, 2014).

The primary obligations of an FFI under FATCA are i) identification and documentation of their account holders, ii) information reporting, and iii) withholding on payments to non-participating FFIs to enforce the aforementioned withholding rules. The FFI must determine if the account holder or payee is a "specified U.S. person", a foreign person, a nonparticipating FFI, or a "recalcitrant account holder." Under FATCA, a "specified U.S. person" is any U.S. individual taxpayer or any U.S. entity that is not a publicly traded corporation, bank, broker dealer, common trust, real estate investment trust, or tax-exempt organization. FFIs are not required to report on accounts owned by individuals with less than \$50,000 at their institution as of the end of the calendar year, but they can elect to forego this exception (and many do). There are additional exceptions to reporting requirements for some types of accounts generally associated with

² This tax applies to "withholdable payments" made to an FFI, defined broadly to include nearly all forms of passive income from U.S. sources. The statute initially also included gross proceeds from the sale or other disposition of property (of a type that can produce U.S. source interest or dividends) in the definition of withholdable payments, which would make withholding far more punitive in a business where net returns are much less than gross returns. Gross proceeds were excluded from withholdable payments by subsequent revisions to the relevant regulations.

³ For the current status of IGA agreements, see <https://home.treasury.gov/policy-issues/tax-policy/foreign-account-tax-compliance-act>.

⁴ Of the 27 countries on our list of tax havens, 22 have an IGA. The exceptions are Aruba, Belize, Cook Islands, Monaco, and Sint Martaan. All of these countries are participating in the CRS.

individual retirement planning (see Regulations Sections 1.1471-5(a)(4) and 1.1471-5(b)(2)). A recalcitrant account holder is an account holder that does not provide sufficient documentation to verify their status.

The FFI must file a Form 8966 with respect to each reportable account, identifying information about the account holder as well as the account balance and four categories of income flowing into the account during the year: interest, dividends, gross proceeds from the sale or redemption of property, and other income. For jointly owned accounts or NFFE account holders, the FFI must report the same information for each account owner, which in the NFFE case means a separate report for each specified U.S. person that owns more than 10% of the stock or interest in the foreign corporation, partnership, or trust that in turn owns the offshore account.⁵

Our main dataset is the full set of Forms 8966 that reported on financial assets held by U.S. persons. We disregard observations pertaining to non-participating FFIs and recalcitrant accounts, as our focus here is on the third-party reporting identifying Americans with offshore assets. The first reports, for the 2014 calendar year, were due March 31, 2015. Relatively few accounts were reported for 2014, presumably because this was the start of the two-year transition period. We start most of our analysis below in the second year of FATCA Form 8966 implementation, for tax year 2015.

Many FFIs purportedly faced difficulties ascertaining information on some account owners, even though in virtually all countries they are required to know about the identity and citizenship of their account owners due to international law enforcing sanctions and anti-money-laundering rules. FFIs often report only partial owner information on the Form 8966, often, notably, not including a Taxpayer Identification Number, i.e. a Social Security Number for an individual or Employer Identification Number for an entity. For Model 1 countries, where the IRS receives information indirectly via foreign governments, the IRS indicated that it would not find FFIs in significant non-compliance (which would trigger withholding) over missing TINs until tax year 2020. The IRS issued periodic guidance about this, essentially requiring that the FFIs make reasonable efforts to procure owner TINs and other information, and that they collect all requisite information for newly opened accounts (see, e.g., IRS, 2017). However, our inability to observe TINs and birth dates when FFIs do not report them causes difficulty in linking a non-trivial share of owners to their tax returns, which is a publicly acknowledged limitation of the early years of FATCA reporting (TIGTA, 2018).⁶ We provide data on this issue in the next session.

Missing TINs are not the only data quality issue we must address. We also exclude records containing no financial information (account balance and income variables), a large number of

⁵ Information about recalcitrant account holders is reported on a pooled basis by type of account identified. In this case, the number of accounts of each type and the aggregate account balance of the pool are reported. We do not use data from reports on recalcitrant accounts in this paper.

⁶ We build on internal work within IRS that uses name matching to assign TINs to Forms 8966 when TINs are missing. We use the TINs obtained by name matching where they are available and do not count these as missing TIN observations in our analysis. Use of name matching reduces the share of accounts with unmatched TINs by about 10 percentage points.

duplicate records and, after careful vetting, a small number of observations where dollar-valued fields contained extremely large and suspicious values that would, if not disregarded, skew aggregate statistics significantly. Appendix B describes the underlying raw data and our data cleaning procedures in more detail.

Prior to FATCA and continuing after it, U.S. individuals were required to self-report certain information about foreign-held assets. Since the 1970s, individuals and certain U.S. entities were required to report on offshore accounts with an aggregate value of over \$10,000 at any time during the year. This reporting happens on FinCEN Form 114, the Report of Foreign Bank and Financial Accounts, commonly referred to as the FBAR. The FBAR was introduced as part of the Bank Secrecy Act of 1970 as a means to discourage and prevent money laundering and tax evasion. Beginning in 2012, “specified U.S. persons” were also required to report “specified foreign assets” to the IRS using the new Form 8938 if the aggregate value of such assets exceeds the reporting threshold of \$50,000 on the last day of the tax year or \$75,000 at any time during the tax year. Specified foreign assets subject to reporting generally include foreign accounts held with FFIs and directly held assets such as stocks or securities issued by a foreign corporation, interest in a foreign partnership, trust or estate, or generally any financial instrument with an issuer that is not a U.S. person. There are a number of instances in which the offshore assets would have already been reportable before the introduction of the Form 8938 under FATCA. These exceptions to reporting requirements are one of many reasons, together with some differences in the set of reportable accounts,⁷ why the aggregate totals from the Form 8966, our main focus here, are not directly comparable to aggregate self-reported information from FBARs or Forms 8938, especially for assets held through a U.S. or foreign company, although we provide some rough aggregate comparisons below.

3. What Do the FATCA Data Show?

3.1 Aggregate

Table 1 provides a broad overview of the data reported by Foreign Financial Institutions (FFIs) on the FATCA Form 8966, from 2015 to 2018.

Overall reporting increased over time along most dimensions. FFIs from 178 countries reported on at least some of their accounts in TY2015, increasing to 190 countries by 2018. The countries driving this increase are small countries and not financial hubs or tax havens. The overall number of reporting FFIs and financial accounts reported increases more rapidly, indicating that reporting expanded within many countries over time. Overall, almost 27,000 institutions reported on just

⁷ Generally speaking, the Form 8938 and the Form 8966 were designed with tax compliance in mind, so the reporting rules usually correspond to the situations in which the filer of the form should report the income derived from the assets for tax purposes, and the 8938 contains further exceptions to avoid double reporting income that would already be reportable to the IRS on other forms under existing rules. In contrast, the FBAR was originally designed to prevent and/or penalize tax fraud, money laundering and other financial crimes. For this reason, the filing requirements are designed to cover essentially all situations in which an American individual or entity controls any offshore financial accounts. The definition of an “account” is also somewhat narrower for FBAR filing requirements.

over 2.0 million accounts in TY2015, increasing by TY2018 to over 45,000 institutions reporting about 4.6 million accounts. Along with the number of reported accounts, the dollar value of total reported assets more than doubled from TY2015 to TY2016, from about \$1.6 trillion to \$3.6 trillion. Total assets changed comparatively little after this, which is consistent with 2015 marking the end of the two-year transition period discussed above. From TY2016 to TY2017, we observe a slight decrease in total assets to \$3.2 trillion, but then total reported assets surged again in 2018 to almost \$4 trillion. In 2018, \$4.0 trillion in assets compared to about \$80 trillion of total household financial assets.

The non-monotonicity in total reported assets from 2016 to 2018 is driven by a small number of large reported accounts at three FFIs in a single country; we were unable to ascertain whether these are valid records (see Appendix B for further discussion). We include these accounts in Table 1, and illustrate in Appendix Table B1 how different ways of handling them would affect account value totals. Regardless of how we handle this issue, we obtain the same qualitative finding: a very large increase in reported wealth in 2016, and another large increase in 2018, with comparatively little change in the intervening year. Both increases are driven by significant increases in reported assets and accounts in a large number of countries, suggesting broad increases in Form 8966 reporting within many countries during these years.

A number of observed accounts contain missing TINs, for reasons discussed above. Virtually all observations with missing TINs come from countries with a Model 1 IGA (99.4% in TY2018). We break down the records with missing TINs further below, and in Appendix B4. For those accounts where the TIN is not missing, we count 790,000 distinct owners in TY2015, increasing to almost 1.5 million U.S. owners by TY2018.⁸ In Appendix B4, we show that missing TINs are especially common in countries where exchanges of information were occurring prior to FATCA under Tax Information Exchange Agreements (TIEA) and/or Double Tax Conventions (DTC); all such countries have Model 1 IGAs. Stricter enforcement since 2018 should reduce the number of missing TINs in future years.

The dollar totals for income reported on Form 8966 also require careful interpretation due to data quality issues. While a positive account balance is reported for over 95% of the accounts in any given year,⁹ in TY2018 just 38% of accounts report any amount of interest income, and an even smaller share report amounts in the other income fields on the form. In TY2018, 45% of accounts have no reported income on any of the four income lines of the form (Part IV Lines 4a-4d), and 41% of total wealth is associated with an account with no income on any of the four lines. Missing observations of variables here include situations where the form was blank and those where the number zero appeared in the field. As we discuss further in Appendix B4, it is difficult to ascertain whether FFIs entered zeros because they were not reporting a value or because there really was no income of the given type and, conversely, whether blank entries

⁸ The researchers worked with masked, i.e., anonymized, TINs throughout the analysis to preserve the anonymity of owners.

⁹ In some cases, e.g. when Form 8966 is filed by a withholding agent, the filer may not know the account balance and may not need to report it. Refer to the Form 8966 instructions for withholding agents in filling out part 4 of the form for details.

should be interpreted as zeros. In Appendix B4, we show that in about 81% of instances where the account balance was missing or zero, the FFI reported income of zero, suggesting that many unobserved account balances are in fact empty accounts. The income variables, in contrast, are far more often blank than zero, and a large share of even very high-dollar accounts have blank income information, leading us to believe that in many cases the assets in the accounts generated income, but the FFIs did not report this on the form 8966. We discussed this issue with experts on the data within the IRS, who agreed with our overall interpretation of these patterns.

Table 1. Overview: Reporting by Financial Institutions on FATCA Form 8966

	Tax Year			
	2015	2016	2017	2018
No. of countries with reporting FFIs	178	179	188	190
No. of reporting FFIs	26,652	36,056	41,829	45,308
No. of accounts	2,008,607	3,703,159	4,225,689	4,566,774
No. of accounts without identified US owners	706,292	1,318,291	1,594,459	1,664,587
No. of identified US owners	789,785	1,223,115	1,296,462	1,477,183
No. of accounts with observed positive balance	1,928,824	3,574,365	4,044,289	4,310,249
Total Account Balance (Billions USD)	1,646	3,648	3,233	3,982
No. of accounts with observed positive income of any kind	1,103,668	1,696,113	2,421,580	2,377,143
Total Interest (Billions USD)	7.9	8.9	13.3	13.2
Total Dividends (Billions USD)	5.3	7.9	26.9	28.5
Total Gross Proceeds/Redemptions (Billions USD)	89.9	218.7	237.4	274.5
Total Other Income (billions USD)	126.7	196.0	290.7	208.0

Notes: This Table presents aggregate statistics from the cleaned Form 8966 data for tax years 2015 to 2018. The top panel reports counts. “Identified owners” refers to those whose accounts we link to their US owners via tax forms. Note that “accounts without identified US owners” counts each account once, while “number of identified US owners” counts each owner once. The bottom panel focuses on the financial information from the form. Along with totals for each dollar-valued field, in billions USD (converted from local currencies where necessary), we report the number of accounts with non-missing, non-zero accounts, and the number of accounts with non-missing or non-zero entries in at least one of the four income fields.

Bearing all this in mind, in TY2018 we observe about \$13.2 billion in total for reported interest income, \$28.5 billion in dividends, \$274 billion in gross proceeds and redemptions, and \$208 billion in “other income.” The gross proceeds amount refers to gross amounts from sales of assets. To translate this gross flow into an amount of taxable capital gains, one would have to subtract the total cost basis for all these sales, which is not reported on the Form 8966. We note, though, that in the most recent public data on the sales of capital assets (SOI, 2012), the total net amount for all capital gains and losses was about 14% of the overall gross sale amount.¹⁰ Multiplying the total for gross proceeds in 2018 by this 14% figure yields a very rough approximation of \$38 billion in net capital gains in 2018. We caution that this and other further

¹⁰ This figure varies significantly across the type of asset. For example, the number is 9.4% for sales of corporate stock, 3% for mutual funds, 28% for interests in partnerships, S corporations, estates, or trusts, 4% for residential real estate, and 41% for sales of farmland. Presumably, much of this is driven by heterogeneous holding periods by asset class. In any case, our use of the overall 14% figure should be considered a very rough approximation.

calculations based on the amount of gross proceeds are uncertain, because the appropriate conversion factor from gross proceeds to net capital gains is unknown. We face similar uncertainty for other income, which can include many types of net and gross income, or even non-taxable flows into the account.¹¹ Given the magnitude of the total, we suppose that most of this income is a gross amount and also scale it by 14% when we approximate a net amount of income, but this is obviously uncertain as well.

What do these data suggest about rates of return on offshore wealth? One reason answering this question is important is that prior estimates of how much tax revenue was lost to offshore tax evasion before FATCA assumed a nominal, taxable rate of return to convert estimated concealed offshore wealth to a taxable income flow (Alstadsæter et al., 2019; Guyton et al., 2021). Understanding rates of return would also help us approximate total income in offshore accounts where income information is missing. However, the data quality issues discussed above make inferring a rate of return from these data more difficult than one might naively expect. We illustrate what the data suggest about rates of return in Table 2.

In Table 2, we estimate rates of return by restricting our attention to specified income flows and to accounts where the specified income flows are reported. To begin, we consider interest income, so we divide total interest income (\$13 billion in 2018) by the total account balance *from observations where interest is present* (33% of \$4.0 trillion, or about \$1.3 trillion in 2018) to obtain what we label the “quasi-rate of return” using interest only. For interest only, the estimates range from 0.6% to 2.2% over the years we consider.¹² Adding dividends into the numerator, we cover only a slightly larger fraction of wealth, 38% in 2018, and we obtain significantly larger returns: the implied return in 2018 increases from 1.0% to 2.8% in 2018. Incorporating capital gains into the income flows and again assuming a 14% conversion factor from gross proceeds to net income, we estimate a 4.6% quasi-return in 2018, based on 44% of total wealth. Incorporating other income, again with the 14% conversion factor, increases the share of wealth covered to 59%, but both the numerator and the denominator of the quasi-rate of return increase proportionally, leaving the rate of return relatively unchanged. These estimated rates of return are close to, but are slightly smaller than, the rough approximations of around 6% used in prior work (Alstadsæter et al., 2019; Guyton et al., 2021). A substantial share of the wealth for which we observe income appears to be generating returns from equity rather than just interest on bank deposits, which would be consistent with portfolio allocations in Swiss data (Zucman 2013, Table II). In the next

¹¹ In many of the observations with very large account value totals, the only type of income present on the 8966 report is “other income,” and interest, dividends, and gross proceeds are blank. This finding is consistent with IRS’s instructions for filling out Form 8966 for assets held through an NFFE, which give the filer of the Form 8966 the option to report in this combined fashion.

¹² The relatively high estimate of 2.2% in 2015 derives from the fact that havens reported a larger share of wealth in 2015, and the quasi-rate of return using interest only is significantly larger in havens. We document and explain these patterns below. The relatively low return in 2016 is partly driven by the large accounts reported in 2016 but not 2017 in a particular country, as described in Appendix Table B1 and the discussion above.

section, we examine further heterogeneity in the totals from Table 1 and in the quasi-rates of return in Table 2.

Table 2. Nominal Rates of Return on Offshore Wealth Implied by FATCA Reports

	Tax year			
	2015	2016	2017	2018
Total reported wealth (billions USD)	1,646	3,648	3,233	3,982
Share of wealth with reported interest	21.4%	42.1%	34.4%	33.3%
Quasi-rate of return: interest only	2.2%	0.6%	1.2%	1.0%
Share of wealth with reported interest or dividends	25.1%	45.1%	40.5%	37.7%
Quasi-rate of return: interest + dividends only	3.2%	1.0%	3.1%	2.8%
Share of wealth with reported int., div., or gross proceeds	43.0%	51.0%	48.2%	44.1%
Quasi-rate of return: int. + div. + 0.14*gross proceeds	3.6%	2.5%	4.7%	4.6%
Share of wealth with reported int., div., G.P., or other income	67.0%	67.9%	69.2%	59.2%
Quasi-rate of return: int. + div. + 0.14*(GP + other)	3.9%	3.0%	5.1%	4.6%

Note: This table summarizes what the FATCA data suggest about rates of return in offshore accounts. Note that we do not observe income for a substantial portion of wealth, as discussed further in the main text and Appendix B4. After the total account balance in the first row, we characterize the subset of observations where we observe interest income, and report the share of wealth for which we observe interest, along with total interest divided by total wealth within this population. We then do the same for observations where we observe either interest or dividends. For gross proceeds, we scale the total by 14%, a conversion factor based on aggregate SOCA data from SOI (2012). We also scale other income by a scaling factor, as its overall magnitude is suggestive of gross flows rather than net income. We caution that, for these reasons, estimates incorporating gross proceeds and/or other income are more uncertain.

3.2 By Owner Type

Table 3 breaks down the offshore wealth reported on Form 8966 by the type of U.S. owner. Refer to Appendix B3 for details on how we match accounts to their owners.

A key finding from Table 3 is that large accounts are disproportionately owned by partnerships. Just 1.4% of accounts are owned by partnerships, but the wealth in these accounts amounts to 32% of the wealth reported on Form 8966 in 2018 (plus whatever share of unmatched wealth belongs to partnerships). Matched individuals, in sharp contrast, own 55% of accounts and just 16% of wealth. Another 14% of the wealth and 1.1% of accounts belong to other entities, mainly C corporations.¹³ The remaining 38% of wealth and 42% of accounts belong to owners whom we cannot link to their tax returns. The largest component of this group represents accounts without a TIN, which contain 26% of the wealth and 37% of accounts. Some unmatched wealth also belongs to owners with TINs that match forms filed by entities with EINs (e.g., Forms 940–945) but not an entity tax return (“unmatched entity”), to owners whose TIN matches a valid SSN but

¹³ Note that publicly traded companies are exempt from FATCA reporting, so these should be private C corporations.

not to a tax return (“unmatched individual”), and to owners with TINs that match to multiple different types of tax returns. We suspect that a large majority of unmatched accounts are owned by individuals, based on hand inspection of a number of records and further results below. The share of the unmatched wealth belonging to individuals versus partnerships or other entities is more difficult to determine with the current data. Table A1 reports analogous data for income fields. The shares of the income flows by owner type are roughly in line with the shares of wealth. Most income types are slightly more concentrated in partnerships than wealth shares; other income, in particular, is very concentrated in accounts owned by partnerships.

Table 3. The Distribution of Owner Types, TY2018

		Account Balance		No. of accounts	
		Total (Billions USD)	Share	Total	Share
Matched	Partnership	1,292	32.4%	62,422	1.4%
	Individual	618	15.5%	2,516,330	55.1%
	C Corporation	401	10.0%	20,345	0.4%
	Tax exempt entity	49	1.2%	9,363	0.2%
	Trust	47	1.1%	9,869	0.2%
	Foreign corporation	21	0.5%	6,491	0.1%
	S corporation	37	0.9%	8,515	0.2%
Unmatched	Missing TIN	1,018	25.5%	1,695,750	37.1%
	Unmatched entity	279	7.0%	13,971	0.3%
	Ambiguous match	154	3.8%	7,338	0.2%
	Unmatched TIN	60	1.5%	66,064	1.4%
	Unmatched individual	7	0.1%	150,316	3.3%

Note: This table reports total wealth (account balances) and the number of accounts disaggregated by owner type. We observe that ownership of accounts is much more concentrated among individual owners, while wealth is more concentrated among entity owners. Unmatched accounts are those that we do not match to their owners’ tax returns. We find that most unmatched wealth comes from observations with missing TINs.

How do the aggregates from Form 8966 compare to what was, and is still, being reported by the taxpayers on Foreign Bank Account Reports (FBARs) and the Form 8938? In the case of direct individual ownership of a depository account, we should expect significant overlap between Form 8966 third-party reports and self-reports on these other two forms. Differences in the requirements for accounts held indirectly through U.S. or foreign entities mean that, in these more complex situations, it is harder to compare across forms. For example, Part IV exceptions on the Form 8938 create situations where values of assets are reportable on Form 8966, but not on Form 8938. These exceptions cause less wealth to be reportable by taxpayers on Form 8938 than is reportable by FFIs on Form 8966. In 2016, we observe \$843 billion in reported wealth on FBARs, of which \$373 billion is reportedly owned by individuals.¹⁴ In that same year we observe \$417 billion in wealth reported on Forms 8938, of which \$386 billion comes from individuals. On

¹⁴ Both of these figures exclude FBARs where the filer has signature authority over, but no financial interest in, the account.

Form 8966, we observe \$3.57 *trillion* in total wealth for 2016, but of this, \$370 billion belongs to matched individuals. The numbers for individuals are comparable, though significant unmatched 8966 wealth likely belongs to US individuals. The FATCA Form 8966 data provide much more comprehensive coverage of offshore wealth, especially for accounts with indirect or complex ownership, which tend to be large accounts.

The next set of results analyzes wealth held in tax havens and in non-tax haven countries separately. Confidentiality rules under tax treaties and other agreements between the U.S. and foreign governments prohibit us from publicly disclosing information on specific countries. We summarize the most important differences across countries by grouping countries into tax havens and non-havens. Separately analyzing haven wealth also allows us to compare what we see in the FATCA data to findings in other research on offshore tax evasion, which often focuses on wealth held in havens. Appendix Table A2 breaks down the dollar totals from Table 1 by haven and non-haven status over time, using the same set of tax havens as Johannesen et al. (2020).¹⁵ We observe \$1.9 trillion in wealth in tax havens in 2018, which is just under half of global reported wealth. For wealth held by individuals and partnerships only, we observe \$1.4 trillion in havens. These numbers for the year 2018 compare to rough estimates of U.S. haven wealth based on macro financial data for the year 2007 of \$1.2 trillion from Zucman (2014), and \$1.0 trillion from Alstadsæter, Johannesen and Zucman (2018). The \$1 trillion is about 7% of GDP in 2007, while our estimate is about 10% of GDP in 2018. This does not necessarily imply the \$1 trillion figure circa 2007 was underestimated: with minimal repatriation/reallocation and a growth rate of 5 to 6% per year (i.e., faster than GDP), \$1 trillion in haven wealth in 2007 would grow to something close to our totals by 2018.

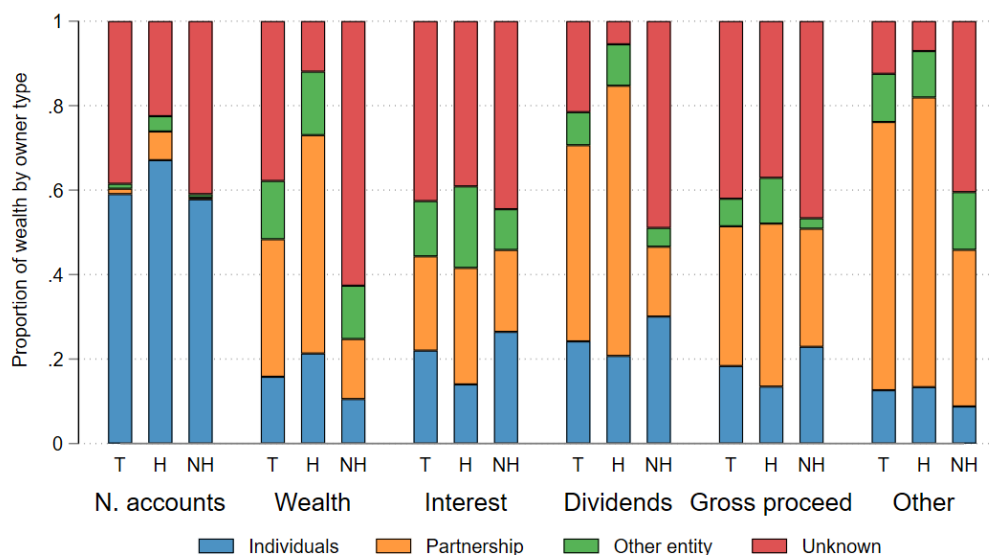
Figure 1 illustrates how the owner type distribution varies between havens and non-havens. For simplicity, we create four owner-type categories, building on Table 3: individuals, partnerships, other matched entities, and unknown owners (i.e., all types below the horizontal line that divides Table 3). In Figure 1, we report the wealth and account shares from Table 3, splitting accounts into accounts in tax havens and accounts elsewhere. We report wealth, account, and income shares for 2018; these shares are similar in earlier years.

We take two key findings away from Figure 1. First, a sizable share of all wealth in tax havens is owned by U.S. partnerships – 52% plus partnerships’ share of unmatched wealth. Matched partnerships own a comparably small share, 14%, of the wealth in non-havens. Following our main finding from Table 1 that the largest accounts are owned by partnerships, here we add that most of this wealth is in tax havens. Second, the accounts with unmatched owners are

¹⁵ The full set of countries is as follows: Anguilla, Antigua and Barbuda, Aruba, Bahamas, Bahrain, Belize, Bermuda, Cayman Islands, Cook Islands, Cyprus, Grenada, Guernsey, Hong Kong, Isle of Man, Jersey, Liechtenstein, Luxembourg, Malta, Mauritius, Monaco, Panama, Saint Kitts and Nevis, Saint Lucia, Singapore, Switzerland, Turks and Caicos Islands, and British Virgin Islands. Additionally, in some instances, the Form 8966 indicated, incorrectly, that offshore assets were located in the U.S. On inspection, the assets for which this occurred are cases where the filer reported a U.S. country code, likely because the investments are managed in the U.S., but the assets are virtually all domiciled in havens. We include these assets in the haven category as well. Using alternative groupings of countries, such as that proposed by Love (2022), leads to very similar findings.

disproportionately located in non-havens, representing 62% of the wealth in non-havens and just 12% in havens; tax havens, perhaps because of the scrutiny they now face, report relatively few unmatched accounts. Everything we find here is consistent with the analysis of FATCA reporting in Belnap, Thornock, and Williams (2019). Appendix Tables A3 and A4 break down the aggregates in Table 3 and Figure 1 further, first dividing non-haven countries into those that did and did not have exchange of information (EOI) agreements with the U.S. under TIEAs and/or DTCs prior to FATCA in Table A3, and then splitting by owner type in Table A4. The main insights that emerge from these additional decompositions are that 1) a large share of non-haven accounts is located in countries where there was some information exchange pre-dating FATCA, and 2) unmatched accounts and wealth are especially concentrated in countries with prior information exchange agreements. Appendix B4 contains some supplemental reporting quality statistics.

Figure 1. Reported Accounts, Wealth, and Income by Owner Type and Haven Status (TY2018)



Notes: This figure illustrates how the distribution of owner varies between havens and non-havens. The first three bars depict the share of accounts in total (T), in haven countries (H), and in non-haven countries (NH) belonging to each of four owner categories. We then illustrate the same shares for dollars of wealth, and for the four income fields. We observe that ownership of accounts is highly concentrated among individuals and unmatched owners. Partnerships own a much larger share of wealth than accounts, especially in tax havens. Like wealth, other dollar-weighted fields are concentrated among partnership owners.

In Table 4, we consider how rates of return vary between owner types and between haven versus non-haven accounts. To do this, we estimate the same quasi-rates of return from Table 2 for these sub-populations, using TY2018 data. We focus on the quasi-rate of return including interest and dividends only and report the other specifications in Table A5. The results suggest that accounts in tax havens have significantly higher returns than accounts in non-havens for all specifications of the quasi-rate of return. Considering interest and dividends, we calculate an implied rate of return of 5.0% in havens and just 1.8% in non-havens. Accounts owned by matched partnerships also have a relatively high rate of return of about 6.4%, and accounts owned by individuals have a return of 4.6%; accounts owned by other entities and unmatched

owners have much smaller returns of 1.1% and 1.6%, respectively. We observe similar comparisons for other definitions of the quasi-rate of return, though the definition of the quasi-rate of return does affect the level of the return.

Table 4. Rates of Return in Selected Subpopulations, TY2018

Sub-population	Total reported wealth (billions USD)	Share of wealth with reported interest or dividends	Quasi-rate of return: interest + dividends only
All accounts	3,982	37.7%	2.8%
Non-haven country	2,042	51.3%	1.8%
Haven country	1,940	23.4%	5.0%
Individual owners	626	33.6%	4.6%
Partnership owners	1,292	19.8%	6.4%
Other entity owners	279	37.7%	1.1%
Unmatched owners	1,510	49.9%	1.6%

Notes: This table reports estimated quasi-rates of return including interest and dividends only in selected subpopulations. We observe significantly higher returns in tax havens, and higher returns among partnership owners and to a lesser extent individual owners. Table A5 reports additional specifications of quasi-rates of return with similar results.

What picture of the offshore world emerges from our results so far? In making sense of these data, it is useful to consider two types of owners of offshore wealth: sophisticated global investors, and Americans with close ties to specific foreign countries (e.g., recent immigrants or those who live or lived abroad). A typical account in these data is likely to be owned by the second type of owner: the median account is a relatively small account directly owned by an individual in a non-tax-haven country. Looking at dollar-weighted statistics rather than account-weighted statistics provides a very different picture, however. A given dollar of wealth in these data is far more likely than a given account to be located in a tax haven and/or owned indirectly by a partnership, suggesting the owner is likely a more sophisticated investor. Next, we provide a descriptive analysis of the tax returns of matched individual and partnership owners, which largely bears out this view.

We note that these findings may not be surprising to experts on international finance. It is common knowledge, for instance, that many hedge funds with U.S. investors operate out of havens, especially the Cayman Islands, to sidestep regulations and facilitate tax avoidance by fund managers.¹⁶ However, administrative data on the full scope of these offshore activities has not been available until now. The data strongly suggest that we need to account for these types of holdings if we want to fully understand the offshore wealth of Americans and related questions

¹⁶ For instance, Brocard and Lhabitant (2016, Figure 1) estimate that 25.5% of hedge funds, and 34% of non-US domiciled funds, are legally domiciled in the Cayman Islands.

about offshore compliance. The degree to which American investors with these types of holdings were tax compliant before FATCA was enacted is unclear.

3.3 Offshore wealth held by partnerships: Details and ownership linkages

In this section, we take a closer look at offshore wealth held by partnerships, motivated by the findings in the previous section on the importance of partnership wealth.

We begin by examining the industries of the partnerships holding foreign accounts. Figure 2 shows the largest industries, in terms of assets. Foreign account and asset holdings are highly concentrated in a small number of industries, mostly within the finance and real estate sectors. The largest industry, “Other Financial Investment Activities”, accounts for 70% of total partnership assets and 60% of partnerships. Other finance and real estate industries account for another 20% of assets and 16% of partnerships. Only a trivial share of assets is held by non-finance industries, with some having missing or invalid industry information on the partnership tax return.¹⁷ Figure 2 and Appendix Table A6 also displays information about where in the world partnerships are holding their assets. Approximately 77% of partnership assets are held in countries often considered to be tax havens. The distribution of assets across industries is very similar between assets held in havens and non-havens; they are each equally concentrated in finance industries.

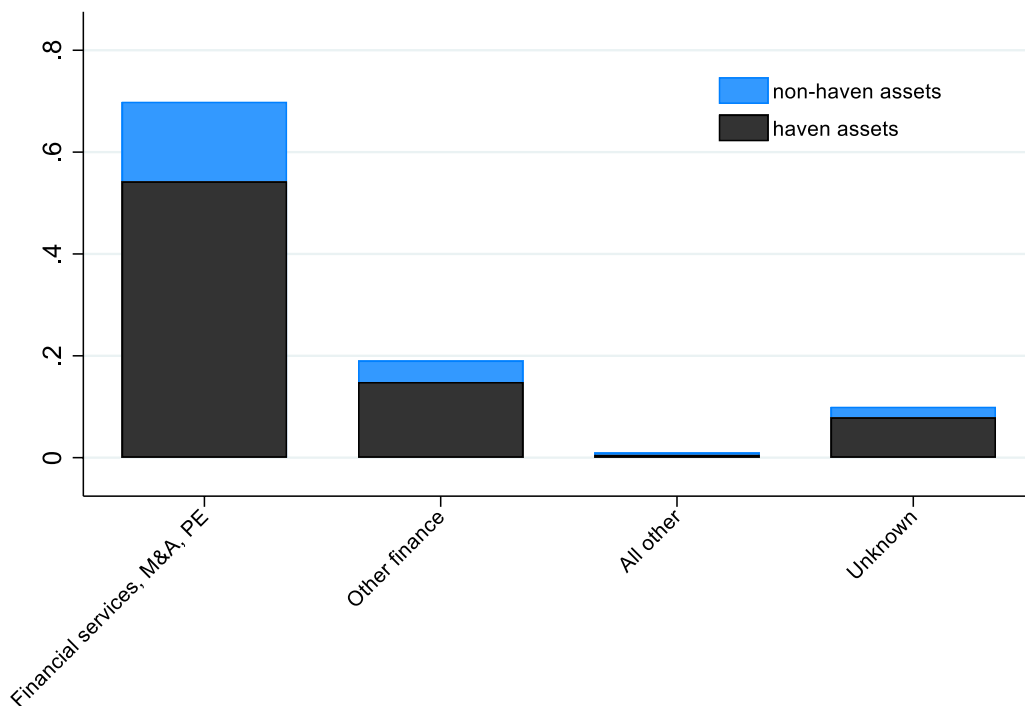
To benchmark the concentration of partnership foreign assets in finance industries, we compare this to population statistics on partnership income. Using the population of U.S. partnerships in 2011, Cooper et al. (2016) find that 70% of partnership income is concentrated in financial industries, compared to 90% of foreign assets being held by partnerships in these industries. While partnership income is generally highly concentrated in the finance and real estate industries, foreign asset holdings of partnership assets are even more concentrated in these industries.

Partnerships are “pass-through” business entities, meaning income and losses are not taxed at the business level but are distributed to and taxed as part of the personal income of the shareholders. Partnerships can have shareholders that are individuals or business entities, and the partnership must issue a Schedule K-1 to each shareholder detailing the income items attributable to the shareholder in that year. Using the K-1 reports, we link partnerships to their shareholders, looking through ownership tiers, to allocate foreign assets owned by partnerships to the ultimate taxable owner. For each pass-through account owner, we allocate the reported foreign assets to shareholders according to the share of that partnership’s reported income allocated to each partner, as reported on the Schedule K-1. If the shareholder is a taxable

¹⁷ Descriptions of these industries can be found at <https://www.naics.com/naics-code-description>. The “Other Financial Investment Activities” industry (code 5239) can include financial services, hedge funds, private equity, mergers and acquisitions and holding companies. Mortgage real estate investment trusts (REITs) are included in “Other Investment Pools and Funds” (code 5259).

individual or entity, they are a terminal node. If, however, the shareholder is another pass-through (an S-corporation or partnership), then the assets allocated to the pass-through shareholder are further allocated to its shareholders according to their income shares. This process is iterated until all assets are allocated to a taxable beneficial owner (i.e., after looking through all pass-through entities in the ownership chain), or until there is no beneficial owner for that portion of the assets to be allocated to. This happens when the owner reported on the K-1 cannot be identified in the tax data, when an invalid Taxpayer Identification Number (TIN) is reported on the K-1, or when the TIN matches with one or multiple owners whose owner type cannot be verified. We call all of these “unclassifiable” owners.

Figure 2: Shares of Partnership Assets by Industry, held in havens and non-havens

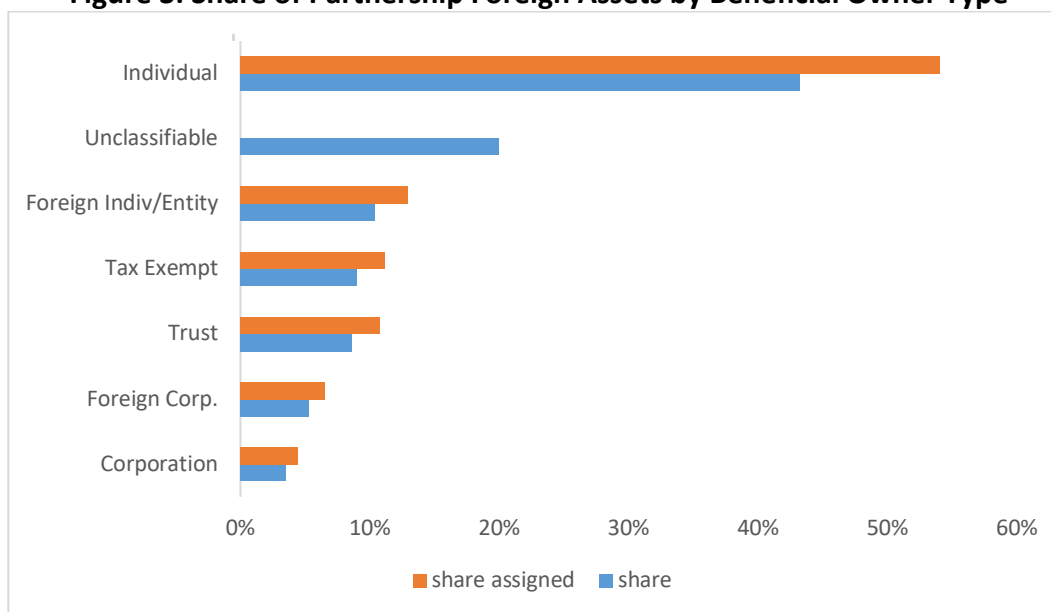


Notes: Figure 2 shows the shares of total assets held by pass-through account owners by the industry of the partnership, split by whether assets are held in haven or non-haven countries. Approximately 70% of foreign assets are held by partnerships involved in Financial Services, Mergers and Acquisitions, Holding Companies and Private Equity (NAICS 5239), and over half of total partnership assets are held in havens by partnerships in these industries. A remaining 20% are held in other finance or real estate industries (NAICS 52, 53, 55) and only 10% are in other industries. Appendix Table A6 shows that the distribution of assets across these industries is very similar for assets held in haven and non-haven countries.

Figure 3 shows the shares of partnership foreign assets allocated to each type of taxable owner. The blue bars show the full allocation of partnership assets. Approximately 43% of assets held by partnerships are owned by U.S. individuals, 10% are owned by foreign individuals or entities, 8% by tax exempt organization and another 8% by trusts. Notably, 20% of assets cannot be linked to beneficial owners. This is similar to the share of partnership income that Cooper et al. (2016)

were not able to allocate to beneficial taxable owners.¹⁸ The orange bars show the distribution of assets that could be allocated to classifiable owners. About 54% of the assets that could be credibly attributed to a taxable owner are owned by U.S. individuals, with over 10% each owned by foreign individual and entities, tax exempt organization and trusts. Appendix Table A7 reports the asset shares along with the share of owners of each type. Almost 90% of beneficial owners are individuals, so the entity owners account for much larger asset shares on average.

Figure 3: Share of Partnership Foreign Assets by Beneficial Owner Type



Note: Figure 3 distributes the total foreign assets held by partnerships to the ultimate taxable owner of the partnerships. The blue bars account for the total distribution of foreign assets, including those going to unclassifiable owners in the tax data, cases where the owner reported on the K-1 cannot be identified in the tax data, an invalid Taxpayer Identification Number (TIN) is reported on the K-1, or the TIN matches with one or multiple owners where the owner type cannot be verified; these represent 20% of total assets. The orange bars show the distribution of ownership for the assets that can be assigned to classifiable taxable owners. Of the assets that can be directly assigned, about 54% belong to individual owners, 12% to foreign individuals and entities, 11% to tax exempt entities, 11% to trusts and the remaining to foreign and domestic corporations. Appendix Table A7 presents asset and owner shares by owner type.

3.4 Combining Individuals and Linked Partnerships: Offshore Wealth Through the Individual Income Distribution

The FATCA data provide a rare opportunity to understand not only how foreign assets are owned, but who owns them. Particularly, we can examine where foreign account owners are in the income distribution and the corresponding concentration of foreign asset holdings. We find that

¹⁸ It is not obvious that these shares should be so similar. Only a subset of partnerships own foreign assets, and assets do not need to be distributed across partnerships the same as income. This does suggest that partnerships that own foreign assets are not particularly more or less likely to have unclassifiable assets relative to the population of partnerships. Perhaps this is not too surprising, given that most unallocated income found in Cooper et al. (2016) is from finance industries and the foreign assets are similarly concentrated in these industries.

foreign account and asset ownership are extremely concentrated at the very top of the income distribution, that assets held indirectly through pass-through entities are substantially more highly concentrated than directly held assets, and that assets held in tax havens are more concentrated than those held in non-havens.

To examine the composition of ownership across the income distribution, we must assign foreign assets to their ultimate individual owners. For individuals that hold the foreign account directly (“individual” account owners from Table 3), this is straightforward. We assign all reported foreign assets and income to the household (Form 1040 filer) holding the account. For pass-through account owners, we look through the entity and assign assets and income to beneficial individual owners, as described in the previous subsection. The statistics that follow are for foreign accounts directly held by individuals and the 44% of partnership assets that can be allocated to individual beneficial owners.¹⁹ Throughout the main text, we use Adjusted Gross Income (AGI) reported on the individual tax return as our preferred definition of income.

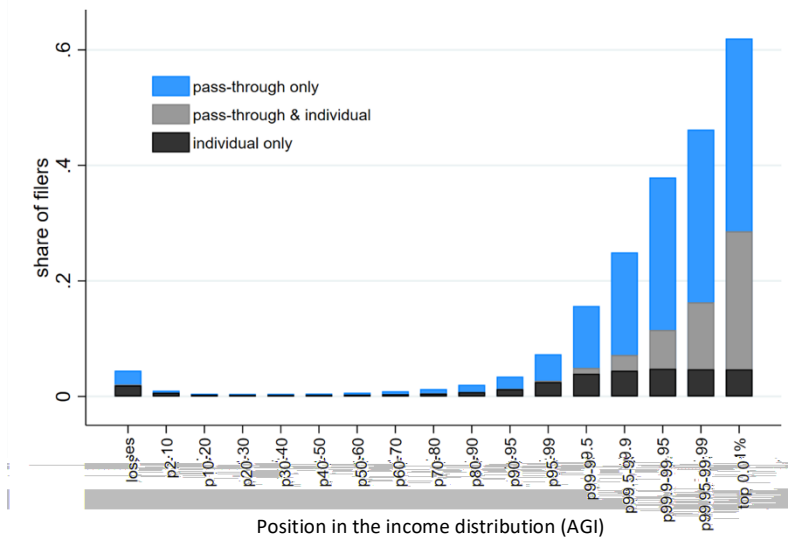
Figure 4 shows the share of taxpayers owning foreign accounts in 2018 either directly or indirectly through pass-through entities by their position in the income distribution. Panel A shows the share of owners by percentile in the distribution of income (AGI) in 2018. The height of the bar represents the total share of taxpayers in that quantile of the distribution that own a foreign account identified on a Form 8966. The blue portion of the bars represents the share holding accounts indirectly through pass-through entities only, the grey portion represents the share that hold some foreign assets directly and some indirectly through pass-through entities and the black bar is the share holding foreign assets directly only. The incidence of foreign account ownership is much higher at the top of the income distribution. Approximately 22% of those in the top 1% hold a reported foreign account, and the incidence sharply increases within the top 1%. About 62% of households in the top 0.01% hold a foreign account identified in FATCA-generated data. The large majority of those at the top that hold foreign assets hold some assets indirectly, and about half also hold some foreign assets directly.

Panel B shows the incidence of holding assets in a tax haven relative to a non-haven. If households have assets in both havens and non-havens, they are categorized as having a haven account. The vast majority of those holding foreign assets at the top of the income distribution have accounts in countries considered to be tax havens. This is true for both accounts held indirectly or directly – 57% of households in the top 0.1% hold some assets in havens. Because a large number of accounts could not be matched to a valid TIN, as discussed in Section 2, these statistics represent a lower bound of the true incidence of foreign account ownership in each quantile.


¹⁹ For the analysis in this subsection, we are only able to examine accounts that were identified on a Form 8966 *and* that could be matched to a valid associated TIN. Unmatched accounts and assets are excluded from this analysis.

Figure 4: Share of taxpayers reported to hold a foreign account by income quantile (TY2018)

Panel A: Share of taxpayers holding accounts directly and/or through pass-through entities



Panel B: By Haven or non-haven



Note: This figure shows the share of taxpayers in each quantile of the income (AGI) distribution that receives a FATCA report (Form 8966) indicating ownership of a foreign account. The share is calculated as the total number of individuals in a quantile listed as an account owner on an 8966 over the total number of tax returns in that quantile. The black bars (individual) represent the share of individual account owners that own a foreign account directly. The blue bars (partnership K-1) represent the share of individual shareholders of partnerships that own foreign accounts. Explicitly, for partnership foreign account owners, we link the shareholders to the partnership through Schedule K-1 and the blue bars represent individual shareholders of these partnerships. Panel A shows foreign account ownership shares by centile of the distribution. Over 20% of those in the top 1% are foreign account owners, with about 14% holding accounts through pass-through entities. Panel B zooms in on the top 1% of the income distribution presenting the shares holding foreign accounts by the top 0.9-0.1% separately, then decomposes the top 0.1% into the 0.09-0.05%, the 0.05-0.01% and the top 0.01%. Foreign account ownership rates increase throughout the top 1%, with approximately 66% of those in the top 0.01% being FATCA account owners, just over half of which hold accounts through pass-through entities.


Next, we examine the distribution of foreign assets as reported on Form 8966. Figure 5 shows the distribution of all foreign assets allocated to households, and the distributions of assets held directly and indirectly and in havens and non-havens. In Panels A and B, the red dashed line shows distribution of all allocated assets. Foreign assets are highly concentrated, with 64% being held by those in the top 1%, and further concentration at the very top with 30% of reported foreign assets being held by households in the top 0.01% of the income distribution. Panel A also shows that assets held indirectly through pass-through entities are substantially more concentrated than directly held foreign assets – 46% of assets held indirectly are owned by those in the top 0.01% of the income distribution (80% in the top 1%). While directly held assets are also concentrated, substantially higher shares are owned by the rich, but not super-rich, meaning those in the 90-99.9th percentiles. Panel B presents the distributions of assets held in haven and non-havens, showing that assets held in havens are more concentrated than those held in non-havens – 35% of assets held in havens are owned by top 0.01% households, compared to about 23% of non-haven assets being owned by this group. Again, while non-haven assets are also quite concentrated, a much larger share are owned by those in the 80-99.5th percentiles.

Panel C decomposes the total shares into the portions held directly and indirectly and in havens and non-havens. About 23% of all assets are held in havens by those in the top 0.01%, implying that 77% of all assets held in the top 0.01% are held in havens (61% of top 1% owned foreign assets are held in havens). Similarly, 77% of assets held by the top 0.01% are held indirectly (74% for the top 1%), and almost 20% of all assets are owned indirectly by those in the top 0.01% (see Appendix Table A8 for associated values).²⁰ While Figure 4 shows a relatively high incidence of ownership in the lowest percentiles and Figure 5 shows substantial asset ownership in these low income percentiles, these are likely driven by high permanent income/wealth owners with low income realizations. Appendix Figures A1 and A2 show the results ranking households by their position in the distribution of total positive income (TPI) and reveal that much of the account ownership and assets at the bottom of the distribution of AGI is associated with households with high total positive income, i.e., households with high income had they not realized large losses in that year. We return to this issue below.


²⁰ Additionally, Foreign Financial Institutions (FFIs) are only required to report on account holders with aggregate asset values over \$50,000. Figure A4 shows an ad hoc adjustment to the distribution of foreign assets to assess the potential importance of this threshold. We arbitrarily assign \$40,000 in foreign assets to 10% of all households in the bottom 90% of the AGI distribution in 2018 (i.e. assume that 10% of households in the bottom 90% have foreign accounts that do not appear in our data because they are slightly below the FATCA reporting threshold). The level and share of assets going to the bottom 90% substantially increases, but the profile remains qualitatively similar; foreign assets would remain highly concentrated at the very top of the income distribution.

Figure 5: Share of assets owned by position in the individual income distribution

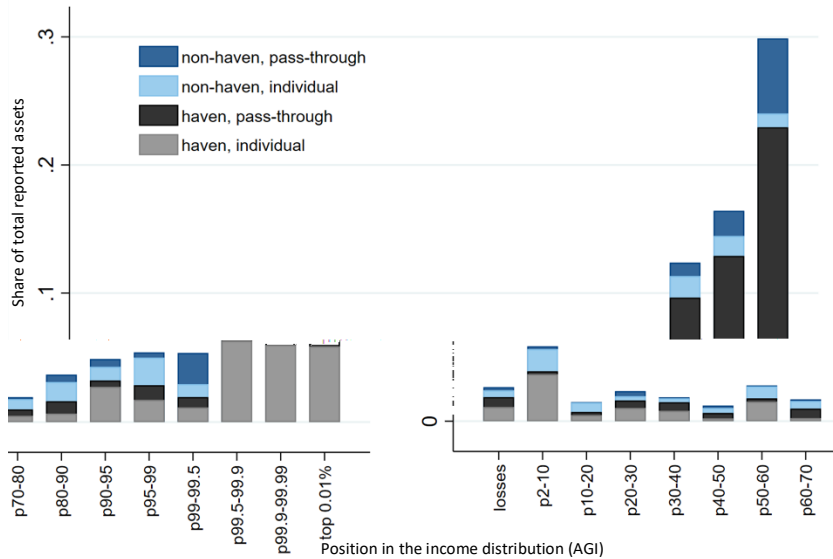
Panel A: Distribution of assets held directly through pass-through entities



Panel B: Distribution of assets held in havens and held in non-havens



Panel C: Distribution of total assets held directly and through pass-through entities, in havens and non-havens

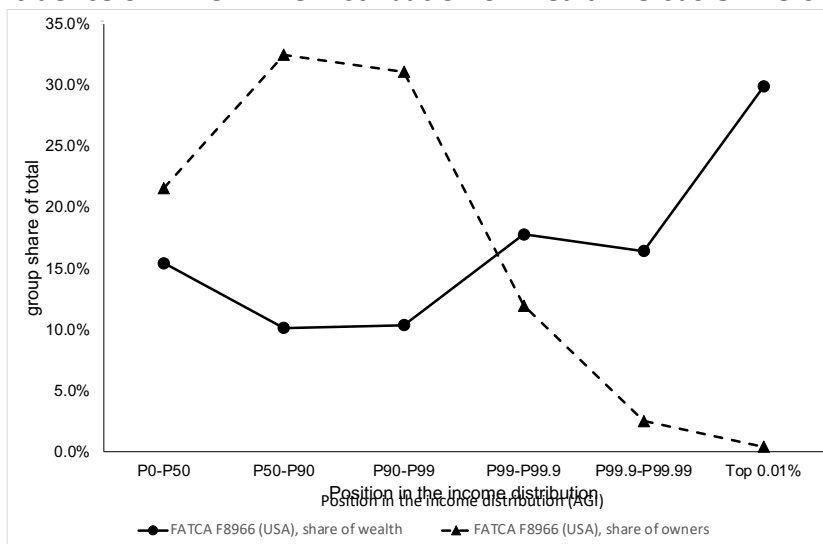


Note: This figure shows the distribution of assets held across the income (AGI) distribution. Panel A depicts the share of assets held by individuals and, separately assets held through pass-through entities (partnerships and S-corporations). Shares are defined as the assets owned by individuals (directly or through pass-throughs) in a given part of the AGI distribution, divided by total assets. Overall assets (red series) are very concentrated with about 28% being held by the top 0.01%. Assets held by pass-through entities (blue bars) are more concentrated, with over 45% owned by the top 0.01%. Directly individually held assets (black bars) are also concentrated, but less so, with about 12% owned by the top 0.01%. Panel B depicts the distribution of assets held in havens versus non-havens. About 33% of assets in havens belong to the top 0.01%. Panel C decomposes the distribution of total assets into amounts held in havens and non-havens, directly and indirectly. Wealth owned by high-income owners is mainly held in havens (grey and black bars), particularly in through pass-through entities (black). Of non-haven assets held at the very top (light and dark blue bars), the majority is held through pass-through entities (dark blue). Appendix Table A8 presents the totals for these groups.

Figures 4 and 5 demonstrate that dollars of offshore wealth are highly concentrated at the top of the income distribution and that ownership of offshore wealth is very common among individuals at the top of the distribution. Nevertheless, most accounts in the FATCA data do not belong to very high-income people. Figure 6 contrasts the overall wealth shares from Figure 5a (combining havens and non-havens) along with the overall share of owners belonging to different parts of the income distribution. The apparent contrast between the small share of accounts belonging to very high-income individuals in Figure 6 and what we see in Figures 5 and 6 reflects the extremely skewed distribution of offshore wealth: a small fraction (3%) of owners in the top 0.1% of the income distribution own a large share (46%) of reported wealth.

We view Figure 6 as informative for two aspects of foreign-held wealth. The wealth shares are likely to be informative about the incidence of the compliance/revenue effect of FATCA, because the income associated with the reported wealth is likely to be similarly concentrated at the top of the distribution. Meanwhile, the share of owners in different parts of the income distribution is informative for the incidence of any compliance costs that must be borne by all owners of offshore accounts as a result of FATCA. Figure 6 suggests that while the tax revenue effects of FATCA may target the very top of the income distribution, any compliance costs attributable to FATCA will fall on a wider set of taxpayers located throughout the income distribution.

Figure 6. The Incidence of FATCA: The Distribution of Wealth versus Owners by Income Rank



Note: This figure contrasts the share of wealth belonging to owners in various parts of the income (AGI) distribution – the same data as the share of total assets from Figure 5 – with the share of owners located in various parts of the income distribution. We observe that while dollars of wealth are highly concentrated among a few high-income owners, most account owners are not extremely high income.

3.5 Comparison to Prior Literature

In this section, we place the results on offshore wealth from the FATCA data within the context of prior literature on offshore wealth and tax enforcement, in order to understand what these results imply about broader questions regarding offshore tax evasion.

As discussed above, prior studies on offshore evasion use data from partial lists of evaders: from taxpayers who either disclosed their offshore account as part of an amnesty or voluntary disclosure programs; from taxpayers implicated in a leak or whistleblowing from a particular financial institution: or, in the case of Johannesen et al. (2020), from those who disclosed an offshore account in the wake of enforcement without participating in the Offshore Voluntary Disclosure (OVD) Program – so-called *quiet disclosures*. We gathered comparable results from these prior studies: the HSBC leak, the Panama Papers and the amnesty disclosures in Scandinavia in Alstadsæter et al. (2019); the Panama Papers and amnesty disclosures in Colombia from Londoño-Vélez and Avila-Mahecha (2021); the amnesty disclosures in the Netherlands from Leeders et al. (2021); and the OVD program participants and first-time FBAR filers with US addresses and tax haven accounts in the United States from Johannesen et al. (2020) and Guyton et al. (2021). We note that the U.S. data pre-date FATCA, and the other datasets pre-date the Common Reporting Standard implemented in other countries, so the consequences of FATCA or CRS may generate differences in the results when we compare across datasets.

Comparing quantitative results across these studies merits caution. The most important methodological difference is that, while the studies on non-U.S. populations rank individuals according to their position in the respective country's wealth distribution, our FATCA data and our pre-FATCA U.S. data rank taxpayers by their position in the *income* distribution. This likely means that some offshore wealth that would be located at the top of the U.S. wealth distribution is located further down in the income distribution--wealthy U.S. taxpayers sometimes have small or negative annual incomes. Subject to these qualifications, we focus on how the concentration of offshore wealth at the top of the distribution described in the previous section compares to findings from other studies. We consider the global FATCA dataset as well as the dataset restricted to accounts in tax havens for the comparisons.

The first and most striking difference between our findings and work on these previous datasets is that the probability that an individual at a given part of the income distribution appears as an owner of offshore wealth is far larger in our data, throughout the distribution. About 60% of individuals in the top 0.01% appear in the FATCA data; for the Colombian amnesty the analogous figure is 41% of the top 0.01% (by wealth), while less than 14% of the top 0.01% appear on every other list. This difference in probabilities is wholly unsurprising, because prior work leverages incomplete lists of owners of offshore wealth, such as lists of participants in an amnesty, while our list in principle includes all owners of offshore wealth.


Setting this level difference aside, we next compare the steepness of the profile of ownership of offshore wealth over the income distribution across datasets. In other words, how much more likely is someone at the very top to appear as an owner, compared to someone near the top? To

answer this, we divide the probabilities of appearing in the dataset at a given part of the distribution by the probability for the P90-P99 income bin, i.e. the top decile excluding the top 1%, in that dataset. Figure 7 Panel A presents the results. We plot the normalized probabilities on a logarithmic scale; to interpret the y-axis, note that if the y-axis takes a value 10 for a given part of the income/wealth distribution, then someone in this part of the distribution is 10 times more likely to appear in the dataset than someone in the P90-P99 part of the distribution.


The results in Figure 7 Panel A suggest that the profile of probabilities of appearing as an owner of offshore wealth in the FATCA data is broadly similar to, but perhaps slightly less steep than, what we might have expected given the profile in most other datasets. Someone in the top 0.01% is just over 10 times as likely as someone in P90-99 of the income distribution to appear in the FATCA data as an owner of an offshore account. This figure is well over 10 in a number of other datasets; the Colombian amnesty stands out as the case with the steepest profile: those in the top 0.01% were over 1,000 times more likely to disclose during this amnesty than those in P90-99. The FATCA profile more closely matches the steepness from the U.S. pre-FATCA datasets, as well as the Scandinavian amnesty. The Dutch amnesty stands out as the dataset with the flattest profile at the top of the distribution. We observe a slightly steeper profile if we restrict the FATCA data to haven accounts, but the difference is very small.

Figure 7 Panel B presents a comparison of the concentration of offshore wealth across datasets. For datasets in which offshore wealth is observed, we plot the fraction of offshore wealth covered by the dataset that belongs to owners in a given part of the income distribution, as in Figure 5 above. All these datasets suggest that ownership of offshore wealth is extremely concentrated, with at least 60% of wealth belonging to the top 1% of the income/wealth distribution in every case. The share of wealth in the top 0.01% in the FATCA data, about 30%, falls in the middle of the wide range of the estimates spanned by earlier studies. Meanwhile, at the bottom of the distribution the shares in the FATCA data are slightly higher than the results of other studies, with 15% of wealth belonging to the bottom 50% by income. Much of this wealth, 8.5% of total matched wealth in the FATCA data, belongs to the bottom 10% by income. We note that the next highest shares at the bottom of the distribution are in the U.S. “New FBAR” data, which also ranks by income. All this suggests that the fact that we can only rank by income with U.S. data decreases the estimated concentration of offshore wealth significantly.

Figure 7: Comparison of FATCA Data with Data on Offshore Wealth from Prior Work
Panel A: Participation Rates



Panel B: Wealth Shares



Note: This figure presents quantitative comparisons of the concentration of ownership of offshore wealth from the FATCA Form 8966 data and several other studies. We and prior studies on the U.S. rank taxpayers by their position in the income (AGI) distribution, while studies on all other countries rank taxpayers by their position in the wealth distribution. In Panel A, we present estimates of the probability of appearing in a dataset of owners of offshore wealth, normalizing probabilities by dividing by the probability in the P90-P99 part of the distribution. We observe that the steepness in the profile of the probability of owning an offshore account in the FATCA data is comparable to or slightly less steep than in most other datasets. In Panel B, for datasets where wealth is observed, we present estimates of the share of offshore wealth belonging to different parts of the distribution. We observe that the FATCA data results are in the middle of the estimates from other datasets when it comes to concentration of ownership at the very top of the distribution, while the share of wealth held by taxpayers at the bottom of the distribution is somewhat larger in the FATCA data than in other datasets.

To explore this further, Appendix Figures A2 and A3 report wealth shares ranking taxpayers by Total Positive Income (TPI) rather than our usual Adjusted Gross Income (AGI). Total positive income is defined as the sum of all positive components of AGI; the main difference between ranking by AGI versus TPI is that, in the latter case, individuals with large business losses – who tend to be wealthy individuals – are ranked further up in the distribution, often moving from the very bottom of the distribution to the very top. We break the results down across havens and non-havens and partnerships and individuals in Figure A2, as in Figure 5B, and we compare the overall shares using TPI or AGI in Figure A3. We do indeed see significant wealth moving further up in the income distribution from the bottom bin in Figure A3. From Figure A2, we observe that ranking by TPI makes a much larger difference for wealth owned by partnerships than for individuals – this makes some sense because partnerships themselves can have pass-through losses. Substantial partnership wealth moves from the very bottom of the distribution to the top 1% when we switch to TPI. Some individual wealth also moves into the top 1%, but compared to partnership wealth, the individual wealth moving away from the very bottom of the distribution relocates more broadly through the distribution. All this suggests that FATCA wealth would be significantly more concentrated at the top of the distribution if we were to rank individuals by wealth instead of income.

We do not estimate the compliance effects of FATCA in this paper, but a number of our findings have implications for ongoing attempts to better understand these compliance effects. Our data pertain directly to owners who held wealth offshore *after FATCA implementation*, so it holds promise for understanding the voluntary compliance effect of FATCA on this group of owners. We discuss the other potential compliance effects below. Similar to Alstadsæter et al. (2019) and Guyton et al. (2021), we can quantify the voluntary compliance effect on this group, in terms of annual tax collections, as

$$\Delta \text{ Tax Collected} = \text{Total Haven Wealth} * \text{Nominal Taxable Rate of Return} * \text{Average Effective Marginal Tax Rate} * \left(\text{Compliance Rate with FATCA} - \text{Compliance Rate without FATCA} \right)$$

This expression describes how our data on total wealth and information on the income generated by this wealth can begin to inform the total compliance effects. The most important unknown in this expression is the causal term in parentheses: how compliant were these owners prior to FATCA implementation, and how compliant are they afterwards? We do not attempt to answer these questions in this paper.

Our descriptive analysis does provide information about the first three terms in the expression, and future work should refine these further. With \$1.9 trillion dollars in total wealth in tax havens,²¹ a 5-7% nominal rate of return in havens (see Table 4 and Table A5), haven wealth would generate \$95 to \$133 billion in income. Though we know these assets are held by those at the

²¹ Including a compliance effect for income derived from *non-haven* wealth would increase the total effect, but calibrating a total effect for non-haven wealth introduces multiple complications: significant unmatched non-haven wealth may not even belong to U.S. taxpayers, and many taxpayers with wealth in non-havens likely remit taxes on the income from that wealth to the foreign country's government and claim Foreign Tax Credits in the U.S.

top of the income distribution, the average effective marginal tax rate on this income is uncertain due to differences in effective marginal tax rates by owner income (including differences induced by offsetting losses) and owner type (some owners are not taxable). With an effective tax rate of 25% and a 35 percentage point change in the compliance rate – roughly the change in mis-reporting rates suggested by moving financial income from offshore wealth from the “Income subject to little or no information reporting” category to the “Income subject to some information reporting” category in IRS (2019) – the total compliance effect on this population would be \$8 billion to \$12 billion annually. With an effective tax rate of 15% and a change in compliance rates of 5 percentage points, the total effect would be \$700 million to \$1 billion. The upper bound for the potential total effect is much higher: for example, with a 90% change in compliance and an effective tax rate of 25%, the implied total effect is \$21 billion to \$30 billion.

We take no stand on which of these illustrative calculations is most accurate. What these calculations nevertheless reveal is that, given the large amount of assets revealed to be held abroad, understanding the change in compliance rates among this group is vitally important for assessing the size of total effect. This might not have been the case *ex ante*: if little wealth remained abroad after FATCA, we could conclude that the total compliance effect for those maintaining accounts abroad is small regardless of the change in compliance for this sub-population. Our findings further highlight that an important component of the total compliance effect is the compliance effect on partnerships in particular. Partnerships own at least 50% of the \$1.9 trillion we observe in havens, so about half of the change in overall compliance rates is determined by partnership behavior.

The voluntary compliance effect on owners who maintained offshore accounts is not the sole component of the total revenue impact of FATCA. If some taxpayers stopped evading offshore and repatriated wealth back to the US in response to FATCA, as suggested by the evidence in De Simone et al. (2020), this would generate an additional revenue effect. Using data on assets that left tax havens around FATCA implementation, De Simone et al. (2020) suggest that repatriated assets could yield another \$1.6 to \$3.3 billion in revenues annually; they note that the effect would be smaller if those who moved financial assets out of havens did not actually repatriate them and start remitting tax on the associated investment income. Additionally, drawing on FATCA information may enable the IRS to collect more tax revenue through audits and collections.

4. Conclusion

Third-party reporting is a key enforcement technology with well-documented effects on tax compliance in many countries and contexts (Kleven et al., 2011; Pomeranz, 2015; Slemrod et al., 2017; Bagchi and Dušek, 2021). Recent policy initiatives in the U.S. and at the global level have attempted to extend third-party reporting to foreign financial institutions with the aim of addressing apparently widespread non-compliance on financial income earned through foreign accounts.

Based on account-level data on foreign assets reported by foreign financial institutions under the U.S. policy initiative, FATCA, we have assessed the overall magnitude of U.S. financial wealth

abroad and its distribution across income groups. Around 1.5 million U.S. taxpayers own foreign accounts with an aggregate value of around \$4 trillion in 2018. Half of these assets are held in jurisdictions usually considered tax havens, which implies a ratio of tax haven assets to GDP of around 10% in 2018, notably higher than a comparable estimate for 2007 based on macro statistics (Alstadsæter et al., 2018). Foreign asset ownership is highly concentrated at the top of the income distribution, with more than 60% of the individuals in the top 0.01% owning a foreign account and around 30% of all assets in foreign accounts belonging to this group. Most related studies based on less comprehensive data sources find even higher concentrations of offshore wealth in the top income and wealth groups (e.g. Alstadsæter et al., 2019; Londoño-Velez and Ávila-Mahecha, 2020, Guyton et al., 2021).

Given the magnitude of offshore wealth we observe under FATCA reporting, an important determinant of the overall effect of FATCA reporting will be the voluntary compliance effect on those owners who maintained their offshore accounts after FATCA implementation. Future research should attempt to quantify this causal effect.

References

- Alstadsæter, A., Johannesen, N., and Zucman, G., 2018. Who owns the wealth in tax havens? Macro evidence and implications for global inequality. *Journal of Public Economics*, 162, pp.89-100.
- Alstadsæter, A., Johannesen, N., and Zucman, G., 2019. Tax evasion and inequality. *American Economic Review*, 109(6), pp.2073-2103.
- Bagchi, S., and Dušek, L., 2021. The effects of introducing withholding and third-party reporting on tax collections: Evidence from the US state personal income tax. *Journal of Public Economics*, 204, 104537.
- Belnap, A., Thornock, J., and Williams, B., 2019. The Long Arm of the U.S. Tax Law: Participation Rates and Costs Related to Mandated Information Sharing. Working Paper.
- Bomare, J., and Herry, S. L. G., 2022. Will We Ever Be Able to Track Offshore Wealth? Evidence from the Offshore Real Estate Market in the UK. EU Tax Observatory Working Paper 4.
- Brocard, M., and Lhabitant, F.-S. 2016. A primer on the tax framework of offshore and onshore hedge funds. EDHEC Business School Working Paper.
- Byrnes, W., and Munro, R., 2017. Background and Current Status of FATCA, Texas U&A University School of Law, Legal Studies Research Paper No. 17–31.
- Cooper, M., McClelland, J., Pearce, J., Prisinzano, R., Sullivan, J., Yagan, D., Zidar, O., and Zwick, E., 2016. Business in the United States: Who Owns It, and How Much Tax Do They Pay? *Tax Policy and the Economy*, 30(1), pp.91-128.
- Casi, E., Spengel, C., and Stage, B.M., 2020. Cross-border tax evasion after the common reporting standard: Game over? *Journal of Public Economics*, 190, 104240.
- De Simone, L., Lester, R., and Markle, K., 2020. Transparency and tax evasion: Evidence from the foreign account tax compliance act (FATCA). *Journal of Accounting Research*, 58(1), pp.105-153.
- Dharmapala, D., 2016. Cross-border tax evasion under a unilateral FATCA regime. *Journal of Public Economics*, 141, pp.29-37.
- Federal Reserve Bank of St Louis, 2022. "Households and Nonprofit Organizations; Total Financial Assets, Level [TFAABSHNO]." Federal Reserve Economic Data (FRED) repository, accessed December 22, 2022. <https://fred.stlouisfed.org/series/TFAABSHNO>
- Guyton, J., Langetieg, P., Reck, D., Risch, M., and Zucman, G., 2021. "Tax Evasion at the Top of the Income Distribution: Theory and Evidence." NBER Working Paper No. 28542.

Internal Revenue Service (IRS), 2017. "Revised Guidance Related to Obtaining and Reporting Taxpayer Identification Numbers and Dates of Birth by Financial Institutions." IRS Notice 2017-46. <https://www.irs.gov/pub/irs-drop/n-17-46.pdf>.

IRS Statistics of Income (SOI)., 2012. "Short-term and Long-term Capital Gains and Losses by Asset type: 2012." *Sales of Capital Assets Reported on Individual Tax Returns*. Data repository, accessed 3 Nov 2020. <https://www.irs.gov/pub/irs-soi/1201insoca.xls>

Johannesen, N., 2014. Tax evasion and Swiss bank deposits. *Journal of Public Economics*, 111, pp.46-62.

Johannesen, N., and Zucman, G., 2014. The end of bank secrecy? An evaluation of the G20 tax haven crackdown. *American Economic Journal: Economic Policy*, 6(1), pp.65-91.

Johannesen, N., Langetieg, P., Reck, D., Risch, M., and Slemrod, J., 2020. Taxing hidden wealth: The consequences of US enforcement initiatives on evasive foreign accounts. *American Economic Journal: Economic Policy*, 12(3), pp.312-346.

Johannesen, N., and Stolper, T. B., 2021. The deterrence effect of whistleblowing. *Journal of Law and Economics*, 64(4), pp.821-855.

Kleven, H.J., Knudsen, M.B., Kreiner, C.T., Pedersen, S., and Saez, E., 2011. Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark. *Econometrica*, 79(3), pp.651-692.

Lane, P.R., and Milesi-Ferretti, G.M., 2007. The external wealth of nations mark II: Revised and extended estimates of foreign assets and liabilities, 1970–2004. *Journal of International Economics*, 73(2), pp.223-250.

Lane, P. R., and Milesi-Ferretti, G. M., 2001. The external wealth of nations: measures of foreign assets and liabilities for industrial and developing countries. *Journal of International Economics*, 55(2), pp.263-294.

Leenders, W., Lejour, A., Rabaté, S., and van't Riet, M., 2020. Offshore Tax Evasion and Wealth Inequality: Evidence from a Tax Amnesty in the Netherlands. CPB Netherlands Bureau for Economic Policy Analysis.

Londoño-Velez, J., and Ávila-Mahecha, J., 2020. Enforcing Wealth Taxes in the Developing World: Quasi-Experimental Evidence from Colombia. Forthcoming in *American Economic Review: Insights*.

Love, M., 2021. Where in the World Does Partnership Income Go? Evidence of a Growing Use of Tax Havens. Evidence of a Growing Use of Tax Havens. Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3985535.

Menkhoff, L., and Miethe, J., 2019. Tax evasion in new disguise? Examining tax havens' international bank deposits. *Journal of Public Economics*, 176, pp.53-78.

Omartian, J.D., 2017. Do banks aid and abet asset concealment: Evidence from the Panama Papers. University of Michigan Working Paper.

O'Reilly, P., Ramirez, K.P., and Stemmer, M.A., 2019. Exchange of information and bank deposits in international financial centres. OECD Taxation Working Papers No. 46.

Piketty, T., Saez, E., and Zucman, G., 2018. Distributional national accounts: methods and estimates for the United States. *Quarterly Journal of Economics*, 133(2), pp.553-609.

Pomeranz, D., 2015. No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, 105(8), pp.2539-2569.

Saez, E., and Zucman, G., 2016. Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *Quarterly Journal of Economics*, 131(2), pp.519-578.

Slemrod, J., Collins, B., Hoopes, J. L., Reck, D., and Sebastiani, M., 2017. Does credit-card information reporting improve small-business tax compliance? *Journal of Public Economics*, 149, pp.1-19.

Treasury Inspector General for Tax Administration (TIGTA), 2018. "Despite Spending Nearly \$380 Million, the Internal Revenue Service Is Still Not Prepared to Enforce Compliance With the Foreign Account Tax Compliance Act." TIGTA report 2018-30-040.

<https://www.treasury.gov/tigta/auditreports/2018reports/201830040fr.pdf>.

Treasury Inspector General for Tax Administration (TIGTA), 2022. "Additional Actions Are Needed to Address Non-Filing and Non-Reporting Compliance Under the Foreign Account Tax Compliance Act." TIGTA report 2022-30-019.

Zucman, G., 2013, The missing wealth of nations: Are Europe and the US net debtors or net creditors? *Quarterly Journal of Economics*, 128(3), pp.1321-1364.

APPENDIX A: Additional Results

Table A1. Aggregate Income by Owner Type

		Interest		Dividends		Gross Proceeds/Redemptions		Other income	
		Total (Billions USD)	Share	Total (Billions USD)	Share	Total (Billions USD)	Share	Total (billions USD)	Share
Matched	Partnership	3.03	23.0%	13.23	46.4%	90.83	33.0%	132.22	63.5%
	Individual	2.80	21.2%	6.83	23.9%	49.51	18.0%	26.02	12.5%
	C Corporation	1.11	8.4%	0.80	2.8%	8.78	3.2%	11.10	5.3%
	Tax exempt entity	0.30	2.3%	0.29	1.0%	2.76	1.0%	4.38	2.1%
	Trust	0.14	1.1%	1.00	3.5%	3.36	1.2%	6.74	3.2%
	Foreign corporation	0.11	0.8%	0.03	0.1%	2.64	0.9%	0.55	0.2%
	S corporation	0.03	0.2%	0.10	0.3%	0.36	0.1%	1.03	0.4%
Unmatched	Missing TIN	3.27	24.8%	3.52	12.3%	91.66	33.3%	13.31	6.4%
	Unmatched entity	0.75	5.6%	0.39	1.3%	11.70	4.2%	4.06	1.9%
	Ambiguous match	1.07	8.1%	1.71	6.0%	6.72	2.4%	2.29	1.1%
	Unmatched TIN	0.48	3.6%	0.51	1.8%	5.29	1.9%	6.16	2.9%
	Unmatched individual	0.05	0.4%	0.05	0.1%	0.82	0.3%	0.11	<0.1%

Note: This Table reports owner shares weighted by income. We observe that most income types are similarly or slightly more concentrated in partnerships, compared to the wealth shares in Table 3. See Table 3 of the main text and the surrounding text for details on the construction of the estimates and group definitions.

Table A2 – Tax Year 2018 Totals: Haven versus Non-Haven Countries

	All	Havens	Non-Havens
No. of accounts	4,566,774	612,406	3,954,368
Total Account Balance (Billions USD)	3,981.8	1,939.8	2,042.0
Total Interest (Billions USD)	13.2	4.7	8.5
Total Dividends (Billions USD)	28.5	18.0	10.5
Total Gross Proceeds/Redemptions (Billions USD)	274.5	132.7	141.8
Total Other Income (billions USD)	208.0	173.7	34.3

Note: This table decomposes the total number of accounts and financial aggregates into totals for haven and non-haven countries. We observe that 14% of accounts and 49% of wealth is located in haven countries.

Table A3 – Total Values by Country Type – TY 2018

	Country type			
	Haven	Non haven	Non-haven – prior EOI	Non-haven – no prior EOI
No. of accounts	612,406	3,954,216	3,222,800	731,416
Total Account Balance (Trillions USD)	1.9	2.0	1.7	.2
Total Interest (Billions USD)	4.6	8.5	5.9	2.5
Total Dividends (Billions USD)	17.9	10.5	9.4	1.1
Total Gross Proceeds/Redemptions (Billions USD)	132.6	141.8	107.0	34.7
Total Other Income (Billions USD)	173.7	34.2	14.6	19.5

Note: This table reports the same statistics as Table A2 and further decomposes non-haven countries based on whether that country's government agreed to exchange information with the US prior to FATCA, under TIEAs and/or DTCs, which we label Prior Exchange of Information (EOI). The full set of country groups is as follows: Haven countries: Anguilla, Antigua and Barbuda, Aruba, Bahamas, Bahrain, Belize, Bermuda, Caribbean Netherlands, Cayman Islands, Cook Islands, Curacao, Cyprus, Dominica, Grenada, Guernsey, Hong Kong, Isle of Man, Jersey, Liechtenstein, Luxembourg, Malta, Mauritius, Monaco, Panama, Saint Kitts and Nevis, Saint Lucia, Singapore, Sint Maarten, Switzerland, Turks and Caicos Islands, United States, and British Virgin Islands. Non haven and prior EOI countries: Australia, Belgium, Canada, China, Czech Republic, Denmark, Estonia, Finland, France, Germany, Hungary, Iceland, India, Indonesia, Italy, Japan, Latvia, Lithuania, Mexico, New Zealand, Norway, Poland, Portugal, Slovakia, Slovenia, Republic of Korea, Spain, Sweden, Trinidad and Tobago, and the United Kingdom. All other countries are non-haven, without prior EOI.

Table A4 – Joint Distribution of Owner Type and Country – TY 2018

Panel A. Number of Accounts by Owner Type and Country

		Country type				Row Total
		Haven	Non Haven	NH – EOI	NH – No EOI	
Owner type	Individuals	403,909	2,262,500	1,705,589	556,911	2,666,409
	Partnership	42,592	19,669	15,024	4,645	62,261
	Other entity	22,241	31,548	24,609	6,939	53,789
	Unknown	143,383	1,639,170	1,477,358	161,812	1,782,553
Column Total		612,125	3,952,887	3,222,580	7,303,07	4,565,012

Panel B. Total Account Balance by Owner Type and Country (Billions USD)

		Country type				Row Total
		Haven	Non Haven	NH – EOI	NH – No EOI	
Owner type	Individuals	412	212	145	67	625
	Partnership	1,001	287	228	58	1,288
	Other entity	291	255	210	44	546
	Unknown	231	1,265	1,196	68	1,497
Column Total		1,937	2,020	1,781	239	3,957

Note: This table decomposes the total number of accounts (Panel A) and total wealth (Panel B) according to the country groups from Table A3 and the owner type categories from Figure 1. To facilitate the computation of shares, we also include row and column totals in each panel.

Table A5. Rates of Return in Selected Sub-Populations: Additional Specifications

	Sample restriction						
	All	Non-haven country	Haven country	Individual owners	Partnership owners	Other entity owners	Unmatched owners
Total reported wealth (billions USD)	3,982	2,042	1,940	626	1,292	279	1,510
Share of wealth with reported interest	33.3%	44.0%	22.0%	27.3%	16.9%	36.9%	43.9%
Quasi-rate of return: interest only	1.0%	0.9%	1.1%	1.7%	1.4%	0.7%	0.8%
Share of wealth with reported interest or dividends	37.7%	51.3%	23.4%	33.6%	19.8%	37.7%	49.9%
Quasi-rate of return: interest + dividends only	2.8%	1.8%	5.0%	4.6%	6.4%	1.1%	1.6%
Share of wealth with reported int., div., or gross proceeds	44.1%	55.0%	32.5%	36.5%	33.3%	38.7%	53.0%
Quasi-rate of return: int. + div. + 0.14*gross proceeds	4.6%	3.5%	6.5%	7.4%	6.7%	2.6%	3.5%
Share of wealth with reported int., div., G.P., or other income	59.2%	57.2%	61.5%	45.5%	67.2%	44.5%	56.3%
Quasi-rate of return: int. + div. + 0.14*(GP + other)	4.6%	3.7%	5.5%	7.2%	5.5%	2.7%	3.7%

Note: This table reports alternative specifications to Table 4 in the main text. The second set of estimates, where the quasi-rate of return includes interest and dividends only, is identical to Table 4. The other rows use alternative definitions of the rate of return. We observe that the higher rate of return for partnerships and individuals, and the higher return in havens than non-havens, both obtain for essentially any definition of the quasi-rate of return.

Table A6 – Partnership account owners, assets and income by industry

	N account owners	Account Balance (Billions)	Interest (Billions)	Dividends (Billions)	Gross Proceeds (Billions)	Other (Billions)
Total						
Financial Services	15,754	910.00	1.90	7.58	58.40	102.50
Other finance	4,255	247.80	0.50	2.63	14.64	12.34
Other industries	1,647	13.27	0.04	0.09	0.50	1.30
Unknown	5,045	130.20	0.64	2.93	17.40	16.61
Non-haven						
Financial Services	4,081	204.00	1.04	0.91	22.40	8.00
Other finance	1,480	54.80	0.30	0.70	11.60	1.74
Other industries	1,155	5.87	0.01	0.02	0.33	0.14
Unknown	2,088	27.20	0.43	0.11	4.70	3.31
Haven						
Financial Services	11,673	706.00	0.86	6.67	36.00	94.50
Other finance	2,775	193.00	0.20	1.93	3.04	10.60
Other industries	492	7.40	0.03	0.08	0.17	1.16
Unknown	2,957	103.00	0.21	2.82	12.70	13.30

Note: This table shows asset and income information for partnerships for year 2018 reported on FATCA Form 8966 by industry and whether the account was held in a haven or non-haven country. It is categorized by the largest industry groups in each country category. The table shows the total number of 2018 partnerships reported on Form 8966 for each industry, total assets held and total income by source.

Table A7 – Allocation of Partnership Foreign Assets to Taxable Owners

	share of owners raw	share of owners assigned	share of assets raw	share of assets assigned
Individual	88.6%	90.1%	43%	54%
Unclassifiable	1.7%	0.0%	20%	0%
Foreign				
Individual/Entity	0.1%	0.1%	10%	13%
Tax Exempt	0.9%	0.9%	9%	11%
Trust	8.0%	8.2%	9%	11%
Foreign Corporation	0.2%	0.2%	5%	7%
Corporation	0.6%	0.6%	4%	4%

Note: Table A7 reports the shares of pass-through owners and allocated assets by owner type, the values underlying Figure 3.

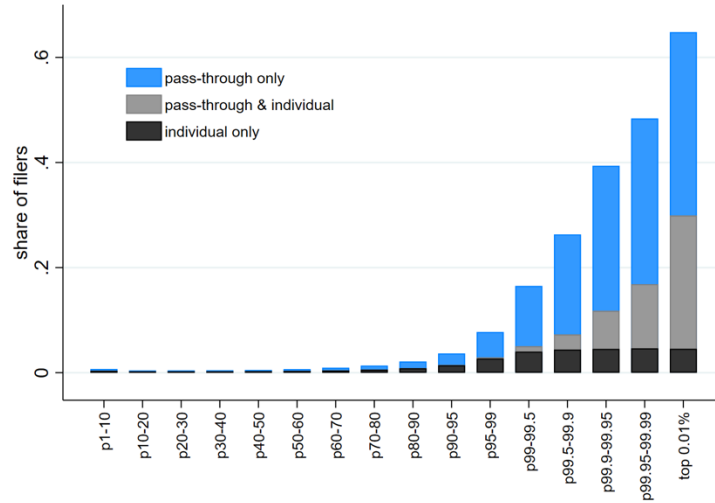
Table A8 – Allocated foreign assets by position in income distribution (2018)

Position in AGI Dist.	Haven Individual (Billion USD)	Haven Pass-through (Billion USD)	Non-haven Individual (Billion USD)	Non-haven Pass-through (Billion USD)	Total (Billion USD)
losses	12.4	8.8	5.9	2.8	29.9
p2-p10	41.2	2.7	19.5	2.5	65.9
p10-p20	5.3	2.8	8.3	0.4	16.8
p20-p30	11.6	6.6	3.7	4.5	26.5
p30-p40	9.1	7.6	3.4	0.9	21
p40-p50	2.5	4.7	4.3	2.2	13.7
p50-p60	17.1	2.9	10.6	0.8	31.5
p60-p70	2.8	8.3	6.8	1.3	19.1
p70-p80	5.5	5.6	9.1	1.8	21.9
p80-p90	7.4	10.9	16.9	6.5	41.7
p90-p95	31	5.6	12.1	6.9	55.5
p95-p99	19.4	13	24.2	4.7	61.3
p99-					
p99.5	12.8	9.2	11.5	27.4	60.8
p99.5-					
p99.9	71.3	37.3	19	11.8	139.4
p99.9-					
p99.99	67.7	77.7	17.5	22.2	185.1
p99.99-					
p100	66.5	192	12.1	66	336.7


Note: Table A8 shows the asset values underlying-Figure 5.

Figure A1: Share of taxpayers with a foreign account reported by FATCA by percentile – Distribution of Total Positive Income (TPI), TY2018

Panel A: Share of taxpayers holding accounts directly and/or through pass-throughs



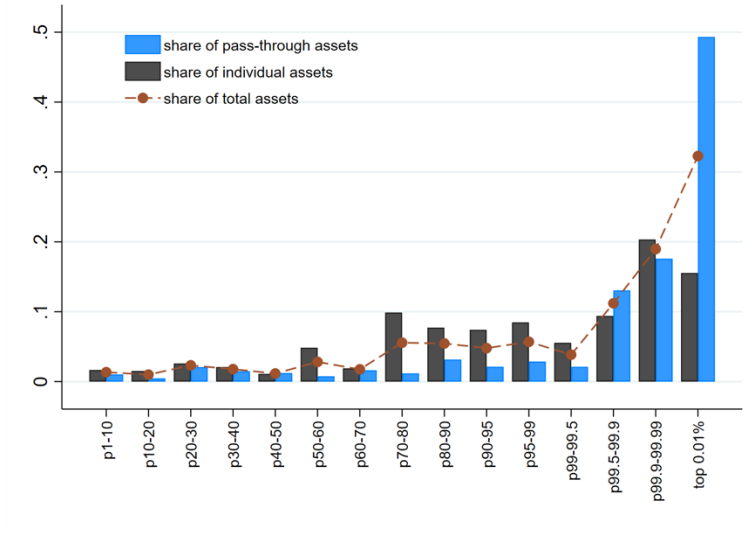
Panel B: By Haven or non-haven




Note: This figure replicates Figure 4, but shows the share of taxpayers in each quantile of the distribution of total positive income (TPI) that receives a FATCA report (Form 8966) indicating ownership of a foreign account. The share is calculated as the total number of individuals in a quantile listed as an account owner on an 8966 over the total number of tax returns in that quantile. The black bars (individual) represent the share of individual account owners that own a foreign account directly. The blue bars (partnership K-1) represent the share of individual shareholders of partnerships that own foreign accounts. Explicitly, for partnership foreign account owners, we link the shareholders to the partnership through form K-1 and the blue bars represent individual shareholders of these partnerships. Panel A shows foreign account ownership shares by centile of the distribution. Over 20% of those in the top 1% are foreign account owners, with about 14% holding accounts through pass-through entities. Panel B zooms in on the top 1% of the income distribution, presenting the shares holding foreign accounts by the top 0.9-0.1% separately, then decomposes the top 0.1% into the 0.09-0.05%, the 0.05-0.01% and the top 0.01%. Foreign account ownership rates increase throughout the top 1%, with approximately 66% of those in the top 0.01% being FATCA account owners, just over half of which hold accounts through pass-through entities.

Figure A2: Share of assets owned by position in the individual income distribution – Distribution of Total Positive Income (TPI), TY2018


Panel A: Distribution of assets held directly and held through pass-through entities



Panel B: Distribution of assets held in havens and held in non-havens




Panel C: Distribution of total assets held directly and through pass-through entities, in havens and non-havens




Note: This figure replicates Figure 5 but shows the share of assets held across the distribution of total positive income (TPI). Panel A looks at the distribution of all assets held in directly by individuals and the distribution of those held through pass-through entities (partnerships and S-corporations), separately. The red series plots total shares across both types. Shares are defined as the assets held by individual (pass-through) account owners in each percentile over the total amount of assets held by individual (pass-through) account owners. Total assets (red series) are very concentrated, with about 28% being held by the top 0.01%. The distribution of assets held through pass-through entities (blue bars), which are extremely concentrated, with over 45% being held by the top 0.01%. Directly held assets (black bars) are also concentrated, but less so, with about 12% being held by the top 0.01%. Panel B looks at the distribution of all assets held in havens and the distribution of those held in non-havens, separately. About 33% of assets in havens are held by those in the top 0.01%. Panel C decomposes the distribution of total assets into amounts held in havens and non-havens, directly and indirectly – 28% of total assets are held by the top 0.01%, the vast majority of which are held in havens (grey and black bars), particularly in through pass-through entities (black). Of non-haven assets held at the very top (light and dark blue bars), the majority is held through pass-through entities (dark blue).

Figure A3 – Comparison of Overall Wealth Shares Ranking by Total Positive Income (TPI) and Adjusted Gross Income (AGI)



Note: This Figure plots the share of total wealth belonging to individuals at different parts of the income distribution when we rank by AGI (as in Figure 5 of the main text) versus when we rank by Total Positive Income (as in Figure A2). We observe that ranking by TPI increases the concentration and especially decreases the ownership share in the bottom 50% of the distribution.

Figure A4 – Robustness of Asset Distribution to FATCA Exemption Threshold



Note: Foreign Financial Institutions (FFIs) are only required to report on account holders with aggregate asset values over \$50,000. Figure A4 shows an ad hoc adjustment to the distribution of foreign assets to assess the potential importance of this threshold. We arbitrarily assign \$40,000 in foreign assets to 10% of all households in the bottom 90% of the AGI distribution in 2018 (i.e. assume 10% of households in the bottom 90% have foreign accounts that do not appear in our data because they are slightly below the FATCA reporting threshold). The distribution of foreign assets across the AGI distribution is re-estimated. The level and share of assets going to the bottom 90% substantially increases, but the profile remains qualitatively similar; foreign assets would remain highly concentrated at the very top of the income distribution.

Appendix B: Cleaning Procedures and Statistics on Reporting Quality for Form 8966 Data

B1. The Form 8966

Form 8966 is submitted to the IRS by either foreign governments or by foreign financial institutions (FFIs) directly, depending on the model of the Intergovernmental Agreement (IGA) that the country has with the United States. See the main text for further details on IGAs. The contents of the form and requirements for what is reportable are generally the same under both IGA models, except that FFIs in Model 1 countries (about 88% of participating countries) were not required to report owner TINs for accounts existing prior to FATCA in the first years of reporting, as discussed in the main text.

The Form 8966 comprises the following five parts:

- **Part I, “Identification of Filer,”** contains information on the FFI or other entity (such as a Sponsoring Entity) that is reporting a U.S. owned account or ownership interest on the form. In addition to the name and address of the filer, this section contains a 2-digit filer category code describing what type of entity is filing the form, and the filing entity’s Global Intermediary Identification Number (GIIN).
- **Part II, “Account Holder or Payee Information,”** contains identifying information about the account holder, including whether the holder is an individual or an entity, what type of entity the holder is if it is an entity, the location of the account holder (address, city, state, and zip code), and the owner TIN. For the majority of 8966s, this section contains the identifying information of the account owner, who is also the account holder. Researchers worked with de-identified records throughout the research process, e.g. a dataset containing masked owner TINs and excluding the name of account holders.
- **Part III, “Owner Information,”** contains identifying information for the account owner when the account owner is different from the account holder. Generally, this part of the form will be populated in relatively infrequent situations, when a U.S. person (the “account owner”) owns an account indirectly through certain kinds of foreign entities (the “account holder”). Specifically, filling out this section is required for substantial U.S. owners of a passive non-financial foreign entity (NFFE), substantial U.S. owners of a Direct Reporting NFFE, substantial U.S. owners of a Sponsored Direct Reporting NFFE, and specified U.S. persons owning certain equity or debt interests in an owner-documented foreign financial institution (ODFFI).
- **Part IV, “Financial Information,”** contains information on the account balance and income for the reportable account in a given tax year. In addition to the account number (masked in the research dataset), currency code, and an indicator for whether the account was closed during that year, the financial account information comprises 5 items: account balance, interest, dividends, gross proceeds/redemptions, and other income.
- **Part V, “Pooled Reporting Type,”** pertains to instances where account information is pooled for reporting purposes. This is used for reporting on recalcitrant accounts (where the account holder is not compliant in providing identifying information to the FFIs). Part V is blank when a given Form 8966 pertains to a specific account owned by a US person, i.e. the accounts in our main analysis sample. When a Form 8966 contains a pooled report with information in Part V, Parts II through IV are blank.

After the introduction of Form 8966 in TY2014, a few changes to the form were made in TY2015. The following table describes these changes, with line number references to the TY2015 version of the Form 8966:

Form Section	TY2014	TY2015-present
Part I, Line 1b: Filer Category Code	Absent	Present
Part II, Line 1b: Indicate whether account holder is an individual or entity	Absent	Present
Part II, Line 5: entity type	Includes Direct Reporting NFFE	Excludes Direct Reporting NFFE
Part IV, Line 3b: Check if account closed during the year	Absent	Present

Following Tax Year 2015, the contents of the form have not changed. We discard TY2014 data because very few accounts were reported for TY2014, so changes to the form over time are irrelevant for our analysis.

B2. Data Cleaning Procedures

The following describes the data cleaning process for the dataset of contents of Forms 8966 used in our analysis. The data were extracted from the IRS database and contains all fields of the Form 8966, as well variables generated when processing the forms. We note that all data cleaning procedures outlined here are done purely for statistical purposes for this paper and do not represent how the IRS processes Form 8966 records internally.

Owner Information

We first create a new set of variables for the beneficial owner of the account (i.e., TIN, owner’s country code and country name). Based on how FFIs are instructed to fill in the form (see above), we pull owner information from Part III of the form when Part III is populated, and otherwise from Part II. In some cases, the beneficial owner TIN reported in Part II or III is missing or is in an invalid format (e.g. a single-digit number); we create a variable that encodes a missing or invalid TIN.

The type of owner for an account, either “Individuals” or “Entity,” is reported in Part II, Line 1b. If “Entity” is selected, Part II, Line 5 is filled out to indicate the type of entity, “US Person Entity,” “Passive NFFE with US Owners,” “Direct Reporting NFFE,” “FFI with US Owners,” and “Nonparticipating FFI.” We combined these five entity categories, along with the “Individual,” to create a categorical variable for owner type. In characterizing owner types in the main text, we generally defer to the owner type variable created during the form match process (see below). Information about the type of owner as reported on the Form 8966 is nevertheless useful for cleaning data and e.g. in disambiguating form matches.

Filer Information

The main variable of interest to us from the Filer information in Part I of the form is the country in which the account is located, but we use other RFI information in data cleaning to check for duplicate records and jointly owned accounts, and to check the validity of very large dollar values (see below). Since most filer information is split between RFI variables and Sponsoring Entity (SPS) variables, these variables are combined to create a single filer variable for all identifying information. Box 1b in Part I of the form contains a 2-digit filer category code with eleven categories. A categorical variable is created by combining

the RFI and SPS filer category codes. Similarly, RFI and SPS GIINs are combined into one variable representing the GIIN of the financial institution where the assets are located.

Currency Conversions

The financial information in Part IV of the form can be reported in a local currency or in U.S. Dollars (USD). All such fields have a corresponding variable containing the three-letter currency code for the currency in which the field is denominated. The raw data contained versions of the account value variables that had already been converted to USD by a team at the IRS. However, to examine the validity of extremely large dollar values and clean data it was useful to convert the raw, local-currency version of these variables to USD ourselves using the same set of exchange rates used by the internal team. The end result is three sets of variables for Part IV of the form: the IRS converted amount, our converted amount, and the amount as reported by the FFI or country.

We generally defer to the IRS converted amount, with two exceptions. First, in some cases, the variable for an unconverted amount is present while the internally converted USD amount is missing. In such cases, we replace the missing internally converted amount with our converted amount. Second, we use originally reported currency amounts to screen for outliers and data issues related to mis-reported currency codes by the filer, for instance when a filer reports an amount in local currency but mistakenly indicates the amount is in USD. When one USD is equivalent to an extremely large nominal amount of local currency, such misreporting can create egregious outliers, as discussed further below.

Sample Restrictions

Our sample restrictions focus on information reported in Part IV (Financial Information) because the key statistics about offshore wealth and related income are drawn from this part of the form. In order to avoid deleting observations that might be useful in other contexts, we create a categorical variable to indicate inclusion in the main analysis sample or, if the observation is excluded from the main analysis sample, to encode the reason why.

We first exclude records that are not useful for our analysis. Specifically, these exclusions are as follows:

- All records related to recalcitrant accounts are removed from the preferred sample (414,450 observations), as these are usually pooled for reporting purposes by FFIs.
- A “Record Status Code” is generated automatically upon submission of each form by filers. Observations with a record data status code of “Bad Data-Void,” “Bad Data-Void Error, or “Bad Data- Record Void Requested” are removed from the sample (2,934,393 observations). Such records contain little to no actual data.
- Observations where the Form 8966 owner type is missing, or the owner type is a Nonparticipating FFI (which also suggests that what is being reported involves recalcitrant accounts), are removed (256,102 observations)
- Finally, we exclude from our sample a large number of observations (6,697,618) in which no financial information reported for any income field or account balance.

The last of these restrictions excludes by far the most observations. We elected to drop these records because our goal in this paper is to characterize the offshore wealth and associated income that is reported by FFIs under FATCA, and in these records no wealth or income is reported. Based on the

instructions for the Form 8966, it is unclear why this information is so frequently missing. The most informative clue we observed is that, in 95% of the records with no financial information, the RFI fills in the account balance field as the number “0” rather than not entering any information at all – meanwhile, the income fields typically contain no information rather than the number 0. These clues suggest that the vast majority of accounts with no information are dormant accounts containing no assets. We consulted with experts on the data who generally agreed with this assessment. Still, some of these accounts with no financial information could be of interest from a tax enforcement or other perspective. We defer a fuller examination of these accounts to other work.

There are also a number of records in which account balance is reported but many or all of the other financial variables are missing or zero. We keep these records in the data and discuss missing income information further below and in the main text.

Following the initial restrictions above, we carefully clean the financial information variables to ensure that the key summary statistics we report in the paper based on these variables are accurate. Our general principle is to assume that FFIs report correctly unless we see a strong reason to believe otherwise. Still, there are a number of highly implausible records that are clearly due to some kind of reporting error. Including these records in the sample would bias our key statistics based on these financial variables.

After these restrictions, the data contain 23,738,383 observations, which break down by tax year as follows:

Tax Year	2014	2015	2016	2017	2018	2019
N. Obs.	1,210,434	2,559,273	4,595,466	5,058,423	5,351,336	4,963,451

Cleaning Financial Variables

One of our goals is to provide the most accurate account possible of the wealth and income reported on Forms 8966. To do this, we carefully screened the financial variables in Part IV of the form, as outliers due to mis-reporting can substantially skew distributional statistics of the kind we seek to calculate.

To begin with, accounts where the account balance is greater than \$100 billion, or any of the income fields are greater than \$10 billion, are removed (399 observations). Inspecting each of these records by hand, we found that all of them appear to be erroneous.

We then excluded observations with suspiciously small income amounts relative to their account balance. Observations where the account balance is greater than \$1 billion, at least one income variable is not missing, and the sum of the income variables is less than \$10 million are removed (386 observations removed).

The above restrictions excluded the most obviously problematic records for tabulating total account balances, but we found that some erroneous outliers remained. We next carefully examined total account balances *by country and year* to identify suspicious observations that warranted further examination. We found that this was a useful way to screen for problematic reporting errors because such errors usually generated inconsistencies within countries across tax years (e.g., a huge spike in a single year in a given

country). Upon finding such an inconsistency, we isolated the records responsible for the inconsistency by breaking down the totals further. In some cases, the anomaly was driven by a very small number of accounts in a given country (e.g., a single account over \$10 billion at an FFI that managed no other accounts over \$1 billion in any other year). In other cases, a particular FFI would appear to report a larger number of accounts inconsistently across time. For example, an FFI might report a few hundred accounts in 2015, 2017 and 2018 but several thousand accounts in 2016. From conversations with experts, one reason we believe we might observe this is if the FFI reports on *all* of the accounts it maintains rather than the accounts owned by U.S. persons. In other cases, the anomaly would be driven by single FFI whose account values spiked massively in a single year – e.g. from a few million to several billion dollars – for a large number of accounts, for unknown reasons. In both of these cases, we typically excluded all accounts or accounts over \$10 million at the problematic FFI. In the latter case, we also checked whether the anomaly could be attributable to a misreported currency code, in which case we would correct the currency code and keep the records in the analysis sample. In yet another case, the anomaly occurred because account balances contained repeated digits (e.g., a string of 9’s) that indicate that they are place holders rather than actual account balances; we excluded such records.

In total, we removed 78,560 observations due to these issues, leaving 23,659,823 observations in the sample, which break down by year as follows:

Tax Year	2014	2015	2016	2017	2018	2019
N. Obs.	1,191,647	2,551,608	4,579,049	5,052,811	5,333,362	4,951,346

Of these, 20,805 were exclusions of a very small number of accounts and 68,528 were broader exclusions implemented at the FFI level. We found obvious currency conversion issues causing anomalies in 33,866 records and corrected them. We are unable to provide further details due to privacy concerns and the data protection provisions of IGA agreements. We note that one implication of these exclusions is that the total wealth figures likely represent a lower bound for the actual total offshore wealth and income of US persons. If some problematic records are still included, this could bias the totals upwards rather than downwards, but we believe that our careful screening ensures any upward bias from this issue is likely minimal.

We next turned to cleaning the financial income variables: interest, dividends, gross proceeds/redemptions, and other income. Observations with over \$10 million in interest income and the interest amount is greater than half of the account balance, or have \$10 million in interest income and the account balance is missing, are dropped from the preferred sample (1,402 observations removed). All observations with over \$1 billion in dividend income are suspicious, as all of these observations have low or no account balance (50 observations removed). For many observations where Gross Proceeds makes up a considerable amount of the total account balance, it appears that the account is new and therefore not suspicious. Accordingly, the restriction for gross proceed amounts is 1.5 times the total account balance. If gross proceed income is greater than \$10 million, and gross proceeds are 1.5 time the account balance, the observations is removed (9,044 observations removed). Accounts with over \$1 billion in gross proceeds and missing account balance are removed (484 observations removed). The restriction of 1.5 time the account balance is also used with other income to discard suspicious observations. If other income is \$1 billion and 1.5 times greater than the account balance, the observation is removed (126

observations removed). Additionally, if other income is greater than \$1 billion and account balance is missing, the observation is removed (52 observations removed).

Following our use of the decomposition of account values by country and year, we followed the same cleaning procedures for the other financial information fields. In other words, we broke down the totals of these fields by country and year, and reviewed this decomposition to identify further anomalies. From this procedure we removed another 17,771 observations, leaving 23,642,052 observations in the sample, which break down by year as follows:

Tax Year	2014	2015	2016	2017	2018	2019
N. Obs.	1,190,809	2,549,525	4,574,280	5,048,854	5,329,536	4,949,048

Going forward, we disregard data from 2014 and 2019. Tax year 2014 contain far fewer records, as we can see in the tabulations above. Tax year 2019 data were incomplete as of the data pull upon which we built our analysis sample, and this is reflected by a drop in the number of observations from 2018 to 2019.

Duplicate Records

The raw data also contain a sizable number of duplicate records, which we turn to next.

We first create a data quality code variable to rank observations by data quality for use in cleaning duplicates. The highest level of data quality are observations that have a “Good Data” Record Data Status code, a positive total amount of income, and a positive account balance (6,717,526 observations). The second highest data quality rank are observations that have a “Good Data” Record Data Status code, but have either a missing, 0, or negative account balance or total income, with the last of these being the most common by far (13,002,262 observations). The next rank are observations that have a “Bad Data-Correction Requested” Record Data Status Code (10,133,629 observations). The records that have this data status code and survive our other cleaning procedures typically contain useful financial information; our understanding is that the requested corrections alluded to in the code pertain to missing owner information (e.g. owner TINs) or parts of the Form 8966 that we do not use in our analysis. The remainder (3,198,766 observations) are ranked as the lowest level of data quality.

To check for duplicate records, we first counted the number of observations per each unique tax year/RFI GIIN/account number combination. We find that 67% of observations are unique to each account number/RFI GIIN/tax year combination, while 33% have more than one observation. We created a “non-duplicate flag” at this point in the data, and assume that the 67% of records with a unique account number, RFI, and tax year are not duplicates. For the remaining 33%, we need to do more work because not all of these records are duplicates. In particular we could observe the same account number, RFI GIIN and tax year for multiple records in two common situations: jointly owned accounts, and accounts where the (masked) account number we observe represents not an actual account number but a placeholder value (e.g. for some types of reportable accounts an account number may not exist or be known to the filer). The next several steps represent our best attempt at separating out the true duplicate records from these other types of situations.

First, we assume if we observe multiple records with the same account number at the same RFI in the same tax year, *but the owner TINs are distinct*, then these do not represent duplicate records. Rather they

are likely to be distinct accounts with a placeholder/missing account number or jointly owned accounts. We handle jointly owned accounts in the next section. Based on this assumption, when we see the same RFI, account number, and *non-missing* owner TIN for multiple records in the same tax year, we keep exactly one of these records. We use the data quality code described above to decide which of the duplicate records to keep – for example, if the Record Data Status Code was “Bad Data, Correction Requested” and the FFI submitted a corrected record, this will ensure that we keep the corrected record. Empirically this happens quite often. In the case where multiple duplicates have the same data quality code, we keep the first record appearing in the dataset.

When owner TINs are not missing, adding the distinct TIN screen is likely all we need to screen out duplicate records. The observations without missing TINs that this procedure flags as duplicates are therefore removed from the sample (2,402,674 observations removed).

Potential duplicates with the same account number at the same RFI in the same year, but with missing owner TINs, are more difficult to screen. We note that in the case of missing owner TINs, keeping one record is sufficient as we have no hope of matching multiple owners, so the main challenge is to ensure that we do not exclude the records with missing TINs and placeholder account numbers. To do this, we screen for distinct account balances, as account balance is rarely missing and truly distinct accounts should have distinct account balances in the vast majority of cases. As such, we encode those records with the same masked account number at the same RFI but a missing owner TIN as a non-duplicate if the account balance is distinct. In the event that we observe multiple records with the same account number and RFI, missing owner TINs, and the same account balance, we regard these as duplicates, as before, we keep the first record with the highest data quality code (469,629 observations removed).

Joint Ownership

Generally, when an account is jointly owned, per the instructions for Form 8966, we expect to see an 8966 for each owner, with the *same account number* and *distinct owner information* on each of these 8966's. Additionally, RFIs are instructed to report the income and asset information in total for the account and file a different Form 8966 for each owner, rather than to divvy these up to various joint owners, so the information on Part IV should also be the *same in the case of joint ownership*.

In principle, we can therefore identify jointly held accounts by looking for cases where the same account number at the same RFI is reported as owned by distinct owner TINs. However, implementing this principle naively would cause us to inadvertently flag accounts with placeholder account numbers as jointly owned accounts. To deal with this issue, we obtained a list of the masked analogue of the 7 most commonly observed “placeholder account numbers” from the IRS (the most commonly observed of which is the masked analogue of the placeholder that RFIs are actually instructed to use when an account number does not exist). We replace the account number to missing when the account number takes one of these placeholder values, and we do not encode any of these as null accounts. Doing so helps us ensure that we do not mistakenly assume that accounts with these placeholder account numbers are jointly held by many owners. If the account number is not missing, the owner TINs are distinct, and the account balance matches another record with the same account number, we assume that these records represent jointly owned accounts.

We next need to devise a way of handling jointly held accounts that 1) avoids double counting the same wealth when we tabulate totals of financial variables, and 2) properly counts the number of owners of

offshore wealth (e.g. to count the fraction of individuals in a given part of the income distribution with an interest in an offshore account). To do this, we keep all of the records for jointly owned accounts (so that owner counts are correct), but we divide the financial variables proportionally among the owners (to avoid double counting dollars of income/wealth). For example, if two individuals own a joint account (by far the modal case), we replace the account balance and income variables by two, allocating half of the wealth and income to each owner. When true ownership is not an equal split among all joint owners, this allocation is conservative in most respects, as e.g. it could reduce the concentration of ownership of offshore wealth.

We next looked in detail at records with large numbers of owners, as these could correspond to errors that are not actually jointly held accounts. On inspection, it appeared that instances where the number of owners is greater than 25 are virtually all due to errors. For example, the distribution of account balances tilts sharply toward very small values above around 25 owners. If the number of owners is greater than 25, we keep the first observation for that account so that total wealth and the total number of accounts are unaffected, but the remaining records are removed from the sample (36,597 observations removed). Overall, 13.41 % of observations in the preferred sample are jointly owned.

Variables for Estimating Rates of Return

We create a few auxiliary variables in order to estimate rates of return. First, we create a series of indicator variables indicating the following: whether interest income is non-missing, whether interest or dividend income is non-missing, whether interest, dividends, or gross proceeds is non-missing, and whether any income is non-missing. We create an analogous set of auxiliary variables equal to the account balance if specified income variables are non-missing, and otherwise the variable is missing.

Unresolved Data Anomaly

While we managed to resolve most obvious data anomalies, we could not resolve one large spike in total wealth in a single country in 2016. Total wealth in that country spikes in 2016, drops in 2017, and then increases again in 2018; these spikes are large enough to matter for the time series of total wealth in the Form 8966 data. We were able to identify three RFIs in this country that were reporting an anomalously large number of accounts with very large total account balances in 2016, but no issue was found either in the currency conversion process or with a specific observation within these RFIs. Meanwhile, all RFIs report a sizable (if inconsistent) number of large accounts in every year, with many more large accounts in 2018 and 2016 than 2017. In other words, it could be that these RFIs reported a number of large accounts unnecessarily in 2016, but it could also be that these RFIs actually neglected to report on some large accounts in 2017. In the main sample, we left these accounts as is, but we report in Table B1 what the total account balances would be if we either 1) exclude all records from these RFIs in 2016 (middle column, our best proxy if 2016 was the problem year) or 2) if we replace total wealth at these RFIs in 2017 with the midpoint between TY2016 and TY2018 (third column, our best proxy if 2017 was the problem year).

Table B1 – Sensitivity analysis for problematic RFIs: Total Wealth Over Time with Various Specifications

SPECIFICATION				
TAX YEAR	Baseline Total Account Balance (Trillions USD)	Drop 2016 records at problematic RFIs (Trillions USD)	Replace 2017 records at problematic RFI with average wealth (Trillions USD)	
2015	1.65	1.65	1.65	1.65
2016	3.65	2.65	3.65	3.65
2017	3.23	3.23	3.74	3.74
2018	3.98	3.98	3.98	3.98

B3. Linking Owners’ Tax Returns

We use standard fuzzy matching procedures to link F8966 owners to tax returns of individuals and entities on the basis of TINs and/or names. The initial output of this process is a nonexclusive categorization of every record as follows:

- Records linked to individuals via Forms 1040 or other records pertaining to individuals (e.g. Social Security records) are classified as **individuals**
- Records linked to partnerships (Forms 1065 and 1066) are categorized as **partnerships**
- Records linked to C-corporation tax returns (the form 1120 family excluding 1120-S and 1120-F) are categorized as **C corporations**
- Records linked to S-corporation tax returns (Form 1120-S) are categorized as **S corporations**
- Records linked to tax returns of foreign corporations (Form 1120-F) are categorized as **Foreign Corporations**
- Records linked to tax exempt entity returns (the Form 990 family) are categorized as **tax exempt**
- Records linked to tax returns of trusts (Form 1040, 1041A, Form 5227) are categorized as **trusts**

For the most part, these describe the owner types for the matched records, see e.g. Table 3 of the main text. A few records belong to more than one of the above categories. One set of cases this occurs is when the entity changed its classification, so we link the 8966 record to one type of entity return in one year and a different type in another year. In these cases we use the most recent tax return we observe to classify the entity. In another set of cases, the TIN matches a unique tax return but the entity name matches another tax return of a different type of entity; in this case we use the TIN match to classify the entity. Third, when we have an ambiguous match but the TIN matches another form filed by just one of the candidate entity types, we disambiguate the owner type on the basis of the other form. For example, a record that matches to a partnership return and a trust return will be classified as a partnership if the owner information also matches a Form 8804. Records that belong to multiple categories after our best attempts at disambiguation are assigned an **ambiguous match** owner type.

Apart from the ambiguous matches, we create a few other categories for records we cannot confidently link to an owner’s tax return. Some records do not match to the tax returns for any of the entity types described above, but the records to match the filer information for non-tax-return forms filed by

businesses or employers (Forms 940, 941, 943, 944, 945, 2290, or 720). In this case, the owner type is classified as an **unknown entity**. (Otherwise the owner type is classified based on the type of tax return it files). Some observations match a valid individual (TIN and/or name) in Social Security administrative records but no tax returns; these individuals are classified as **unknown individual**.

Records where the owner TIN was missing and we were unable to match to the owner’s tax return with confidence based on name are classified as **missing TIN**. Records where an owner TIN was present on the Form 8966 but matched none of these other data, and a name match was not possible, are classified as **Unmatched TIN**. We suspect many of these correspond to incorrect or invalid TINs.

B4. Missing Data

This section contains supplemental analysis of reporting quality issues in the main sample.

Table B2 reports the share of observations (accounts) in the main sample in 2018 for whom a given variable is missing, breaking observations down by country type. These rates of missing data are relatively stable over time; we focus on 2018 information because this data are most used in the main text. We observe that Missing TINs are especially concentrated in non-havens and in particular non-havens with prior automatic exchange of information. For the financial variables, the rate at which variables are missing does not vary systematically by owner type. Just over 5% of observations have no account balance but some reported income (recall that observations without account or income information are excluded from the sample, see above). Income information is far more frequently missing, with around 60% of observations having no reported interest and over 90% having no information for the other three income variables.

Table B2 - Share of accounts with missing information by country type (TY2018)

	Country type			
	Haven	Non-haven	NH - EOI	NH - No EOI
Missing Tin	20.8	38.8	43.3	19.0
Missing Account Balance	6.4	5.4	5.5	5.4
Missing Interest	59.6	61.5	62.7	56.1
Missing Dividends	93.1	92.1	91.2	96.4
Missing GP	93.6	93.6	93.8	92.6
Missing Other	89.1	96.7	97.1	95.1
Number of observations	612,406	3,954,216	3,222,800	731,416

In the previous table and in the main Table figures concerning missing data (e.g. Table 2), we count as missing instances where the variable in question was reported as the number zero and those where the reported dollar amount is blank. In Table B3 below, we separately tabulate the frequency of these two cases. We observe that in general, blank variables are more common than zeros for the income fields,

while account balances, in contrast, are more often zero than blank. This suggests that while many “missing” account balances might represent accounts that had no balance at the end of the year, the income variables are often missing because the FFI did not report this information rather than because the FFI indicated that the account received none of the given type of income.

We also report dollar-weighted figures for the income variables, i.e. the share of account balances where an income variable was blank or zero. Observations are zero interest income are noticeably less common for large accounts, so the dollar-weighted zero share is much smaller for interest, suggesting that when we observe zero interest, in many cases, this is because the account received no interest income. However, this pattern does not appear so strongly for other income variables, making it less clear what to make of zeros for these. Dividends are almost always blank rather than zero. Likewise, given that a large share of observations and dollars of wealth are associated with blank income fields, we find it plausible that usually, when these variables are missing, it occurs because the filer did not report income, not because there was no income to report.

In Figure B1, we examine how the likelihood of a blank or zero income variable varies by the account balance. In Panel A, we observe that the probability of a blank income variable is decreasing as we move to very large accounts, but the share blank remains high even for very large accounts. Even at the very top of the account balance distribution, about of accounts have no reported income information of any kind. The probability of observing zeros for income variables is lower. For interest, this probability drops sharply with the account balance. For other income fields, the probability is surprisingly flat by account balance – we might have expected larger accounts to be associated with riskier and/or more sophisticated investments that would generate dividend and gross proceeds income, making zeros less likely.

From these statistics and supplemental analysis of the blank versus zeros question, we concluded that we cannot be certain that when the number zero appears, this is an actual report of zero dollars of income/account balance rather than that the information is not reported by the filer. Likewise, blank income fields are far to prevalent even for very large accounts that almost certainly should receive some type of income for us to regard blank observations as zeros. As such, we included both cases in our missing data definition in the main text. We have discussed this issue with experts on the data, who indicated they were aware of the prevalence of missing income data but unaware of how we might further disambiguate true zero amounts from non-reported information.

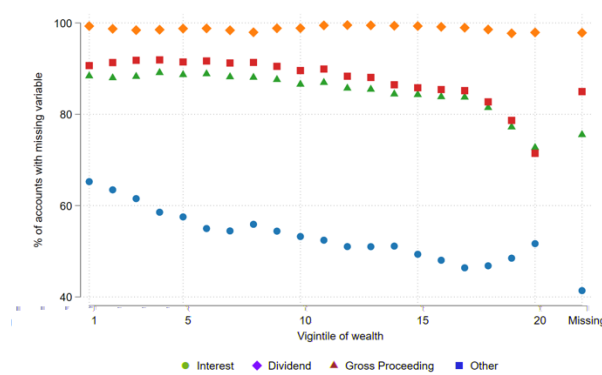
Table B3 – Blank and zero probabilities for Dollar-Valued Variables (TY2018)

A. Pr(variable is blank), in %					
	Account Balance	Interest	Dividends	Gross Proceeds	Other
Account weighted	1.1	53.2	98.8	84.9	87.6
Weighted by Acct Bal	--	62.9	99.3	78.0	70.5
B. Pr(variable is zero), in %					
	Account Balance	Interest	Dividends	Gross Proceeds	Other
Account weighted	4.5	14.8	0.1	8.9	8.2
Weighted by Acct Bal	--	3.9	0.0	7.8	10.0
C. Pr(variable is blank variable is blank or zero), in %					
	Account Balance	Interest	Dividends	Gross Proceeds	Other
Account weighted	19.3	78.3	99.9	90.5	91.4
Weighted by Acct Bal	--	94.1	100.0	90.9	87.5

Note: This Table reports the share of account (unweighted and weighted by account balance) with blank (panel A) or zero (panel B) in one of the main variables on all accounts and share of blank on all accounts with blank or zero value (panel C).

Figure B1. Probability of Blank versus Zero Income Variables by Quantiles of Account Balance

Panel A. Prob Income Variable is Blank



Panel B. Prob Income Variable is Zero

