

Homes Incorporated: Offshore Ownership of Real Estate in the U.K.*

Niels JOHANNESSEN (University of Copenhagen, CEBI and CEPR)
Jakob MIETHE (University of Munich)
Daniel WEISHAAR (University of Munich)

2 December 2022

Abstract

Ownership of real estate through corporations in offshore tax havens creates opportunities for tax evasion and money laundering and may have undesirable effects in housing markets. In this paper, we study offshore ownership of real estate in the United Kingdom by combining several data sources: administrative data from the land register, a comprehensive transaction database, a propriety database on corporate ownership links, and a handful of offshore data leaks. Our descriptive analysis shows that the market share of offshore corporations has increased over time and varies strongly across market segments: It currently stands at 1.25% in the overall residential market and around 15% for top-end properties. When data leaks allow us to trace ownership through offshore corporations to the beneficial owners, we find that around half have ties to Africa, Asia and the Middle East, but that the largest 'foreign' investor is the United Kingdom itself. Turning to causal evidence, we show that changes in tax incentives and ownership transparency induce strong responses in patterns of offshore ownership, suggesting that both taxation and secrecy are important motives for the beneficial owners. Finally, we show that the Brexit referendum was followed by a sharp increase in property sales by offshore owners and a large differential decrease in property prices in local areas with more offshore ownership, conditional on area and property characteristics. This suggests that the reduction in demand from offshore investors triggered by Brexit had a negative causal effect on property prices and, more broadly, that offshore ownership can have significant real effects in housing markets.

*Niels Johannesen: niels.johannesen@econ.ku.dk; Jakob Miethe: jakob.miethe@econ.lmu.de; Daniel Weishaar: daniel.weishaar@econ.lmu.de. We are grateful to Clemens Fuest, Dominika Langenmayr, Andreas Peichl, Kurt Schmidheiny, Juan Carlos Suárez Serrato, Gabriel Zucman, and to several seminar and conference audiences, especially at UC Berkeley, NHH, University of Copenhagen, and the University of Munich for excellent comments and suggestions. All remaining errors are our own. Johannesen gratefully acknowledges financial support from the Danish National Research Foundation (grant DNRF134), from the Danish Finance Institute (DFI), from the European Commission through the EU Tax Observatory and from the United Nations University World Institute for Development Economics Research (UNUWIDER) through the project on illicit financial flows. Weishaar gratefully acknowledges financial support from the Joachim Herz Foundation. Declarations of interest: none

1 Introduction

Anecdotal evidence suggests that top-end real estate in large cities like London, New York and Miami is increasingly held through opaque offshore structures. For instance, an investigative news story documents that 59 apartments out of 76 in One Hyde Park, an exclusive residential development on one of London’s most fashionable addresses, belong to corporations in offshore tax havens like the Cayman Islands, the British Virgin Islands and the Isle of Man with no trace of the beneficial owners in the land register (Shaxson, 2013).

Anonymous holding of real estate through offshore corporations is a cause of concern for several reasons. First, it may play a role in money laundering and tax evasion. While cross-border exchange of bank information has made it increasingly difficult to hide financial assets on offshore accounts (Guyton et al., 2022), it remains simple to launder money and evade capital taxes in the home country through investments in foreign real estate markets as there is no systematic cross-border reporting of real estate assets. Second, it may facilitate avoidance of local tax liabilities arising around property transactions. In many countries, taxes on transaction values and realized capital gains constitute an important source of revenue but such taxes are difficult to enforce when properties can change hands through the exchange of shares in offshore corporations with no change of ownership recorded in the land registry. Third, the broader effects on real estate prices and urban development may be socially undesirable. The construction of exclusive properties for foreign investors motivated by tax and secrecy may crowd out scarce housing for middle-class dwellers and contribute to the increasing trend in real estate prices (Knoll et al., 2017).

In this paper, we provide new evidence on offshore ownership of residential real estate in the United Kingdom. Combining data from many sources, we first address a range of *descriptive questions*: What share of the residential properties is owned through corporations in offshore tax havens? How does offshore ownership vary across price segments and local areas, and how has it evolved over time? Who are the ultimate owners behind the anonymous offshore structures that own real estate? Exploiting different sources of policy variation, we also address several *causal questions*: Are decisions about offshore ownership shaped by tax advantages and secrecy for the beneficial owners? Does offshore ownership affect prices in housing markets?

Our analysis combines micro-data from multiple sources. First, to delimit the set of properties with offshore corporate owners, we rely primarily on the land register. For recent years, this set of properties is directly observable in datasets published at regular intervals by the land register; however, to explore developments over longer time horizons, we need to make inferences

based on transaction data from the same source, obtained through a Freedom of Information request. Second, to determine the value of the properties with offshore corporate owners as well as their share of the total market, we add transaction prices from a database covering the entire residential property market. This step is far from trivial as it involves unbundling both transactions and land titles, which may comprise multiple properties in the raw data, to the level of single properties with unique postal addresses using natural language processing techniques. Third, to identify properties that are controlled indirectly through an offshore corporation, e.g. a Jersey corporation holding a U.K. corporation that owns a property, we employ a global firm database that allows us to unwind corporate ownership structures. Fourth, to learn about the beneficial owners behind the offshore corporations controlling U.K. properties, we search a database containing information from five major leaks from offshore tax havens, e.g. Panama Papers and Paradise Papers.

The first part of the paper provides rich descriptive evidence. We show that corporations in offshore tax havens – most commonly registered in the British Virgin Islands, Jersey and Guernsey – own residential properties in England and Wales worth around £50 billion. This is around five times more than the combined holdings of corporations registered in foreign non-haven countries, such as the United States, Japan and France. There are systematic differences across market segments in the importance of ownership through offshore tax havens: Their market share stands at around 1.25% in the overall residential market but reaches 15% in the most expensive market segments.¹ Relatedly, there is barely any offshore ownership in rural areas and smaller towns but a high market share in affluent parts of large cities like Manchester, Leeds and most notably London. Using transaction data to track offshore ownership over a period of 25 years, we are able to describe the long-run development in the market share of offshore tax havens and document that the offshore boom in property markets mainly occurred in the period 2005-2015. Finally, exploiting the data leaks, we identify the beneficial owners for a subsample of properties held through offshore tax havens. Around half of the properties have ties to Africa, Asia and the Middle East and much fewer to North America and neighboring European countries. Residents of the United Kingdom account for around 15% of the property investments in the sample, suggesting an important role for round-tripping, i.e. domestic investors buying domestic assets through offshore corporate vehicles.²

¹This finding also resonates with the well-documented pattern that offshore financial assets are highly concentrated in the hands of the wealthiest (Alstadsæter et al., 2019; Londoño-Vélez and Ávila-Mahecha, 2021; Guyton et al., 2021; Leenders et al., 2022; Guyton et al., 2022; Londoño-Vélez and Tortarolo, 2022).

²Existing studies have identified round-tripping in the context of financial investment (Hanlon et al., 2015; Damgaard et al., 2019).

We compare these patterns to offshore ownership in financial markets and obtain three interesting insights. First, while offshore ownership is generally less pervasive in residential property markets than in financial markets, the opposite is true at the top of the property price distribution: a larger share of the money invested in exclusive residential properties in the U.K. comes from offshore tax havens than money sitting in U.K. bank accounts. Second, the offshore boom is a much more recent phenomenon in property markets than in financial markets: While offshore ownership of funds in U.K. bank accounts has been roughly stable since the financial crisis, offshore ownership of residential properties increased more than fivefold in the period 2005-2015 and continue to grow although at a slower pace. Third, money invested in residential properties through offshore structures is much more likely to originate in developing countries than financial investments through offshore accounts: We show that beneficial owners in Asia, Africa and the Middle East account for around 50% of the properties in the U.K. with nominal offshore owners, much higher than what other sources suggest for offshore financial assets.

The next set of results concerns the causal determinants of offshore ownership, in particular the role of taxation and secrecy for the ultimate owners. We exploit two sources of U.K. policy variation for identification. First, a reform that tightened the capital gains taxation of foreign corporations with property in the U.K., but temporarily created a loophole for corporations registered in Luxembourg. Second, a government announcement mandating a group of offshore tax havens known as the Overseas Territories (e.g. Bermuda) to fight financial secrecy by setting up public corporate registers with information about beneficial ownership, but not extending the mandate to another group of offshore tax havens known as the Crown Dependencies (e.g. Jersey). Our results document that both events caused sharp behavioral responses with offshore ownership migrating to jurisdictions with low taxation and low transparency.

The final set of results concerns the real effects of offshore ownership in local property markets. Focusing on London where offshore ownership is particularly widespread, we use the surprising “leave” outcome of the referendum on European Union membership in 2016 as a natural experiment that triggered a sharp increase in property sales by offshore owners. In a model that uses repeat sales for identification and allows price dynamics to vary flexibly across price segments, we show that local areas with a high ex ante share of offshore ownership experienced larger price decreases after the referendum. Consistent with a causal interpretation, local areas with different offshore ownership were on parallel trends prior to the referendum and controls for potential confounders, such as the share of foreign residents and the share of total corporate ownership in the local area, have little impact on the estimates. The implied economic

effects are relatively large with a differential decrease in property prices of around 1.5% at the 75th percentile of offshore ownership shares, around 5% at the 90th percentile, and around 7.5% at the 95th percentile.

Employing the same empirical framework, we find evidence suggesting that the effect of offshore ownership on prices is not limited to the expensive properties most frequently held through offshore corporations. Specifically, we show that local areas with a higher ex ante share of offshore ownership in the high-price segment experienced larger price decreases after the referendum in other price segments too. This result highlights that middle-class home owners may also be exposed to the consequences of offshore property ownership due to spillover effects across price segments.

Our paper contributes to the broad descriptive literature on household wealth held through offshore tax havens (Zucman, 2013; Alstadsæter et al., 2019). While this literature focuses almost exclusively on financial wealth, we add to the scant evidence on offshore ownership in property markets. By focusing on properties held through corporations in offshore tax havens, our object of interest is distinct from a recent paper that measures all ownership by foreign nationals, including local residents, in the Dubai property market (Alstadsæter et al., 2022). We complement existing descriptive analysis of offshore ownership in the United Kingdom (e.g. Bomare, 2019) in two ways. First, by integrating comprehensive price data in the analysis, we are able to compute the share of all residential real estate assets owned through offshore tax havens, a key metric that we denote the offshore market share. Second, by exploiting transaction data, we can track this metric back to the 1990s and, thus, document the rise of offshore ownership in the property market over several decades. We are not aware of existing papers that produce long time series of offshore property ownership. Our efforts to pierce the veil of offshore secrecy and assign real estate assets held through opaque ownership structures to their ultimate owners relates to similar efforts for corporate profits (Tørsløv et al., 2022); portfolio investment (Zucman, 2013; Coppola et al., 2021); bank deposits (Alstadsæter et al., 2018; Collin, 2021); and direct investment (Damgaard et al., 2019).

We also contribute to a small existing literature studying the causal drivers of offshore property ownership. Two studies report conflicting evidence on the role of secrecy. One of them shows that increased transparency in property investments – i.e., U.S. rules requiring disclosure of the ultimate owners when corporations make all-cash acquisitions of real estate – is associated with *no change* in behavior (Collin et al., 2022). The other one shows that increased transparency in financial investments – i.e., global rules requiring banks to share account information with

the home countries of their foreign customers – is associated with *more* corporate investment in properties (Bomare and Le Guern Herry, 2022). Our finding that offshore ownership shifts away from offshore tax havens that improve transparency toward those that do not is consistent with a secrecy motive and resonates with earlier evidence on shifting of financial assets across offshore tax havens in response to enforcement efforts (Johannesen and Zucman, 2014; Johannesen, 2014; Menkhoff and Miethe, 2019). Our finding that offshore ownership in property markets are highly responsive to tax incentives does not, to the best of our knowledge, have antecedents in the literature.

Finally, we contribute by showing that ownership through offshore tax havens has real consequences in housing markets. We show that a negative shock to offshore demand for U.K. property, triggered by the outcome of the Brexit referendum, was associated with a significant drop in prices. This relates to earlier work showing that positive shocks to foreign demand for property, triggered by political risk in investor countries (Badarinza and Ramadorai, 2018) or by new taxes in other investment countries (Gorback and Keys, 2020), can drive up prices. Our analysis differs from this earlier literature by focusing on ownership through offshore corporations, as opposed to ownership by foreigners per se, and by identifying spillover effects from the top-end markets most directly affected by lower offshore demand to other market segments.

The paper proceeds in the following way. Section 2 describes the data. Section 3 provides descriptive evidence on offshore ownership. Sections 4 and 5 report results on the causes and consequences of offshore ownership respectively. Section 6 concludes.

2 Data

This section describes the data sources we use in the paper and explains how we combine them to obtain a database of offshore ownership of real estate in England and Wales. In brief, we obtain comprehensive information about corporate acquisitions of property from the land register (section 2.1); add information about prices and transaction dates from an administrative database of real estate transactions in the residential market (section 2.2); and add information about holding companies and ultimate owners from a global firm database and from various offshore leaks (section 2.3). We provide an overview of the database (section 2.4) and report technical details about how the different data sources are combined to the Online Appendix B.

2.1 Corporate Ownership

Our main data source is two publicly available datasets from HM Land Registry with comprehensive administrative information about land and property in England and Wales held by legal entities. The Commercial and Corporate Ownership Data (CCOD) covers titles held by entities incorporated in the UK whereas the Overseas Company Ownership Data (OCOD) covers titles held by entities incorporated outside of the United Kingdom. Both datasets are organized as snapshots taken at regular intervals: The CCOD is available quarterly from March 2014 and monthly from December 2017 whereas OCOD is available quarterly from October 2015 and monthly from December 2017. Both datasets contain information about properties (e.g. address, tenure), about owners (e.g. company name, country of incorporation) and the date on which the proprietor was added to the land register.

We complement these relatively recent snapshots of the land register with information from two alternative sources that cover a longer time horizon. First, through a Freedom of Information (FOI) request, we have obtained data on all purchases of real estate in England and Wales by foreign corporations since 1990. Second, we use information from the current affairs magazine *Private Eye* (PE) that has collected property and ownership information about property titles in England and Wales acquired by foreign companies between 1999 and 2014. No similar data on purchases by domestic corporations is available.³

By harmonizing property addresses and company names across these four sources using natural language processing techniques, we are able to combine them to a single database of corporate acquisitions of property in the U.K. For each acquisition, we record the property, the corporation and the month in which the corporation is added to the land register as the owner of the property. For domestic corporations, the only data source is CCOD, which means that our database includes all acquisitions after March 2014, but only acquisitions before March 2014 to the extent that the corporation remained the registered owner by that date. For foreign corporations, we rely on OCOD as well as FOI and PE, which means that our database includes all acquisitions since 1990.

2.2 Transaction Prices and Dates

We combine the database of corporate acquisitions with the Price Paid Data (PPD) from the HM Land Registry, which covers all residential property transactions in the United Kingdom since January 1995. This step is complicated by the fact that both land titles in CCOD/OCOD and

³HM Land Registry confirmed this in a response to our Freedom of Information request.

transactions in PPD not rarely lump together multiple properties. We overcome this challenge by developing an algorithm that breaks down all three data sources to the level of single properties with unique addresses and merge them at this level (see Online Appendix B for details).

Combining with transaction data from PPD allows us to add important information for the subset of acquisitions that involves residential properties. First, we obtain the transaction prices. While the most recent transaction price is often available directly for the properties in the CCOD/OCOD snapshots, this is not the case for properties that were acquired by corporations and sold again prior to the first snapshots. For instance, in the case an offshore corporation acquires a property in July 2010 and sells it in May 2012, the property is not in the CCOD/OCOD snapshots; however, we can identify the acquisition in the FOI data and retrieve the transaction prices in the PPD data. Second, we obtain the transaction month. This does not generally coincide with the month in which the new owner is recorded in the land register because of administrative delays. For instance, in the case an offshore corporation acquires a property in July 2010, it is typically added to the land register in October or November the same year, but we can determine the actual transaction month in the PPD data. Third, we obtain the month in which corporate ownership ends. Again, this can often be inferred approximately from CCOD/OCOD, as the month when a corporate owner drops out of the land registry, but the PPD allows us to determine the timing more precisely, as the month in which a property previously acquired by a corporation is transacted, and - perhaps more importantly - we can also apply this approach to the period not covered by the CCOD/OCOD snapshots. For instance, in the case an offshore corporation acquires a property in July 2010 and sells it in May 2012, the property is not in the CCOD/OCOD snapshots. However, we observe that ownership is transferred to a foreign corporation in the FOI data and we observe that the property is sold to another owner in May 2012 in the PPD data.

2.3 Upstream Owners

We take two distinct approaches to tracing the owners behind the corporations recorded as proprietors of real estate in the land register. The first approach uses ORBIS, a proprietary dataset with accounting information and ownership links for a global sample of firms. For each of the domestic corporate owners identified in the CCOD dataset, we search for upstream ownership links in Orbis. This allows us to identify properties whose direct owner is a domestic corporation, which is in turn controlled by a holding company in a tax haven. Such structures allow the ultimate owners to benefit fully from the advantages offered by tax haven ownership in terms of

secrecy and tax advantages while not exposing this directly in the land register.

The second approach uses leaked information from corporate service providers and corporate registers in offshore tax havens published by the International Consortium of Investigative Journalists (ICIJ). The leaks pierce the secrecy emblematic for tax havens by identifying the individuals who are the ultimate beneficial owners of the corporations in the leaked files. For each tax haven corporation in CCOD and for each upstream tax haven owner of non-haven corporations identified in ORBIS, we thus search for information about the ultimate beneficial owners in all of the five available leaks: Pandora Papers, Paradise Papers, Bahamas Leaks, Panama Papers and Offshore Leaks. This results in a sample of properties held through offshore corporations, for which we are able to determine the home country of the individuals who are the ultimate beneficial owners. To be clear, in most of the analysis, we are interested in offshore ownership per se and do not restrict the sample to properties, for which we can identify the ultimate beneficial owners.

2.4 Data Overview

Table 1 provides an overview of the data sources and how we use them. First, the various datasets deriving from the land register identify more than 2.4 million properties that are currently owned by corporations or have been owned by corporations at some point since 1990 (Panel A). While most of the corporate owners are domestic, almost 185,000 properties are associated with corporations registered outside of the United Kingdom. In around half of these cases, the foreign corporate owners are recorded in the OCOD snapshots of the land register available since 2015. In the other half of the cases, they are included in the FOI transaction data, implying that the properties no longer had foreign corporate owners at the time of the first snapshots. Second, around 300,000 of the corporate-owned properties are controlled directly or indirectly through offshore tax havens (Panel B).⁴ This is the case for around 180,000 of the properties with a foreign corporate owner identified in OCOD/FOI, but also for an additional 120,000 properties with a domestic corporate owner identified in CCOD, which is itself owned by a corporation in a tax haven identified in ORBIS. Third, for just over 7,000 properties controlled through offshore tax havens, we are able to retrieve information about the beneficial

⁴Our list of offshore tax havens combines the lists from Johannesen and Zucman (2014) and Gravelle (2015): Andorra, Anguilla, Antigua-Barbuda, Aruba, Austria, Bahamas, Bahrain, Barbados, Belgium, Belize, Bermuda, British Virgin Islands, Cayman Islands, Chile, Cook Islands, Costa Rica, Curacao, Cyprus, Dominica, Gibraltar, Grenada, Guernsey, Hong Kong, Ireland, Isle of Man, Jersey, Jordan, Lebanon, Liberia, Liechtenstein, Luxembourg, Macao, Malaysia, Maldives, Malta, Marshall Islands, Mauritius, Monaco, Montserrat, Nauru, Niue, Panama, Samoa, San Marino, Seychelles, Singapore, Sint Maarten, St. Kitts-Nevis, St. Lucia, St. Vincent-Grenadines, Switzerland, Tonga, Trinidad-Tobago, Turks-Caicos Islands, U.S. Virgin Islands, Uruguay, Vanuatu.

owners by searching the ICIJ leaks (Panel C). The Panama Papers contribute most of the links to beneficial owners, but the more recent Paradise Papers and Pandora Papers also make valuable contributions. Fourth, in each stage, a significant share of the properties we identify are residential, which allows us to obtain transaction prices by combining with the PPD data (Column 2). In the parts of the analysis that require price information, we will rely on this residential sample only.

3 Descriptive Evidence

This section provides descriptive evidence on offshore ownership of residential properties in England and Wales. Our key metric is the *offshore market share*, which we define as the share of the property value in a specific market - the entire residential market or a submarket - that is owned directly or indirectly through corporations in offshore tax havens. Our starting point for estimating offshore market shares is the database created in the previous section.

Below, we present estimates of the offshore market share in the overall residential property market (section 3.1); break it down by jurisdictions (section 3.2); show how it varies cross-sectionally across price segments and local areas (section 3.3); and illustrate how it has evolved over time (section 3.4). Finally, we present evidence on the beneficial ownership of properties held through offshore corporations (section 3.5); and compare to patterns of beneficial ownership in financial markets (section 3.6).

3.1 The Offshore Market Share

To estimate offshore market shares in recent years where the land register snapshots are available, we face two empirical challenges. First, we need the values of all properties. As we only observe the market price of properties directly at the time they are transacted, we impute property values at all other times by applying changes in the district-level house price index to the most recent transaction prices. Second, we need to delimit the set of properties in the market of interest. While the land register snapshots only cover corporate-owned properties, the PPD covers all residential properties in England and Wales that have been transacted at least once since 1995. Hence, we can estimate the denominator of the offshore market share by summing imputed values over the set of properties in the market and the numerator by summing imputed values over the subset of these properties with offshore owners.

Table 2 shows the relevant magnitudes for the overall residential property market in England and Wales at the end of 2019. We estimate that the aggregate value of all properties in the

market is around £4,000 billion.⁵ The estimated value of the properties with offshore corporate owners is around £50 billion, with the bulk of the offshore ownership being direct and only a smaller fraction indirect through domestic corporations. These estimates imply a modest offshore market share in the overall residential market of around 1.25%.

By comparison, the aggregate value of residential properties with foreign non-haven corporate ownership is only around £10 billion. It follows that the foreign corporations owning properties in England and Wales are overwhelmingly registered in offshore jurisdictions. This striking pattern suggests that secrecy and tax evasion may be important motives for these investments, a hypothesis that we explore below. Finally, the value of the properties with domestic corporate ownership and no foreign corporations among the upstream owners is around £180 billion, three times more than properties with foreign corporate owners. To be clear, the individuals who are the ultimate owners of these properties may be domestic as well as foreign.

3.2 Foreign Corporate Owners

Having shown above that the bulk of foreign corporate owners are registered in tax havens, we provide more detail by breaking down foreign corporate ownership by jurisdictions in Figure 1. The top owners are the British Virgin Islands and Jersey whose corporations jointly own residential property in England and Wales of more than £25 billion. The vast majority of those properties are held directly (dark shading) whereas properties held indirectly through U.K. corporations only contribute around £1 billion to the aggregate (light shading). The next jurisdictions are Guernsey, Luxembourg and the Isle of Man whose corporations together control properties worth around £10 billion. Hence, all jurisdictions in the top-5 are small tax havens with a legal infrastructure that makes them attractive places to establish shell corporations. Only six countries in the top-20 are not tax havens, e.g. United States, Japan and Germany.

3.3 Price Segments and Local Areas

We explore the heterogeneity in the offshore market share across price segments and local areas. As shown in Figure 2, the offshore market share in the overall market of around 1.25% conceals a strong positive price gradient. While less than 0.5% of the residential properties in low-price and mid-price segments have offshore owners, the increases monotonically with the price level and exceeds 10% in the most expensive price segment with prices above £5 million. This pattern

⁵Figure A2 in the Online Appendix shows that our estimates of the aggregate market value are consistent with national accounts, both in terms of levels and time series dynamics.

is consistent with anecdotal accounts that associate offshore ownership with the most exclusive residential properties.

As shown in Figure 3, offshore ownership is highly concentrated in urban centers such as London, Manchester and Leeds and in selected medium-sized cities such as Exeter, Nottingham and Durham while it is virtually non-existent in rural areas (Panel A). This is consistent with the price gradient in offshore ownership documented above, as most top-end residential properties are located in cities. Moreover, there is considerable variation in the offshore market share across neighborhoods within the same urban centers. Specifically, offshore ownership in London is extremely high in areas such as Shoreditch, Finsbury Park and Somers Town with market shares exceeding 25% and much lower in suburban areas such as Bexley, Havering and Redbridge (Panel B).

3.4 Long-Run Dynamics

We produce a time series for the offshore market share of residential property that goes as far back in time as the data allows.⁶ For recent years where the OCOD/CCOD snapshots are available, we can employ the methodology developed in section 3.1 for each time period separately. For earlier years, we face one additional challenge: In the FOI/PE data, we observe directly when foreign corporations buy properties and, combining with PPD data, we can also infer when they subsequently sell. Hence, we can track whether a property is owned by a foreign corporation at any date *after* the first time it is transacted within the period with data coverage, but not *before* the first transaction.⁷ We address this challenge by assigning to each of these “first transactions” a probability that the seller was a foreign corporation based on what we learn from other transactions in the dataset. Specifically, the assigned probability is the share of foreign corporate sellers across all transactions in the same year, for which we know the seller because there is an earlier transaction in the dataset.⁸ Figure A2 in the Online Appendix shows

⁶We provide more details on the methodology in Online Appendix C.

⁷For example, if a property is transacted in June 1997 according to PPD, we can determine whether the buyer is a foreign corporation by checking if the transaction is included in the FOI/PE datasets. If the property is transacted again later on, say in May 2003, we can do the same. Hence, we can determine precisely when the property has a foreign corporate owner and when it does not at each point in time since 1997. However, if the transaction in 1997 is the first one within the period with data coverage, we cannot infer from the transaction data whether the seller in 1997 was a foreign corporation or not. Hence, in this case, we cannot make inference about ownership at points in time before 1997.

⁸Consider again the example where a property’s first observed transaction is in 1997. What is the probability that the seller was a foreign corporation? We identify all the transactions in 1997, for which there exists an earlier transaction in the dataset and for which we therefore know whether the seller in 1997 was a foreign corporation or not. For this set of transactions, we compute the share with a foreign corporate seller and we assume this to be the probability of a foreign corporate seller in the case at hand.

that this method delivers accurate imputations in the recent period where we can check them against the OCOD/CCOD dataset.

As shown in Figure 4, the resulting time series reveals a striking pattern. Over the full period, the offshore market share has grown very significantly, from around 0.1% in the middle of the 1990s to the level of around 1.25% in 2020. There are notable differences across time periods: the pace of the increase was slow in the first decade (1995-2005), extremely fast in the second decade (2005-2015) and then again slow in the last part of the sample period (2015-2020).⁹

3.5 Beneficial Owners

Finally, we shine light on the individuals who own properties through offshore corporations by computing the distribution of beneficial ownership across countries for the set of properties that can be linked to their beneficial owners through the ICIJ leaks. As shown in Figure 5, the United Kingdom is by far the largest beneficial owner country with a market share of almost 15%.¹⁰ Most of the properties that are truly foreign-owned have beneficial owners in the Middle East (e.g. United Arab Emirates and Saudi Arabia), Africa (e.g. South Africa and Nigeria) and Asia (e.g. China and Malaysia) and a relatively small fraction belongs to beneficial owners in Europe and North America. Strikingly, the two largest beneficial owners in Europe (i.e. Monaco and Switzerland), are both popular residence countries for global economic elites and their shares of beneficial ownership therefore do not necessarily reflect beneficial ownership by European nationals. It is interesting that Russia appears to account for a relatively small share of the beneficial ownership given the anecdotal evidence that Russian oligarchs are heavily invested in London properties (Guardian, 2022) and the massive focus on the seizure of such properties following the invasion of Ukraine.

It is important to emphasize that the results are based on a relatively small subsample, just over 5%, of all the residential properties that are owned through offshore corporations. This implies that the ownership shares are estimates associated with uncertainty: individual countries may have a smaller or larger weight in the universe of beneficial owners than in the leaked subsample. The bootstrapped standard errors used to construct the confidence intervals in the figure quantify the sampling uncertainty under the assumption that beneficial owners from all countries have the same ex ante probability of being disclosed in a leak.¹¹

⁹Figure A1 in the Online Appendix shows the evolution in the aggregate value of properties held through corporations in offshore tax havens and foreign non-havens, i.e. without scaling by the value of the total market.

¹⁰This is consistent with other evidence on round-tripping investment in other economic domains (Hanlon et al., 2015; Damgaard et al., 2019).

¹¹This follows the approach taken by Alstadsæter et al. (2019). In our setting, the risk of selection bias is

3.6 Comparison with Financial Assets

It is illustrative to compare the patterns in offshore ownership in property markets that we have identified above to offshore ownership in financial markets, which has been the focus of a number of earlier studies. In the Online Appendix, we document how the share of deposits in U.K. banks owned through offshore tax havens has evolved over time (Figure A3) and illustrate how the beneficial ownership of financial assets in offshore tax havens may be distributed across countries (Figure A4). The former analysis uses comprehensive data on U.K. banks from the BIS Locational Banking Statistics. The latter draws on leaked customer data from the HSBC Switzerland, a private bank in Switzerland managing around 2% of global offshore assets at the time of the leak (Alstadsæter et al., 2019).¹²

Three interesting insights emerge from the analysis. First, the tax haven boom in property markets occurred much later than in the financial market. While we established above that offshore ownership in property markets built up relatively recently (2005-2015 as shown in Figure 4), offshore ownership in deposit market increased most rapidly in the decade before the financial crisis (1997-2008) and stayed roughly constant, if not falling slowly, in the remainder of the sample period (2008-2020). Second, the property market is generally less affected by offshore ownership than financial markets except at the top of the property price distribution. Specifically, the share of offshore ownership in the overall property market (around 1.25% as shown in Table 2) is significantly lower than in the deposit market (around 10%); but the opposite is true when considering only the most exclusive properties (around 15% as shown in Figure 2). Third, in property markets, the individuals behind the offshore ownership structures are to a much larger extent residents of developing countries than in financial markets. For instance, beneficial owners from Africa and Asia account for a large share of the properties held through offshore corporations (around 20% and 30% respectively) whereas their share of financial assets in offshore banks is much smaller (only around 5% and 15% respectively).¹³

4 Evidence on Causes

This section investigates the underlying motives for using offshore corporations as vehicles for owning real estate. The literature generally highlights two advantages associated with investment

smaller because we rely on data from five distinct data leaks and more than five distinct data sources, as some of the leaks include data from multiple sources.

¹²The leaked data was published by the ICIJ under the name *Swiss Leaks*.

¹³The comparison of beneficial ownership should be interpreted with some caution, as the data come from leaks with incomplete coverage.

through tax havens, *lower taxation* and *secrecy*, and we provide evidence on each of these motives in turn. We first study the role of taxation by estimating the behavioral response to a reform that eliminated the favorable treatment of some but not all offshore structures in terms of capital gains taxation (section 4.1). We then study the role of secrecy by estimating the behavioral response to a reform intended to enhance transparency for some but not all offshore structures (section 4.2).

4.1 The Tax Motive

Generally, owning real estate through an offshore shell corporation can legally reduce the effective taxation in different ways. Most prominently, it may allow owners to avoid transaction taxes and capital gains taxes in the country where the property is located, as they can trade the shares in the offshore shell corporation rather than the property itself.

We consider a reform in the United Kingdom that targets this type of tax avoidance by extending the taxation of capital gains to "indirect disposals", i.e. the trading of shares in companies whose main assets are property in the United Kingdom. The reform was announced in November 2017 and took effect from April 2019. It allowed owners within the scope of the new rules to rebase property values so that only capital gains accruing after April 2019 would be subject to taxation.

While the reform aimed to remove one of the main advantages of holding real estate through offshore shell corporations, its effectiveness was limited by an existing double tax treaty with Luxembourg, which prevented the United Kingdom from taxing capital gains made on the disposal of shares in Luxembourg companies. The provision thus created a loophole by effectively exempting capital gains on U.K. properties, realized indirectly through the sales of Luxembourg holding companies, from capital gains taxation in the United Kingdom. The loophole was eventually eliminated with Luxembourg joining the OECD multilateral convention amending tax treaties to prevent base erosion and profit shifting and with the United Kingdom and Luxembourg signing a new treaty that explicitly covers the tax treatment of indirect disposals of properties.¹⁴

These events create useful variation in the tax incentives to hold property through offshore corporations. Before April 2019, offshore corporations in any tax haven would allow owners to avoid taxation of capital gains, but only corporations in Luxembourg could serve this purpose after April 2019. To the extent that tax avoidance is a motive for holding property through

¹⁴We provide more details on the policy reform, the treaty amendment and their implications for taxation of capital gains in Online Appendix D.

offshore corporations, we should therefore expect an increase in holdings through Luxembourg relative to holdings through other tax havens. Specifically, we should expect owners of property to switch to Luxembourg holding structures at some time between November 2017 and April 2019, since this would allow them to continue to avoid capital gains taxes after April 2019.¹⁵

To test these predictions, we use the database on corporate ownership constructed above to compare the evolution of property titles in England and Wales held by corporations in Luxembourg and corporations in other tax havens. We make the scope of the analysis as broad as possible by including both residential and non-residential properties and estimate the following model:

$$\log(y_{it}) = \sum_{t \neq \text{Oct.2017}} \beta^t d_t \times d_i^{Lux} + \gamma_t + \alpha_i + \epsilon_{it} \quad (1)$$

where y_{it} is the number of property titles held by corporations in country i in month t , d_i^{Lux} is a dummy indicating whether country i is Luxembourg and d_t captures monthly time dummies where the omitted category corresponds to one period before the announcement of the policy in November 2017. The equation includes country fixed effects, α_i , that absorb all cross-sectional variation in holding patterns as well as time fixed effects, γ_t . The estimates of β_t have the flavor of difference-in-differences estimates capturing the percentage change in the number of properties held through Luxembourg since the reference period over and above the percentage change in properties held in other tax havens.

Consistent with tax avoidance being an important motive for holding real estate through offshore corporations, the number of properties held by corporations in Luxembourg increased significantly relative to corporations in other tax havens through 2018 and early 2019, as illustrated in Figure 6. Importantly, the parallel trends prior to the policy announcement in late 2017 strengthens the case for a causal interpretation of the divergence in early 2018.

In the Online Appendix, we provide additional evidence on the cross-border flows of property titles underlying these results in Figure A6. We show that the increase in properties owned through Luxembourg corporations primarily reflects shifting of property titles from corporations in the United Kingdom and nearby tax havens like Jersey and Guernsey; however, also properties not previously held through a corporation contribute significantly to the increase. We also report the results from additional robustness analysis in Figure A5. While the baseline uses a logarithmic transformation of the outcome, we show that the results remain similar when

¹⁵To the extent that owners anticipated that the Luxembourg loophole would eventually be closed, it may plausibly have muted the behavioral response.

we replace it with the inverse hyperbolic sine transformation (Panel A) and when we use no transformation at all (Panel B). Likewise, the units in the baseline analysis are property holdings, i.e. we count the number of postal addresses, but we show that the results are robust to alternative units: property titles as they are defined in the UK land register (Panel C) or offshore companies (Panel D).

4.2 The Secrecy Motive

Besides legal tax advantages, owning real estate through an offshore corporation may create additional advantages for the owners by allowing them to conceal their identity. Generally, shareholders in corporations registered in tax havens benefit from de facto anonymity, which may serve several purposes. First, it protects owners against governments or private creditors trying to seize their real estate assets. A recent example is the efforts of European governments to seize the assets of Russian oligarchs after the onset of the Russian-Ukrainian war, which were frustrated by offshore ownership structures blurring the links between owners and assets (The Atlantic, 2022). Second, it creates opportunities for tax evasion. Foreign real estate and the income that it generates are often taxable in the owner's home country, but ownership through offshore structures makes it much more difficult to detect non-compliance.

To investigate the importance of the secrecy motive, we consider a policy initiative targeting the secrecy provided by offshore corporations. In May 2018, the U.K. government announced that it would mandate the Overseas Territories - including well-known tax havens such as Bermuda, the Cayman Islands and the British Virgin Islands - to set up public corporate registers with information about beneficial ownership. Combined with the existing public land registers, the corporate registers would make it possible to trace the ultimate owners of real estate in England and Wales through corporate structures in these jurisdictions and thus make such structures ineffective as vehicles of asset protection and tax evasion. Importantly, the mandate did not affect the Crown Dependencies - another group of well-known tax havens comprising Jersey, Guernsey and the Isle of Man - whose constitutional status is different and grants them a higher degree of autonomy. In June 2019, the Crown Dependencies announced that they had decided to set up corporate registers voluntarily.¹⁶

This policy process creates useful variation in the incentives to use corporations in otherwise similar tax havens to conceal the ultimate ownership of real estate. By May 2018, it was clear that the Overseas Territories would expose shareholders in a public register whereas no such

¹⁶We provide more details on the policy process leading to the creation of corporate registers in the Overseas Territories and the Crown Dependencies in Online Appendix E.

decision was made by the Crown Dependencies before June 2019. To the extent that secrecy is a motive for holding property through offshore corporations, we should therefore expect an increase in holdings of properties through the Crown Dependencies relative to holdings through Overseas Territories over this period.¹⁷

As the nature of the variation in secrecy provisions is similar to the variation in tax provisions studied above, we can use the same empirical approach to study its effect on behavior. In this case, however, it is natural to restrict the sample to the Crown Dependencies and the Overseas Territories so that identification comes from a comparison of tax havens with highly similar legal institutions affected differently by the mandate of the U.K. government. We estimate the following model:

$$\log(y_{it}) = \sum_{t \neq \text{Mar.2018}} \beta^t d_t \times d_i^{OT} + \gamma_t + \alpha_i + \epsilon_{it}, \quad (2)$$

where the notation follows equation (1) except that d_i^{OT} is a dummy indicating whether country i is an Overseas Territory. As the government’s decision to mandate the register was widely covered in the press already in April 2018, we choose March 2018 to be the omitted time period.

The estimates imply that the number of properties held by corporations in the Crown Dependencies started increasing relative to corporations in the Overseas Territories shortly after the mandate to set up corporate registers in May 2018, as illustrated in Figure 7. Consistent with a causal effect of the mandate, the trends in the two outcomes were parallel prior to the mandate. Moreover, there are no clear signs of differential trends after the decision by the Crown Dependencies in June 2019 to set up similar corporate registers. The finding that ownership structures respond strongly to differential changes in transparency across highly similar tax havens suggests that secrecy is an important motive for holding real estate through offshore corporations.

In the Online Appendix, we report additional results from a robustness analysis in Figure A7. The tests are analogous to the ones we conducted for the tax shock: We show that the results are robust to alternative transformations (Panels A-B) and units (Panel C-D).

5 Evidence on Consequences

This section studies the real effects of offshore ownership in housing markets while exploiting the surprising decision to leave the European Union as a natural experiment that deterred foreign

¹⁷If there was some anticipation that Crown Dependencies would eventually agree to improve transparency for corporate ownership, it may plausibly have muted the behavioral response.

investment in U.K. properties. We first develop an empirical design that exploits variation in offshore ownership shares across local areas in London to estimate how offshore ownership contributed to falling housing prices after the referendum (section 5.1) and then report the results (section 5.2).

5.1 Empirical Design

Two features of the referendum on European Union membership on 25 June 2016 are important for our empirical design. First, the narrow “leave” outcome was generally unexpected, as evidenced by the strong reactions on financial markets.¹⁸ Second, it made it less attractive for foreigners to own real estate in the United Kingdom, as uncertainty about the conditions under which foreign nationals could live, work and invest in the United Kingdom was mounting and strong declines in property prices were anticipated.¹⁹

Consistent with these two conjectures, we show empirically that Brexit was associated with a sudden increase in property sales by owners in offshore tax havens. For the full sample of transactions of residential real estate in London, we regress a dummy indicating that the sale is from an offshore seller to a domestic buyer on a set of time dummies. The regression also includes a set of property fixed effects ensuring that the time dummies are identified exclusively from within-property variation. As shown in Figure 8, the probability that offshore ownership ends jumps sharply in 2016 and stays elevated through the rest of the post-Brexit period relative to the level of the pre-Brexit period 2010-2015 (full line). This pattern suggests that the decision to leave the European Union induced a sharp increase in the propensity of investors holding properties through offshore corporations to liquidate these investments. It did not reflect a mere reshuffling of the offshore portfolio, as the probability that offshore ownership starts, estimated in an analogous way, decreased significantly over the same period (dashed line).

To investigate how the liquidations by offshore owners affected property prices, we exploit that they mechanically affected local housing markets *differentially* due to the pre-referendum cross-sectional heterogeneity in offshore ownership. Intuitively, when offshore owners sell their properties in response to a negative shock, it has a larger impact in local areas where offshore ownership initially accounts for a large share of the properties than in areas where the offshore-owned share of the market is negligible. The tight link between the observable ex ante distribution

¹⁸The British pound initially plunged by 8% against the U.S. dollar whereas the London blue-chip index FTSE 100 fell by 7% (BBC, 2016).

¹⁹The UK government projected that Brexit would cause an overall drop in housing prices of 10-18% over the medium term (Guardian, 2016a) whereas analysts projected price drops of up to 50% in the most expensive parts of London (Guardian, 2016b).

and the size of the differential impact makes negative shocks particularly useful for identification.

In our main analysis, we thus estimate to what extent the Brexit referendum was associated with a differential decline in property prices in local areas with a higher ex ante share of offshore ownership while controlling for other area characteristics. The controls serve to address the confounding effect of area characteristics correlating with offshore ownership and mediating the effect of the Brexit referendum on prices. One salient example is ex ante property prices. Our earlier analysis showed that offshore ownership correlates strongly with the price level, and at the same time the adverse price effects of Brexit may have been different in high-end property markets regardless of offshore ownership. Another important example is local geography. We showed earlier that offshore ownership is concentrated in London and other urban centers and it is possible that Brexit, irrespective of the responses of offshore owners, affected property prices differently in precisely these regions. Hence, our empirical specification needs to ensure that we effectively compare local property markets within the same price segment and within the same geographical region that vary with respect to offshore ownership.

Indexing individual properties by i , local geographical areas by a and calendar years by t and restricting the sample to residential properties in London and the period 2010-2019, we estimate the following transaction-level empirical model:

$$\log(\text{price}_{it}) = \mu_i + \sum_t \Gamma^t d_t \times \mathbf{x}_i + \sum_t \beta^t d_t \times \text{Offshore}_a + \varepsilon_{it} \quad (3)$$

where the dependent variable is the transaction price for property i observed in year t in logs; μ_i denotes property fixed effects; d_t denotes time dummies with the omitted category 2015, i.e. the year before the Brexit referendum; Offshore_a is the share of the residential real estate market in a local area owned through corporations in tax havens before the Brexit referendum; and \mathbf{x}_i denotes a vector of time-invariant controls at the property-level. In the baseline specification, \mathbf{x}_i includes 100 dummies capturing the percentiles of the national distribution of property prices prior to the Brexit referendum.²⁰ In robustness tests, we augment the model with more property-level as well as area-level controls (interacted with time dummies).

The first and second terms of the model capture the estimated price dynamics for properties in areas with no offshore ownership. It is effectively identified by comparing transaction prices for the same property at different points in time and allows flexibly for heterogeneity across properties belonging to the 100 different price bins. The third term captures how deviations

²⁰Specifically, we consider each property’s last transaction price prior to the Brexit referendum and determine the position of this price in the national distribution of transaction prices in the same quarter. Properties not transacted before the Brexit referendum fall out of the sample.

from this price dynamics correlate with offshore ownership. Specifically, β measures how much an additional 1% share of offshore ownership in the local area adds to the expected change in property prices since 2015. It is identified by comparing within-property price changes across local areas that belong to the same geographical region (i.e. London) and the same narrow price segment (i.e. percentile of the national distribution of transaction prices) but have different ex ante penetration of offshore ownership. The main objects of interest are the post-2015 elements in β , which we interpret as the effect on property prices of the decrease in offshore investors' demand for London real estate triggered by the outcome of the Brexit referendum. This causal interpretation assumes that price trends would have been parallel across areas with high and low offshore ownership absent the Brexit shock. While this assumption concerns a counterfactual and is therefore not directly testable, the pre-2015 elements in β informs us about its plausibility by showing how price trends correlated with offshore ownership before the Brexit referendum.

We address a number of remaining identification concerns by modifying and augmenting the baseline model. First, we add controls for the ex ante share of foreigners *living* in local areas to prevent that the estimated effect of offshore ownership is confounded by migration responses to Brexit. Second, we add a term capturing *domestic* corporate ownership to ascertain whether the estimated effect of offshore ownership reflects a broader difference between corporate and private owners. Third, we investigate whether the results are sensitive to the precise treatment of outliers, the granularity of the price controls, additional property-level controls and the geographical level at which offshore ownership is measured. We conduct all these additional tests in a compact framework where d_t is replaced with a simple post-2015 indicator.

Before presenting the results, we characterize the identifying variation in the baseline model: the pre-Brexit distribution of offshore ownership shares across the 983 local areas in London. Figure A8 in the Online Appendix illustrates this distribution for the 676 local areas with non-zero offshore ownership. In most of these areas, offshore ownership accounts for less than 2% of the local market, but in a handful of areas this metric exceeds 10%.

5.2 Results

Main Results

Figure 9 illustrates the results from the baseline model: the estimated percentage change in local property prices relative to the level in 2015 associated with an additional 1% ex ante market share of offshore owners. The estimates are negative and statistically significant for all years after the Brexit referendum, suggesting that offshore owners selling properties in the wake of the

referendum had a negative impact on prices. Consistent with this interpretation, the estimates are generally much smaller and far from statistical significance in all years before the Brexit referendum, which implies that local areas with different offshore ownership shares were generally on parallel trends in this period (conditional on the price controls).

The size of the estimates suggest that the effect is economically significant. The point estimate for 2018 is just around -2, suggesting that liquidations by offshore owners caused property prices to drop by around 2 percentage points over the period 2015-2018 for each additional 1% ex ante market share of offshore owners. This implies that prices were around 1.5% lower in 2018 in local areas at the 75th percentile of offshore ownership shares, around 5% lower at the 90th percentile, and around 7.5% lower at the 95th percentile.

Robustness

We conduct a series of robustness tests and report the results in Table 3. First, we re-estimate the baseline model in the compact framework to obtain a benchmark. The estimates suggest that an additional 1% ex ante market share of offshore owners lowered property prices by -1.91% between the pre-referendum period (2013-2015) and the post-referendum period (2016-2018), as shown in Column (1). Second, we add interactions between time dummies and a variable capturing the ex ante share of the local property market with corporate owners, domestic or foreign ones. The results indicate that a higher share of corporate ownership is not in itself associated with a differential price decrease, except when the corporate owners are offshore, as shown in Column (2). Third, we add interactions between time dummies and a variable expressing the ex ante share of foreigners in the local population interacted with time dummies. The results suggests that migration responses to Brexit may have depressed property prices; however, the estimated effect of offshore ownership barely changes, as shown in Column (3).²¹ Fourth, we winsorize offshore ownership shares to investigate the impact of extreme observations. Columns (4)-(5) show that the main estimate tends to be moderately larger when we winsorize. Fifth, we employ a more conservative clustering scheme. Standard errors are clustered at the level of individual properties in the baseline, but Column (6) shows that the estimates remain highly significant when moving the clustering at the level of local areas. Finally, we add interactions between time dummies and two additional property characteristics: tenure type (e.g. freehold, leasehold) and property type (e.g. detached, semi-detached, terraced). In these specifications, the main estimate is somewhat lower, but remains highly statistically significant (t-statistic > 5).

²¹Information on the country of origin of the population is available from the office of national statistics at the level of 33 districts within London. We therefore re-calculate the offshore market share at the level of these districts and normalize both measures to have comparable variation.

Mechanism

Our interpretation of these results is that Brexit, by making it less attractive for foreigners to own properties in the United Kingdom, constituted a negative shock to offshore ownership, which depressed property prices. A key implication of this mechanism is that the price effect was larger in areas where the beneficial owners were foreigners and smaller where they were locals. Testing this implication is not straightforward, as we only observe beneficial owners in the subsample of offshore corporations matched to the leaks from tax havens (see section 2.3). However, under the assumption that selection into the leaks is random within each offshore jurisdiction (Alstadsæter et al., 2019), we can use the leaks to assign probabilities of foreign and local beneficial ownership in the rest of the sample. For instance, we observe that the beneficial owners are foreign for around 50% of the Jersey corporations and around 95% of the Singapore corporations in the subsample that is in the leaks; hence, assuming random selection, the estimated probability of foreign ownership is 50% for Jersey corporations and 95% for Singapore corporations in the subsample that is not in the leaks.

Using these estimated probabilities, we re-estimate the baseline model while decomposing the offshore ownership share into expected foreign and expected domestic beneficial ownership. The results are reported in Table 4, Column (2). Consistent with our hypothesis mechanism, the estimated price effects of offshore ownership primarily come from properties with foreign beneficial owners with a much smaller contribution from properties with local beneficial owners. However, the estimates are imprecise, notably the contribution from local beneficial owners, and they are not statistically distinguishable.²²

Spill-Over effects

Finally, we ask whether the price effects only affect the specific market segments with significant offshore ownership, largely high-end properties with presumably wealthy owners, or whether the price effects spill over on middle-class households with less expensive properties. To address this question, we construct separate offshore market shares for two market segments within each local area: properties with transaction prices above and below £500,000 respectively (before the Brexit referendum). We then re-estimate the baseline model while restricting the sample to properties below £500,000 and including both of the price-specific offshore market shares in the model.

The results are reported in Table 4, Column (3). Consistent with the mechanism explored

²²The reported standard errors should be interpreted with some caution, as they do not account for the fact that the two dependent variables are themselves estimates associated with uncertainty.

above, the estimates show that a higher ex ante offshore market share in the low-price market was associated with larger price decreases for low-price properties after the Brexit referendum with a semi-elasticity of around -2.2. However, the estimates also show that a higher ex ante offshore market share in the high-price market was associated with larger price decreases for low-price properties with a slightly smaller semi-elasticity of around -2. This is consistent with spill-overs across different price segments and suggest that large parts of the population are exposed to the price volatility created by property investment through offshore structures.

6 Conclusion

Our descriptive analysis documents that the market share of offshore corporations in the property market has increased dramatically over the two most recent decades. It currently stands around 1.25% in the overall residential market - and around 15% for the most expensive properties, which implies that money invested in exclusive properties in the U.K. is more likely to come from tax havens than money sitting in U.K bank accounts. We estimate that individuals in Africa, Asia and the Middle East are the beneficial owners of around half of the properties with nominal offshore ownership; however, the United Kingdom itself is the largest beneficial owner suggesting an important role for round-tripping investment.

Turning to questions about causes and consequences, we show that offshore ownership in the property market is highly responsive to changes in capital gains taxes as well as changes in ownership transparency. This is suggestive that both taxation and secrecy are important motives for funnelling property investments through offshore tax havens. We also show that the Brexit referendum triggered a sharp increase in property sales by offshore owners, which was accompanied by differential price decreases in areas with higher ex ante offshore ownership rates. The analysis suggests that offshore ownership has significant consequences for real outcomes in housing markets and that spill-overs from the top-end markets where offshore ownership is most pervasive may create exposure for middle-class home owners in other markets.

References

- ALSTADSÆTER, A., N. JOHANNESSEN and G. ZUCMAN (2018). “Who Owns the Wealth in Tax Havens? Macro Evidence and Implications for Global Inequality”. *Journal of Public Economics* 162, pp. 89–100.
- (2019). “Tax Evasion and Inequality”. *American Economic Review* 109 (6), pp. 2073–2103.
- ALSTADSÆTER, A., G. ZUCMAN, B. PLANTEROSE and A. ØKLAND (2022). “Who Owns Offshore Real Estate? Evidence from Dubai”. EU Tax Observatory Working Paper 1.
- ATLANTIC (2022). “Seize the Oligarchs’ Wealth”. 26 February 2022.
- BADARINZA, C. and T. RAMADORAI (2018). “Home Away from Home? Foreign Demand and London House Prices”. *Journal of Financial Economics* 130 (3), pp. 532–555.
- BBC (2016). “Stock Markets Tumble after Leave Vote”. 24 June 2016.
- BOMARE, J. (2019). “Tax Evasion and British Real Estate: An Estimation of the Offshore Wealth Held through British Real Estate”. Paris School of Economics Master Thesis.
- BOMARE, J. and S. LE GUERN HERRY (2022). “Will We Ever Be Able to Track Offshore Wealth? Evidence From the Offshore Real Estate Market in the UK”. Sciences Po Economics Discussion Paper 2022-10.
- BOURNE, J., A. INGIANNI and R. MCKENZIE (2022). “What’s in the Laundromat? Mapping and Characterising Offshore Owned Domestic Property in London”. arXiv Preprint.
- COLLIN, M. (2021). “What Lies beneath: Evidence from Leaked Account Data on How Elites Use Offshore Banking”. Brookings Global Working Paper 156.
- COLLIN, M., F. M. HOLLENBACH and D. SZAKONYI (2022). “The Impact of Beneficial Ownership Transparency on Illicit Purchases of U.S. Property”. Brookings Global Working Paper 170.
- COPPOLA, A., M. MAGGIORI, B. NEIMAN and J. SCHREGER (2021). “Redrawing the Map of Global Capital Flows: The Role of Cross-Border Financing and Tax Havens”. *Quarterly Journal of Economics* 136 (3), pp. 1499–1556.
- DAMGAARD, J., T. ELKJAER and N. JOHANNESSEN (2019). “What Is Real and What Is Not in the Global FDI Network?” IMF Working Paper 19/274.
- GORBACK, C. S. and B. J. KEYS (2020). “Global Capital and Local Assets: House Prices, Quantities, and Elasticities”. NBER Working Paper 27370.

- GRAVELLE, J. G. (2015). “Tax Havens: International Tax Avoidance and Evasion”. Congressional Research Service 7-5700.
- GUARDIAN (2016a). “House Prices Could Fall by 18% If Britain Quits EU, Says George Osborne”. 20 May 2016.
- (2016b). “Brexit Could Cut London House Prices by More than 30%, Says Bank”. 18 July 2016.
- (2022). “How London Became the Place to Be for Putin’s Oligarchs”. 6 March 2022.
- GUYTON, J., P. LANGETIEG, D. RECK, M. RISCH and J. SLEMRD (2022). “The Offshore World According to FATCA: New Evidence on the Foreign Wealth of U.S. Households”. Working Paper.
- GUYTON, J., P. LANGETIEG, D. RECK, M. RISCH and G. ZUCMAN (2021). “Tax Evasion at the Top of the Income Distribution: Theory and Evidence”. NBER Working Paper 28542.
- HANLON, M., E. L. MAYDEW and J. R. THORNOCK (2015). “Taking the Long Way Home: US Tax Evasion and Offshore Investments in US Equity and Debt Markets”. *Journal of Finance* 70 (1), pp. 257–287.
- JOHANNESSEN, N. (2014). “Tax Evasion and Swiss Bank Deposits”. *Journal of Public Economics* 111, pp. 46–62.
- JOHANNESSEN, N. and G. ZUCMAN (2014). “The End of Bank Secrecy? An Evaluation of the G20 Tax Haven Crackdown”. *American Economic Journal: Economic Policy* 6 (1), pp. 65–91.
- KNOLL, K., M. SCHULARICK and T. STEGER (2017). “No Price like Home: Global House Prices, 1870–2012”. *American Economic Review* 107 (2), pp. 331–53.
- LEENDERS, W., A. M. LEJOUR, S. RABATÉ and M. van ’T RIET (2022). “Offshore Tax Evasion and Wealth Inequality: Evidence from a Tax Amnesty in the Netherlands”. *Journal of Public Economics* (forthcoming 2022).
- LONDOÑO-VÉLEZ, J. and J. ÁVILA-MAHECHA (2021). “Enforcing Wealth Taxes in the Developing World: Quasi-Experimental Evidence from Colombia”. *American Economic Review: Insights* 3 (2), pp. 131–148.
- LONDOÑO-VÉLEZ, J. and D. TORTAROLO (2022). “Revealing 21% of GDP in Hidden Assets: Evidence from Argentina’s Tax Amnesties”. UNU WIDER Working Paper 103/2022.
- MENKHOFF, L. and J. MIETHE (2019). “Tax Evasion in New Disguise? Examining Tax Havens’ International Bank Deposits”. *Journal of Public Economics* 176, pp. 53–78.

SHAXSON, N. (2013). “A Tale of Two Londons”. April 2013.

TØRSLØV, T. R., L. S. WIER and G. ZUCMAN (2022). “The Missing Profits of Nations”. *Review of Economic Studies* (forthcoming 2022).

ZUCMAN, G. (2013). “The Missing Wealth of Nations: Are Europe and the U.S. Net Debtors or Net Creditors?” *Quarterly Journal of Economics* 128 (3), pp. 1321–1364.

Table 1 – Data Sources

	(1)	(2)
	All properties	Residential
(A) Total with corporate ownership	2,469,283	501,286
- Commercial and Corporate Ownership Data (CCOD)	2,283,212	471,522
- Overseas Company Ownership Data (OCOD)	95,847	21,503
- Freedom of Information (FOI)	90,224	8,261
(B) Offshore corporate ownership	302,532	34,613
- Direct owner is corporation in tax haven (OCOD + FOI)	181,463	27,375
- Indirect owner is corporation in tax haven (Orbis)	121,069	7,238
(C) Ultimate owners (natural persons)	7,310	1,918
- Pandora Papers	1,086	283
- Paradise Papers	1,211	348
- Panama Papers	3,817	968
- Offshore Leaks	306	77
- Bahamas Leaks	890	242

Notes: This table shows the number of post-able addresses after the data preparation effort described in the main text. Our data is conservatively limited to addresses that have an identifiable postcode, street, and house number. They can, however, have more information about units, unit numbers, or potentially numbered localities. Panel A shows first the number of addresses we identify in the corporate ownership data. In the second column, it shows the number of addresses matched with the price paid data while maintaining a conservative timing assumption of at most 4 month distance between purchase (the price paid date) and registration (the corporate register date) of a property. Panel B shows information on offshore ownership: The first column collects all properties for which we observe an offshore corporate owner as the immediate investor (OCOD + FOI) and as indirect investors through the Orbis match outlined in the main text (Orbis). Column 2 again shows the number of addresses matched with the price paid data. Finally, Panel C summarizes the match with the large data leaks. When reducing to natural persons and to addresses that fulfil the above criteria, column 1 shows the addresses to which we can assign at least one natural person. Column 2 again shows this number in the matched data.

Table 2 – Aggregate Ownership Statistics (Dec. 2019)

(A) All residential properties in England and Wales (£ billion)	4,027.91
(B) Owned directly or indirectly by offshore corporations (£ billion)	49.8
- of which: offshore corporation is immediate owner	46.1
- of which: offshore corporation is ultimate holding company	3.8
Offshore market share (= B/A)	1.24%
(C) Owned directly or indirectly by other foreign corporation (£ billion)	9.9
- of which: foreign corporation is immediate owner	3.6
- of which: foreign corporation is ultimate holding company	6.3
Other foreign corporate market share (= C/A)	0.25%
(D) Owned exclusively by domestic corporations (£ billion)	184.0
Domestic corporate market share (= D/A)	4.57%

Notes: This table shows aggregate ownership statistics of the stock value of the residential real estate market as of December 2019. The total residential real estate stock value (A) is split up between the value owned by offshore tax haven corporations (B), other foreign corporations (C), and domestic corporations. Both foreign markets are further differentiated by direct ownership based on the ownership information in the land register, and indirect ownership through a UK corporation visible in ORBIS. For details on the ownership assignment and the estimation of market shares, see Appendix B and Appendix C.

Table 3 – Offshore Ownership and Property Prices, Robustness

	(1) Baseline	(2) Additional area controls	(3) -1.85*** (0.216)	(4) Alternative winsorization -2.34*** (0.266)	(5) -2.94*** (0.376)	(6) Alternative clustering -1.91*** (0.376)	(7) Alternative property controls -1.16*** (0.162)	(8) -1.26*** (0.221)
Post × Offshore	-1.91*** (0.216)	-1.91*** (0.221)	-1.85*** (0.216)	-2.34*** (0.266)	-2.94*** (0.376)	-1.91*** (0.376)	-1.16*** (0.162)	-1.26*** (0.221)
Post × Corporate		-0.013 (0.109)						
Post × Foreign Population			-0.039** (0.019)					
Property FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
100 price bins × year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property type × year FE							Yes	
Tenure type × year FE								Yes
Clustering	Property	Property	Property	Property	Property	983 areas	Property	Property
Winsorization				2x0.5%	2x2%			
Observations	99,565	99,565	99,565	99,197	97,395	99,565	97,692	99,565
Adjusted R ²	0.9356	0.9356	0.9356	0.9353	0.9355	0.9356	0.9788	0.9366


Notes: This table shows robustness exercises for our baseline specification. The sample is limited to 3 years before and after Brexit (2013:2018) with the difference-in-differences coefficient taking value 1 for years 2016:2018 and 0 before that. Column 1 shows the baseline effect interacting the did dummy with the offshore penetration measure introduced in the main text. Column 2 adds the total corporate market share as a control variable (domestic and foreign corporations). Column 3 introduces the foreign population share as a control variable defined at the level of the local authority area. Columns 4 winsorizes the sample eliminating the top and bottom 0.5% of our offshore market share. Column 5 does the same for the top and bottom 2%. Column 6 moves the clustering to the level of the treatment variation, namely the 983 middle super output areas in London at which we calculate the offshore market share. Column 7 adds property type times year effects as additional property level controls (Detached, Semi-Detached, Terraced, Flats/Maisonettes). Properties characterized as “other” are excluded from the sample here due to a reporting change in the sample period. Column 8 instead introduces tenure type times year fixed effects (freeholds and leaseholds). Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Table 4 – Offshore Ownership and Property Prices, Mechanism and Spill-overs

	(1) Baseline	(2) Mechanism	(3) Spill-overs
Post x Offshore	-1.91*** (0.216)		
Post x Offshore, Expected British		-1.01 (2.01)	
Post x Offshore, Expected Foreign		-2.10*** (0.545)	
Post x Offshore, Low Price			-2.23*** (0.610)
Post x Offshore, High Price			-1.98*** (0.257)
	99,565	99,565	68,868
Property FE	Yes	Yes	Yes
100 price bins x year FW	Yes	Yes	Yes
observations	99,565	99,565	68,868
Adjusted R ²	0.9356	0.9356	0.9421

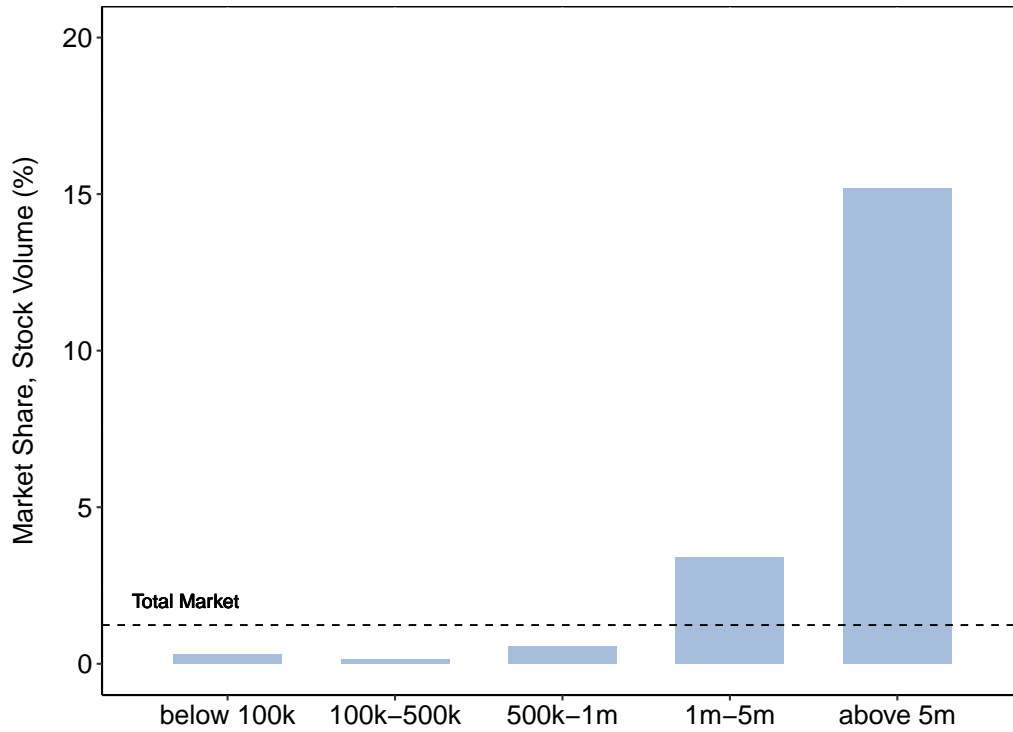
Notes: This table shows robustness exercises for our baseline specification. The sample is limited to 3 years before and after Brexit (2013:2018) with the difference-in-differences coefficient taking value 1 for years 2016:2018 and 0 before that. Column 1 shows the baseline effect interacting the DiD dummy with the offshore market share. Column 2 shows the effect of true foreign investment versus British round-tripping. Based on the ultimate owners found using the large data leaks, we assign to each investing country a probability that the ultimate owners is British or a true foreigner. When constructing the offshore market share, the price of a property held from that country is multiplied with this probability. We assign a third category for countries that have either no British or no non-British ultimate owners. This category is omitted from the table but included in the regression. Column 3 shows spillovers across price segments. We calculate two offshore market shares: One using only properties with a pre-Brexit price below £500,000 (Low Price) and one using those above (High Price). The sample is limited to properties with a pre-Brexit price below £500,000 to show price effects on this market. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Figure 1 – Offshore Market Share by Investing Country (Dec. 2019)



Notes: The figure shows the stock volume of residential real estate that is held from foreign corporations in different jurisdictions as of December 2019. Jurisdictions are differentiated between tax havens (blue bars) and non havens (green bars). Stock volumes both include direct ownership (dark color) and indirect ownership links through a UK corporation (light color). For details on the ownership assignment and the estimation of stock volumes, see Appendix B and Appendix C.


Figure 2 – Offshore Market Share by Price Segment (Dec. 2019)




Notes: The figure shows the offshore tax haven market share in the residential real estate market across different price segments as of December 2019. Ownership from offshore tax havens both includes direct ownership and indirect ownership links through a UK corporation. Allocation into price segments are based on property values as of December 2019. The dashed line indicates the offshore tax haven market share across all price segments. For details on the ownership assignment and the estimation of market shares, see Appendix B and Appendix C.

Figure 3 – Geography of Offshore Market Shares (Dec. 2019)

(a) England and Wales




(b) London




Notes: The figure shows the geographical dispersion of the offshore tax haven market share in the residential real estate market as of December 2019. Ownership from offshore tax havens both includes direct ownership and indirect ownership through a UK corporation. Figure 3a displays market shares across local authority districts (LAD) in England and Wales while Figure 3b displays market shares across Middle Layer Super Output Areas (MSOA) in London. For details on the ownership assignment and the estimation of market shares, see Appendix B and Appendix C.

Figure 4 – Offshore Tax Haven Market Share over Time




Notes: The figure shows the offshore tax haven market share in the residential real estate market between January 1995 and December 2019. Ownership from offshore tax havens both includes direct ownership and indirect ownership links through a UK corporation. The dashed line indicates the start of the real time snapshots for direct ownership of offshore companies (OCOD). For details on the ownership assignment and the estimation of market shares, see Appendix B and Appendix C.

Figure 5 – Beneficial Ownership by Home Country




Notes: These figures shows the 20 most important countries from our ultimate ownership match. The sample is reduced to properties which we could match to corporate owners and their ultimate owners. The vertical axis therefore shows the share of each country in percent of the total “ultimate owner market”. To determine when the position of countries can be statistically differentiated from one another, we draw 1000 bootstrap samples from this sample. For each bootstrap sample, we calculate the market share of ultimate owners from each country. From the distribution of these 1000 market shares for each country, we then plot their mean as the gray bars. The error bars in both plots the 25th percentile and the 75th percentile.

Figure 6 – The Tax Motive




Notes: The figure plots the coefficients from the baseline specification of the tax shock that compares the evolution of property holdings held by companies in Luxembourg to property holdings held by companies in other tax havens. Confidence bands show heteroskedasticity and autocorrelation robust standard errors around the point estimate. We normalize effects in October 2017, the month before the extension of CGT coverage was announced (first dashed line). The second dashed line refers to the implementation of the CGT extension. For details on the political process of the CGT extension, see Appendix D.

Figure 7 – The Secrecy Motive




Notes: The figure plots the coefficients from the baseline specification of the secrecy shock that compares the evolution of property holdings held by companies in the Overseas Territories to property holdings held by companies in the Crown Dependencies. Confidence bands show heteroskedasticity and autocorrelation robust standard errors around the point estimate. We normalize effects in March 2018, two months before the SAML A amendment passed parliament (first dashed line). The second dashed line refers to the announcement of Crown Dependencies to voluntarily establish public ownership registries. For details on the political process of the SAML A amendment, see Appendix E.

Figure 8 – Flight of Offshore Capital Triggered by Brexit.



Notes: This figure plots the result of two separate regressions of the following type: $y_{it} = \mu_i + \sum_t \beta d_t + \varepsilon_{it}$. For the solid black line, y_{it} is a dummy variable taking value 1 if a transaction of property i at time t was sold from an offshore seller to a domestic owner, zero elsewhere. We plot the β coefficients and associated 95% error bounds of the time dummies excluding the pre-Brexit year 2015. Property fixed effects are included in μ_i , the vertical axis therefore shows the relative selling probability compared to 2015. The dotted line replaces the outcome variable y_{it} with a dummy taking value 1 if a property was sold from a domestic owners to an offshore buyer, zero elsewhere.

Figure 9 – Offshore Ownership and Property Prices



Notes: This figure plots the coefficients from the baseline specification. We control for 100 national price bins interacted with year dummies to ensure that our effects are not driven by the differential investment patterns of foreign companies outlined in the descriptive section. Our variable of interest is the offshore penetration, calculated in January 2016, before Brexit and plotted in figure A8. We normalize effects in the year before Brexit (2015). The Brexit referendum was announced in February 2016 and carried out in June 2016. The sample is limited to properties sold both before and after Brexit. We therefore use within property variation and compare developments across the 983 administrative middle super output areas.


Online Appendix

“Homes Incorporated: Offshore
Ownership of Real Estate in the U.K. ”


A Extra Results

Figure A1 – Stock Value of Foreign Residential Real Estate

(a) Current Values




(b) CPI Adjusted (Dec. 2019)




Notes: This figure displays the stock volume of residential real estate in England and Wales held by foreign corporate entities. Ownership from offshore tax havens both includes direct ownership and indirect ownership links through a UK corporation. Figure A1a displays the stock volume in current prices while Figure A1b displays the series adjusted for CPI-Inflation. For details on the ownership assignment and the estimation of market shares, see Appendix B and Appendix C.

Figure A2 – Validation




Notes: This figure displays results from two validation exercises that verify our strategy for the estimation of long-run time series of real estate volumes. Figure A2a compares the total stock value of residential real estate in England and Wales according to our estimation (black line) with adjusted values from the National Accounts (red points). Figure A2b compares imputed (y-axis) and actual (x-axis) month-to-month growth rates of the offshore tax haven market share for the time between October 2015 and December 2019 when reliable snapshot data is available. For details on the estimation of market shares and the validation exercises, see Appendix C.

Figure A3 – Tax Haven Share - British Real Estate vs. British Bank Deposits



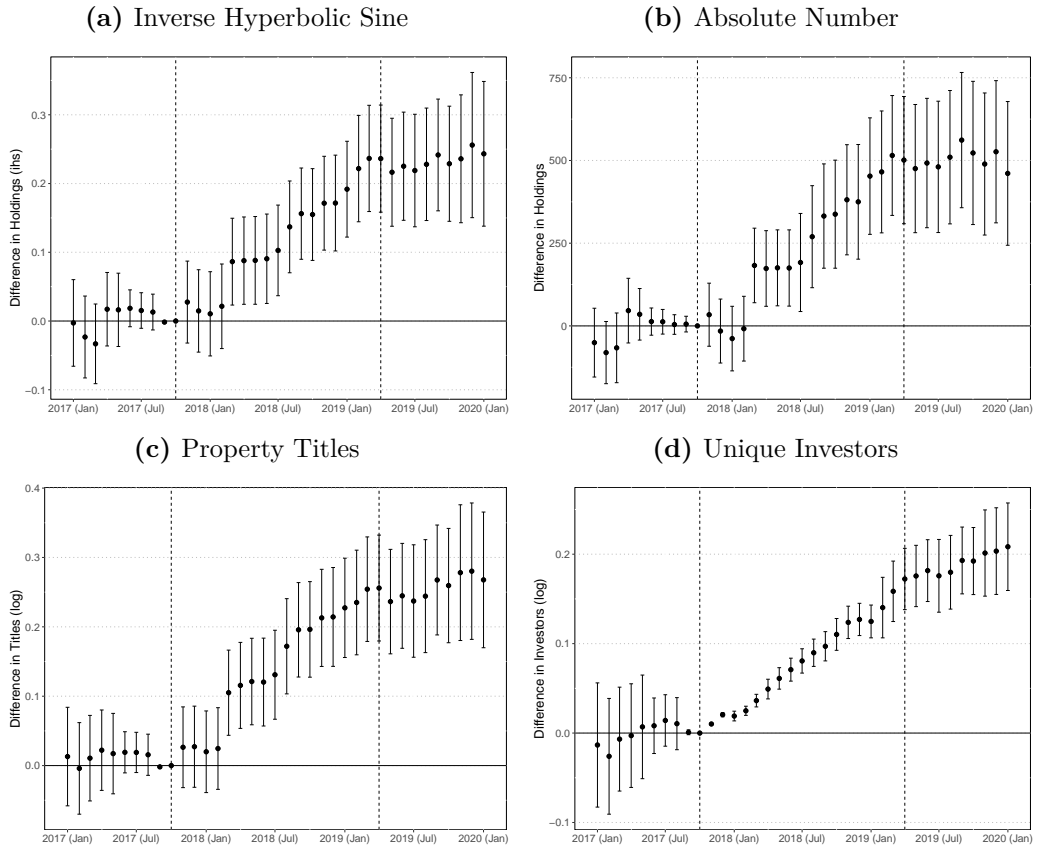
Notes: The figure compares the development of offshore held real estate in % of the total residential market (solid line, left horizontal axis) to the tax haven deposit share in British banks (dashed line, right horizontal axis). The numerator of the latter adds the non-bank deposits from tax havens in British banks. The denominator adds all counterparties (non-havens + tax haven non-banks) but also adds domestic deposits which are available through the BIS. The series can therefore be interpreted as the tax haven share in % of total bank deposits in Britain.

Figure A4 – Beneficial Ownership by Region - British Real Estate vs. Swiss Private Banking



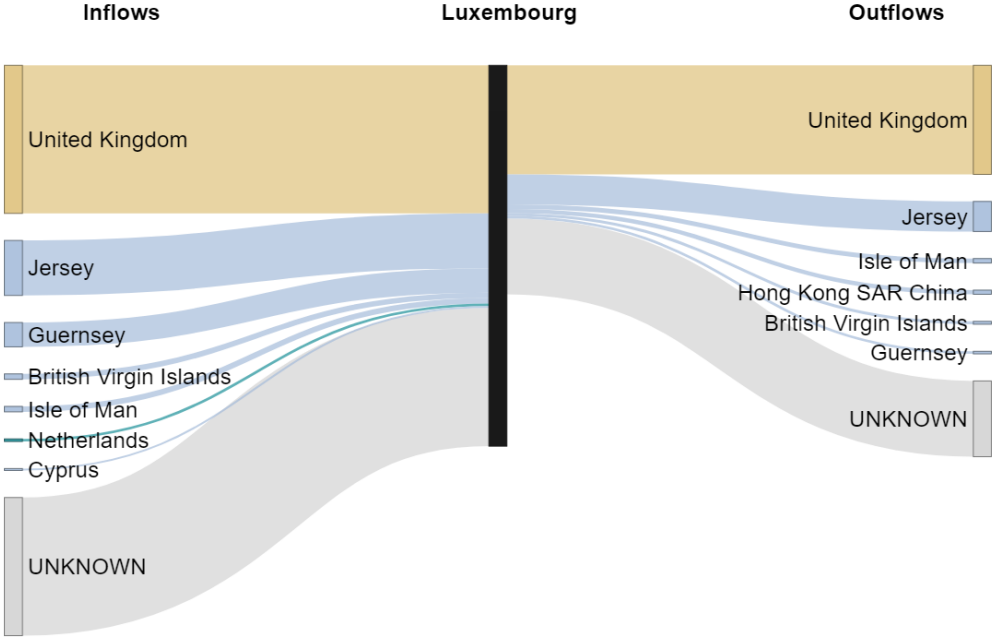
Notes: The figure compares the distribution of beneficial ownership across geographical regions for British real estate (dark gray bars) and for financial assets managed in Switzerland (light gray bars). The latter set of bars is based on the leaked customer records from HSBC Switzerland published by the ICIJ as the *Swiss Leaks*.

Figure A5 – Tax Shock - Robustness



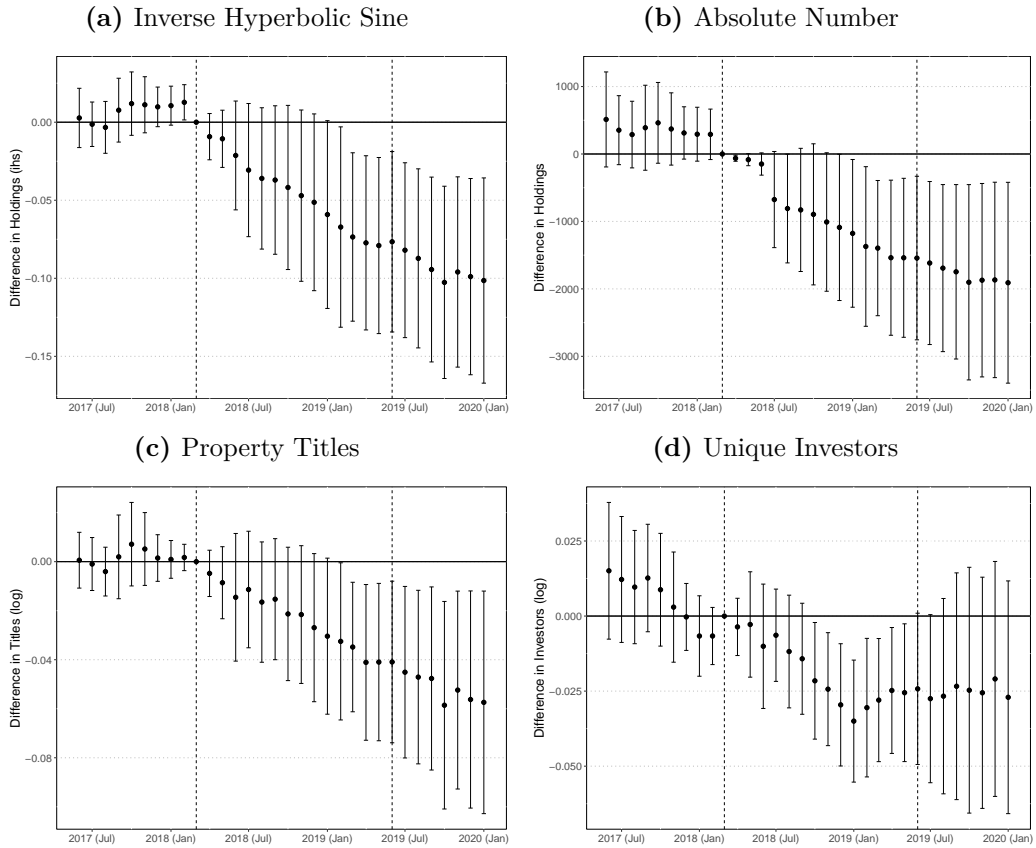
Notes: The first row shows the results for the tax shock (analogous to Figure 6) with an alternative transformation of the outcome variable, the inverse hyperbolic sine transformation (Panel A), and without any transformation (Panel B). The second row shows the results with alternative units of observations: the number of property titles (Panel C) and the number of unique investors (Panel D).

Figure A6 – Tax Shock - Flows (Oct. 2017 - Apr. 2019)




Notes: The figure visualizes the flows of property holdings to (from) companies incorporated in Luxembourg between October 2017 and April 2019. Country assignments in flows to (from) Luxembourg are based on the owner that precedes (succeeds) the Luxembourg corporation within at most a time window of 6 months. If no direct predecessor (successor) is visible in the data, we assign this predecessor to the “Unknown” category that consists of private owners, deconstructed properties and newly built properties.

Figure A7 – Secrecy Shock - Robustness



Notes: The first row shows the results for the secrecy shock (analogous to Figure 7) with an alternative transformation of the outcome variable, the inverse hyperbolic sine transformation (Panel A), and without any transformation (Panel B). The second row shows the results with alternative units of observations: the number of property titles (Panel C) and the number of unique investors (Panel D).

Figure A8 – Distribution of Offshore Ownership Shares in London



Notes: The figure shows the distribution of offshore ownership across the 983 super middle output areas in London (omitting the 307 areas with no offshore ownership).

B Combining Data Sources

This part of the appendix provides details on the combination of price, purchase date, and ownership information to create a comprehensive dataset of British real estate transactions. In Section B.1, we describe in detail, how address string information from various sources is harmonized. Section B.2 discusses, how ownership information from the land register is combined with information about ownership chains from ORBIS and ultimate ownership details from the ICIJ leak data. Finally, we provide insights on the conditions that are utilized to match price and ownership information (Section B.3).

B.1 Preparing British Addresses

Understanding the relative importance of foreign held residential real estate requires information that is not readily available from British sources. The Price Paid data records sales price information for the entire residential market, the different ownership datasets contain the immediate ownership information for the entire corporate market. Our goal is to match these sources to obtain combined price-ownership information for the residential market that is owned by corporate investors. Unfortunately, the UK registers do not publish identifiers linking the Price Paid data and the ownership data. Even if these were published, they would most likely link the data specific identifiers which are of limited use. The main identifier in the Price Paid data is a “transaction id” that can include many properties sold in one transaction. The main identifier in the ownership data is the “land title registration number” that can include dozens of addresses that have the same owner and are registered as one land title. To be able to combine the data, we move all datasets to the level of a postal address. An example for a simple data entry is:

“FLAT 5, 213 SUSSEX GARDENS, LONDON (W2 2RJ)”

This entry can easily be moved to the postal address level:

Postcode	Street	House No.	Unit	Unit No.	Locality	Locality No.
W2 2RJ	SUSSEX GARDENS	213	FLAT	5		

An example for a complicated data entry is:

“FLATS 1-27 WALLACE COURT, 54 TIZZARD GROVE, LONDON (SE3 9EE), FLATS 103-128 WALLACE COURT, 44 TIZZARD GROVE, LONDON (SE3 9EQ) AND FLATS 129-157 WALLACE COURT, 52 TIZZARD GROVE, LONDON (SE3 9FE), 1-48 GRAYSTON HOUSE, 21 ASTELL ROAD LONDON (SE3 9FN), 49-110 GRAYSTON HOUSE, 1 OTTLEY DRIVE, LONDON (SE3 9FP), 1-62 MALTBY HOUSE, 2 OTTLEY DRIVE,

LONDON (SE3 9FJ), 63-105 MALTBY HOUSE, 18 TUDWAY ROAD, LONDON (SE3 9FL), 5-12 OTTLEY DRIVE, LONDON (SE3 9FT), 2-16 (EVEN) TUDWAY ROAD, LONDON (SE3 9FR)”

This entry is separated by our code into multiple postal addresses:

Postcode	Street	House No.	Unit	Unit No.	Locality	Locality No.
SE3 9EE	TIZZARD GROVE	54	FLAT	1	WALLACE COURT	
SE3 9EE	TIZZARD GROVE	54	FLAT	2	WALLACE COURT	
...
SE3 9EE	TIZZARD GROVE	54	FLAT	27	WALLACE COURT	
SE3 9EQ	TIZZARD GROVE	44	FLAT	103	WALLACE COURT	
...
SE3 9FP	OTTLEY DRIVE	1			GRAYSTON HOUSE	94
...


This process is carried out for the following datasets:

- the Price Paid database that collects all residential transactions
- the CCOD database that contains all domestic corporate investment
- the historical CCOD database
- the OCOD database that contains all foreign corporate investment
- the historical OCOD database
- the freedom of information request data containing foreign transactions since 1990
- the Private Eye data containing a 2014 snapshot of foreign ownership

Detecting Postcodes

Our code starts with basic string cleaning exercises such as: Moving everything to uppercase letters, adjusting typos such as “Esat” instead of “East”, eliminating apostrophes, aligning cardinal points (“Notherly = North”, or “North-West = North West”), aligning punctuation, removing trailing or leading whitespaces as well as unnecessary punctuation.

We then start the address preparation by separating data entries along the information on postcodes. We exploit the fact that postcode assignment in the UK follows regularized patterns and contains different layers which are shown in the following exemplary postcode:



Given that the different layers allow for a limited number of patterns, we can use regular expressions to identify the postcode. Since postcodes usually appear at the end of an address entry, any auxiliary information that appears after the postcode is dropped.²³ In very rare cases, there is only one postcode for a land title that includes several streets and it is only specified for the first street. If we find any additional street information after a postcode, we assign the same postcode.²⁴

Detecting Streets

Having separated addresses by postcodes, we next turn to street names. Since a data entry can contain multiple streets under the same postcode, we first use regular expressions to identify, how many different streets are included in the entry of the land registry.²⁵

Depending on the number of matches, we then repeatedly extract streets from the data entry following a two step procedure. The first step makes use of the list of street names provided by the national mapping agency of the UK.²⁶ In particular, we extract streets in the data entry by looking through all street names associated with the respective outward postcode. The identification of street names via location-specific lists has the advantage that it reduces false positives associated with terms that can be a street name in one place, but are a locality or property name in others. In a second step, if we have not identified a street using the lists above, we extract remaining street names based on street regular expressions using general terms commonly associated with streets.

For each extracted address, we also obtain a leading string that contains house numbers, unit information (e.g. “FLAT 10”, “APARTMENT 5”), locational information (e.g. “LAND AT”, “BUILDINGS LYING TO THE SOUTH OF”), and locality information (e.g. “WALLACE COURT”). In the next sections, we discuss how we separate these elements.

Separating House Numbers, Units, and Localities

In the data, house number information is placed directly in front of the street while unit and locality information is located in front of the house number (e.g. “FLAT 7, DORCHESTER

²³For instance, this auxiliary information includes strings like “and adjacent parking lot” that are not relevant for the preparation of the address.

²⁴In particular, we look for the appearance of any of the following street regular expressions: STREET, ROAD, ARCADE, VALE, LANE, AVENUE, CLOSE, WHARF, ACRE, MEWS, WALLS, CHASE, SQUARE, CRESCENT, PLACE, ROW, TERRACE, WAY, QUAY, BOULEVARD, GATE, CLOSE, DRIVE, GROVE, GARDENS, GARDEN, CIRCUS, WALK, EMBANKMENT, HEIGHTS, GARTH, PARKWAY, COURT, YARD, VIEW, HILL, PARADE, APPROACH, RISE, PATH, GLADE, and TERRACE.

²⁵We use the same list of street pointers used in the postcode preparation above.

²⁶These data are available from the Ordnance Survey.

COURT, 77-81 SLOANE STREET, LONDON (SW1X 9SE)”). This order allows us to identify the house number as a regular expression that directly precedes the street name and contains only non-alphabetical strings.²⁷

Once the house number is extracted, we detect units and their corresponding number(s) by searching the remaining entry for elements of an extensive list of unit types.²⁸ The non-alphabetical regular expressions following the extracted unit types are recorded as the respective unit number(s). Again, unit numbers are allowed to follow the same patterns as house numbers.

The land registers are full of special property names such as “Wallace Court” or “Grayston House”. We opt for treating these localities different from the units described above, though they often fulfill the same role in identifying addresses. The most important difference is that unlike flats or apartments, these localities often appear without a number associated with them. We detect localities using regular expressions which is simpler than detecting units after the latter have been taken out of the running string.

Treating House, Unit and Locality Numbers

Having identified the strings that include information about the house, unit and locality numbers, we transform them into a harmonized numeric appearance. For instance, we create integer sequences of simple numeric sequences (“1-3” becomes “1,2,3”). Thereby, we pay special attention to ranges of numbers that also include letters (e.g. “3A-4C”) and make sure that all combinations (e.g. “3A, 3B, 3C, 4A, 4B, 4C”) are created for the final data.

Whenever one of the above steps leads to more than one postcode, street, house number, unit number, or locality number, we reshape the dataset into long format. This process results in a dataset of postal addresses that are uniquely identified. Whenever data quality concerns prohibit precise recognition, the entire land register entry is dropped from the dataset.

Comparison to Alternative Address Classification Methods

The classification of address information is a challenge that is also discussed in the field of computer science. For instance, Bourne et al. (2022) apply a very general model of natural language processing (spaCy) to structure the address string snippets from the UK land register. Such models have the advantage that they can be applied very broadly in various situations and

²⁷We allow for one-digit alphabetical strings attached to numbers and for particular alphabetical strings such as “ODD” or “EVEN”. Thereby, our code is also able to identify more complicated house number combinations like “1-11 (ODDS)”, “16-20 (EVEN)”, or “14, 17A-18C”.

²⁸Among others, unit types include “APARTMENT”, “FLAT”, “UNIT”, “BUILDINGS”, “BUNGALOWS”, “SUITE”, “ROOM”, “FLOOR”.

that they require less context-specific manual cleaning of the data.

We decided to not follow a machine learning approach in our address classification, because the address information in the land register is very unstructured. This implies that an algorithm like the spaCy model that is trained with very clean data can produce systematic errors. For instance, in the address preparation by Bourne et al. (2022), streets can be classified as units, and units are not differentiated from localities. This can be particularly problematic, if information needs to be connected across different data sources that do not necessarily record information in a harmonized way. By manually inspecting the different data sources (FOI, PPD, OCOD, CCOD), our approach allows a characterization that is tailored to capture the heterogeneity in the address recordings in our various data sources. In addition, our classification is very transparent while machine learning algorithms are often a black box, whose classifications can change in an unpredictable way when the pre-trained model is updated.

B.2 Preparing Ownership Information

HM Land Registry differentiates its registers for commercial real estate between real estate titles held by domestic companies (CCOD) and titles held by foreign companies (OCOD).²⁹

The Commercial and Corporate Ownership Data (CCOD) contains all freehold and leasehold registered titles to land or property in England and Wales which are held by non-private entities incorporated in the UK. Examples of non-private entities include public limited companies, unlimited companies and limited liability partnerships.³⁰ While land registration in the UK exists for more than 150 years, is only publicly available and updated on a monthly basis from November 2017 onward. In addition, using the legal procedure of a Freedom of Information (FOI) request, we obtained quarterly data for the time period between March 2014 and September 2017. The latest version of the data contains more than three million property titles including both property-related information (e.g. geographical information, tenure) and owner-related information (e.g. name, company registration number, proprietor type).³¹

The Overseas Company Ownership Data (OCOD) is the corresponding administrative register for property titles registered to companies outside of the UK. OCOD was made publicly available in 2017 as part of the government’s strategy to enhance the understanding of the UK housing.³²

²⁹Detailed information on CCOD and OCOD is provided by HM Land Registry.

³⁰The data does exclude Charitable Incorporated Organizations. Land register data for private individuals is also maintained by HM Land Registry, but not available for the public.

³¹For some properties, the land register also contains information about the price paid. However, since the recordings of the prices is incomplete, we ignore the price information from the land register, since we have comprehensive price information in the transaction data (Price Paid data).

³²This process had been pushed by a data publication of the current affairs magazine *Private Eye* in 2014. It

Apart from information recorded in the CCOD, the OCOD contains the country of incorporation of the proprietor(s). Although the country of incorporation has been routinely recorded in the register since 1999, OCOD is also only available from November 2017. In the process of our FOI request, we could extend the coverage obtaining quarterly OCOD data which dates back to October 2015. As of June 2020, OCOD counts around 100,000 property titles which we separate into the addresses described above and in the main text.

Preparing Companies in the British Registers

In order to link ownership information across different data sources, we also harmonize the details about the corporate owners in three steps. First, we make a general alignment of the string entries by treating leading and trailing white spaces, caps, or special characters. Second, we extract company types from the strings based on an extensive list of incorporation types and separate them from the information about the company name.³³ Finally, we harmonize company types and reclassify them in a unified way given that a limited liability company for instance could be recorded as “Limited”, “Ltd.”, or “Limited Liability Company”. For companies incorporated abroad, we use a fourth step to extract and harmonize information about the country of incorporation. This harmonization is not only necessary due to frequent misspellings of country names, but also because the records sometimes present regional or local information instead of the country information (e.g. “Tortola” instead of “British Virgin Islands”, or “Abu Dhabi” instead of “United Arab Emirates”).

Adding Orbis Ownership Information

While the British land register data separated between real estate held by foreign and domestic companies, a domestic company can directly be owned from abroad. This would disguise foreign ownership as domestic ownership. In order to address this issue, we employ corporate ownership information from the Orbis dataset by Bureau van Dijk (BvD). This proprietary dataset includes data on 400 million companies and their ownership structure. Starting from the British company, we first harmonize company incorporation types to create a 1:1 match with the Orbis dataset and assign a BvD ID to the British companies that appear in the British registers. In the simplest cases, Orbis already has information on the global ultimate owner of a company, defined as holding more than 50% of the voting shares (GUO 50 in Orbis) of a company. In that case, we

contains property and ownership information of around 100,000 property titles in England and Wales acquired by foreign companies between 1999 and 2014. For details on the data and related news coverage, see information by the Private Eye Magazine.

³³Examples of company types include “LTD”, “SARL”, and “PTC”.

can simply re-assign the country of origin from “United Kingdom” to that owner for the company in question. If this is not the case, we exploit information on shareholders. If a company in Britain for example has one shareholder in France who is not the ultimate owner, we move to that shareholder and check if it has an ultimate owner. The country that has the majority ultimate owner share is then designated as the country of origin of an investment. Whenever we cannot identify global ultimate owners, we refrain from reclassifying the country of origin. This is a conservative approach once the shareholder structure has moved offshore but avoids assigning the shareholder country when no global ultimate owner is identified. Therefore, in our dataset, this data effort re-codes the country of origin from “United Kingdom” to the respective country of the GUO 50 we find in Orbis.

Matching ICIJ Ownership Information

Shell companies in tax havens are routinely used to hide ownership. A regular stream of large data leaks has been collected into a consistent database by the International Consortium of Investigative Journalists (ICIJ). The last installment of this dataset was published in May 2022, adding further data from the *Pandora Papers* leak. The dataset used here includes all available information from the *Offshore Leaks* (2013), the *Panama Papers* (2016), the *Bahamas Leaks* (2016), the *Paradise Papers* (2017), and the *Pandora Papers* (2021). Together with the files on Mossack Fonseca in the Paradise Papers, parts of the corporate registers of six well known offshore jurisdictions were leaked with much less media attention but more data.³⁴ In total, the data contains information on more than 810,000 offshore entities and the respective managerial structures. Most importantly, the data identifies the persons connected to an offshore identity (‘officers’) and reveals their respective role in the company (e.g. beneficial owner, nominee director, nominee shareholder). After several updates over the last years, this database has now available online as a network database.

Preparing Offshore Entities In order to match the information in the leaked databases to the British land registers, we first prepare the company name and company type information using the approach described above for British companies. The list of company types, including misspelled versions is longer here (298 company types, only including those that exist in the offshore leaks database) because the data spans multiple corporate registers. For example, limited liability companies could be a Malaysian SDN BLD, a Luxembourgian SARL, or a British LTD.

³⁴These six jurisdictions are: Aruba, the Cook Islands, the Bahamas, Barbados, Malta, Nevis [part of St. Kitts and Nevis], and Samoa.

Having aligned company names, we match both data sources using name, type, and country of incorporation. This conservative approach can lead to false positive matches only if two firms with the identical name and the same incorporation type exist in the corporate register of one country. This is legally prohibited in the corporate register of the UK and we have not found such a duplicate in the tax havens for which we have data.

Since the country information is therefore important for the match with the leak data and the registers, we go through the country information in the offshore leaks database in detail. In some cases, the description of the jurisdiction is sub-national. We update the sub-national jurisdictions (such as Delaware, Texas, As Al Khaimah, or Dubai) to the country level (USA and United Arab Emirates for these cases). Instead of creating a synthetic Netherlands Antilles, we assign historical Netherlands Antilles cases to Curacao. Other updates concern spelling mistakes usually due to the mistreatment of special characters at some point of the data construction which we re-assign using regular expressions.

Preparing Offshore Officers Officers are the shareholders, directors, and beneficial owners that the offshore leak database identifies. These officer nodes can be linked to entity nodes with undirected edges supplied in the offshore leaks database. This is a many to many match as one officer can be connected to more than one entity and one entity can have more than one officer.

We add one important step to this link: The ‘officer’ assignment in the offshore leak database includes legal entities if they are entered as shareholders, beneficial owners, or directors. In order to distinguish those from legal entities, we employ our company type recognition code and classify any officer name in which we find a company type (such as “Ltd.”) as a legal entity. Whenever we do not, we classify it as a natural person. Visual inspection suggests that this assignment is close to perfect.

Matching Linked Entity-Officer Data with British Register Data The resulting combination of offshore shell companies with their beneficiaries can then be matched to our previously prepared address data. In order to do so, we rely on three prepared pieces of information: The name of a company, its incorporation type, and the country of incorporation. Our matching first uses the original capitalized names and types and the country of incorporation (49% of all matches). We then proceed to match remaining data using the prepared company types, their name stump without a company type, and the country of incorporation as three variables (29% of all matches). Finally, we discard the company type and only match using the name stump and the country of incorporation (22% of all matches). At the current stage, we use all these

uncertainty levels in the results of the paper and we have not detected non-randomness of the resulting country distributions across these three uncertainty levels.

Note that we do not use the ICIJ leak data to update our main dataset. A property will still be counted as tax offshore held for our main statistics even if we could identify the beneficial owner in the ICIJ leak. We analyze the ICIJ match in the respective section in detail.

B.3 Matching Price and Ownership Data

Having access to prepared postal address information for the transaction price dataset (Prices Paid data), as well as prepared addresses and company information for the corporate ownership datasets (CCOD, historical CCOD, OCOD, historical OCOD, our freedom of information request, and the Private Eye data) we then proceed to match the two.

Timing of Transactions and Owner Registrations Beyond address information, the second ingredient of our match is information on the timing of real estate transactions. The price paid data registers the time that the transaction was recorded while the “date proprietor added” in the ownership data registers the date of entry of the new owner in the land register. When we match addresses unconditionally and eyeball the time difference between both dates, we see peaks in the same month. We also see mass with at most four months difference: That is the owner being recorded up to four month after the transaction took place. Correspondence with the British land register confirmed this expected time lag. We therefore only match transaction and ownership data when there is at most a four month lag between both. This is a conservative approach because longer lags do exist but at much lower probabilities. All matches described below take place using this timing tolerance.

We start by matching simple addresses: identifiable by a postcode, a street, a house number, and information on the tenure (freehold versus leasehold) with no further information. We then match properties where this information was identified but also units, for example flats, and their unit numbers were identified. This code is also extended to properties with more than one unit (such as “Block 1, Flat 3”). Finally, we match on localities and locality numbers instead of units. However, we still only use properties where information on postcode, street, house number, and tenure is available.

Having matched each of our ownership datasets to the price paid data, it is possible that one transaction was matched in several data sources. For example, if a property was sold to a foreign company in 2010 and not sold since then, this sale will show up in the freedom of information

data as a transaction. It will also show up in the historical OCOD stock data since the entry still captures the current owner. It will also show up a few months later in the OCOD dataset. We introduce a data hierarchy going from the last available dataset to the first (OCOD \rightarrow OCOD hist \rightarrow FOI \rightarrow Private Eye; CCOD \rightarrow CCOD hist) and take all other auxiliary information from the data source of the highest point in the hierarchy where the match appears. This is the point where the Private Eye matches drop out of our dataset: For the properties we matched, it is dominated by OCOD and its historical version. This is expected and a useful sanity check as the Private Eye data was originally based on a freedom of information request on foreign property purchases: The same data that was used to construct OCOD in the first place. We do retain matched from the FOI request though which is crucial: This shows that it can be employed to update the information that OCOD, being a stock database, does not provide. Namely, we can see transactions of a property that took place before the current owner at the start of the real time database was entered into the register.

This iterative process makes sure that we only match on well identified addresses and within a small time tolerance. When we match a foreign owner to a transaction, we can be certain that it is correct. Any deviation in data quality will prohibit a match. False positives are therefore very unlikely. Our match should therefore be interpreted as a lower bound estimate.

C Long-run Dynamics of the Offshore Market Share

To estimate the stock volume of foreign held residential real estate, we combine transaction information on all *residential* real estate from the Price Paid Data (PPD) with ownership information on *all corporate held* real estate from the land register (LR) and ORBIS. The confidence of this stock volume estimation varies over time. While after October 2015, regular snapshots of the land register are available, we only have limited information on purchases by foreign companies between 1995 and October 2015. We thus discuss the estimation strategy for these two periods, i.e. 1995 to October 2015, and after October 2015 separately.

Estimation after October 2015

Given that snapshots of the land register are available after October 2015 on a regular basis, we can estimate the foreign held stock volume based on the information in each respective snapshot. For this purpose, we first merge the land register snapshot to PPD and supplement it with ownership information from ORBIS.³⁵

For the estimation of the aggregate offshore market share at different points in time, we improve upon this data combination by making two adjustments. First, we also include transactions from the land register data that do not appear in the PPD, but whose transacted properties appeared at any point in the PPD. We thereby assume that a property is always either residential or non-residential but does not switch between these categories. Second, we impute the real estate stock value between the land register snapshots by using the transaction dates instead of the dates of the land register snapshot. Purchase prices are adjusted using district level house price indices (HPI). By using the transaction date information separately from the actual snapshot time, we correct for lagged appearances of properties in the land register.³⁶

Estimation before October 2015

Before October 2015, we no longer have regular LR snapshots at our disposal. We therefore use all available purchase information before October 2015 to estimate the stock volume and shares of foreign held residential real estate between January 1995 and October 2015.³⁷

³⁵For details on the combination of land register data, Price Paid Data, and ORBIS, see Appendix B.

³⁶For example, a property might be registered in LR with a date proprietor added (DPA) of August 2018, but only appears in the snapshot of December 2018. Disentangling snapshot time from transaction time might thus increase the accuracy of our stock volume measure.

³⁷As above, PPD information is supplemented with land register information on properties that appear at any time in PPD and are thus considered residential.

Stocks For every year-month t , we are interested in the stock value of real estate that is held by private owners (S_t^P), domestic corporate owners (S_t^D), foreign tax haven corporate owners (S_t^{TH}), foreign non haven corporate owners (S_t^{NH}), and foreign corporate owners without country information (S_t^F). Stock estimates at time t refer to the end-of-month stock.

Flows While we do not observe stocks directly, we observe at every year-month t the flow value F_t^{o-d} from origin o to the destination d . Consistent with the stock owner types, origin and destination can be P, D, TH, NH, F . In addition, the origin of a flow can also be New . Flows from New to one of the other owner type segments are identified from PPD information on whether a property is newly built or an established residential building.

Linking Stocks and Flows The following equation shows exemplary for the stock of real estate held by tax havens, how stocks and flows can be linked. Note that we link stocks and flows separately for every district i .

$$\begin{aligned}
S_{t,i}^{TH} = & \underbrace{(1 + g_{t-1,i})}_{\text{HPI growth adjustment}} \underbrace{S_{t-1,i}^{TH}}_{\text{Stock end of previous period}} \\
& - \underbrace{F_{t,i}^{TH-P} - F_{t,i}^{TH-D} - F_{t,i}^{TH-NH} - F_{t,i}^{TH-F}}_{\text{Flows out of tax havens into other owner type segments}} \\
& + \underbrace{F_{t,i}^{P-TH} + F_{t,i}^{D-TH} + F_{t,i}^{NH-TH} + F_{t,i}^{F-TH}}_{\text{Flows into tax havens from other owner type segments}} \\
& + \underbrace{F_{t,i}^{New-TH}}_{\text{New Properties}}
\end{aligned} \tag{4}$$

Importantly, we observe foreign ownership of purchases only since 1995. As a consequence, we do not observe all flows with their origin o . Flows thus consist of an observed and an unobserved component. For instance, the flow from tax haven to private reads

$$F_{t,i}^{TH-P} = \underbrace{\hat{F}_{t,i}^{TH-P}}_{\text{observed}} + \underbrace{\tilde{F}_{t,i}^{TH-P}}_{\text{unobserved}} . \tag{5}$$

The unobserved part of the flow can be estimated, since we know if a property has existed already before the purchase. For instance, denote $F_{t,i}^{?-P}$ the inflow of property value into the private sector that has no previous owner and does not consist of newly built properties.

In order to assign the origin of these flows into destination d , we assume that the origins are proportional to the origins of flows into destination d for which we know the owner. For the example above, it follows that

$$\begin{aligned}
\tilde{F}_{t,i}^{TH-P} &= \underbrace{\hat{F}_{t,i}^{?-P}}_{\text{flows to P, unobserved origin}} \times r_{y(t)}^{TH-P} \\
&= F_{t,i}^{?-P} \times \frac{\hat{F}_{y(t)}^{TH-P}}{\hat{F}_{y(t)}^{TH-P} + \hat{F}_{y(t)}^{NH-P} + \hat{F}_{y(t)}^{F-P} + \hat{F}_{y(t)}^{P-P}}
\end{aligned} \tag{6}$$

Note that the ratio $r_{y(t)}^{TH-P}$ is calculated at the country level and at the year level $y(t)$ to have a reasonably large sample size for the estimation of the shares.

Combination and Inversion We estimate the stocks at every point in time by inverting equation (4) and solving for the stocks going backward in time. This is possible, because we have a correct estimate of the stock of real estate with corresponding owner types in October 2015. Combining the inverted equation (4) with equations (5) and (6), we obtain

$$\begin{aligned}
S_{t-1,i}^{TH} &= \frac{1}{(1 + g_{t-1,i})} \left[S_{t,i}^{TH} \right. \\
&\quad + \hat{F}_{t,i}^{TH-P} + \hat{F}_{t,i}^{TH-D} + \hat{F}_{t,i}^{TH-NH} + \hat{F}_{t,i}^{TH-F} \\
&\quad + \tilde{F}_{t,i}^{TH-P} + \tilde{F}_{t,i}^{TH-D} + \tilde{F}_{t,i}^{TH-NH} + \tilde{F}_{t,i}^{TH-F} \\
&\quad - \hat{F}_{t,i}^{P-TH} - \hat{F}_{t,i}^{D-TH} - \hat{F}_{t,i}^{NH-TH} - \hat{F}_{t,i}^{F-TH} \\
&\quad - \tilde{F}_{t,i}^{P-TH} - \tilde{F}_{t,i}^{D-TH} - \tilde{F}_{t,i}^{NH-TH} - \tilde{F}_{t,i}^{F-TH} \\
&\quad \left. - F_{t,i}^{New-TH} \right].
\end{aligned} \tag{7}$$

Aggregation In a final step, we aggregate the district level stock estimates and calculate the stock shares by ownership with respect to the total residential market denoted by s . For instance, the aggregate offshore tax haven share is calculated as

$$s_t^{TH} = \frac{\sum_i S_{t,i}^{TH}}{\sum_i S_{t,i}^{TH} + \sum_i S_{t,i}^{NH} + \sum_i S_{t,i}^F + \sum_i S_{t,i}^D + \sum_i S_{t,i}^P}. \tag{8}$$

Validation

To validate our results, we employ two validation exercises. First, we benchmark our estimate of the total volume of residential real estate against information from the national accounts. Second, we validate our imputation strategy by applying it to the period after October 2015 for which reliable estimates of the offshore market share growth rates exist.

Residential Market Volume The national accounts³⁸ provide since 1995 at the national level a yearly estimate of the capital stock value for dwellings (residential properties), other buildings and structures, and land. We make two assumptions to make the time series of the national accounts comparable to what our estimate captures. First, since land values are not separately shown for dwellings and other buildings and structures, we allocate land value proportionally to dwellings and other structures. Second, since the national accounts capture the whole UK, we scale series by the land share of England and Wales. Figure A2a compares the values from the adjusted national accounts with our own estimation at the monthly level.

Offshore Market Share Growth To validate our imputation strategy, we focus on the two most recent years in our data (October 2015 to December 2019). In this time period, we have credible longitudinal information on the offshore market share growth rates to which we can compare results from our imputation. We take the December 2019 value of real estate as given. We then deliberately ignore information about previous owners at the property level. Instead, we follow the imputation strategy above by classifying previous owners based on the year-specific transition probabilities calculated at the aggregate level. Figure A2b compares the resulting month-to-month growth rates for the tax haven share in residential real estate to its actual values.

³⁸see UK National Accounts.

D Taxation of Capital Gains

Capital gains tax (CGT) in the UK was introduced in 1965 and applies to profits resulting from selling, swapping or transferring an asset which increased in value. With respect to real estate holdings, the capital gains tax applies to all real estate property that is (i) not the primary residence of the taxpayer, or (ii) is the primary residence of the taxpayer, but is let out, used for business, or is larger than 5,000 square meters. While capital gains valuation are usually based the difference between the purchase and sales price, market valuations can be applied if assets are transferred at prices below their market value. Capital gains tax applies to net gains across all taxable assets, i.e. losses from one asset can be deducted from the gains of another asset. Both the tax rates and the tax-free allowances for the capital gains tax vary with the type of investor and the income level of the taxpayer.³⁹

Historically, foreign investors—both individuals and corporate—were mostly exempt from capital gains tax.⁴⁰ In April 2015, CGT rules were extended to foreign individuals, close companies (companies controlled by five individuals or less), trusts and personal representatives disposing UK residential property. However, both the disposal of UK residential property by diversely held foreign companies and the disposal of non-residential property remained exempt.

The 2017 Policy Change In November 2017, the government announced to remove these remaining advantages for non-resident investors:

To align the UK with other countries and remove an advantage which non-residents have over UK residents, all gains on non-resident disposals of UK property will be brought within the scope of UK tax. This will apply to gains accrued on or after April 2019.

Autumn Budget 2017, p. 35

The details of this policy change were further specified during a consultation process between November 2017 and July 2018. The final draft legislation included provisions to extend CGT coverage to both direct disposals of non-residential properties and indirect disposals of residential property through “property-rich” corporations. In particular, indirect disposals through shares became subject to CGT if the company derives at least 75 percent of its gross asset value from UK real estate and if the person making the disposal held at anytime at least a 25 percent

³⁹See detailed information on the capital gains tax by the UK government.

⁴⁰From 2013 to 2015, CGT did apply to foreigners subject to the Annual Tax on Enveloped Dwellings (ATED).

Table 5 – Capital Gains Tax Rates for Foreigners, April 2019

Investor Type	Non-Residential Property	Residential Property
Individuals	10% / 20%	18% / 28%
Companies	19%	19%
Trustees	20%	28%

Notes: The table displays capital gains tax rates applicable to non-resident individuals, companies, and trustees applicable to capital gains as of April 2019. Non-residential properties include commercial property, for example shops or offices, agricultural land, forests, and any other land or property which is not used as a residence. Residential properties include buildings used or suitable for use as a dwelling as well as the underlying garden or grounds associated with such buildings. Marginal tax rates for individuals differ, since individuals pay lower rates within the basic income tax band and a higher rate on any amount above the basic income band. *Source:* HM Revenue & Customs.

investment in the company during the last two years.⁴¹ Table 5 summarizes the tax rates applicable to foreigners on gains accrued after April 2019.

The Luxembourg Exemption At the time of announcement, the new regulation interacted with existing double tax treaties, most notably with the double tax treaty between the UK and Luxembourg. While taxing rights regarding the direct disposal of property are usually not affected by double tax treaties, the UK was restricted in taxing the gains made by foreign investors via shares in UK property rich entities. Although the UK government was aware of this loophole and enacted general anti-forestalling rules, it was and is still unclear, how the exact coverage and the enforcement of these rules will look like.⁴²

As an additional measure against “treaty shopping”, the UK implemented the OECD multilateral convention to implement tax treaty related measures to prevent base erosion and profit shifting in October 2018. This multilateral amendment included a particular clause on indirect disposals of “property-rich” entities:

For purposes of a Covered Tax Agreement, gains derived by a resident of a Contracting Jurisdiction from the alienation of shares (...) may be taxed in the other Contracting Jurisdiction if (...) these shares or comparable interests derived more than 50 per cent of their value directly or indirectly from immovable property.

OECD Multilateral Amendment, Article 9, Paragraph 4

This multilateral amendment allowed the UK to regain the rights to capital gains through

⁴¹While the UK government explained that both the asset value test and the ownership test aggregate disposals of multiple entities and ownership shares held by certain persons through a series of entities, it is unclear how this provision is implemented, if beneficial owners are concealed through offshore structures. Additional exemptions to this rule include UK property used by the UK branch or agency of a foreign company, property held by overseas pension schemes, sovereign wealth funds and charities. For details, see information on the consultation process and the regulatory framework on the capital gains tax for non-residents.

⁴²See for instance a comment by the law firm Goodwin Procter, but also a more recent assessment by the law firm DLA Piper stating that “UK anti-forestalling rule (...), depending on the facts, could be circumvented.”

the indirect disposal of immovable property. However, Luxembourg ratified the amendment only in April 2019 and the corresponding changes on capital gains tax formally apply only to capital gains after April 2020.⁴³ In mid 2022, the UK and Luxembourg signed a new double taxation treaty that directly integrates this treatment of the indirect disposals of “property-rich” entities:

Gains derived by a resident of a Contracting State from the alienation of shares or comparable interests, such as interests in a partnership or trust, deriving more than 50 per cent of their value directly or indirectly from immovable property (...) situated in the other Contracting State may be taxed in that other State.

2022 UK-Luxembourg Double Taxation Convention, Article 13, Paragraph 2

While this tax treaty will only come into force in 2023, it might lead to a new wave of restructurings away from Luxembourg entities that will no longer offer opportunities for capital gains tax avoidance on real estate holdings in the UK.

⁴³See information by HM Revenue & Customs on Tax Treaties between the UK and Luxembourg.

E Public Ownership Registers

At least since 2016, the installation of ownership registers and the exchange of the corresponding beneficial ownership information has become a core priority of governments around the world to combat illicit financial flows including terrorist financing, money laundering and tax evasion.⁴⁴

While ownership registers have been in place in some of the Overseas Territories (OT) and the Crown Dependencies (CD) for a long time, the access to these registers (if available) was historically very restricted. Even the UK government could not access the recorded information until a series of bilateral agreements was negotiated between the UK government and the OT and CD authorities in 2016.⁴⁵ These agreements were supposed to give law enforcement authorities fast and secure access to beneficial ownership information on companies incorporated in OT and CD countries. Importantly, however, OT and CD did not plan to make the ownership information *publicly* available at that time.

The 2018 Policy Announcement The announcement of public corporate registers has its roots in an amendment to the *Sanctions and Anti-Money Laundering Act* (SAMLA) that entered the parliamentary process already at the beginning of 2018.⁴⁶ While the government first opposed the amendment, it became apparent in April 2018 that it had enough supporters to defeat the government.⁴⁷ The amendment was approved in the House of Commons on May 1st 2018.⁴⁸

According to the amendment, the UK Secretary of State shall provide assistance to the Overseas Territories in setting up a publicly accessible register of the beneficial ownership of companies registered in each government's jurisdiction. In the case that OT governments fail to install public registers until the end of 2020, the UK government is required by law to formally mandate such a public register:

The Secretary of State must, no later than 31 December 2020, prepare a draft Order in Council requiring the government of any British Overseas Territory that has not introduced a publicly accessible register of the beneficial ownership of companies within its jurisdiction to do so.

Sanctions and Anti-Money Laundering Act 2018, Section 51, Subsection 4

⁴⁴See for instance the joint statement by multiple countries on the systematic sharing of beneficial ownership information and the implementation Financial Action Task Force standards.

⁴⁵See the collection on the bilateral agreements between the UK and different OT and CD jurisdictions.

⁴⁶See the proposal of Amendment 73 to SAMLA at the report stage of the bill.

⁴⁷See for instance a Guardian newspaper article on the SAMLA amendment.

⁴⁸See details on the legislative implementation procedure of SAMLA.

While the law itself does not specify a particular date at which the registers are expected to be accessible, a Joint Ministerial Council explained in December 2018 that the mandate would request the registers to be operational at the end of 2023.

The Exclusion of Crown Dependencies While OT were supposed to be directly mandated to set up public beneficial ownership registers in the future, CD were excluded from SAMLA, because their constitutional status implies a higher degree of autonomy. In particular, Guernsey, Jersey and the Isle of Man elect their own legislatures and run their own domestic legislation. Although a formal mandate is not possible, CD have announced in June 2019 to set up public registers voluntarily:

Within 12 months of [the publication of the Implementation Review of the 5th AMLD in January 2022], we will each bring forward to our own parliament legislative proposals to establish public access to beneficial ownership data of companies held on a central register.

Joint commitment by Guernsey, Jersey and the Isle of Man

At the end of 2020, when the OT mandate was supposed to be issued, the government published the draft Order in Council.⁴⁹ It is noteworthy, however, that the UK government did not see the need to officially bring this mandate into force given the “firm commitments from all of the inhabited Overseas Territories to adopt publicly accessible registers”.⁵⁰

⁴⁹See draft Order in Council.

⁵⁰See Ministerial Statement on publicly accessible registers in the UK Overseas Territories.