Journal of Housing Economics 62 (2023) 101964

Contents lists available at ScienceDirect

Journal of Housing Economics

journal homepage: www.elsevier.com/locate/jhec

The heterogeneous relationship of owner-occupied and investment property with household portfolio choice☆

Marco Felici, Franz Fuerst *

University of Cambridge, Department of Land Economy, United Kingdom

ARTICLE INFO

JEL classification:
G11
G51

Keywords:
Household finance
Property
Portfolio choice

ABSTRACT

The dual nature of property as both a consumption and investment good presents a challenge for household portfolios. Prior theoretical literature predicts a constraint imposed by property on investment decisions, and empirical studies support this notion. However, previous research often overlooks investigating the heterogeneity of this constraint and fails to differentiate between owner-occupied and investment property. Building on a stochastic control model, we analyse the UK's Wealth and Assets Survey panel and find that along the distribution of how household allocate their wealth a one percentage point (pp) increase in the share of owner-occupied property in the total portfolio is associated with a 0.07 pp decrease in the share of stocks in liquid assets. However, this association varies significantly based on the value of the owner-occupied property share. For low values of the owner-occupied property share, the association with stockholdings is negligible. As the share of owner-occupied property increases, the negative association with stockholdings becomes more pronounced: when the owner-occupied property share reaches 90%, a further 1 pp increase corresponds to a 0.14 pp decrease in the share of stocks in liquid assets. By contrast, buy-to-let property shows no significant relationship with stockholdings, supporting the idea that the constraint on portfolio decisions is primarily driven by the role of property as a consumption good.

1. Introduction

The dual nature of property, serving both as a consumption and investment good, has intricate effects on household investment decisions. Renters typically hold fewer property assets in their portfolios compared to homeowners and property owners in general. Additionally, the composition of their liquid assets, not including property, may differ, particularly concerning risk. This divergence could be influenced by individual preferences, income, age, or specific characteristics related to property, such as its relatively low liquidity.[1]

Consensus exists that owner-occupied property has an impact on financial portfolios, particularly on the allocation to stocks (Grossman and Laroque, 1990; Brueckner, 1997; Fratantoni, 2001; Cocco, 2005; Hu, 2005; Chetty and Szeidl, 2007; Pelizzon and Weber, 2008; Stokey, 2009). However, the direction of this effect depends on various factors, such as adjustment costs, committed expenditure risk, or correlation with returns from other assets. Empirical evidence tends to support a negative relationship between property and stockholdings (Yao and Zhang, 2005; Chetty et al., 2017; Vestman, 2019). On the other hand, when it comes to stock market participation, estimates are often null or positive (Vestman, 2019; Kong et al., 2021).

Nonetheless, two essential aspects have been either ignored or overlooked in existing research: the distinct relationship between owner-occupied and investment property with stockholdings and the potential heterogeneity of these relationships among households. While most studies concentrate on owner-occupied property, some suggest that renters' portfolio choice should not be affected by property at all (Brueckner, 1997), and very few explicitly consider investment property for either renters or homeowners (such as Yao and Zhang (2005)). Yao and Zhang (2005) is also one of the few studies that include non-linear terms in their estimations, allowing for differential effects of property relative to wealth based on its size.

☆ We thank the editor and three anonymous referees for insightful comments that greatly improved the manuscript. We also thank Rebecca Tunstall, Li Wan as well as seminar participants at the ENHR (2019) and ERES (2019, 2021) Annual Conferences and at a 2020 Savills internal seminar for helpful discussions. Funding from the Cambridge Humanities Research Grant Scheme is gratefully acknowledged. Franz Fuerst gratefully acknowledges the support of the Cambridge University Land Society (CULS).

* Corresponding author.
E-mail address: ff274@cam.ac.uk (F. Fuerst).

[1] For a comprehensive review of household portfolios and finance, refer to Gomes et al. (2021).

https://doi.org/10.1016/j.jhe.2023.101964
Received 29 November 2022; Received in revised form 5 August 2023; Accepted 21 August 2023
Available online 9 September 2023

Our paper addresses these two aspects, focussing on the theoretical implications and empirical evidence of how investment property (including non-owner-occupied and buy-to-let property) relates to stockholdings differently from owner-occupied property, and whether these relationships vary across households. We base our hypotheses on the stochastic control model with fixed adjustment costs by Stokey (2009), which indicates how adjustment costs in changing owner-occupied property can result in state-dependent risk aversion and thus different risky asset holdings. Since this behaviour is linked to the consumption nature of owner-occupied property, it should not apply to investment property. Consequently, we predict that the relationship between the risky asset and owner-occupied property changes based on the relative size of the latter, while no such relationship exists with investment property.

While most empirical literature suggests an overall negative effect of property on stockholdings, our analysis reveals a significantly heterogeneous distribution. For households with owner-occupied property comprising a relatively small portion of their portfolios (e.g., 10 or 20%), an increase in the share of owner-occupied property is positively associated with changes in the share of stocks, albeit with large standard errors that make it indistinguishable from 0. However, the association becomes less positive as the share of property in the portfolio increases and eventually turns negative for households with almost their entire portfolio composed of property.

Our results are consistent with a portfolio choice model where a property's effect on stockholdings is non-linear due to risk aversion and the presence of adjustment costs for owner-occupied property. The adjustment costs should matter only for owner-occupied property and not for investment property. This is confirmed in the empirical analysis only for a strict definition of investment property: it holds when we consider buy-to-let property, whose acquisition is clearly not motivated by a consumption motive but not as much for all investment properties which are defined more broadly as any property that is not owner-occupied.

While studies on this topic explore other countries apart from the US (e.g., Vestman (2019) for Sweden), there is limited evidence specific to the United Kingdom, and our paper contributes to this aspect as well. The subsequent sections detail the theoretical framework (Section 2), describe the Wealth and Assets Survey and relevant data patterns (Section 3), outline the empirical strategy to identify the property–stockholdings relationship (Section 4), present and discuss the results, emphasizing the heterogeneity and significance of owner-occupied property (Section 5), and conclude (Section 6).

2. Theoretical framework

The fundamental portfolio choice of a household involves these two categories:

Definition 2.1 (*Illiquid Share*)**.** The *Illiquid Share* is the share of property in **total** assets.

Definition 2.2 (*Risky Share*)**.** The *Risky Share* is the share of stocks in **liquid** assets.

The *Illiquid Share* can be further divided into two components: owner-occupied and investment. By contrast, the *Risky Share* is typically defined in terms of non-property or financial assets (for a recent example see Chang et al. (2022)).

The theoretical significance of defining the *Illiquid Share* in this manner becomes apparent in Stokey (2009), a stochastic control model that we adopt and qualitatively discuss in this section to highlight the theoretical aspects supporting our empirical approach. Fig. 1 provides a graphical representation of the main features of this framework. This model, which extends the model of Grossman and Laroque (1990), maintains that when there are adjustment costs associated with changing the amount of property, the relationship between holdings of risky assets and property, relative to wealth, is not constant. Instead, it depends on the size of property relative to wealth. Notably, this relationship remains constant or zero for owner-occupiers in the absence of adjustment costs, and it is zero for renters irrespective of the presence of such costs. This behaviour arises from the fact that, with adjustment costs, risk aversion becomes non-linearly dependent on the relative size of property, whereas without such costs, the relationship remains constant or zero.

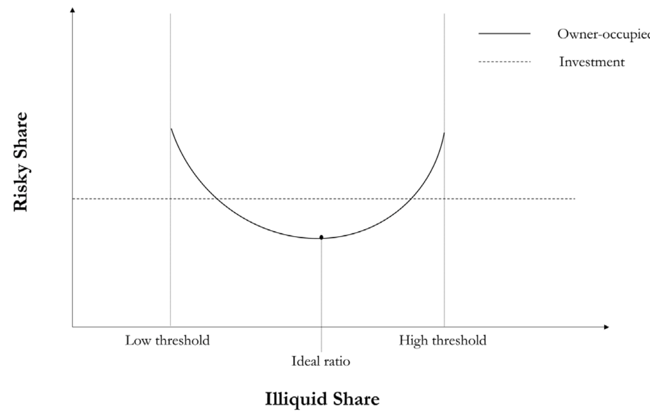

Fig. 1. The theoretical relationship between the *Illiquid Share* and the *Risky Share*, by type. Note: Diagram showing the main features of the theoretical model, after Figure 1b in Stokey (2009).

Let us consider a scenario where a household has an ideal ratio of owner-occupied property to wealth, and wealth fluctuates randomly based on their holdings of the risky asset. Changing the main residence involves considerable costs in terms of time and money spent on finding a new place and moving, leading to infrequent moves that only occur when convenient or forced by sudden events. The household deems it convenient to move only when the ratio of owner-occupied property to total wealth deviates significantly from their ideal ratio in either direction. Thus, there are high and low thresholds that trigger a move.

This situation gives rise to an inaction region between the two thresholds, where the ratio changes but no move is triggered, and an action region below the low threshold or above the high threshold, which initiates a move to return to the ideal ratio. Changes in the ratio within the inaction region cause fluctuations in the distance between the actual ratio and the thresholds, altering the household's risk aversion, which, in turn, impacts their stock holdings (in this sense, we effectively consider the *Illiquid Share* as the state variable and therefore exclude simultaneity with the stockholding decision as in Beaubrun-Diant and Maury (2016)).

Numerical simulations by Stokey (2009) demonstrate that the model results in a U-shaped portfolio policy when adjustment costs are present and binding. The household exhibits lower risk aversion when the *Illiquid Share* is close to the thresholds, becoming most risk-averse towards the centre of the inaction region (consistent with Grossman and Laroque (1990)). However, for the "frictionless consumer", which we frame as the case of the investment *Illiquid Share*, risk aversion remains constant.

We chose this model as the basis for our hypotheses because its central feature – the presence of adjustment costs – directly reflects the consumption motive as the factor driving the difference in the effects of owner-occupied and investment property on the financial portfolio. Additionally, the model enables us to study the non-linear association between the *Risky Share* and the *Illiquid Share* when adjustment costs are binding. While other existing models focus on factors like price risk effects of housing (Chetty et al., 2017) or life-cycle considerations (Vestman, 2019), they are less concerned with the distinction between owner-occupied and investment property and do not fully account for adjustment costs. Based on this framework, we formulate two hypotheses:

Table 1
Overview of summary statistics.

	Renting (N = 14,364)		Owning (N = 51,448)	
	Mean	St. Dev.	Mean	St. Dev.
Overall *Illiquid Share* (%)	3.64	17.43	84.59	16.65
Illiquid Share for investment property (%)	3.64	17.43	4.44	12.89
Illiquid Share for buy-to-let property (%)	0.99	9.25	2.15	9.20
Risky share (%)	1.71	10.15	5.57	16.04
Total property value (£100,000)	0.08	0.91	3.14	3.96
Investment property value (£100,000)	0.08	0.91	0.41	2.65
Buy-to-let property value (£100,000)	0.03	0.35	0.20	1.25
Liquid assets (£100,000)	0.13	0.60	0.89	7.31
Mortgage debt (£100,000)	0.02	0.20	0.62	1.13
Home equity (£100,000)	0.07	0.85	2.52	3.72
Income (£100,000)	0.23	0.18	0.38	0.38
Number of children	0.66	1.04	0.52	0.89
Below age 35 (%)	23.45	42.37	9.56	29.40
In employment (%)	53.74	49.86	66.00	47.37
Education at degree level or above (%)	16.30	36.94	33.51	47.20

Note: Summary statistics for the first six waves (2006–2018) of the Wealth and Assets Survey, pooling all cross-sections together. Sample weights are used throughout and nominal values have July 2015 as reference period.

Hypothesis 1. The relationship between the *Risky Share* and the owner-occupied *Illiquid Share* changes based on the size of the *Illiquid Share*.

Simultaneously, when adjustment costs are less relevant or not applicable (indicating a motive solely focused on investment), there should be no relationship. In such cases, property that is not owner-occupied loses its consumption nature. While Stokey (2009) interprets this scenario as that of a renter, we extend it to encompass all investment property, whether owned by renters or homeowners. It is crucial to clarify that assuming negligible adjustment costs for property that is not owner-occupied does not imply the absence of transaction costs altogether. All types of property involve transaction costs, such as search expenses, broker fees, and due diligence. However, we contend that such adjustment costs are less impactful on household decisions regarding stockholdings compared to the costs involved when the property is owner-occupied.

In fact, our concept of adjustment costs goes beyond mere transaction costs. In the words of Stokey (2009): "Moving typically entails substantial adjustment costs. These include direct financial costs, such as agents' commissions, legal fees, transfer taxes, and shipping/transportation costs, as well as the time cost of searching, transacting, and executing the move, and the psychic cost of changing school districts, broken emotional ties, and other disruptions.". Consequently, while transaction costs may vary in different ways between investment and owner-occupied property (for instance, rental property might entail informational costs due to lease considerations), the emotional costs are likely negligible for investment property but significant for owner-occupied property. The case of vacation homes is a special one. While transacting vacation homes may entail relatively lower transaction costs since neither owners nor tenants need to vacate the dwelling, there might still be emotional costs associated with selling a second home.

Hence, we assume that adjustment costs do not apply, and there are no action and inaction regions in the relationship between the relative share of property used for investment purposes and the relative share of stockholdings. In other words, the relative allocation to investment property does not significantly impact the allocation to stocks in this scenario. Our second hypothesis is then:

Hypothesis 2. There is no relationship between the *Risky Share* and the investment *Illiquid Share*.

Although the role of mortgage debt is not explicitly addressed in this framework, it is worthwhile to consider what its theoretical implications may be. Fratantoni (2001) compares a Consumption-Capital Asset Pricing Model without housing to one with housing, noting that the additional uncertainty stemming from the "committed-expenditure risk", the risk coming from holding a mortgage, can explain a large drop in the predicted holdings of a risky asset. In fact, the effect of the mortgage commitment is more important than that of house price risk, and is in addition to the uncertainty stemming from labour income. Empirically, we are including a role for debt holdings by controlling for net worth, which indirectly accounts for mortgage debt, as well as estimating a specification where we introduce an indicator variable distinguishing between outright homeowners and mortgagors (Section 4.3).

3. Patterns and data description

Our primary data source consists of the first six waves of the Wealth and Assets Survey (WAS), spanning 2006 to 2018 (Office for National Statistics, 2020). The WAS is a repeated cross-section of British households and incorporates a panel component. It offers valuable insights into asset distribution within households and includes information on various demographic variables.

To enhance sample efficiency, considering that wealth is not distributed uniformly, addresses likely to correspond to wealthier households are sampled at a higher rate. The estimation sample we use comprises approximately 66,000 observations, representing around 18,500 unique households. On average, households are tracked for 3.6 waves, with the minimum being one wave and the maximum six waves. In Table 1, we present summary statistics for the main variables of interest, differentiating between renters and homeowners (including mortgagors). To compile the table, we pool data from all waves and households, treating each occurrence of the same household in different waves as separate observations.

The table demonstrates a strong correlation between property ownership and various demographic and socio-economic factors. As expected, homeowners have a significantly higher *Illiquid Share*, which includes owner-occupied property, resulting in a higher absolute property value. On the other hand, renters and homeowners have similar values for the *Illiquid Share* of investment properties (excluding owner-occupied property), but not for the *Illiquid Share* of buy-to-let properties.

Homeowners also exhibit a substantially higher *Risky Share*, along with more investment and buy-to-let properties in absolute value. They tend to have higher income, fewer children, older age, a larger proportion in employment, and higher educational levels.

The *Risky Share* is calculated as the absolute value of stocks in £, divided by the sum of the absolute value of stocks, bonds, accounts, insurance, and a residual component comprising additional financial assets outside these categories, also in £. Stocks include UK, overseas,

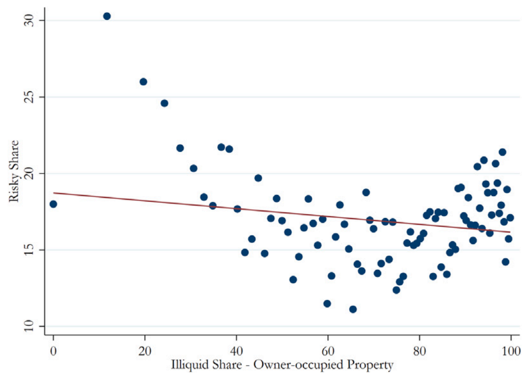

Fig. 2. *Risky Share* against *Illiquid Share* for owner-occupied property. Note: Fig. 2 is a binned scatterplot showing the relationship between the *Risky Share* and the owner-occupied *Illiquid Share*, for participants in the stock market. The binned scatterplots are built by dividing each variable in equally sized bins, computing the average within each bin and then plotting those against each other.

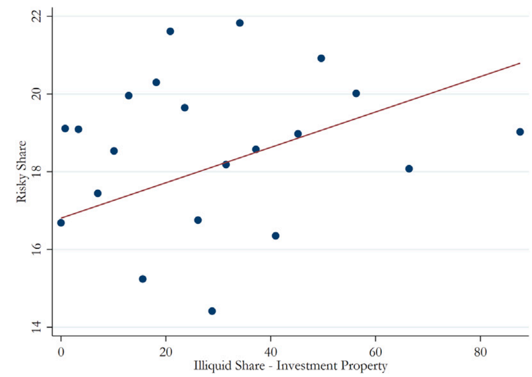

Fig. 3. *Risky Share* against *Illiquid Share* for investment property. Note: Fig. 3 is a binned scatterplot showing the relationship between the *Risky Share* and the investment *Illiquid Share*, for participants in the stock market. The binned scatterplots are built by dividing each variable in equally sized bins, computing the average within each bin and then plotting those against each other. The small number of bins for Fig. 3 are due to the distribution of the *Illiquid Share* for investment property being concentrated at 0.

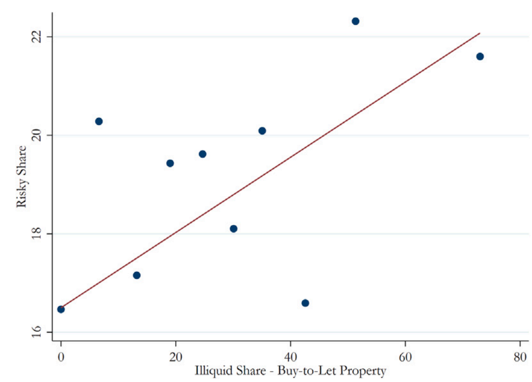

Fig. 4. *Risky Share* against *Illiquid Share* for buy-to-Let property. Note: Fig. 4 is a binned scatterplot showing the relationship between the *Risky Share* and the buy-to-Let *Illiquid Share*, for participants in the stock market. The binned scatterplots are built by dividing each variable in equally sized bins, computing the average within each bin and then plotting those against each other. The small number of bins for Fig. 4 are due to the distribution of the *Illiquid Share* for buy-to-let property being concentrated at 0.

and employees' stocks; bonds consist of UK, overseas bonds/gilts, and fixed-term investment bonds; accounts encompass current and savings accounts, Individual Savings Accounts, and National Savings Products; insurance represents the sum of insurance products with cash value; and the residual component covers various formal and informal financial assets, such as trust funds and endowments, or amounts privately loaned to others.

The *Illiquid Share* measures in its numerator the value of property, including the main residence, other houses, buy-to-let houses, UK and overseas land, and any residual property, in £. It is divided by the sum of property and liquid assets, as previously defined, in £. Throughout the paper, variations of this share are used, including owner-occupied *Illiquid Share*, which considers only the main residence; investment property *Illiquid Share*, including all properties except the main residence; and buy-to-let *Illiquid Share*, which includes only buy-to-let houses.

Notably, retirement accounts are not part of the theoretical framework or the estimation for similar reasons as those cited in Chetty et al. (2017): the lack of detailed information on the portfolio allocation within pension accounts, which makes it challenging to separate holdings by their level of risk. Additionally, pension accounts are generally not meant for withdrawal before retirement, making them less flexible than property as a form of wealth.

Based on the theoretical framework, we hypothesize a non-linear relationship between the *Risky Share* and the *Illiquid Share* for owner-occupied property, while no such relationship is expected for investment property. Three binned scatterplots (Figs. 2, 3, and 4) depict the relationship between these variables for participants in the stock market, categorized by types of *Illiquid Share*. As predicted by Hypothesis 1, Fig. 2 illustrates a non-linear pattern, while Figs. 3 and 4 show a positive, linear relationship, which contradicts Hypothesis 2.

These scatterplots lack consideration of the selection bias resulting from stock market participation. If non-participants were included, it would introduce numerous $0s$ for the *Risky Share*, potentially distorting the results. However, solely focussing on stock market participants also fails to offer a comprehensive view as it overlooks systematic differences between those who participate and those who do not. Another critical aspect is the absence of control for various factors that could influence both shares. To address these issues, the following section outlines the empirical design, aimed at mitigating these sources of bias.

4. Empirical strategy

The insights from the Stokey (2009) model suggest that we test for the presence and linearity of the relationship between the *Risky Share* and the *Illiquid Share* across owner-occupied and investment property. The identification strategy consists of first accounting for selection into the stock market and then use variations of a fixed effects model for the different types of *Illiquid Share*.

4.1. Accounting for selection into the stock market

As a classical problem of selection (Heckman, 1979), we can observe stockholdings only for stock market participants. Estimating our model only on stock market participants would therefore lead to a bias, specifically:

$$\mathbb{E}(s_{i,w}|X_{i,w}, D_{i,w} \geq 0) = \beta_x' X_{i,w} + \mathbb{E}(v_{i,w} > -\beta_h' H_{i,w}) \tag{1}$$

where $s_{i,w}$ is the *Risky Share* for household i in wave w, $X_{i,w}$ is the vector of explanatory variables for stock market participants, $D_{i,w}$ is the latent variable such that $s_{i,w}$ is observed if $D_{i,w} \geq 0$ (0 is used as a threshold without loss of generality), $v_{i,w}$ is the error term in the selection equation and $H_{i,w}$ the vector of explanatory variables in the selection equation. The bias term exists as long as $Cov(\varepsilon_{i,w}, v_{i,w}) \neq 0 \Leftrightarrow \mathbb{E}(v_{i,w} > -\beta' H_{i,w}) \neq 0$ (with $\varepsilon_{i,w}$ being the error term in the selected model), that is, if stock market participants are different from non participants, which is the empirical consensus (Campbell, 2006; Gomes et al., 2021).

To account for selection, we first model this as a probit describing the probability of participating in the stock market as a function of educational attainment, age and whether the household has positive or negative net worth (these characteristics are represented by $H_{i,w}$):

$$\mathbb{P}(S_{i,w} = 1 | H_{i,w}) = \Phi(\beta_h' H_{i,w}) \quad (2)$$

where $S_{i,w}$ indicates participation in the stock market and $\Phi(\bullet)$ is the standard normal cumulative distribution function. Participation is defined for those households and periods in which the household held stocks, or a period in which they held no stocks but following a period in which they held stocks (so that 0, in such a case, does not represent non-selection but an actual value of 0 for stockholdings). Notice that this is a pooled regression, so a same household is treated as a different observation in each time period, relying on the assumption that selection into the stock-market can be predicted without relying on time-invariant characteristics; a fixed effects panel probit regression would also suffer from the incidental parameters problem and lead to biased estimation (Greene, 2004). Our selection variables are three that find strong support in the literature. Higher educational attainment, as well as correlated concepts such as financial literacy, have been associated to a higher likelihood of stock market participation (Black et al., 2018; van Rooij et al., 2011; Calvet et al., 2007). Age is also associated with participation, although the direction of association is not established in a conclusive manner (Ameriks and Zeldes, 2004; Fagereng et al., 2017), while net worth is associated with higher participation and captures both wealth effects (wealth is intimately connected to stock market participation, supporting a financial cost barrier explanation, see Campbell, 2006; Guiso and Sodini, 2013) and effects related to being in debt, such as committed expenditure risk (Fratantoni, 2001; Cocco, 2005; Yao and Zhang, 2005). As a next step, we recover the inverse Mills ratio as:

$$\lambda_{i,w} = \frac{\phi(\beta_h' H_{i,w})}{\Phi(\beta_h' H_{i,w})} \quad (3)$$

with $\phi(\bullet)$ the standard normal density function. The inverse Mills ratio will be used to correct the selection bias in the main estimation equation.

Households facing liquidity constraints are both less likely to invest in stocks and less able to obtain a mortgage to buy a property. Because of this, many of these households are likely to be included in the subset with values of zero both for the Risky Share and for the Illiquid Share, resulting in a selection effect in both shares. Although it is not part of our analysis, one could correct for the probability to participate in stock and property markets based on an observed measure of liquidity constraint, such as a wealth of zero or a very small wealth to income ratio (Zeldes, 1989).

4.2. Testing the relationship between the Risky Share and the Illiquid Share

Our baseline specification tests the linearity of the relationship between the *Risky Share* and the *Illiquid Share*:

$$s_{i,w} = \beta_p p_{i,w} + \beta_{p^2} p_{i,w}^2 + \underline{\beta}_P P_{i,w} + \beta_{FTSE} FTSE_t + w_i + f_i + \lambda_{i,w} + \varepsilon_{i,w} \quad (4)$$

That is, each combination of household i in wave w constitutes an observation, and the *Risky Share* for the household in a specific wave ($s_{i,w}$) is a function of a second order polynomial for the *Illiquid Share* ($p_{i,w}$), the vector of household characteristics ($P_{i,w}$), the close value of the FTSE 100 on the first day of the month of survey ($FTSE_t$), the wave fixed effects (w_i) and the household fixed effects (f_i). $\lambda_{i,w}$ is the inverse Mills ratio computed in Section 4.1 and $\varepsilon_{i,w}$ is the error term. The vector of household characteristics $P_{i,w}$ includes household net worth, as well as educational level and age band of the Household Representative Person (HRP). The educational level is a binary variable indicating whether the HRP has a degree or not, while there are seven household age bands: 16–24, 25–34, 35–44, 45–54, 55–64, 65–74, 75+.

Incorporating household fixed effects allows us to account for time-invariant characteristics of households that are not directly observable, such as constant risk preferences. On the other hand, certain non-time-invariant variables, such as age, education, and net worth, could influence both the *Risky Share* and the *Illiquid Share*, as observed in the case of participation in Section 4.1. Net worth is expressed in July 2015 £, rather than being represented as a binary variable indicating positive or negative net worth (as in the selection Eq. (2)). This approach better captures the wealth effects on the intensive margin of stockholdings. Additionally, the wave of the survey in which the response was obtained is included to account for generic time-specific events.

Stock market performance is included to directly capture its impact on the *Risky Share*. Although assumed not to affect the *Illiquid Share*, including this variable can enhance the precision of the estimate.[2] Certain variables, such as measures of housing market performance or the place of residence, are excluded from the model. These are believed to influence the *Illiquid Share* but not the *Risky Share*.[3] These considerations are summarized in the Directed Acyclical Graph (DAG)[4] presented in Fig. 5.

As discussed in Section 2, there are theoretical grounds to believe that the impact of property on portfolio choice may differ depending on whether a household resides in it or not. The consumption motive is present when a property is owner-occupied, but it is eliminated or relaxed (in the case of holiday homes) when it is not lived in. To investigate this difference, we decouple the share of owner-occupied and investment property, and we additionally focus on a subset of non-owner-occupied property that is strictly exempt from the consumption motive—buy-to-let property. In Eq. (4), the terms $\beta_p p_{i,w}$ and $\beta_{p^2} p_{i,w}^2$ will refer in turn to the owner-occupied *Illiquid Share*, to the investment *Illiquid Share*, and to the buy-to-Let *Illiquid Share*.

In all cases, the presence of selection would lead to inconsistent standard errors (Heckman, 1979), so we resort to bootstrapping, clustered at the household level (200 repetitions). To exploit as much information as possible, the panel we use is unbalanced (18,509 households with on average 3.6 observations per household in the full sample and 7536 households, with on average 2.8 observations per household, in the sample of stock market participants). To ensure that the imbalance does not affect estimation, we test for the randomness of household attrition (Verbeek and Nijman, 1992), as described in Wooldridge (2010): by adding a lagged or leading selection indicator to model (4) estimated on the full sample, therefore without the inverse Mills ratio.

Survey weights are only used in robustness checks but not in the main model estimates.[5] This was motivated by the fact that the weighting aimed primarily at accounting for oversampling of wealthy households, but we consider this bias already by controlling for net worth (Solon et al., 2015).

[2] Note that for a same wave, the month of the survey varies across households, so that the variation in stock returns cannot be fully captured by wave fixed effects.

[3] To test the robustness of our findings concerning the possibility that regional variations in house prices might lead to greater appreciation and subsequently more liquidity to invest, we estimate the primary model while also controlling for the regional house price index (HPI) in the survey month (see Appendix D). The inclusion of the HPI has minimal impact on the estimates.

[4] For an introduction to DAGs, refer to Cunningham (2021). For a comprehensive discussion, see Pearl (2009).

[5] We ran two different robustness checks to ensure that weighting does not affect our main results. Since we need a stable weight throughout the waves for each household, the first check uses the weight assigned to the household in its first wave, while the second check uses the weight assigned to the household in its last wave. Result tables and charts are reported in Appendix E.

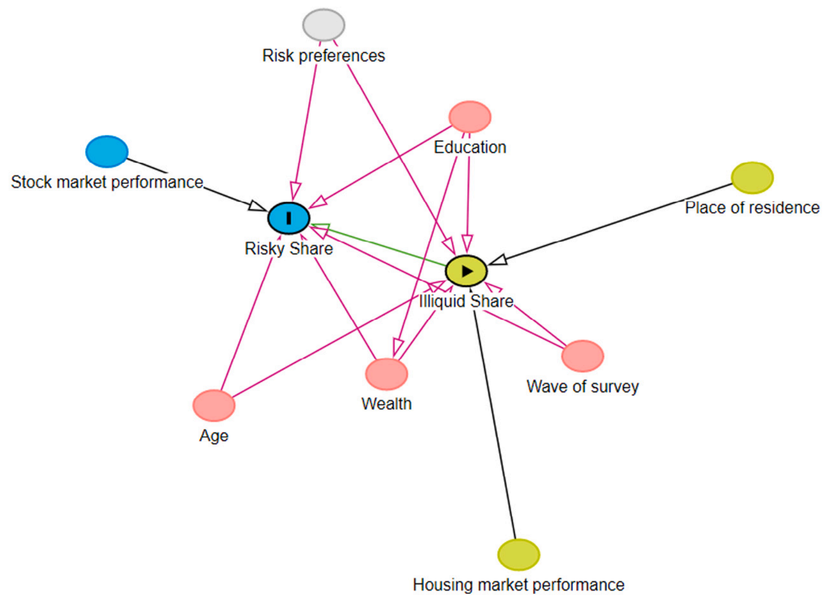

Fig. 5. Directed Acyclical Graph for the baseline model. Note: Directed Acyclical Graph describing the model in Eq. (4), created using DAGitty. The green arrow identifies the relationship of interest. Purple arrows represent confounding effects, while black arrows refer to effects that do not affect our relationship of interest. Pink nodes are observed confounders, while grey nodes unobserved ones. Finally, blue nodes and yellow nodes are parent variables of the *Risky Share* and of the *Illiquid Share* respectively that do not affect the relationship between the two. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.3. Intrroducing homeowners and renters

Another way to formulate the model is to interact the overall *Illiquid Share* with an indicator variable that distinguishes three categories: homeowners without investment property, homeowners with investment property and renters with investment property (excluding renters without investment property as the base category in the estimation). This can be formulated as:

$$s_{i,w} = \beta_p p_{i,w} + \beta_{p^2} p^2_{i,w} + (\beta_{pI} p_{i,w} + \beta_{p^2 I} p^2_{i,w}) I_c + \underline{\beta}_P P_{i,w} + \beta_{FTSE} FTSE_t + w_i + f_i + \lambda_{i,w} + \varepsilon_{i,w} \quad (5)$$

where $p_{i,w}$ in this case represents the overall *Illiquid Share* and I_c is the indicator variable for the category of *Illiquid Share*. Framing the model in this fashion allows to adopt a different angle: rather than looking at how much the *Risky Share* increases or decreases for a $1pp$ increase in the share for owner-occupied or investment property across their distribution (as in Section 4.2), we look at the variation in *Risky Share* for a $1pp$ increase in the *Illiquid Share*, across its distribution, for those who rent but still have property, for those who owner-occupy without owning additional property and for those who owner-occupy and own investment property. In other words, the perspective shifts from the type of investment to the type of investor. For homeowners, we further distinguish between outright homeowners and mortgagors, to check if the presence of debt has any implications. The same econometric considerations made in Section 4.2 also apply here.

5. Results and discussion

This section addresses first the results from the model presented in Section 4.2, which looked at the different types of *Illiquid Share*, to then move on to an alternative specification that considers the type of investor instead and matches the model in Section 4.3. Appendix A reports the coefficients of controls for the main estimation, and compares the results to the case excluding controls. We do not report the results for Section 4.1 separately, which accounts for selection into the stock market, and only include the coefficient for $\lambda_{i,w}$, the inverse Mills ratio in the tables. An estimation including those with 0 stockholdings is reported in Table 6 in Appendix B. Robustness checks using different weighting specifications are reported in Appendix E and referred to in the text.

5.1. The heterogeneous relationship of property and stockholdings

Table 2 presents the results for model (4), where the *Illiquid Share* corresponds to owner-occupied (1), investment (2), and buy-to-Let (3) property. Upon examination, we observe that the first-order term is hardly distinguishable from zero for all three types of the illiquid portion of household portfolios. However, the square term exhibits a negative value and is clearly distinguishable from zero for owner-occupied property, to a lesser extent for investment property, and once again, it cannot be differentiated from zero in the case of buy-to-let property.

These findings provide evidence of a concave pattern in the relationship between the *Illiquid Share* level and the *Risky Share* for owner-occupied property. This concave relationship is also broadly observed for investment property but it does not apply to buy-to-let property. These results align with our Hypothesis 1 and, to some extent, with Hypothesis 2. Notably, a $1pp$ increase in the *Illiquid Share* for owner-occupied property does not yield a constant effect across the distribution of the *Illiquid Share*, leading to variations in the marginal effect (Fig. 6), and an n-shaped trajectory for the *Risky Share* (Fig. 7).

Fig. 6 and Table 7 (Appendix C) reveal how the estimated marginal effect hovers around zero for low levels of the owner-occupied share and then steadily becomes more negative as the share increases. While on average a $1pp$ increase in the owner-occupied *Illiquid Share* is associated with a $0.07pp$ decrease in the *Risky Share*, when the owner-occupied *Illiquid Share* reaches 90%, a $1pp$ increase corresponds to a $0.14pp$ decrease in the *Risky Share*. This finding aligns with the prediction that risk aversion in the value function is contingent on the ratio of owner-occupied property to overall wealth. Once the model is calibrated, both (Stokey, 2009) and earlier work by Grossman and Laroque (1990) find that the *Risky Share*, as a function of the *Illiquid Share*, is higher when the property-to-wealth ratio approaches the thresholds that trigger adjustments and lower when it deviates from those thresholds (i.e., immediately after a property adjustment, when it is closer to the household's preferred ratio). The thresholds and ideal ratio are endogenously determined in the model and are influenced, among other factors, by the coefficient of relative risk aversion and by the size of the adjustment costs, which are likely to vary across households. Consequently, these thresholds and ideal ratios are expected

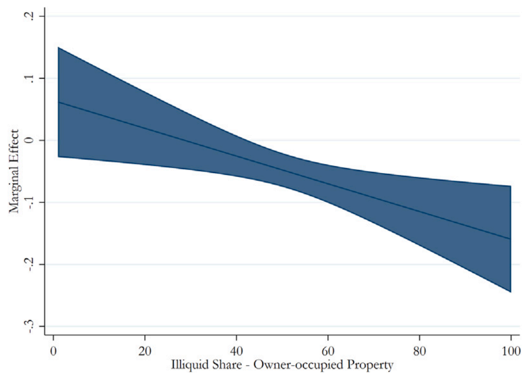

Fig. 6. Marginal effect of the owner-occupied *Illiquid Share* on the *Risky Share*, with 95% CI. Note: Fig. 6 plots the marginal effects and 95% confidence intervals at each percentage level of the owner-occupied *Illiquid Share*, based on model specification (4).

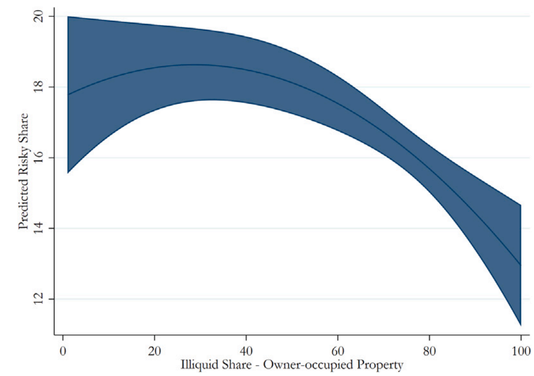

Fig. 7. Predicted *Risky Share* by level of the owner-occupied *Illiquid Share*, with 95% CI. Note: Fig. 7 plots the fitted values and 95% confidence intervals for the *Risky Share* at each percentage level of the owner-occupied *Illiquid Share*.

Table 2
Association of the owner-occupied, investment and buy-to-let *Illiquid Share* with the *Risky Share*.

	(1)	(2)	(3)
Owner-occupied Share	0.0641 (0.0452)		
Investment Share		0.0820* (0.0454)	
Buy-to-let Share			−0.0351 (0.0666)
Owner-occupied Share Sq.	−0.00112*** (0.000408)		
Investment Share Sq.		−0.00144** (0.000651)	
Buy-to-let Share Sq.			0.000523 (0.00101)
Inverse Mills Ratio	−4.174 (3.577)	−2.485 (3.106)	0.335 (5.854)
Observations	21 041	21 041	19 075
Adj. R-squared	0.3612	0.3594	0.3862
Household FE	YES	YES	YES
Wave FE	YES	YES	YES

Note: All three models refer to specification (4), varying the type of *Illiquid Share*. Additional controls include household net worth, HRP's age band and educational attainment, and the FTSE 100 monthly performance. Standard errors (in parentheses) are bootstrapped, clustered at the household level (200 repetitions). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

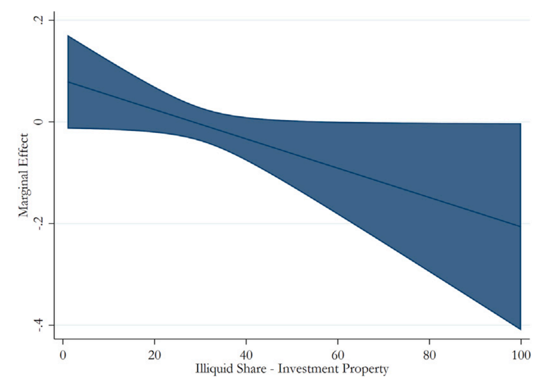

Fig. 8. Marginal effect of the investment *Illiquid Share* on the *Risky Share*, with 95% CI. Note: Fig. 8 plots the marginal effects and 95% confidence intervals at each percentage level of the investment *Illiquid Share*, based on model specification (4).

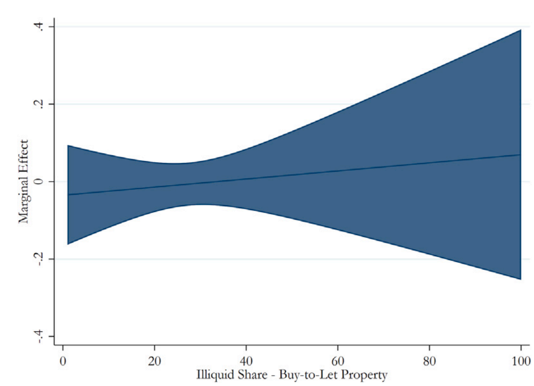

Fig. 9. Marginal effect of the buy-to-let *Illiquid Share* on the *Risky Share*, with 95% CI. Note: Fig. 9 plots the marginal effects and 95% confidence intervals at each percentage level of the buy-to-let *Illiquid Share*, based on model specification (4).

to differ among households. Although we cannot directly observe the exact pattern described theoretically in Stokey (2009) and Grossman and Laroque (1990), we nonetheless observe a non-linear relationship.

If we assume that the results of those studies hold, that is that risk aversion is highest when the ratio of property to wealth is close to the ideal and lower when the ratio moves towards the thresholds, then we could conclude that the ideal ratio in the data is to be found above 40%, and that the higher above 40%, the higher the share of households who have that as an ideal ratio, or alternatively that the higher the ratio, the closer we get to the ideal ratio on average, and therefore the higher is the risk aversion. While our expectation from the theoretical framework was that no relationship exists between investment property and stockholdings, the answer from the model is more complex. When we take a broad definition of investment property, that is everything outside owner-occupied property, we see a similar relationship as for owner-occupied property (Fig. 8), but the marginal effect cannot be as precisely estimated and might in fact be 0 all throughout.

Limiting the definition of investment property to buy-to-let (Fig. 9) changes the picture: the marginal effect becomes more uncertain, and a clear relationship cannot be established. The dataset still contains a considerable number of observations, 19,075 in total, of which 11.13% have a non-zero value for the buy-to-let share (compared to 24.18% for the investment share). While our theoretical expectation suggested no relationship between the *Risky Share* and the *Illiquid Share* for investment property, and an initial examination of the joint distribution of the data in Figs. 3 and 4 hinted at a linear relationship, it appears that no detectable relationship exists for the strict definition of investment property (buy-to-let), but that in its looser definition, investment property is somewhat related to stockholdings and appears to behave in a similar, albeit weaker, fashion as compared to owner-occupied property.

Table 3
Tests for imbalance in the sample.

Illiquid Share	Indicator	χ^2	*p*-value
Owner-occupied	Lag	2.03	0.1538
	Lead	0.01	0.9056
Investment	Lag	2.57	0.1088
	Lead	<0.01	0.9962
Buy-to-let	Lag	2.53	0.1120
	Lead	2.01	0.1559

Two possible explanations for this observation arise. Firstly, the looser definition of investment property includes properties, such as second homes, that have both consumption and investment aspects, possibly showing intermediate properties between the two aspects. In contrast, buy-to-let properties have a pure investment nature. Secondly, agents may adjust their stockholdings based on the covariance structure of expected returns, as suggested by Pelizzon and Weber (2008). Their extension of the classic mean–variance optimality condition introduces an additional linear term that adapts risky shares based on the correlation between property and risky financial assets. When the correlation between property and risky financial assets is zero, this additional term vanishes. Testing this formally would require considering the joint correlation structure of asset returns. However, this explanation seems less plausible due to the similarity in patterns between the looser definition of investment property and owner-occupied property (which suggests a connection to consumption), and no relationship in the case of buy-to-let, where the investment nature is more pronounced. Additionally, most households may lack the ability or time to consider covariance structures in expected returns, as evidenced by the general lack of sound financial knowledge found in the financial literacy literature, including in the UK (Gomes et al., 2021; Hastings et al., 2013; Nicolini et al., 2013).

The robustness checks presented in Appendix E, which estimate the same models as those in Table 2 using different weighting strategies, confirm the general patterns of the baseline results. They also suggest the possibility of a positive relationship between the owner-occupied *Illiquid Share* and the *Risky Share* at low levels of the *Illiquid Share*. Our findings align with previous empirical literature while contributing additional insights by highlighting the importance of investigating the heterogeneity of the relationship between property and stockholdings. We confirm the significance of owner-occupied property in stockholdings decisions, while also revealing a richer distribution for this association that potentially spans from positive to negative.

The coefficient on the inverse Mills ratio cannot be distinguished from zero in any of the models, indicating that the covariates already account for selection into the stock market (especially net worth). To rule out potential bias from the unbalanced panel, we estimated model (4) on the full sample, excluding the inverse Mills ratio, and incorporating lagged and leading selection indicators following Wooldridge (2010). For all variations of model (4) using different specifications of the *Illiquid Share*, we could not reject the hypothesis that the coefficients on the lagged or leading indicators are equal to zero (Table 3). Thus, imbalance in the sample is unlikely to be a source of bias in the estimates.

5.2. Does investor type matter?

An alternative approach to exploring the same issue involves analysing the marginal effect of the overall *Illiquid Share* for different types of investors. Instead of breaking down the share into owner-occupied, investment, or buy-to-let components, we focus on how different groups of investors respond to an increase in the *Illiquid Share*. Specifically, we examine homeowners who solely own their main residence, homeowners who own both their main residence and investment property, and renters who own investment property.

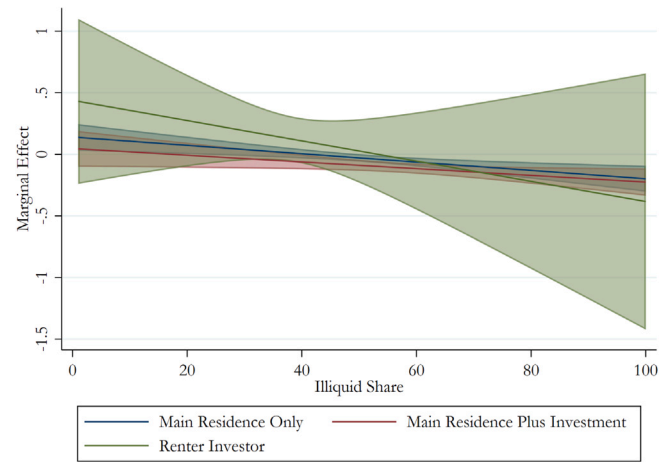

Fig. 10. Marginal effect of the *Illiquid Share* on the *Risky Share* by tenure and investor type, with 95% CI. Note: Fig. 10 plots the marginal effects and 95% confidence intervals at each percentage level of the *Illiquid Share*, by tenure and investor type, based on model specification (5).

To distinguish these three groups more clearly, we classify a household as belonging to the renters group if it was consistently a renter throughout the sample period and held non-owner-occupied property in at least one period. The homeowner without property investment category includes households that owner-occupied for at least one period but never held non-owner-occupied property. Finally, the homeowner with investment property category comprises households that owner-occupied for at least one period and also held non-owner-occupied property for at least one period. Fig. 10 and the first three columns of Table 8 (Appendix C) illustrate the marginal effects for model (5) across these three groups.

While the estimate for renters is challenging to pin-point (likely due, in part, to statistical power, as the renter group consists of only 613 observations, with 324 having a non-zero value for the *Illiquid Share*), the results for the two types of homeowners are similar to each other and comparable to those of the owner-occupied share in Fig. 6. These findings further support the hypothesis that the consumption motive drives the relationship between property and stockholdings.

A potentially pertinent distinction within the category of homeowners is that between those with and those without a mortgage on their main residence (whilst allowing a household to change categories between waves). This additional analysis aims to determine if the relationship between the *Risky Share* and the *Illiquid Share* depends on whether the household is leveraged, beyond what is accounted for by net worth (included as a covariate). Fig. 11 and the last two columns of Table 8 illustrate that the *Risky Share* exhibits a similar relationship with the *Illiquid Share* regardless of whether the household has a mortgage on its main residence. This finding aligns with the theoretical framework of Section 2 and provides support for an adjustment cost explanation for the relationship between the *Risky Share* and the *Illiquid Share*, which remains independent of the committed expenditure implied by holding a mortgage.

6. Conclusion

The portfolio choice model incorporating adjustment costs predicts a non-linear relationship between the *Risky Share* and the *Illiquid Share* when the adjustment costs for housing bind, and no relationship when they do not. Based on this framework, we hypothesised that the adjustment costs should only bind for owner-occupied property, resulting in a heterogeneous relationship between its share and that of stockholdings, depending on the level of the share. Conversely, for investment property, where adjustment costs are less significant, we expect no

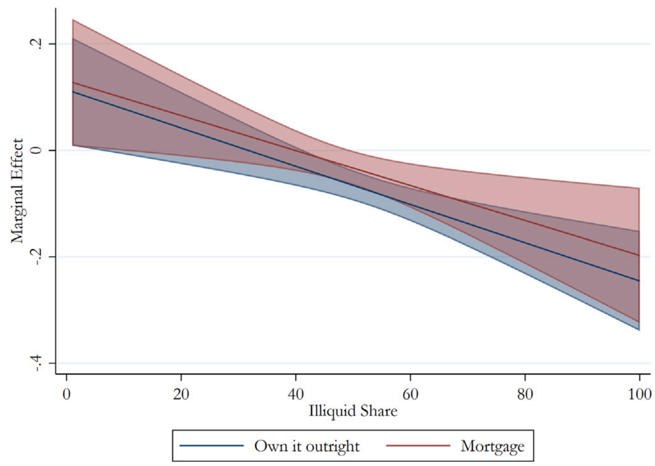

Fig. 11. Marginal effect of the *Illiquid Share* on the *Risky Share* by homeownership type, with 95% CI. Note: Fig. 11 plots the marginal effects and 95% confidence intervals at each percentage level of the *Illiquid Share*, by type of homeownership, based on model specification (5).

relationship. To test these hypotheses while accounting for selection into the stock market, we employed linear fixed effects models at the household level, incorporating a squared term to capture potential curvatures in the relationship between the *Risky Share* and the *Illiquid Share*. The analysis was performed on data from the Wealth and Assets Survey for the UK, which tracks households over time and provides detailed information on their asset distribution, including property.

Our findings suggest that the relationship between the *Risky Share* and the *Illiquid Share* for owner-occupied property is not constant and varies depending on the level of the *Illiquid Share*. At low levels of the *Illiquid Share*, the relationship is close to zero or slightly positive, transitioning to a more negative relationship at higher levels. These results hold true for both homeowners holding only a main residence and for those holding an additional non-main residence property. Furthermore, the relationship remains consistent for both homeowners with and without a mortgage. However, there seems to be no discernible relationship between the *Risky Share* and the *Illiquid Share* for a strict definition of investment property, namely buy-to-let property. For the broader definition of investment property (excluding owner-occupied property), a similar relationship to owner-occupied property exists, although it is weaker.

One limitation of the present study is that the value of the house is self-reported which might not correspond accurately to the true market value of the property. Nevertheless, this should not significantly affect the results since household decisions are arguably driven to a larger degree by perceptions of worth than by the intrinsic or true market value of the property. Furthermore, this study refrains from making causal claims since we cannot rule out additional confounding or endogeneity problems in the *Illiquid Share*. While our results may not establish a causal effect of the *Illiquid Share* on the *Risky Share*, they do provide evidence, consistent with theory, regarding the heterogeneous nature of the relationship and the significance of owner-occupied property in comparison to investment property.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The authors report financial support by the School of Humanities and Social Sciences, University of Cambridge to pay for research assistance and dissemination.

Data availability

The Wealth and Assets Survey is provided by the UK Data Service. The close values for the FTSE 100 are provided by Yahoo! Finance. Data on the regional House Price Index are provided by HM Land Registry. Data on the Consumer Prices Index including owner occupiers' housing costs (CPIH), used to adjust nominal values, are provided by the Office for National Statistics.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jhe.2023.101964.

References

Ameriks, J., Zeldes, S.P., 2004. How do Household Portfolio Shares Vary with Age. Technical report, Columbia University.
Beaubrun-Diant, K.E., Maury, T.-P., 2016. Home tenure, stock market participation, and composition of the household portfolio. J. Hous. Econ. 32, 1–17.
Black, S.E., Devereux, P.J., Lundborg, P., Majlesi, K., 2018. Learning to take risks? The effect of education on risk-taking in financial markets. Rev. Financ. 22 (3), 951–975.
Brueckner, J., 1997. Consumption and investment motives and the portfolio choices of homeowners. J. Real Estate Financ. Econ. 15 (2), 159–180.
Calvet, L.E., Campbell, J.Y., Sodini, P., 2007. Down or out: Assessing the welfare costs of household investment mistakes. J. Polit. Econ. 115 (5), 707–747.
Campbell, J.Y., 2006. Household finance. J. Finance 61 (4), 1553–1604.
Chang, Y., Hong, J.H., Karabarbounis, M., Wang, Y., Zhang, T., 2022. Income volatility and portfolio choices. Rev. Econ. Dyn. 44, 65–90.
Chetty, R., Sándor, L., Szeidl, A., 2017. The effect of housing on portfolio choice. J. Finance 72 (3), 1171–1212.
Chetty, R., Szeidl, A., 2007. Consumption commitments and risk preferences. Q. J. Econ. 122 (2), 831–877.
Cocco, J.F., 2005. Portfolio choice in the presence of housing. Rev. Financ. Stud. 18 (2), 535–567.
Cunningham, S., 2021. Causal Inference: The Mixtape. Yale University Press.
Fagereng, A., Gottlieb, C., Guiso, L., 2017. Asset market participation and portfolio choice over the life-cycle. J. Finance 72 (2), 705–750.
Fratantoni, M.C., 2001. Homeownership, committed expenditure risk, and the stockholding puzzle. Oxf. Econ. Pap. 53 (2), 241–259.
Gomes, F., Haliassos, M., Ramadorai, T., 2021. Household finance. J. Econ. Lit. 59 (3), 919–1000.
Greene, W., 2004. Fixed effects and bias due to the incidental parameters problem in the tobit model. Econometric Rev. 23 (2), 125–147.
Grossman, S., Laroque, G., 1990. Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. Econometrica 58 (1), 25–51.
Guiso, L., Sodini, P., 2013. Household finance: An emerging field. In: Handbook of the Economics of Finance, Vol. 2. Elsevier, pp. 1397–1532.
Hastings, J.S., Madrian, B.C., Skimmyhorn, W.L., 2013. Financial literacy, financial education, and economic outcomes. Annu. Rev. Econ. 5 (1), 347–373.
Heckman, J.J., 1979. Sample selection bias as a specification error. Econometrica 47 (1), 153–161.
Hu, X., 2005. Portfolio choices for homeowners. J. Urban Econ. 58 (1), 114–136.
Kong, D., Cheng, Y., Liu, S., 2021. Unexpected housing wealth appreciation and stock market participation. J. Hous. Econ. 52, 101768.
Office for National Statistics, S.S.D., 2020. Wealth and Assets Survey, Waves 1-5 and Rounds 5-6, 2006–2018. [Data Collection], eleventh ed. UK Data Service, SN: 7215.
Nicolini, G., Cude, B.J., Chatterjee, S., 2013. Financial literacy: A comparative study across four countries. Int. J. Consum. Stud. 37 (6), 689–705.
Pearl, J., 2009. Causality. Cambridge University Press.
Pelizzon, L., Weber, G., 2008. Are household portfolios efficient? An analysis conditional on housing. J. Financ. Quant. Anal. 43 (2), 401–431.
Solon, G., Haider, S.J., Wooldridge, J.M., 2015. What are we weighting for? J. Hum. Resour. 50 (2), 301–316.
Stokey, N.L., 2009. Moving costs, nondurable consumption and portfolio choice. J. Econom. Theory 144 (6), 2419–2439.
van Rooij, M., Lusardi, A., Alessie, R., 2011. Financial literacy and stock market participation. J. Financ. Econ. 101 (2), 449–472.
Verbeek, M., Nijman, T., 1992. Testing for selectivity bias in panel data models. Int. Econ. Rev. 681–703.
Vestman, R., 2019. Limited stock market participation among renters and homeowners. Rev. Financ. Stud. 32 (4), 1494–1535.
Wooldridge, J.M., 2010. Econometric Analysis of Cross Section and Panel Data. MIT Press.
Yao, R., Zhang, H.H., 2005. Optimal consumption and portfolio choices with risky housing and borrowing constraints. Rev. Financ. Stud. 18 (1), 197–239.
Zeldes, S.P., 1989. Consumption and liquidity constraints: an empirical investigation. J. Polit. Econ. 97 (2), 305–346.