BANK OF ENGLAND

Staff Working Paper No. 939

Preferred habitat investors in the UK government bond market

Julia Giese, Michael Joyce, Jack Meaning and Jack Worlidge

December 2023

This is an updated version of the Staff Working Paper originally published on 10 September 2021

BANK OF ENGLAND

Staff Working Paper No. 939

Preferred habitat investors in the UK government bond market

Julia Giese,[1] Michael Joyce,[2] Jack Meaning[3] and Jack Worlidge[4]

Abstract

Most tests of preferred habitat theory are indirect; they infer the existence of preferred habitat behaviour in financial markets by examining the behaviour of asset prices. We instead identify preferred habitat behaviour directly from whether investors show a preference towards a particular duration habitat. We do so by making use of a newly available and highly granular dataset on the UK government bond (gilt) market, which allows us to examine investors' gilt transactions and their daily stock of gilt holdings during 2016 and 2017. Using cluster analysis, we find that investors can be classified into distinct groups, some of which more closely display the behavioural properties that theory associates with preferred habitat investors. We find that these groups of investors are less sensitive to price movements than other investor groups and include institutional investors, like life insurers and pension funds, which are typically associated with preferred habitat behaviour. Evidence from the Bank of England's QE4 purchase programme during August 2016 to March 2017 suggests that these investor groups sold relatively more of their gilt holdings to the Bank than other groups of investors.

Key words: Preferred habitat, gilt market, yield curve, cluster analysis.

JEL classification: E43, E52, G11, G12.

(1) Bank of England. Email: julia.giese@bankofengland.co.uk
(2) Bank of England. Email: mike.joyce@bankofengland.co.uk
(3) Bank of England. Email: jack.meaning@bankofengland.co.uk
(4) Bank of England. Email: jack.worlidge@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful to Bank of England colleagues and especially Jeremy Franklin and Peter Spencer for helpful discussions at various stages of this research. We would also like to thank Benoit Nguyen and Marcello Pericoli for comprehensive comments on earlier versions of the paper and Dimitri Vayanos for insightful discussions throughout. We would like to thank Euroclear UK and International for providing us with their data. We are also grateful to Northern Trust for helpful discussions on the structure of the market and the CREST accounts system.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

"Whether central banks' large-scale asset purchases succeed in reducing term premiums hinges upon whether the preferred habitat hypothesis holds."

Haruhiko Kuroda, Governor of the Bank of Japan

1.Introduction

The origins of market segmentation and preferred habitat theories go back to the 1950s and 1960s and the seminal papers of Culbertson (1957) and Modigliani and Sutch (1966). The key idea is that investors value certain assets for reasons other than expected return or risk, so that they have a preference for certain asset classes or, in the case of bonds, maturities. This has the general implication that local demand and supply conditions can become important in price determination, as investors require compensation for changing their asset holdings in response to shocks.[1] In recent years, this theoretical result has become a central part of the narrative around the transmission of central bank quantitative easing (QE) policies through the so-called portfolio balance channel and the view that this mainly works by reducing term premia. More broadly, preferred habitat behaviour has important implications for a number of other literatures where the investor structure matters, including debt management (Andritzky, 2012), the role of different investors in absorbing or amplifying price shocks (Timmer, 2018) and in the price formation process (Koijen, Richmond and Yogo, 2020).

In this paper we examine preferred habitat behaviour in the UK government bond (gilt) market using a newly available, highly granular dataset provided by Euroclear UK and International, which contains detailed account-level information on gilt transactions and portfolio holdings. The gilt market is often thought of as being influenced by preferred habitat investors because of the large gilt holdings of institutional investors like pension funds and life insurers, so it provides a natural laboratory for examining this kind of behaviour.[2] Since we do not have information on other asset holdings outside the gilt market, preferred habitat is defined by an investor exhibiting a specific maturity preference for gilts. We interpret this in terms of the investor's preferred portfolio duration habitat, given that nearly all investors hold a portfolio of different bonds and that gilts are typically coupon bonds, so that duration provides a more appropriate measure of interest rate risk. A focus on duration is also consistent with the desire of some investors to match the long duration of their liabilities.

This paper makes three important contributions to the literature. The first is to classify investors according to their investment behaviour, based on their gilt transactions and gilt portfolio holdings during 2016 and 2017. Using cluster analysis applied to the average portfolio duration of individual investor accounts and the observed range of that average

[1] This requires there to be some market imperfection that prevents the activities of arbitrageurs offsetting the effects, see discussion below.

[2] Greenwood and Vayanos (2010) discuss how pension reforms in the UK induced a large increase in pension fund demand for long-dated gilts, which led to dramatic effects on the term structure of gilts yields in 2005 and 2006.

portfolio duration over the sample, we find that investors can be classified into distinct groups, some of which display the behavioural properties that theory would associate with preferred habitat investors. Moreover, these preferred habitat investors are present across the yield curve rather than being limited to any specific segment. By matching our data with other information sources, we also show that our identified groups of preferred habitat investors include institutional investors, like life insurers and pension funds, which are commonly thought of as exhibiting preferred habitat behaviour. This acts as a cross-check on our identification methods, but also provides some support for studies in the literature that use sectoral data to examine the impact of preferred habitat demand, by making the assumption that certain sectors can act as proxies for preferred habitat investors.

Our second contribution is to examine the other behaviours associated with preferred habitat investors. We estimate how sensitive the demand of our investor groups is to price changes, finding evidence that the price elasticities of our identified preferred habitat investors are lower than for other investor groups, as theory would suggest. We also find evidence that our preferred habitat investor groups turn over their security holdings less than other investor groups, which is another implication of the theory. As well as being interesting in their own right, both findings support our identification methods.

The third contribution of the paper is to use our investor group classifications to examine a case study; namely, the effects of the Bank of England's £60 billion of gilt purchases during August 2016 to March 2017, so-called QE4. Using our identified investor groups, it appears that groups identified as preferred habitat investors sold a greater proportion of their holdings to the Bank than other groups. These results appear consistent with the operation of the so-called portfolio balance channel but, because the holdings of our sample of identified preferred habitat investors represent a relatively small share of outstanding government debt, we cannot reliably infer how important the incidence of preferred habitat was in explaining the aggregate impact of QE4 across the yield curve.

The rest of this paper is structured as follows. Section 2 discusses how preferred habitat investors are described in theory and tested for in practice. Section 3 describes the properties of our main dataset and the sample we use in subsequent analysis. Section 4 discusses our clustering methodology and the main criterion we use to examine preferred habitat behaviour. Section 5 provides our empirical results, which identify four preferred habit investor clienteles from the data and, using other data sources, attempts to match which sectors they come from. Section 6 looks at the price sensitivity of the different investor categories we identify. Section 7 examines the Bank of England's gilt market purchases in August 2016 to March 2017 and contemporaneous changes in the holdings of different investor groups. Section 8 summarises our main conclusions.

2. Preferred habitat investors in theory

In order to understand preferred habitat investors in theory, the natural starting point is the classic papers by Culbertson (1957) and Modigliani and Sutch (1966), referred to in Section 1 above.

Culbertson's paper was intended to offer "a first approximation to a realistic theory of the behaviour of the term structure of interest rates", where "substitutability between short-term and long-term debt on the part of both borrowers and lenders … is limited in extent". This has the consequence that "when the maturity structure of debt supplied to the economy undergoes a substantial short-run change … this is reflected in the rate structure". Culbertson doesn't refer to preferred habitat investors explicitly, although his theory clearly relies on lenders and borrowers preferring securities with specific maturities.

In their paper on the Kennedy Administration's "Operation Twist", Modigliani and Sutch (1966) develop a model, referred to as Preferred Habitat theory, which is a modified version of the Market Segmentation Hypothesis due to Culbertson (1957). As well as incorporating risk premia (as in Hicks (1939)), Modigliani and Sutch allow for the fact that different agents may have different habitats, rather than just wanting to turn their portfolio into cash in the short term, raising the possibility of preferred habitats existing in different parts of the yield curve. Risk aversion leads investors to prefer to stay in their maturity habitat, unless other maturities compensate for the risk and cost of moving out of it. In this model, the risk premium can be negative or positive depending on supply and demand in the particular yield curve segment. As a result, the spread between long and short rates will be influenced by "the supply of long- and short-term securities by primary borrowers (i.e., by borrowers other than arbitrageurs, relative to the corresponding demand of primary lenders, to an extent reflecting prevailing risk aversion, transactions costs, and facilities for effective arbitrage operations)".

Although popular with market practitioners, preferred habitat theories were often criticised for their lack of micro foundations, and mainstream finance theory with its faith in the role of arbitrage and market efficiency tended to dismiss their relevance for asset pricing.[3] As a result, the topic of preferred habitat had largely been banished from serious academic research by the time of the global financial crisis in 2007-2009.[4]

[3] The empirical evidence for the hypothesis has also been mixed, with the perceived failure of the Fed's attempt to twist the shape of the yield curve in the 1960s ("Operation Twist") often cited as evidence against the theory (see Modigliani and Sutch (1966)). More recent analysis of this intervention has been more favourable (see Swanson (2011)).

[4] The literature on the optimal portfolio choice problem of a long horizon investor (e.g. Merton (1969, 1971 and 1973) and Stiglitz (1969)) can also be thought of as providing a theoretical grounding for long-term preferred habitats. This is made explicit by Wachter (2003), who shows that as risk aversion approaches infinity the optimal portfolio reduces to a bond maturing at the terminal point, which she claims formalizes the preferred habitat intuition of Modigliani and Sutch (1966). See also Riedel (2006) for another theoretical attempt to justify preferred habitat behaviour.

With that crisis, however, these theories were dusted off and often took centre stage when analysing the impact of central bank asset purchases (Quantitative Easing, aka QE), as alluded to by Kuroda (2015). In order for QE, or more generally innovations in the relative stocks of financial assets, to affect financial prices through a so-called portfolio balance effect (Tobin (1961, 1963, 1969) and Brunner and Meltzer (1973)) there must be some degree of imperfect substitutability between financial assets. Preferred habitat behaviour, whether motivated by regulation or investment mandates, provides an underpinning for expecting assets to be imperfectly substitutable and therefore has become closely associated with the view that asset purchases by central banks affect the term premia on government bonds. Taking this as its starting point, a large empirical literature has emerged looking at the financial market effects of QE and debt management shocks (for recent surveys see e.g. Bhattarai and Neely (2016) and Haldane et al (2016)).

At more or less the same time, an emerging theoretical literature was making progress in incorporating preferred habitat behaviour into equilibrium macroeconomic models. Although the authors do not explicitly reference preferred habitat, an early influential paper by Andres, Lopez-Salido and Nelson (2004) incorporates similar behaviour into a dynamic stochastic general equilibrium model by assuming there is asset market segmentation. More specifically, their model includes two sets of households: households who are restricted to only invest in long-term bonds, who can be thought of as preferred habitat investors, and unrestricted households who can invest in both short and long-term securities. With the important addition of portfolio adjustment costs, which provides a role for real money balances to matter, the model embodies an additional channel for the central bank to affect the real economy besides changing the expected path of future short rates. Similar approaches are developed in papers by Chen et al (2012), Harrison (2012), Sudo and Tanaka (2018) and Ray (2019).

In another influential paper with a modern finance perspective, Vayanos and Vila (2009, 2021) embed preferred habitat behaviour into a no-arbitrage model with heterogeneous agents. In their model, preferred habitat investors demand bonds with specific maturities and do not arbitrage across the term structure. That role is played by arbitrageurs whose carry trades across the term structure smooth through interest rate and demand shocks. Depending on the degree of risk aversion of arbitrageurs,[5] shocks to demand are either smoothed across the term structure or exhibit more localised effects. But, unless arbitrageurs are unconstrained and risk neutral, demand shocks from preferred investors can have an impact on yields through local supply effects and an additional duration channel, which stems from arbitrageurs requiring compensation for the amount of duration risk they face. If central banks are thought of as preferred habitat investors then the model has a simple read across for looking at the impact of QE. If central banks target their purchases at the holdings of

[5] The Vayanos and Vila paper is part of the larger 'limits to arbitrage' literature, which looks more generally at how real-world arbitrage may be constrained in practice by lack of capital and its implications for explaining various financial market anomalies (see the review in Gromb and Vayanos (2010)).

preferred habitat investors, there will need to be a significant fall in yields in order to encourage them to invest in other assets.

Despite the renewed interest in preferred habitat behaviour and the importance of it for understanding the effects of unconventional policy, there is little hard evidence on its prevalence in different financial markets. The origins of the theory of preferred habitat behaviour were largely theoretical, corroborated if at all with only indirect evidence from the behaviour of market prices, rather than with direct evidence on investor behaviour. More recent analysis, often examines the relationship between yields and aggregate measures of net debt supply (e.g. Greenwood and Vayanos (2014), Li and Wei (2014), Strohsal (2017)). Sometimes the demand of preferred habitat investors is proxied by the sectoral asset holdings of specific financial institutions (e.g. Zinna (2013), Fukunaga, Kato, and Koeda (2015), Boermans and Vermeulen (2018), Greenwood and Vissing-Jorgensen (2018), Kaminska and Zinna (2019), Koijen et al (2021)), but the assumption that these institutions - typically overseas official investors, insurance companies and pension funds - exhibit preferred habitat behaviour is not directly tested. One attempt to identify habitat preference from investor behaviour can be seen in Koijen et al (2017) who investigate the existence of home-bias in the Euro Area government bond market. This is a different form of habitat preference to the duration preference we will consider in this paper, however, and Koijen and co-authors only have access to data at the sectoral level. The availability of micro data on individual investor behaviour makes it possible for us to go further.

3. Data

As discussed above, in order to examine the behaviour of individual investors in the gilt market, we make use of a newly available dataset provided by Euroclear UK and International (henceforth EUI). As the central securities depository for all UK government securities, EUI has information on both transactions and stocks at a highly granular level. Specifically, we are able to observe the end of day stock holdings of each account in the CREST system[6] at the individual ISIN level and all settled transactions between accounts, again at the level of individual ISINs. When aggregated, the daily stock holdings covered by our dataset represent more than 99% of the total stock outstanding, on any given day.

The accounts that make up the CREST system are classified in one of three ways; proprietary accounts, segregated client accounts or omnibus accounts. In the first two instances, the account represents a single underlying investor. The distinction between them is that, in a proprietary account, the underlying owner of the securities also manages the account in the CREST system. In a segregated account, the underlying owner of securities employs a custodian to interact with the CREST system on their behalf, although they remain the owner

[6] The CREST system settles exchange-traded and Over-The-Counter (OTC) securities transactions for UK, Irish, Jersey, Guernsey and Isle of Man equities, warrants and covered warrants and for UK government bonds and money market instruments. It has been owned and operated by Euroclear UK and International since 2002.

of the securities and make the related investment decisions. In the case of an omnibus account, a single custodian will pool the securities of multiple clients into a single account, usually for reasons of efficiency. Within our dataset, 58% of the conventional gilt stock, by par value, is held in accounts that are either proprietary or client segregated and can therefore be safely associated with a single underlying investor.[7]

By combining the stock and transaction information for these accounts with publicly available data on the ISINs themselves, we are able to construct a range of portfolio characteristics for each investor through time, such as their weighted-average duration, the size and diversity of their portfolio and how frequently they trade. Summary statistics of a number of these variables are presented in Table 1.

Table 1: Summary statistics on conventional gilts in EUI dataset, Jan 2016-Dec 2017

	Full sample	**Excluding omnibus accounts**	**Excluding omnibus accounts and APF**	**Omnibus accounts**
# unique ISINs	48	48	48	48
		Data on end of day stocks		
# accounts	7,329	6,650	6,649	679
Average stock held per account (£mn)	159	119	66	553
Average # ISINs held by each account	3.7	3.1	3.1	9.2
Average weighted average duration	6.6	6.3	6.3	9.2
		Data on settled transactions		
# accounts that trade	4,761	1,917	1,916	2,844
Average # trades per day[8]	6,638	6,552	6,542	2,948
Quantity traded per day (£bn)	109	108	106	44

Our sample covers a two-year period between 4 January 2016 and 31 December 2017 and focuses on the conventional gilt market, which had a nominal value of just over £1 trillion at the end of 2017 and an average maturity of 13.75 years. Over the period we consider, we have 9.8 million day-ISIN-account observations and 3.4 million trades.

4. Identifying investor groups

4.1 A metric for duration preference

To assess whether or not there exist investors in our dataset with a preferred duration habitat, we need a metric with which to measure this preference. In the theoretical literature,

[7] While some institutions hold multiple accounts it seems safe to conclude that they follow different investment mandates and can therefore be viewed effectively as independent investors.

[8] Only trades for which a qualifying account is on at least one side of the trade are included.

preferred habitat investors are often defined in terms of their preference for holding bonds with specific maturities, which could be anywhere along the yield curve. In the highly stylised model of Vayanos and Vila (2009, 2019), this is taken to the extreme – by assuming that preferred habitat investors only demand the bond corresponding to their maturity habitat. But, as the authors acknowledge, this is solely for analytical convenience.

In reality, most investors are likely to hold a portfolio of bonds, as is confirmed by our data on individual accounts. It therefore makes sense to think of preferred habitat as relating to the maturity of the investor's portfolio, rather than that of a single bond. At the same time, gilts typically pay coupons, so that duration is a better measure of interest rate risk than maturity. This, together with the fact that some investors may be aiming to match the duration of their gilt portfolio to their liabilities, would suggest using average portfolio duration as the benchmark measure of preferred habitat.[9] It follows that those investors who target a particular maturity habitat will attempt to minimise fluctuations in their portfolio's average duration.

In our benchmark analysis, we measure the duration habitat preference of each investor by the weighted average duration of their gilt portfolio and the strength of this preference by the 10-90 percentile range of the weighted-average duration of the same portfolio over the sample. Conceptually, if the latter statistic is low then the investor does not allow the duration of their portfolio to vary much and could therefore be said to have a preferred habitat. If an investor allows the average duration of their portfolio to vary significantly from month to month then it would seem self-evident that they do not have a preference for achieving a specific time invariant duration.

It might be objected that we would naturally expect some investors to exhibit less variation in the duration of their portfolios than others, but this would not necessarily indicate the existence of preferred habitat behaviour. Rather it could be a result of a typical distribution of values. Figure 1 plots the frequency distribution of the weighted average duration range of each portfolio[10] and brings out the point that the data would be poorly described by fitting a single distribution; a point that the more formal analysis that follows confirms. The investors we identify as having a narrow range for their duration habitat are not simply those who would naturally appear in the tail of any single distribution, but rather are separable and distinct from the other investors. More formally, the methodology we use to identify investor groups does not impose multiple distributions, so if the data was most accurately represented by a single Gaussian distribution, it would return that result.

The results we will go on to present are robust to using alternative measures of the stability of the weighted average duration of an investor's portfolio, including the stricter 0-100

[9] We measure duration using modified duration, but the results are robust to using Macauley duration.
[10] For this analysis we exclude a number of accounts, as discussed below in Section 4.3.

range and the standard deviation, and also to using average portfolio maturity rather than duration.

Figure 1: Frequency distribution of weighted average duration range

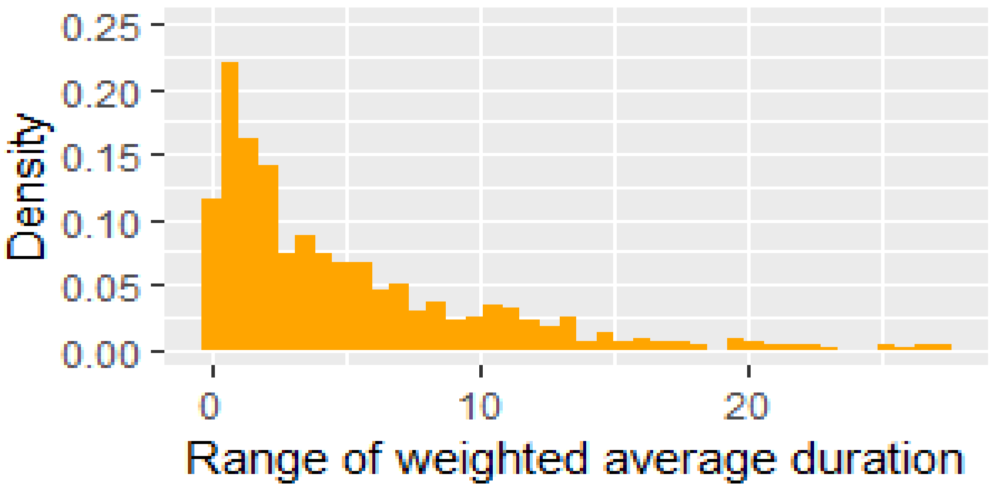

Notes: refers to 10-90 percentile range of investor accounts defined in text.

4.2 What habitats are investors targeting?

While the range of average duration can tell us whether or not an investor is operating within a narrow duration habitat, it does not tell us what that habitat is. For this, we need to augment our metric to include a location variable; for instance, the mean weighted average duration of each investor, around which the range is formed. Figure 2 plots each investor's mean weighted average duration against the range over which it varies. From this, it is clear that there is a wide range of target habitats, with investors collectively targeting all parts of the yield curve and to a more or less narrow extent. For instance, investor A (shown in the figure in red) has a mean weighted-average duration of 2 years and a range around this of a little less than 6 months. This means their implied preferred habitat is between 21 and 27 months. Alternatively, investor B has a mean weighted-average duration of 20 years and a range around that of approximately 2 years implying their preferred habitat is between 19 and 21 years. Investor C has the same mean weighted-average duration as B, but allows their portfolio duration to vary significantly more around that mean, ranging over the sample from between around 9 years and almost 30 years. This suggests investor C is happy for their portfolio to have a relatively flexible habitat.

Figure 2: Differing preferred habitats of investors

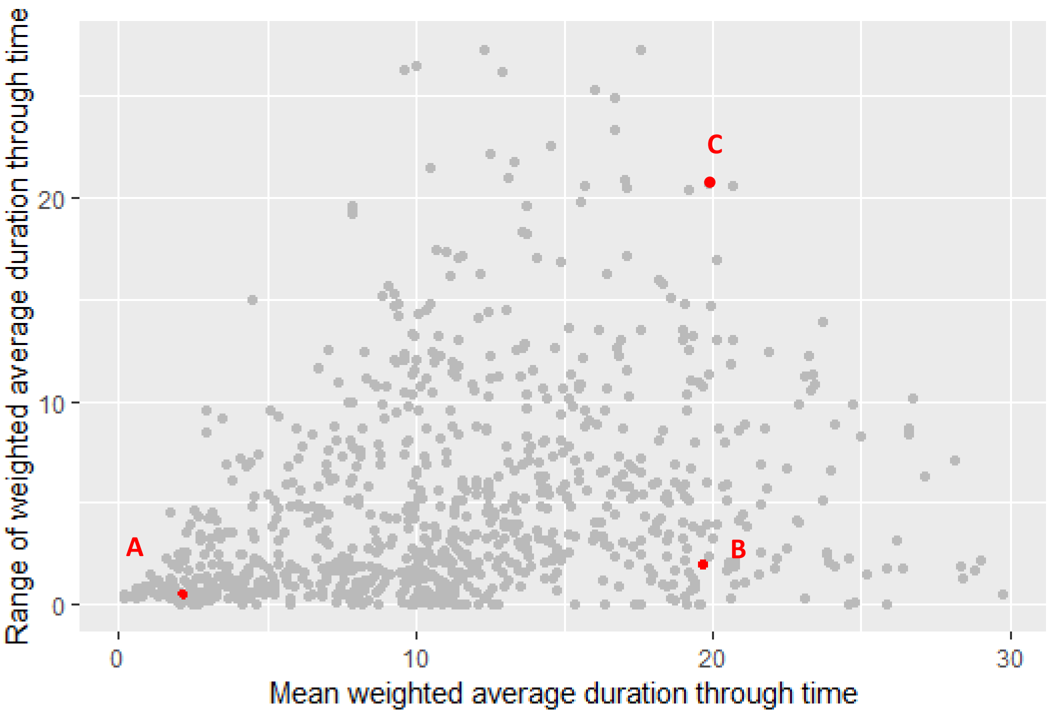

4.3 Identifying investor groups: Gaussian mixture model

To identify groups of 'similar' investors within our sample we employ a clustering algorithm: a Gaussian Mixture Model (GMM). This procedure models the data on the assumption that observations are generated from one of J underlying multivariate normal distributions. We chose GMM due to its relative flexibility compared to other popular clustering algorithms such as K-means. In particular, it allows for more flexible estimation by allowing for different variances across groups. While this procedure allows for the possibility of multiple groups, there is nothing inherent in the methodology that requires there to be multiple distributions in the data.

For each CREST account, we define a vector y_i of length D of the account's portfolio characteristics (D=2 in our benchmark case). Assuming that y_i is drawn iid from one of the set of J multivariate normal distributions, we pose the problem as a model (as in Lucas, Schaumburg and Schwaab, 2016) with latent discrete indicator variables z_j, where z_j takes the value one in the case of membership of an account with normal distribution j and zero otherwise. The model can then be written:

$$\underbrace{y_i}_{\substack{\text{vector of}\\ \text{D characteristics}}} = \underbrace{z_1}_{\substack{\text{Indicator if i is member}\\ \text{of cluster 1}}} . \left(\underbrace{\mu_1}_{\substack{\text{Mean}\\ \text{characteristics}\\ \text{cluster 1}}} + \underbrace{\Sigma_J^{-\frac{1}{2}}}_{\substack{\text{Covariance of}\\ \text{characteristics}\\ \text{cluster 1}}} \varepsilon_{i1} \right) + \dots + z_J . \left(\mu_J + \Sigma_J^{-\frac{1}{2}} \varepsilon_{iJ} \right)$$

where $\varepsilon_{ij} \sim N(0,1)$.

In order to estimate the mean and variance-covariance parameters, we maximise the log likelihood given the account's realised characteristics. Defining $\pi_j = \Pr(z_j)$, the log likelihood is:

$$\log(L(\Theta)) = \sum_{i=1}^{n} \log \sum_{j=1}^{J} \pi_j . \phi(y_i | \mu_j; \Sigma_j)$$

Subject to $\sum_{j=1}^{J} \pi_j = 1$ and where ϕ denotes the standard normal density. Because this is complex to solve numerically, we employ the expectation maximisation (EM) algorithm until the solution converges. As Dempster et al (1977) show, this algorithm strictly improves the log likelihood at each iteration. Because this can converge to a local optimum, it is common practice to run the algorithm many times with different starting values to gauge the stability of the optimum. To avoid a local solution, we repeated the exercise with different and random initial mean and variance-covariance parameters. To estimate the optimal number of clusters J, we employed the use of the Bayesian information criterion (BIC), such that our optimal number of groupings, J, maximises the likelihood function subject to a penalty for the possibility of overfitting, although we apply some judgement in considering the optimal suggestion of other information criterion.

We estimated the model for accounts based on the vector of characteristics described above: each account is described by the weighted average duration of the portfolio and the 90-10 quantile range of the weighted average duration over the sample period. We estimated the model many times and found that our results were not dependent on initial conditions. In addition, we found that the optimal number of groups J - that maximises the BIC - is 7. These groups are described in more detail below.

Before fitting our benchmark model we first removed a number of accounts from our sample. Most importantly, we exclude three groups of well-identified investors that have specific and unique motivations for their behaviour in the gilt market: the Bank of England (which is dominated by the Asset Purchase Facility), the Debt Management Office and Gilt-Edged Market Makers (GEMMs). We also remove omnibus accounts as the diversity of the underlying multiple investor base means that these accounts' aggregate portfolio behaviour cannot be interpreted safely as that of a single, representative investor.[11] Ultimately, we fit

[11] We also remove a tail of small and inactive accounts that hold fewer than 3 ISINs in sum across the entire sample and rarely trade (fewer than 4 trades in the entire 2 years). These accounts are small in terms of the stock holdings they represent. Our key results are robust to their inclusion, but they add additional noise. We also ran the clustering algorithm with varying degrees of strictness for removing inactive accounts, or particularly noisy and volatile accounts, as an additional check to confirm the robustness of the central grouping.

our GMM to data on 926 accounts, which hold around 13% of outstanding stock, on average. When combined with the holdings of the Bank of England, the Debt Management Office and the GEMMs, this means our classified investor groups hold, on average, 58% of the total stock of conventional gilts.

5. Identified groups of investors

5.1 Preferred habitat investors

The resulting clusters are shown in Figure 3 below, for our benchmark estimation with seven groups. We find that investors can be classified into distinct groups, some of which more closely display the behavioural properties theory associates with preferred habitat investors. Four groups of investors exhibit varying degrees of preferred habitat behaviour. One distinct group (labelled STPHI, for short-term preferred habitat investors) focuses on the short end of the yield curve (less than 5 years duration). Two other groups focus more on the medium term segment of the yield curve: less than 15 years duration (labelled MIDPHI, for medium-term preferred habitat investors) and duration in a very narrow arrange around 10-12 years (labelled MIDPHI2, for medium-term preferred habitat investors). There is one group of preferred habitat investors that mainly target a duration at the long end of the yield curve (on average greater than 15 years, labelled LTPHI, for long-term preferred habitat investors). The three other investor groups identified exhibit much larger variation in their portfolio durations, consistent with "arbitrageur" behaviour, which we denote STARB, MIDARB and LTARB to reflect their relative preferences in terms of duration space (short-term, medium-term and long-term).

Table 2 shows summary statistics for each of these seven groups. Preferred habitat investors at the short end of the curve (STPHI) maintain a very narrow weighted-average duration over the sample period. This would suggest an extreme preferred habitat from which they do not deviate by more than plus or minus 3 ½ months. There are 115 investors classified as most likely belonging to this group and they hold the largest average stock of gilts among the 7 identified groups. Investors in the MIDPHI group allow the range of their portfolio's duration to vary by slightly more, by plus or minus 5 ½ months. This suggests a slightly less extreme form of preferred habitat, but is still remarkably narrow and far below the full sample average. The third group of preferred habitat investors (MID2PHI) also exhibits a much more stable portfolio duration than the average investor in our full sample, with a range of plus or minus 6 ½ months. What is striking about this group is the tight concentration of the portfolios of these investors just above 10-years duration on average. The final group of preferred habitat investors we identify (LTPHI) keep their portfolio duration in a larger but still relatively narrow range of about 2 years, but collectively their preferred duration habitat targets a much larger range of the yield curve.

Investors in the final three groups make up about a third of the investors in our analysis by number, but hold relatively little in terms of end of day stocks on average. They allow the duration of their portfolios to fluctuate dramatically over the sample period with ranges of around +/- 3, +/- 4 and +/- 5 years respectively. It would be hard to interpret this observed behaviour as in any way consistent with having a preferred duration habitat and it seems more consistent with arbitrageur behaviour.

It is noticeable that as well as holding larger amounts of conventional gilts on their balance sheets on average, the preferred habitat investor groups conduct fewer trades and churn their portfolios less than the other investors, something that would also be consistent with theoretical preferred habitat behaviour (see Table 2).

Our first result is therefore that there are identifiable groups of different investors within the gilt market, and that some of these exhibit behaviour consistent with having a preferred duration habitat.

Figure 3: Clustering of investors based on the 10-90 range of portfolio duration (Y-AXIS) and mean portfolio duration (X-AXIS)

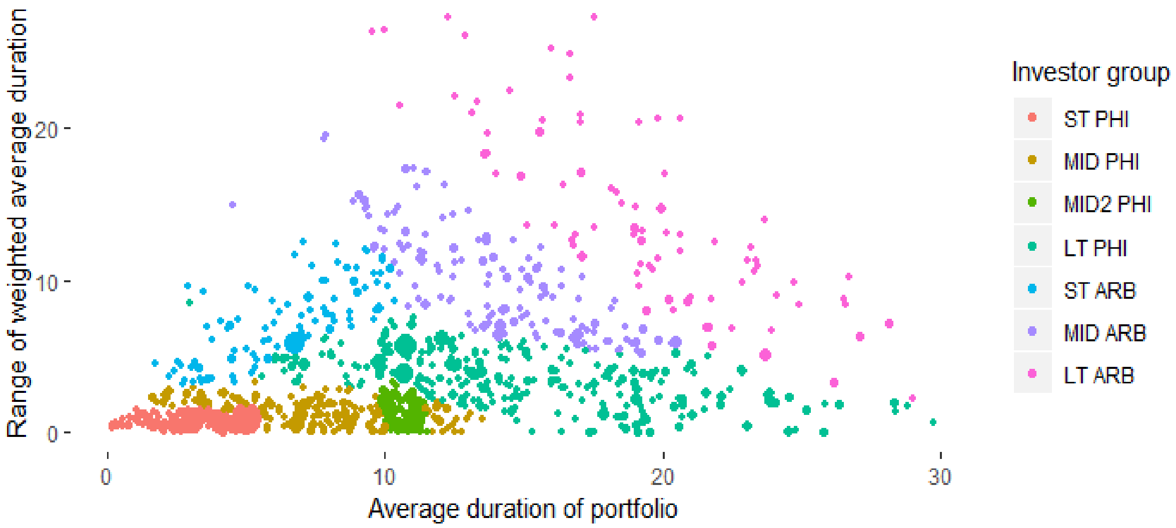

Notes: Results from GMM algorithm estimated over 2016-2017. Point size is scaled by average quantity of investor gilt holdings.

Table 2: Summary statistics on identified clusters

	STPHI	MIDPHI	MID2PHI	LTPHI	STARB	MIDARB	LTARB
# accounts	115	146	95	263	89	131	87
Mean weighted average duration	2.7	7.1	10.5	14.5	6.5	14.6	19.6
Mean weighted average duration range	0.7	1.1	1.3	3.9	6.2	8.1	10.3
Mean stock held (£bn)	50.2	18.2	19.7	35.5	7.6	7.1	3.3
Average trades over sample	147	535	214	597	1556	986	422
Average churn*	2.9	18.6	7.8	30.11	104.3	43	43.2

* Measured by turnover as a multiple of the investor's maximum balance sheet size over the sample.

5.2 Who are preferred habitat investors?

It is often assumed in the related literature that investors in certain economic sectors behave as preferred habitat investors. For instance, pension funds and insurance companies and domestic and foreign official investors are often assumed to behave in this way (see e.g., Zinna (2013), Fukunaga, Kato, and Koeda (2015), Boermans and Vermeulen (2018), Greenwood and Vissing-Jorgensen (2018), Kaminska and Zinna (2019)). At best this may be motivated in terms of investment mandates or regulatory pressures, but typically it is taken as self-evident. Our data and analysis allow us to test the validity of this assumption more formally.

One limitation of the EUI data we use is that it contains limited counterparty information on the CREST account. Accounts are identified with a unique alpha-numeric code and, while we know whether or not that code is associated with a single investor, for many we are unable to go beyond this. To uncover counterparty information, we therefore match this dataset to the regulatory, transaction-level ZEN database, maintained by the Financial Conduct Authority (FCA). The ZEN database contains details on all secondary market trades of UK-regulated firms, or branches of UK firms regulated in the European Economic Area. This dataset covers the same time-period and provides information on the trade: the instrument traded, counterparty identification of the buyer and the seller, the (clean) price and the notional amount traded.

We proceeded to match identical transactions in the EUI data and ZEN data. This allows us to uncover the identity of the counterparties that are uniquely identified and that belong to a CREST account associated with a single investor. Not all investor accounts are uniquely identified, although in some cases we were able to apply judgement to identify the end-investor. Using this matching exercise and data the Bank of England holds as a regulator, we reliably uncover the underlying investor behind 440 accounts in the EUI data (48% of the individual investor accounts).

At the sectoral level there are some clear patterns. Pension funds, insurers and foreign central banks are all largely identified by our behavioural classification as preferred habitat investors. This goes some way to validating the commonly made assumption to proxy preferred habitat investors with these sectors. However, there are also some definite sectoral differences, most notably where they sit on the yield curve. Figure 4 shows that foreign central banks are present at the shorter end of the yield curve, largely targeting duration habitats of 5 years or less. Pension funds on the other hand tend to target duration habitats of 15 years or greater. Insurance companies sit somewhere between the two. This may seem intuitive when one considers the differing motivations and business models of these sectors. The implication is that while using certain sectors as proxies for preferred habitat investors is in general valid, one should be careful about how those sectors are chosen and which part of the market is being evaluated. Holdings by foreign central banks, for instance, may have less relevance to supply and price dynamics at the long end of the curve than holdings by pension funds.

Figure 4: Sectoral mapping of investor groups

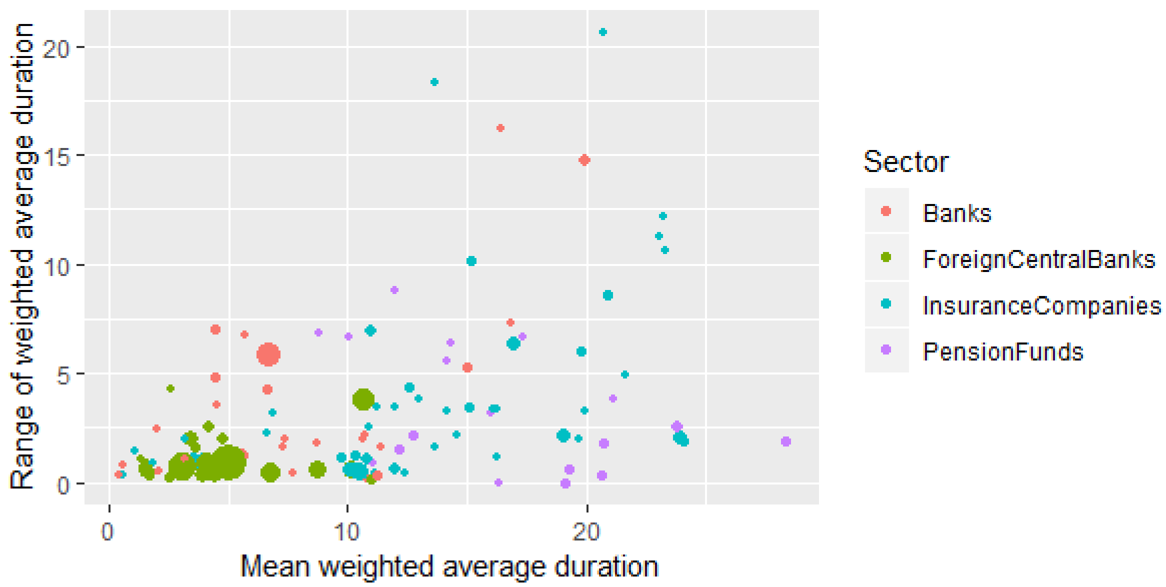

Notes: Point size is scaled by average quantity of investor gilt holdings.

6. Preferred habitat investors and price sensitivity

6.1 A stylised motivation for reduced price sensitivity

It is possible to illustrate how a habitat preference can lead to reduced price sensitivity by making a simple modification to standard portfolio theory to allow deviations from the preferred asset mix to impose a loss of utility.

We begin with the generic 2-asset case, where the assets are denoted A and B. We assume an investor maximises a standard quadratic utility function in the mean and variance of the portfolio return, but has a target weight of asset A in their portfolio, $\overline{w}_A$, so that deviations from this target impose a loss of utility.

We can then write their utility function as

$$E(U) = E(R^p) - \frac{k}{2}Var(R^p) - \Gamma[w_A - \overline{w}_A]^2$$

where $\boldsymbol{E(R^p)}$ is the expected return on the investor's portfolio, $\boldsymbol{k}$ is the parameter that represents risk aversion, $\boldsymbol{Var(R^p)}$ is the variance of the portfolio, $\boldsymbol{w_A}$ is the portfolio share of asset A, $\overline{w}_A$ is the optimal share of asset A in the portfolio and $\boldsymbol{\Gamma}$ denotes the cost the investor faces in deviating from their optimal asset share.

Assuming the two asset weights sum to one, we can expand out the expression for utility to give

$$E(U) = w_A.E(R_A) + (1 - w_A).E(R_B) - \frac{k}{2}[w_A^2\sigma_A^2 + (1 - w_A)^2\sigma_B^2 + 2w_A(1 - w_A)\,Cov(R_A, R_B)] - \Gamma[w_A - \overline{w}_A]^2$$

Differentiating this expression with respect to w_A and setting to zero, we get

$$\frac{dE(U)}{dw_A} = E(R_A) - E(R_B) - \frac{k}{2}[2w_A\sigma_A^2 + (1 - w_A)^2 2\sigma_B^2 + 2(1 - 2w_A)\,Cov(R_A, R_B)] + 2\Gamma[w_A - \overline{w}_A] = 0$$

Rearranging and solving for w_A, we have

$$w_A = \frac{E(R_A) - E(R_B) + 2\Gamma\overline{w}_A + k\left(\sigma_B^2 - Cov(R_A, R_B)\right)}{k\left(\sigma_A^2 + \sigma_B^2 - 2\,Cov(R_A, R_B)\right) + 2\Gamma}$$

As in the standard case (as outlined by Markowitz (1959, 1991) among others), the optimal weight of asset A in the investor's portfolio is a positive function of the spread in the expected return of asset A over asset B, a positive function of the unhedged risk in asset B (the last term in the numerator) and a negative function of the risk in the resulting portfolio, scaled by the degree of risk aversion (the first term in the denominator). However, there are now some additional terms which depend on the cost the investor places on deviating from their optimal asset share.

If we now differentiate with respect to the return spread, we have

$$\frac{\partial w_A}{\partial(E(R_A) - E(R_B))} = \frac{1}{k\left(\sigma_A^2 + \sigma_B^2 - 2\,Cov(R_A, R_B)\right) + 2\Gamma}$$

It follows that the weight of asset A increases as the expected rate of return on A rises relative to asset B. But the introduction of preferred habitat behaviour results in the additional term in the denominator of this expression, which means that the asset share responds less to changes in relative returns as the preferred habitat of the investor towards a particular asset increases (ie for larger values of Γ), ie preferred habitat investors are less price sensitive.

If we assume the investor can only invest in gilts then the two assets would represent different maturity bonds. While highly stylised, the desire for holding a specific weight of asset A could be motivated by a wide range of objectives and this format allows us to abstract from the specifics of each and analyse the core consequences. A strong habitat preference for a particular maturity would be reflected in a high value of Γ and lower sensitivity to expected gilt price movements. It would also be reflected in lower volatility in the maturity, or duration, of the corresponding portfolio. As an example, it could be that asset A is more liquid, and so for balance sheet management reasons, or regulation, the investor wishes to hold a given fraction. Alternatively, our framework also encompasses the case in which the investor is looking to match the characteristics of their liabilities (duration, for example) and the weight, $\overline{w}_A$, is the asset mix which best does that.

6.2 Testing price sensitivity

We test the hypothesis that investors with a stronger habitat preference are less sensitive to (relative) price changes by estimating the following specification on daily data for our 2-year sample (2016-2017).

$$Q_{j,i,t} = \alpha + \beta_i D_i F_{j,t-1} + \delta_j + \varepsilon_t$$

where $Q_{j,i,t}$ is the absolute net change in investor i's holdings of bond j as a fraction of their holdings on the previous day and $F_{j,t-1}$ is the lagged absolute fitting error of bond j, based on daily spot yield curves fitted using the variable roughness penalty (VRP) methodology.[12] The fitting error is intended to capture the relative cheapness or dearness of the bond on a given day and as such can be viewed as a relative price. We look at both the quantity and fitting errors in absolute space because our interest is the extent to which investors change their portfolios, not necessarily the direction in which they change them. Absolute values also enable us to overcome the issue that, absenting a change in the total supply of a bond, a positive change in the holdings of one investor must be offset by an equally negative change in the holdings of other investors in the market.

[12] For details on the VRP methodology, see Anderson and Sleath (2001).

Importantly, we allow β_i to vary by investor group by interacting $F_{j,t-1}$ with a set of dummies indicating whether or not a particular investor belongs to each of our 7 previously identified groupings. Estimating the group coefficients in this way is superior to aggregating at the group level for a number of reasons. First, it allows us to capture intra-group trading, which could be an important part of rebalancing. Second, it appears more defensible to assume that investors are price takers at the individual level, rather than at a higher, aggregate level. Along with the lagged nature of our explanatory variable, this implies the risk of endogeneity from trading positions driving fitting errors is negligible.

Lastly, we control for bond fixed effects, δ_j.

This specification therefore amounts to testing the average relative price elasticities of each group of investors identified in our earlier analysis.

The results of the regression are presented in Table 3. They show that, as a bond becomes cheaper or dearer relative to the curve, investors respond by changing their holdings of it by more; the coefficient on F, β_i, is positive and significant. Of more interest to our question though is the pattern of coefficients on the interactive terms. What they show is that investors that are in groups that our cluster analysis identifies as having tight preferred habitats are significantly less sensitive to the relative cost of the bond than investors in groups identified as arbitragers; the interacted coefficients are significantly negative and significantly different to those of the arbitrageur groups. In each of the preferred habitat groups, the sum of the benchmark coefficient and the interacted term is negative which implies that these groups actually reduce the amount they change their portfolio as the fitting error increases. To be clear, this is not suggesting they move against the market in a countercyclical manner, but rather that, as a bond becomes either cheaper or more expensive, these investors become less sensitive to changes in its relative price. Within the arbitrageur groups, the degree of price sensitivity varies, but notably, those that have the least stable portfolio durations are the most sensitive to changes. This supports the conclusion that they are reacting to smaller fluctuations in price differentials and seeking to arbitrage them where possible, while our preferred habitat groups require a larger degree of compensation to make the same adjustments to their portfolios.

Table 3: Regression results

(Intercept)	0.495***
	(0.01)
F	0.04***
	(0.002)
STPHI *F	-0.075***
	(0.002)
MIDPHI *F	-0.06***
	(0.002)
MID2PHI*F	-0.047***
	(0.002)
LTPHI*F	-0.051***
	(0.002)
STARB*F	-0.029***
	(0.003)
MTARB*F	0.016***
	(0.005)
# obs: 142,050	
Robust standard errors clustered at the bond-investor group level.	
Fixed effect coefficients not shown.	

Source: Data on yield errors based on Bloomberg Finance L.P., Tradeweb Data and Bank Calculations

7. Preferred habitat investors and portfolio rebalancing: QE4 in the UK

Following the UK referendum on leaving the EU in June 2016, the Bank of England announced a package of monetary policy actions on 4 August 2016 to stimulate the economy, including a fourth round of government bond purchases (QE4). Between August 2016 and March 2017 the Bank of England purchased £60bn of conventional gilts as part of this new round, taking the total stock of QE purchases to £425bn. It also purchased a further £8bn as part of its policy of reinvesting the proceeds of maturing gilts purchased in previous rounds. These purchases were implemented through a programme of reverse auctions.

The Bank's QE4 gilt purchases provide an interesting case study for understanding the investment behaviour of preferred habitat investors in response to a shock to net bond supply. In an accounting sense, the Bank's purchases would have been matched by sales from other agents in the economy, or an increase in the total stock of gilts outstanding. Past communications by Bank MPC members (e.g. Fisher (2010)) have stressed that the intention of QE was to buy assets from the non-bank sector, and specifically from pension and insurance funds. If the Bank's purchases came from relatively price insensitive investors with a preferred habitat, it seems possible that they will have viewed the bank deposits they received in exchange as an imperfect substitute and looked to rebalance their portfolios into assets closer to those bonds. This 'portfolio rebalancing' will have led to an increase in the demand for other assets and thus a more generalised increase in asset prices and reduction in yields.

We can examine this episode using our estimates of the gilt holdings of different investor groups to produce a simple accounting of the counterparts to the Bank's purchases. The behaviour around the Bank's auctions is shown in Figure 5, which applies a same-day accounting decomposition to the net change in conventional gilts for the 81 reverse auctions the Bank ran during August 2016 to March 2017. We include all gilts so that we capture the net sellers of all gilts on the day of each auction. During the period, the Bank of England bought on average £0.8bn on each auction date. Comparing the observed changes in gilts holdings, by what might have been expected had the reaction been proportionate to the relative stock holdings of each investor group, suggests that the Bank's purchases seem to have come to a much larger extent than expected from the MID2PHI category of preferred habitat investors. This analysis obviously abstracts from the behaviour of omnibus accounts, which combine the behaviour of a number of individual investors. This explains why the sum of the observed change in gilt holdings across all investor groups (£0.33 billion) is considerably smaller than the average auction size of £0.8billion.

Figure 5: The same day reaction to QE auctions (averaged across all QE4 auctions), all conventional gilts

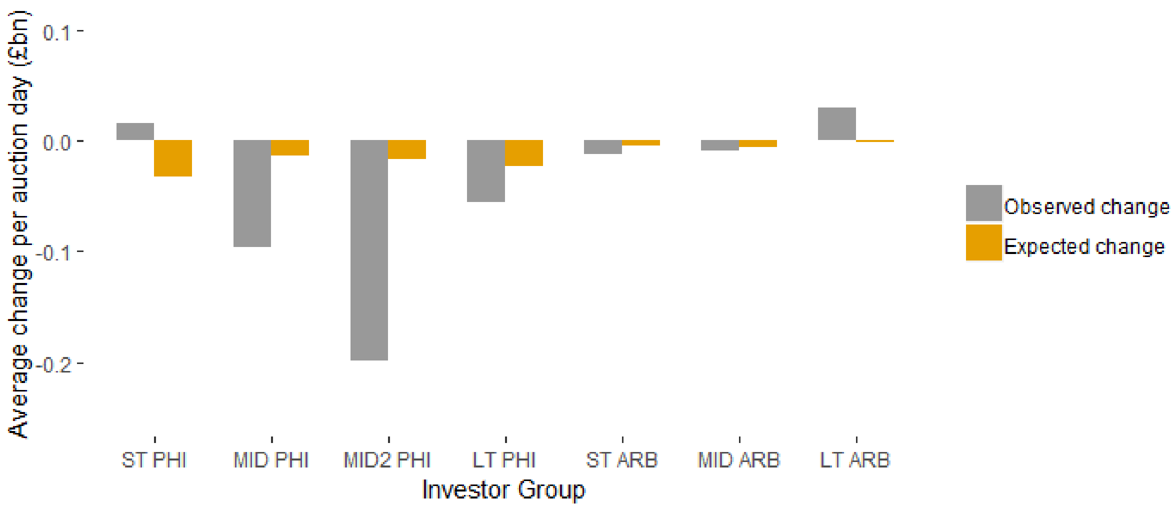

This result is also apparent when repeating the analysis for net change over the full window of the purchase programme shown in Figure 6. The conventional gilt holdings of the MIDPHI2 group of investors seem to have declined disproportionately, suggesting they sold a relatively large share of the gilts bought by the Bank over the period.

The decline in holdings of preferred habitat investors seems consistent at face value with a wider portfolio balance channel (such as found in earlier QE episodes, see e.g. Joyce, Liu and

Tonks, 2017), although information on where these investors invested instead and a plausible counterfactual would be necessary for a full assessment. The fact that purchases appear to have come relatively more from preferred habitat investors rather than other investors suggests their impact on the yield curve may have been larger than otherwise, given their lower price elasticity of demand. However, a fuller empirical analysis of the relationship between investor holdings and the impact on yields across the term structure is not possible, as our data does not allow us to identify the gilt holdings of those individual investors whose gilt holdings are aggregated together in omnibus accounts, which are a significant part of the market.[13]

Figure 6: Change in stocks of investor groups over the QE4 period (3 Aug 2016 – 13 Mar 2017), all conventional gilts

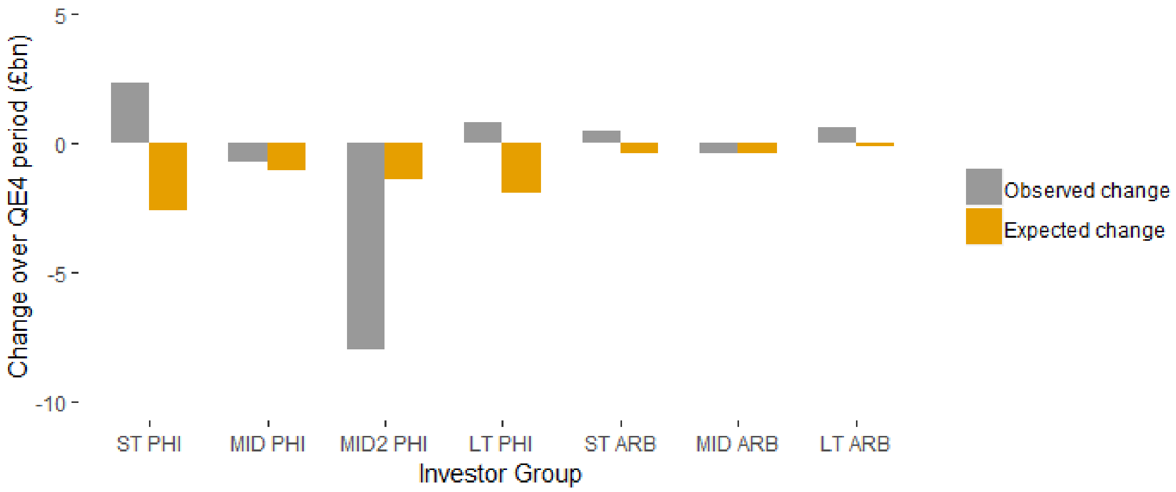

8. Conclusions

Using cluster analysis and a newly available and highly granular dataset on investor behaviour in the gilt market, we are able to show that there are distinct clienteles of investors in the market for UK government securities. Some of these investors display the behavioural properties theory would associate with preferred habitat investors. We find that these investors show a portfolio duration habitat and that they are less sensitive to price movements than the other investor groups. Moreover, by matching our main dataset with other sources, we find that they include institutional investors, like life insurers and pension funds, which are typically thought to exhibit preferred habitat behaviour.

In a separate exercise, we use a simple accounting decomposition to show that one of the main counterparts to the Bank of England's gilts purchases, during its QE4 purchase programme over August 2016 to March 2017, was a reduction in gilt holdings of the groups

[13] This is not an issue for our earlier analysis that looks at classifying and analysing the impact of individual investors, and it still allows us to make confident statements on the portion of the market that we are able to see in our data.

of investors we classify as exhibiting preferred habitat behaviour. The fact that some of these investors appear to have rebalanced out of the securities purchased by the Bank is consistent with QE purchases operating to some extent through the so-called portfolio balance channel, although more data would be necessary to make a comprehensive assessment.

References

Anderson, N. and Sleath, J. (2001), “New estimates of the UK real and nominal yield curves”, Bank of England Working Paper no. 126.

Andritzky, J. R. (2012), “Government bonds and their investors: What are the facts and do they matter?”, IMF Working Paper Series.

Andres, J. J., Lopez-Salido, D. and Nelson, E. (2004), “Tobin’s Imperfect Asset Substitution in Optimizing General Equilibrium.” *Journal of Money, Credit and Banking*, 36, 665–90.

Bhattarai, S., and Neely, C. J., (2016), “A Survey of the Empirical Literature on U.S. Unconventional Monetary Policy.” Federal Reserve Bank of St. Louis Working Paper 2016-021A.

Boermans, M. and Vermeulen, R. (2018), “Quantitative easing and preferred habitat investors in the euro area bond market”, DNB Working Paper No.586.

Chen, H., Cúrdia V., and Ferrero, A. (2012), “The Macroeconomic Effects of Large-scale Asset Purchase Programmes.” *The Economic Journal*, 122(564), F289–F315.

Culbertson, J. M. (1957), “The Term Structure of Interest Rates.” *Quarterly Journal of Economics*, 71, 485–517.

Dempster, A.P., Laird, M. and Rubin, D.B. (1977), “Maximum Likelihood from Incomplete Data via the EM algorithm”, *Journal of the Royal Statistical Society*, Vol. 39, No. 1.

Fisher, P., (2010) “An Unconventional Journey: The Bank of England’s Asset Purchase Program.” Speech to contacts of the Bank’s Agency for the South-West, October. http://www.bis.org/review/r101020e.pdf

Fukunaga, I., Kato, N. and Koeda, J. (2015),“Maturity Structure and Supply factors in Japanese Government Bond Markets’, Monetary and Economic Studies, November, 45-95.

Greenwood, R., and Vayanos, D. (2010), “Price pressure in the government bond market," *American Economic Review: Papers and Proceedings* 100, 585-590.

Greenwood, R. and Vayanos, D. (2014), “Bond Supply and Excess Bond Returns”, *Review of Financial Studies*, 2014, 27, 663-713.

Greenwood, R. and Vissing-Jorgensen, A. (2018), "The Impact of Pensions and Insurance on Global Yield Curves", Harvard Business School, Working Paper 18-109.

Gromb, D and Vayanos, D. (2010), “Limits of Arbitrage: The State of the Theory”, *Annual Review of Financial Economics*, 2, 251-275.

Haldane, A. G., Roberts-Sklar, M., Wieladek, T. and Young, C. (2016), “QE: the story so far,” Bank of England Staff Working Paper No. 624.

Harrison, R. (2012), “Asset purchase policy at the effective lower bound for interest rates”, Bank of England Working Paper No. 444.

Hicks, J. (1939), Value and Capital, Oxford University Press.

Joyce, M. A. S., Liu, Z. and Tonks, I. (2017), “Institutional investors and the QE portfolio balance channel”, *Journal of Money, Credit & Banking*, 49(6), 1225-46.

Kaminska, I. and G. Zinna (2019), “Official Demand for U.S. Debt: Implications for U.S. Real Rates," Bank of England Working Paper No. 796.

Koijen, R. S. J., Koulischer, F., Nguyen, B., and Yogo, M. (2017), "Euro-Area Quantitative Easing and Portfolio Rebalancing," *American Economic Review*, 107(5), 621-27.

Koijen, R. S. J., F. Koulischer, B. Nguyen, and M. Yogo (2021), “Inspecting the mechanism of quantitative easing in the euro area”, *Journal of Financial Economics*, 140(1), 1-20.

Koijen, R. S. J., Richmond, R. J. and Yogo, M. (2020), “Which Investors Matter for Equity Valuations and Expected Returns?”, NBER Working Papers 27402.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012), “The aggregate demand for Treasury debt”, *Journal of Political Economy*, 120(2), 233–67.

Kuroda, H, (2015), “How unconventional monetary policy stimulates demand - theory and practice”, Remarks for the Panel Discussion at the Bank for International Settlements Annual General Meeting, Basel, 28 June 2015. https://www.bis.org/review/r150702c.pdf

Li, C. and Wei, M. (2013), “Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs”, *International Journal of Central Banking*, 9(1), 3-39.

Lucas, A., Schaumburg, J. and Schwaab, B. (2018), “Bank Business Models at Zero Interest Rates”, *Journal of Business and Economic Statistics*, 1-14.

Markowitz, H. M. (1959,1991), *Portfolio Selection*, Blackwell, Oxford.

Merton, R. (1969): “Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case,” *Review of Economics and Statistics*, 51, 247–257.

Merton, R. (1971): “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”, *Journal of Economic Theory*, 3, 373–413.

Merton, R. (1973): “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 41, 867–87.

Modigliani, F. and Sutch, R. C., (1966) “Innovations in Interest Rate Policy.” *American Economic Review*, Papers and Proceedings, 56, 178–97.

Strohsal, T. (2017), “Bond yields and debt supply: new evidence through the lens of a preferred habitat model”, *Quantitative Finance*, 17(10), 1509-22.

Sudo, N. and Tanaka, M. (2018), “Do Market Segmentation and Preferred Habitat Theories Hold in Japan? : Quantifying Stock and Flow Effects of Bond Purchases”, Bank of Japan Working Paper Series No.18-E-16.

Ray, W. (2019), “Monetary Policy and the Limits to Arbitrage: Insights from a New Keynesian Preferred Habitat Model”, mimeo.

Riedel, F. (2004), “Heterogeneous Time Preferences and Humps in the Yield Curve: The Preferred Habitat Theory Revisited”, *European Journal of Finance*, 10, 3–23.

Stiglitz, J. (1969): “A Consumption-Oriented Theory of the Demand for Financial Assets and the Term Structure of Interest Rates,” *Review of Economic Studies*, 37, 321 – 351.

Swanson, E. (2011), “Let’s Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2,” *Brookings Papers on Economic Activity*, Spring, 151–188.

Timmer, Y. (2018), “Cyclical investment behavior across financial institutions”, *Journal of Financial Economics*, 129(2), 268-286.

Tobin, J. (1961) “Money, capital and other stores of value.” *American Economic Review, Papers and Proceedings,* 51(2), 26-37.

Tobin, J. (1963), “An essay on the principles of debt management.” *Fiscal and Debt Management Policies*, Englewood Cliffs, N.J.; Prentice Hall, 143-218.

Tobin, J. (1969), “A general equilibrium approach to monetary theory.” *Journal of Money, Credit and Banking,* 1(1), 15-29.

Vayanos, Dimitri and Jean-Luc Vila. (2009) “A Preferred-Habitat Model of the Term Structure of Interest Rates.” NBER Working Paper No. 15487.

Vayanos, Dimitri and Jean-Luc Vila. (2021) “A Preferred-Habitat Model of the Term Structure of Interest Rates”, *Econometrica*, 89, 77-112.

Wachter, Jessica (2003), “Risk aversion and allocation to long-term bonds”, *Journal of Economic Theory*, 112, 325–33.

Zinna, G. (2016), “Price Pressures on UK Real Rates: An Empirical Investigation”, *Review of Finance*, 1587-1630.