Essays in Household Finance

Erkki Vihriälä
St Catherine's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics
April 2020

Contents

Abstract **3**

Introduction **4**

Limited demand for flexibility – commitment, inertia and inattention in debt repayment **7**

Untangling the credit card debt puzzle **54**

Sources of solace — household heterogeneity and macroeconomic insurance mechanisms **108**

Conclusions **158**

Abstract

This thesis contains three essays on household finance focusing on household consumption and savings decisions. The first paper studies household demand for flexibility in debt repayment. Based on evidence from a natural experiment, I find that few households apply for flexibility in debt repayment, although all households weakly prefer flexibility in a standard consumption-savings model. The paper studies how self-control problems and a desire for commitment, as well as inattention and inertia can explain observed household behavior. The second paper untangles reasons for the credit card debt puzzle describing the co-holding of high-cost debt and low-yield assets. I test existing theories for co-holding and provide evidence on the contribution to co-holding of intra-household financial pooling and of anchoring to minimum credit card payments. The final paper studies household heterogeneity in insuring against cohort-level earnings shocks using: (i) investment income flows; (ii) the tax-and-transfer-system; and (iii) credit markets. The insurance mechanisms vary by household socio-economic characteristics captured by housing tenure.

Acknowledgements I am indebted to the faculty and fellow graduate students at the Department of Economics for many essential comments on the work included in this thesis. I am especially grateful to my supervisor Andrea Ferrero for his advice and support. I also want to thank Hamish Low for his generous help. I acknowlege funding from Yrjö Jahnsson Foundation, Osk. Huttunen Foundation, Emil Aaltonen, Finnish Cultural Foundation, and the Savings Banks Research Foundation without which the completion of the thesis would not have been possible. Finally, I want to thank Leena and our son for supporting me during my studies and making the transition to the next stage of my life full of intrigue.

Introduction

My thesis belongs in the field of household finance. The field studies how households use different financial products, for instance mortgages and credit cards, and how those choices affect outcomes such as consumption and savings. Household finance aims to understand to what extent household behavior corresponds to predictions of classical economic models, and what extensions to the standard framework may further refine our understanding. To the extent that researchers worry about households making suboptimal financial decisions, research in household finance can provide lessons on how to (i) improve household financial literacy, (ii) develop financial products that adequately address the empirical evidence on how households make financial decisions.

This thesis contains three essays. The essays are independent and each essay includes an introduction containing a literature review, conclusions, and references. Tables and figures pertaining to a given essay follow the main text of the essay.

The first essay, *Limited demand for flexibility – commitment, inertia and inattention in debt repayment*, studies household demand for debt-repayment flexibility. On the one hand, debt-repayment flexibility should help households with low liquidity smooth consumption. On the other hand, households with self-control problems may worry that debt-repayment flexibility would lead to overconsumption. I use a natural experiment to study whether households apply for debt-repayment flexibility, and I try to understand which theories can explain the observed behavior. In the experiment, mortgage holders of a large Finnish bank received a free flexibility option to reduce minimum principal payments to zero for 1 to 12 months. Because the offer relaxes the constraint on minimum payments but does not cap payments, a standard consumption-savings model would imply that all households weakly prefer flexibility to no

flexibility. Yet, only one in five households apply for flexibility. I document additional key empirical findings and consider which mechanisms can explain the observed behavior. First, inattention could rationalize limited demand for flexibility if households are unaware of the offer. Second, inertia because of a time or effort application cost could rationalize limited demand. Finally, households with self-control problems may commit to limiting the amount of flexibility to avoid overconsumption. I reject the null hypothesis that inattention and inertia are sufficient to explain the observed behavior, because inattention and inertia fail to explain two key empirical findings. First, many applicants prefer short to maximum flexibility. Second, consumption drops for applicants who are liquidity-constrained before the offer at the predictable end of flexibility. The two key findings are consistent with a model where some households have self-control problems and a desire for commitment. That said, my results do not imply that inattention and inertia could not explain part of the observed behavior, and I do not make a quantitative statement on the relative importance of the different mechanisms.

The second essay, *Untangling the credit card debt puzzle*, studies why households co-hold high-cost debt and low-yielding assets. My first contribution is to document, for the first time, the persistence of co-holding at the monthly frequency. Because of missing the imperfect persistence, the prevailing method of extrapolating from a cross-section gives inaccurate results on the true costs of co-holding. Second, I test whether liquidity-based theories are sufficient to explain co-holding. Liqudity-based theories suggest that co-holding arises because households need liquid assets for certain expenses or that households worry that banks will cut their credit lines. Yet, even households with a large liquidity buffer co-hold in my data, in line with earlier results. What is more, co-holding occurs in Finland although negative credit risk is negligible compared to studies based on US data. In addition, liquidity-based theories struggle to explain why most co-holding households do not take up an offer that provided liquidity at a low interest rate

compared to existing credit card debt. Third, I study if the lack of intra-household financial pooling can explain part of co-holding. I find that the prevalence and persistence of significant co-holding is less likely at the individual than at the household level. What is more, the distribution of intra-household assets and liabilities impacts the take up of the liquidity offer. Therefore, lack of intra-household financial pooling can explain part of co-holding in households with multiple members. Finally, I document evidence consistent with anchoring to the minimum credit card payment contributing to co-holding. I find that many households regularly make minimum or near-minimum credit card payments even if they have significant liquidity. The low credit card payments are consistent with other studies that suggest that the minimum credit card payment acts as a downward anchor to household repayment choices. I bridge the literature on anchoring in credit card payments and the credit card debt puzzle literature by showing that households who regularly make low payments are more likely to co-hold.

The third essay, *Sources of solace — household heterogeneity and macroeconomic insurance mechanisms*, studies how the insurance mechanisms that households use against cohort-level earnings shocks vary by household socio-economic characteristics. I consider three insurance mechanisms studied in the literature: (i) investment income flows, (ii) the tax-and-transfer-system, and (iii) credit markets. I proxy socio-economic characteristics by housing tenure. I find that households in social housing rely on fiscal smoothing of cohort-level earnings shocks, whereas wealthier households who own their house drive overall credit smoothing. Studying the heterogeneity in insurance mechanisms matters, because aggregate results give more weight to wealthier households and can mask the insurance mechanisms relevant to less well-off households.

Limited demand for flexibility – commitment, inertia and inattention in debt repayment

Erkki Vihriälä

Abstract

Debt-repayment flexibility should help households with low liquidity smooth consumption, but households with self-control problems may limit flexibility to avoid overconsumption. This paper studies household response to an offer of debt-repayment flexibility with a natural experiment. The experiment provided Finnish mortgage holders a free option to reduce minimum principal payments to zero for 1 to 12 months. Only one in five apply for this flexibility, although all households weakly prefer flexibility in a standard consumption-savings model. I consider three main mechanisms for household behavior: (i) commitment because of self-control problems, (ii) inertia because of time or effort application costs, and (iii) inattention to the offer. Although inertia and inattention can explain part of the behavior, inertia and inattention are inconsistent with two key empirical findings consistent with self-control problems and a desire for commitment. First, many households prefer short to maximum flexibility. Second, consumption drops at the predictable end of flexibility for households who are liquidity constrained before the offer.
Keywords: debt-repayment flexibility, commitment, inertia, inattention
JEL codes: D14, D15, D91, E21

1 Introduction

Debt-repayment flexibility has received attention as a means for households to smooth income shocks. For instance, in relation to the coronavirus pandemic, multiple countries have promoted mortgage payment holidays that allow households to temporarily stop mortgage payments. Debt-repayment flexibility can have significant benefits, because many indebted households are liquidity constrained even if they may have significant illiquid wealth. On the other hand, households with self-control problems may prefer to commit to limiting debt-repayment flexibility if they worry about overconsumption. Yet, empirically we do not know enough about demand for debt-repayment flexibility, or the reasons for why households may or may not take up flexibility.

This paper provides evidence on demand for debt-repayment flexibility and considers different mechanisms to explain household behavior based on a large natural experiment in Finland. In the experiment, mortgage customers of a large Finnish bank received a free flexibility option to reduce minimum principal payments to zero for 1 to 12 months. Before the experiment a mortgage holder needed to make strictly positive minimum principal payments every month. If a household applied for flexibility, the default principal payment decreased to zero, but the household could make extra principal payments for free.

I find that only one in five households apply for flexibility, although all households would weakly prefer flexibility in a standard consumption-savings model with exponential discounting. Households should weakly prefer flexibility, because flexibility relaxes the minimum principal payment constraint, but flexibility does not cap principal payments. Furthermore, I find that most liquidity-constrained households who have negative net liquid assets and pay interest on their credit card do not apply for flexibility. Hence, indifference between flexibility and no flexibility because of high liquidity does not sufficiently explain the limited demand.

I consider three main mechanisms to understand the observed behavior. First, inattention can explain limited demand for flexibility if households are unaware of the offer. Second, a time or effort application cost can generate inertia and explain limited demand for flexibility if the cost is larger than the potential benefits of flexibility. Third, households may commit to positive minimum principal payments if self-control problems would lead them to overconsume during flexibility.

I reject the null hypothesis that inattention and inertia are sufficient to explain household behavior. Contrary to the inattention hypothesis, I find that most households who arguably know about the flexibility offer – households that include an employee of the bank responsible for the offer – do not apply for flexibility. Second, almost half of applicants restrict the amount of flexibility to a shorter period than the maximum. A time or effort application cost cannot rationalize preference for short flexibility, because in a standard model the benefits of flexibility are weakly increasing in the length of flexibility. Third, applicants who are liquidity constrained before the offer do not smooth consumption at the predictable end of flexibility. An exponential discounting model predicts that households would smooth consumption at the end of flexibility despite ex ante constraints, because smoothing only requires households to save during flexibility.

Allowing for households with self-control problems because of present bias in the form of quasi-hyperbolic discounting (Laibson, 1997) can help explain additional features of household behavior. First, self-control problems can rationalize why some households commit to short flexibility. A household with self-control problems can both benefit from immediate flexibility but want to constrain its flexibility after an initial period to smooth future consumption. Second, quasi-hyperbolic discounting can explain the large consumption drop at the predictable end of flexibility for households who are liquidity constrained before the offer. In the quasi-hyperbolic model, the present bias discount factor β modifies the standard Euler equation to rationalize significant

month-on-month consumption drops. Although self-control problems can help understand the data, I do not argue that inattention or inertia do not matter, or that self-control problems are more important in understanding household behavior. My paper does not provide a quantitative comparison of the importance of the different mechanisms.

My paper relates, firstly, to other studies on demand for flexibility in debt repayment. The closest paper is Cho and Rust (2017), who show that Korean households often forego interest-free installment loans or commit to quick repayment. I similarly find that many households do not apply for payment flexibility, or prefer only short flexibility in a setting with higher stakes.[1] My natural experiment also precludes that households' need for liquidity explains the preference for short flexibility.[2] In addition, a number of experimental papers in development economics study demand for flexibility versus commitment. For instance, Afzal et al. (2019) randomly vary credit contracts of a micro-finance institution and find that flexibility does not increase take up. By contrast, Barboni and Agarwal (2018) document positive demand for flexibility particularly among time-consistent and financially disciplined borrowers.

Because appying for flexibility requires an active choice, my paper also relates to the´household finance literature on default option effects (Beshears et al., 2018; Choi et al., 2002; see, for instance, Madrian and Shea, 2001). Default option effects in debt repayment can reflect, for instance, anchoring (Keys and Wang, 2019; Navarro-Martinez et al., 2011; Stewart, 2009) or inattention and inertia (Andersen et al., 2015; Bajo and Barbi, 2015). I provide evidence consistent with the hypothesis that sticking to a restrictive default option can also provide commitment.

[1]The average installment loan in Cho and Rust (2017) is around $300, and the maximum interest-free repayment period is 12 months. Therefore, savings from using the interest-free period to pay down other credit or to invest in positive return assets are small. Consequently, the authors note that "transaction and mental accounting costs are the most likely additional or alternative explanation for the behavior we find."

[2]Installment loans count against the credit limit in Cho and Rust (2017). Therefore, households may prefer quick repayment of interest-free installment loans to ensure sufficient liquidity at the point of purchase.

My paper also relates to the literature on liquidity constraints among households with illiquid wealth. A household can be wealthy-hand-to-mouth (WHtM), because the long-run benefits from illiquid assets outweigh the short-run costs of liquidity constraints (Kaplan and Violante, 2014; Kaplan et al., 2014). Alternatively, a household with self-control problems can be WHtM because the illiquid asset serves as a commitment (Angeletos et al., 2001; Kovacs and Moran, 2019; Laibson et al., 1998; Schlafmann, 2016). WHtM households without self-control problems should value debt-repayment flexibility, because flexibility would alleviate short-run liquidity constraints without sacrificing the long-run benefits from illiquid assets. My findings are consistent with commitment contributing to the prevalence of WHtM households, because few liquidity-constrained homeowners apply for maximum flexibility.

Finally, my paper relates to the literature on consumption behavior in response to predictable decreases in discretionary income. Unlike with a predictable income increase, liquidity constraints cannot explain a consumption discontinuity after a predictable income decrease, because a household can smooth consumption by saving.[3] Yet, Ganong and Noel (2019) show that households decrease consumption after the predictable end of unemployment benefits. I find a similar result using the increase in minimum principal payments as a predictable decrease in discretionary income. My findings strengthen the interpretation in Ganong and Noel (2019) that consumption decreases in response to predictable income decreases can reflect present bias, because overoptimism about the job finding rate is not a confounder in my setup. My results also complement the finding by Jørring (2018) that households with interest-only home equity lines of credit decrease consumption after the interest-only period ends. Interestingly, Jørring (2018) finds that households who make financial mistakes in other domains drive the consumption drop. The interest-only periods in the setting of Jørring (2018) typically last five to ten years, and

[3]Papers that study the response of consumption to predictable income increases include, for instance, Souleles (1999), Hsieh (2003), Johnson et al. (2006), Shapiro and Slemrod (2009), and Parker et al. (2013).

hence inattention correlated with financial mistakes could explain the consumption drop. By constrast, inattention is unlikely to explain the consumption drop in my setting where households explicitly choose a flexibility length of a much shorter duration.

The paper proceeds as follows. The next section describes the natural experiment on demand for debt-repayment flexibility. Section 3 presents the theoretical framework to understand the flexibility offer and distinguish different mechanisms of household behavior. Section 4 presents the proprietary bank data on which I base my empirical results. Section 5 presents the key empirical findings and discusses the ability of the different mechanisms to explain the findings. Section 6 studies the robustness of the key findings and considers alternative mechanisms for the findings. Section 7 concludes.

2 Natural experiment of demand for flexibility

In February 2015, a Finnish bank offered its mortgage customers a free flexibility option to reduce minimum mortgage principal payments to zero for 1 to 12 months. Households would continue to pay interest during the flexibility period. Households in severe payment difficulties could not apply for the offer, but almost all mortgage customers were eligible. The offer lasted from early February to the end of June 2015.

The offer provided households an unambiguous increase in flexibility. After application, the default principal payment decreased to zero and mortgage maturity was extended by the flexibility length. But, mortgage holders can make extra principal payments for free by making a bank transfer from their deposit account to their loan account. Extra principal payments are free on adjustable-rate mortgages, which account for 97 percent of mortgages in my data. Therefore, the concrete effect of the offer was to temporarily reduce the minimum principal payment from

a positive amount to zero.

The flexibility offer was a major public campaign by the bank and application for flexibility was easy. Application was free and took a few minutes to complete online. In addition to the bank's direct marketing initiatives, the national media also reported on the offer. For instance, the largest national newspaper, Helsingin Sanomat, had the flexibility offer as its lead article (February 6, 2015, edition). The term for the policy ("lyhennysvapaa") returns 26 articles published during the application window in Helsingin Sanomat's digital archive.[4] The national broadcasting company Yle (equivalent to BBC) also covered the offer.

3 Theoretical framework

This section describes the flexibility offer in a stylized consumption-savings model. I first consider the conditions under which flexibility is beneficial for households with or without self-control problems. I also consider the choice of flexibility length and the implications of temporary flexibility on consumption smoothing at the end of flexibility.

3.1 Two-period problem with no self-control problems

Consider a two-period consumption and debt repayment problem with no self-control problems and exponential discount factor δ. The household starts with no liquid assets but has a mortgage M_1. The household receives endowment income y_i in both periods. The household chooses consumption C_i and debt repayment Δ_i in both periods to maximize total utility:

$$\max_{\{C_1,C_2,\Delta_1,\Delta_2\}} u(C_1) + \delta \mathbb{E}[u(C_2)] \tag{1}$$

[4] https://www.hs.fi/haku/. The search was made on the 20th of March 2019.

so that,

$$\Delta_1 \geq \tau M_1, \tag{2}$$

$$Y_1 - \Delta_1 - C_1 \geq -\kappa, \tag{3}$$

$$C_2 = Y_2 + R_{liq}(Y_1 - C_1 - \Delta_1) - R_{mort}M_2, \tag{4}$$

$$\Delta_2 = M_2 = M_1 - \Delta_1. \tag{5}$$

In the second (final) period, the household consumes everything after repaying the mortgage and any debt on the liquid asset (eq. 4). First period consumption and principal payment need to respect two constraints: a constraint on the minimum principal payment (eq. 2), and a borrowing constraint on the liquid asset (eq. 3), where $\kappa > 0$.

The household weakly prefers the mortgage flexibility offer in the model. The flexibility offer relaxes the first period principal payment constraint from $\Delta_1 \geq \tau M_1$ to $\Delta_1 \geq 0$. Therefore, the flexibility offer only gives the household more options. The weak preference of flexibility by a household without self-control problems is not dependent on a particular utility function, income process, number of periods, or mortgage and liquid asset interest rates. In the empirical section, I consider indifference between flexibility and no flexibility.

3.2 Self-control problems and flexibility

To consider why a household with self-control problems may not weakly prefer flexibility, I make two extensions to the previous two-period model. First, I model self-control problems with quasi-hyperbolic discounting (Laibson, 1997), which introduces an additional discount factor β between the present and future periods. Therefore, the problem for the household in period 1 is:

$$\max_{\{C_1, C_2, \Delta_1, \Delta_2\}} u(C_1) + \beta\delta\mathbb{E}[u(C_2)]. \tag{6}$$

In addition, I introduce a pre-period t_0 when the household starts with a mortgage M_0 and makes a decision regarding the flexibility option on first period principal payment Δ_1. Crucially, the decision maker in t_0 is patient between consumption in t_1 and t_2, because both periods occur in the future, and the quasi-hyperbolic discount factor β only distinguishes between the *present* period and all future periods. Consequently, in evaluating the benefits of flexibility, the t_0 decision maker solves a modified version of the problem in equation 1 without the present bias discount factor β. The modification is that the t_0 decision maker is sophisticated and understands that the t_1 decision maker chooses C_1 as a solution to the problem in equation 6. For simplification, I assume no income uncertainty, and that the household does not consume or make principal payments in t_0.

With no income uncertainty, the flexibility offer matters if the constraints under no flexibility bind. Consumption in t_1 under no flexibility is $C_1 \leq Y_1 - (R_{mort} - 1)M_0 - \tau M_0 + \kappa$. The household needs to pay interest on the mortgage inherited from t_0 and respect the principal payment and liquid asset constraints. If t_1 optimal consumption under flexibility is $C^{opt}_{1,flex} \leq Y_1 - (R_{mort} - 1)M_0 - \tau M_0 + \kappa$, the flexibility choice coincides with the no flexibility choice, $C^{opt}_{1,flex} = C^{opt}_{1,noflex}$. Therefore, the flexibility option changes household decisions only if the constraint under no flexibility binds and limits t_1 consumption: $C^{opt}_{1,flex} > Y_1 - (R_{mort} - 1)M_0 - \tau M_0 + \kappa$.

With a binding constraint in t_1, $C^{opt}_{1,flex} > C^{opt}_{1,noflex}$. For simplification I assume that the household is sufficiently impatient in t_1 that it consumes as much as it can under both no flexibility and flexibility. Consequently, the choice about whether to opt for flexibility in t_0 is a comparison of the the two utility flows (LHS: flexibility; RHS: no flexibility):

$$u[Y_1 - (R_{mort} - 1)M_0 + \kappa] + \delta u[Y_2 - \kappa R_{liq} - R_{mort}M_0] \lesseqgtr$$
$$u[Y_1 - (R_{mort} - 1)M_0 - \tau M_0 + \kappa] + \delta u[Y_2 - \kappa R_{liq} - R_{mort}(M_0 - \tau M_0)]. \tag{7}$$

Applying for flexibility increases C_1 by τM_0 and decreases C_2 by $R_{mort}\tau M_0$. Whether the t_0 household prefers flexibility depends on the initial state variables, the relative income level in t_1 and t_2, the minimum principal payment under no flexibility, as well as interest rate and preference parameters. An intuitive case where flexibility may not be beneficial is if income is considerably higher in t_1 than in t_2. If $Y_1 \gg Y_2$, then $C_1 \gg C_2$, even without flexibility. Because the t_0 decision maker wants to smooth consumption between t_1 and t_2, a relatively large income level in t_1 leads to a bad inter-temporal allocation. Flexibility would only make things worse by further accentuating the inter-temporal consumption discrepancy. Yet, flexibility can also be beneficial for a household with self-control problems, for instance, if $Y_2 \gg Y_1$. Without flexibility, C_1 would be suboptimally low relative to the total endowment. Flexibility helps the household with low Y_1 smooth consumption by relaxing the constraint on maximum C_1.

3.3 Confounders of inertia and inattention

The result that a household with no self-control problems weakly prefers flexibility abstracts from inertia and inattention. To understand how inertia and inattention can lead to not applying for flexibility, I first define the pure flexibility benefit: $U_{pure} = U_{flex} - U_{noflex}$. For households with no self-control problems, $U_{pure} \geq 0$, based on weak preference. Yet, a household that is inattentive to the flexibility offer may simply not observe the pure benefit of flexibility and hence not apply. I consider inattention by a "sleep" probability p. With probability p, the household

is unaware of the offer and does not apply for flexibility.

On the other hand, inertia refers to a fixed application cost ψ. The application cost reflects the time or effort cost of application. With inertia, the total benefit of flexibility is:

$$U_{total} = U_{pure} - \psi. \tag{8}$$

Therefore, even with no self-control problems, inertia can cause households with low pure benefits of flexibility or a high effort cost to strictly prefer no flexibility ($U_{total} < 0$). Finally, also households with self-control problems can experience inattention or inertia.

3.4 Choice of flexibility length

Models of household behavior with or without self-control problems give different predictions regarding the choice of the flexibility length. With no self-control problems, the pure benefits of flexibility are non-decreasing with the length of flexibility. The two default options for the length of flexibility were six and twelve months. Therefore, for a household without self-control problems $U_{total,12-month-flex} \geq U_{total,6-month-flex}$. The effort cost can be larger than the benefits of maximum flexibility, $U_{total,12-month-flex} < 0$, and hence some households may prefer no flexibility. But, households without self-control problems weakly prefer, conditional on applying for flexibility, maximum flexibility. In the empirical section, I consider indifference between short and long flexibility.

On the other hand, households with self-control problems may prefer short flexibility to either no flexibility or maximum flexibility. I extend the model to three periods where the household starts with a mortgage and chooses consumption and principal payment already in period 0. For

simplification I assume no income in t_0. The maximization problem is now:

$$\max_{C_0,\Delta_0} u(C_0) + \beta\mathbb{E}[\delta u(C_1) + \delta^2 u(C_2)]. \tag{9}$$

As before, the household in t_0 is patient between consumption in t_1 and t_2. Yet, the t_0 household discounts heavily future consumption relative to present consumption C_0. The household now has a choice of the length of flexibility, where the household can reduce the constraint on minimum principal payment to zero for t_0 and t_1, only for t_0, or for neither period. The constraints on the principal payment under no flexibility are $\Delta_i \geq \tau_i M_i$, $i = 0, 1$.

For simplification, I again consider the case of no income uncertainty. I assume that the household is sufficiently impatient that the principal payment constraints are binding in t_0, t_1, and that the household borrows up to the liquid asset limit in both periods under both short and long flexibility. Consequently, total utility in t_0 from short (LHS) versus long (RHS) flexibility is:

$$u(\kappa) + \beta\{\delta u[Y_1 - (R_{mort} - 1)M_0 - (R_{liq} - 1)\kappa - \tau_1 M_0] + \delta^2 u[Y_2 - R_{liq}\kappa - R_{mort}(M_0 - \tau_1 M_0)]\} \lesseqgtr$$
$$u(\kappa) + \beta\{\delta u[Y_1 - (R_{mort} - 1)M_0 - (R_{liq} - 1)\kappa] + \delta^2 u[Y_2 - R_{liq}\kappa - R_{mort}M_0]\}. \tag{10}$$

The household consumes the maximum in t_0 under both short and long flexibility, but C_1 is lower under short versus long flexibility and C_2 is higher. By subtracting $u(\kappa)$ from both sides and dividing by $\beta\delta$, the problem simplifies to resemble the earlier problem on flexibility or no flexibility in equation 7. Therefore, the decision on short versus long flexibility can go either way. For instance, the household may prefer only short flexibility if income in t_1 is high relative to t_2,

because restricting flexibility in t_1 helps the decision maker in t_0 smooth consumption between t_1 and t_2. Conversely, the household prefers maximum flexibility, for instance, if income in t_2 is high relative to t_1, because otherwise consumption in t_1 would be suboptimally low compared to consumption in t_2.

A takeaway from the previous analysis is that the duration of the present affects the predictions of the sophisticated quasi-hyperbolic discounting model. At the one extreme, if the present period lasts for one month, and flexibility is available only starting from next month, the flexibility period occurs fully in the future. Consequently, a present-biased household may reject all flexibility. At the other extreme, if the present period lasts for a year or more, the maximum flexibility of twelve months occurs fully in the present. Consequently, a present-biased household is likely to prefer maximum flexibility. In between the two extremes, a present-biased household may consider that part of the flexibility is available to increase present consumption and that part of flexibility occurs in the future when the household wants to smooth consumption. Households with such an intermediate duration of the present are relatively likely to prefer short flexibility to either no flexibility or maximum flexibility.

3.5 Flexibility and consumption smoothing

Models of household behavior with or without self-control problems can also generate different predictions on consumption smoothing at the predictable end of flexibility. Whereas a household with no self-control problems should smooth consumption at the end of flexibility, the high short-run discount factor for households with self-control problems can lead to a discontinuous consumption drop at the end of flexibility.

The flexibility offer provides a rare test of consumption smoothing of a fully predictable income

decrease. A predictable discretionary income decrease, such as the end of the flexibility period, should not lead to a consumption discontinuity for households without self-control problems. When the flexibility period ends, the constraint on the minimum principal payment reverts from zero to a strictly positive amount. But, a household can always save in the liquid asset during the flexibility period to smooth consumption once flexibility ends. Hence, contrary to a predictable income increase, ex ante liquidity constraints do not prevent consumption smoothing at the end of flexibility.

Because the end of flexibility period is perfectly predictable, the Euler equation determines the consumption path:

$$u'(C_{flex}) \geq \delta R_{liq} \mathbb{E}[u'(C_{post-flex})] \tag{11}$$

, where C_{flex} refers to consumption in the last month of flexibility and $C_{post-flex}$ refers to consumption in the first month after flexibility. Under no income uncertainty, zero interest rate, and CRRA utility:

$$C_{post-flex}/C_{flex} \geq \delta^{1/\rho} \tag{12}$$

, where ρ is the coefficient of relative risk aversion in the CRRA utility function $\frac{C^{1-\rho}}{1-\rho}$.

Table 1 presents the optimal relative consumption $C_{post-flex}/C_{flex}$ at the end of the flexibility period at the monthly level for different levels of the annualized exponential discount factor and the coefficient of relative risk aversion. Consumption drops at the end of flexibility are small for common parameter values in the literature, because even a relatively high annual exponential discount factor corresponds to limited discounting at the monthly level. For instance,

with log-utility and an annualized $\delta = 0.95$, the month-on-month consumption drop is less than 0.5 percent. Furthermore, the consumption ratios in Table 1 overestimate the expected consumption drop if we allow for positive interest rates and income uncertainty as in equation 11. Uncertainty in $Y_{post-flex}$ generates a precautionary savings motive which reduces the expected consumption drop after the end of flexility relative to no uncertainty. A positive interest rate increases $C_{post-flex}$ because saving becomes more valuable.

On the other hand, models with self-control problems because of quasi-hyperbolic discounting can rationalize larger consumption drops. Following Harris and Laibson (2001), and abstracting from uncertainty and interest rates, the Euler equation of a sophisticated quasi-hyperbolic discounter is:

$$C_{flex}^{-\rho} \geq [C'_{post-flex}(W_{post-flex})\beta\delta + (1 - C'_{post-flex}(W_{post-flex}))\delta]C_{post-flex}^{-\rho} \tag{13}$$

, where $C'_{post-flex}(W_{post-flex})$ refers to the marginal propensity to consume out of liquid wealth. The additional discount factor β between the present period and all future periods allows the quasi-hyperbolic model to match significant short-run consumption decreases if the household is close to its borrowing constraint and has a high marginal propensity to consume. At the extreme, if $C'_{post-flex}(W) = 1$, then $C_{post-flex}/C_{flex} \geq (\beta\delta)^{1/\rho}$. With log-utility, $\delta = 1$ and $\beta = 0.9$, the consumption drop at the end of flexibility is 10 percent. In earlier work, Shapiro (2005) shows that a sensible calibration of the quasi-hyperbolic model can match a 10 to 15 percent decline in caloric intake within a month unlike the exponential discounting model. Similarly, Ganong and Noel (2019) explain consumption drops at the predictable end of unemployment benefits with quasi-hyperbolic preferences, although Ganong and Noel (2019) consider the case where households are naive about their present bias.

4 Data

Data for this paper come from the bank responsible for the flexibility offer. The bank is among the largest in Finland, with branches across the country. Therefore, the bank's customers are a broad sample of the population. I aggregate nominal variables to the household level because I consider that households make mortgage decisions jointly. In robustness checks (section 6), I discuss the relevance of intra-household strategic behavior for the decision to apply for flexibility.

Demographics and grouping variables: I observe individuals' age, gender, and the number of adults and children in a household. I identify employees of the bank.

Debt: I observe debts at the bank including their purpose (for instance mortgage, unsecured credit), balance, interest rate, and payment schedule. I cannot observe debts at other banks.

Assets: I observe assets at the bank but not outside the bank. I observe a property value estimate equal to the property price at mortgage origination multiplied by a house price index.

Consumption: I observe purchases with debit and credit cards issued by the bank by store category (for instance grocery store, gas station). I use the categorization to create measures of non-durables expenditure as proxies for consumption.

Income: I derive a disposable income measure based on administrative information on taxable labor income in 2015. I observe a second income measure based on monthly flows to bank accounts.

Sample selection

To reduce measurement error, I focus on households who conduct most of their finances at the bank. The bank usually requires new mortgage customers to transfer all finances to the bank. The bank also financially encourages customers to concentrate their finances. To further reduce

measurement error, I filter households with the following criteria:

1. No household member states another bank as their main bank.
2. The household receives income payments to their bank accounts in 2014.
3. Household after-tax income is at least 500 EUR/month in 2015.
4. The household has a deposit account with the bank.
5. The household has a credit card issued by the bank.
6. The household makes average monthly card purchases of at least 200 EUR in 2014.
7. The mortgage has a property collateral.
8. I omit entrepreneurs because of uncertainty about business and personal accounts.
9. I omit households with fixed-rate mortgages because they do not allow free extra principal payments.

I also omit households in severe payment difficulties to ensure eligibility for the offer. After the restrictions, my main sample includes 204,246 households. Table 2 provides descriptive statistics on the key variables for this population.

5 Empirical results

This section first documents the overall popularity of the flexibility offer and how applicants differ from households who do not apply. Finally, I present key empirical findings on household behavior and discuss if inattention, inertia, and self-control problems/commitment can explain the key empirical findings.

5.1 How many and what types of households apply for flexibility?

Demand for flexibility is limited but heterogeneous by household characteristics. Out of 204,246 households in my data, 44,292 apply for flexibility, that is, 21.7 percent. Unconditionally, application correlates positively with liquidity constraints proxied by low deposit balances, high unsecured credit balances, and high minimum principal payments (Table 3). Applicants also have

a higher loan-to-value ratio. Liquidity constraints remain among the best application predictors in a linear probability model that measures the marginal impact of household characteristics on application probability (Table 4).

On the other hand, applicants are not particularly poor or (ex-ante) risky. Application propensity does not vary strongly by expenditure or property value, either unconditionally or conditionally. Applicants have lower mortgage interest rates than non-applicants. Consequently, the bank did not consider the applicants (ex-ante) particularly risky. The negative connection between application and interest rates is understandable, because a higher interest rate decreases the benefits of flexibility.

The relationships of age and income to application need care. Unconditionally, applicants are younger than non-applicants (Table 3). Yet, the marginal predictive power of age is low after controlling for differences in other characteristics (Table 4). Conversely, average disposable income is similar in the two groups, although the marginal impact of extra income is negative. But, disposable income is the only variable partly determined after the February 2015 offer, because the variable measures income for the calendar year 2015. Therefore, the negative correlation can reflect reverse causality if households worked less after application. Unfortunately, I lack administrative income data for 2014. If I replace the administrative income measure with income flows to the bank account in the six months *before* the offer announcement, income does not have a marginal effect on application. The conflicting results can reflect reverse causality or measurement error in the bank-based income measure.

5.2 Key empirical findings

This section details key empirical findings to evaluate different explanations for limited demand for flexibility. The key findings are that: (i) a minority of attentive households apply for flexibility, (ii) a minority of liquidity-constrained households apply for flexibility, (iii) almost half of applicants restrict the length of flexibility to a shorter period than the maximum, and (iv) consumption drops at the end of the flexibility period for applicants who are liquidity constrained before the offer.

Finding 1: A minority of attentive households apply for flexibility

The simplest explanation for why households do not apply for flexibility is unawareness about the offer. To evaluate whether inattention is sufficient to explain non-application, I isolate a subset of mortgage holders who must have known about the offer: households who included an employee of the bank responsible for the offer (N = 7,439). Only 34 percent of bank-employee households with a mortgage apply. Consequently, inattention is an insufficient explanation for limited demand for flexibility.

Yet, inattention can still contribute to the low application rate. I study the link between general attention to household finances and application with two attention proxies. The first proxy is the number of logins to online or mobile banking, which has been used previously in the literature (Jørring, 2018; Olafsson and Pagel, 2017). The login data unfortunately exist only from 2017 after the flexibility policies had ended. Consequently, a connection between logins and application can reflect reverse causality if application for flexibility increases attention to household finances. Therefore, I use the share of messages sent by the bank opened by the customer in online or mobile banking as a second attention proxy. The message data are from

2014, the year before the offer announcement.

Both attention proxies correlate positively with application in a linear probability model when controlling for other household characteristics (Table 5). A household in the top quartile by online logins is 8 percentage points (35 percent of the mean probability) more likely to apply than a household in the bottom quartile. A household who read all online messages sent by the bank in 2014 is 4 percentage points (20 percent of the mean probability) more likely to apply than a household who read no messages. Both specifications control for household characteristics present in the baseline linear probability model without attention proxies (Table 4).[5] In addition, the regression based on online messages controls for the total number of messages. Limited use of online banking is not a concern for the representativeness of the results, because over 98 percent of households received at least one online message in 2014 and logged in at least once to online or mobile banking in 2017.

Finding 2: Most liquidity-constrained households do not apply

Although all households without self-control problems weakly prefer flexibility, households with high liquidity may be close to indifferent between flexibility and no flexibility. For instance, for a household with 100,000 EUR of liquid savings, a flexibility option on a mortgage with a 500 EUR minimum principal payment is unlikely to change the household's consumption-savings decisions. Indeed, households who do not apply for flexibility have more liquid savings than applicants (Tables 3 and 4). Therefore, small or non-existent benefits could explain non-application for certain households.

[5]The controls include: the average age of adults in the household, the number of adults and children, the municipality, deposits, unsecured credit, credit card utilization, disposable income, card purchases, principal payment and outstanding mortgage, property value, mortgage age, and for whether the household includes a bank employee.

Nevertheless, only 38 percent of liquidity-constrained households apply for flexibility, although the potential benefits are large. For these households, flexibility allows to replace expensive unsecured credit with cheap mortgage debt. I define a household as liquidity constrained if net liquid assets are negative at the end of the month and the household pays interest on unsecured debt.[6] According to my definition, 20 percent of households with a mortgage are liquidity constrained before the offer in January 2015. The size of the group is similar to the share of wealthy-hand-to-mouth households with illiquid wealth but little or no liquid assets in a cross-section of advanced economies (Kaplan et al., 2014).[7] Because most liquidity-constrained households do not apply for flexibility, exact indifference between flexibility and no flexibility does not explain why so few households apply.

Transitory liquidity constraints or inattention do not explain the low application rate among liquidity-constrained households. If households know that constraints are transitory, for instance because of a predictable future windfall, households may be indifferent between flexibility and no flexibility. Yet, 71 percent of liquidity-constrained households who do not apply remain liquidity constrained after 12 months (Figure 1). Inattention is an insufficient explanation, because only 55 percent of liquidity-constrained (and attentive) bank-employee households (N = 1,195) apply.

Finding 3: Many apply for short flexibility policies

The distribution of flexibility policy lengths reveals whether applicants restrict the amount of debt-repayment flexibility. The two default lengths for flexibility were 6 and 12 months, but a household could choose an idiosyncratic length shorter than 12 months.

[6]Net liquid assets equal the sum of deposit balances, mutual fund holdings and stock holdings minus interest-paying unsecured credit.

[7]In my sample, 91 percent of liquidity-constrained households are wealthy-hand-to-mouth because they have positive net housing equity. In Kaplan et al. (2014), the share of wealthy-hand-to-mouth households is typically around 20 percent of all households, but their sample does not include Finland. I only study households with a mortgage.

Many households choose a flexibility period shorter than the maximum. Conditional on application, 24 percent choose 6 months and 21 percent choose a length other than 6 but shorter than 12 months (Figure 2). These households know about the offer, choose to pay any time or effort application cost, and yet forego maximum flexibility.

I study whether an improvement in household finances explains the preference for short flexibility. If households have currently low liquidity but expect a positive income event during the next year, they can be essentially indifferent between short and maximum flexibility. I study the behavior of households who are liquidity constrained before the offer in January 2015 and apply for a short flexibility policy of six months. I find that 74 percent of these households remain liquidity constrained in January 2016 (Figure 3). Therefore, transitory liquidity constraints do not explain the preference for short flexibility.

Finding 4: Consumption drops at the end of flexibility for constrained applicants

Consumption behavior around the flexibility period relates to how households discount utility intertemporally. I focus on whether households smooth consumption at the end of flexibility when the resumption of positive minimum principal payments causes a predictable decrease in discretionary disposable income.

I study consumption at the end of flexibility with an event-study regression, in which I compare the consumption change for applicants with non-applicants. Non-applicants are a valid control group under the assumption that the end of flexibility is the only major change with heterogeneous impacts on the consumption of applicants and non-applicants around the end of flexibility for applicants. In robustness checks (section 6), I re-estimate the event-study regression with a matched-on-observables sample of non-applicants as controls.

The event study regression is:

$$y_{i,c,t,p} = \alpha_i + \lambda_t + \mathbf{1}_{c=apply} \sum_{p=e-3}^{e+5} \delta_p \mathbf{1}_{period=p} + \delta_{early} \mathbf{1}_{c=apply} \mathbf{1}_{period<e-3} + \delta_{late} \mathbf{1}_{c=apply} \mathbf{1}_{period>e+5} + \epsilon_{i,c,t,p} \tag{14}$$

, where y is logged consumption, c refers to the cohort (apply, do not apply), α_i are household fixed effects, and λ_t are common time effects. The symbol e refers to the last month of flexibility, which is the baseline month (period dummy omitted from the regression). The period dummies interacted with the application dummy capture the consumption evolution for applicants relative to non-applicants around the end of flexibility. I estimate the equation separately for applicants who are liquidity constrained before the offer and for applicants unconstrained before the offer.

Applicants who are liquidity constrained before the offer reduce consumption by approximately 5–6 percent two months after flexibility ends compared with the last month with zero minimum principal payment (Figure 4 and Table 6). Conversely, unconstrained applicants smooth consumption at the end of flexibility. The results do not depend on the consumption measure: total non-durables expenditure or expenditure on essentials including groceries, energy, public transportation, and gasoline. For completeness, Figure A.2 describes the full consumption path before, during, and after flexibility. The increase in consumption at the beginning of flexibility can be explained by flexibility relaxing ex ante constraints.

My setup provides evidence against consumption smoothing in a special setting, because unlike with a predictable positive income event such as a pre-announced tax rebate, ex-ante constraints do not necessitate the consumption drop at the end of flexibility. Households can smooth consumption at the end by saving during flexibility. In a different setting, Ganong and Noel

(2019) find that consumption drops at the predictable end of unemployment benefits. My results complement their findings in a setting where possible optimism about the job finding rate is not a confounder.

5.3 Discussion of empirical results

This section discusses the ability of inattention, inertia, and commitment to explain the key empirical findings.

5.3.1 Inattention

Inattention fails to explain a number of key empirical findings. First, inattention cannot explain that most attentive employee households do not apply. Second, inattention cannot explain the preference for short flexibility. Finally, inattention cannot explain the consumption drop at the predictable end of flexibility. Yet, the evidence is consistent with inattention contributing to the low application rate. Households who pay less attention to household finances in general are less likely to apply for flexibility.

5.3.2 Inertia

Inertia can explain low application rates with no restriction on the application time or effort cost. If the cost ψ is free to vary in equation 8, inertia can explain any overall application rate. Whether the cost required for inertia to explain that only one in five apply is reasonable would require structural estimation of a rich life-cycle model.

Yet, inertia cannot explain the preference for short flexibility, because without self-control problems, the benefits of flexibility are non-decreasing in flexibility length, but the time or effort

cost is independent of flexibility length. A caveat would be that households who prefer short flexibility improve their financial situation quickly, and hence are essentially indifferent between short and maximum flexibility. Consequently, a tiny cost of making an extra principal payment, a different form of inertia, would rationalize the preference for short flexibility. Yet, I did not find that liquidity-constrained households who apply for short flexibility policies quickly improve their financial situation (Figure 3).

Finally, inertia cannot explain the consumption drop at the end of flexibility, because a time or effort application cost is irrelevant to the consumption decision at the end of flexibility.

5.3.3 Self-control problems and commitment

Self-control problems can create a desire for commitment, which can reduce demand for flexibility overall. Yet, the overall application rate is not sufficient evidence of a desire for commitment, because also inertia and inattention could rationalize why a household does not apply. Again structural estimation of a life-cycle model could help gauge to what extent overall application rates reflect a desire for commitment versus inertia or inattention.

A desire for partial commitment can explain the preference for short flexibility policies. Short flexibility policies may represent the optimum, for instance, for households with quasi-hyperbolic discounting who consider that a part the flexibility period occurs in the present and a part in the future. These households may use short-term flexibility to increase present consumption, but may reject long-term flexibility to smooth future consumption.

Self-control problems can also explain the consumption drop at the end of flexibility. The quasi-hyperbolic Euler equation and results by Shapiro (2005) and Ganong and Noel (2019) show that quasi-hyperbolic discounting can rationalize consumption drops of the magnitude

found in my event study regressions.

5.3.4 Summary

The key empirical findings suggest that inattention and inertia are insufficient to explain the data. The popularity of short flexibility policies and the consumption drop at the predictable end of flexibility are consistent with self-control problems. That said, I do not argue that inattention and inertia explain none of the findings or only a negligible share of household behavior. My paper is not able to estimate quantitatively to what extent inertia, inattention, self-control problems/commitment explain the data. A quantitative breakdown of the different mechanisms would require the estimation of a rich life-cycle model.

6 Robustness

This section studies the robustness of the empirical findings, and alternative mechanisms for the empirical findings.

6.1 Robustness of empirical findings

Consumption smoothing event study with matched sample

In the event-study regression on consumption smoothing at the end of flexibility, I used the full sample of non-applicants as controls to sweep out monthly variation in consumption. Yet, because non-applicants are systematically different from applicants, the consumption profile for applicants could reflect a confounding event, for instance, a legislative change, that affects applicants and non-applicants differentially and coincides with the end of the flexibility policy.

Another event is unlikely to explain the consumption profile at the monthly frequency though, particularly because flexibility policies do not all end in the same month, because the sign-up window was five months.

To further address the concern of a confounding event, I rerun the event-study regression with a matched-on-observables group of non-applicants as controls. For each applicant household, I first find an exact match from the pool of non-applicants by (i) the number of adults and children, (ii) the age decile of adults, (iii) the urbanization category of the municipality (urban/semi-urban/rural), (iv) bank-employee status, and (iv) an indicator variable of liquidity constraints. Within the group of potential non-applicant matches, I use nearest-neighbor matching based on application propensity score to further balance covariates. Matching is with replacement; that is, some non-applicants are matched to multiple applicants.

Figure A.1 and Table A.1 present results on consumption behavior at the end of flexibility with the matched sample of non-applicants as controls. The qualitative pattern is the same as when using the full sample of non-applicants as controls, even though the consumption drop for constrained applicants is somewhat smaller than in the baseline estimates. The conclusion remains that liquidity-constrained applicants reduce consumption discontinuously at the end of flexibility.

Figure A.3 depicts the evolution of consumption over the full flexibility window with the matched sample. The consumption path during flexibility is relatively stable for applicants who are liquidity constrained before the offer, unlike the upward trend in results with the full sample of non-applicants as controls (Figure A.2). The stable consumption path during flexibility combined with the drop at the end of flexibility is consistent with liquidity-constrained applicants being close to their borrowing constraint around the flexibility window.

6.2 Alternative mechanisms for empirical findings

Low financial literacy

Low financial literacy could explain why households do not apply for flexibility or the preference for short flexibility policies if households do not understand the terms of the offer.

Yet, the behavior of bank-employee households suggests low financial literacy is an insufficient explanation for the key findings. As shown earlier, most employee households do not apply, even though bank employees should understand the offer made by their employer. In addition, a significant share (32 percent) of employee applicants prefer short flexibility policies (43 percent among all applicants).

The preference for short flexibility policies could also reflect that households with low financial literacy exhibit random noise in their selection of the length of the flexibility policy. The two default lengths for the flexibility policy were 6 and 12 months. Noise could explain why some choose 6-month policies, because the 6-month option was as prominently displayed as the 12-month option and hence salient to households. But households were free to choose a non-default length other than 6 and shorter than 12 months. Because 18 percent of applicants actively choose such a non-default length, random noise is unlikely to explain the preference for short flexibility.

Intra-household bargaining

A household may not apply for flexibility, because household members have different preferences or do not trust each other (see, for instance, Ashraf, 2009; Chiappori and Mazzocco, 2017). Intra-household bargaining because of conflicts of interest could also explain the preference for short flexibility policies as a compromise.

I show in a companion paper that intra-household bargaining does influence decisions in certain households (Vihriälä, 2020). In particular, an individual weighs more his/her own financial situation than the situation of his/her partner in the decision to apply for flexibility. Yet, this finding only holds for households in which members have individual mortgages. Conversely, in households with joint mortgages (the norm), the distribution of financial assets and debts within the household does not influence the application probability.

If intra-household bargaining explained limited demand for flexibility, we would expect single-adult households to be more likely to apply and less likely to prefer short flexibility policies. Yet, the application rate is 21 percent among single-adult households versus 22 percent in the full sample. In addition, the share of shorter than maximum flexibility policies among applicants is 45 percent in single-adult households versus 43 percent in the full sample. Because application patterns are similar in households with and without potential conflicts of interest, intra-household bargaining is unlikely to explain why most households do not apply for flexibility.

Application stigma

Another hypothesis for the low application rate is that households worry about a stigma if they apply for flexibility. For instance, a household could worry that application would lead to worse terms on future borrowing, because the bank would consider the household as a less reliable borrower.

Although private considerations of households are unobservable, the stigma hypothesis does not fit the public portrayal of the flexibility offer. First, the bank explicitly stated that flexibility would not have any future repercussions in terms of higher interest rates or other borrowing costs. Furthermore, the bank portrayed applying for the offer as a prosocial endeavor. At the time of the offer, the Finnish economy had been stagnant for several years. The bank framed the

offer as a joint endeavor to boost the economy by encouraging households to spend. Therefore, the stigma hypothesis is opposite to the actual framing of the offer.

Finally, even if worries about a stigma may have deterred some households from applying, stigma struggles to explain the preference for short flexibility or, in particular, the consumption drop at the end of flexibility.

7 Conclusions

This paper has studied demand for flexibility in debt repayment with a natural experiment that gave households a free option to reduce minimum mortgage principal payment to zero for 1 to 12 months. A small minority apply for debt-repayment flexibility, although a model without self-control problems suggests that all households should weakly prefer flexibility. Although inertia and inattention can explain part of household behavior, inattention and inertia fail to explain two key findings. First, many households prefer short to maximum flexibility. Second, applicants who are liquidity constrained before the offer do not smooth consumption at the predictable end of flexibility. A model where the household has self-control problems because of quasi-hyperbolic discounting can explain the two key findings. On the one hand, self-control problems can motivate certain households to commit to short flexibility. On the other hand, quasi-hyperbolic discounting can rationalize a large consumption drop at the end of flexibility.

My paper argues that self-control problems can help understand additional features of household behavior, but I do not claim that self-control problems matter more than inattention or inertia. A quantitative comparison would require the estimation of a rich life-cycle model. An estimated life-cycle model could help gauge, for instance, the size of the time or effort application cost consistent with the observed application rates. In addition, a life-cycle model could yield

estimates of the duration of the present in the quasi-hyperbolic model and the present bias discount factor consistent with observed application rates and the consumption drop at the end of flexibility.

My paper provides new evidence that households often stick to the default option in financial decisions. Understanding the reasons for inaction matters for policy. If inaction reflects informational frictions such as inattention or inertia, policymakers should make the option of active choice salient and easy. Yet, policymakers should not interfere with inaction that reflects the commitment value of the default option for households with self-control problems.

References

Afzal, U., D'Adda, G., Fafchamps, M., Quinn, S.R., Said, F., 2019. Implicit and Explicit Commitment in Credit and Saving Contracts: A Field Experiment (Working Paper No. 25802). National Bureau of Economic Research. https://doi.org/10.3386/w25802

Andersen, S., Campbell, J.Y., Nielsen, K.M., Ramadorai, T., 2015. Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market (Working Paper No. 21386). National Bureau of Economic Research. https://doi.org/10.3386/w21386

Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., Weinberg, S., 2001. The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation. Journal of Economic Perspectives 15, 47–68. https://doi.org/10.1257/jep.15.3.47

Ashraf, N., 2009. Spousal Control and Intra-household Decision Making: An Experimental Study in the Philippines. American Economic Review 99, 1245–1277.

Bajo, E., Barbi, M., 2015. Out of Sight, Out of Mind: Financial Illiteracy, Inattention, and Sluggish Mortgage Refinancing (SSRN Scholarly Paper No. ID 2531366). Social Science Research Network, Rochester, NY.

Barboni, G., Agarwal, P., 2018. Knowing what's good for you: Can a repayment flexibility option in microfinance contracts improve repayment rates and business outcomes? Job Market Paper.

Beshears, J., Choi, J.J., Laibson, D., Madrian, B.C., 2018. Behavioral Household Finance (Working Paper No. 24854). National Bureau of Economic Research. https://doi.org/10.3386/w24854

Chiappori, P.-A., Mazzocco, M., 2017. Static and Intertemporal Household Decisions. Journal of Economic Literature 55, 985–1045. https://doi.org/10.1257/jel.20150715

Cho, S., Rust, J., 2017. Precommitments for Financial Self-Control? Micro Evidence from the 2003 Korean Credit Crisis. Journal of Political Economy 125, 1413–1464. https://doi.org/10.

1086/693037

Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A., 2002. Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance. Tax Policy and the Economy 16, 67–113. https://doi.org/10.1086/654750

Ganong, P., Noel, P., 2019. Consumer Spending during Unemployment: Positive and Normative Implications. American Economic Review 109, 2383–2424. https://doi.org/10.1257/aer.20170537

Harris, C., Laibson, D., 2001. Dynamic Choices of Hyperbolic Consumers. Econometrica 69, 935–957. https://doi.org/10.1111/1468-0262.00225

Hsieh, C.-T., 2003. Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund. American Economic Review 93, 397–405. https://doi.org/10.1257/000282803321455377

Johnson, D.S., Parker, J.A., Souleles, N.S., 2006. Household Expenditure and the Income Tax Rebates of 2001. American Economic Review 96, 1589–1610. https://doi.org/10.1257/aer.96.5.1589

Jørring, A., 2018. The Cost of Financial Mistakes: Evidence from U.S. Consumers. Job Market Paper.

Kaplan, G., Violante, G.L., 2014. A Model of the Consumption Response to Fiscal Stimulus Payments. Econometrica 82, 1199–1239. https://doi.org/10.3982/ECTA10528

Kaplan, G., Violante, G.L., Weidner, J., 2014. The Wealthy Hand-to-Mouth. Brookings Papers on Economic Activity 48, 77–153.

Keys, B.J., Wang, J., 2019. Minimum payments and debt paydown in consumer credit cards. Journal of Financial Economics 131, 528–548. https://doi.org/10.1016/j.jfineco.2018.09.009

Kovacs, A., Moran, P., 2019. Temptation and commitment: Understanding the demand for illiquidity. IFS Working Paper W19/18. https://doi.org/10.1920/wp.ifs.2019.1819

Laibson, D., 1997. Golden Eggs and Hyperbolic Discounting. The Quarterly Journal of Economics 112, 443–478.

Laibson, D.I., Repetto, A., Tobacman, J., Hall, R.E., Gale, W.G., Akerlof, G.A., 1998. Self-Control and Saving for Retirement. Brookings Papers on Economic Activity 1998, 91–196. https://doi.org/10.2307/2534671

Madrian, B.C., Shea, D.F., 2001. The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. The Quarterly Journal of Economics 116, 1149–1187.

Navarro-Martinez, D., Salisbury, L.C., Lemon, K.N., Stewart, N., Matthews, W.J., Harris, A.J., 2011. Minimum Required Payment and Supplemental Information Disclosure Effects on Consumer Debt Repayment Decisions. Journal of Marketing Research 48, S60–S77. https://doi.org/10.1509/jmkr.48.SPL.S60

Olafsson, A., Pagel, M., 2017. The Ostrich in Us: Selective Attention to Financial Accounts, Income, Spending, and Liquidity (Working Paper No. 23945). National Bureau of Economic Research. https://doi.org/10.3386/w23945

Parker, J.A., Souleles, N.S., Johnson, D.S., McClelland, R., 2013. Consumer Spending and the

Economic Stimulus Payments of 2008. The American Economic Review 103, 2530–2553.

Schlafmann, K., 2016. Housing, Mortgages, and Self Control (SSRN Scholarly Paper No. ID 2862206). Social Science Research Network, Rochester, NY.

Shapiro, J.M., 2005. Is there a daily discount rate? Evidence from the food stamp nutrition cycle. Journal of Public Economics 89, 303–325. https://doi.org/10.1016/j.jpubeco.2004.05.003

Shapiro, M.D., Slemrod, J., 2009. Did the 2008 Tax Rebates Stimulate Spending? American Economic Review 99, 374–379. https://doi.org/10.1257/aer.99.2.374

Souleles, N.S., 1999. The Response of Household Consumption to Income Tax Refunds. The American Economic Review 89, 947–958.

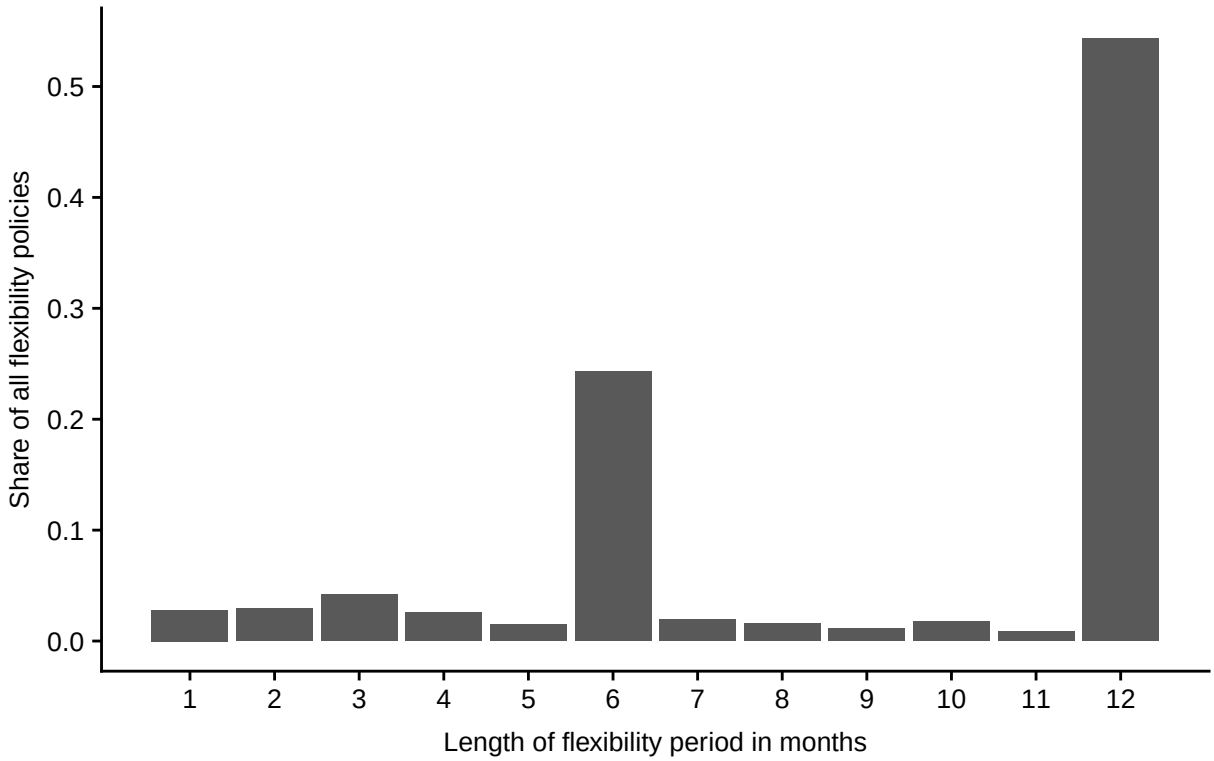
Stewart, N., 2009. The Cost of Anchoring on Credit-Card Minimum Repayments. Psychological Science 20, 39–41. https://doi.org/10.1111/j.1467-9280.2008.02255.x

Vihriälä, E., 2020. Untangling the credit card debt puzzle. Mimeo.



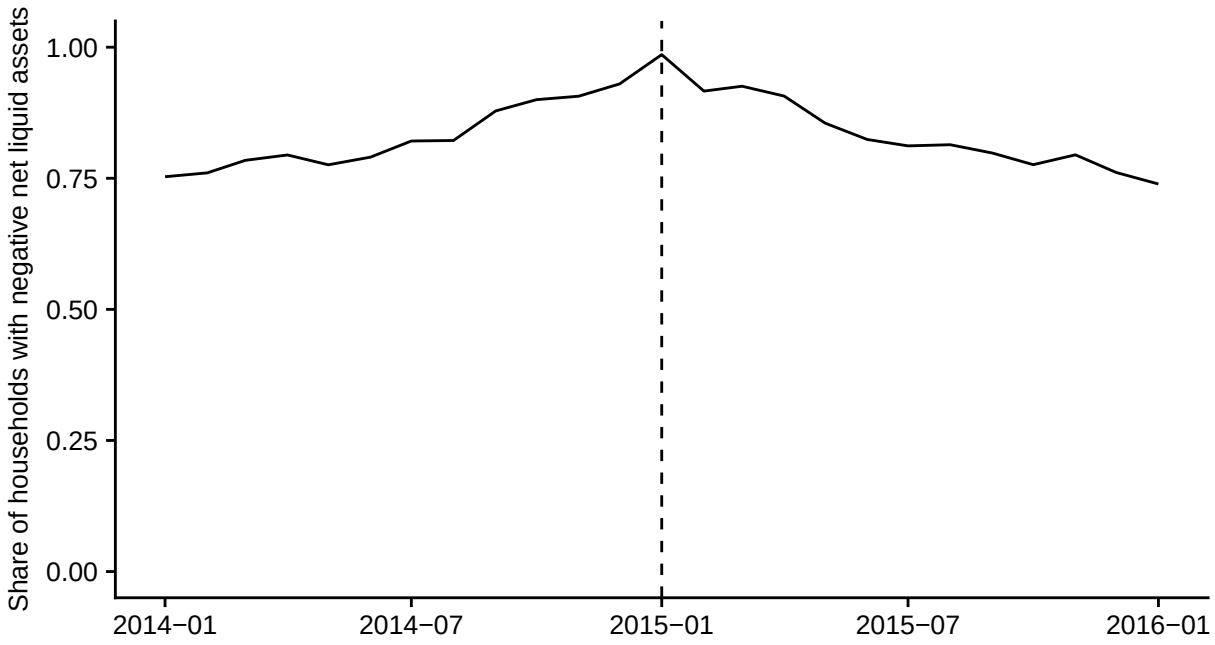
Note: The figure plots the persistence of liquidity constraints among households classified as liquidity constrained in the baseline month of January 2015 who do not apply for flexibility. I define a household as liquidity constrained if the household has negative net liquid assets, where net liquid assets equal the sum of deposit balances, mutual fund holdings and stock holdings minus interest−paying unsecured credit.

Figure 1: Liquidity constraints are persistent among non-applicants



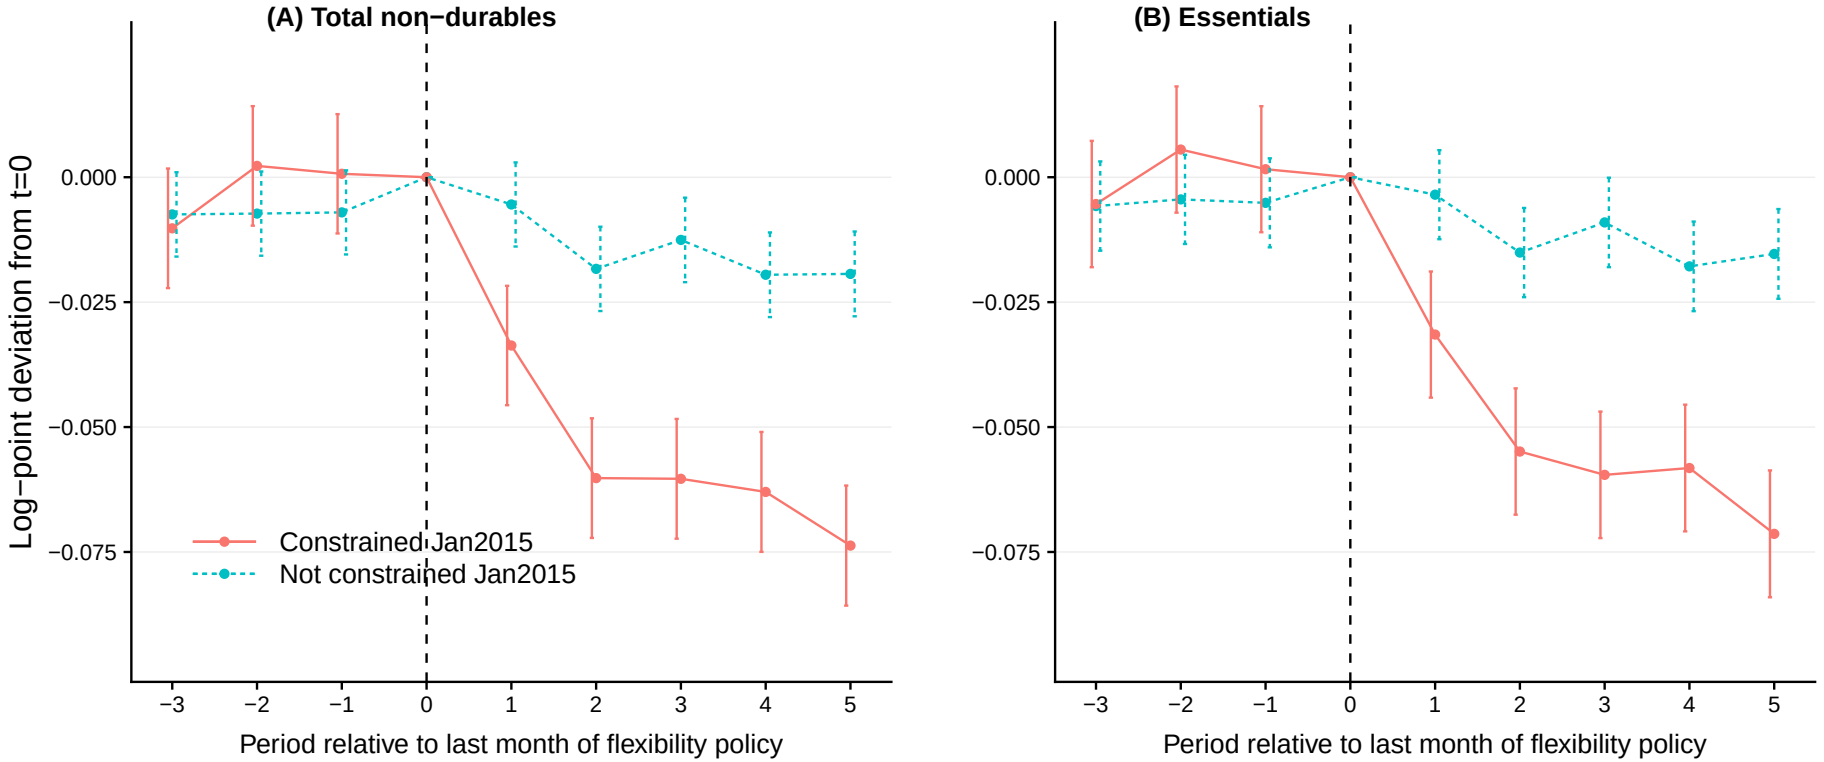
Note: The figure plots the popularity of flexibility policies of different lengths among applicants. The share of applications for policies shorter than maximum is 46 percent.

Figure 2: Almost half of applicants choose a flexibility policy shorter than maximum



Note: The figure plots the persistence of liquidity constraints among households who are liquidity constrained in January 2015 and apply for a six−month flexibility policy. I define a household as liquidity constrained if the household has negative net liquid assets, where net liquid assets equal the sum of deposit balances, mutual fund holdings and stock holdings minus interest−paying unsecured credit.

Figure 3: Liquidity constraints are persistent among applicants for six-month flexibility



Note: The figure plots estimated delta coefficients from equation 14 that measure the deviation of consumption from the level in the last month of flexibility. Panel A depicts results from regressions with logarithm of total non−durable expenditure as the dependent variable. Panel B represents results with logarithm of essentials expenditure (groceries, energy, public transportation, gasoline) as dependent variable. For each dependent variable I estimate equation 14 separately for households who are liquidity constrained before offer announcement and for unconstrained households. Applicant households include those who apply for a 12−month flexibility policy so that time from the start of the flexibility policy is the same for all applicants. I remove monthly variation in consumption by using non−applicants as controls. Period 0 refers to the last month of the flexibility policy. Vertical bars represent 95 percent confidence intervals.

Figure 4: Event study: liquidity-constrained applicants do not smooth consumption at the end of flexibility

Table 1: Predicted monthly consumption drop at the end of flexibility for exponential discounter

	ρ							
	0.5	1 (log-utility)	1.5	2	2.5	3	3.5	4
Annualized δ								
0.800	0.963	0.982	0.988	0.991	0.993	0.994	0.995	0.995
0.850	0.973	0.987	0.991	0.993	0.995	0.995	0.996	0.997
0.900	0.983	0.991	0.994	0.996	0.996	0.997	0.997	0.998
0.950	0.991	0.996	0.997	0.998	0.998	0.999	0.999	0.999
0.975	0.996	0.998	0.999	0.999	0.999	0.999	0.999	0.999

[a] The table presents the predicted monthly consumption drop at the end of flexibility as a ratio ($C_{post-flex}/C_{flex}$) based on the Euler equation for a household with CRRA utility for a range of coefficients of relative risk aversion ρ and exponential discount factor δ. For ease of interpretation, tabulated values of δ represent annualized values. The Euler equation abstracts from uncertainty and interest rates, which overestimates the expected consumption drop.

Table 2: Descriptive statistics for main variables

	N=204,246							
	Mean	SD	Q0.01	Q0.25	Median	Q0.75	Q0.99	Unit
Apply for flexibility (0/1)	0.22	0.41	0.00	0.00	0.00	0.00	1.00	Share
Demographics								
Adults	1.85	0.63	1.00	1.00	2.00	2.00	4.00	Number
Children	0.73	1.08	0.00	0.00	0.00	1.00	4.00	Number
Age	45.34	12.18	25.00	35.50	44.00	54.00	74.50	Years
Female	0.51	0.27	0.00	0.50	0.50	0.50	1.00	Share
Mortgage/housing								
Outstanding mortgage	93,792	77,105	2,000	35,959	77,742	131,967	342,118	EUR
Principal payment	634	417	75	364	545	805	2,060	EUR/month
Interest rate	1.33	0.52	0.25	0.98	1.25	1.64	2.57	Percent
Margin rate	0.90	0.47	0.00	0.55	0.80	1.20	2.01	Percent
Mortgage age	4.71	3.38	0.00	2.00	4.00	7.00	14.00	Years
Remaining maturity	11.15	6.18	0.00	6.00	11.00	16.00	24.00	Years
Property value	195,013	132,953	36,039	112,101	164,282	242,147	669,549	EUR
Net housing equity	101,220	105,227	-35,308	36,029	79,309	138,667	466,788	EUR
Income and expenditure								
Disposable income	3,647	1,676	825	2,404	3,517	4,506	9,270	EUR/month
Card expenditure	1,804	1,024	276	1,040	1,683	2,375	4,798	EUR/month
Financial assets								
Deposits	13,298	30,164	-1,355	1,939	5,697	15,008	107,364	EUR
Other assets	7,157	197,738	0	0	0	2,066	111,669	EUR
Other credit								
Total unsecured credit	1,916	2,666	0	87	874	2,637	12,020	EUR
Interest-paying unsecured credit	1,465	2,558	0	0	0	1,895	11,452	EUR
Credit card utilization	0.37	0.36	0.00	0.03	0.26	0.71	1.00	Pct of limit

[a] All variables except disposable income are determined before offer announcement in February 2015. Most variables refer to end-of-January 2015 values, but card expenditure is a monthly average over 2014.

[b] Disposable income is average monthly labor income after taxes based on administrative tax information for 2015.

[c] The principal payment measure is the default minimum principal payment due after the announcement of the flexibility offer. The measure is a monthly average of minimum principal payments due from February 2015 to June 2016.

[d] Total unsecured credit includes both interest-paying unsecured credit and the part of unsecured credit with an interest-free period.

Table 3: Characteristics of households by offer application

		Apply (N=44,292)			Do not apply (N=159,954)		
	Unit	Mean	SD	Median	Mean	SD	Median
Demographics							
Adults	Number	1.87	0.62	2.00	1.85	0.63	2.00
Children	Number	0.94	1.19	0.00	0.67	1.04	0.00
Age	Years	43.09	11.03	41.00	45.96	12.41	45.00
Female	Share	0.51	0.26	0.50	0.51	0.27	0.50
Mortgage/housing							
Outstanding mortgage	EUR	118,832	81,293	104,356	86,859	74,430	70,173
Principal payment	EUR/month	735	428	644	606	409	517
Interest rate	Percent	1.25	0.51	1.20	1.35	0.52	1.28
Margin rate	Percent	0.84	0.46	0.75	0.92	0.48	0.80
Mortgage age	Years	4.79	3.11	4.14	4.69	3.45	4.00
Remaining maturity	Years	12.66	5.50	13.00	10.74	6.29	10.47
Property value	EUR	207,040	133,652	176,076	191,682	132,566	161,014
Net housing equity	EUR	88,208	98,006	67,891	104,823	106,860	82,573
Income and expenditure							
Disposable income	EUR/month	3,621	1,584	3,519	3,654	1,701	3,517
Card expenditure	EUR/month	1,852	971	1,757	1,791	1,037	1,661
Financial assets							
Deposits	EUR	7,022	14,883	2,919	15,035	32,963	6,917
Other assets	EUR	3,906	26,721	0	8,058	222,994	0
Other credit							
Total unsecured credit	EUR	2,939	3,195	1,902	1,633	2,425	664
Interest-paying unsecured credit	EUR	2,420	3,138	1,271	1,200	2,304	0
Credit card utilization	Pct of limit	0.53	0.36	0.57	0.33	0.34	0.20

[a] All variables except disposable income are determined before offer announcement in February 2015. Most variables refer to end-of-January 2015 values, but card expenditure is a monthly average over 2014.

[b] Disposable income is average monthly labor income after taxes based on administrative tax information for 2015.

[c] The principal payment measure is the default minimum principal payment due after the announcement of the flexibility offer. The measure is a monthly average of minimum principal payments due from February 2015 to June 2016.

[d] Total unsecured credit includes both interest-paying unsecured credit and the part of unsecured credit with an interest-free period.

Table 4: Predicting application: linear probability model
(application rate = 0.22; N = 204,246; R^2 = 0.12)

Predictor	Group	Estimate	T-statistic	P-value
Age	Q2	0.00	0.93	0.35
	Q3	0.00	1.65	0.10
	Q4	-0.01	-2.78	0.01
Deposits	Q2	-0.05	-19.14	0.00
	Q3	-0.11	-38.86	0.00
	Q4	-0.17	-58.46	0.00
Total unsecured credit	Q2	0.02	5.66	0.00
	Q3	0.05	10.80	0.00
	Q4	0.10	18.49	0.00
Credit card utilization	Q2	0.01	1.98	0.05
	Q3	0.02	3.85	0.00
	Q4	0.08	13.87	0.00
Disposable income	Q2	-0.03	-11.79	0.00
	Q3	-0.07	-21.71	0.00
	Q4	-0.11	-28.51	0.00
Card expenditure	Q2	0.00	1.86	0.06
	Q3	0.01	2.68	0.01
	Q4	0.01	3.58	0.00
Principal payment	Q2	0.05	20.78	0.00
	Q3	0.08	30.29	0.00
	Q4	0.14	40.73	0.00
Property value	Q2	0.00	-0.17	0.86
	Q3	-0.01	-4.01	0.00
	Q4	-0.03	-8.59	0.00
Outstanding mortgage	Q2	0.05	21.41	0.00
	Q3	0.08	26.69	0.00
	Q4	0.13	34.62	0.00
Mortgage interest rate	Q2	-0.03	-10.57	0.00
	Q3	-0.05	-17.30	0.00
	Q4	-0.05	-16.24	0.00
Household includes bank employee		0.11	19.94	0.00

[a] The table presents estimates from a linear probability model in which the dependent variable is 1 if household applies for flexibility and 0 otherwise.
[b] The omitted category for binned variables is the bottom quartile.
[c] The regression includes additional controls for number of adults and children, municipality and mortgage age but I do not report the coefficients for brevity.
[d] The R^2 of the linear probability model is the difference in the mean predicted application probability by application status.
[e] Standard errors robust to heteroskedasticity.
[f] Total unsecured credit includes both interest-paying unsecured credit and the part of unsecured credit with and interest-free period.

Table 5: Attention to household finances correlates positively with application

	Dependent variable:	
	Apply for flexibility (0/1)	
	(1)	(2)
Share of online messages read	0.04***	
	(0.002)	
Login Quartile 2		0.03***
		(0.002)
Login Quartile 3		0.05***
		(0.003)
Login Quartile 4		0.08***
		(0.003)
Household controls	Yes	Yes
Application rate	0.22	0.22
Observations	201,164	201,803
R^2	0.12	0.12

Note: *p<0.1; **p<0.05; ***p<0.01

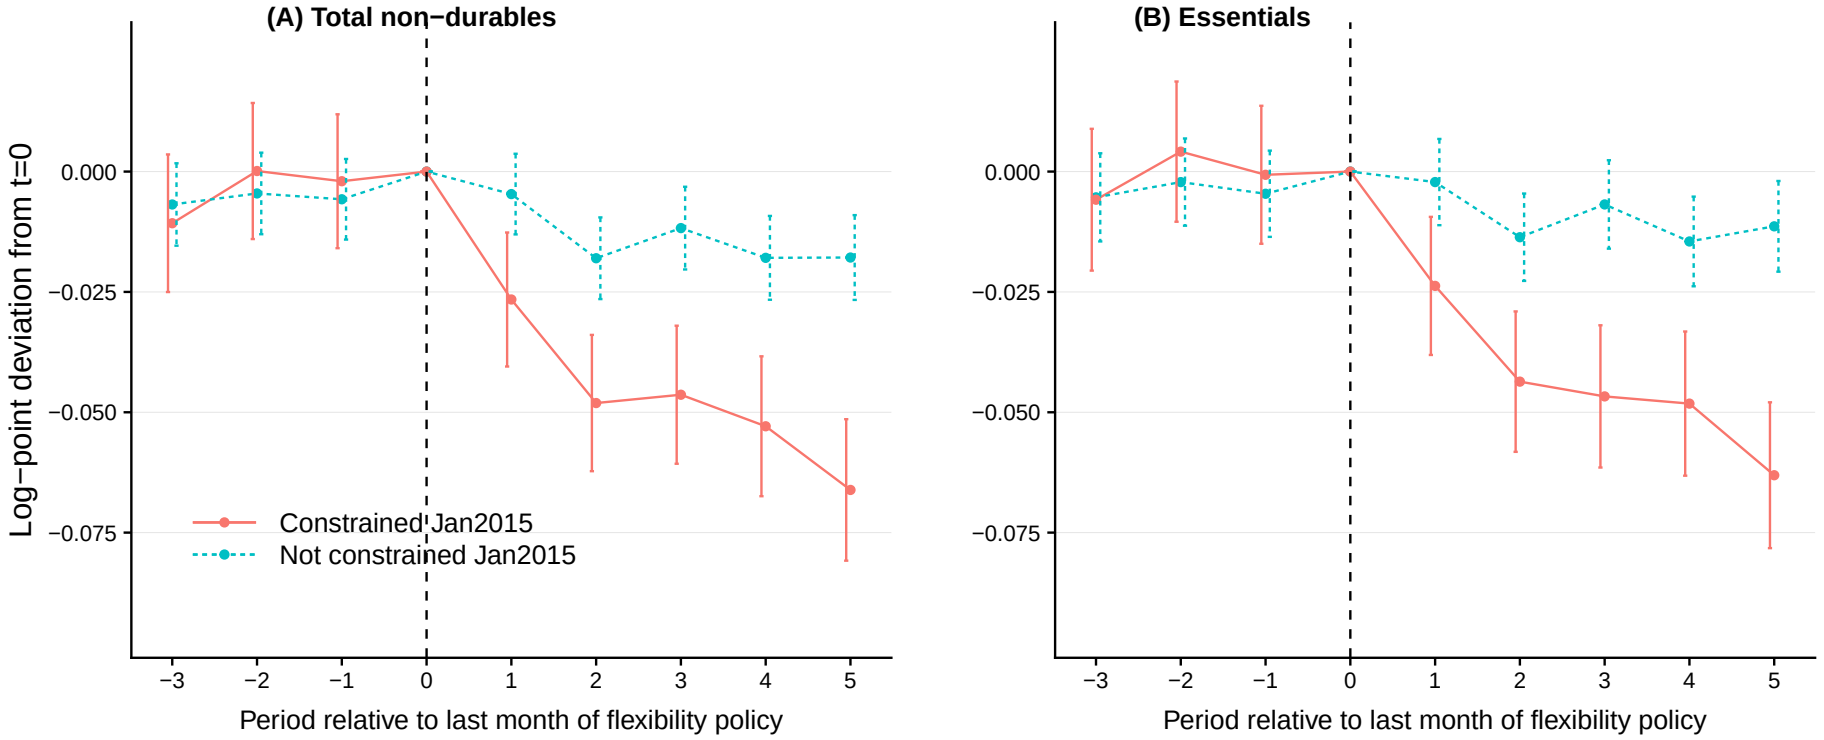
The table presents estimation results from two linear probability models in which the dependent variable is 1 if the household applies for flexibility and 0 otherwise. Model (1) uses the share of online messages read prior to the offer as an attention proxy. Model (2) uses the frequency of logins to online and mobile banking in 2017 (data prior to offer not available). Standard errors robust to heteroskedasticity in parentheses. Both specifications control for household characteristics present in the baseline linear probability model without attention proxies (Table 4). In addition, the regression based on share of online messages read controls for the total amount of messages received. In Model (1) I include households who received at least one online message in 2014. In Model (2) I include households who logged in online at least once in 2017. These conditions explain the small discrepancy in the number of observations relative to the full sample (N = 204,246). The R^2 of the linear probability model is the difference in the mean predicted application probability by application status.

Table 6: Event study: liquidity-constrained applicants do not smooth consumption at the end of flexibility

	Dependent variable:			
	Total non-durables	Total non-durables	Essentials	Essentials
	(1)	(2)	(3)	(4)
9th flexibility month	-0.007^{*}	-0.010^{*}	-0.006	-0.005
	(0.004)	(0.006)	(0.005)	(0.006)
10th flexibility month	-0.007^{*}	0.002	-0.004	0.006
	(0.004)	(0.006)	(0.005)	(0.006)
11th flexibility month	-0.007	0.001	-0.005	0.002
	(0.004)	(0.006)	(0.005)	(0.006)
1st post-flexibility month	-0.005	-0.034^{***}	-0.003	-0.031^{***}
	(0.004)	(0.006)	(0.005)	(0.006)
2nd post-flexibility month	-0.018^{***}	-0.060^{***}	-0.015^{***}	-0.055^{***}
	(0.004)	(0.006)	(0.005)	(0.006)
3rd post-flexibility month	-0.013^{***}	-0.060^{***}	-0.009^{**}	-0.060^{***}
	(0.004)	(0.006)	(0.005)	(0.006)
4th post-flexibility month	-0.020^{***}	-0.063^{***}	-0.018^{***}	-0.058^{***}
	(0.004)	(0.006)	(0.005)	(0.006)
5th post-flexibility month	-0.019^{***}	-0.074^{***}	-0.015^{***}	-0.071^{***}
	(0.004)	(0.006)	(0.005)	(0.006)
Household sample	Unconstrained	Constrained	Unconstrained	Constrained
N(Applicants)	16,784	8,408	16,780	8,408

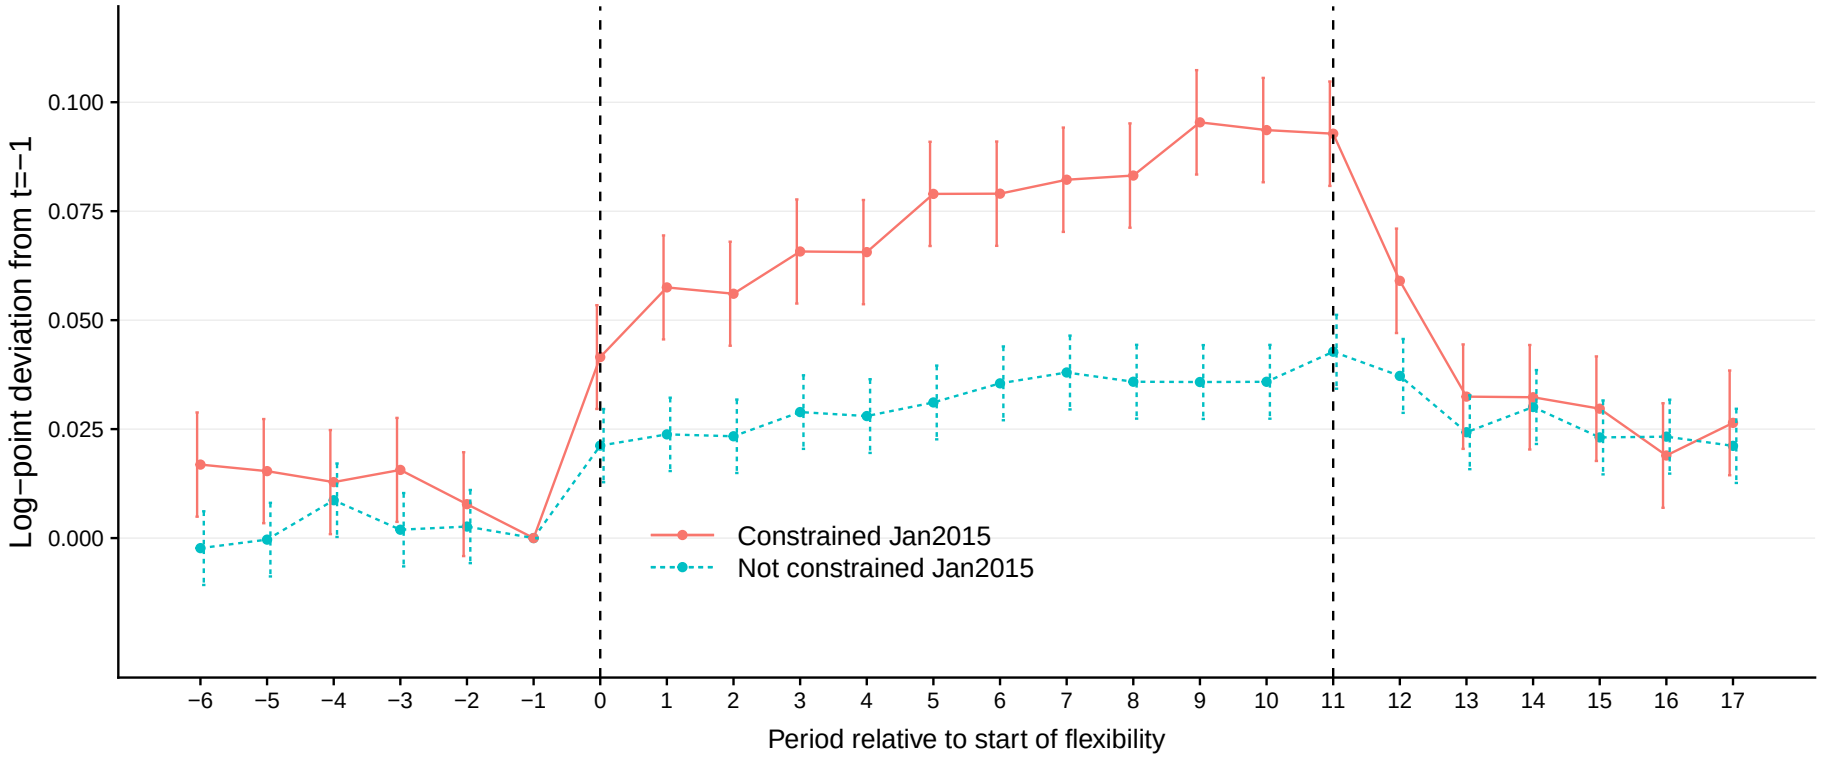
Note: *p<0.1; **p<0.05; ***p<0.01

The table presents estimates of the delta coefficients in equation 14. The omitted baseline month is the last flexibility month. Models 1 and 2 use the logarithm of total non-durable expenditure as the dependent variable. Models 3 and 4 present results with logarithm of essentials expenditure (groceries, energy, public transportation, gasoline) as the dependent variable. Standard errors clustered by household in parentheses. For each dependent variable I estimate equation 14 separately for households who are liquidity constrained before offer announcement and for unconstrained households. Applicant households include those who apply for a 12-month flexibility policy so that time from the start of the flexibility policy is the same for all applicants. I present the results graphically in Figure 4.



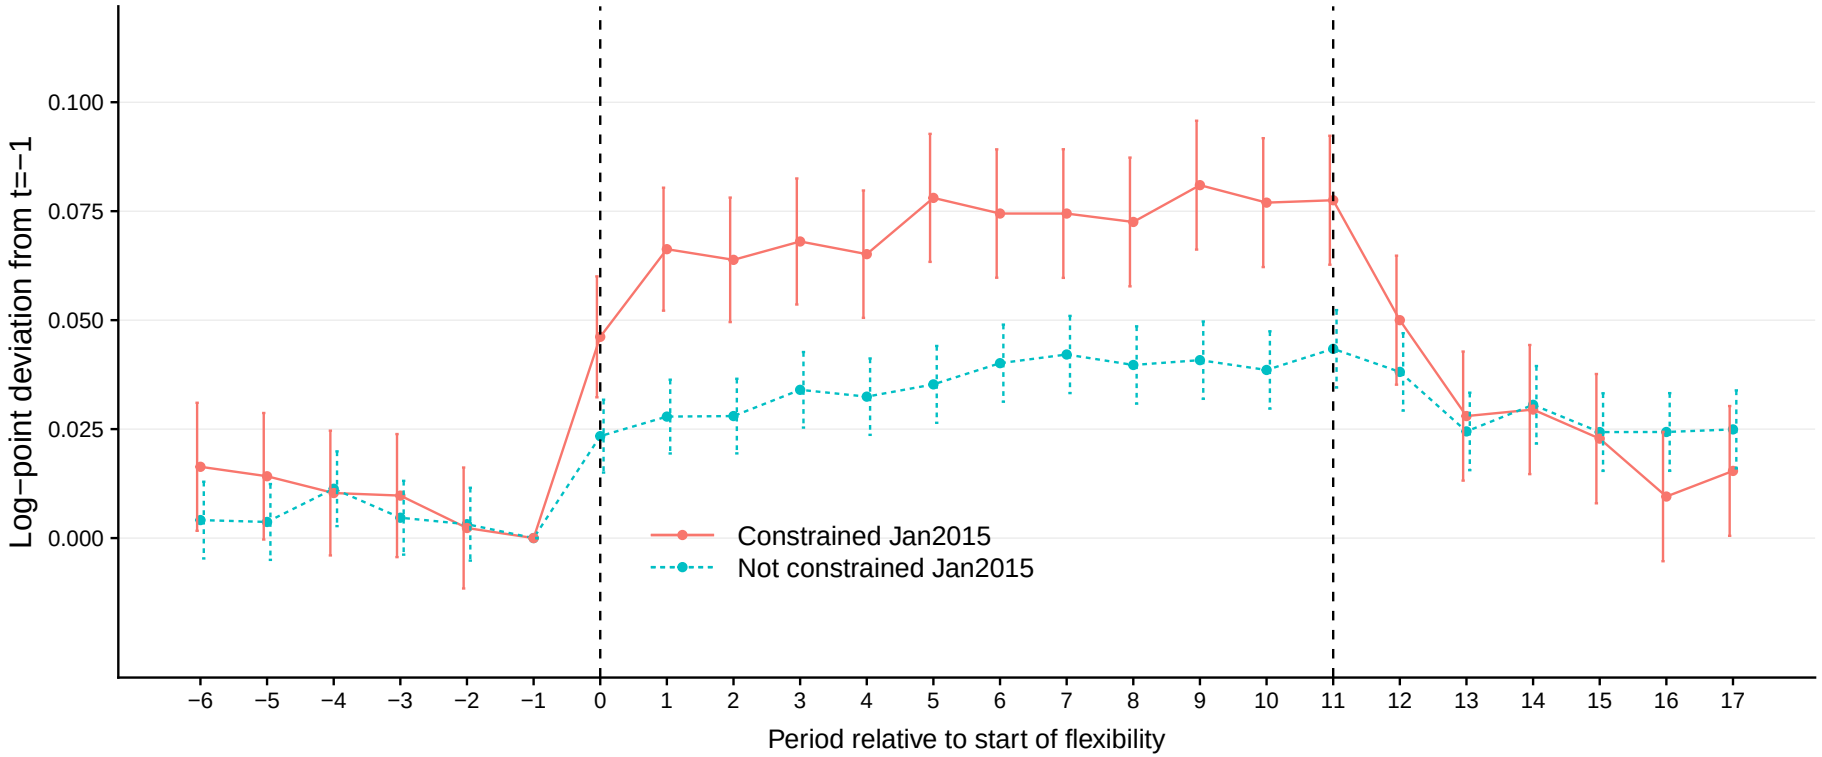
Note: The figure plots estimated delta coefficients from equation 14 that measure the deviation of consumption from the level in the last month of flexibility. Panel A depicts results from regressions with log of total non−durable expenditure as the dependent variable. Panel B represents results with log of essentials expenditure (groceries, energy, public transportation, gasoline) as dependent variable. For each dependent variable I estimate equation 14 separately for households who are liquidity constrained before offer announcement and for unconstrained households. Applicant households include those who apply for a 12−month flexibility policy so that time from the start of the flexibility policy is the same for all applicants. I remove monthly variation in consumption by using matched non−applicants as controls. I use propensity score matching within a particular household structure, age bin, urbanization category, bank−employee status, and liquidity constraint dummy. Period 0 refers to the last month of the flexibility policy. Vertical bars represent 95 percent confidence intervals.

Figure A.1: Event study: liquidity-constrained applicants do not smooth consumption at the end of flexibility (matched sample)



Note: The figure plots the consumption path for households who apply for a twelve−month flexiblity policy before, during, and after the flexibility policy. The dependent variable is log of total non−durable expenditure. I estimate the consumption path separately for households liquidity constrained and not constrained in January 2015. The figure complements the figure focusing on consumption smoothing at the end of flexibility. I estimate consumption deviation relative to the month before the start of the flexibility policy (period −1). I remove monthly variation in consumption by using non−applicants as controls. Period 0 refers to the last month of the flexibility policy. Vertical bars around point estimates represent 95 percent confidence intervals. Dashed lines at period 0 and 11 refer to the first and last month of the flexibility policy.

Figure A.2: Consumption before, during, and after flexibility



Note: The figure plots the consumption path for households who apply for a twelve−month flexiblity policy before, during, and after the flexibility policy. The dependent variable is log of total non−durable expenditure. I estimate the consumption path separately for households liquidity constrained and not constrained in January 2015. The figure complements the figure focusing on consumption smoothing at the end of flexibility. I estimate consumption deviation relative to the month before the start of the flexibility policy (period −1). I remove monthly variation in consumption by using a matched−on−observables sample of non−applicants as controls. Period 0 refers to the last month of the flexibility policy. Vertical bars around point estimates represent 95 percent confidence intervals. Dashed lines at period 0 and 11 refer to the first and last month of the flexibility policy.

Figure A.3: Consumption before, during, and after flexibility (matched sample)

Table A.1: Event study: liquidity-constrained applicants do not smooth consumption at the end of flexibility (matched sample)

	Dependent variable:			
	Total non-durables	Total non-durables	Essentials	Essentials
	(1)	(2)	(3)	(4)
9th flexibility month	−0.007	−0.011	−0.005	−0.006
	(0.004)	(0.007)	(0.005)	(0.008)
10th flexibility month	−0.005	0.0001	−0.002	0.004
	(0.004)	(0.007)	(0.005)	(0.007)
11th flexibility month	−0.006	−0.002	−0.005	−0.001
	(0.004)	(0.007)	(0.005)	(0.007)
1st post-flexibility month	−0.005	−0.027***	−0.002	−0.024***
	(0.004)	(0.007)	(0.005)	(0.007)
2nd post-flexibility month	−0.018***	−0.048***	−0.014***	−0.044***
	(0.004)	(0.007)	(0.005)	(0.007)
3rd post-flexibility month	−0.012***	−0.046***	−0.007	−0.047***
	(0.004)	(0.007)	(0.005)	(0.008)
4th post-flexibility month	−0.018***	−0.053***	−0.015***	−0.048***
	(0.004)	(0.007)	(0.005)	(0.008)
5th post-flexibility month	−0.018***	−0.066***	−0.011**	−0.063***
	(0.004)	(0.007)	(0.005)	(0.008)
Household sample	Unconstrained	Constrained	Unconstrained	Constrained
N(Applicants)	16,712	8,293	16,708	8,293

Note: *p<0.1; **p<0.05; ***p<0.01

The table presents estimates of the delta coefficients in equation 14 that measure the deviation of consumption from the level in the last month of flexibility. The omitted baseline month is the last flexibility month. Models 1 and 2 use the log of total non-durable expenditure as the dependent variable. Models 3 and 4 present results with log of essentials expenditure (groceries, energy, public transportation, gasoline) as the dependent variable. Standard errors clustered by household in parentheses. For each dependent variable I estimate equation 14 separately for households who are liquidity constrained before offer announcement and for unconstrained households. Applicant households include those who apply for a 12-month flexibility policy. I remove monthly variation in consumption by using matched non-applicants as controls. I use propensity score matching within a particular household structure, age bin, urbanization category, bank-employee status, and liquidity constraint dummy. The results are presented graphically in Figure A.1.

Untangling the credit card debt puzzle

Erkki Vihriälä

Abstract

This paper documents new facts and untangles various mechanisms regarding the credit card debt puzzle describing the co-holding of high-cost debt and low-yield assets. First, I document the persistence of co-holding at the monthly frequency. The prevailing method of extrapolating from a cross-section overestimates the costs of co-holding for the typical household, because extrapolation misses the imperfect persistence of co-holding. Second, liquidity-based theories are insufficient to explain co-holding because many households with plenty of liquidity and low negative credit limit risk co-hold. Furthermore, a natural experiment reveals that few co-holding households take up an offer of low-cost liquidity. Third, my results suggest that intra-household dynamics can significantly impact the magnitude of the puzzle. I reject the hypothesis that households pool finances, because the intra-household distribution of assets and liabilities influences the take up of the liquidity offer. Finally, my results suggest that anchoring to the minimum credit card payment contributes to co-holding.

Keywords: household finance, credit card debt puzzle, intra-household financial pooling
JEL codes: D13, D14, D91, E21

1 Introduction

Many households borrow and save at the same time by holding both high-cost debt and low-yield liquid assets. Across vintages of the Survey of Consumer Finances (SCF), around 30 percent of households in the US co-hold as noted first by Gross and Souleles (2002). The decision not to use liquid assets to repay high-cost debt to economize on interest costs has been labeled the *credit card debt puzzle* or the *co-holding puzzle*[1].

Despite earlier explanations for the puzzle, we still do not understand many features of co-holding because of both data limitations and a lack of policy experiments. For instance, we do not know (i) how persistent co-holding is because of the lack of high-frequency panel data; (ii) whether households would stop co-holding if they had access to sufficient liquidity; (iii) how the intra-household distribution of assets and liabilities impacts co-holding; (iv) or whether anchoring of credit card payments to the contractual minimum contributes to co-holding. My paper contributes on these fronts with a new Finnish dataset with information on hundreds of thousands of households at a monthly frequency combined with evidence from a policy experiment.

First, I show that a cross-section, such as the SCF used in much of the literature, does not give an accurate picture of the persistence of co-holding. Households classified as co-holders in one month on average reduce co-holding over the subsequent year. Therefore, extrapolation from a cross-section tends to overestimate the costs of co-holding for the typical household. The mean discrepancy in extrapolated versus actual costs is not large, with actual costs on average 90 percent of extrapolated costs. Yet, the spread in the ratio of actual to extrapolated costs across households is significant (45 percent in the 10th percentile versus 130 percent in the 90th

[1]Co-holding puzzle is a more accurate description when not restricting attention to only credit card debt but also other high-cost borrowing (Gathergood and Weber, 2014).

percentile).

Second, I test whether the two liquidity-based theories are sufficient to explain co-holding. The *liquidity premium hypothesis* suggests that cash is more liquid than unused credit, and hence co-holding arises because a cash buffer is essential for certain expenses (Telyukova, 2013; Zinman, 2007). Yet, I find that co-holding often coincides with high liquidity well above a modest cash buffer, in line with earlier results (see, for instance, Gross and Souleles, 2002). By contrast, the *precautionary borrowing hypothesis* suggests that co-holding arises because households fear that banks would reduce their credit lines after a negative shock. The risk of, for instance, a job loss coinciding with a credit reduction creates an incentive to borrow and co-hold in advance (Druedahl and Jørgensen, 2018; Fulford, 2015; Gorbachev and Luengo-Prado, 2016). Yet, co-holding occurs also in my data despite the low negative credit limit risk in Finland. Finally, I test the liquidity premium and the precautionary borrowing hypotheses with a natural experiment that provided certain co-holding households access to new low-cost liquidity. According to liquidity-based theories, co-holding households should find an offer of low-cost liquidity beneficial. However, only one quarter of co-holding households take up the extra liquidity, and I find no evidence that access to additional liquidity would have reduced co-holding.

Third, I provide evidence on the effect of intra-household dynamics on co-holding, which has been understudied previously because of data limitations. The standard approach of defining co-holding at the household level contains an implicit assumption of perfect intra-household pooling of assets and liabilities. However, the accountant-shopper models of Bertaut and Haliassos (2001) and Bertaut et al. (2009) used to explain co-holding can reflect a strategic game between two separate individuals. I provide supporting empirical evidence that in multi-adult households (i) co-holding is less prevalent at the individual than at the household level; (ii) the persistence

of co-holding is lower at the individual than at the household level; (iii) the intra-household distribution of assets and liabilities impacts the take up of the bank's liquidity offer. Therefore, part of the co-holding in multi-member households is not puzzling if we relax the assumption of perfect intra-household pooling. My results on intra-household dynamics complement the findings of Choi and Laschever (2018) that income inequality and agreeableness within the couple impact co-holding.

Fourth, I providence evidence consistent with the hypothesis that achoring (Tversky and Kahneman, 1974) to the minimum credit card payment contributes to co-holding. Various studies have shown that credit card debt payments are consistent with anchoring (Guttman-Kenney et al., 2018; Hershfield and Roese, 2015; Keys and Wang, 2016; Navarro-Martinez et al., 2011; Stewart, 2009). My paper corroborates, in particular, the findings of Keys and Wang (2016) regarding the pervasiveness of low credit card payments in a similar real-world setting. I further support the conclusion of Keys and Wang (2016) that liquidity constraints (instead of anchoring) are unlikely to explain low credit card payments, because I observe directly the (often high) liquidity associated with low payments. Consequently, my paper bridges the literature on anchoring in credit card payments and the literature on the credit card debt puzzle by showing that anchoring households are more likely to co-hold.

Finally, the earlier literature and my setup allow to rule out other mechanisms for co-holding. For instance, strategic bankruptcy motives (Lehnert and Maki, 2002) are unlikely to explain co-holding in Finland where households find it much harder to restructure/discharge their unsecured debts than in the US. Self-control problems at the individual level are unlikely to explain co-holding in my data, because the high share of electronic transactions and ubiquitous debit cards make deposits liquid in Finland. The liquidity of deposits in my data does not imply, though, that self-control problems could not explain co-holding as an equilibrium in

multi-member households or in circumstances where deposits are less liquid. Indeed, Gathergood and Weber (2014) show that impulsiveness predicts co-holding with UK household data. Finally, financial illiteracy is unlikely to explain the puzzle, because earlier research has documented that co-holding households have relative high financial literacy (Gathergood and Weber, 2014; Gorbachev and Luengo-Prado, 2016).

The paper is structured as follows. In Section 2, I define the credit card debt puzzle and introduce the data. In Section 3, I document the persistence and costs of co-holding. In Section 4, I show that liquidity-based theories are insufficient to explain co-holding. In Section 5, I show that intra-household dynamics can contribute to co-holding in multi-member households. In Section 6, I provide evidence consistent with anchoring to the minimum credit card payment influencing co-holding. Section 7 concludes.

2 Definitions and data

In this section I first define the credit card debt puzzle. Second, I present the Finnish proprietary bank data. Finally, I document the sample selection procedure and provide descriptive statistics on the sample households.

2.1 Definition of the credit card debt puzzle

The simplest definition would classify any household with simultaneous positive holdings of liquid assets and interest-accruing credit card debt as a "puzzling" co-holder. Yet, low levels of co-holding carry negligible costs and can merely reflect differential timing of expenditure and income flows. Therefore, to classify a household as a co-holder, I require simultaneous holding of

(i) non-negligible interest-accruing credit card debt and (ii) liquid assets above a buffer.

Another issue is the definition of the asset and borrowing subsets underlying the puzzle. I opt to consider all balances in checking and savings accounts with no or limited withdrawal restrictions as liquid assets. In terms of the puzzle, credit card debt is not special in comparison to other unsecured debt at similar interest rates. Therefore, in defining the puzzle, I sum all unsecured debt.

Consequently, the baseline puzzle criteria at the household level are:

1. Have interest-accruing unsecured debt above 500 EUR
2. Have deposits greater than 1,500 EUR

The two criteria allow me to focus on non-negligible co-holding while allowing for a liquidity buffer as is standard in the literature.

2.2 Data

The data used in this study comes from one of the largest banks in Finland. I have customer-level data on, for instance, assets, debts, income and card expenditure. Most of the data is available at a monthly frequency for 2014–2017. I can link customers belonging to the same household using a household identifier.

The main variables to analyze co-holding are:

Unsecured debt: I observe interest-accruing credit card debt and other unsecured debt at the bank. In addition to credit cards, the bank offers two different unsecured borrowing products. The first of these is a flexible credit line for repeated use. The second is a traditional credit contract designed for a single purchase. Nominal interest rate margins on both credit cards and

these two unsecured borrowing products are similar (in the 5 to 10 percent range). The most common product has a median interest rate margin of 7 percent. The total nominal interest rate on all the products is the interest rate margin rate + the 3 month Euribor reference rate, but the Euribor rate was close to zero during the sample period. None of the debt products preclude full or partial prepayment. In calculating co-holding, I sum all interest-accruing unsecured debt.

Liquid assets: I include in liquid assets balances on all checking and savings accounts that do not have any or only limited withdrawal restrictions[2]. These deposits are available flexibly to repay unsecured debt. The nominal yield on these assets during the sample period (2014–2017) was essentially zero.

2.3 Sample selection

The pool of customers at the bank is large and diverse, which matters for the representativeness of the results. For instance, in mortgages the market share of the bank is around 1/3. The bank also has branches in all parts of the country.

A concern with data from a single bank is that I do not observe households' financial activity at other banks. One mitigating factor is that the bank has a loyalty scheme that encourages households to concentrate their full financial portfolio (assets, debts, insurance) at the bank. Still, a significant chunk of customers are only loosely tied to the bank. Therefore, the criteria to be included in the sample population are:

1. Household needs to have at least one person over 25 years old
2. Household needs to have a credit card issued by the bank

[2]Checking accounts have no restrictions. Some of the savings accounts have restrictions on the number of free withdrawals per year (typically four). However, these savings accounts also have yields close to zero percent over the sample period. Therefore, keeping money in savings accounts is arguably 'puzzling' if the household faces a non-negligible probability to use unsecured debt.

3. No household member has stated another main financial service provider
4. All over 25 year olds have a checking account at the bank
5. All over 25 year olds regularly use cards issued by the bank during 2014–2017
6. Household disposable income is over 500 EUR per month in 2015
7. Information on deposits is not missing in any month in 2014–2017
8. Household uses one of the two main languages in the country (Finnish/Swedish)
9. I drop entrepreneur households not to mix personal and business accounts
10. Household has at most two members over 20 years (to omit young flatshares)

Based on these criteria, the main sample consists of 357,950 households.

2.4 Sample descriptives

Table 1 presents descriptive statistics on the key co-holding variables both for the full sample and the 20 percent of households that satisfy the co-holding criteria in the baseline month of January 2015. The median amount of co-holding among co-holders is about 1,800 EUR, slightly more than 50 percent of monthly disposable income. The share of households who co-hold rises somewhat during the sample period from 2014 to 2017 and ranges from 17 percent to 23 percent (Figure 1).

Table 2 provides more information on how co-holding households compare to other households. The table classifies all households as borrowers, co-holders, savers, or in a 'neutral' category similarly to Gathergood and Weber (2014) and Gorbachev and Luengo-Prado (2016). 'Borrower' households have less than 1,500 EUR of deposits and more than 500 EUR of unsecured debt. 'Co-holder' households have more than 1,500 EUR of deposits and more than 500 EUR of unsecured debt. 'Saver' households have more than 1,500 EUR of deposits and less than 500

EUR of unsecured debt. 'Neutral' households have less than 1,500 EUR of deposits and less than 500 EUR of unsecured debt.

Co-holding households are relatively large, with above average disposable income and expenditure levels. Co-holding households are also relatively likely to have a mortgage. Unsurprisingly, co-holding households are considerably better off than borrowers but have fewer assets than savers. Relative to savers, co-holding households tend to be also younger. My findings on the relative characteristics of co-holders are in line with, for instance, Gathergood and Weber (2014).

Although co-holders have distinctive features relative to other households, co-holding is not restricted to a particular subgroup. Table 3 documents that co-holding occurs in all age and income quintiles, even if co-holding is most common among middle-aged households with relatively high income. The positive relationship with income is to a certain effect driven by the puzzle criteria where I only consider co-holding above a certain buffer of liquid assets. Therefore, households with modest deposits, correlated with low income, are too 'poor' to satisfy the co-holding threshold of at least 1,500 EUR in deposits.

3 Persistence and costs of co-holding

This section contributes to the literature on the credit card debt puzzle by studying for the first time the persistence of co-holding at the monthly level. Studying the persistence of co-holding informs researchers about how reliably a cross-section indentifies co-holding households. Persistence of co-holding also impacts the cost of co-holding.

3.1 How persistent is co-holding?

Existing evidence on the persistence of co-holding is scant. The only hard evidence is from Gorbachev and Luengo-Prado (2016) who present transition probabilities between waves 2004, 2008 and 2012 of the National Longitudinal Survey of Youth (NLSY). According to their 'strict' puzzle criteria around 40 percent of co-holding households continue to co-hold in two subsequent waves. In addition to Gorbachev and Luengo-Prado (2016), the SCF includes a survey question on whether households typically pay off their credit card balance. The high share of respondents that typically do not pay off their balance has been used as evidence that co-holding tends to be persistent. Yet, this conclusion does not necessarily follow from the survey answers, because even households who do not regularly pay off their balance may cease to co-hold if they reduce liquid assets.

I first study whether a cross-section gives an accurate picture of households who tend to co-hold. Table 2 shows that information on co-holding in January 2015 neatly identifies households who tend to co-hold for a significant amount of time in 2015. A household who co-holds in January 2015 co-holds on average in about nine months during 2015. By contrast, households classified as 'savers', 'borrowers' or 'neutrals' co-hold on average in less than two months in 2015. Therefore, the cross-section provides a fairly accurate picture of which households tend to co-hold.

Although the cross-section neatly distinguishes co-holding between the groups overall, many households identified as co-holders in the cross-section do not co-hold consistently over the rest of the year. Figure 2 panel A depicts the survival share of households that satisfy the puzzle conditions in a given baseline month in the subsequent months. To measure an average survival share, panel B plots the results for all households satisfying puzzle conditions at some point during January 2014 to December 2016. The survival share here includes households who do not

meet the puzzle criteria in one month but do so again later. The average survival share is 77 percent after 3 months, 73 percent after 6 months and 69 percent after 12 months.

A different way to measure persistence is to calculate the distribution of co-holding months over the next year conditional on satisfying the puzzle criteria in a given month. Figure 3 provides these results for the households who co-hold in January 2015. The mean number of co-holding months is 9.1 and the median 11. About 44 percent of households that co-hold in January 2015 have over 500 EUR of unsecured debt and 1,500 EUR of deposits in all twelve months of 2015. By contrast, about 25 percent of households that co-hold in January 2015 co-hold for six months or less in 2015.

The persistence results so far have used a binary measure of co-holding with the somewhat arbitrary criteria of at least 500 EUR of unsecured debt and 1,500 EUR of deposits. Therefore, I look separately at the persistence of co-holding at the intensive margin. First, Table 2 shows that for households defined as co-holders in January the mean amount of co-holding is 2,302 EUR in January 2015 versus 1,987 EUR (or 86 percent of the baseline amount) over the rest of 2015.

Figure 4 provides a fuller picture of the persistence of co-holding at the intensive margin. The figure plots the average persistence of co-holding using a continuous EUR measure of co-holding for all households that safisfy the puzzle criteria at any point during January 2014 to December 2016. In addition, the figure also depicts the differential persistence by initial level of co-holding (initial co-holding below or above the median). The results show that whereas small amounts of co-holding are persistent, relatively large initial amounts of co-holding tend to decline over time.

3.2 Effect of persistence on costs of co-holding

This section calculates how persistence affects the estimated costs of co-holding. I calculate yearly costs for the baseline group satisfying the puzzle conditions in January 2015 in two scenarios: (i) if yearly costs are proportional to the costs in January 2015; (ii) by considering the actual co-holding during the rest of 2015.

I calculate the costs of co-holding as the foregone savings in interest costs from using liquid assets to pay down unsecured debt following Zinman (2007). The interest rate on unsecured debt differs somewhat by product (for instance, type of credit card), but overall unsecured interest rates are homogeneous across households in Finland (see Table 2). In addition to the nominal interest rate spread betwen unsecured debt and deposits, I include avoidable invoicing fees in the cost of co-holding. I assume that a household can avoid invoicing fees if the household has enough deposits to pay off all unsecured debt.

$$\text{costs} = \min(\text{deposits}, \text{unsecuredDebt}) * (r_{\text{unsecuredDebt}} - r_{\text{deposits}}) + \theta \tag{1}$$

, where θ refers to the avoidable invoicing fees.

Overall, the yearly costs are fairly low (Table 4). The median cost is 162 EUR if I extrapolate from January 201, and the actual cost for the median household is 144 EUR. At the 90th percentile these costs are 345 EUR (extrapolated) and 315 EUR (actual). Although avoidable yearly costs are fairly low, the total costs of unsecured debt can cumulate significantly over time as shown by Gorbachev and Luengo-Prado (2016).

Two caveats pertain to the level of costs. First, I do not observe assets or liabilities at other banks. Although the sample restrictions aim to identify households who have the bank for which

I have data as their main bank, my results are still a lower bound for the total costs. Second, account fees can form a substantial fraction of the true costs of unsecured debt, especially for fairly small balances. However, I do not assume that co-holding households would close their unsecured credit accounts to avoid these costs. Therefore, I do not calculate the avoidable costs based on the total annual percentage rate (APR) of unsecured debt, and instead only consider avoidable invoicing fees and the nominal interest rate wedge.

Although the exact level of avoidable costs is uncertain, the relative difference depending on whether I consider extrapolated or actual avoidable costs is likely to be less affected by the concerns above[3]. The extrapolation method overestimates the costs of co-holding over 2015 for roughly 70 percent of households that co-hold in January 2015[4]. That said, the overestimation of the extrapolation method is on average relatively small because actual costs are about 90 percent of extrapolated costs. Yet, the spread in the ratio of actual to extrapolated costs across households is significant (45 percent in the 10th percentile versus 130 percent in the 90th percentile). Therefore, extrapolating costs of co-holding gives an inaccurate picture of the actual costs for many households.

4 Liquidity-based hypotheses for coholding

Prior literature has emphasised how the differential liquidity of deposits versus unused credit can explain co-holding. In particular, if credit cards cannot pay for certain expenses, for instance rent, some degree of co-holding is necessary and reflects a *liquidity premium* (Telyukova, 2013;

[3]Relative costs are not sensitive to unobserved information if unobserved unsecured debt and liquid assets are proportional to the activity at the bank for which I have data.

[4]Conversely, the extrapolation method underestimates the costs of co-holding for about 30 percent of households classified as co-holders in January 2015 who increase co-holding relative to January 2015. In addition, the extrapolation method underestimates the costs for households who are not classified as co-holders in January 2015 but co-hold over the rest of 2015. Yet, Table 2 showed that on average households not classified as co-holders in January 2015 do not co-hold much over the rest of the year.

Zinman, 2007). A second liquidity-based candidate for co-holding is the *precautionary borrowing hypothesis* (Druedahl and Jørgensen, 2018; Fulford, 2015). This hypothesis explains co-holding as a response to negative credit limit risk. Precautionary borrowing can make sense, because in bad times banks can cut credit limits but banks cannot demand immediate repayment of existing debt.

This section argues that liquidity-based hypotheses are insufficient to explain co-holding. First, even households with high liquidity co-hold, which poses a challenge to the liquidity premium hypothesis. Secondly, negative credit limit risk in Finland is low, which limits the explanatory potential of the precautionary borrowing hypothesis in my context. Finally, both liquidity-based hypotheses would predict strong demand for lower-cost additional liquidity among co-holders. I use a natural experiment that provided a subset of co-holders access to new liquidity at a considerably lower interest rate than the unsecured debt rate. Yet, only about 30 percent of co-holders take up the option of new low-cost liquidity. In addition, I find no evidence that the liquidity offer decreases co-holding.

4.1 Liquidity distribution among co-holders

Table 5 presents the distribution of liquid assets associated with co-holding in January 2015. The median co-holding household has over 4,500 EUR of deposits at the end of the month. Because households with higher income need more liquidity, Table 5 also tabulates the distribution of the deposits-to-income ratio associated with co-holding. Almost 40 percent of co-holding households have deposits worth more than twice the monthly disposable income. Although the liquidity premium hypothesis may explain why households with few deposits co-hold, co-holding among households with considerably liquidity is a challenge for the liquidity premium hypothesis. My

results are in line with earlier findings on the often high liquidity among co-holders (see, for instance, Gross and Souleles, 2002). In the portfolio choice model of Telyukova (2013), the liquidity premium channel explains well liquidity balances of the median co-holder. By constrast, model-predicted mean liquidity is lower than in the data, because the model struggles to match the upper tail of liquidity holdings among co-holders.

4.2 Negative credit limit risk in Finland

This section calculates two types of negative credit risk in my data to gauge whether the precautionary borrowing hypothesis could explain co-holding in Finland. The first type of negative credit risk is that the bank decreases the credit limit. Table 6 documents the risk of a credit limit decrease during 2015–2017. The results show that only 0.2 to 0.3 percent of credit accounts observe a decrease in the limit in a given year conditional on the account not being closed. What is more, some limit decreases were probably voluntary, although I cannot distinguish in my data whether the bank or the customer iniatiates the credit limit decrease. Even if all limit decreases are involuntary, though, the risk of a limit decrease is low compared to the 12 percent of credit accounts that experience a decrease in the US *per quarter* (Fulford, 2015)[5].

The second type of negative credit risk is that the bank closes the credit account. Table 6 documents the probability of account closure over 2015–2017 and distinguishes between closures initiated by the customer and the bank. The yearly risk of a bank closure of the credit line is only 0.6–0.7 percent; again only a fraction of the 5.4 percent reported by Fulford (2015) in the US[6]. Yet, co-holding occurs also in Finland in an environment with low negative credit risk.

[5]Fulford (2015) reports that "In any given quarter, two thirds of accounts had no change in credit limit. ... 20 percent of accounts in any given quarter had an increase in limit, while 12 percent had a decrease.".

[6]Fulford (2015) reports that "In the Consumer Finance Monthly (...) 5.4 percent (of consumers) report their

4.3 Low-cost liquidity offer as test of liquidity-based hypotheses

This section provides further evidence that liquidity-based hypotheses are insufficient to explain co-holding, because few co-holders take up an option of additional low-cost liquidity.

The liquidity-based theories suggest that the benefits of liquid assets outweigh the interest costs of unsecured debt for co-holders. Because co-holders are willing to pay the unsecured interest rate to maintain liquidity, they should value additional liquidity at a lower interest, either to further increase liquidity or decrease the interest costs of their liquidity.

I use a natural experiment to test demand for additional low-cost liquidity. In February 2015, the bank made a surprise offer to all of its mortgage customers. Mortgage customers received a free option to reduce the minimum mortgage principal payment to zero for up to 12 months without fees or conditions. For instance, a household paying 500 EUR/month in principal could obtain an additional 6 000 EUR of liquidity by taking up this offer. If the household took up the offer, the mortgage maturity would be extended by 12 months. Hence, the liquidity offer provided co-holders with a mortgage a long-term source of liquidity.

Crucially, the experiment provided households liquidity at a lower cost than the unsecured debt interest rate. Among co-holders eligible for the offer, the median nominal interest rate on unsecured debt was 7.1 percent, versus 1.2 percent on the mortgage. Mortgage and unsecured debt have similar interest rate risk because both unsecured debts and mortgages in Finland are almost always adjustable-rate loans[7]. Therefore, a lower interest rate risk does not compensate the higher interest rate margin on unsecured debt.

bank closed an account (during the previous year).".

[7]Total interest rate on mortgages and unsecured debt is the interest rate margin + a Euribor reference rate. If anything, mortgage interest rate risk is probably somewhat lower because unsecured debt rates are typically tied to the more volatile 3-month Euribor rate, whereas mortgages are often tied to 6 or 12-month Euribor rates.

4.3.1 Demand for low-cost liquidity by co-holders

Table 7 presents the take-up rate of the new liquidity offer by different subgroups of the January 2015 puzzle population holding a mortgage at the bank. Out of the 72,172 co-holders, 51,964 have a mortgage at the bank.

Overall, the take-up rate is around 30 percent. Predictably, take up increases with unsecured debt but is only 37 percent even in the top quartile. I do not find a connection between take up and the total months a household co-holds in 2015. Therefore, the low take up is not a result of most households stopping co-holding quickly and having no need for additional liquidity.

The low take up of the liquidity offer by co-holders seems at odds with the liquidity preference and precautionary borrowing hypotheses. Both theories argue that households value liquidity sufficiently to borrow at the unsecured debt interest rate. Therefore, co-holders should value additional low-cost liquidity, if only to repay high-cost debt whilst keeping equal or greater total liquidity[8].

A potential caveat to the low take up would be that households were inattentive to the offer. Yet, the liquidity offer was salient and the bank marketed it widely. The offer was the headline national news on the day of the launch, and media interest in the offer persisted during the five month application window. For instance, the main national newspaper, Helsingin Sanomat, had the liquidity offer as their lead article the day after the offer launch (Helsingin Sanomat, February 6th 2015 edition). The term for the liqudity offer ("lyhennysvapaa") returns 26 articles published during the application window in Helsingin Sanomat's digital archive[9]. Visibility was not restricted to print media because the offer was widely covered also by the national

[8]Low take up cannot be explained by the offer forcing "too much" liquidity on the households, because the overwhelming majority of mortgages are adjustable-rate and have no restrictions on extra principal payments.

[9]https://www.hs.fi/haku/. The search was made on the 20th of March 2019.

broadcasting company Yle (equivalent to BBC).

To further address the worry that despite the major advertising campaign some households may still have been inattentive, I focus on the behavior of households that almost certainly knew about the liquidity offer. I identify eligible co-holding households in which at least one household member was an *employee* of the bank that provided the liquidity offer (N = 2,099). These households arguably know about the offer their employer was promoting. Yet, only 45 % of co-holding bank-employee households with a mortgage take up the liquidity offer[10].

4.3.2 Difference-in-difference analysis of effects of liquidity offer

In addition to overall take up, the liquidity offer allows me to test if access to new liquidity had an effect on co-holding. If liqudity needs drive co-holding, co-holding should decrease if households had sufficient liquidity. Therefore, I use an intention-to-treat framework to compare the evolution co-holding by household eligibility for the liquidity offer.

The identification strategy is to compare co-holding households *with* a mortage at the bank (=treated) to co-holding households *without* a mortgage at the bank (=control). A caveat is that some competitor banks also made similar offers to their mortgage customers in 2015. Therefore, some households in my control group may have had a mortgage at one of the competitor banks and would have been also 'treated'. However, as detailed in the sample selection section, I focus on households that used the bank that I have data for as their main bank. Therefore, it is unlikely that a substantial portion of control households were eligible for a liquidity offer from another bank.

First, Figure 5 plots the overall co-holding rate by mortgage status before and after the liquidity

[10]Employee households are about equally likely to co-hold than regular customers. For instance, in January 2015 the employee household co-holding rate is 24 % versus 20 % among all customers.

offer. As a reminder, the liquidity offer was introduced in February 2015 so it is exogenous from the perspective of decisions made up to January 2015. Although co-holding is more common among mortgage households, the trends in both groups are similar. In particular, we do not observe a discontinuity in co-holding among households eligible for the liquidity offer after January 2015.

The unconditional comparison in Figure 5 does not account for characteristics that differ systematically between mortgage and non-mortgage households and influence co-holding. Therefore, I create a matched sample of households where I use propensity score matching to find for each household with a mortgage a non-mortgage match that is similar in terms of observable characteristics. The household characteristics that I use for matching are the number of adults and children, deposits and unsecured debt in the baseline month of January 2015, average card expenditure over 2014 and the average monthly disposable income over 2015.

I study the relationship between liquidity offer eligibility and co-holding while accounting for differences in household characteristics using a a difference-in-difference regression:

$$\begin{aligned}\mathbb{1}_{\text{puzzle},h,m} =& \delta_h + \gamma_{\text{pre}}\mathbb{1}_{m<\text{Jan2015}} + \beta_{\text{pre}}\mathbb{1}_{m<\text{Jan2015}} * \mathbb{1}_{\text{mortgage}} + \\ & \gamma_{\text{post}}\mathbb{1}_{m>\text{Jan2015}} + \beta_{\text{post}}\mathbb{1}_{m>\text{Jan2015}} * \mathbb{1}_{\text{mortgage}} + \epsilon_{h,m}\end{aligned} \tag{2}$$

, where h refers to the household, m to the month and δ_h refers to household fixed effects. The regression tests whether co-holding decreases disproportionately after the offer for households eligible for additional liquidity ($\beta_{\text{post}} < 0$). The regression also tests for equality of pre-trends ($\beta_{\text{pre}} = 0$) as a validity check of the difference-in-difference approach. The baseline month to which I compare the change in co-holding is January 2015. I estimate the regression with a panel

that spans from January 2014 to December 2016.

Table 8 presents the estimation results for co-holding by liquidity offer eligibility where I measure co-holding by a dummy variable. The first two columns present results for specifications where I use the full sample of households without matching. Columns 3–4 provide estimation results with the matched sample. No specification finds that households eligible for the liquidity offer reduced their co-holding propensity relative to non-eligible households after the liquidity offer. The results from the full sample specification with household fixed effects (column 2) create some concern about the validity of the constant pre-trends assumption because $\hat{\beta}_{\text{pre}}$ is significantly different from 0 although quantitatively small. Yet, the common pre-trends assumption holds in the matched sample specification with household fixed effects (column 4), and the results imply that households eligible for the liquidity offer do not decrease co-holding relative to non-eligible households after the offer.

Even if the liquidity offer did not affect the probability to co-hold at the extensive margin, the liquidity offer could reduce co-holding at the intensive margin. Therefore, I estimate a second difference-in-difference regression where I change the dependent variable from a binary co-holding dummy to a continuous EUR variable. Results for this specification are presented in Table 9. Again, I find no evidence for mortgagors reducing co-holding after the liquidity offer.

In summary, (i) only a minority of eligible co-holding households take up the offer of additional low-cost liquidity, (ii) I find no evidence that co-holding decreases for eligible households after the liquidity offer. Consequently, the results of this section signal that liquidity-based hypotheses are unlikely to sufficiently explain the credit card debt puzzle. That said, my results do not imply that liquidity-based explanations do not contribute to co-holding. These mechanisms can matter particularly for households with limited liquidity and in countries with high negative

credit limit risk (Druedahl and Jørgensen, 2018; Fulford, 2015; Gorbachev and Luengo-Prado, 2016).

5 Intra-household financial pooling and co-holding

The prevailing method of defining the puzzle at the household level implicitly assumes perfect pooling of household finances. This section studies whether the intra-household distribution of assets and liabilities matters. I find that co-holding is considerably less likely at the individual than at the household level in households with multiple adults. Furthermore, the persistence of co-holding is lower at the individual level. The intra-household distribution of assets and liabilities also influences the take up of new lower-cost liquidity. Therefore, part of co-holding reflects imperfect intra-household financial pooling. These results are consistent with the (multi-person) accountant-shopper model à la Bertaut et al. (2009).

5.1 Co-holding at individual vs. household level

This section documents that the prevalence and persistence of co-holding is lower at the individual than at the household level in multi-adult households. These findings are consistent with the hypothesis that imperfect intra-household financial pooling sustains co-holding at the household level.

I start this section by comparing the prevalence of co-holding at the household and individual level in households with multiple adults. The purpose is to study in how large a proportion of households classified as co-holders the puzzle criteria are satisfied only if I pool household assets and liabilities. Conversely, the results reveal in how many co-holding households at least one

member co-holds a significant amount individually.

I find that the co-holding rate is on average 11.9 *percentage points* lower at the individual than at the household level (Figure 6). In relative terms, the co-holding rate is 42 *percent* lower. Therefore, in almost half of multi-adult households classified as co-holders the puzzle criteria are only met if I pool household assets and liabilities. Figure A.1 complements the results on co-holding rates by presenting the EUR amount of co-holding at the household level and the total EUR amount of co-holding by individuals in those households. On average, total individual co-holding accounts for 69 percent of total household level co-holding in multi-adult households. Consequently, if we relax the assumption of intra-household financial pooling, a large share of co-holding in multi-adult households is no longer a puzzle[11].

Next, I study whether the persistence of co-holding differs at the household and individual level. Figure 7 panel A depicts the average persistence of co-holding conditional on co-holding at some point in January 2014 through December 2015. The survival rate for *households* after 12 months is approximately 75 percent whereas it is barely over 50 percent for *individuals*. Another way to study the difference in persistence is to compare the distribution of total co-holding months in 2015 conditional on co-holding in January 2015 (panel B). A household (individual) satisfies the puzzle conditions on average in 9.5 (7.4) months over 2015.

The impact of intra-household distribution of assets and liabilities on co-holding has not been documented previously, because earlier studies have relied on data that aggregates assets and liabilities at the household level. That said, Choi and Laschever (2018) also consider intra-household dynamics and co-holding. They find that the "Big Five" psychological characteristics of both members of the couple impact the probability of co-holding. In addition, they document

[11]Though I focus here on the intra-household behavior of multi-adult households, the co-holding rate is also lower in single-adult (11 percent) than multi-adult (28 percent) households in January 2015.

that the income inequality and the level of agreeableness within the couple impact the likelihood of co-holding. My quantitative breakdown of co-holding at the individual versus household level provides further evidence in favor of intra-household dynamics impacting co-holding.

5.2 Intra-household financial pooling and take up of liquidity offer

This section documents that the intra-household distribution of assets and liabilities impacts the take up of the low-cost liquidity offer in households where individuals do not have to reach consensus over the decision to take up the liquidity offer. This finding is consistent with households not pooling their financial resources.

A final test of intra-household financial pooling is whether the distribution of assets and liabilities impacts the household decision to take up the liquidity offer. I compare the take-up rate between households with similar amounts of total deposits and unsecured debt but varying distribution across household members. If the household pools its resources, only total assets and liabilities, but not their intra-household distribution, should influence take up of additional liquidity. Conversely, if the household does not pool its resources, the saver may prevent the take up of the liquidity offer to preclude the spender from overconsuming.

The aim is to test whether the personal financial situation of one of the partners affects disproportionately the take up of the liquidity offer. I operationalise the test as follows. For each individual mortgage holder belonging to a multi-adult household, I derive information on both individual characteristics of the mortgage holder, individual characteristics of the partner, and common household characteristics (for instance, number of children). Then I estimate a linear probability model where the dependent variable takes the value 1 if the mortgage holder takes up the liquidity offer on at least one of her mortgages and 0 otherwise.

To study how the distribution of assets and liabilities within the household affects the take up of the liquidity offer, I need to define the primary and secondary partner in the household. With individual mortgages, it is straightforward to regress liquidity offer take up on individual financial variables of both the mortgage holder (primary partner) and her co-habitant (secondary partner). The definition of the primary and secondary partner is somewhat fuzzier with joint mortgages for which both are liable. I follow the bank classification, whereby the bank registers one of the individuals as the primary mortgage holder. In the case of joint mortgages, the primary mortgage holder is more likely to be male (69 percent of cases), has higher deposits (median 1149 EUR versus 910 EUR for the partner) and more unsecured debt (median 997 EUR versus 698 EUR for the partner)[12].

To take up the liquidity offer on a joint mortgage, both debtors need to agree to it. On the other hand, for decisions regarding individual mortgages the partners do not need to consult each other. Therefore, my hypothesis is that the distribution of assets and liabilities would impact less the take up of the liquidity offer in households with a joint mortgage than in households with individual mortgages.

The main regression specification is:

$$\begin{aligned}\mathbb{1}_{\text{liquidity},p,h} = &\sum_{b=1}^{10} \gamma_b \text{totalUnsecuredDebtBin}_{\text{h,b}} + \sum_{b=1}^{10} \theta_b \text{totalDepositsBin}_{\text{h,b}} + \\ &\beta_1 \text{unsecuredDebtShare}_{\text{p,h}} + \beta_2 \text{depositShare}_{\text{p,h}} + \\ &\Omega \mathbb{X}_h + \epsilon_{p,h}\end{aligned} \tag{3}$$

The subscript p refers to a given partner with a mortgage in household h. The key explanatory

[12]These figures pertain to the sample of multi-adult households with positive deposit and unsecured debt in January 2015.

variables are *depositShare* and *unsecuredDebtShare*, which measure the share of total household deposits and unsecured debt held by the (primary) mortgage holder. Under financial pooling, the distribution of assets and liabilities does not affect the take up of the liquidity offer, and hence $\beta_1 = \beta_2 = 0$. On the other hand, if the (primary) mortgage holder has more weight in the household decision, then $\beta_1 > 0$ and $\beta_2 < 0$. That is, under imperfect financial pooling, the take-up rate is higher if the primary mortgage holder accounts for a larger share of total household unsecured debt or has a smaller share of total household deposits. All regressions control for total household deposits and unsecured debt as well as other household characteristics $\mathbb{X}_h$[13].

I reject the hypothesis of perfect intra-household financial pooling when estimating equation (3) using my baseline sample (Table 10, column 1), which includes all two-adult households with positive deposits and unsecured debt in January 2015. Yet, quantitively the effect on liquidity offer take up is moderate. For instance, if the partner accounts for all unsecured debt in the household, the take-up rate is 2.8 percentage points lower (or 10 percent in relative terms) than if the primary mortgage holder accounts for all unsecured debt. Results are similar if I concentrate on two-adult households that satisfied the co-holding criteria in January 2015 (column 2).

Households with individual mortgages drive the overall rejection of intra-household financial pooling. Among households with only individual mortgages, the take-up rate is 7 percentage points lower (or 30 percent in relative terms) if the partner accounts for all unsecured debt than if the primary mortgage holder accounts for all unsecured debt (column 4). The effect of the deposit distribution is of a similar size. On the other hand, coefficient estimates for households with only joint mortgages (column 3) are statistically insignificant[14].

[13]Number of adults and children, average age of adults, municipality, mortgage value, size of principal payment, mortgage interest rate, property value, disposable income level, card purchase level.

[14]A small minority of households have both joint and individual mortgages. I exclude these households from models 3 and 4. Therefore, the sum of observations in models 3 and 4 is smaller than in model 1.

The differential effects of the intra-household distribution of assets and liabilities on liquidity offer take up by mortgage contract type probably reflect multiple factors. First, the consensus requirement to take up the liquidity offer on joint mortgages can promote co-operation. Second, couples can self select into joint or individual mortgages depending on the desired level of financial pooling. Therefore, I am not able to say to what extent the null result for households with joint mortgages reflect the consensus requirement versus a mutual desire for financial pooling.

6 Anchoring and co-holding

Until now, I have focused on strategic motives for co-holding such as liquidity concerns or (lack of) intra-household financial pooling. This section provides evidence consistent with anchoring, a behavioral mechanism, influencing co-holding.

The anchoring hypothesis posits that a salient cue can anchor household choice even if the cue should be irrelevant to the decision problem (Tversky and Kahneman, 1974). With regards to the credit card debt puzzle, the salient cue is the contractual minimum payment. If the contractual minimum payment anchors regular debt payments at or near the contractual minimum, co-holding will arise almost mechanically as long as households use credit cards.

Various studies have shown that the prevalence of low credit card payments is consistent with anchoring (Guttman-Kenney et al., 2018; Hershfield and Roese, 2015; Keys and Wang, 2016; Navarro-Martinez et al., 2011; Stewart, 2009)[15]. The paper closest to mine, Keys and Wang (2016), estimates that 29 percent of credit card accounts in the US regularly make only minimum or near-minimum payments. The authors use a clever identification mechanism based on changes

[15]Although my paper considers anchoring to the minimum credit card payment, households can also follow other heuristics. Gathergood et al. (2019) show that the repayment behavior of households with multiple credit cards is consistent with balance matching.

in minimum payment rules to argue that the behavior of near-minimum payers reflects anchoring instead of liquidity constraints. If the contractual minimum payment rises, households that were originally making smaller payments than the new minimum because of liquidity constraints should bunch at the new higher minimum payment. Yet, the authors show that, consistent with anchoring, a significant fraction of households round up their regular payments both above the original (lower) contractual minimum and the new (higher) contractual minimum.

Keys and Wang (2016) argue that although anchoring is the likely explanation for low payments, it is difficult to distinguish anchoring and "near-rational" rounding: "Some consumers may both be liquidity constrained and have an intrinsic preference for round numbers, either for budgeting or aesthetic purposes. If liquidity-constrained consumers round up from the minimum, then their responsiveness to changes in the minimum would be observationally similar to what we interpret as anchoring, but would have different welfare implications." The authors find anchoring a more likely explanation, though, because near-minimum payments are common across the age and income distribution that they use to proxy liquidity.

I provide additional evidence to distinguish anchoring from liquidity constraints as the driver of low credit card payments. Crucially, I can observe liquidity directly, which allows me to show that low payments often coincide with high liquidity. Finally, I bridge the literature on anchoring in credit card payments and the literature on the credit card debt puzzle by showing that anchoring households are more likely to co-hold.

I study anchoring among the subgroup of households with electronic debt invoicing for whom I observe both the minimum payment and the actual payment. Roughly 2/3 of households use electronic debt invoicing and 1/3 mail invoicing. For households with mail invoicing, I do not observe all the required information to study anchoring. I also omit from the analysis the 12

percent of electronic invoicing households with autopay, because anchoring concerns an active choice.

First, I characterize each individual payment observation relative to the minimum payment in 2015. I focus on the payment behavior of the individual with the highest credit card limit within the household. I classify an account-invoice observation as *minimum payment* if the payment equals the contractual minimum. The contractual minimum payment equals $\max[\tau * \text{creditLimit}, x]$, where τ measures the percentage of the total credit limit, and x is a fixed EUR amount[16]. I classify a payment as *near-minimum payment* if the payment is within 50 EUR of the minimum payment. I classify a payment as *maximum payment* if the payment equals the total invoice balance (sum of interest-free and interest-accruing credit card balance)[17].

After classifying each payment observation, I classify individuals according to their regular payment type in 2015 following Keys and Wang (2016). If an individual makes the minimum, near-minimum, or maximum payment in over 50 percent of account-invoice observations, I classify the individual as a *minimum, near-minimum or maximum payer*. If none of the three main payment types accounts for over 50 percent of account-invoice observations, I classify the individual as *mixed payer*.

Table 11 documents the share of different payer types. A bit less than a third of individuals are classified as minimum or near-minimum payers. By contrast, about 40 percent of individuals regularly pay the full invoice balance. The share of minimum and near-minimum payers is similar to Keys and Wang (2016), although my distribution tilts more toward exact minimum payers. One potential reason is that in my setting the minimum payment is a percentage of the

[16]The contractual minimum payment is a share of the credit limit, and not of the monthly balance, as for instance in the data of Keys and Wang (2016).

[17]I classify a payment as *maximum payment* if the payment equals the invoice balance and the invoice balance is at most the contractual minimum payment + 50 EUR to distinguish maximum payments from "overlapping" near-minimum payments.

credit limit, and not account balance, as in Keys and Wang (2016). Therefore, the minimum payment in my setting is already a round number, unlike in Keys and Wang (2016), and hence individuals in my data could feel less urge to round up the minimum payment.

Next, I study if liquidity constraints explain low payments by presenting the distribution of individual deposits associated with minimum or near-minimum payments in 2015 (Table 12). Because I only observe end-of-month deposits, I present two liquidity measures: 'pre-invoice deposits-to-income' uses deposits in the last day of the month preceding the invoice month; "post-invoice deposits-to-income" uses deposits in the last day of the invoice month. Panels A and B reveal that in over a third of low payment observations, individuals have deposits greater than one month's disposable income both before and after receiving the invoice[18]. Therefore, liquidity constraints are unlikely to sufficiently explain the propensity of low payments.

Yet, even if liquidity constraints do not solely explain the propensity of low payments, low payments do not necessarily reflect anchoring. Instead, low payments can reflect the differential attention or effort required to make low versus larger payments. The attention or effort hypothesis is relevant, because I study individuals with electronic invoicing. With electronic invoicing, individuals observe the default invoice amount, which often coincides with the contractual minimum payment, and can accept the suggested payment without manual changes and without opening the invoice pdf statement with information on the total invoice balance[19]. By contrast, to pay an amount greater than the default invoice amount, the individual needs to manually change the payment amount. Therefore, the propensity of low payments can reflect a default option effect because of limited attention or effort.

[18]I focus on individual instead of household deposits to avoid the assumption of intra-household financial pooling.

[19]The default invoice amount is by default the contractual minimum payment, but individuals can choose a higher default invoice amount. In principle, setting a near-minimum default invoice amount can reflect anchoring, but for households who have set a default invoice amount close to the contractual minimum, payment choices do not distinguish anchoring from limited attention or effort.

Therefore, to isolate anchoring I study the subset of *active* near-minimum payments when the default invoice amount equals the contractual minimum payment but the individual actively makes a larger payment by manually changing the suggested payment. Table 12, panel C documents that the deposit distribution is similar in the case of active near-minimum payments as in the case of minimum payments and all near-minimum payments. The propensity to actively make low credit payments despite high liquidity provides evidence consistent with anchoring in a setting where limited attention or effort are not confounders.

Finally, low regular payments on electronic credit card invoices are associated with higher co-holding (Table 13). Households that regularly make minimum or near-minimum payments on their credit card invoice co-hold more than households that regularly make larger electronic invoice payments. The association is not a mechanical consequence, because households that pay a low amount on their monthly invoice can still make additional manual payments from their deposit account to their credit card account. I do not observe in my data additional manual transfers, but results on co-holding reveal that manual transfers do not undo the positive correlation between low invoice payments and co-holding.

In summary, individuals often make low payments on their credit card debt despite high liquidity. The observation holds in active choices by individuals where limited attention or effort are not confounders. The behavior is consistent with anchoring where the minimum payment anchors actual payments, even though the individual would have sufficient liquidity to make higher payments. Finally, the propensity to make low credit card debt payments correlates positively with co-holding.

7 Conclusions

I have studied the credit card debt puzzle of why households co-hold high-cost debt and low-cost assets and do not economize on interest costs. I have used a high-frequency Finnish dataset on hundreds of thousands of households to both document new facts and test existing and new hypotheses for co-holding.

The first main finding of the paper is that a cross-section used in most of the existing literature does not give an accurate picture of the persistence of co-holding. Because of missing the imperfect persistence, extrapolating the costs of co-holding from a cross-section only seldomly matches the actual costs for any given household. Future research should study more what drives increases and decreases in co-holding to further understand the reasons for co-holding.

Second, liquidity-based theories are insufficient to explain co-holding, although my findings do not imply that liquidity concerns do not contribute to co-holding. But, liquidity-based theories struggle to explain why households with plenty of liquidity and low negative credit limit risk co-hold; or why most co-holding households do not take advantage of additional liquidity at considerably lower interest rates than the credit card rate.

Third, the assumption of perfect intra-household pooling of assets and liabilities significantly impacts the magnitude of the puzzle. However, it is debatable why pooling of finances should be the benchmark or why the lack of pooling would be *puzzling* (see, for instance, Chiappori and Mazzocco, 2017). My paper adds to the evidence that intra-household distribution of financial assets and liabilities matters.

Finally, my results suggest that the tendency to anchor credit card payments to the contractual minimum contributes to co-holding. The results are consistent with the burgeoning literature

on debt-repayment behavior that suggests that households often use heuristics or focus on particularly salient information instead of solving a complex optimization problem. Further research should study how to reduce the costs of anchoring without complicating the debt repayment choice for households with limited attention.

References

Bertaut, C.C., Haliassos, M., 2001. Debt Revolvers for Self Control.

Bertaut, C.C., Haliassos, M., Reiter, M., 2009. Credit Card Debt Puzzles and Debt Revolvers for Self Control. Review of Finance 13, 657–692. https://doi.org/10.1093/rof/rfn033

Chiappori, P.-A., Mazzocco, M., 2017. Static and Intertemporal Household Decisions. Journal of Economic Literature 55, 985–1045.

Choi, H.-s., Laschever, R.A., 2018. Credit Card Debt Puzzle and Noncognitive Ability. Review of Finance 22, 2109–2137.

Druedahl, J., Jørgensen, C.N., 2018. Precautionary Borrowing and the Credit Card Debt Puzzle. Quantitative Economics forthcoming.

Fulford, S.L., 2015. How important is variability in consumer credit limits? Journal of Monetary Economics 72, 42–63. https://doi.org/10.1016/j.jmoneco.2015.01.002

Gathergood, J., Mahoney, N., Stewart, N., Weber, J., 2019. How Do Individuals Repay Their Debt? The Balance-Matching Heuristic. American Economic Review 109, 844–875. https://doi.org/10.1257/aer.20180288

Gathergood, J., Weber, J., 2014. Self-control, financial literacy & the co-holding puzzle. Journal of Economic Behavior & Organization 107, 455–469. https://doi.org/10.1016/j.jebo.2014.04.018

Gorbachev, O., Luengo-Prado, M.J., 2016. The Credit Card Debt Puzzle: The Role of Preferences, Credit Risk, and Financial Literacy (SSRN Scholarly Paper No. ID 2826962). Social Science Research Network, Rochester, NY.

Gross, D.B., Souleles, N.S., 2002. Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data. The Quarterly Journal of Economics 117, 149–185.

Guttman-Kenney, B., Leary, J., Stewart, N., 2018. Weighing anchor on credit card debt. Financial Conduct Authority Occasional Paper 43.

Hershfield, H.E., Roese, N.J., 2015. Dual payoff scenario warnings on credit card statements elicit suboptimal payoff decisions. Journal of Consumer Psychology 25, 15–27. https://doi.org/10.1016/j.jcps.2014.06.005

Keys, B.J., Wang, J., 2016. Minimum Payments and Debt Paydown in Consumer Credit Cards. NBER Working Paper 22742.

Lehnert, A., Maki, D.M., 2002. Consumption, debt and portfolio choice: Testing the effect of

bankruptcy law (Finance and Economics Discussion Series No. 2002-14). Board of Governors of the Federal Reserve System (U.S.).

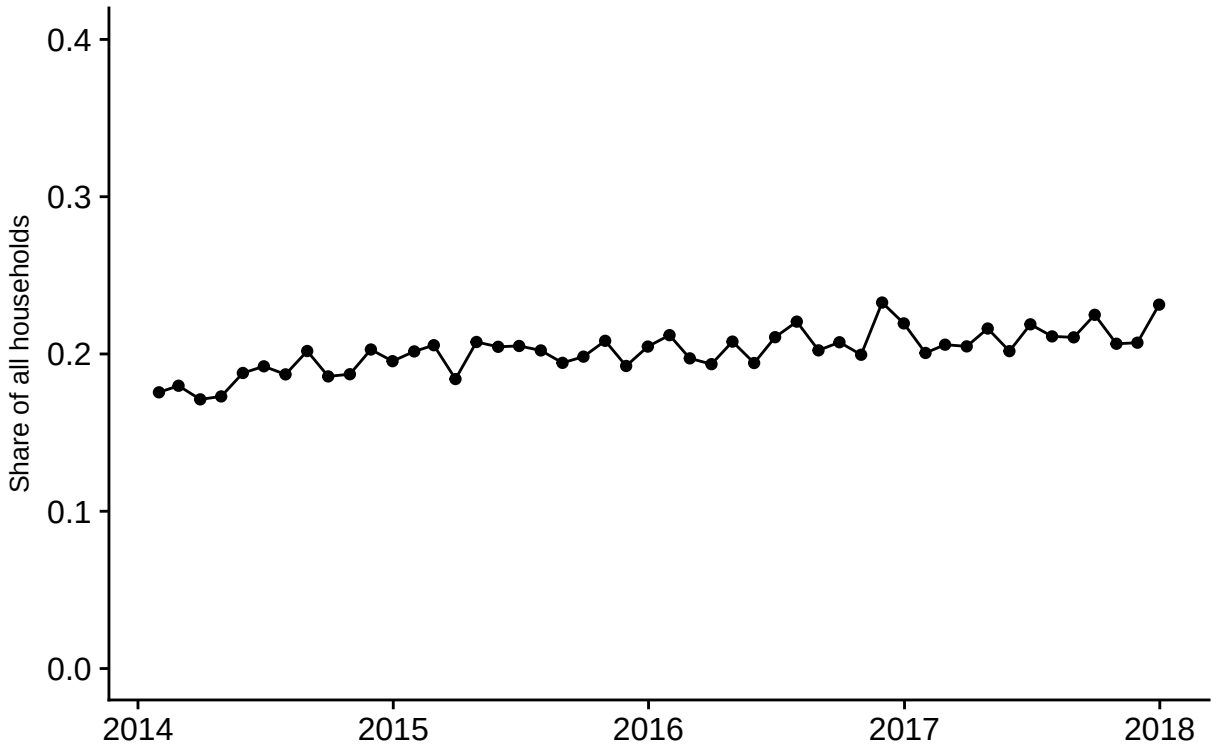
Navarro-Martinez, D., Salisbury, L.C., Lemon, K.N., Stewart, N., Matthews, W.J., Harris, A.J., 2011. Minimum Required Payment and Supplemental Information Disclosure Effects on Consumer Debt Repayment Decisions. Journal of Marketing Research 48, S60–S77. https://doi.org/10.1509/jmkr.48.SPL.S60

Stewart, N., 2009. The Cost of Anchoring on Credit-Card Minimum Repayments. Psychological Science 20, 39–41. https://doi.org/10.1111/j.1467-9280.2008.02255.x

Telyukova, I.A., 2013. Household Need for Liquidity and the Credit Card Debt Puzzle. The Review of Economic Studies 80, 1148–1177. https://doi.org/10.1093/restud/rdt001

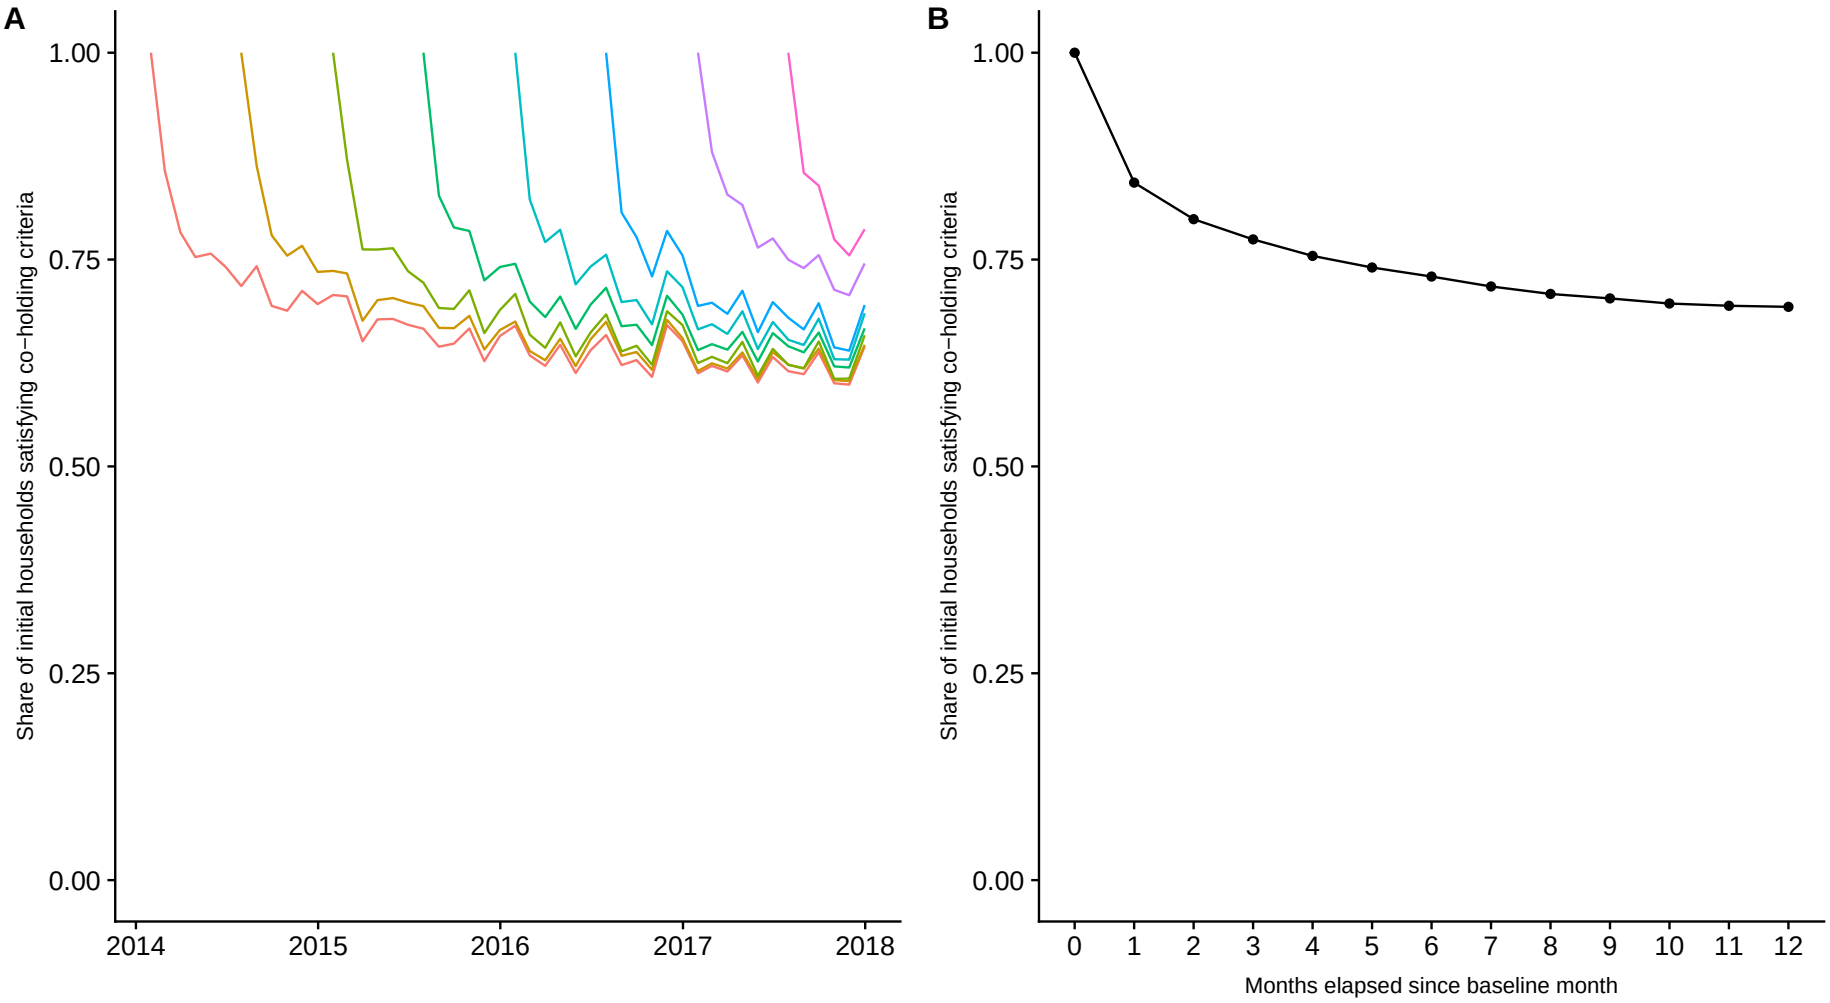
Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science 185, 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Zinman, J., 2007. Household Borrowing High and Lending Low Under No-Arbitrage. Mimeo.



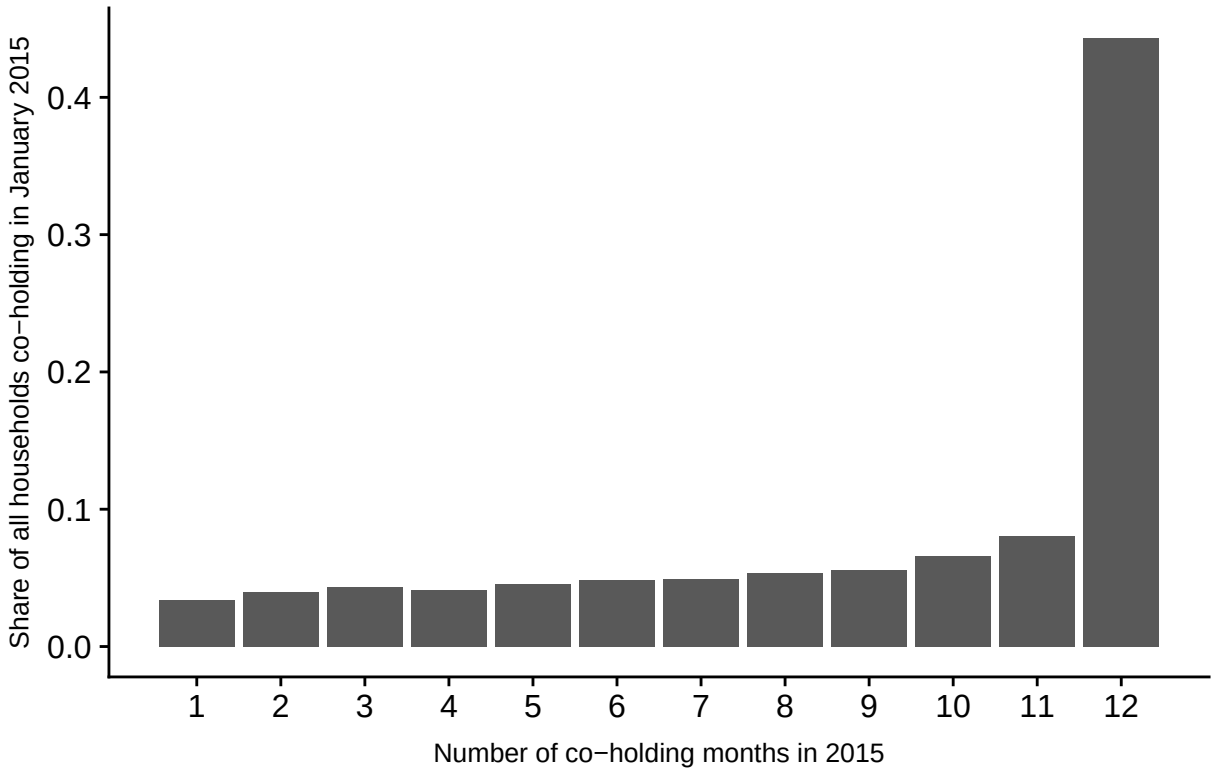
Note: The two co−holding criteria are: (i) interest−accruing unsecured debt above 500 EUR; (ii) deposits above 1,500 EUR.

Figure 1: Share of households satisfying co-holding criteria by month



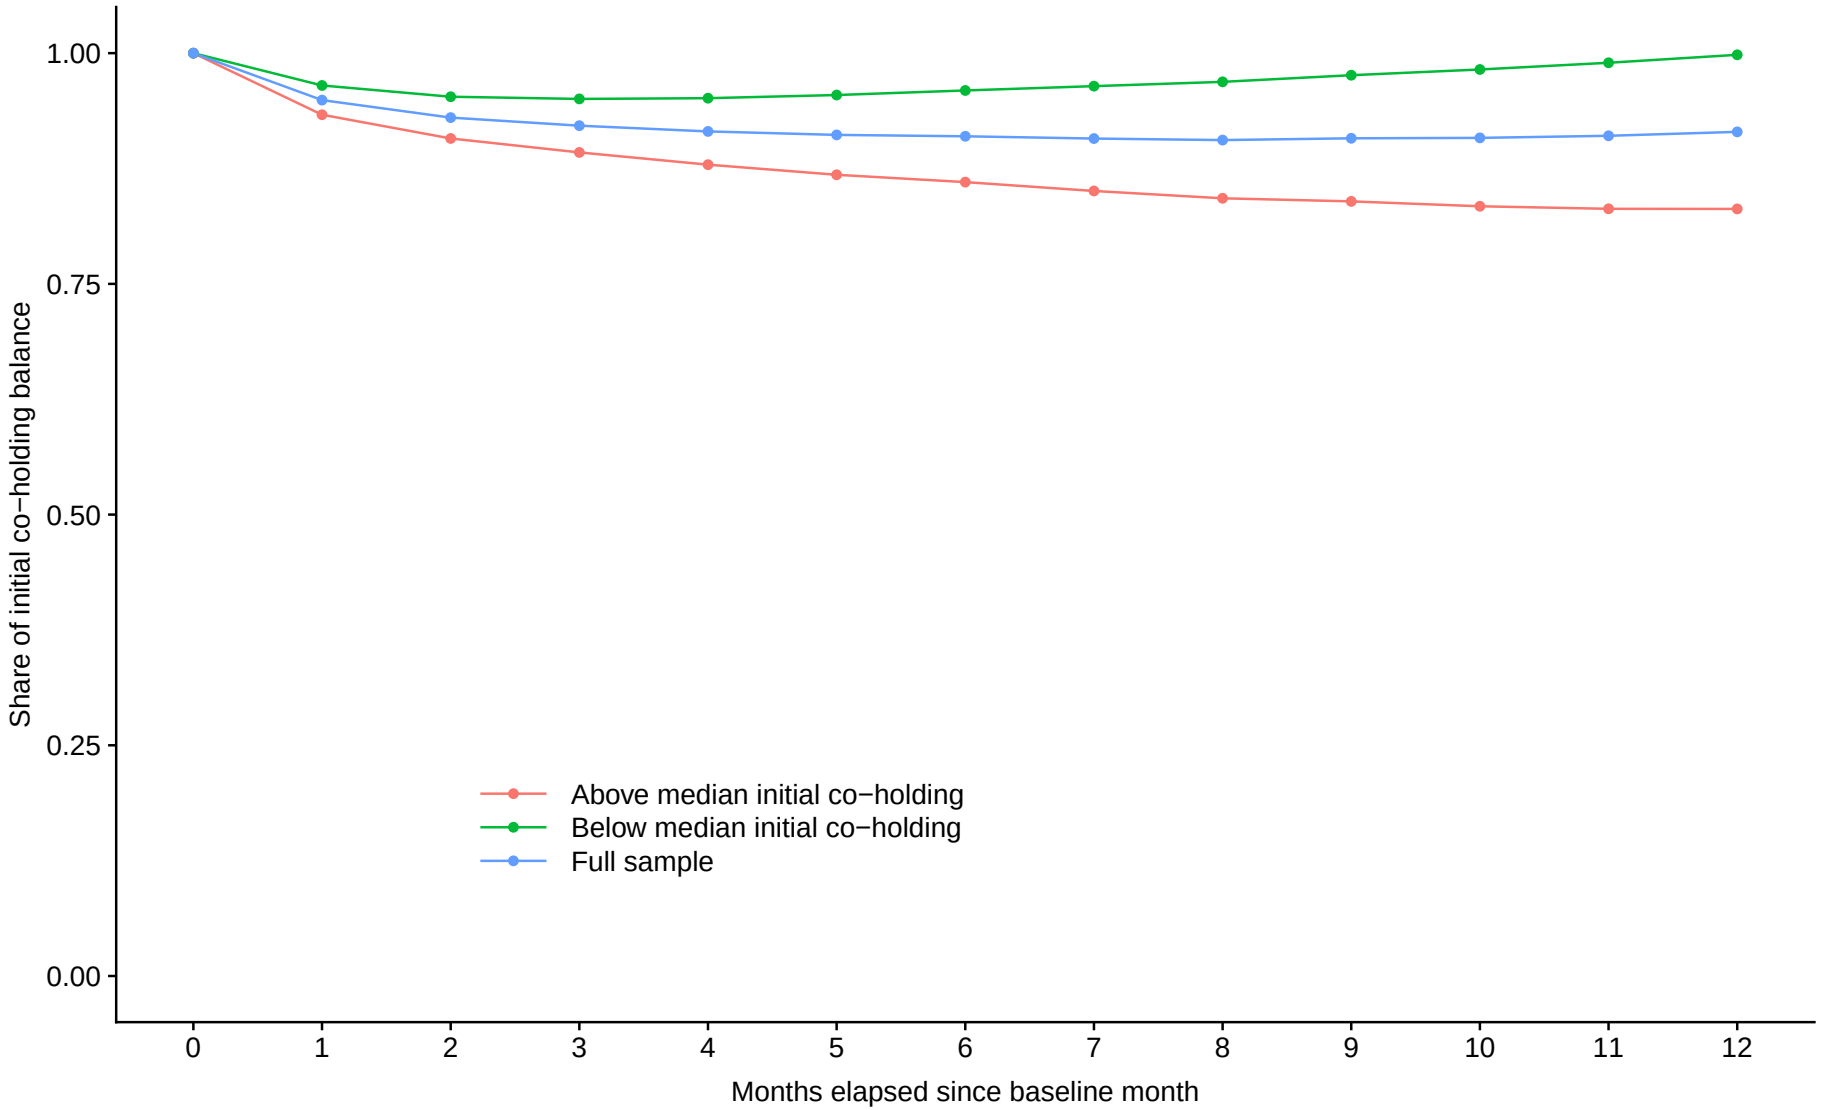
Note: Panel A presents the survival share of households that satisfy the co−holding criteria (unsecured debt over 500 EUR and deposits over 1,500 EUR) in the baseline months of January and July in 2014, 2015, 2016, and 2017. Panel B depicts the average survival share for households that satisfy the co−holding criteria at some point during January 2014 to December 2016.

Figure 2: Persistence of co-holding (binary co-holding measure)



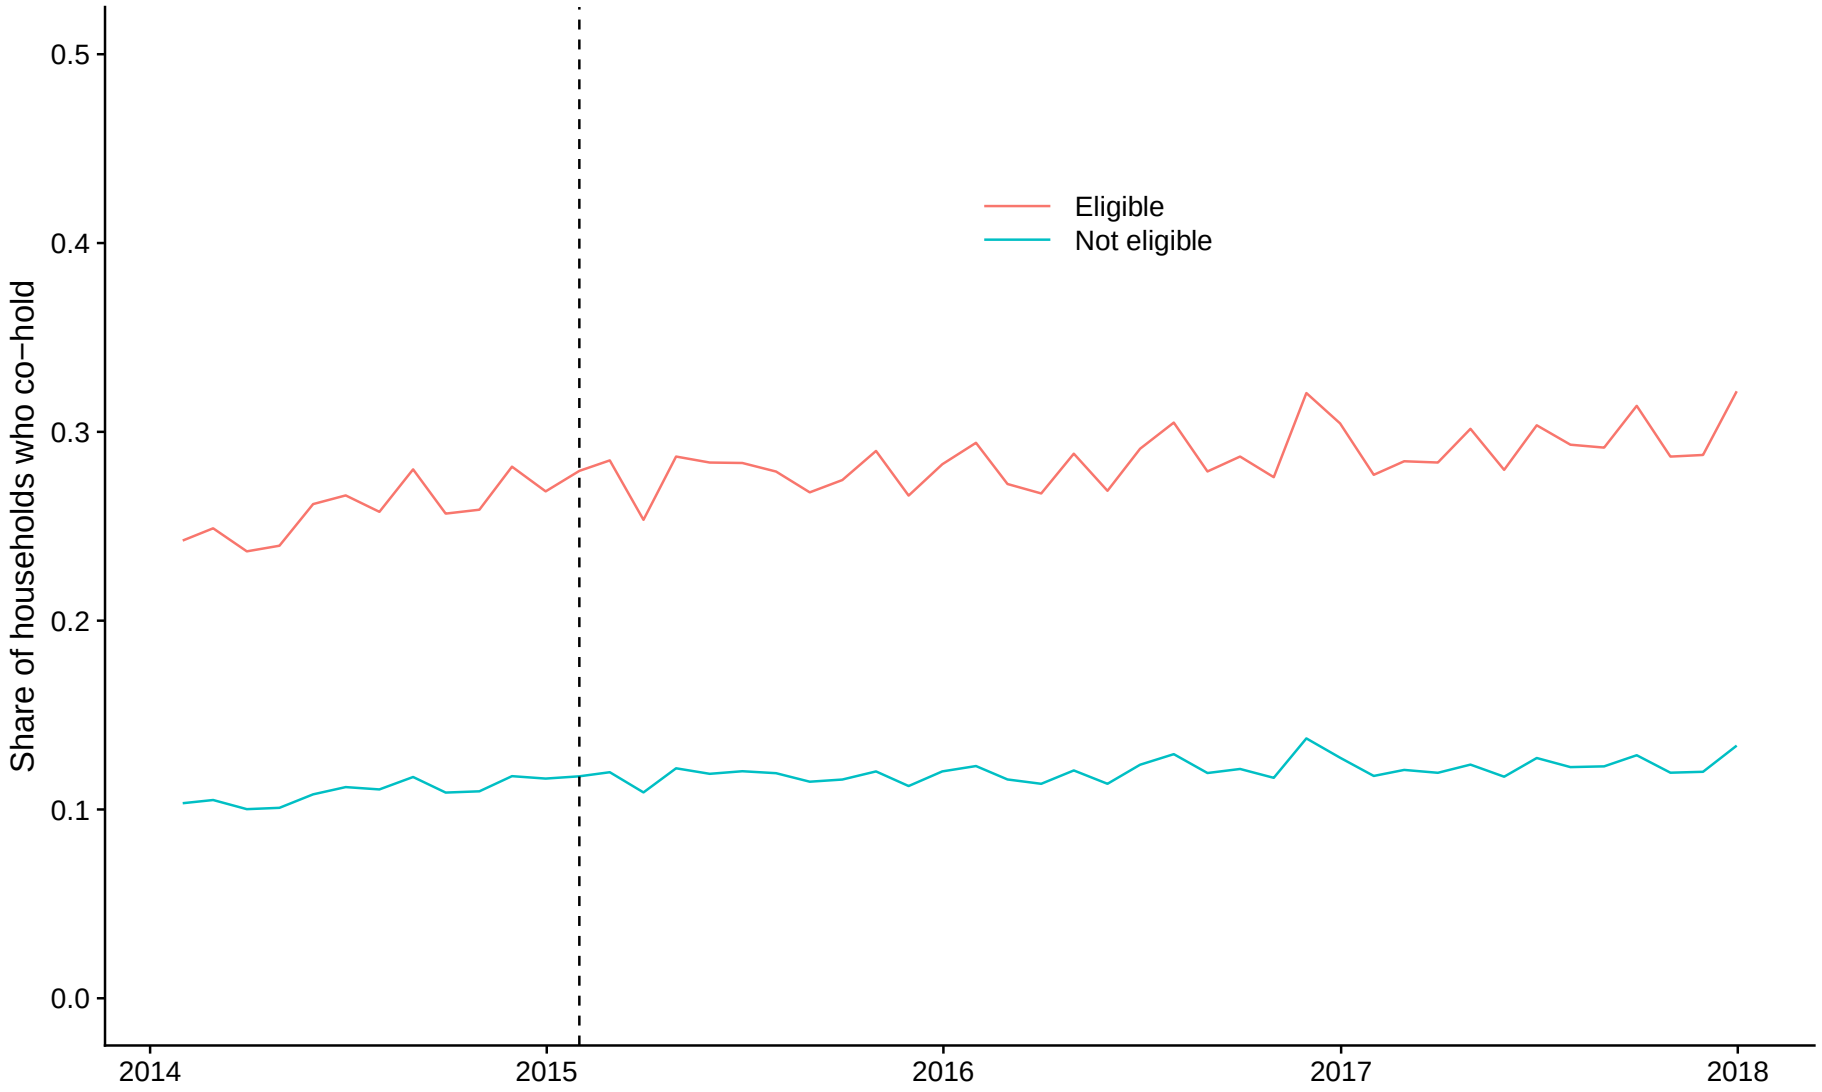
Note: The figure plots the distribution of co−holding months in 2015 among households who satisfy the co−holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) in January 2015.

Figure 3: Distribution of co-holding months in 2015 for January 2015 co-holders



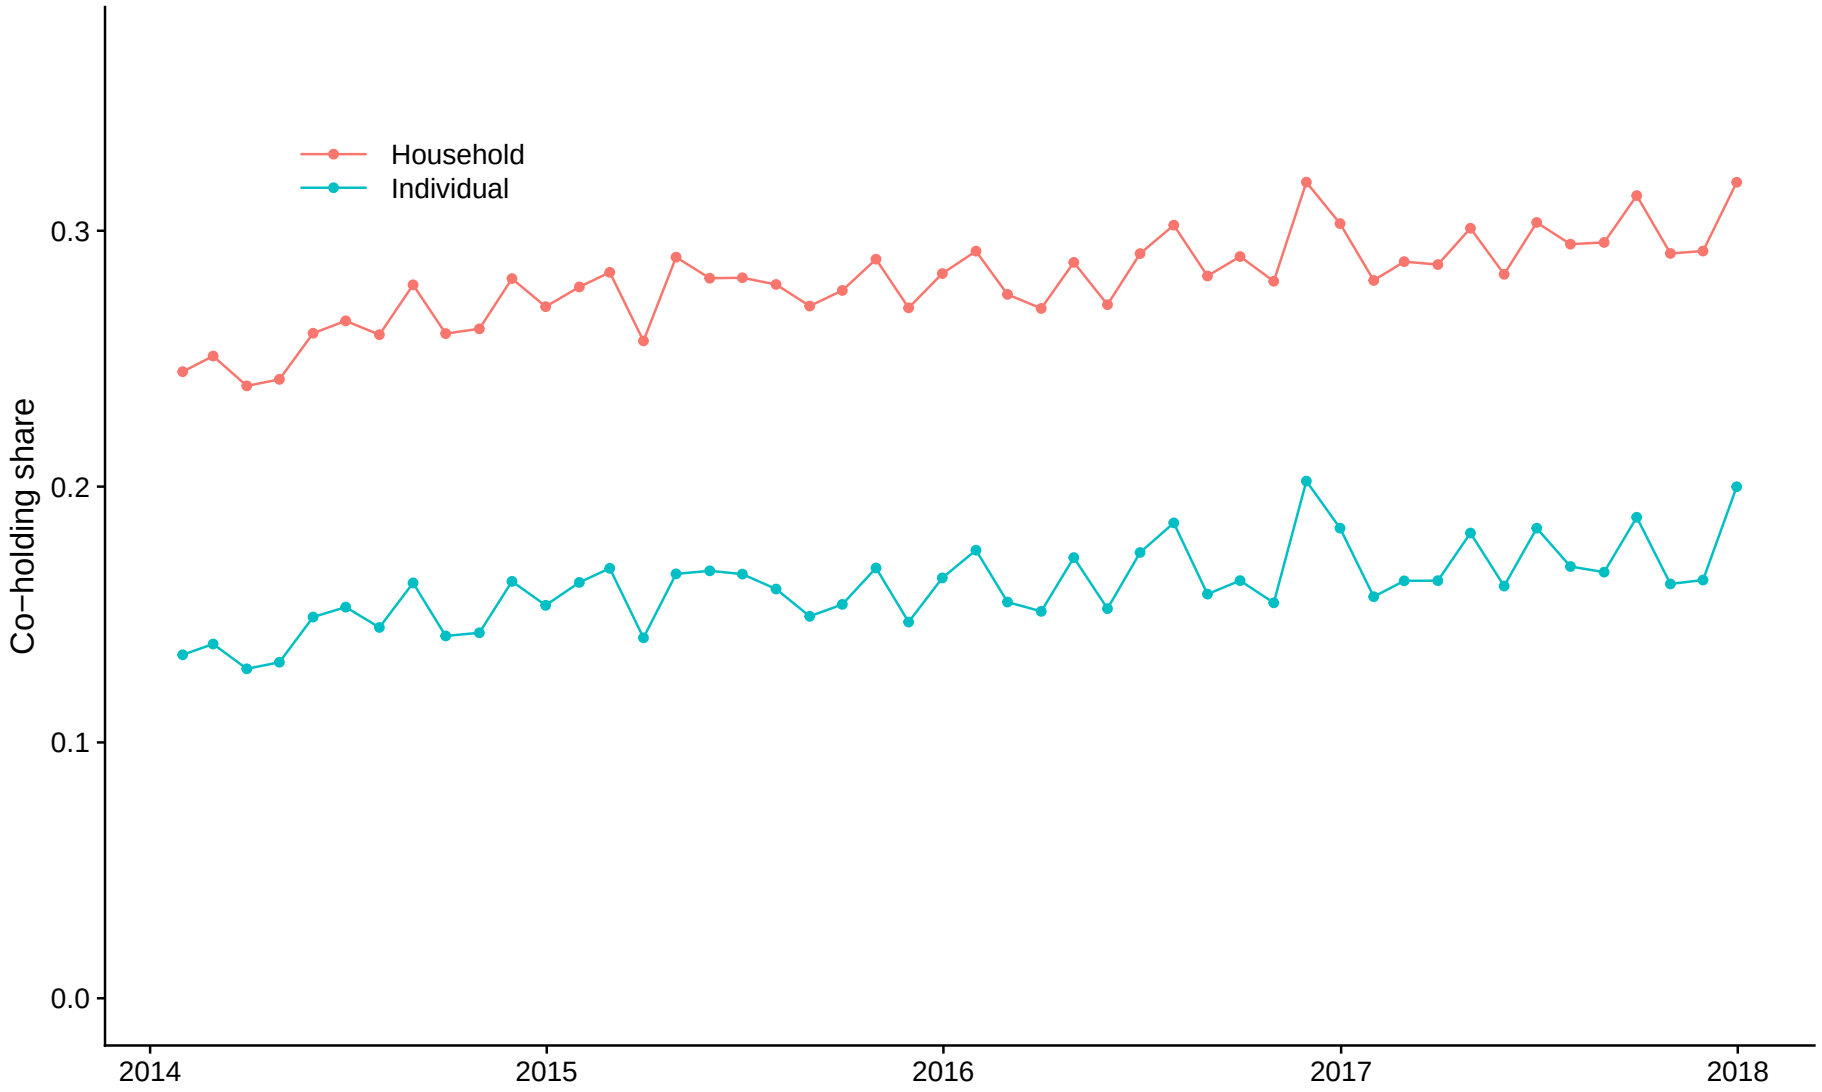
Note: The figure depicts the average persistence of co−holding for households who co−hold at some point during January 2014 to December 2016 where I measure co−holding as a continuous EUR variable. In addition to presenting the average persistence by all households who co−hold at some point in time, I present results separately depending on whether the initial level of co−holding is below or above median co−holding in the baseline month.

Figure 4: Co-holding is less persistent for larger initial co-holding balances (continuous measure of co-holding)



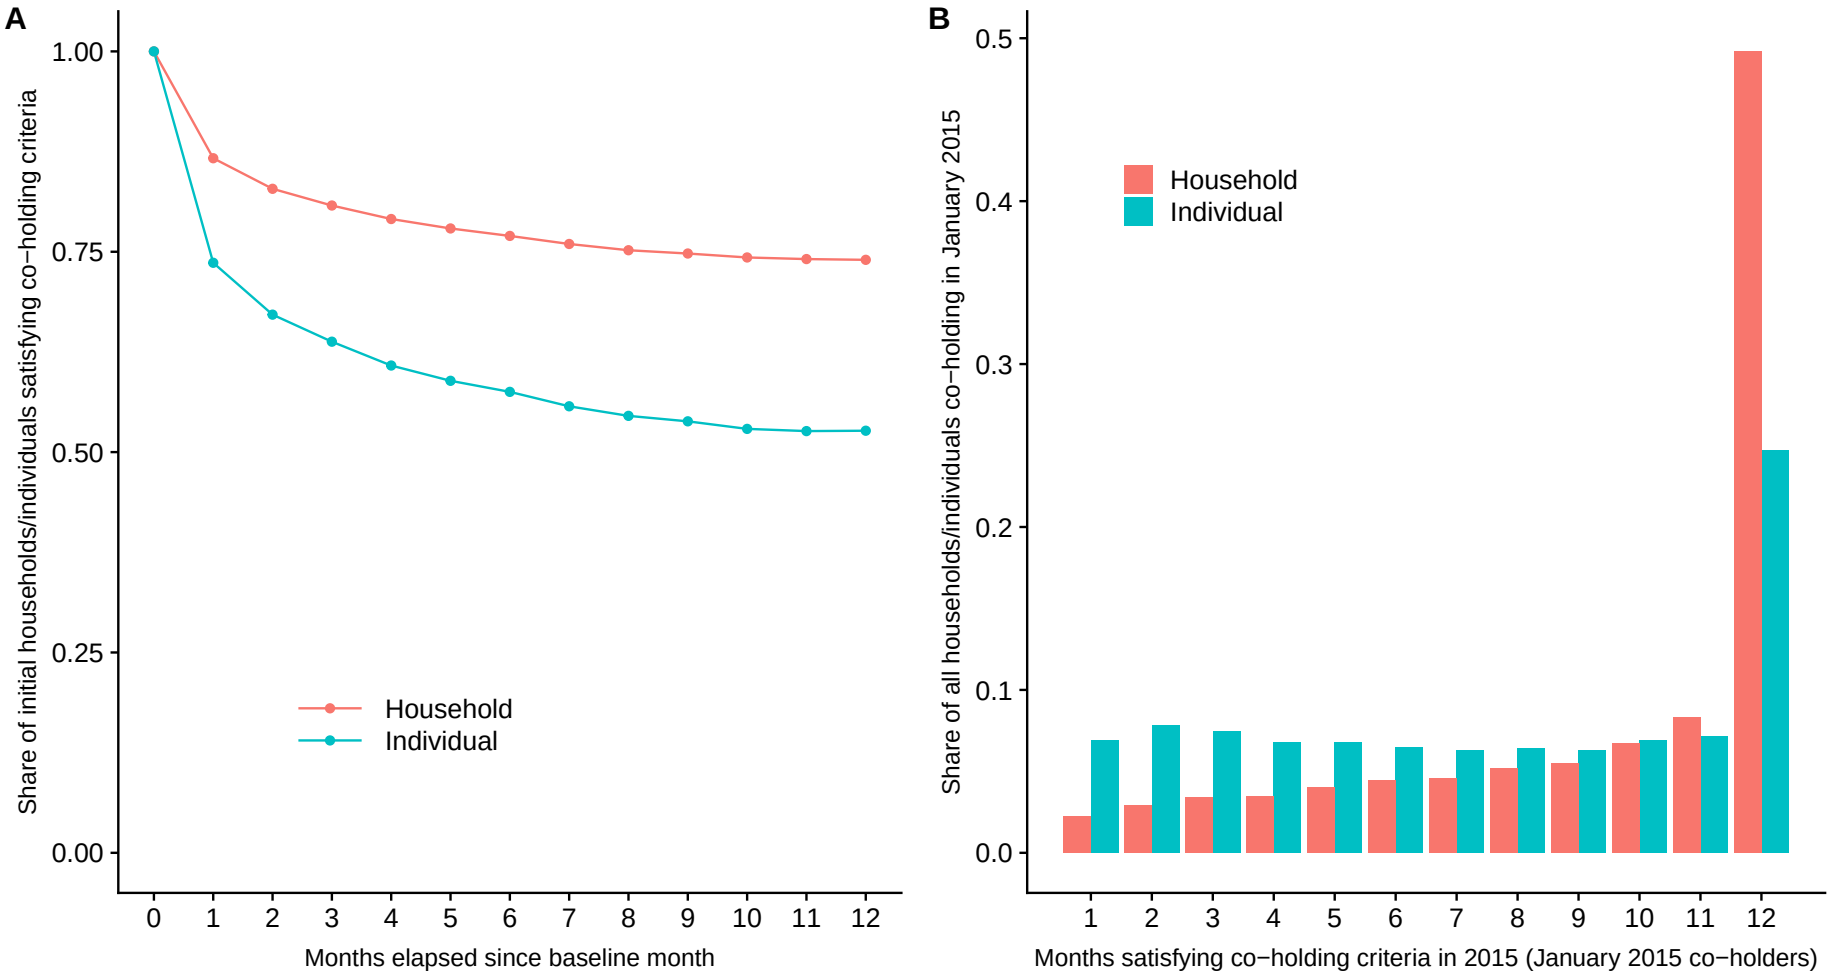
Note: The figure plots the share of households satisfying the co−holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) before and after the liquidity offer made by the bank in February 2015 (denoted by the dashed line) depending on eligibility for the liquidity offer. Households with a mortgage at the bank were eligible for the liquidity offer.

Figure 5: Probability of co-holding by eligibility before and after liquidity offer



Note: The figure plots the share of multi−adult households satisfying the co−holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) at the household level (red line). It also plots the proportion of these same households where at least one individual satisfies the co−holding criteria at the individual level (blue line).

Figure 6: Prevalence of co-holding at individual vs. household level in multi-adult households



Note: Panel A depicts the average survival share for households/individuals that satisfy the co−holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) at some point during January 2014 to December 2016. Panel B calculates the number of months that households/individuals co−hold in in 2015 conditional on satisfying the co−holding criteria in January 2015. Both figures are based on the sample of multi−adult households.

Figure 7: Persistence of co-holding: individual versus household

Table 1: Sample descriptives, January 2015

	N	Mean	Q0.1	Q0.25	Q0.5	Q0.75	Q0.9
Full sample							
Deposits	357,950	18,812	525	1,991	6,920	20,432	46,730
Unsecured debt	357,950	1,109	0	0	0	1,282	3,788
Co-holding	357,950	555	0	0	0	606	1,837
Disposable income	357,950	2,957	1,156	1,753	2,662	3,926	4,981
January 2015 puzzle							
Deposits	72,172	10,463	1,876	2,580	4,669	10,214	22,247
Unsecured debt	72,172	3,068	725	1,117	2,093	4,065	6,852
Co-holding	72,172	2,300	725	1,117	1,828	2,854	4,423
Disposable income	72,172	3,555	1,700	2,420	3,518	4,395	5,333

[a] Co-holding equals the minimum of deposits and unsecured debt.

[b] January 2015 puzzle households refer to households with over 500 EUR of unsecured debt and 1,500 EUR of deposits in January 2015.

Table 2: Household characteristics by household type

	Mean value				Median value			
	Borrower	Co-holder	Neutral	Saver	Borrower	Co-holder	Neutral	Saver
Group size	50,463	72,112	22,953	212,422	50,463	72,112	22,953	212,422
Demographics								
Adults	1.41	1.80	1.24	1.57	1.00	2.00	1.00	2.00
Children	0.41	0.81	0.22	0.44	0.00	0.00	0.00	0.00
Age of oldest adult	47.65	46.15	49.11	51.87	47.00	45.00	49.00	52.00
Co-holding								
Deposits	336.44	10,448.50	432.67	28,025.92	482.23	4,665.59	683.86	14,234.38
Interest-accruing unsecured debt	3,356.41	3,070.31	64.54	21.91	2,458.50	2,096.01	0.00	0.00
Unsecured credit limit	4,625.45	5,748.43	2,390.65	3,293.51	4,000.00	5,000.00	2,000.00	2,700.00
Co-holding (Jan 2015)	535.25	2,301.88	51.16	21.91	482.23	1,828.31	0.00	0.00
Co-holding (Feb–Dec 2015)	698.06	1,986.74	137.78	87.70	527.81	1,573.71	14.41	0.00
Co-holding months in 2015	1.78	9.07	0.37	0.52	1.00	11.00	0.00	0.00
Income/expenditure								
Disposable income	2,404.08	3,555.04	1,977.26	2,991.79	2,106.71	3,517.78	1,783.34	2,673.78
Card expenditure	1,163.46	1,918.47	1,014.82	1,523.88	994.39	1,837.56	849.20	1,322.25
Other debts and assets								
Mortgage indicator variable	0.56	0.72	0.44	0.45	1.00	1.00	0.00	0.00
Mortgage balance	49,365.77	82,484.44	32,611.81	42,739.06	16,034.45	64,435.35	0.00	0.00
Other financial assets	803.87	4,118.56	3,087.00	18,864.07	0.00	0.00	0.00	0.00
Other debt	5,554.09	6,117.29	2,770.14	2,031.44	0.00	0.00	0.00	0.00
Borrowing rates								
Unsecured debt interest rate	6.94	6.95	7.10	7.08	7.08	7.08	7.08	7.08
Mortgage interest rate	1.44	1.35	1.44	1.31	1.35	1.26	1.35	1.24

[a] The table presents household characteristics by household type. 'Borrower' households have less than 1,500 EUR of deposits and more than 500 EUR of unsecured debt. 'Co-holder' households have more than 1,500 EUR of deposits and more than 500 EUR of unsecured debt. 'Saver' households have more than 1,500 EUR of deposits and less than 500 EUR of unsecured debt. 'Neutral' households have less than 1,500 EUR of deposits and less than 500 EUR of unsecured debt. Values refer in general to January 2015 values, but disposable income is a monthly average over 2015, and card expenditure is a monthly average over 2014. Interest rates are nominal rates (not APRs) and do not include non-interest credit fees.

Table 3: January 2015 co-holding share by group

	Disposable income quintile					
	All	Q1	Q2	Q3	Q4	Q5
Age quintile						
All	0.20	0.08	0.13	0.19	0.30	0.30
Q1	0.22	0.10	0.14	0.21	0.32	0.29
Q2	0.28	0.13	0.17	0.23	0.35	0.34
Q3	0.25	0.12	0.15	0.21	0.32	0.33
Q4	0.18	0.09	0.13	0.18	0.25	0.25
Q5	0.09	0.06	0.09	0.12	0.14	0.16

[a] The table tabulates the share of co-holders (unsecured debt above 500 EUR and disposable income above 1,500 EUR) in each age and income quantile in January 2015

Table 4: Cost of co-holding at household level for January 2015 co-holders

	N	Mean	Q0.1	Q0.25	Q0.5	Q0.75	Q0.9
Extrapolated yearly costs	72,171	197.19	93.39	116.55	161.89	236.96	344.87
Actual yearly costs	72,171	172.33	62.66	95.75	144.39	217.67	315.12
Actual to extrapolated ratio	72,171	0.90	0.45	0.68	0.90	1.05	1.29

[a] The table presents two estimates of the yearly costs of co-holding for households satisfying the co-holding criteria (unsecured debt above 500 EUR and deposits over 1,500 EUR) in January 2015. Both estimates are based on the interest rate difference between liquid assets and unsecured debt and the amount of liquid assets available to reduce debt [equation (1)]. The first estimate, 'Extrapolated yearly costs', assumes that yearly costs are proportional to co-holding in January 2015. The second estimate, 'Actual yearly costs', uses data on actual co-holding in each month during 2015. The final row, 'Actual to extrapolated ratio', calculates the ratio of the two cost measures for each household.

Table 5: Co-holding often associated with high liquidity

Liquidity measure	N	Q25	Q50	Q60	Q70	Q80	Q90	Q95
Deposits in EUR	72,172	2,580.43	4,668.94	6,145.64	8,509.49	12,644.99	22,247.21	35,818.78
Deposits-to-income ratio	72,172	0.85	1.46	1.90	2.61	3.90	7.06	11.75

[a] The table presents the distribution of liquid assets associated with co-holding observations in January 2015. I define a household as a 'co-holder' if the household has at least 500 EUR of unsecured debt and 1,500 EUR of deposits. The two measures of liquidity are end of month deposits in EUR, and deposits scaled by average monthly disposable income over 2015.

Table 6: Credit limit risk in bank data

Year	Number of accounts	P(limit increase)	P(limit decrease)	P(Customer closure)	P(Bank closure)
2015	1,098,608	2.8	0.3	3.0	0.7
2016	1,098,488	2.9	0.2	1.9	0.7
2017	1,121,244	3.0	0.3	1.6	0.6
Average		2.9	0.2	2.2	0.7

[a] The table presents the yearly risk of credit limit changes over 2015–2017. Values refer to percentage points. I do not have the data for 2014.

Table 7: Take up of liquidity offer among January 2015 co-holders

	January 2015 unsecured debt quartile				
	All	Lowest quartile	Second quartile	Third quartile	Top quartile
Co-holding months in 2015					
All	0.29	0.23	0.27	0.31	0.37
1 to 3	0.28	0.21	0.28	0.32	0.42
4 to 6	0.30	0.23	0.29	0.33	0.40
7 to 9	0.31	0.25	0.29	0.35	0.38
10 to 12	0.29	0.23	0.26	0.30	0.35

[a] The table presents the take up rate of the liquidity offer by households with a mortgage at the bank (=eligible) by initial unsecured debt and the number of months that the household satisfies the co-holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) in 2015.

Table 8: Effect of liquidity offer on co-holding (binary dep. variable)

	Dependent variable:			
	Co-holding (0/1)			
	(1)	(2)	(3)	(4)
$\mathbb{1}_{m<\text{Jan2015}} * \mathbb{1}_{\text{mortgage}}$	0.15^{***}	-0.01^{***}	0.02^{***}	0.004
	(0.002)	(0.002)	(0.01)	(0.005)
$\mathbb{1}_{m>\text{Jan2015}} * \mathbb{1}_{\text{mortgage}}$	0.16^{***}	0.002	0.03^{***}	0.01^{**}
	(0.002)	(0.002)	(0.01)	(0.004)
Sample	Full	Full	Matched	Matched
Household FEs	No	Yes	No	Yes
Co-holding share	0.2	0.2	0.26	0.26
Observations	3,221,568	3,221,568	2,303,424	2,303,424

Note: *p<0.1; **p<0.05; ***p<0.01

The table presents estimation results for different specifications of equation (2). The data is a panel from January 2014 to December 2016. The dependent variable takes the value 1 if the household co-holds in a given month (unsecured debt over 500 EUR and deposits over 1,500 EUR) and 0 otherwise. January 2015 is the baseline month because it prececed the bank's liquidity campaign for mortgage households that started in February 2015. Standard errors are clustered at the household level. The specification in column 1 does not include household fixed effects and hence the coefficients estimate the overall difference in co-holding probability by mortgage status before and after the offer. The specification in column 2 includes household fixed effects and hence the coefficients estimate the difference in the trend of co-holding by mortgage status before and after the offer. Columns 3–4 provide equivalent estimates using the matched sample where I find for each household eligible for the liquidity offer a similar-on-observables non-eligible household using propensity score matching.

Table 9: Effect of liquidity offer on co-holding (continuous dep. variable)

	Dependent variable:			
	co − holding			
	(1)	(2)	(3)	(4)
$\mathbb{1}_{m<\text{Jan2015}} * \mathbb{1}_{\text{mortgage}}$	416.98***	−40.25***	52.05***	7.29
	(6.60)	(4.11)	(18.19)	(14.98)
$\mathbb{1}_{m>\text{Jan2015}} * \mathbb{1}_{\text{mortgage}}$	482.03***	24.80***	70.96***	26.20*
	(7.02)	(4.70)	(20.04)	(15.42)
Sample	Full	Full	Matched	Matched
Household FEs	No	Yes	No	Yes
Mean co-holding (EUR)	555.3	555.3	739.9	739.9
Observations	3,221,568	3,221,568	2,303,424	2,303,424

Note: *p<0.1; **p<0.05; ***p<0.01

The table presents estimation results for a variation of equation (2) where I change the dependent variable to the EUR amount of co-holding. The data is a panel from January 2014 to December 2016. January 2015 is the baseline month because it prececed the bank's liquidity campaign for mortgage households that started in February 2015. Standard errors are clustered at the household level. The specification in column 1 does not include household fixed effects and hence the coefficients estimate the overall difference in co-holding by mortgage status before and after the offer. The specification in column 2 includes household fixed effects and hence the coefficients estimate the difference in the trend of co-holding by mortgage status before and after the offer. Columns 3–4 provide equivalent estimates using the matched sample where I find for each household eligible for the liquidity offer a similar-on-observables non-eligible household using propensity score matching.

Table 10: Liquidity offer take up and the intra-household distribution of assets and liabilities

	Dependent variable:			
	Liquidity offer take up (0/1)			
	(1)	(2)	(3)	(4)
depositShare	-0.016^{***}	-0.011	0.006	-0.059^{***}
	(0.006)	(0.007)	(0.007)	(0.010)
unsecuredDebtShare	0.028^{***}	0.031^{***}	0.007	0.070^{***}
	(0.004)	(0.005)	(0.005)	(0.008)
Sample	All two-adult	Puzzle	Joint mortgage	Individual mortgage
Take-up rate	0.28	0.28	0.29	0.23
Unique individuals	59,156	39,622	41,577	14,832
Unique households	51,706	34,416	38,747	11,702
Household controls	Yes	Yes	Yes	Yes
Observations	58,191	38,973	40,926	14,554

Note: *p<0.1; **p<0.05; ***p<0.01

The table presents estimation results of equation (3) with varying samples. Model (1) includes all two-adult households with positive deposits and unsecured debt and a mortgage in January 2015. Model (2) includes the subset of households in Model (1) that satisfied the co-holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) in January 2015. Model (3) includes the subset of households in Model (1) with only joint mortgage(s). Finally, Model (4) includes the subset of households in Model (1) with only individual mortgage(s). The dependent variable takes the value 1 if the primary mortgage holder takes up the liquidity offer. The unit of observation is an individual primary mortgage holder within a two-adult household. All models control for the total amount of deposits and unsecured debt, the number of adults and children, average age of adults, municipality, mortgage value, size of principal payment, mortgage interest rate, property value, disposable income level, card purchase level.

Table 11: Payer type distribution

Payer type	N	Share
Maximum	56,367	0.37
Minimum	30,235	0.20
Mixed	53,155	0.35
Near minimum	12,903	0.08

[a] The table tabulates the distribution of payer types based on the 2015 payment behavior of the individual with the highest credit card limit within the household. A minimum payer makes contractual minimum payments on the card in at least 50 percent of account-invoice observations. A near-minimum payer makes payments within 50 EUR of contractual minimum payments on the card in at least 50 percent of account-invoice observations. A max payer pays the total invoice balance in at least 50 percent of account-invoice observations. The residual cases are categorized as mixed payers.

Table 12: Low credit card debt payments coincide often with significant liquidity

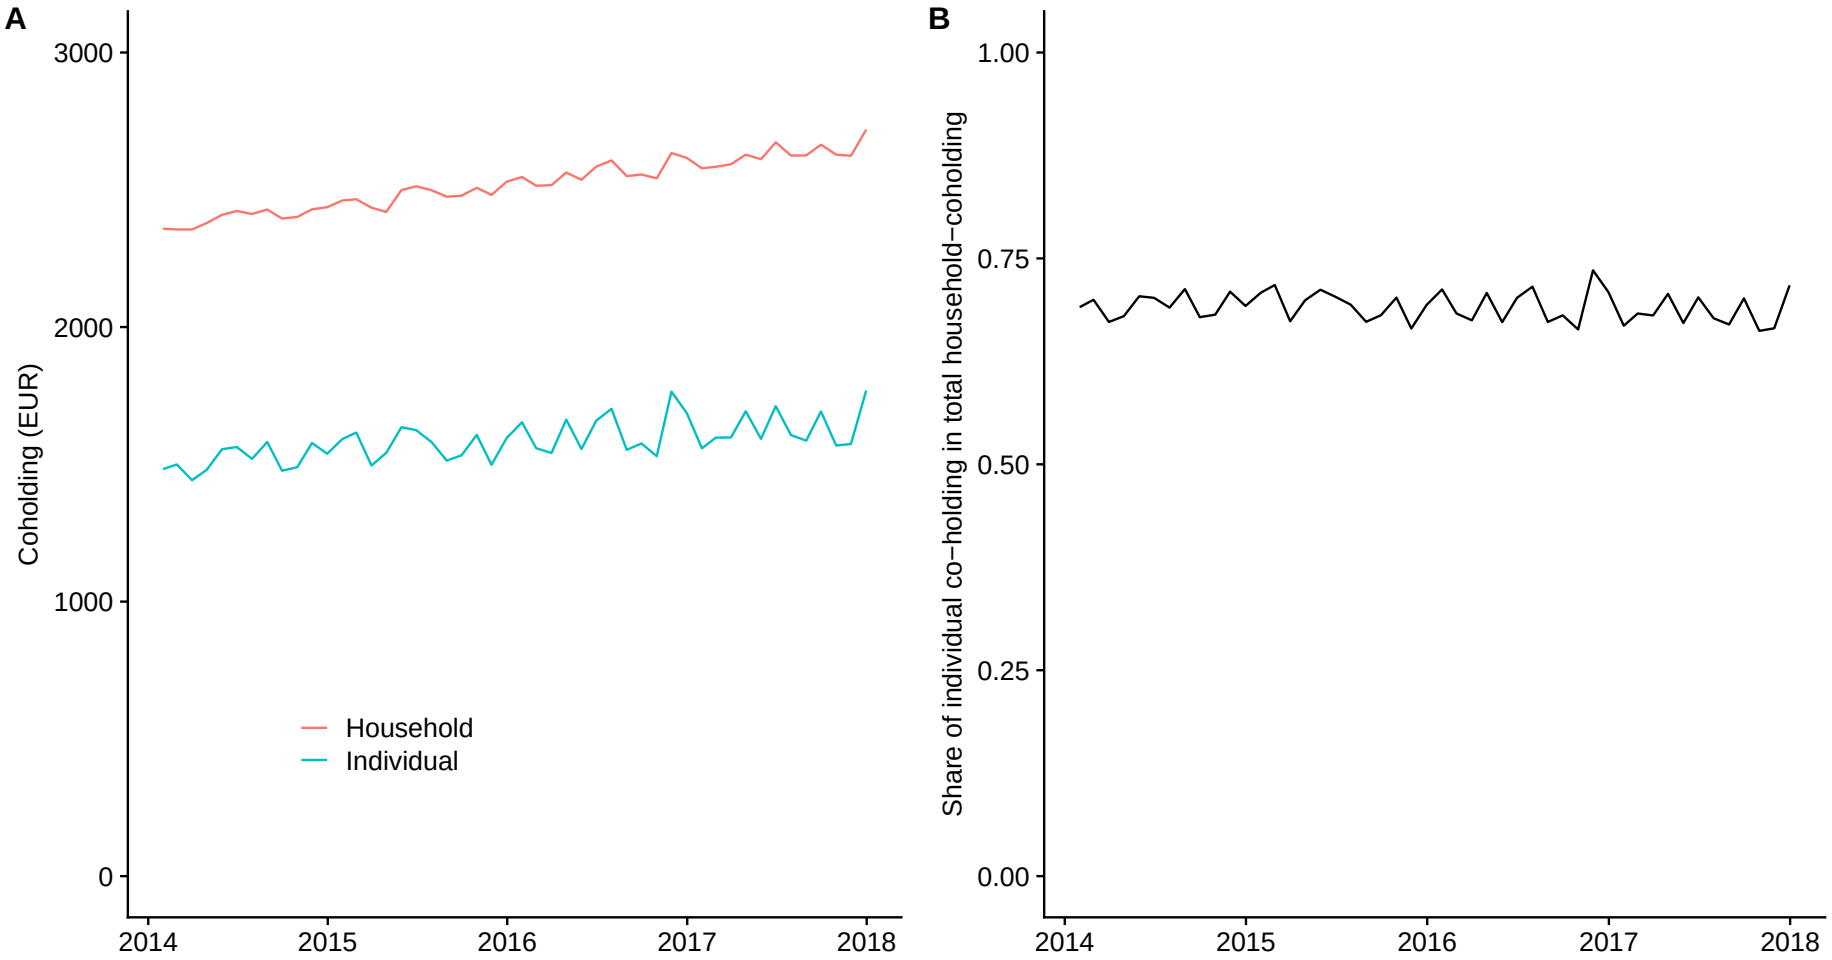
	N	Q25	Q50	Q60	Q70	Q80	Q90	Q95
A) Minimum payments								
Pre-invoice deposits-to-income	321,024	0.16	0.55	0.81	1.17	1.92	4.27	8.32
Post-invoice deposits-to-income	321,024	0.15	0.55	0.81	1.17	1.91	4.28	8.33
B) Near-minimum payments								
Pre-invoice deposits-to-income	172,698	0.15	0.53	0.77	1.12	1.80	3.97	7.56
Post-invoice deposits-to-income	172,698	0.15	0.53	0.78	1.13	1.81	3.99	7.61
C) Near-minimum payments (active)								
Pre-invoice deposits-to-income	50,631	0.19	0.59	0.83	1.17	1.80	3.63	6.57
Post-invoice deposits-to-income	50,631	0.20	0.60	0.85	1.18	1.83	3.64	6.57

[a] The table presents the distribution of (individual) deposits to monthly disposable income ratios associated with minimum or near-minimum credit card debt payments in 2015 by the individual with the highest credit card limit within the household. Panel A presents the distribution of the deposit-to-income ratio associated with account-invoice observations when the payment equals the contractual minimum. Panel B presents the distribution of deposit-to-income ratio associated with account-invoice observations when the payment is within 50 EUR of the contractual minimum. Panel C presents the distribution of the deposit-to-income ratio associated with account-invoice observations when the payment is within 50 EUR of the contractual minimum and the default invoice amount equals the contractual minimum. "Post-invoice deposits-to-income" refer to deposits measured in the last day of the invoice month. "Pre-invoice deposits-to-income" refer to deposits measured in the last day of the month preceding the invoice month. The income measure is fixed at the average monthly disposable income over 2015.

Table 13: Households with low regular credit card invoice payments co-hold more

Payer type	P(co-hold)	Mean co-holding
Maximum	0.05	127.59
Minimum	0.41	1,262.34
Mixed	0.31	823.88
Near minimum	0.40	1,189.63

[a] The table presents the probability of satisfying the co-holding criteria (unsecured debt above 500 EUR and deposits above 1,500 EUR) and the average level of co-holding over 2014–2017 depending on the credit card payer type of the individual with the highest credit limit within the households. A minimum payer makes contractual minimum payments on the card in at least 50 percent of account-invoice observations. A near-minimum payer makes payments within 50 EUR of contractual minimum payments on the card in at least 50 percent of account-invoice observations. A max payer pays the total invoice balance in at least 50 percent of account-invoice observations. The residual cases are categorized as mixed payers.



Note: Panel A presents average co−holding at the household level and the average of the sum of co−holding by individuals in multi−adult households for households who satisfy the co−holding criteria (unsecured debt over 500 EUR and deposits over 1,500 EUR). Panel B depicts the share of household co−holding accounted for by individual co−holding.

Figure A.1: Individual co-holding as share of household co-holding in multi-adult households

Sources of solace — household heterogeneity and macroeconomic insurance mechanisms

Erkki Vihriälä

Abstract

This paper extends the framework of Asdrubali et al. (1996) to study household heterogeneity in insuring against earnings shocks using: (i) investment income flows; (ii) the tax-and-transfer-system; and (iii) credit markets. The results imply strong variation depending on household socioeconomic characteristics captured by housing tenure. Low-income households rely on fiscal smoothing whereas aggregate credit smoothing is driven by wealthier households. The complementarity of (i)–(iii) has implications for the design of optimal insurance mechanisms against macroeconomic shocks.

JEL Codes: E21, E24, E32, E60, D12, D14, D31, D91

1 Introduction

Economic research has studied both how income shocks affect households at the microeconomic level, and how larger entities such as regions or countries are insured against output volatility. However, there is a disconnect between the two fields. Policy conclusions drawn from the aggregate literature do not reflect the importance of household heterogeneity, potentially leading to inappropriate recommendations. This paper attempts to bridge the gap.

In a seminal paper, Asdrubali et al. (1996) (henceforth ASY) quantified the extent to which US states smooth idiosyncratic output shocks using (i) factor income flows; (ii) the tax-and-transfer-system; and (iii) credit markets. They found that 39 % of the shocks are smoothed by factor income, 13 % through federal fiscal channels and 23 % via credit markets. Significantly, market insurance mechanisms (factor income and credit markets) are more important in the aggregate than fiscal smoothing.

ASY spawned comparative research on how other regions share risk. Particularly, smoothing of country-specific shocks in Europe lags behind that of US states (Sorensen and Yosha, 1998; Becker and Hoffmann, 2006; Sorensen et al., 2007; Demyanyk et al., 2008; Balli et al., 2012). Furthermore, smoothing of country-specific shocks falters in recessions (Furceri and Zdzienicka, 2013; Kalemli-Ozcan et al., 2014)[1]. Reasons for incomplete risk sharing in the euro area are multiple: factor income flows are inhibited by home bias; fiscal integration is almost non-existent; and credit market smoothing is constrained by incomplete financial integration.

Consequently, there have been calls for both a banking union and a euro-area fiscal capacity

[1]Hoffmann and Shcherbakova-Stewen (2011) detect the same effect in the US and find it stronger in states where small businesses are more important. However, the pro-cyclicality of risk sharing in the US has decreased since the banking deregulation of the 1980s.

(Pisani-Ferry et al., 2012; Wolff, 2012; Pisani-Ferry et al., 2013). Commentators have, however, differed on priorities. Based on the US results, some have favored further integration of capital and credit markets instead of mutual fiscal policy (see e.g. Hoffmann and Sorensen, 2012).

In contrast to papers in the ASY tradition, the microeconomic consumption smoothing literature (for a review, see e.g. Guvenen, 2011) has uncovered significant heterogeneity in how different types of households are able to insure against income shocks[2]. In summary, lower income households have more difficulties in smoothing shocks and fiscal taxes and transfers play a large role for them (Blundell et al., 2008, 2015, 2016). However, these considerations are absent from the aggregate risk sharing studies. The risk sharing literature implicitly assumes that what matters is overall smoothing, not for whom a particular insurance method is relevant. Therefore the different smoothing mechanisms are portrayed as substitutes although, if household heterogeneity is important, they should be considered as complements.

My paper identifies the importance of the factor income, fiscal, and credit market smoothing channels by household socioeconomic characteristics. Using British survey data from 1968 to 2015, I find that low-income households rely primarily on fiscal smoothing to stabilize their consumption. Overall credit smoothing is driven by wealthier households. There is therefore complementarity across the different insurance mechanisms. I do not find evidence for factor income smoothing but this can be influenced by under-reporting of factor income.

The analysis is based on housing tenure pseudo-cohorts formed from yearly cross-section surveys[3]. Housing tenure captures relevant variation in socioeconomic characteristics (income, wealth, education). Its slow-moving nature allows to identify a clearly defined group

[2]For the theory of partial insurance of income shocks due to incomplete markets see e.g. Aiyagari (1994); Heathcote et al. (2014).

[3]For grouping estimators, see Browning et al. (1985).

of households from year-to-year (Cloyne and Surico, 2017). I complement the main analysis based on actual housing tenures by testing the sensitivity of the results to constructing the cohorts by predicted tenure in the spirit of Attanasio et al. (2002).

To the best of my knowledge, the only papers employing the ASY decomposition to household-level data are Balli et al. (2016) and Asdrubali et al. (2015). Balli et al. use panel data to estimate the ASY equations by different subgroups (including income quartiles). Savings account for the bulk of smoothing for all income levels. However, their focus is not on how changes in market-based earnings for the working-age population are smoothed. Therefore they do not measure the insurance role of government benefits. What is more, if pensioners have relatively low current income but significant savings, measuring their smoothing will increase credit and factor smoothing estimates for low-income groups. The household-level data also does not account for private information in deriving the measure of income shock. This is much less of a problem with the tenure cohorts used in this paper. Finally, the data covers mostly food consumption, likely overstating smoothing.

Asdrubali et al. use Italian biannual survey data from 2008–2012 to estimate household-level smoothing. They capture intra-family insurance by studying the co-movement of individual and household income. On average Italian households were able to smooth around $\frac{4}{5}$ of changes in the income of the household head, mostly arising from within-family insurance as well as savings adjustment. There is little insurance on average from public taxes and transfers or investment income. Aggregate smoothing is lower for households with little net wealth. The major difference in approach is that I do not study insurance at the level of individual households (including within the household) with respect to individual income shocks. Instead, I measure smoothing of macroeconomic shocks by households differing in socioeconomic characteristics captured by housing tenure. Furthermore, I construct a long

time-series (1968–2013) of housing tenure groups and detect heterogeneity in the relevance of fiscal and credit smoothing between socioeconomic groups.

The organisation of the rest of the paper is the following. In the next two sections, I discuss the methodology before introducing the data. The fourth section presents the main empirical results. The final section concludes.

2 Methodology

2.1 Model

The classic full risk sharing studies assess whether consumption reacts to idiosyncratic income shocks, measured here by 'private earnings' (defined below). Full risk sharing implies $\beta = 0$ in:

$$\Delta \log C_{it} = \alpha + \beta \Delta \log \mathrm{PE}_{it} + \epsilon_{it} \tag{1}$$

In the case of incomplete risk sharing, the ASY methodology decomposes the smoothed part $(1 - \beta)$ into different channels. This is performed by starting from the following identity:

$$\mathrm{PE} = \frac{\mathrm{PE}}{\mathrm{PI}} \frac{\mathrm{PI}}{\mathrm{DI}} \frac{\mathrm{DI}}{\mathrm{C}} \mathrm{C} \tag{2}$$

, where

- PE (private earnings) refers to market-based income before taxes except investment

income; i.e wages and self-employment income

- PI (private income) = PE + investment income
- DI (disposable income) = PI − net taxes
- C (consumption)

Investment income only includes interest, rental and dividend income and does therefore not account for capital gains/valuation effects[4].

Taking a variance decomposition of the log-difference of (2) yields:

$$\begin{aligned}\text{Var}[\Delta \log \text{PE}] = \quad & \text{Cov}[\Delta \log \text{PE}, \Delta \log \text{PE} - \Delta \log \text{PI}] \\ & + \text{Cov}[\Delta \log \text{PE}, \Delta \log \text{PI} - \Delta \log \text{DI}] \\ & + \text{Cov}[\Delta \log \text{PE}, \Delta \log \text{DI} - \Delta \log \text{C}] \\ & + \text{Cov}[\Delta \log \text{PE}, \Delta \log \text{C}]\end{aligned}$$

Dividing by $\text{Var}[\Delta \log \text{PE}]$ yields:

$$1 \equiv \beta_K + \beta_F + \beta_C + \beta_U$$

This identity simply means that a shock is either smoothed fully via the three channels (factor income, fiscal, credit) or part of the shock remains unsmoothed, i.e. expenditure changes.

The β coefficients (covariance divided by variance) can be estimated in least squares regressions. Therefore, the main econometric specification is:

[4]My estimate of credit smoothing captures the financing of consumption via assets sales (see below).

$$\Delta \log(\text{PE})_{c,t} - \Delta \log(\text{PI})_{c,t} = \alpha_{K,t} + \delta_{K,c} + \beta_K \Delta \log(\text{PE})_{c,t} + \epsilon_{K,c,t} \tag{3}$$

$$\Delta \log(\text{PI})_{c,t} - \Delta \log(\text{DI})_{c,t} = \alpha_{F,t} + \delta_{F,c} + \beta_F \Delta \log(\text{PE})_{c,t} + \epsilon_{F,c,t} \tag{4}$$

$$\Delta \log(\text{DI})_{c,t} - \Delta \log(\text{C})_{c,t} = \alpha_{C,t} + \delta_{C,c} + \beta_C \Delta \log(\text{PE})_{c,t} + \epsilon_{C,c,t} \tag{5}$$

$$\Delta \log(\text{C})_{c,t} = \alpha_{U,t} + \delta_{U,c} + \beta_U \Delta \log(\text{PE})_{c,t} + \epsilon_{U,c,t} \tag{6}$$

The subscript c refers to the household cohort. Because the time dummies α account for common shocks affecting all cohorts, my β coefficients measure smoothing of cohort-specific shocks. In the empirical section, I estimate either the aggregate importance of a particular smoothing channel, or heterogeneity in the importance of a particular smoothing channel by household tenure group. In the latter case, I interact the explanatory private earnings growth variable with a dummy for the specific household tenure group.

The equations measure, respectively, smoothing via factor income flows, the fiscal authority and the credit markets. The last equation quantifies the unsmoothed residual and is analoguous to equation (1). As an intuitive example, if in equation (4) β_F is large, shocks to market earnings translate into little change in income after taxes and transfers, indicating fiscal smoothing. Conversely, if β_F is small, the growth rate of market earnings and income after taxes and transfers matches closely (little/no fiscal smoothing).

2.2 Relationship with ASY

Although the general empirical structure is similar to ASY, my approach differs in two respects.

Firstly, ASY use changes in an output measure (gross state product) as their exogeneous variable whereas I use the previously defined measure of private earnings at the cohort level. Shocks to output are naturally not equivalent to shocks to market earnings. For instance, smoothing within the firm (Guiso et al., 2005) can drive a wedge between the two. On the one hand, my approach allows to study household heterogeneity in smoothing mechanisms. On the other hand, I cannot account for the output shocks that are not reflected in private earnings growth.

Secondly, ASY measure risk sharing between regions whereas I estimate heterogeneity in smoothing by socioeconomic groups of both national and regional shocks. In order to show the connection between the approaches, I estimate separately (a) aggregate smoothing of regional shocks (ASY approach); (b) smoothing of national shocks by household socioeconomic group (proxied by housing tenure); and (c) smoothing of regional shocks by household socioeconomic group (dissaggregated ASY). Contrasting (a) and (c) distinguishes between aggregate and heterogenous smoothing estimates. Comparing (b) and (c) reveals whether there are differences in how national and regional shocks are smoothed by socioeconomic groups.

2.3 Definition of shocks

ASY considers cross-sectional deviations of gross state product growth from average growth as their exogeneous shock. In a similar vein, my shock measure is the difference between cohort growth from average cohort growth. Because I use cohort-aggregated data, the problem of disentangling predicted income changes from true "shocks", particularly problematic with household-level panel data, is attenuated. Furthermore, synthetic cohorts constructed

from yearly cross-sections are robust to attrition, which allows for a longer time dimension than most panels (Deaton, 1985).

3 Data

The British Family Expenditure Survey has data on household incomes and consumption available annually online from 1968 to 2015[5]. The yearly surveys contain data on approximately 5000–7500 households per year. The quality of the data is in general good, although investment and self-employment income are prone to under-reporting (see Appendix A.2).

To focus on the effects of changes in labour earnings, I limit my sample to households whose head is aged between 25 and 60. I divide the sample into: 'social housing', 'private rent', 'owned with mortgage', 'owned outright'[6]. For each year and tenure group I drop observations with consumption in the bottom or top 1 % to limit the effects of outliers on the estimation results. The ensuing sample sizes for each tenure and year are presented in Table I. In some of the regressions I use quarterly and/or regional data as further detailed in section 4.

Income variables were presented in the previous section[7]. I measure consumption by expenditure on non-durables, although the results are consistent when using non-housing expenditure. I construct cohort variables by averaging over households in a particular tenure group using sampling weights. Nominal values are deflated using a retail price index ex-

[5]The UK budget survey has been renamed twice. The original 'Family Expenditure Survey' (1957—2001) was replaced by the 'Expenditure and Food Survey' (2001—2007), which was transformed into the 'Living Costs and Food Survey' in 2008. All the raw data is available from the UK Data Service website (https://www.ukdataservice.ac.uk/) after registration.

[6]I exclude a few categories with very few observations from the analysis, e.g. 'rent free'

[7]To account for the small share of private pension income and social security pension income (given my age criterion), I classify these subsets into private earnings and social benefits respectively

cluding housing. Changes in household composition are taken into account by equivalising income and consumption variables.

The validity of housing tenure pseudo-cohorts relies on identifying similar groups of households over adjacent years. I follow Cloyne and Surico (2017) in arguing that this is the case given the slow-moving nature of changes in tenure composition (Figure I). Furthermore, using the British Household Panel Survey, I am able to track actual tenure changes for households observed in multiple periods. The probability of staying in a given tenure class the following year (Table II) is very high (over 95 % for all but the 'private rent' category). Moreover, the flows across tenures do not distort the composition of groups. For a given year and tenure, households who were part of the same tenure group the previous year form an overwhelming share of the total. I also peform sensitivity analyses by running the regressions with a post-1990 sub-sample when the tenure shares have been particularly stable. Finally I re-do the analysis by using propensity score methods to predict tenure based on fully predictable observables (Attanasio et al., 2002).

3.1 Descriptive statistics on household tenure cohorts in the UK

The classification by housing tenure relies not only on its validity for the formation of pseudo-cohorts, but also on its ability to proxy for relevant heterogeneity in the socioeconomic characteristics of households.

Firstly, Figure II shows that social renters and mortgagors come in all ages whereas private renters and outright owners tend to be young and old respectively. Social renters have the lowest income followed by private renters, whereas the income distribution of mortgagors and outright owners is more similar (Figure III). Income is dominated by wages, although

government benefits are significant particularly for social renters (Figure IV). Investment income is most important in relative terms for outright owners but still quite low (6.5 %). Education levels have been rising across all tenure groups, but social renters remain relatively uneducated (Figure V).

The yearly cross-section surveys have very limited balance sheet information. However, The British Household and Panel Survey can be used to derive measures of net worth by tenure for years 1995, 2000 and 2005 (Table III). Housing forms a bulk of the total net worth of the household sector. Financial net worth, as reported, provides little cushion for most renters. Accumulation of assets is not needed, though, if credit is available. Nevertheless, self-reported credit constraints are again most prevalent for renters (Figure VI)[8].

Figures VII–X show how the different income subsets and consumption have evolved. Juxtaposing the tenure-specific and aggregate growth rate, it is clear that the aggregate quantities are driven primarily by mortgagors and to a lesser extent outright owners. The relatively large income volatility of the renters is presumably due to them bearing the brunt of the adjustment in total working hours. The large variation for private renters around 1990 can also be influenced by sampling error caused by relatively few observations (see Table I).

The equivalent scale in all the figures allows a visual examination of the importance of different smoothing channels. Private income behaves in a very similar manner to private earnings for all tenures, indicating little insurance from factor income flows. On the other hand, disposable income is significantly less volatile than market income for the (social) renters, pointing to fiscal smoothing. Finally, for all except social renters, consumption varies considerably less than disposable income, indicating credit smoothing.

[8]The data is from the Bank of England/NMG household survey available at http://www.bankofengland.co.uk/research/Pages/onebank/datasets.aspx.

4 Estimation and results

In this section I present quantitative estimates of the importance of different smoothing channels. I begin by estimating equations (3)–(6) using data dissaggregated only by region to make the connection with the ASY approach of measuring geographical risk sharing. Next, I dissaggregate the data by housing tenure to measure risk sharing of national shocks between socioeconomic groups. Lastly, I combine the two approaches by aggregating data into tenure-region cells. This last step estimates socioeconomic heterogeneity in smoothing channels vis-à-vis regional shocks.

In all cases, I include cohort fixed effects in addition to time fixed effects to allow for differences in average growth. I control for sampling variation by including a control for change in average age between adjacent years. Standard errors allow for auto-correlation within a cohort and contemporaneous correlation between cohorts. In the region-only and tenure-only dissaggregations, I divide the data into quarterly cells. However, due to the limited number of observations, I use annual data when studying housing tenure cells subdivided by region. I present the distribution of number of households in group-averaged cells in Table VII. Clearly, the results for private renters should be considered quite unreliable.

I first provide results on the overall smoothing of regional shocks, where I do not dissaggregate households by housing tenure. The regional set-up splits Britain into quarterly observations for 'South England', 'North England' and 'Rest of the UK'. I find that credit markets smooth 44 % of regional shocks, and the tax-and-transfer system smooths an additional quarter of shocks (Table IV). Factor income smoothing is negligible, but as mentioned before, the result can be due to underreporting of such income. I also do not capture factor income smoothing that happens at the firm level, unlike ASY, because I do not study output shocks.

Next, in the intermediate specification, I study risk-sharing of national shocks by tenure group (Table V). I find quantitatively large credit smoothing for houseowners for whom fiscal smoothing is only half as large. However, the picture is very different for social renters for whom the tax-and-transfer system smooths around 50 % of shocks whereas credit smoothing is insignificant. Private renters have the most balanced mix of public vs. private insurance, although inferences regarding them should take into account the small sample sizes.

Finally, Table VI presents the results for my main set-up, where I study smoothing of regional shocks by housing tenure. The results are quite similar to the intermediate set-up based on national shocks, with credit smoothing being particularly important for houseowners and fiscal smoothing for social renters. The outright owners have the highest credit smoothing estimates vis-à-vis regional shocks, also relative to mortgagors, although the difference is of a degree than order of magnitude. The larger (financial) net worth of outright owners could be the explanation (Kaplan et al., 2014)[9]. Unsurprisingly, the fraction of shocks that are not smoothed tends to be lower for regional than national shocks.

4.1 Sensitivity analysis

In this section I take the results regarding risk sharing of tenure-specific regional shocks (Table VI) as the baseline and study their sensitivity to relevant modifications.

I test if the results change if the early years of the sample are excluded given the relatively large changes in tenure shares. Using a post-1990 sample, the results do not materially change (Table A1). I choose 1990 as the cutoff year because it is after the peak in the 'Right to Buy'-acqusitions of social housing at a discount. However, results do not materially change

[9]On the other hand, in the case of national shocks I do not detect any difference in credit smoothing among houseowners depending on debt status.

when using an alternative cut-off (1986, 1988 or 1992)[10].

For completeness, Table A2 provides estimates when using only the first part of the sample until 1989. This excludes the 2000s when the private earnings of the social renters as measured in the survey have been particularly volatile. Fiscal smoothing remains the primary insurance mechanism for the (social) renters, even though quantitatively smaller than in later years. The credit smoothing estimates for houseowners and private renters are also somewhat lower than in the more recent period. Fiscal smoothing is estimated to have played a larger role for mortgagors in the early period.

Figures VII–X reveal some observations with high yearly growth rates. The fiscal smoothing estimate for social renters decreases somewhat but results stay qualitatively similar if I exclude cases when the absolute value of private earnings growth was above 10 % (Table A3). The results also do not change qualitatively if I measure consumption by total non-housing expenditure instead of non-durable consumption (Table A4).

I have argued above for the validity of (observed) housing tenure as a grouping variable. However, as a final robustness check, I re-do the analysis by classifying households based on their predicted probability of belonging to a particular tenure group based on fundamentals (Attanasio et al., 2002). The technique aims to guard against compositional change among a particular tenure group by conditioning classification into a group on equivalent observables in adjacent years.

The previous results show the clearest contrast between social housing tenants and houseowners. Therefore I follow Cloyne and Surico (2017) by dropping private renters from the sample and classify household into 'unlikely' and 'likely' social tenants by fully predictable

[10]These results are not separately reported.

observables. The tenure predictors in the probit regression are polynomials in education level, age of household head and a time trend as well as their interactions (Table A5). For each pair of adjacent years, I calculate income and and consumption growth for the two groups. As an intuitive explanation, let us suppose that households with a head of 45 or less with only compulsory education are *likely* to live in social housing in 2002. In that case, the consumption growth for the 'likely social housing'-group in 2003 is the difference in consumption among households with a head of 46 or less with only compulsory education in 2003 and the equivalent cohort in 2002.

Formally:

$$\Delta \log C_{t+1}^{social} = [\overline{\log C_{t+1}} \mid p(\text{social})_t > p_t] - [\overline{\log C_t} \mid p(\text{social})_t > p_t] \tag{7}$$

The notation $[\bar{x} \mid A]$ refers to the sample mean of the relevant income/consumption variable for the households that satisfy condition A. Crucially, the condition in the two right-hand side terms is the same: for households observed in year $t+1$, I calculate the probability that they lived in social housing in year t.

In order to classify households into likely and unlikely social renters, I use two different cut-offs p_t. The first is time-varying and ensures that the share of households classified in social housing *in the base period* is equal to the actual share of households in social housing in that period[11]. The second choice is a fixed cut-off that classifies *over the whole sample* an equal share of households in the 'likely social housing'-group as there are actual households in social housing. Attanasio et al. argue that the time-varying cut-off minimizes the effects of misclassification and the fixed cut-off those from compositional change.

[11] I.e. when I calculate the growth rate for year t, the base period is $t-1$.

For this excercise, I need to use data starting in 1978 because consistent information on education is not available prior to that. The results using either cut-off reveal that fiscal smoothing remains relatively important for households that are more likely to live in social housing (Tables A6 and A7). The differences among cohorts in point estimates are considerably smaller, though, than in the main results where I use observed housing tenure. Furthermore, credit smoothing is now important for both groups and only the time-invariant cut-off creates a higher point estimate for the unlikely social housing group.

The usefulness of the Attanasio et al. technique rests on the ability of the predictors (education level and age) to sufficiently account for differences in tenure status. The social housing probability density is multimodal (Figure A1). Some households (those with A-levels or college education) have very low probability of living in social housing. However, observables are not able to clearly differentiate among the large number of households with only compulsory education. Therefore, a significant mass of households with similar predicted tenure probabilities are classified into the 'likely' and 'unlikely' groups. If enough households with high income and consumption are *misclassified* to the 'likely' social housing group, they can significantly bias the smoothing estimates. In particular, this can lead to a downward estimate of the importance of fiscal smoothing for social renters (and vice versa for credit smoothing).

5 Discussion and conclusions

My results support the hypothesis that fiscal smoothing is important for low-income households. The evidence also indicates that aggregate credit smoothing is driven by the relatively well-off. I do not find any evidence for factor income smoothing although this can be influ-

enced by under-reporting of income from assets.

Furthermore, factor income smoothing could play an important role in the link between output shocks and private earnings and for agents not studied in this paper (pensioners, firms, governments) with concomitant effects on the economy as a whole. Therefore, my results do not provide a global assessment of the importance of different smoothing mechanisms, even though the heterogeneity in respect to private earnings shocks is important. Further mapping of output and earnings shocks by different socioeconomic groups would be worthwhile.

My paper stems from the observation that information on important subsections of the population can be lost with aggregate data. My tenure-specific approach is able to reveal some of the heterogeneity. However, there can be significant variation in the importance of different smoothing methods within (and not just across) tenure groups. For instance, 'mortgagors' can include a non-negligible subgroup for whom fiscal smoothing matters even though it would not be identified with tenure-aggregated data.

From a policy perspective, it is not only the degree of aggregate insurance that matters, but also the types of shocks that are smoothed. Firstly, it is disconcerting that smoothing falters especially in the case of large shocks (see Furceri and Zdzienicka, 2013)[12]. Secondly, the welfare effects of incomplete smoothing are qualitatively different depending on the baseline level of material well-being. However, the aggregate smoothing estimates largely reflect the experiences of the relatively well-off.

As an example of policy relevance, the euro area is in the process of developing its cross-border insurance mechanisms. There is a particular need given that national fiscal flexibility

[12]Global shocks naturally cannot be smoothed via risk sharing – here I refer to particularly adverse shocks to individual regions or demographic groups.

is limited by common political agreements and market panics. Although there are efforts to increase the integration of capital and credit markets, these are not necessarily sufficient to provide consumption insurance for all households. Inclusive consumption smoothing is unlikely to be achieved without effective national (or international) fiscal support.

References

Aiyagari, S. Rao (1994) "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, Vol. 109, pp. 659–84, URL: https://ideas.repec.org/a/tpr/qjecon/v109y1994i3p659-84.html.

Asdrubali, Pierfederico, Bent E. Sorensen, and Oved Yosha (1996) "Channels of Interstate Risk Sharing: United States 1963-1990," *The Quarterly Journal of Economics*, Vol. 111, pp. 1081–1110, URL: http://ideas.repec.org/a/tpr/qjecon/v111y1996i4p1081-1110.html.

Asdrubali, Pierfederico, Simone Tedeschi, and Luigi Ventura (2015) "Household Risksharing Channels," MPRA Paper 65906, University Library of Munich, Germany.

Atkinson, A. B. and J. Micklewright (1983) "On the Reliability of Income Data in the Family Expenditure Survey 1970- 1977," *Journal of the Royal Statistical Society. Series A (General)*, Vol. 146, pp. 33–61, URL: http://www.jstor.org/stable/2981487, DOI: http://dx.doi.org/10.2307/2981487.

Attanasio, Orazio P., James Banks, and Sarah Tanner (2002) "Asset Holding and Consumption Volatility," *Journal of Political Economy*, Vol. 110, pp. 771–792, URL: https://ideas.repec.org/a/ucp/jpolec/v110y2002i4p771-792.html.

Balli, Faruk, Sebnem Kalemli-Ozcan, and Bent E. Sørensen (2012) "Risk sharing through capital gains," *Canadian Journal of Economics/Revue canadienne d'économique*, Vol. 45, pp. 472–492, URL: http://onlinelibrary.wiley.com/doi/10.1111/j.1540-5982.2012.01702.x/abstract, DOI: http://dx.doi.org/10.1111/j.1540-5982.2012.01702.x.

Balli, Faruk, Filippo Maria Pericoli, and Eleonora Pierucci (2016) "Channels of Risk Sharing at Micro Level: Savings, Investments and Risk Aversion Heterogeneity," *International Journal of Finance & Economics*, Vol. 21, pp. 90–104, URL: http://onlinelibrary.wiley.com/doi/10.1002/ijfe.1538/abstract, DOI: http://dx.doi.org/10.1002/ijfe.1538.

Banks, J and P Johnson (1998) *How reliable is the Family Expenditure Survey? Trends in incomes and expenditures over time* in , IFS Reports, No. R57.

Becker, Sascha O. and Mathias Hoffmann (2006) "Intra- and international risk-sharing in the short run and the long run," *European Economic Review*, Vol. 50, pp. 777–806, URL: http://ideas.repec.org/a/eee/eecrev/v50y2006i3p777-806.html.

Blundell, Richard, Michael Graber, and Magne Mogstad (2015) "Labor income dynamics and the insurance from taxes, transfers, and the family," *Journal of Public Economics*, Vol. 127, pp. 58–73, URL: https://ideas.repec.org/a/eee/pubeco/v127y2015icp58-73.html.

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008) "Consumption Inequality and Partial Insurance," *The American Economic Review*, Vol. 98, pp. 1887–1921, URL: http://www.jstor.org/stable/29730156.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten (2016) "Consumption Inequality and Family Labor Supply.," *American Economic Review*, Vol. 106, pp. 387–435, URL: https://www.aeaweb.org/articles.php?doi=10.1257/aer.20121549.

Browning, Martin, Angus Deaton, and Margaret Irish (1985) "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, Vol. 53, pp. 503–43, URL: https://ideas.repec.org/a/ecm/emetrp/v53y1985i3p503-43.html.

Cloyne, James S. and Paolo Surico (2017) "Household Debt and the Dynamic Effects of Income Tax Changes," *The Review of Economic Studies*, Vol. 84, pp. 45–81, URL: https://academic.oup.com/restud/article/84/1/45/2669946, DOI: http://dx.doi.org/10.1093/restud/rdw021.

Deaton, Angus (1985) "Panel data from time series of cross-sections," *Journal of Econometrics*, Vol. 30, pp. 109–126, URL: http://ideas.repec.org/a/eee/econom/v30y1985i1-2p109-126.html.

Demyanyk, Yuliya, Charlotte Ostergaard, and Bent E. Sorensen (2008) "Risk sharing and portfolio allocation in EMU," European Economy - Economic Paper 334, Directorate General Economic and Monetary Affairs (DG ECFIN), European Commission.

Furceri, Davide and Aleksandra Zdzienicka (2013) "The Euro Area Crisis: Need for a Supranational Fiscal Risk Sharing Mechanism?" IMF Working Paper 13/198, International Monetary Fund.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi (2005) "Insurance within the Firm," *Journal of Political Economy*, Vol. 113, pp. 1054–1087, URL: http://ideas.repec.org/a/ucp/jpolec/v113y2005i5p1054-1087.html.

Guvenen, Fatih (2011) "Macroeconomics With Heterogeneity: A Practical Guide," Working Paper 17622, National Bureau of Economic Research.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2014) "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, Vol. 104, pp. 2075–2126, URL: https://ideas.repec.org/a/aea/aecrev/v104y2014i7p2075-2126.html.

Hoffmann, Mathias and Iryna Shcherbakova-Stewen (2011) "Consumption Risk Sharing over the Business Cycle: The Role of Small Firms' Access to Credit Markets," *Review of Economics and Statistics*, Vol. 93, pp. 1403–1416, URL: http://dx.doi.org/10.1162/REST_a_00123, DOI: http://dx.doi.org/10.1162/REST_a_00123.

Hoffmann, Mathias and Bent E. Sorensen (2012) "Don't expect too much from EZ fiscal union – and complete the unfinished integration of European capital markets!," URL: http://goo.gl/nF7322.

Johnson, Paul and Julian McCrae (1998) "Robustness of FES Income Data, 1985–92," in *How reliable is the family expenditure survey? Trends in incomes and expenditures over time*.

Kalemli-Ozcan, Sebnem, Emiliano Luttini, and Bent Sørensen (2014) "Debt Crises and Risk-Sharing: The Role of Markets versus Sovereigns," *The Scandinavian Journal of Economics*, Vol. 116, pp. 253–276, URL: http://onlinelibrary.wiley.com/doi/10.1111/sjoe.12043/abstract, DOI: http://dx.doi.org/10.1111/sjoe.12043.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner (2014) "The Wealthy Hand-to-Mouth," Working Paper 20073, National Bureau of Economic Research.

Pisani-Ferry, Jean, André Sapir, Nicolas Véron, and Guntram B. Wolff (2012) "What kind of European banking union?" Policy Contribution 2012/12, Bruegel.

Pisani-Ferry, Jean, Erkki Vihriälä, and Guntram B. Wolff (2013) "Options for a Euro-area fiscal capacity," Policy Contribution 2013/01, Bruegel.

Sorensen, Bent E., Yi-Tsung Wu, Oved Yosha, and Yu Zhu (2007) "Home bias and international risk sharing: Twin puzzles separated at birth," *Journal of International Money and Finance*, Vol. 26, pp. 587–605, URL: http://ideas.repec.org/a/eee/jimfin/v26y2007i4p587-605.html.

Sorensen, Bent E. and Oved Yosha (1998) "International risk sharing and European monetary unification," *Journal of International Economics*, Vol. 45, pp. 211–238, URL: http://ideas.repec.org/a/eee/inecon/v45y1998i2p211-238.html.

Tanner, Sarah (1998) "How much do consumers spend? Comparing FES and national accounts," in *How reliable is the family expenditure survey? Trends in incomes and expenditures over time*.

Wolff, Guntram B. (2012) "A budget for Europe's monetary union," Policy Contribution 2012/22, Bruegel.

Table I: Number of observations by year and tenure group

Year	Mortgage	Owned outright	Private rent	Social housing
1968	1605	630	812	1498
1969	1566	620	746	1281
1970	1451	479	658	1301
1971	1703	599	677	1436
1972	1772	531	603	1458
1973	1749	580	618	1323
1974	1681	596	489	1308
1975	1893	549	511	1354
1976	1777	573	512	1365
1977	1963	531	431	1408
1978	1915	534	429	1356
1979	1925	529	334	1276
1980	2062	505	277	1358
1981	2199	561	304	1527
1982	2341	562	266	1372
1983	2269	567	221	1128
1984	2369	543	240	1110
1985	2400	540	234	1029
1986	2508	494	253	973
1987	2559	535	242	1040
1988	2666	593	202	935
1989	2682	590	217	999
1990	2627	488	249	858
1991	2606	522	227	791
1992	2726	507	299	978
1993	2602	494	292	888
1994	2609	478	333	814
1995	2550	453	334	871
1996	2441	405	330	854
1997	2439	459	356	798
1998	2421	496	392	842
1999	2550	600	382	880
2000	2370	573	359	813
2001	2675	646	425	894
2002	2579	628	369	787
2003	2560	640	403	774
2004	2440	633	422	781
2005	2338	685	464	745
2006	2316	649	451	648
2007	2075	595	456	632
2008	1958	552	407	609
2009	1870	584	473	580
2010	1588	475	500	548
2011	1737	526	566	601

Table I: Number of observations by year and tenure group

Year	Mortgage	Owned outright	Private rent	Social housing
2012	1639	492	589	519
2013	1401	477	523	503
2014	1401	493	561	517
2015	1377	500	598	520

Table II: Tenure changes

Tenure	P(Same tenure next year)	Share of 'old' HHs
Owned outright	0.97	0.93
Mortgage	0.95	0.96
Social rent	0.96	0.96
Private rent	0.82	0.86

Note: These probabilities are calculated using the British Household Panel Survey.

Table III: Net wealth by tenure and asset class (current pounds)

Wealth	Year	Tenure	p25	median	p75	Obs.
Housing	1995	owned outright	48500	70000	106446	351
		mortgage	13500	34000	63500	1895
		social rent	0	0	0	498
		private rent	0	0	0	265
Housing	2000	owned outright	65000	107015	200000	448
		mortgage	23803	50000	93852	1823
		social rent	0	0	0	461
		private rent	0	0	0	258
Housing	2005	owned outright	140000	200000	300000	451
		mortgage	70196	118000	187589	1770
		social rent	0	0	0	368
		private rent	0	0	0	263
Financial	1995	owned outright	0	998	21243	351
		mortgage	-1200	0	4400	1895
		social rent	-360	0	0	498
		private rent	-1038	0	842	265
Financial	2000	owned outright	0	3501	25130	448
		mortgage	-1992	100	8002	1823
		social rent	-504	0	100	461
		private rent	-2600	0	300	258
Financial	2005	owned outright	0	2000	29964	451
		mortgage	-3000	0	5500	1770
		social rent	-1000	0	0	368
		private rent	-4916	0	151	263

Note: Derived using the British Household Panel Survey

Table IV: Smoothing of regional shocks – aggregate results

	Factor	Fiscal	Credit	Not smoothed
$\Delta \log(\text{Private earnings})$	-0.02	0.24^{***}	0.44^{***}	0.33^{***}
	(0.02)	(0.02)	(0.03)	(0.04)
Observations	567	567	567	567
Adjusted R^2	0.02	0.66	0.48	0.47

Note: *p<0.1; **p<0.05; ***p<0.01

The columns correspond, in order, to equations (3)–(6), where the cohorts are the three regions (South England, North England, rest of the UK). All regressions include quarter and region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for autocorrelation within region and contemporaneous correlation between regions. Data is quarterly from 1968 to 2015.

Table V: Smoothing of tenure shocks – national aggregation

	Factor	Fiscal	Credit	Not smoothed
Mortgage	0.05***	0.18***	0.45***	0.33***
	(0.01)	(0.05)	(0.08)	(0.03)
Owned Outright	−0.02	0.23***	0.46***	0.32***
	(0.02)	(0.02)	(0.06)	(0.05)
Private rent	−0.004	0.23***	0.28***	0.50***
	(0.02)	(0.02)	(0.05)	(0.04)
Social rent	0.02**	0.54***	0.05	0.39***
	(0.01)	(0.02)	(0.03)	(0.03)
Observations	756	756	756	756
Adjusted R^2	0.03	0.67	0.26	0.40

Note: *p<0.1; **p<0.05; ***p<0.01

The columns correspond, in order, to equations (3)–(6), where the cohorts are the four tenure groups (mortgage, owned outright, private rent, social rent). Covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include quarter and tenure fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for auto-correlation within tenure and contemporaneous correlation between tenures. Data is quarterly from 1968 to 2015.

Table VI: Smoothing of tenure shocks – regional aggregation

	Factor	Fiscal	Credit	Not smoothed
Mortgage	0.02***	0.18***	0.47***	0.33***
	(0.001)	(0.04)	(0.06)	(0.06)
Owned Outright	−0.003	0.24***	0.62***	0.15**
	(0.04)	(0.06)	(0.06)	(0.07)
Private rent	−0.005	0.29***	0.29***	0.43***
	(0.01)	(0.03)	(0.06)	(0.03)
Social rent	−0.0004	0.60***	0.16***	0.24***
	(0.01)	(0.02)	(0.05)	(0.07)
Observations	564	564	564	564
Adjusted R^2	-0.003	0.76	0.30	0.31

Note: *p<0.1; **p<0.05; ***p<0.01

The columns correspond, in order, to equations (3)–(6), where the cohorts are the four tenure groups (mortgage, owned outright, private rent, social rent) subdivided by region. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for auto-correlation within tenure-region groups and contemporaneous correlation between tenure-region groups. Data is annual from 1968 to 2015.

Table VII: Number of households in group-averaged cells

region	min	p25	median	p75	max	mean
englandNorth	280	393	441	470	530	427
englandSouth	257	402	425	447	496	413
otherUK	105	160	174	197	271	179

Region-quarter regression (Table IV)

tenure	min	p25	median	p75	max	mean
mortgage	230	444	552	622	715	535
owned outright	84	124	136	149	194	137
private rent	41	72	100	128	213	105
social rent	107	185	224	323	426	243

Tenure-quarter regression (Table V)

tenure	min	p25	median	p75	max	mean
mortgage	143	415	772	1016	1202	715
owned outright	55	125	197	230	296	183
private rent	17	75	122	198	355	139
social rent	73	232	310	398	649	326

Annual tenure-region regression (Table VI)

Figure I:
Tenure composition

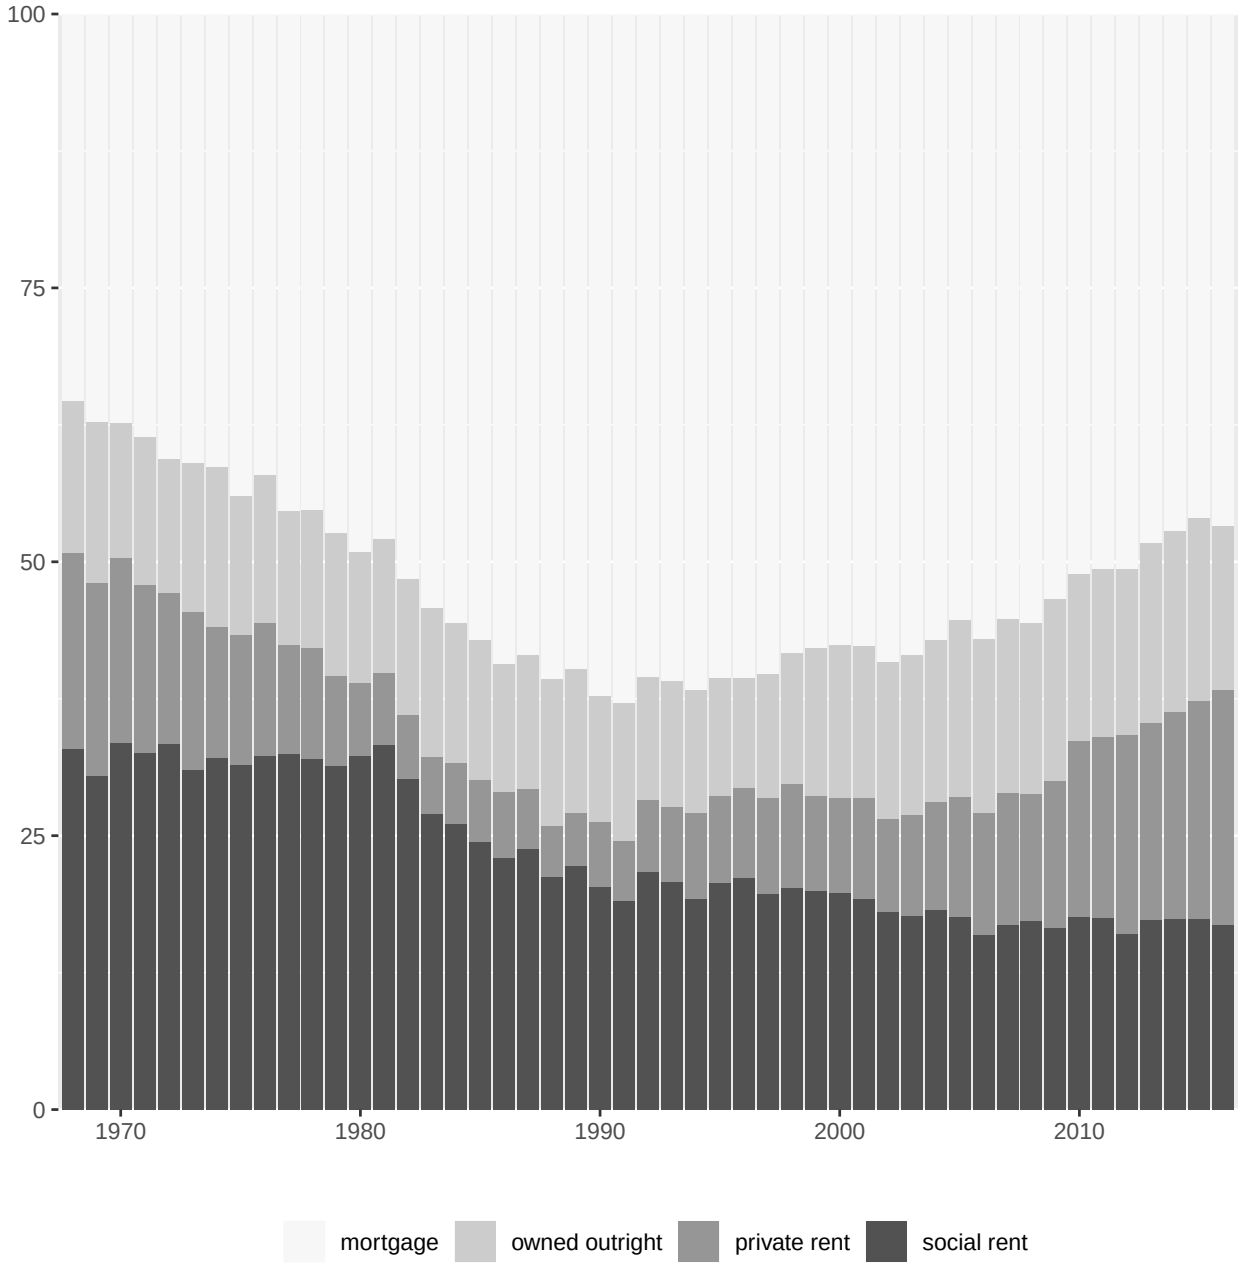


Figure II:
Age distribution by tenure

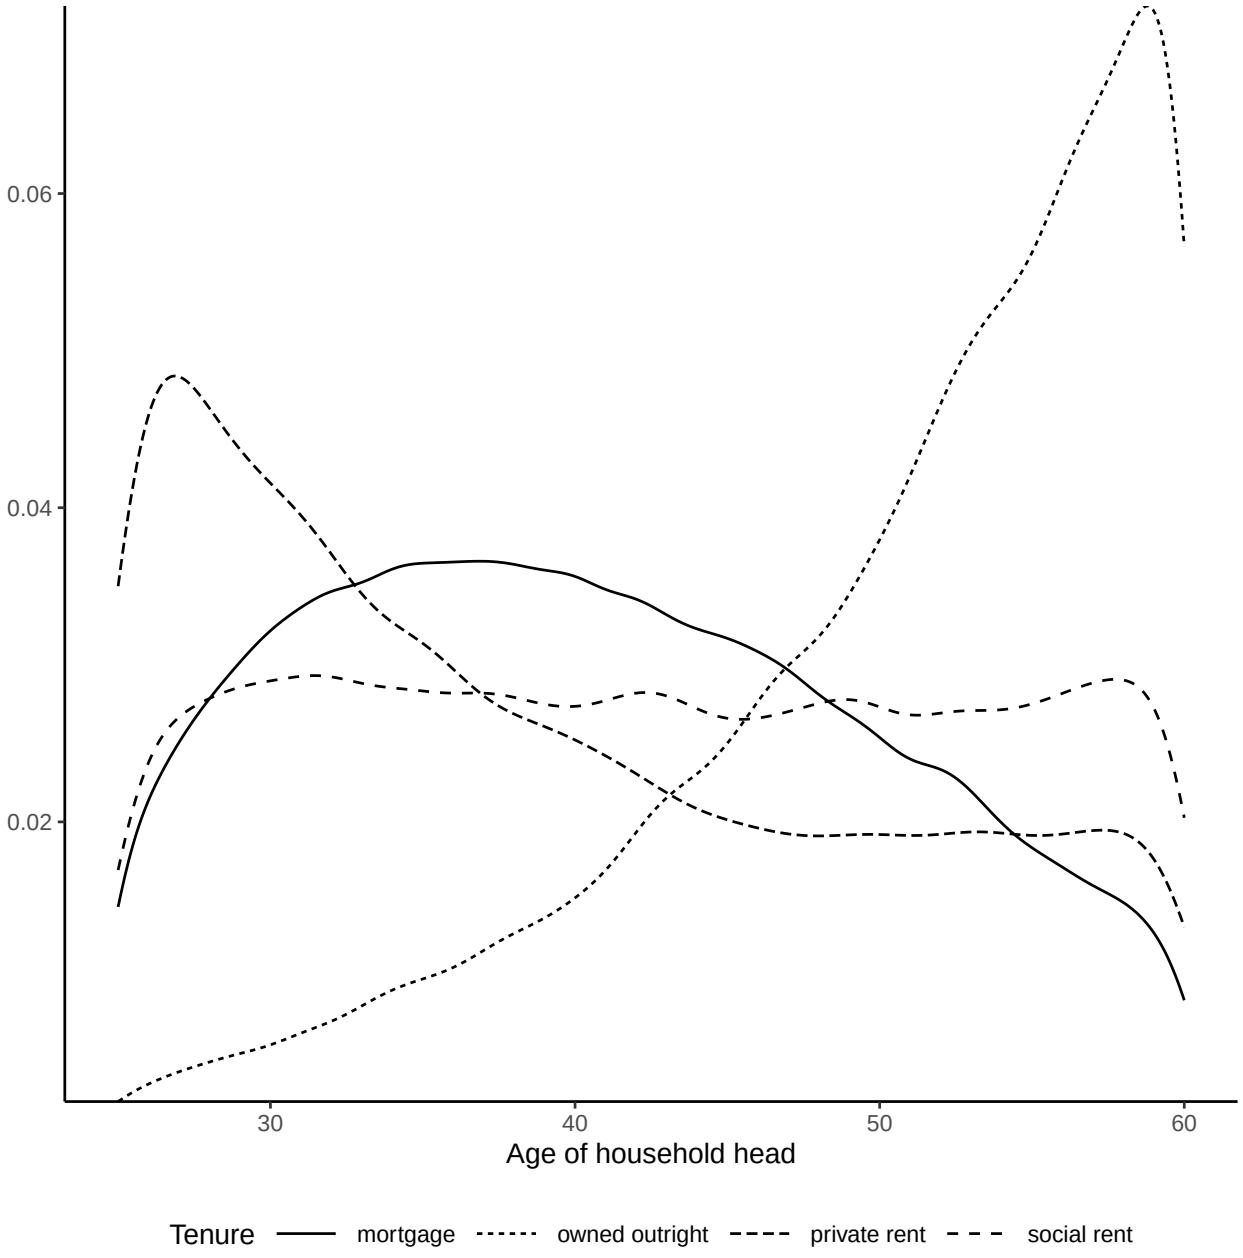


Figure III:
Relative income distribution by tenure

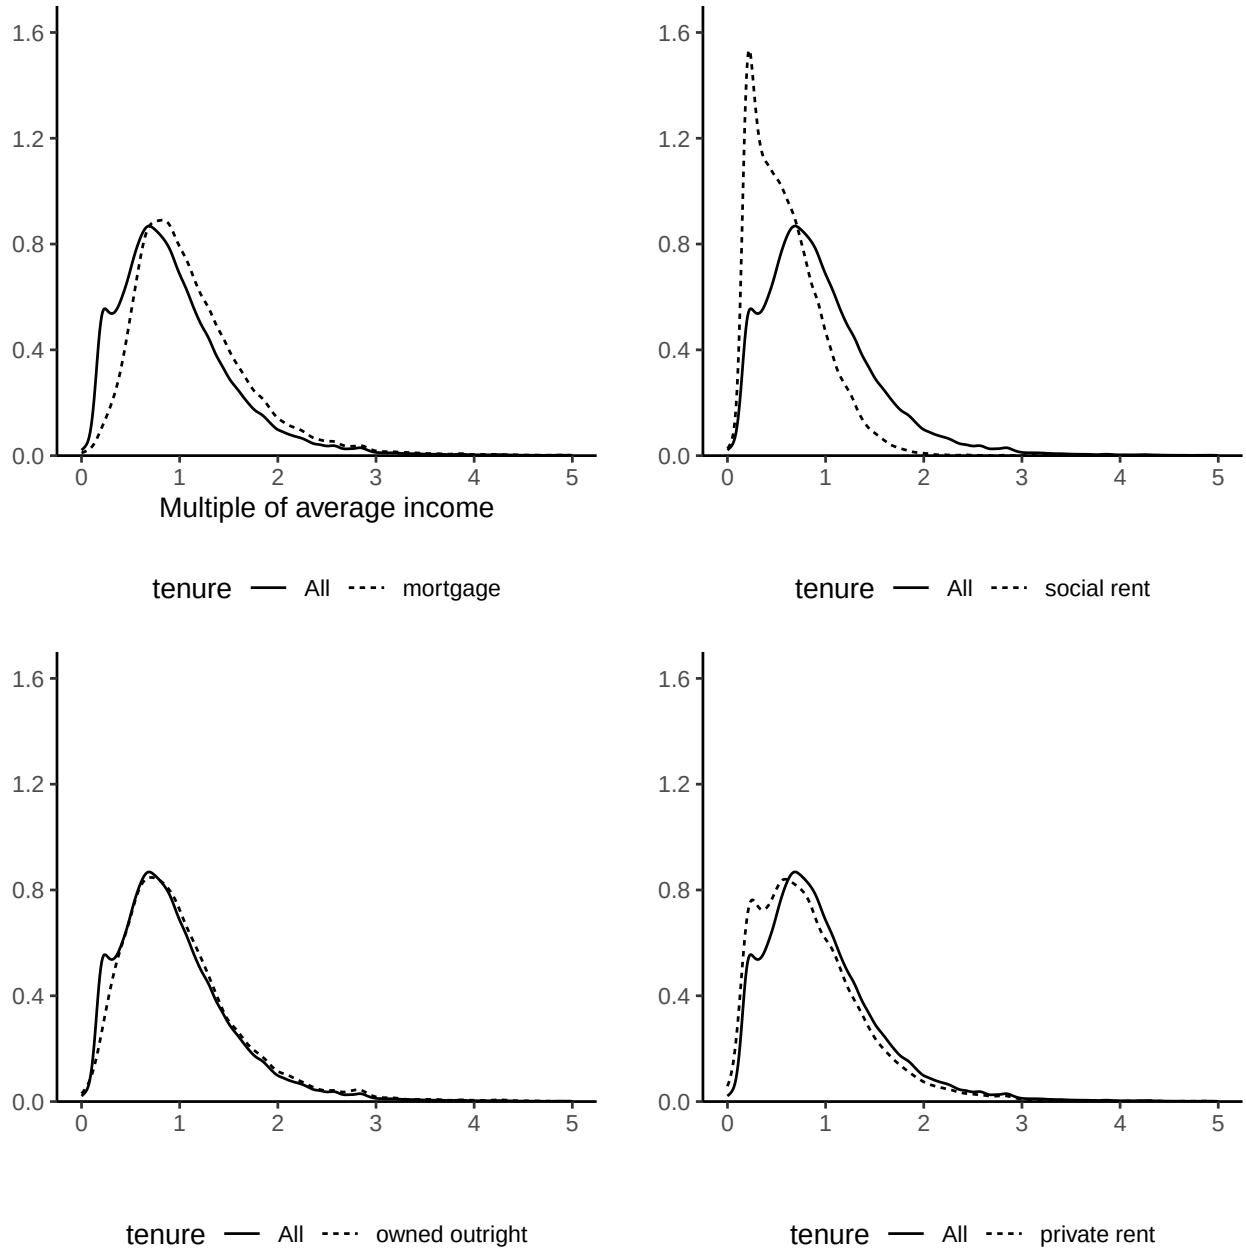


Figure IV:
Income composition by tenure

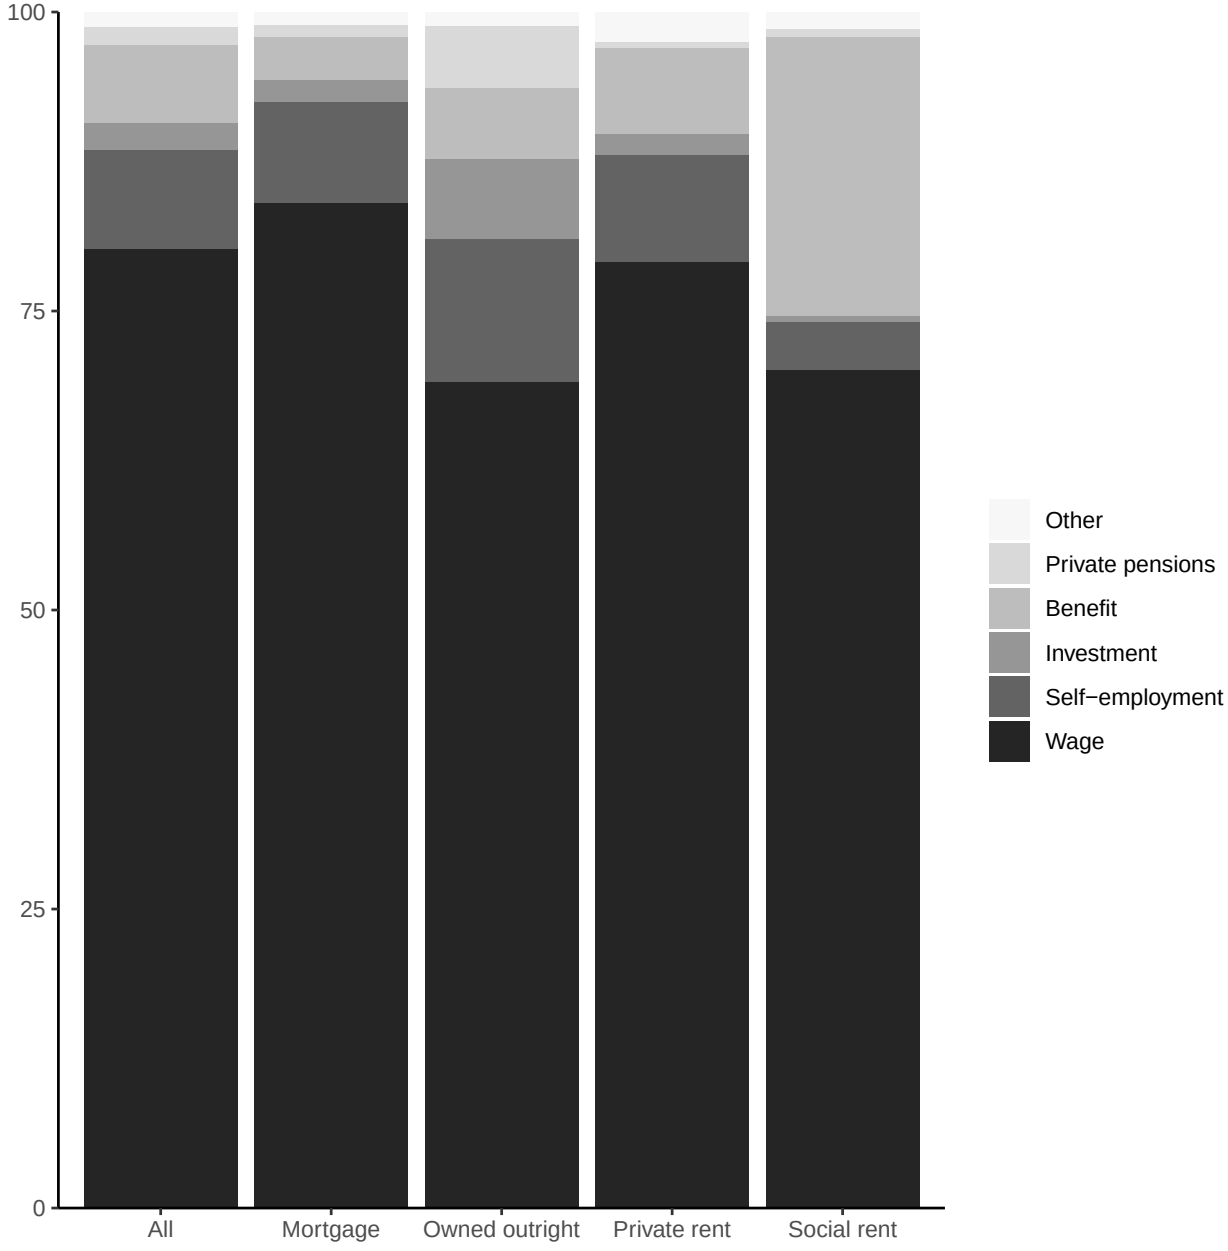


Figure V:
Post-compulsory schooling by tenure (smoothed)

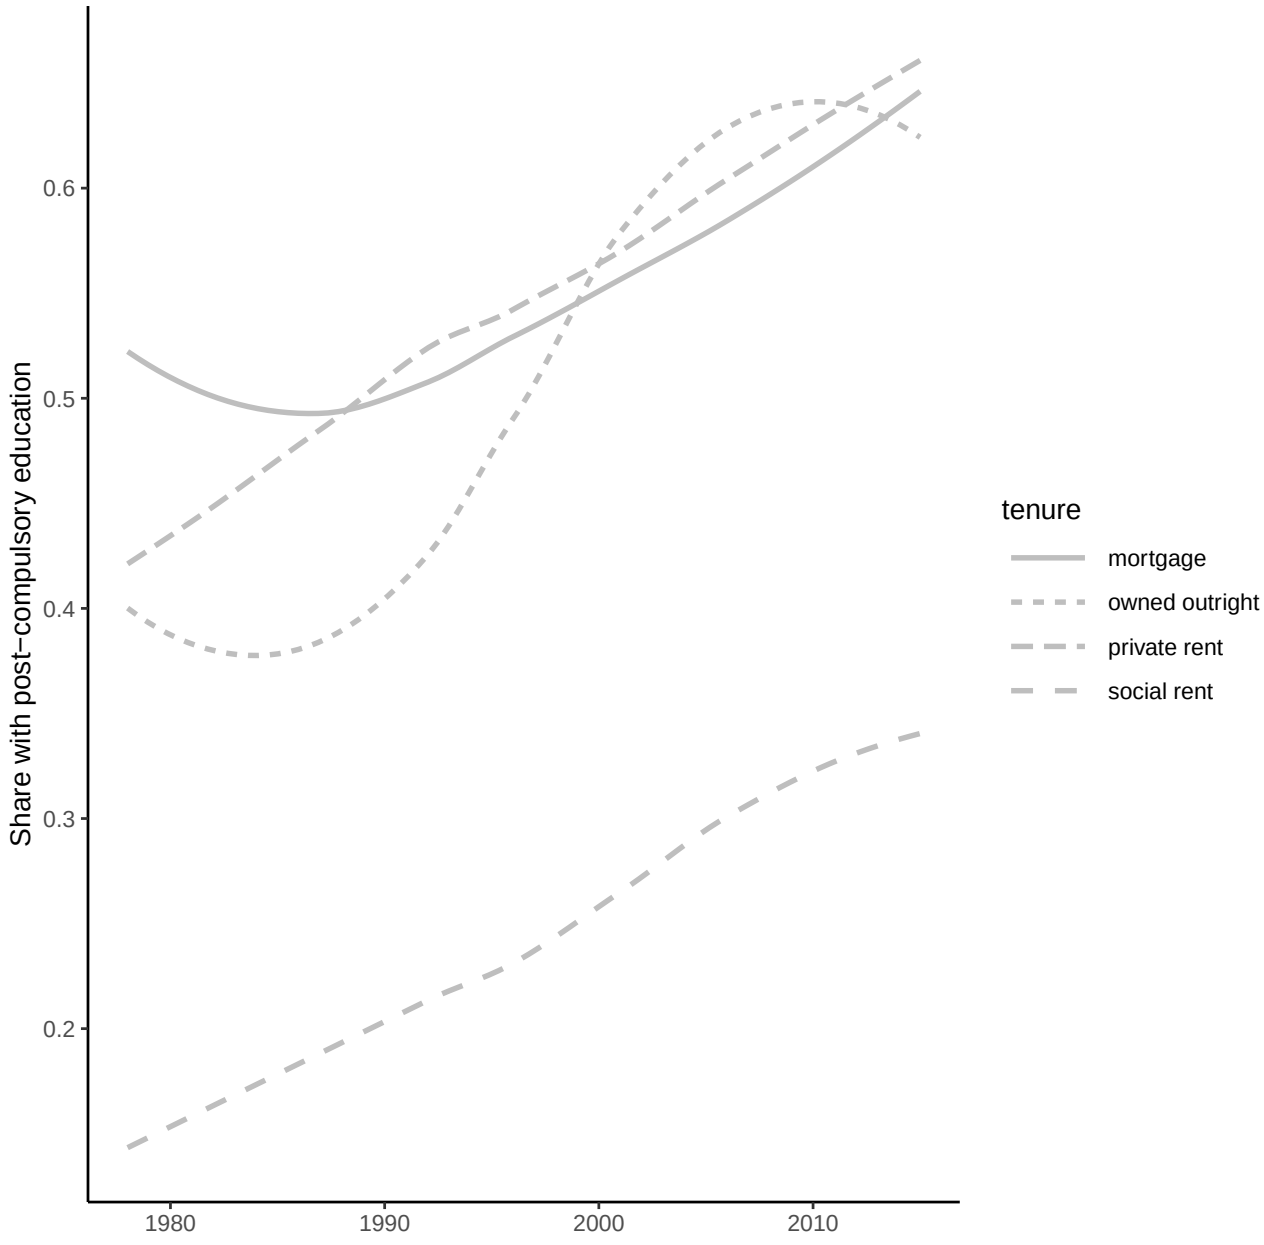


Figure VI:
Share of households eschewed spending because concern for credit availability

Bank of England/NMG survey 2006-2013

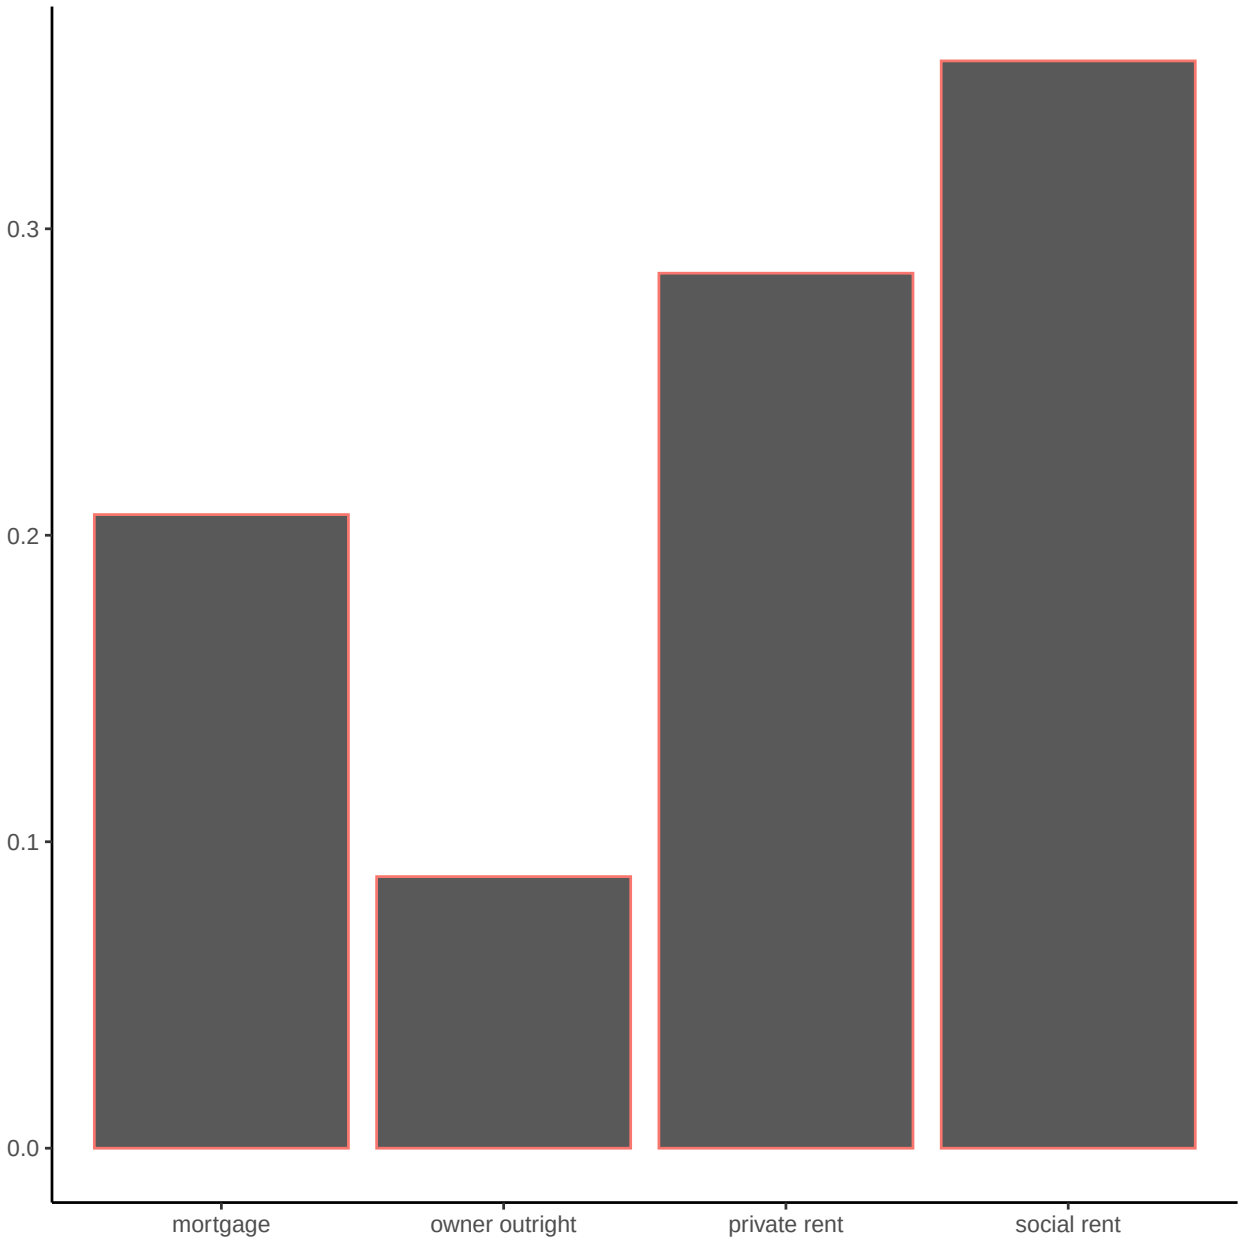


Figure VII:
Real equivalized private earnings growth by tenure

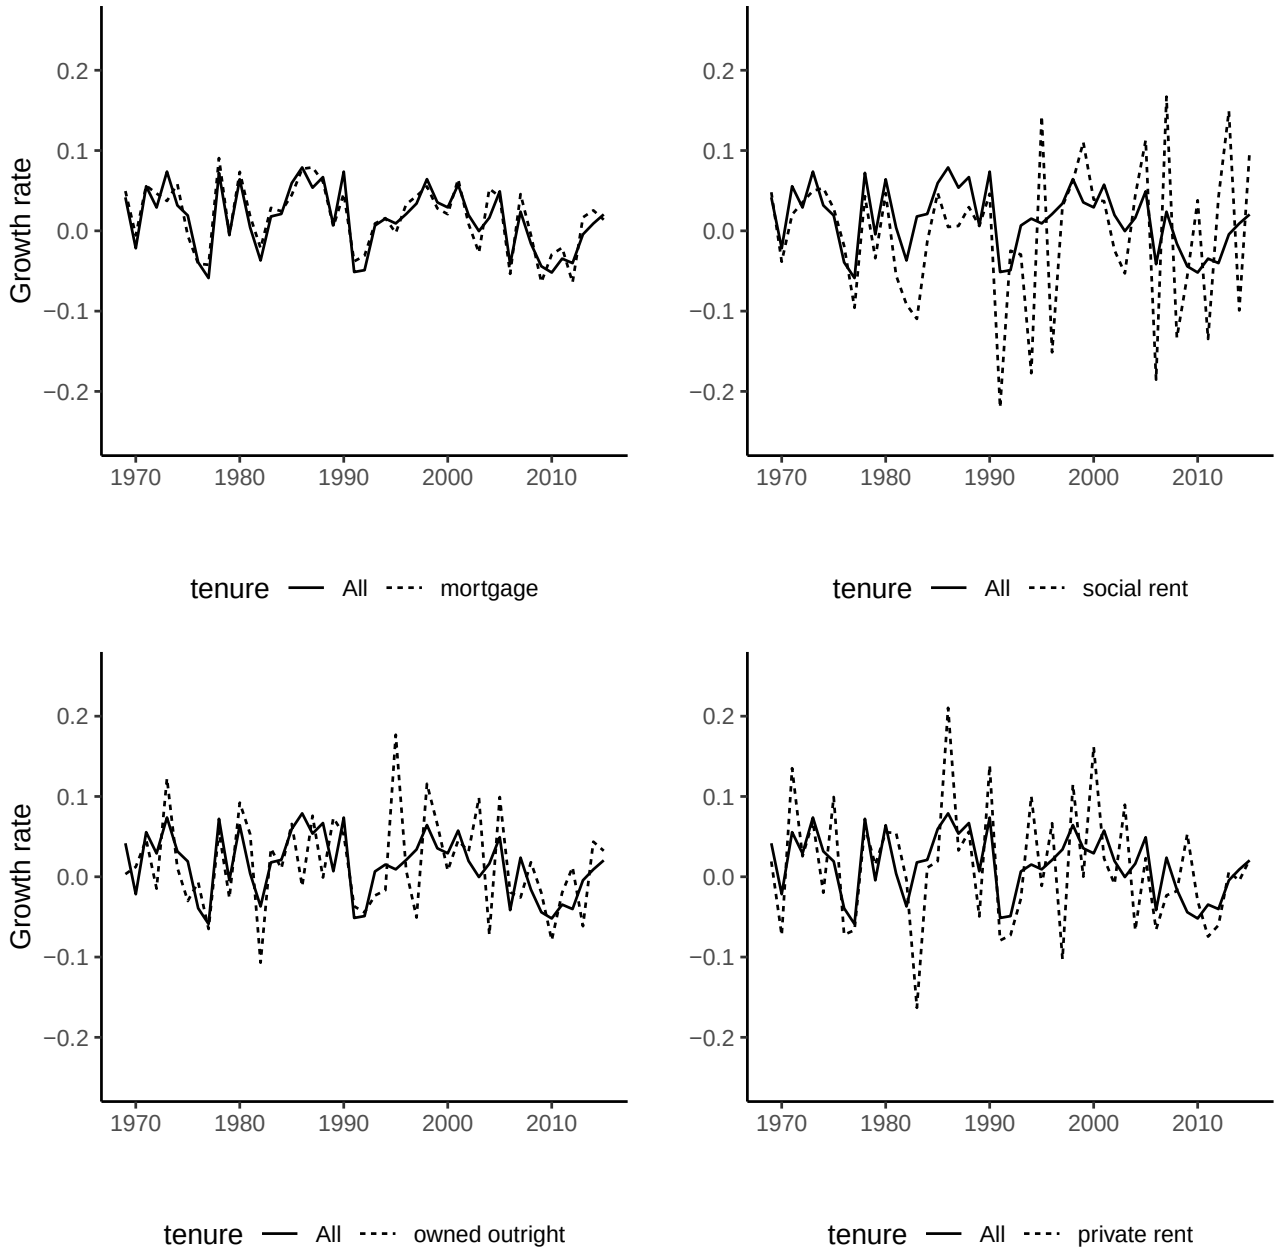


Figure VIII:
Real equivalized private income growth by tenure

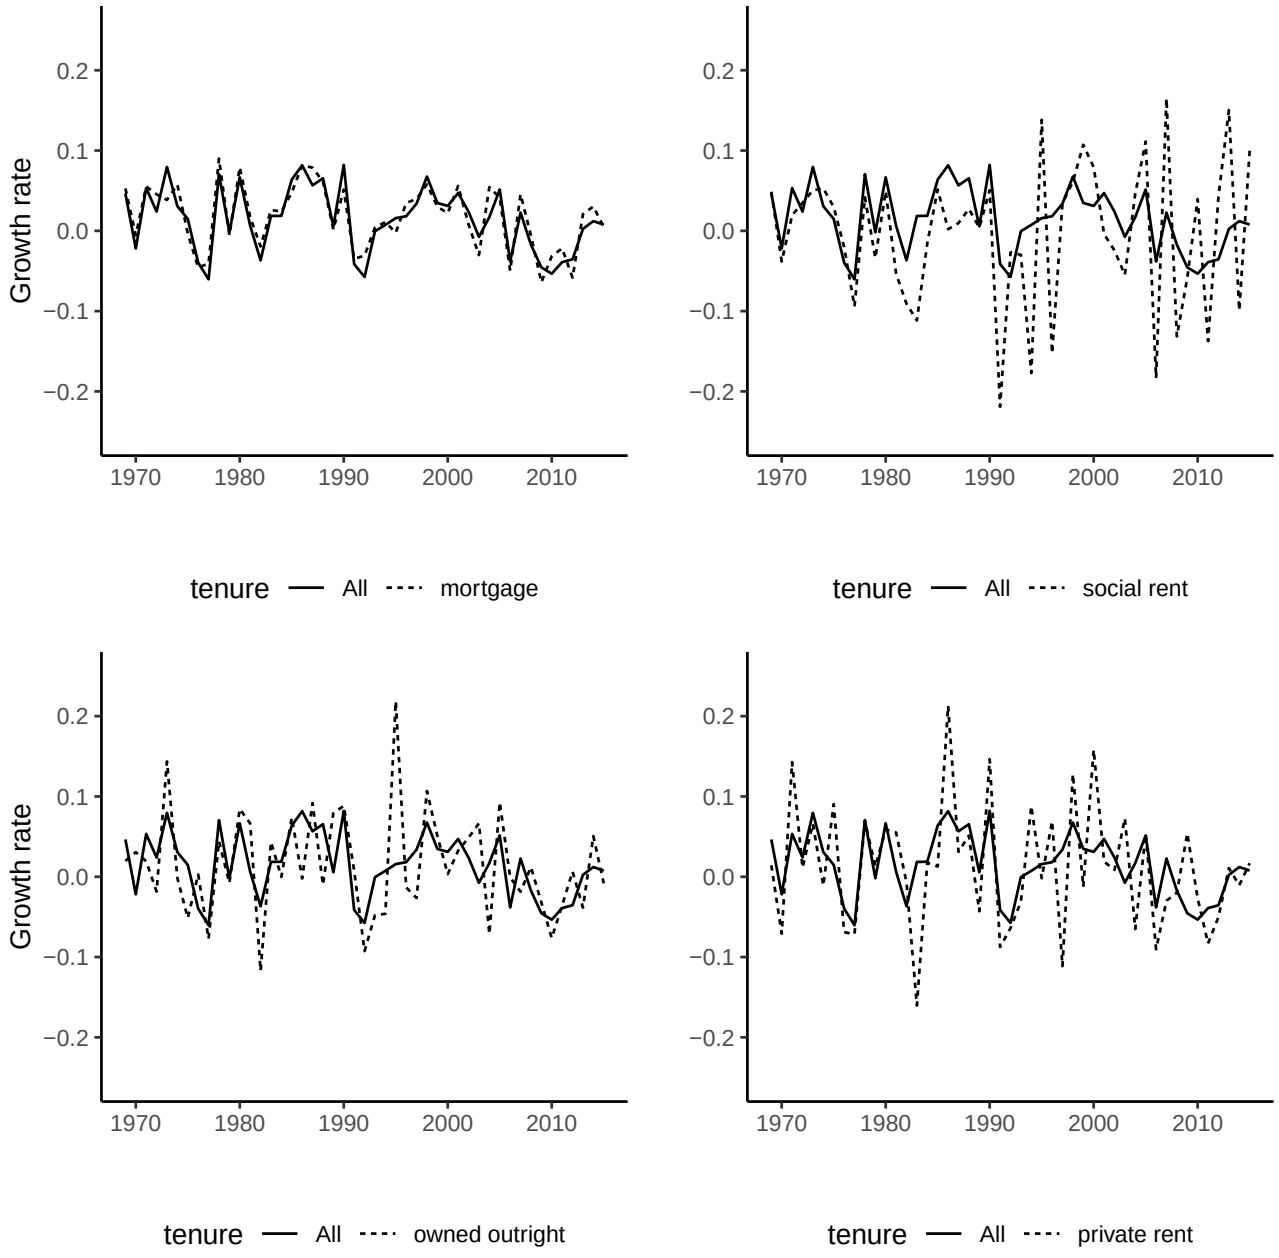


Figure IX:
Real equivalized disposable income growth by tenure

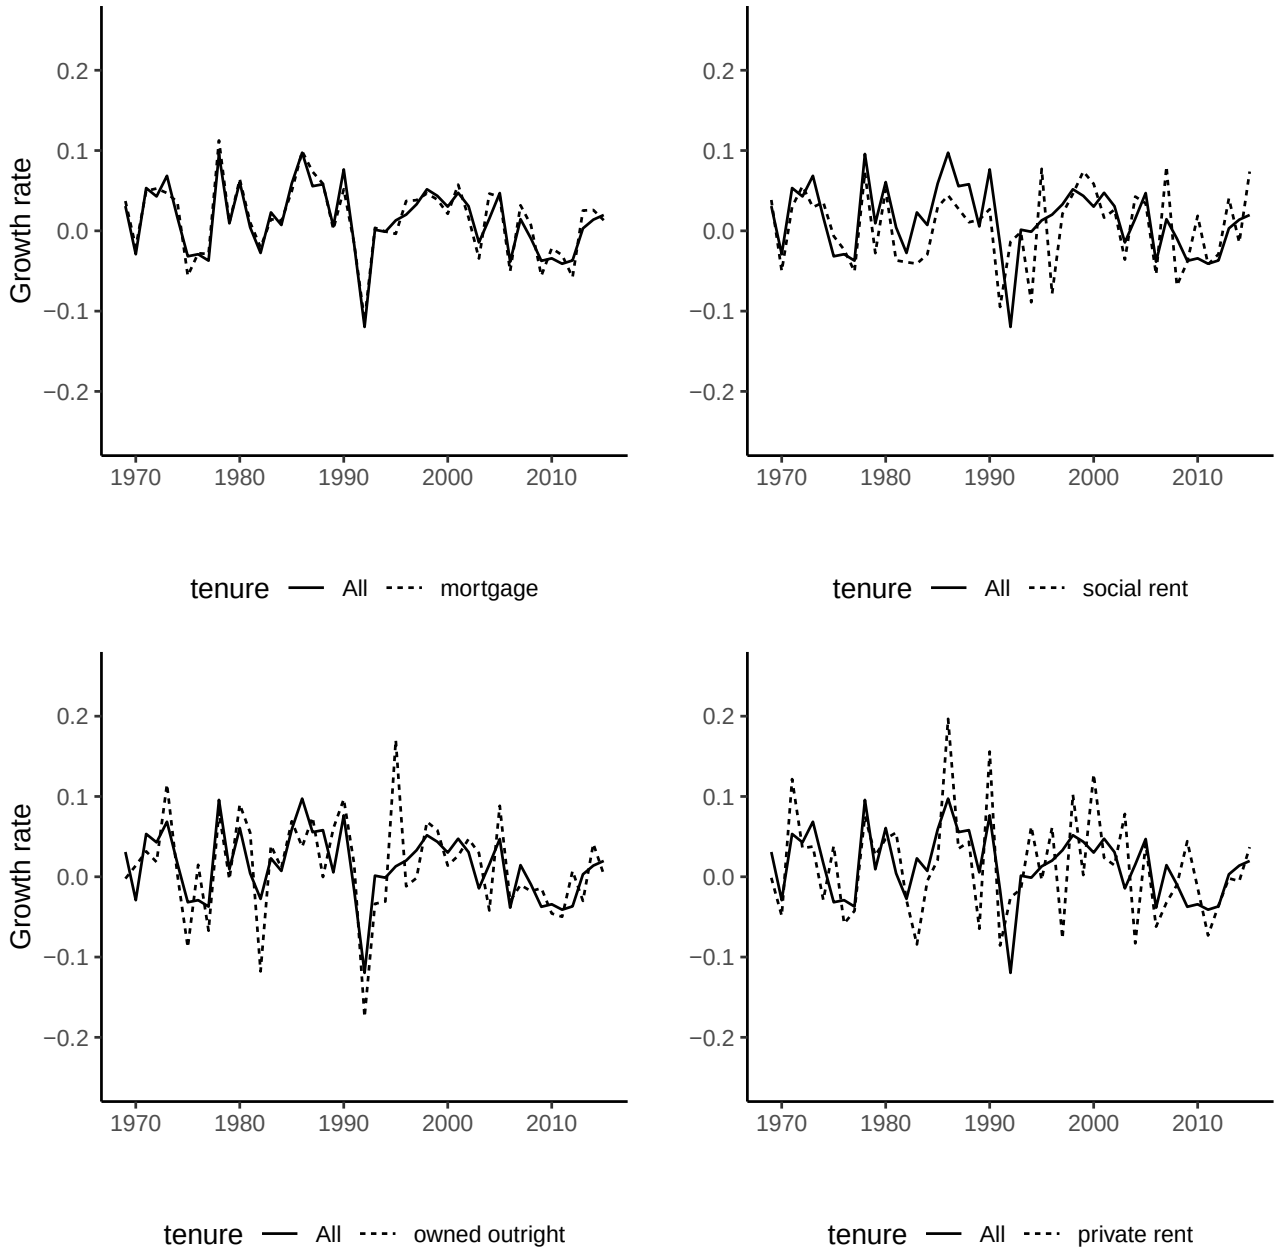
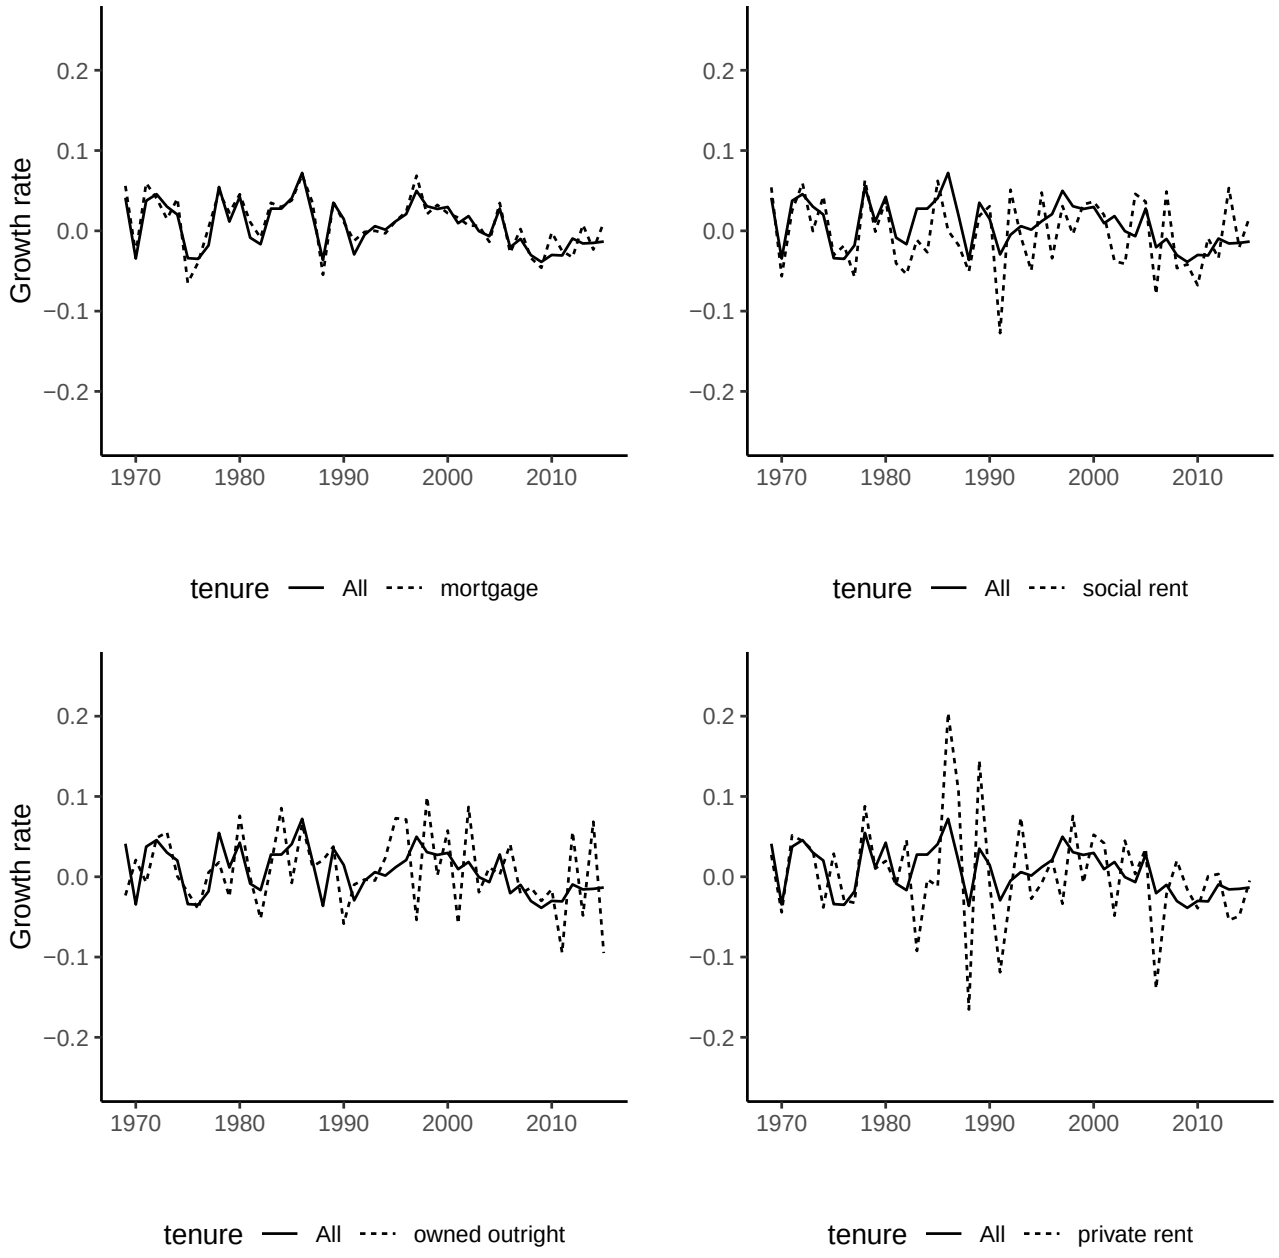


Figure X:
Real equivalized consumption growth by tenure



A Appendices

A.1 Dataset details

The raw data is available from the UK Data Service website (https://www.ukdataservice.ac.uk/) after registration.

Weighting

The Data Archive contains household weights for the yearly cross-section surveys starting in 1996. Weights prior to that were derived using data provided generously by the staff at the Institute for Fiscal Studies.

Equivalisation

Income and consumption variables take into account household size by dividing nominal values by modified household size. The formula is $\text{Size} = 1 * \text{FirstAdult} + 0.5 * \text{OtherAdults} + 0.4 * \text{Children}$.

Converting nominal into real values

The price index used for deflating nominal equivalized values is the Office for National Statistics 'Retail price index excluding housing' (CZBI).

A.2 Evidence on the representativeness of the UK data

There have been some studies concerning the aggregate representativeness of the UK survey data although they are already somewhat dated (Atkinson and Micklewright, 1983; Banks and Johnson, 1998). Johnson and McCrae (1998) found that survey income data matched quite well with corresponding national account aggregates. For the period 1985–1992, the FES survey caught about 90 % of the total gross income recorded in national accounts. Particularly the largest items, wage and benefit income, are fairly consistent with comparable national account data (95 % cover ratio).

Smaller income classes are less representative though. Self-employment income was on average 73 % of the comparable national account aggregate in 1985–1992. For investment income the ratio was only 52 %. Johnson and McCrae argue that the biggest reason is the undersampling of the very wealthy. There is also a timing issue. In particular, investment income seems to be reported with a lag of one year. This is probably because at the time of the survey respondents know last year's income but do not yet know their investment income for the current year. This problem is more prevalent with investment income (and to a lesser degree self-employment income) than other subsets because these income classes are more volatile and less amenable to forecasting.

In addition to the representativeness of levels, Johnson and McCrae contrast yearly growth rates of survey and national account aggregates. Correlation of growth rates for total gross income was 55 % for the study period. For benefits, the correlation was a very high 95 % and for earnings and self-employment income it was above 70 %. For investment income, however, the correlation was only 16 %. This is probably because of the lag effect discussed earlier – indeed, correlation with the one-year lead of survey-based investment income growth and the national account counterpart was 77 %.

Tanner (1998) examines the representativeness of the expenditure data. She finds that the survey-based measure represented on average 86 % of spending recorded in the national accounts in 1974–1992. The correlation in the growth rates of the two aggregates was 74 %.

Table A1: Smoothing of tenure shocks – regional aggregation
Late sample: 1990–2015

	Factor	Fiscal	Credit	Not smoothed
Mortgage	0.03***	0.12**	0.51***	0.34***
	(0.005)	(0.05)	(0.04)	(0.06)
Owned Outright	−0.02	0.24***	0.66***	0.12***
	(0.06)	(0.07)	(0.05)	(0.04)
Private rent	−0.02	0.27***	0.32***	0.43***
	(0.01)	(0.04)	(0.09)	(0.06)
Social rent	−0.0004	0.63***	0.16***	0.21***
	(0.01)	(0.03)	(0.05)	(0.06)
Observations	312	312	312	312
Adjusted R^2	-0.01	0.81	0.33	0.25

Note: *p<0.1; **p<0.05; ***p<0.01
The columns correspond, in order, to equations (3)–(6), where the cohorts are the four tenure groups (mortgage, owned outright, private rent, social rent) subdivided by region. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for auto-correlation within tenure-region groups and contemporaneous correlation between tenure-region groups. Data is annual from 1990 to 2015.

Table A2: Smoothing of tenure shocks – regional aggregation
Early sample: 1968–1989

	Factor	Fiscal	Credit	Not smoothed
Mortgage	0.01	0.23***	0.36**	0.40***
	(0.01)	(0.05)	(0.14)	(0.13)
Owned Outright	0.03	0.23***	0.53***	0.21*
	(0.02)	(0.04)	(0.09)	(0.12)
Private rent	0.004	0.30***	0.25***	0.44***
	(0.01)	(0.04)	(0.07)	(0.05)
Social rent	−0.001	0.39***	0.07	0.54***
	(0.01)	(0.07)	(0.12)	(0.07)
Observations	252	252	252	252
Adjusted R^2	-0.06	0.59	0.23	0.35

Note: *p<0.1; **p<0.05; ***p<0.01
The columns correspond, in order, to equations (3)–(6), where the cohorts are the four tenure groups (mortgage, owned outright, private rent, social rent) subdivided by region. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for auto-correlation within tenure-region groups and contemporaneous correlation between tenure-region groups. Data is annual from 1968 to 1989.

Table A3: Smoothing of tenure shocks – regional aggregation
Excluding potential outliers

	Factor	Fiscal	Credit	Not smoothed
Mortgage	0.01	0.12**	0.43***	0.44***
	(0.02)	(0.05)	(0.05)	(0.06)
Owned Outright	0.03	0.19***	0.57***	0.21*
	(0.04)	(0.04)	(0.12)	(0.11)
Private rent	−0.02	0.34***	0.38***	0.30**
	(0.03)	(0.06)	(0.10)	(0.13)
Social rent	0.01	0.44***	0.17*	0.38***
	(0.02)	(0.07)	(0.10)	(0.09)
Observations	414	414	414	414
Adjusted R^2	0.01	0.41	0.16	0.14

Note: *p<0.1; **p<0.05; ***p<0.01

The columns correspond, in order, to equations (3)–(6), where the cohorts are the four tenure groups (mortgage, owned outright, private rent, social rent) subdivided by region. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for auto-correlation within tenure-region groups and contemporaneous correlation between tenure-region groups. Data is annual from 1968 to 2015 where I drop observations with growth in private earnings larger than 10 percent.

Table A4: Smoothing of tenure shocks – regional aggregation
Using total non-housing expenditure

	Factor	Fiscal	Credit	Not smoothed
Mortgage	0.02***	0.18***	0.40***	0.40***
	(0.001)	(0.04)	(0.06)	(0.07)
Owned Outright	−0.003	0.24***	0.57***	0.20***
	(0.04)	(0.06)	(0.06)	(0.07)
Private rent	−0.005	0.29***	0.32***	0.39***
	(0.01)	(0.03)	(0.06)	(0.05)
Social rent	−0.0004	0.60***	0.17***	0.22***
	(0.01)	(0.02)	(0.05)	(0.07)
Observations	564	564	564	564
Adjusted R^2	-0.003	0.76	0.28	0.26

Note: *p<0.1; **p<0.05; ***p<0.01
The columns correspond, in order, to equations (3)–(6), where the cohorts are the four tenure groups (mortgage, owned outright, private rent, social rent) subdivided by region. I change the expenditure measure from non-durable expenditure to total non-housing expenditure. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for auto-correlation within tenure-region groups and contemporaneous correlation between tenure-region groups. Data is annual from 1968 to 2015.

Table A5: Probit for social housing tenure (1978–2015)

	Social housing
Constant	1.4050***
	(0.0921)
Age	−0.0792***
	(0.0038)
Age^2	0.0010***
	(0.0000)
A-levels	−0.7713***
	(0.0410)
College	−1.4049***
	(0.0626)
Trend	0.0348**
	(0.0154)
$Trend^2$	−0.0071***
	(0.0022)
$Trend^3$	0.0006***
	(0.0001)
$Trend^4$	−0.0000***
	(0.0000)
$Trend^5$	0.0000***
	(0.0000)
Age:A-level	−0.0006
	(0.0009)
Age:College	−0.0010
	(0.0015)
Age:Trend	−0.0012***
	(0.0001)
$Age:Trend^2$	0.0000***
	(0.0000)
A-level:Trend	0.0097***
	(0.0009)
College:Trend	0.0140***
	(0.0013)
AIC	137207.4699
BIC	137364.9033
Log Likelihood	−68587.7349
Deviance	137175.4699
Num. obs.	138633

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A6: Smoothing of tenure shocks – regional aggregation
Attanasio et. al. classification by time-*varying* cut-off

	Factor	Fiscal	Credit	Not smoothed
Unlikely social housing	−0.003	0.22***	0.33***	0.45***
	(0.04)	(0.03)	(0.06)	(0.06)
Likely social housing	0.01	0.32***	0.32***	0.34***
	(0.01)	(0.04)	(0.04)	(0.05)
Observations	222	222	222	222
Adjusted R^2	0.09	0.70	0.42	0.40

Note: *p<0.1; **p<0.05; ***p<0.01

The columns correspond, in order, to equations (3)–(6), where the cohorts are "likely" and "unlikely" social housing cohorts using the Attanasio et. al. classification method with a time-varying cut-off for likely housing tenure. The annual sample is 1979–2015 because I lack information on education for earlier years. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects and control for change in sample mean age by cohort in subsequent years. Standard errors allow for autocorrelation within tenure-region groups and contemporaneous correlation between tenure-region groups.

Table A7: Smoothing of tenure shocks – regional aggregation
Attanasio et. al. classification by time-*invariant* cut-off

	Factor	Fiscal	Credit	Not smoothed
Unlikely social housing	−0.005	0.19***	0.35***	0.47***
	(0.04)	(0.03)	(0.08)	(0.10)
Likely social housing	−0.001	0.36***	0.20***	0.45***
	(0.004)	(0.03)	(0.05)	(0.04)
Observations	222	222	222	222
Adjusted R^2	0.05	0.74	0.32	0.49

Note: *p<0.1; **p<0.05; ***p<0.01

The columns correspond, in order, to equations (3)–(6), where the cohorts are "likely" and "unlikely" social housing cohorts using the Attanasio et. al. classification method with a time-invariant cut-off for likely housing tenure. The annual sample is 1979–2015 because I lack information on education for earlier years. The covariates labeled by tenure refer to the coefficient on the interaction of the specific tenure dummy with private earnings growth. All regressions include year and tenure-region fixed effects with a time-varying cut-off for likely housing tenure. Standard errors allow for auto-correlation within tenure-region groups and contemporaneous correlation between tenure-region groups.

Figure A1:
Probability density of predicted social housing tenure

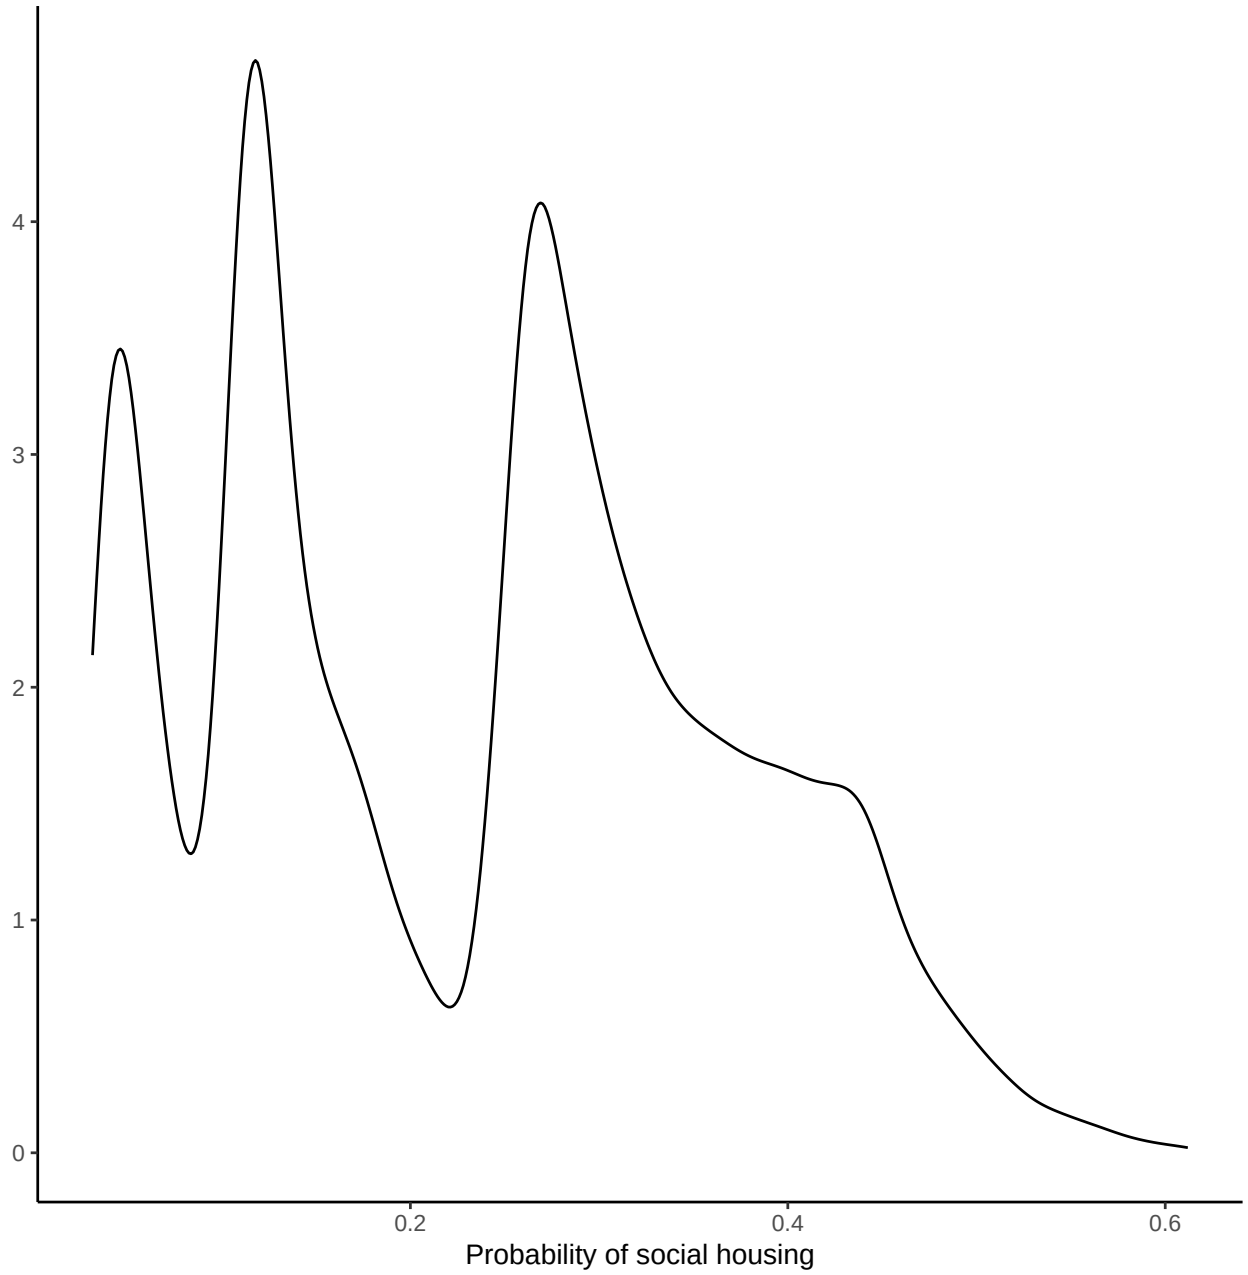
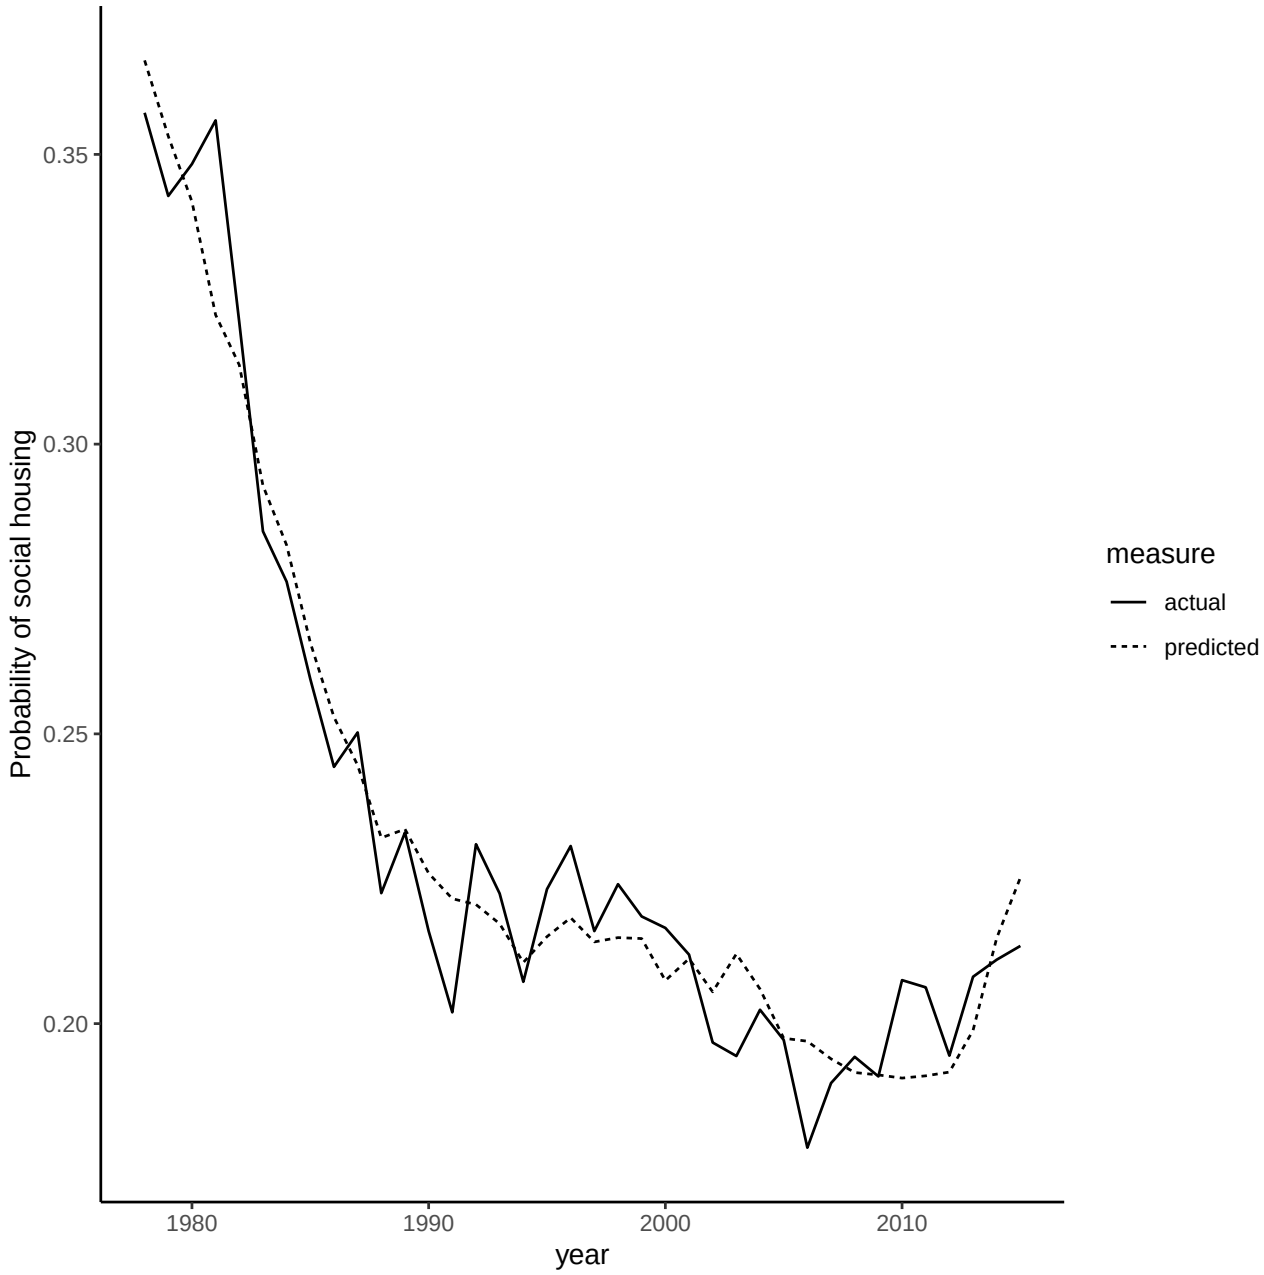


Figure A2:
Social housing share by year – probit model vs. data



Conclusions

This thesis has contained three essays on household finance topics. Because each essay contains their own conclusions that summarize the main findings, I conclude here with more general remarks on the topics in this thesis, how my thesis ties to the field of household finance, and further directions for research.

A general theme in my research is that household heterogeneity matters. My findings suggest that households differ in their attention to household finances, the degree of intra-household financial pooling, the degree of self-control problems, and in the insurance mechanisms they use. My findings complement earlier research in household finance that has documented significant heterogeneity in various dimensions. Household heterogeneity poses a challenge for interventions that policymakers may wish to undertake to promote desireable household financial outcomes, because one-size-fits-all policies are unlikely to work. Therefore, the design of, for instance, appropriate default options on financial products and financial literacy education need to account for household heterogeneity to be effective.

A second intriguing theme is possible heterogeneity in the decision-making processes within individual households over time or in different contexts. For instance, regarding the issue of self-control problems, is the degree of present bias fixed or can changes in the context or framing of the problem have an impact on the activation of a short-term or long-term mindset? Further research in household finance will advance our understanding of the heterogeneity in household decision making with, hopefully, lessons on how to improve financial well-being.