Essays in Social Interactions and Financial Decision-Making

ENGELS, CHRISTIAN

Doctoral Thesis

Essays in Social Interactions and Financial Decision-Making

Author:
Christian Engels

Supervisors:
Prof. Dennis Philip
Prof. Abderrahim Taamouti

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

in the

Department of Economics and Finance
Durham University Business School
Durham University

November 2020

Essays in Social Interactions and Financial Decision-Making

Christian Engels

Abstract

This thesis contributes to the literature on social interactions and financial decision-making. The first essay builds on previous research that highlights how subjective well-being deteriorates with low financial resources by showing that financial knowledge intensifies this relationship. More specifically, the evidence suggests that the financially literate, with the skills and abilities to manage finances, are subject to disproportionate deterioration in subjective well-being levels when financially distressed. In this respect, the essay suggests for public policy to consider wider outcomes negatively influenced by financial literacy when implementing measures to increase its overall levels in society. The second essay contributes to the economics literature on social identity by highlighting the relationship between social norms and mental health. It shows that individuals with highly positive views on the welfare state and benefit recipients report a higher prevalence of mental health problems. In societies where work is valued as a normative good, and in which benefit recipients are characterised as benefit scroungers, expressing such positive welfare attitudes can constitute deviations from social norms. Theoretical analysis suggests that social sanctioning for expressing these welfare attitudes can explain this relationship and implies that policies to destigmatise welfare can improve public health. The third essay builds on evidence that rates of investment participation are generally low and inspects the role of peers. Empirical models of peer effects are used to quantify the contribution of the propensity to choose peers like oneself. Contrary to what is intuitively expected, the results indicate that the influence of peers depresses the overall participation rate. Overall, the thesis brings to light the different mechanisms that govern social interactions and illuminates new dimensions to understanding financial decision-making. The findings contribute to recent discussions on financial well-being and financial inclusion.

Supervisors: Prof. Dennis Philip and Prof. Abderrahim Taamouti

Contents

Declaration

The work in this thesis is based on research carried out in regular consultation with academic supervisors at the Department of Economics and Finance of Durham University Business School, Durham University, UK. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is the sole work of the author unless referenced to the contrary in the text.

List of Figures

List of Tables

Acknowledgements

I could not have completed the journey towards this work without the unfailing support by my supervisors, Prof. Dennis Philip and Prof. Abderrahim Taamouti. My deepest gratitude goes to both for sharing their insights and expertise and time with me – and for encouraging me during those stretches I found hardest.

In particular, I would like to thank Prof. Dennis Philip, who consistently has gone beyond the call of duty in my supervision, and who empowered me, with patience, graciousness and the occasional push to excel, to shape my future. It has been a gift given freely that I cannot repay, but one that I will always remember.

My final words of gratitude go to Ms. Annabel Logan. "Is this it, you writing essays until the end of time?" you ask in part disbelief, part jest. It might just be, and I am blessed that you love me all the same. I dedicate this work to you.

This work was further supported by the Economic and Social Research Council [NEDTC training grant number ES\J500082\1].

Introduction

This thesis presents three essays on social interactions and financial decision-making. This thesis's first essay is titled "Financial Literacy and Subjective Well-Being" and relates to the academic literature that documents a preponderance of benefits of financial literacy. For example, financial literacy enables the better management of planning for retirement and greater effectiveness in accumulating wealth (Klapper & Panos, 2011; Lusardi & Mitchell, 2007a, 2007b, 2011); it increases engagement with financial markets (Balloch, Nicolae & Philip, 2015; Van Rooij, Lusardi & Alessie, 2011; Yoong, 2011); and it empowers individuals to overcome the negative effects of economy-wide shocks (Klapper & Panos, 2013). However, in the US context of stagnating household income growth (Piketty, Saez & Zucman, 2018) and declining absolute income mobility (Chetty et al., 2017), this chapter hypothesises that financial knowledge makes the lower chances of financially distressed individuals to realise their financial aspirations in life more apparent, thereby reducing individual subjective well-being, also known as happiness or life satisfaction.

Accordingly, this first essay establishes that financial literacy can impose significant costs on the financially distressed. In this essay, financial distress is defined as lacking the resources to meet basic consumption needs, relating to food, health care, housing and utilities. Throughout the analysis, we draw on a representative survey of the US population from 2016 and show that for individuals with higher levels of financial literacy, financial distress is strongly associated with disproportionate deterioration in subjective well-being. The relationships we observe further intensify when individuals perceive that they are falling short of their financial goals. This is associated with depletions in subjective well-being of 35% relative to population-wide baseline levels. The findings suggest that a failure to achieve greater financial aspirations that come with financial literacy can be strongly detrimental to individual well-being. Our findings imply important considerations for financial education programmes when employed for poverty alleviation, as deteriorations in subjective well-being that can arise from higher financial knowledge may influence individual motivations to better their circumstances.

This thesis's second essay is titled "Social Norm Enforcement and Mental Health". Therein, we add to the body of research that shows the influence of individual characteristics, circumstances and wider economic conditions on mental health (Bridges & Disney, 2010; Currie, Duque & Garfinkel, 2015; Gathergood, 2012; Grip, Lindeboom & Montizaan, 2011). We do so by investigating the effect of social norm enforcement (Bernhard, Fehr & Fischbacher, 2006; Goette, Huffman & Meier, 2006) on mental health. As Posner (1997) notes: "Norms are also enforced by expressions of disapproval, by ridicule, and in extreme cases, by ostracism. The efficacy of the milder 'sanctions' lies in their implicit threat of ostracism, that is, of refusal of advantageous transactions." (p. 366) This essay therefore investigates the potential psychological repercussions these enforcement means can have and asks whether the enforcement of social norms harm mental health.

This essay, then, provides evidence that deviating from social norms can indeed harm mental health. In societies that assign a normative value to work (entailing prescriptions such as"you *should* work" or "life means standing on your own two feet"), expressing positive attitudes to the welfare state or benefit recipients (such as "the government should spend more money on benefit") can constitute a deviation from social norms that prevail in society. Thus, such expressions can result in social sanctioning, which can induce deterioration in mental health. Using a set of representative surveys that measure such welfare attitudes in the British population, we show that more positive attitudes are strongly associated with a higher prevalence of mental health problems. In order to provide evidence on the mechanism of social norm enforcement, we exploit changes in the UK's political parties elected into office and show that the relation intensifies under conservative governments, known to be tough on welfare. Such changes in political environments can increase the degree of social sanctioning of positive welfare attitudes by reducing its costs. Further, we find that the relationship is more pronounced for women than for men, especially under conservative rule. Our findings suggest a real societal cost to deviating from mainstream social norms and imply that destigmatising welfare receipt in societies in which engaging in paid employment is a social norm can function as a policy measure to improve public health.

This thesis's third essay is titled "Endogenous Peer Choice and Investment Participation". The motivation for this third essays lies in the findings that societal differences in financial inclusion have significant implications for wealth inequality (Favilukis, 2013), as well as labour market and economy-wide stability (Epstein & Shapiro, 2020). As individuals with greater incomes and education are more likely to marry one another (Nakosteen, Westerlund & Zimmer, 2004) and, at the same time, are more optimistic about future macroeconomic conditions (Das, Kuhnen & Nagel, 2020), the possibility exists that financial behaviours that build wealth, such as the use of investment products, perpetuate differently in groups of societies with homogeneous intra-group characteristics, with subsequent differential impacts on financial inclusion.

Accordingly, the third essay in this thesis studies how the endogenous choice of peers affects investment participation through the social interactions channel. We use data from the UK's Understanding Society and British Household Panel Survey to uncover expansive household networks. By observing with whom respondents live throughout the sample period, we are able to obtain precise measurements of individual characteristics and social ties. We find that peer choice accounts for approximately 25% of the social interactions channel of investment participation. We provide evidence that this finding arises from significant tendencies of individuals to choose peers similar to themselves. In a counterfactual exercise, we shut down the peer effect in our Bayesian econometric model and observe that investment participation increases by up to 2.5%, suggesting that non-invested individuals reinforce their inclinations of non-participation among themselves. The results of this essay call for further policy efforts to promote the known benefits of investments for wealth accumulation in the cross-section of society that holds none, in order to counteract the self-sustaining group tendencies of not investing.

From here on, the thesis proceeds by presenting the essay "Financial Literacy and Subjective Well-Being" in Chapter 1, "Social Norm Enforcement and Mental Health" in Chapter 2 and "Endogenous Peer Choice and Investment Participation" in Chapter 3. The final section offers concluding remarks.

1

Financial Literacy and Subjective Well-Being

Highlights

∗ Financial distress is associated with an approximately 13% drop in aggregate subjective well-being (SWB) against predicted population-wide levels obtained from the empirical model. The relationship between financial distress and SWB intensifies for high financial knowledge individuals. For instance, for the highly knowledgeable, financial distress is associated with a drop of 18% in aggregate SWB versus predicted overall levels.

∗ When high financial knowledge individuals report that they fail to follow through with their own financial goals, the relationship intensifies further to a drop of 35% in aggregate subjective well-being. Controlling for a wide range of individual- and household-level characteristics, this suggests the presence of higher financial aspirations for these individuals.

∗ These results can be interpreted through the lens of reference-dependent preferences, such that the decision to invest in the acquisition of financial knowledge was taken given prospects of greater financial well-being. Financial distress is then evaluated against this prospect, resulting in increased losses in subjective well-being.

1.1 Introduction

> People may adapt to misery and hardship, and cease to see it for what it is. People do not necessarily perceive the constraints caused by their lack of freedom; the child who is potentially a great musician but never has a chance to find out will not express a lack of life satisfaction. Whole groups can be taught that their poor health or their lack of political participation are natural or even desirable aspects of a good world. (Deaton, 2008, p. 69)

Recent years have brought sobering developments to the financial prospects of US households. Income growth has stagnated from 1980 to 2014 for the bottom 50% (Piketty et al., 2018); 90% of children born in 1940 earn more than parents versus 50% in 1980 (Chetty et al., 2017). Households financial distress is prevalent, with approximately 40% of households experiencing one or more lack of basic consumption necessities (Karpman, Zuckerman & Gonzalez, 2018). At the same time, in the last two decades, happiness and life satisfaction in the US have consistently declined (Helliwell, Layard & Sachs, 2019). Given the well-documented finding in the literature that lower incomes are related to decreased subjective well-being (Deaton, 2008; Kahneman & Deaton, 2010; Sacks, Stevenson & Wolfers, 2012), we investigate the role of financial literacy in this context. In particular, we ask: does higher financial literacy decrease individual life satisfaction in the face of adverse financial circumstances, as the lower chances of realising one's financial aspirations in life are felt more acutely?

Specifically, we investigate whether greater financial knowledge is associated with lower subjective well-being, popularly referred to as happiness or satisfaction, for individuals facing financial distress. We propose that higher levels of financial knowledge can further erode subjective well-being when struggling to pay bills and everyday expenses, as greater financial knowledge could have been acquired in response to higher financial aspirations in individuals. Financial distress then results in a widened gap between desired financial prospects and the actual realities of the individual financial situation, thereby disproportionally eroding individuals' life satisfaction, as well as wider dimensions of how respondents evaluate their lives.[1]

Studying these relationships, our analyses draw on rich cross-sectional data representative of the US adult population, obtained from the 2016 National Financial Well-Being Survey fielded and made available by the Consumer Financial Protection Bureau (2017). The data elicit in detail respondents' levels of subjective well-being, financial literacy and financial distress, as well as a preponderance of demographic characteristics. Our main outcome variables relate to subjective well-being. Three survey questions measure the subjective well-being dimensions of 1) life satisfaction, 2) optimism about the future and 3) the belief that work yields success. Respondents convey their answers on a 7-point Likert scale, which for construction of an aggregate scale, we map to the integers from 0 to 6, where greater values indicate higher levels of well-being. For each respondent, we sum all three answers and thus obtain the aggregate subjective well-being score, ranging from 0 to 18. Our main explanatory variable of interest is the financial knowledge score by Knoll and Houts (2012), a composite measure capturing respondents' understanding of a wide range of financial concepts. We further construct an indicator variable capturing financial distress in the last 12 months from questions about experiences of worrying about running out of food or actually doing so; being unable to afford medical treatment; and not being able to afford housing or paying for utilities. Financial distress is then defined as answering "often" or "sometimes" to at least one of these questions. Further, we construct variables capturing the respondents' marital status, health, education, financial depend-

[1]This is mechanism is consistent with the recent research that distinguishes two aspects of subjective well-being. Kahneman and Deaton (2010), for instance, show that high income improves evaluation of life but not emotional well-being.

ants, age, gender and household income for inclusion as controls in our empirical analysis. Exact variable definitions are provided in the Appendix. Our analysis sample contains 6394 observations.

In our empirical analysis, we first estimate multivariate linear and ordered probit regressions that establish the baseline that experiences of financial distress are associated with decreased aggregate subjective well-being levels of approximately 13% versus predicted overall levels, the average of the model-fitted subjective well-being levels of respondents. The next step involves separating respondents by high (low) financial knowledge levels, defined as scores greater (lower) than the sample average, which reveals substantial heterogeneity across financial knowledge levels: the negative association changes to -18% and -11%, respectively, for the high and low financial knowledge group. The effects are found to be statistically different. Importantly, these associations are stable across all individual dimensions of subjective well-being: life satisfaction, optimism about the future and the belief that work today yields success in the future. All these measures deplete more strongly for high financial knowledge individuals in financial distress.

Next, we test for evidence that greater financial aspirations can provide an explanation for the negative relation between financial knowledge and subjective well-being for financially distressed individuals. To do so, we rely on two survey questions that elicit how well respondents feel that the statements "I follow-through on financial goals I set for myself" and "I follow-through on my financial commitments to others", respectively, apply to themselves. The possible answer choices are on a 5-point Likert scale ranging from "completely" to "not at all". To identify greater financial aspirations, we derive two further variables that capture which respondents indicate that the respective statements apply "not at all" or "very little". Ceteris paribus, controlling a wide range of important individual characteristics, failing to meet one's own financial goals is indicative of these goals being greater to begin with. Analogous reasoning applies to financial commitments towards others, and together these variables capture internal (towards oneself) and external (towards others) dimensions of financial aspirations.

Then, to test the notion that financial aspirations can help explain why high financial knowledge intensifies experiences of financial distress, we regress subjective well-being levels on the interaction between the high/low financial knowledge indicator and the variables capturing internal and external dimensions of financial aspirations, respectively. Two important observations emerge from this analysis. First, our findings reveal that for the highly financially literate, failure to meet internal financial aspirations is associated with a decrease in aggregate subjective well-being of 35% against population-wide baseline levels – an approximately twofold increase relative to the baseline estimate regarding the mediating role of financial knowledge. All these associations are highly statistically significant, vis-a-vis the baseline where high financial knowledge individuals do not fail to meet their internal aspirations. Second, in the case of external financial aspirations, failure to follow through does not intensify the relationship between financial knowledge, as indicated by p-values consistently greater than 0.1 for the corresponding tests.

In a set of additional analyses, we address endogeneity concerns relating to potentially omitted drivers introducing bias in the observed relationships. First, we employ an instrumental variable approach using the method by Lewbel (2012). The results corroborate the conclusions drawn in our empirical analysis that financial distress disproportionally decreases subjective well-being levels for the highly financially literate. To reduce potential biases in estimating the association between financial knowledge and subjective well-being stemming from the endogenous acquisition of financial knowledge (Lusardi, Michaud & Mitchell, 2017), we further use a sample obtained through a standard propensity score matching approach, balancing respondents' covariates across high and low levels of financial knowledge. We test the stability of our baseline results and the proposed mechanism and corroborate the conclusions drawn in the empirical analysis for the case of individual life satisfaction. Confirming that the observed associations hold in this group of respondents is crucially important, as respondents outside of financial distress will be significantly less likely to select into financial education programmes designed for poverty alleviation. In this way, we directly establish the relevance of our findings for financial educators who likely will be in close contact with such individuals and may consider the implications of the associations we uncover for curriculum design.

Our analysis proceeds as follows. Section 1.2 discusses related literature. Section 1.3 describes our data and variables; Section 1.4 presents the results of our empirical analysis; Sections 1.5 and 1.6 provide the results of our additional analyses; Section 1.7 discusses the limitations of this study and possibilities for future research, while Section 1.8 concludes.

1.2 Related Literature

Subjective well-being is distinguished into a component that represents life evaluation (how well am I doing in life generally) and another that captures emotional well-being (how pleasant my day-to-day experiences are) (Deaton, 2008; Kahneman & Deaton, 2010). This distinction is particularly pertinent for the study of income on subjective well-being: while income consistently affects one's evaluation of life, emotional well-being is only influenced to a certain income threshold ($75,000 in the US), and then tapers off with further income increases (Kahneman & Deaton, 2010). This thesis chapter contributes to the determinants of the evaluative dimension of subjective well-being: life satisfaction and its related dimensions.

By now, a large literature exists that documents a positive relationship between income and wealth on the one hand, and subjective well-being on the other (Deaton, 2008; Sacks et al., 2012). One study serves particularly well to illustrate important findings of this literature: Sacks et al. (2012) show in their re-examination of previous studies using comprehensive global data on subjective well-being that richer individuals report greater subjective well-being than poorer ones; richer countries have higher individual subjective well-being than poorer ones; economic growth is related to rising well-being; there is no satiation point of subjective well-being with respect to income; and the magnitude of these relationships, whatever the domain, are approximately equal. The recent influential study by Lindqvist, Östling and Cesarini (2020) suggests that the main mediator of sustained increases in subjective well-being from higher incomes arises from greater financial satisfaction.

In contrast to financial knowledge, an important mediating characteristic between income and subjective well-being that is difficult to influence through policy measures is brought to light in the findings of Grace, Lee, Sirgy and Bosnjak (2019), who show that happiness materialism – understood as the belief that higher incomes lead to greater happiness – is associated with lower subjective well-being levels as income rises. They empirically test, and find evidence in favour of, the explanation that the moderating effect arises from happiness materialism inducing greater frequencies with which individuals evaluate their standard of living. Resulting from this, individuals, to their detriment, compare their actual financial circumstances with idealised, unattainable ones, which then leads to lower levels of subjective well-being. This thesis chapter relates to the idea that financial aspirations encapsulate desired standards of living. It is intuitive that reference points and financial aspirations are closely linked, as aspirations by definition imply a desired but unfulfilled state of existence. Financial aspirations thus constitute another mediator intractable to policymakers in the income and subjective well-being relationship, studied in the context of developing economies and proposed as a mechanism explaining deteriorations in subjective well-being as income rises (Graham & Pettinato, 2002).

This thesis chapter is further related to the literature on financial literacy. Studies therein have uncovered a host of benefits of being financially literate. For instance, the financially literate are more likely to participate in financial markets (Balloch et al., 2015; Van Rooij et al., 2011; Yoong, 2011), fare better in response to aggregate shocks (Klapper & Panos, 2013) and are better at detecting fraud in their financial accounts (Engels, Kumar & Philip, 2020). However, of particular interest in the context of this thesis chapter is the financial literacy literature on retirement planning (Klapper & Panos, 2011; Lusardi & Mitchell, 2007a, 2007b). Planning inherently involves expectation formation and entails the possibility of disappointment. This literature indicates potential avenues in which financial aspirations can operate when influencing individual reference points. While planning means maximising the likelihood of aspirations materialising in the future, episodes of financial distress can make their fulfilment appear more and more unlikely, thereby deteriorating life satisfaction and happiness.

1.3 Data and variables

We use the rich cross-sectional data from the National Financial Well-Being Survey, made publicly available by the Consumer Financial Protection Bureau (2017). The responses were collected in 2016, and the data are representative of the US adult population. The data contain detailed information on the respondents' levels of subjective well-being, financial literacy and financial distress, as well as their demographic characteristics, incomes and childhood background variables. Engels et al. (2020) use this data set to study the effect of financial literacy on fraud detection, and some definitions of control variables used in this article overlap with those used in their study. The total survey contains responses of 6394 individuals, and we use all respondent observations in our analyses.

1.3.1 Subjective well-being

Our main outcome variables relate to subjective well-being. Respondents are invited to state how well a set of subjective well-being statements apply to themselves. These include “I am satisfied with my life”, “I am optimistic about my future” and “If I work hard today, I will be more successful in the future”.

The available answers lie on a 7-point Likert scale ranging from “strongly disagree” to “strongly agree”. To construct an aggregate subjective well-being variable, we map the supplied assessments to the integers from 0 to 6, where greater values indicate higher levels of well-being. For each respondent, we sum all three answers to obtain an aggregate subjective well-being score, ranging from 0 to 18.

1.3.2 Financial knowledge

To make financial literacy operational in our analyses, we employ the respondents’ financial knowledge as our explanatory variable of interest. We utilise the financial knowledge score by Knoll and Houts (2012) that is provided in the data. The Consumer Financial Protection Bureau (2017) derives the score from a two-parameter item response model using nine survey questions that test the respondents’ understanding of a wide range of financial concepts. The financial knowledge questions used relate to 1) savings accounts, bonds and stock return characteristics, 2) savings accounts, bonds and stock risk characteristics, 3) risk diversification, 4) possibilities of stock market losses, 5) life insurance savings features, 6) house price losses, 7) credit card repayments, 8) the relation between interest rates and bond prices and 9) the relation between mortgage term and interest.

Based on a set of binary variables, one for each question, indicating whether respondents supplied the correct answer to the respective question, financial knowledge scores using the item response model are estimated, jointly with parameters allowing for the various questions' differing levels of difficulty and discriminatory power. Estimates are obtained through Maximum Likelihood Estimation, in which respondents' financial knowledge scores are assumed to follow a standard normal distribution (Consumer Financial Protection Bureau, 2017). Therefore, financial knowledge can be thought of in terms of z-scores: a respondent with average financial knowledge receives a score of zero, while all other scores can be interpreted as distances from the financial knowledge average, normalized by the financial knowledge standard deviation.

1.3.3 Financial distress

The survey captures self-assessments of respondents regarding how frequently they have experienced adverse financial states in the last 12 months, indicative of severe financial distress. Specifically, questions that were asked capture whether 1) the respondent worried food might run out before having money to buy more, 2) whether any household member couldn't afford to see a doctor or to go to a hospital, 3) food didn't last, and there wasn't any money to get more, 4) any household member stopped or reduced medication due to cost, 5) the respondent couldn't afford a place to live and 6) utilities were shut off due to non-payment. The exact question wordings are reported in the Appendix.

Possible answers to these question include "often", "sometimes" and "never", in addition to an option of declining to provide an answer. We construct an indicator variable that takes the value one if respondents choose "often" or "sometimes" in response to at least one experience of financial distress, and zero otherwise. This variable constitutes our key variable indicating respondent-level financial distress.

1.3.4 Summary statistics

Table 1.1 reports summary statistics for our baseline data sample. The columns show the mean, minimum, maximum, standard deviation and number of observations for each variable used. The variables not discussed in the preceding sections are self-explanatory, though detailed variable definitions are provided in the Appendix.

Inspecting the summary statistics reveals that subjective well-being levels appear high: the mean of aggregate subjective well-being equals 13.30, corresponding to 74% of the maximum attainable score of 18. For the subjective well-being questions "Satisfied with my life", "Optimistic about my future" and "Work hard today, successful in future" the mean values correspond to 4.39 (73% of maximum), 4.42 (74%) and 4.50 (75%), respectively.

Turning to financial knowledge, the values of its summary statistics reflect the standard normal parametric specification of the scores. The mean of the respondents' estimated financial knowledge scores is -0.028, with a standard deviation of 0.80 and a range from -2.05 to 1.27.

The financial distress indicator variable reveals that approximately 29% of respondents experienced an adverse financial state indicative of severe financial distress in the last 12 months. This is consistent with 2017 population estimates suggesting that a large part of the US population is subject to financial distress (Karpman et al., 2018).

With respect to other demographic and financial characteristics that enter our empirical analysis as control variables, the summary statistics reveal that the majority of respondents is married or lives with their partner (66%) and has good health (85%). Slightly more than one-third of respondents have completed university education (38%) and a similar fraction report not being financially responsible for any dependent children (34%). The mean respondent age is 51.26; approximately 48% respondents identify as female. The mean household income of respondents is roughly $75K.

1.4 Empirical analysis

1.4.1 Descriptive evidence

We begin our analysis with a visual inspection of the relationship between financial distress, financial knowledge and subjective well-being in Figure 1.1 and provide the results from statistical difference testing in Table 1.2. Figure 1.1 shows the proportions of individuals with subjective well-being greater than the sample average. For the purpose of these investigations, the ordered variables retain their mapping to the integers 0 to 6. Panel (a) shows the values for aggregate subjective well-being; Panel (b) for the dimension "Satisfied with my life"; Panel (c) for "Optimistic about my future"; and Panel (d) for "Work hard today, successful in the future". The horizontal axis separates individuals by their levels of financial knowledge. Financial knowledge is defined as low (high) for scores below (above) the sample average. The white and grey bars indicate proportions for individuals without and with recent experiences of financial distress, respectively.

Three observations summarise the figure characteristics. First, across Panels (a) to (d) and financial knowledge levels, reports of financial distress are associated with lower subjective well-being levels. The aggregate subjective well-being score, in particular, reveals a stark difference between well-being levels inside and outside of financial distress. Second, for respondents in financial distress, a visible difference in subjective well-being can be noted when comparing the proportions for low and high financial respondents: across all panels, subjective well-being levels are lower for high financial knowledge individuals. Third, in particular for aggregate subjective well-being levels, financial knowledge levels do not appear to be strongly related to to subjective well-being levels.

In Table 1.2, we further inspect the subsample of respondents in financial distress and test for the univariate statistical significance of differences in subjective well-being levels across low and high financial knowledge levels. Specifically, the table reports mean values that capture the proportions of respondents with subjective well-being greater than the sample average along the respective dimensions. The columns are split by financial knowledge levels, reporting the respective respondent counts and means. The final set of columns shows the difference in proportions (Δ) between low and high financial knowledge levels, in addition to t-tests for the statistical significance in the differences in means. Overall, the t-tests indicate that the null hypothesis of equality in means across financial knowledge levels is strongly rejected, with the high financially literate exhibiting lower subjective well-being levels across all measured dimensions.

1.4.2 Financial distress and subjective well-being

We begin by establishing the magnitude of the relationship between financial distress and subjective well-being, controlling for a wide range of potential confounding influences. In this way, we establish a baseline for our sample of respondents to judge the size of the moderating effect of financial knowledge explored in the later sections. Though it is known that economic hardships influence subjective well-being (Reeskens & Vandecasteele, 2017), the granularity of individual experiences captured by our financial distress variable, as well as the breadth of subjective well-being dimensions in our data, will allow for a more precise measurement of the association.

We do so by first examining the relationship between financial distress ($FinDistress$) and aggregate subjective well-being ($SubjWellBeing$) in an ordinary least squares (OLS) framework:

$$SubjWellBeing_i = \alpha + \gamma(FinDistress_i) + X_i'\theta + \varepsilon_i, \tag{1.1}$$

where i denotes the individual-level identifier. The financial distress variable takes the value one if respondents indicate experiencing any financial distress in the last 12 months. The vector of controls, X, captures important covariates, namely whether respondents are married or cohabiting, have good health, University education, no children to financially support, as well as the respondents' ages, genders and household incomes. Detailed variable descriptions are provided in the Appendix. Robust standard errors control for arbitrary heteroskedasticity of the error term. The parameter γ in Equation (1.1) captures the association of interest.

The variables capturing the individual dimensions of subjective well-being can take ordered values form "strongly disagree" to "strongly agree" on a 7-point scale. In order to account for the limited-dependent variable nature of these variables, we employ an ordered probit estimation methodology. The regression equation takes the following form:

$$SubjWellBeing_i^d = g(p_i^*) \tag{1.2}$$

$$p_i^* = \alpha^d + \gamma^d(FinDistress_i) + X_i'\theta^d + \varepsilon_i \tag{1.3}$$

$$g(p_i^*) = \begin{cases} \text{“Strongly disagree” if } -\infty < p_i^* \leq c_1 \\ \quad \vdots \\ \text{“Strongly agree”} \quad \text{ if } \quad c_6 < p_i^* < \infty \end{cases} \tag{1.4}$$

where the superscript d indicates the different subjective well-being variables. Specifically, the dependent variable in Equation (1.2) is either the respondents' degree of agreement with the statements “I am satisfied with my life”, “I am optimistic about my future” or “If I work hard today, I will be more successful in future”, respectively. The function $g(\cdot)$ in Equation (1.4) is the ordered probit link (Roodman, 2011), mapping the latent subjective well-being propensities in Equation (1.3), p_i^*, to the observed outcomes. The cutoff points c_1 to c_6 are quantities to be estimated. The regression error term, ε_i follows a standard normal distribution in accordance with the ordered probit normalisation assumption. The remainder of the regression setup follows Equation (1.1).

Table 1.3 reports the estimation results. Column (1) shows the results for aggregate subjective well-being, while Columns (2) to (4) show those for the respective individual dimensions of subjective well-being. Column (1) reports OLS parameter estimates; by contrast, Columns (2) to (4) show ordered probit coefficient estimates. We observe that in aggregate, being in financial distress is associated with a highly significant decrease in subjective well-being of 1.735 points on the aggregate subjective well-being scale. The average subjective well-being levels predicted from Equation (1.1) equal 13.298, which implies that financial distress is associated with deterioration in subjective well-being levels of -1.735/13.298 =

-0.130 or, equivalently, -13%.

Regarding the individual subjective well-being dimensions, the probit coefficient estimates indicate that the greatest association (-0.560) is observed for the life satisfaction measure, followed by those capturing optimism about the future (-0.403) and the belief that working hard today, will increase chances of future success (-0.297). All estimates are statistically significant at the 1% level. In order to provide economic significance regarding these estimates, we compute the average marginal effects for the ordered probit outcome “strongly agree”. In all cases, we observe that respondents in financial distress become less likely to report strong agreement, for which the corresponding AMEs are -0.161, -0.124 and -0.101, implying a reduction in the probability to strongly agree with the respective subjective well-being statements of 16.1, 12.4 and 10.1 percentage points, respectively, with significance at the 1% level throughout.

The associations of the control variables with subjective well-being are intuitive. Respondents who are married or live with their partner report greater subjective well-being, as do those with good health, no children to financially support and higher household incomes. The relations between age and being female, respectively, to subjective well-being are mixed, while having completed a University education infers no relationship of statistical significance at the 10% level with subjective well-being.

Overall, we conclude that the relationship between financial distress and subjective well-being, in aggregate and with respect to its individual dimensions, is economically meaningful and statistically significant in our sample of US respondents. These findings are consistent with those by Reeskens and Vandecasteele (2017).

1.4.3 The mediating role of financial knowledge

Having quantified the relationship between financial distress and subjective well-being in our sample of US respondents, we now turn our investigation to testing for the mediating role of financial knowledge while controlling for a preponderance of individual- and household-level confounding influences. Presupposing that financial knowledge does indeed mediate the association between financial distress and subjective well-being, given the descriptive evidence presented above, the magnitude and significance of the relationship remains an empirical question.

To test for the existence and the nature of the mediating role of financial knowledge with respect to aggregate subjective well-being, we modify the regression framework of Equation (1.1) to include the interaction of financial distress with an indicator variable capturing financial knowledge levels greater than the sample average:

$$\begin{aligned} SubjWellBeing_i = \alpha &+ \gamma_1(HighFinKnow_i) + \gamma_2(FinDistress_i) \\ &+ \gamma_3(FinDistress_i \times HighFinKnow_i) \\ &+ X_i'\theta + \varepsilon_i, \end{aligned} \tag{1.5}$$

The financial distress variable ($FinDistress$) again takes the value one if respondents report having experienced any financial distress in the last 12 months. Financial knowledge is measured through scale by Knoll and Houts (2012), and the high financial knowledge variable ($HighFinKnow$) takes the value one for scores greater than the sample average, and zero otherwise. The remainder of the specification stays unaltered.

To also account, in this context, for the limited-dependent variable nature of the respondent's degree of agreement with the statements "I am satisfied with my life", "I am optimistic about my future" or "If I work hard today, I will be more successful in future", respectively, we again adopt an ordered probit regression approach for these variables. The regression equation takes the following form:

$$SubjWellBeing_i^d = g(p_i^*) \tag{1.6}$$

$$\begin{aligned} p_i^* = \alpha^d &+ \gamma_1^d(HighFinKnow_i) + \gamma_2^d(FinDistress_i) \\ &+ \gamma_3^d(FinDistress_i \times HighFinKnow_i) \\ &+ X_i'\theta^d + \varepsilon_i \end{aligned} \tag{1.7}$$

where the superscript d indicates the different subjective well-being variables. Specifically, the dependent variable in Equation (1.6) is either the variable capturing the respondents' degree of agreement with the statements "I am satisfied with my life", "I am optimistic about my future" or "If I work hard today, I will be more successful in future", respectively. The function $g(\cdot)$ is the ordered probit link introduced in Equation (1.4). The regression error term, ε_i follows a standard normal distribution in accordance with the ordered probit normalisation assumption. The remainder of the regression setup follows Equations (1.2) and (1.3).

Table 1.4 displays the relevant coefficient estimates. Column (1) shows OLS coefficient estimates for aggregate subjective well-being; whereas Columns (2) to (4) show ordered probit estimates for the respective individual dimensions of subjective well-being. Two findings stand out. First, high financial knowledge individuals outside of financial distress do not exhibit lower, strongly significant subjective well-being levels than the base group of respondents with no financial distress and low levels of financial knowledge. Second, across all measures of subjective well-being, respondents that are in financial distress show strongly significant deteriorations in their subjective well-being levels if their financial knowledge levels are high. In all cases, the point estimates for those with low financial knowledge

(γ_2) are lower than those with high levels (γ_3). The reported results of Wald tests for coefficient equality ($H_0 : \gamma_2 = \gamma_3$) reject the null hypothesis that the association between financial distress and subjective well-being is constant across groups of individuals with low and high financial knowledge.

To assign an economic interpretation to the ordered probit estimation results, Figure 1.2 reports the high financial knowledge average marginal effects (AMEs) for individuals inside of financial distress (solid line) and outside of it (dashed line), together with their 95% confidence intervals. Each panel indicates the results for one individual dimension of subjective well-being. The horizontal axis shows the corresponding ordered probit outcome, ranging from 1 (strongly disagree) to 7 (strongly agree), to which the estimates refer.

It stands out that the high financial knowledge AMEs for individuals not in financial distress are economically negligible and largely statistically insignificant. However, the estimates for individuals in financial distress reveal a starkly different picture: across subjective well-being measures, outcomes associated with positive subjective well-being become less likely, as indicated by negative AMEs, while those for negative subjective well-being states become more likely. In terms of economic magnitudes, Panel (a) shows that the probabilities with respect to high subjective well-being outcomes 5 to 7 decrease jointly by approximately 11 percentage points. The non-overlapping confidence intervals between individuals inside and outside of financial distress indicate that the differences in economic magnitudes of the relationship between financial knowledge and subjective well-being differ in in strongly significant ways across financial distress states.

The findings of this section establish strong evidence for a mediating role of financial knowledge in the relationship between financial distress and subjective well-being. Subjective well-being deteriorates more severely for respondents with high financial knowledge when they face financial distress across all considered measures of subjective well-being. In this way, we provide first evidence on a financial knowledge channel in the relationship between a lack of financial resources and deteriorations in subjective well-being (Deaton, 2008; Kahneman & Deaton, 2010).

1.4.4 The mechanism of internal and external financial aspirations

The previous section establishes a strongly negative association between financial knowledge and subjective well-being for individuals in financial distress; however, it remains silent on deeper mechanisms underlying this relationship. In this section, we explore one possible mechanism. This includes tests for the possibilities that financial knowledge influences the degrees of internal and external financial aspirations felt by individuals, namely the following through on individual financial goals (internal dimension) or the meeting of financial commitments towards others (external dimension). For illustration, individuals with greater financial knowledge may have greater personal financial aspirations, so that if they face hardships, ceteris paribus, the gap between the actual and desired reality is greater, leading to a stronger deterioration of subjective well-being. Similarly, financial knowledge can induce a greater sense of aspiring to meet financial responsibilities towards others that the presence of financial distress makes more difficult to attain.

To test this mechanism of financial aspirations, we rely on two questions in the 2016 National Financial Well-Being Survey that elicit how well respondents feel that the statements “I follow-through on my financial commitments to others” and “ I follow-through on financial goals I set for myself”, respectively, apply to themselves. The possible answer choices to each statement are “completely”, “very well”, “somewhat”, “very little” and “not at all”. We derive two further variables, one from each statement, that take the value one if respondents respond with “not at all" or "very little", and zero otherwise. By fixing the frequency with which respondents follow through on their financial goals and commitments to infrequent levels, in our analysis that controls the specific respondent characteristics, we likely capture the strength or degrees that respondents feel with respect to these internal and external dimensions of financial aspirations.

To test this notion empirically, beginning with aggregate subjective well-being, we modify the regression framework described via Equation (1.5) to include the triple interaction of financial distress, high financial knowledge and our constructed measures of financial aspiration:

$$\begin{aligned}
SubjWellBeing_i = \alpha &+ \gamma_1(Aspiration_i) + \gamma_2(HighFinKnow_i) \\
&+ \gamma_3(HighFinKnow_i \times Aspiration_i) + \gamma_4(FinDistress_i) \\
&+ \gamma_5(FinDistress_i \times Aspiration_i) \\
&+ \gamma_6(FinDistress_i \times HighFinKnow_i) \\
&+ \gamma_7(FinDistress_i \times HighFinKnow_i \times Aspiration_i) \\
&+ X_i'\theta + \varepsilon_i, \qquad (1.8)
\end{aligned}$$

where the financial aspirations variable ($Aspiration$) is equal to either the derived variable "Cannot meet financial goals" or "Cannot meet financial commitments" described above. The financial distress variable ($FinDistress$) takes the value one if respondents report having experienced any financial distress in the last 12 months. Financial knowledge is measured through the scale by Knoll and Houts (2012), and the high financial knowledge variable ($HighFinKnow$) takes the value one for scores greater than the sample average, and zero otherwise. The remainder of the specification stays unaltered.

For the variables "I am satisfied with my life", "I am optimistic about my future" or "If I work hard today, I will be more successful in future", respectively, we again adopt an ordered probit regression approach. The regression equation takes the following form:

$$
\begin{aligned}
SubjWellBeing_i = \ & g(p_i^*) && (1.9)\\
p_i^* = \ & \alpha + \gamma_1^d(Aspiration_i) + \gamma_2^d(HighFinKnow_i)\\
& + \gamma_3^d(HighFinKnow_i \times Aspiration_i) + \gamma_4^d(FinDistress_i)\\
& + \gamma_5^d(FinDistress_i \times Aspiration_i)\\
& + \gamma_6^d(FinDistress_i \times HighFinKnow_i)\\
& + \gamma_7^d(FinDistress_i \times HighFinKnow_i \times Aspiration_i)\\
& + X_i'\theta^d + \varepsilon_i, && (1.10)
\end{aligned}
$$

where the superscript d indicates the different subjective well-being variables. The function $g(\cdot)$ is again the ordered probit link introduced in Equation (1.4). The regression error term, ε_i, follows a standard normal distribution in accordance with the ordered probit normalisation assumption. The remainder of the regression setup follows Equation (1.2).

Table 1.5 reports the estimation results. Column (1) shows the OLS estimation results for aggregate subjective well-being, while Columns (2) to (4) show the ordered probit results for the respective individual dimensions of subjective well-being. Panel A shows those for the financial aspiration variable "Cannot meet financial goals", while Panel B shows those for "Cannot meet financial commitments others". In Panel A, the effect of the triple interaction on subjective well-being when the financial aspiration variable is equal to zero (γ_6) is negative and strongly significant across all subjective well-being measures. However, the effect more than doubles when respondents report an inability to follow through with their own financial goals (γ_7). The reported results of Wald tests for coefficient equality ($H_0 : \gamma_6 = \gamma_7$) reject the null hypothesis that the association between financial distress and subjective well-being is constant across high financial knowledge individuals with differing levels of personal financial aspirations. In contrast, conducting the analogous Wald tests for the results in Panel B reveals that the observed differences in coefficients therein are not statistically significant at the 10% level.

Figure 1.3 shows the financial aspirations average marginal effects (AME) for each ordered probit regression outcome. Panel (a) to (c) display the results for the variables “Satisfied with my life”, “Optimistic about my future” and “Work hard today, succeed in future,” respectively. The AMEs are calculated for the subsample of high financial knowledge respondents in financial distress. The white and grey bars, respectively, refer to financial aspirations relating to own financial goals and financial commitments towards others. Three observations stand out.

First, the economic magnitudes of internal financial aspirations (financial goals) increase substantially versus those for high financial knowledge in Figure 1.2. For instance, cumulatively the probabilities to achieve SWB outcomes 5 to 7 in Figure 1.2 in Panel (a) sum to -11 percentage points, while those in Figure 1.3 sum to approximately -20 percentage points. Second, while the error bars indicate that the AMEs for financial goals are precisely estimated, implying strong statistical significance, those for financial commitments are not: most error bars include the horizontal axis at zero. Third, all three panels provide a consistent picture of the effects of financial aspirations regarding financial goals: the inability to meet these aspirations when in financial distress is associated with strong deterioration in subjective well-being.

Overall, the results provide strong evidence that internal financial aspirations provide the mechanism through which financial knowledge influences subjective well-being, while the channel of external financial aspirations appears to be insignificant.

1.5 Instrumental variable approach

We begin addressing causal identification challenges, such those induced by omitted variable bias, in an instrumental variable approach. Specifically, individuals with more materialistic attitudes might seek to acquire more financial knowledge in order to achieve their desires to consume more. Further, materialism is associated with lower negative subjective well-being levels (Grace et al., 2019) and can thus be an important potential omitted variable, introducing bias in the OLS estimate of the interaction between financial distress and high financial knowledge on subjective well-being.

We use the method by Lewbel (2012) to achieve identification of these interaction terms. Specifically, Lewbel (2012) shows that it is possible to identify the effect of potentially endogenous regressors, in the absence of strong candidates for instrumental variables, if exogenous variables (Z) and heteroskedastic errors are present. We describe the mechanics of the method below. In the context of our application, the full set of control variables serves as our exogenous variables Z for construction of heteroskedasticity-based instruments: married or cohabitating, good health, University education, no dependent children, age, female and household income. We further include an additional set of control variables that capture childhood interactions regarding financial education between the respondents and their parents for the purpose of this instrumental variable analysis.

This identification method can be combined with the traditional instrumental variable approach. In two specifications, we complement the generated heteroskedasticy-based instruments, whose construction is described in the below, with an instrumental variable that captures whether the respondents' parents were University educated.

In order to qualify as a valid instrument, this external instrumental variable ought to satisfy the requirements to (i) be strongly related to the potentially endogenous variable, financial knowledge, and (ii) to influence subjective well-being only through its effect on financial knowledge. University educated parents are likely to have higher degrees of financial knowledge (Lusardi et al., 2017) and to consequently to pass this knowledge on to their children, in turn raising their children's financial knowledge levels, together with their motivations to remain financially literate as adults. Thus, we expect a strong relationship between the instrument and the endogenous regressor, thereby satisfying condition (i), instrumental relevance. Further, given the inclusion of the additional child-parent background variables, parents' University education is unlikely to exert a direct influence on the respondents' subjective well-being levels; thus meeting condition (ii), the exclusion restriction. Overall, it appears likely that the criteria for valid instruments are satisfied in our setting. Given that the inclusion of additional heteroskedasticity-based instruments will provide more instruments than endogenous regressors, we will, in addition to the instrumental relevance condition, be able to provide evidence in favour of the exclusion restriction being satisfied by testing the validity of the overidentifying restrictions.

We use the Stata implementation ivreg2h (Schaffer & Baum, 2012) of the Lewbel (2012) two-stage estimator. To construct the heteroskedasticity-based instruments, each endogenous variable is regressed on all of the control variables in the subjective well-being equation (denoted below by vector X) along with the vector Z, in following first-stage regression:

$$\mathbb{1}(FinDistress_i = 1 \text{ and } HighFinKnow_i = 0) = \alpha_1 + X_i'\beta_1 + Z_i'\gamma_1 + \varepsilon_{i,1} \quad (1.11)$$

$$\mathbb{1}(FinDistress_i = 1 \text{ and } HighFinKnow_i = 1) = \alpha_2 + X_i'\beta_2 + Z_i'\gamma_2 + \varepsilon_{i,2} \quad (1.12)$$

$$\mathbb{1}(FinDistress_i = 1 \text{ and } HighFinKnow_i = 1) = \alpha_3 + X_i'\beta_3 + Z_i'\gamma_3 + \varepsilon_{i,3} \quad (1.13)$$

where $\mathbb{1}(\cdot)$ take the value one if the argument is true and zero otherwise. The residuals $\hat{\varepsilon}_{i,1}, \hat{\varepsilon}_{i,2}$ and $\hat{\varepsilon}_{i,3}$ are then retrieved in order to create the heteroskedasticity-based instruments as follows:

$$(Z_i - \bar{Z})\hat{\varepsilon}_{i,s} \quad \text{for } s = 1, 2, 3 \tag{1.14}$$

where $\bar{Z}$ is the vector of sample means of $Z_i, \forall i$. As Lewbel (2012) shows, identification requires that the error terms in the first-stage regressions in Equations (1.11) to (1.13) are heteroskedastic; we therefore follow the recommendations in Baum and Lewbel (2019) and use the Breusch-Pagan test of heteroskedasticity. In all cases, the results show that the null of homoskedastic errors is clearly rejected in each first-stage regression with a p-value effectively equal to zero. We then use the set of instruments obtained through Equation (1.14) for each first-stage regression in (1.11) to (1.13) in a standard two-stage instrumental variable approach to estimate the causal effects of the interaction between financial distress and high financial knowledge on aggregate subjective well-being.

Table 1.6 reports the regression results. Columns (1) to (4) all show the results from the second-stage equation in which aggregate subjective well-being is the dependent variable. Column (1) shows the results in which the control variables denoted by (i) are used in the construction of heteroskedasticity-based instruments, and where the external instrument (Unversity-educated parents) is excluded from the first-stage regression. In Column (2), this external instrument is included, keeping the remainder of the Column (1) specification unaltered. Columns (3) and (4) use all controls, denoted (i) and (ii), in the construction of the heteroskedasticity-based instruments. Columns (3) and (4) differ in the inclusion of the external instrument: Column (3) excludes this instrument; Column (4) does not.

Importantly, two tests suggest that the instrumental variables in the approaches in Columns (1) to (4) identify the causal effects of the interactions between financial distress and high financial knowledge on aggregate subjective well-being. First, the first-stage F statistics in all columns exceed the 10% critical value by Stock and Yogo (2005), suggesting that weak identification is of little concern and, specifically, that the instrumental relevance condition is met. In particular, the F statistics in Columns (1) and (2) are greater than the corresponding 5% critical value, suggesting that the heteroskedasticity-based instruments generated from the controls (i) are more strongly related to the endogenous regressors than those in (ii). Further, in all columns, the p-values for the Hansen (1982) J statistic indicate failure to reject the null that the overidentifying restrictions are violated, suggesting no evidence that the exclusion restriction is not met.

Table 1.6 thus shows the effect of the interaction between financial distress and high financial knowledge on aggregate subjective well-being. The main focus in this table is on the moderating effect of high financial knowledge in the relationship between financial distress and subjective well-being. The corresponding estimates can be interpreted as heterogeneous treatment effects of financial distress with respect to financial knowledge levels. It emerges that the interactions involving cases of financially distressed are all highly significant ($p < 0.01$). The corresponding effect sizes for low financial knowledge respondents are stable across all columns with magnitudes of approximately -0.7; whereas the magnitudes for high financial knowledge respondents vary between -1.438 (Column (1)) to -1.346 (Column (4)). This shows that the negative effect of financial distress almost doubles as financial knowledge increases from low to high.

The findings of these sections provide causal evidence that financial knowledge acts as a mediating variable for the effect of financial distress on subjective well-being; thus, we contribute evidence for financial knowledge as a channel in the relationship between financial resources and subjective well-being (Deaton, 2008; Kahneman & Deaton, 2010).

1.6 Propensity score matching

In order to reduce potential biases in estimating the association between financial knowledge and subjective well-being stemming from the endogenous acquisition of financial knowledge, we use a standard propensity score matching approach (Caliendo & Kopeinig, 2008; Dehejia & Wahba, 2002). For instance, Lusardi et al. (2017) show that individuals face differing incentives to acquire financial knowledge when smoothing consumption over their life cycles. For some individuals, this can lead to a need to maximise returns on investments, thereby forming a strong motive to increase financial knowledge. In other words, these individuals can select into the high financial group of respondents given their observable characteristics, potentially introducing self-selection bias into the estimation of the treatment effect of financial knowledge on subjective well-being.

Specifically, given this heterogeneity in incentives to acquire financial knowledge, we face the concern that for some individuals, in light of their own characteristics and circumstances, the possibility exists that it is never optimal to acquire high financial knowledge. This zero probability of being equipped with financial knowledge, formally expressed as the violation of the Common Support assumption, disqualifies these respondents as valid members of the control group in the estimation of the treatment effect of financial knowledge. Intuitively, identifying the treatment effect requires that the subjective well-being levels of the control group are a good substitute for the counterfactual outcomes of the highly knowledgeable individuals had they instead been endowed with low financial knowledge. Individuals that are too different and will never acquire high financial knowledge are clearly unsuitable for approximation of these counterfactual outcomes and ought to be discarded from analysis. Propensity score matching makes this notion operational, and we describe the estimation details in the next section.

In the estimation of the propensity scores, we restrict our sample to those respondents in financial distress. Following from the above, our approach involves discarding from the analysis those individuals in financial distress that have low estimated likelihoods of acquiring financial knowledge levels given their observed characteristics. Our subsample of interest in this exercise is the group of financially distressed respondents with high financial knowledge and the 1-to-1 nearest neighbour matched low financial knowledge respondents. Aside from the associated econometric considerations, it is of direct economic interest to inspect the relationship between financial knowledge and subjective well-being directly in the subset of respondents facing financial distress. Confirming that the observed relationships hold in this group of respondents is crucially important, as respondents outside of financial distress will be significantly less likely to select into financial education programmes designed for poverty alleviation. In this way, we directly establish the relevance of our findings for financial educators who likely will be in close contact with such individuals and may consider the implications of the associations we uncover for curriculum design.

1.6.1 Obtaining the matched sample

In order to obtain a matched sample of respondents, we employ the high financial knowledge indicator variable introduced in the preceding sections, $HighFinKnow_i$, that takes the value one if the respondents' financial knowledge levels are greater than the sample average, and zero otherwise. We subset the full data for only those respondents experiencing financial distress and then estimate a probit model with $HighFinKnow_i$ as the dependent variable. From the estimated model, we obtain the predicted propensity scores for each respondent, i, given their observed characteristics, X_i. The included observable characteristics capture whether respondents are married or cohabiting, have good health, a University education, no children to financially support, as well as the respondents' ages, genders and household incomes.

The probit model to obtain the predicted propensity scores takes the following form:

$$HighFinKnow_i = \mathbb{1}(ps_i^* > 0), \tag{1.15}$$

$$ps_i^* = X_i'\theta + \varepsilon_i, \quad \forall i \in \{i \mid FinDistress_i = 1\}, \tag{1.16}$$

where ps_i^* denotes the latent propensity score for individual i, while $\mathbb{1}(\cdot)$ is the indicator function, taking the value one if the argument is true and zero otherwise, and ε_i is a standard normal error term. The predicted latent propensities from this probit regression are then used to construct a nearest-neighbour matched sample of respondents; where those with high financial knowledge ($HighFinKnow_i$ equal to unity) are matched, without replacement, to those with the closest propensity score chosen from the group of respondents with low financial knowledge levels ($HighFinKnow_i$ equal to zero).

The assess whether our matching procedure results in balancing the observed characteristics between high and low financial knowledge respondents and succeeds in constructing a sample in which high financial knowledge occurs as-if randomly between respondents, we follow the approach by Rodnyansky and Darmouni (2017) and conduct the following test: we compare the estimates from the probit model in Equation (1.16) that generates the propensity scores with the results of the equivalently specified probit model estimated using the sample of matched respondents only. After matching, none of the employed covariates should be predictive of high financial knowledge, and the predicted baseline probabilities of the re-estimated probit model should indicate high financial knowledge occurring with 50% probability.

Table 1.7 reports the results. Compared to the propensity score model in Columns (1), the magnitudes of the probit regression coefficients decline when compared to the probit model using the matched sample in Column (2). Importantly, none of the sources of heterogeneity continue to play any role in explaining high financial knowledge levels, whereas the variables University education, age, female and household income were statistically significant in the pre-match model of Column (1). Furthermore, in Column (2), the χ^2 test for overall model fit shows that one cannot reject the null hypothesis that all coefficient estimates are zero: the p-value equals 0.625. Further, the baseline predicted probability using the matched sample equals 50%. Overall, we conclude from this test results that the matching process removes meaningful differences along observable dimensions between the two groups of respondents with high and low financial knowledge.

Moreover, Table 1.8 reports the covariate differences in means across low and high financial knowledge levels for the full and matched samples. The columns are split by financial knowledge (FK) levels, reporting the respective respondent means together with the differences in means (Δ) between low and high financial knowledge levels, in addition to t-tests for their statistical significance. Overall, the t-tests indicate that in the full sample the respondent characteristics across low and high financial knowledge levels are strongly significant, except for the variable no dependent children. This is in contrast to the matched sample obtained through propensity score matching: the differences between financial knowledge groups are insignificant at the 10% level and economically negligible. These results strengthen the case that the subjective well-being outcomes of the low financial knowledge group constitute reasonable approximations for the high financial knowledge respondents' counterfactual outcomes.

1.6.2 Results using the matched sample

We now turn to estimating the relationships between financial knowledge, financial aspirations and subjective well-being using information on respondents in financial distress that were matched in the process of balancing the respondents' covariates across low and high financial knowledge groups. In this way, the influence of the endogenous acquisition of financial knowledge is attenuated.

Employing this propensity-score-matched subset of respondents, we begin by regressing our aggregate subjective well-being variable ($SubjWellBeing_i$) on the high financial knowledge indicator variable, $HighFinKnow_i$:

$$\begin{aligned} SubjWellBeing_i = \alpha + \gamma(HighFinKnow_i) + X_i'\theta + \varepsilon_i, \\ \forall i \in \{i \mid FinDistress_i = 1 \text{ and } Matched_i = 1\}, \end{aligned} \tag{1.17}$$

where X_i captures the respondents' characteristics used in the matching procedure and ε_i is a normal error term. The parameter γ is the key association between financial knowledge and aggregate subjective well-being.

In order to account for the limited-dependent nature of the additional subjective well-being variables, in line with the discussions in the preceding sections, we again adopt an ordered probit regression approach. The regression equation takes the following form:

$$SubjWellBeing_i^d = g(p_i^*) \tag{1.18}$$

$$\begin{aligned} p_i^* = \alpha^d + \gamma^d(HighFinKnow_i) + X_i'\theta^d + \varepsilon_i^d, \\ \forall i \in \{i \mid FinDistress_i = 1 \text{ and } Matched_i = 1\}, \end{aligned} \tag{1.19}$$

where the superscript d indicates the different subjective well-being variables ("I am satisfied with my life", "I am optimistic about my future" or "If I work hard today, I will be more successful in future", respectively). As before, the function $g(\cdot)$ is the ordered probit link introduced in Equation (1.4). The regression error term, ε_i follows a standard normal distribution in accordance with the ordered probit normalisation assumption.

Table 1.9 reports the estimation results. Column (1) shows the results for aggregate subjective well-being, while Columns (2) to (4) show those for the respective individual dimensions of subjective well-being. Column (1) shows that the significance of high financial knowledge for aggregate subjective well-being has dropped to marginal significance at the 10% level. From Columns (2) to (4) it appears that this drop in significance can be attributed to the insignificance of high financial knowledge for the variables "Optimistic about my future" and "Work hard today, successful in future"; however, high financial knowledge remains strongly significant for the variable "Satisfied with my life".

While these results corroborate the conclusion drawn from our earlier sections – that life satisfaction deteriorates for high financial knowledge individuals in adverse financial circumstances – these results are strongly encouraging for interventions seeking to improve the financial prospects of the financially distressed. Though life satisfaction deteriorates, optimism about the future and the belief that hard work pays off are not affected. Consequently, this suggests the possibility that the motivations of these respondents to better their circumstances are unaffected, potentially providing the opportunity for public policy interventions to be positively received by these individuals.

We further investigate the importance of financial aspirations for the subsample of the financially distressed as a mediating variable between high financial knowledge and subjective well-being. In order to do so, we slightly adapt the regression frameworks described in Section 1.4.4. In the case of the linear specification for the aggregate subjective well-being variable, this takes the form of an interaction between financial aspirations ($Aspiration_i$) and high financial knowledge ($HighFinKnow_i$) on the subset of financially distressed respondents:

$$\begin{aligned} SubjWellBeing_i = \alpha &+ \gamma_1(Aspiration_i) + \gamma_2(HighFinKnow_i) \\ &+ \gamma_3(HighFinKnow_i \times Aspiration_i) \\ &+ X_i'\theta + \varepsilon_i, \\ &\forall i \in \{i \mid FinDistress_i = 1 \text{ and } Matched_i = 1\}, \end{aligned} \tag{1.20}$$

Whereas the ordered probit specification for the individual dimensions of subjective well is adjusted accordingly and is expressed as follows:

$$SubjWellBeing_i^d = \ g(p_i^*) \tag{1.21}$$

$$\begin{aligned} p_i^* = \alpha^d &+ \gamma_1^d(Aspiration_i) + \gamma_2^d(HighFinKnow_i) \\ &+ \gamma_3^d(HighFinKnow_i \times Aspiration_i) \\ &+ X_i'\theta^d + \varepsilon_i^d, \\ &\forall i \in \{i \mid FinDistress_i = 1 \text{ and } Matched_i = 1\}, \end{aligned} \tag{1.22}$$

All other features of the regression framework remain the same.

Table 1.10 reports the estimation results. Panel A displays the results for the aspirations relating to the individuals' own financial goals; Panel B shows those for financial commitments towards others. Column (1) displays the regression results for the aggregate subjective well-being measures, and Columns (2) to (4) show those for the individual subjective well-being dimensions.

Two findings stand out in Panel A. The interaction between high financial knowledge (for brevity, denoted in this table by FK) and the inability to meet financial goals is highly significant ($p < 0.01$) across all columns. While the interaction for low financial knowledge individuals that cannot meet their financial goals is also significant, the coefficient estimates in all cases approximately double in magnitude for those with high financial knowledge. Second, for high financial individuals that can meet their financial goals, the effect on subjective well-being is insignificant, with the exception for the variable “Satisfied with my life”, though for this case statistical significance is marginal at the 10% level. These findings together corroborate that the inability to meet financial goals constitutes a significant mediator between high financial knowledge and subjective well-being.

Panel B does not show clear evidence that the inability to meet financial commitments towards others is a mediator between high financial knowledge and subjective well-being. In Columns (1) and (2), both interaction terms involving high financial knowledge are statistically significant, while only one is so in Column (3), and none in Column (4). This indicates that it is indeed not the inability to meet financial commitments towards others that moderates the relationship between financial knowledge and subjective well-being but that to meet one’s own financial goals.

1.7 Limitations and future research

In the light of their possible relevance for financial educators, our findings identify a need for further research. It is important to understand whether lower subjective well-being levels for high financial knowledge individuals in financial distress strengthen or depress their motivations to return to more positive financial outcomes. For example, in the case of the link between financial literacy and fraud detection, Engels et al. (2020) show that lower levels of subjective well-being diminish the abilities of the highly financially literate to detect fraud. The case of depleting motivations could suggest an additional channel through which the link between financial education and financial outcomes could be put at risk and would complement the findings of Carpena, Cole, Shapiro and Zia (2019) who explore non-financial personal constraints that can impede the transmission of financial education to financial outcomes.

In addition, a limitation of this study is the reliance on the cross-sectional nature of our sample. While we attenuate endogeneity concerns through an instrumental variable and propensity score matching approach, we cannot provide conclusive evidence that an intervention that raises financial literacy levels will unambiguously lead to a decrease in subjective well-being levels for financially distressed individuals. Future work could exploit time-ordered data to examine changes in financial distress status on subjective well-being levels across financial knowledge groups and over time or test for the presence of causal effects in a randomised control trial approach.

1.8 Conclusion

This thesis chapter provides evidence for how financial literacy influences subjective well-being for respondents in financial distress using a representative survey from the year 2016 of 6394 US respondents. The empirical analysis reveals two important findings. First, individuals in financial distress suffer higher deterioration in their life satisfaction when their financial knowledge is high versus corresponding decreases for low financial knowledge respondents. Second, a significant driver of deteriorating life satisfaction for these respondents is the inability to meet the financial goals they set for themselves. Interestingly, not being able to meet financial commitments towards others does not affect the subjective well-being of high and low financial knowledge respondents differently. In order to assign a causal interpretation to these findings, we implement two approaches to mitigate endogeneity concerns. First, we instrument the interactions of financial distress and high financial knowledge with heteroskedasticity-based instruments using the method by Lewbel (2012). Second, we test the observed relationships for the subsample of financially distressed respondents, matching high financial knowledge respondents 1-to-1 to low knowledge respondents in a propensity score matching approach. Both approaches support that financial literacy is inversely related to subjective well-being for respondents in financial distress.

These results have relevance for policy in light of two parallel developments in public policy. First, as higher subjective well-being has been found to positively influence a variety of desirable outcomes, such as health, life expectancy, productivity and the quality of relationships, systematic efforts to measure subjective well-being at the national and local level are underway; in many instances informing public policy decisions (Diener, Oishi & Tay, 2018). Second, financial distress, as measured by problems fulfilling basic consumption needs relating to housing, utilities, health care or food, is prevalent in the US population. In 2017, approximately 40% of US adults across the income distribution experienced at least one of these issues, and, among those affected, 60.2% experienced problems in two or more of these domains (Karpman et al., 2018). As a consequence, poverty alleviation programmes to reduce financial distress, such as financial education programmes, have attracted considerable interest (Brown & Robinson, 2016; McKernan, Ratcliffe & Iceland, 2018). These two policy agendas risks being at odds with one another if the interplay between financial distress, financial knowledge and subjective well-being is not reflected in their designs. Our findings indicate that, in order to avoid such a clash, public policy can combine the provision in financial education programmes with financial support on preferential terms that empowers financially distressed individuals to attain their immediate financial goals.

Figure 1.1

Financial distress, subjective well-being and financial knowledge

This figure shows the proportions of individuals with subjective well-being greater than the sample average. Panel (a) shows the values for aggregate subjective well-being; Panel (b) for the dimension “Satisfied with my life”; Panel (c) for “Optimistic about my future”; and Panel (d) for "Work hard today, successful in the future”. The horizontal axis separates individuals by their levels of financial knowledge. Financial knowledge is defined as low (high) when it is below (above) the sample average. The white and grey bars indicate proportions for individuals without and with recent experiences of financial distress, respectively. The exact variable definitions can be found in the Appendix.

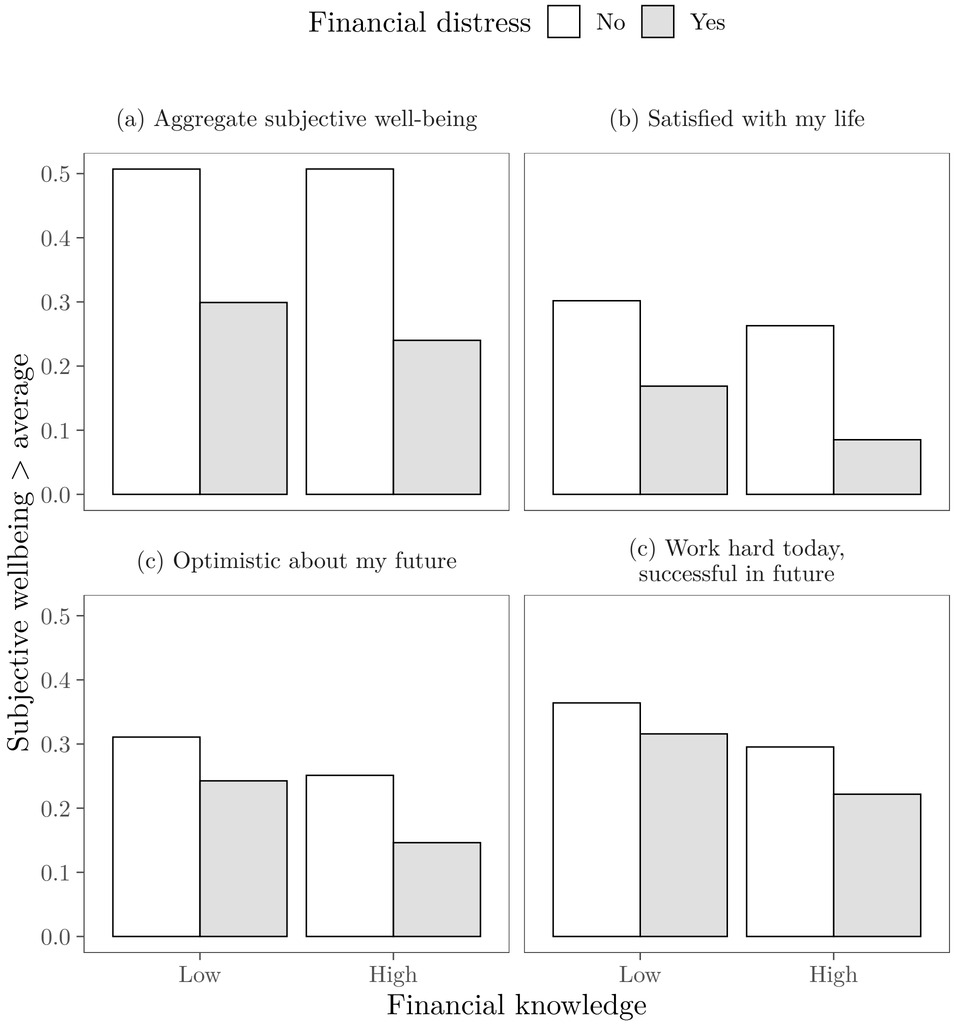

Figure 1.2

High financial knowledge average marginal effects

This figure shows financial knowledge average marginal effects (AMEs) (vertical axis) for each ordered probit regression outcome (horizontal axis) with respect to the subjective well-being variables indicated in Panels (a) to (c). The error bars display the corresponding 95% confidence intervals. Black triangles show the high financial knowledge AMEs for respondents inside of financial distress; white triangles for those outside of financial distress. Panel (a) shows the AMEs for “Satisfied with my life”; Panel (b) for “Optimistic about my future”; and Panel (c) for "Work hard today, successful in the future”. The exact variable definitions can be found in the Appendix.

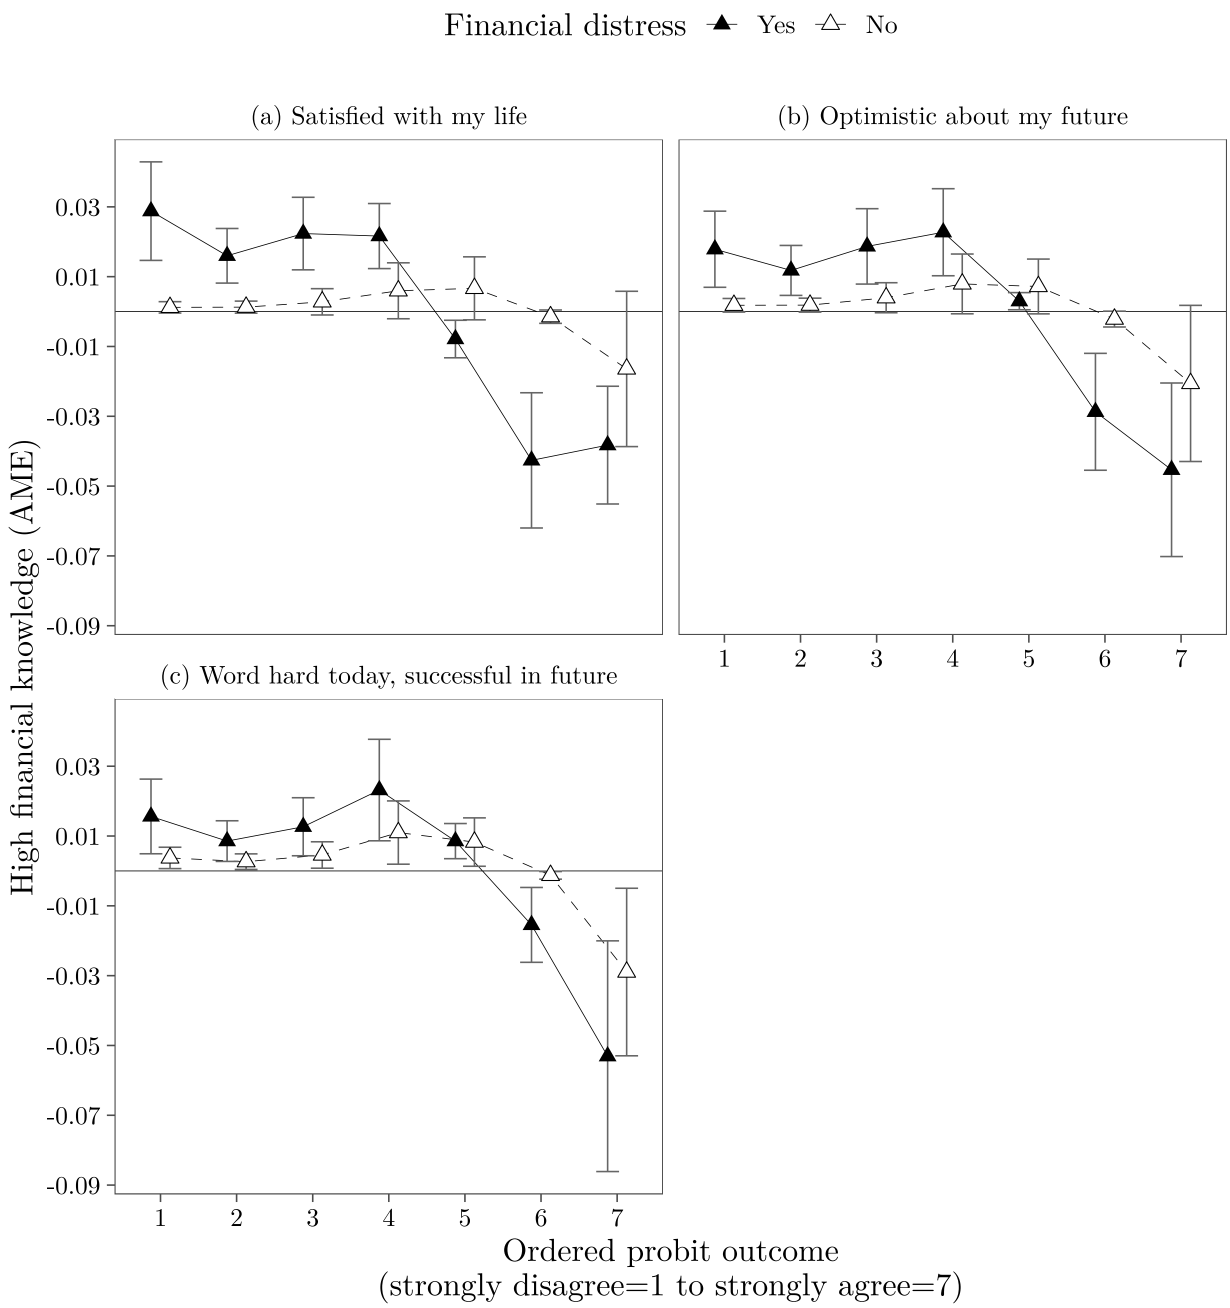

Figure 1.3

Financial aspirations average marginal effects

This figure shows financial aspirations average marginal effects (AMEs) (vertical axis) for each ordered probit regression outcome (horizontal axis) with respect to the subjective well-being variables indicated in Panels (a) to (c). The AMEs are calculated for the subsample of high financial knowledge respondents in financial distress. The error bars display the corresponding 95% confidence intervals. White bars show AMEs for respondents that report an inability to meet their own financial goals; grey bars for those that report an inability to meet financial commitments to others. Panel (a) shows the AMEs for "Satisfied with my life"; Panel (b) for "Optimistic about my future"; and Panel (c) for "Work hard today, successful in the future". The exact variable definitions can be found in the Appendix.

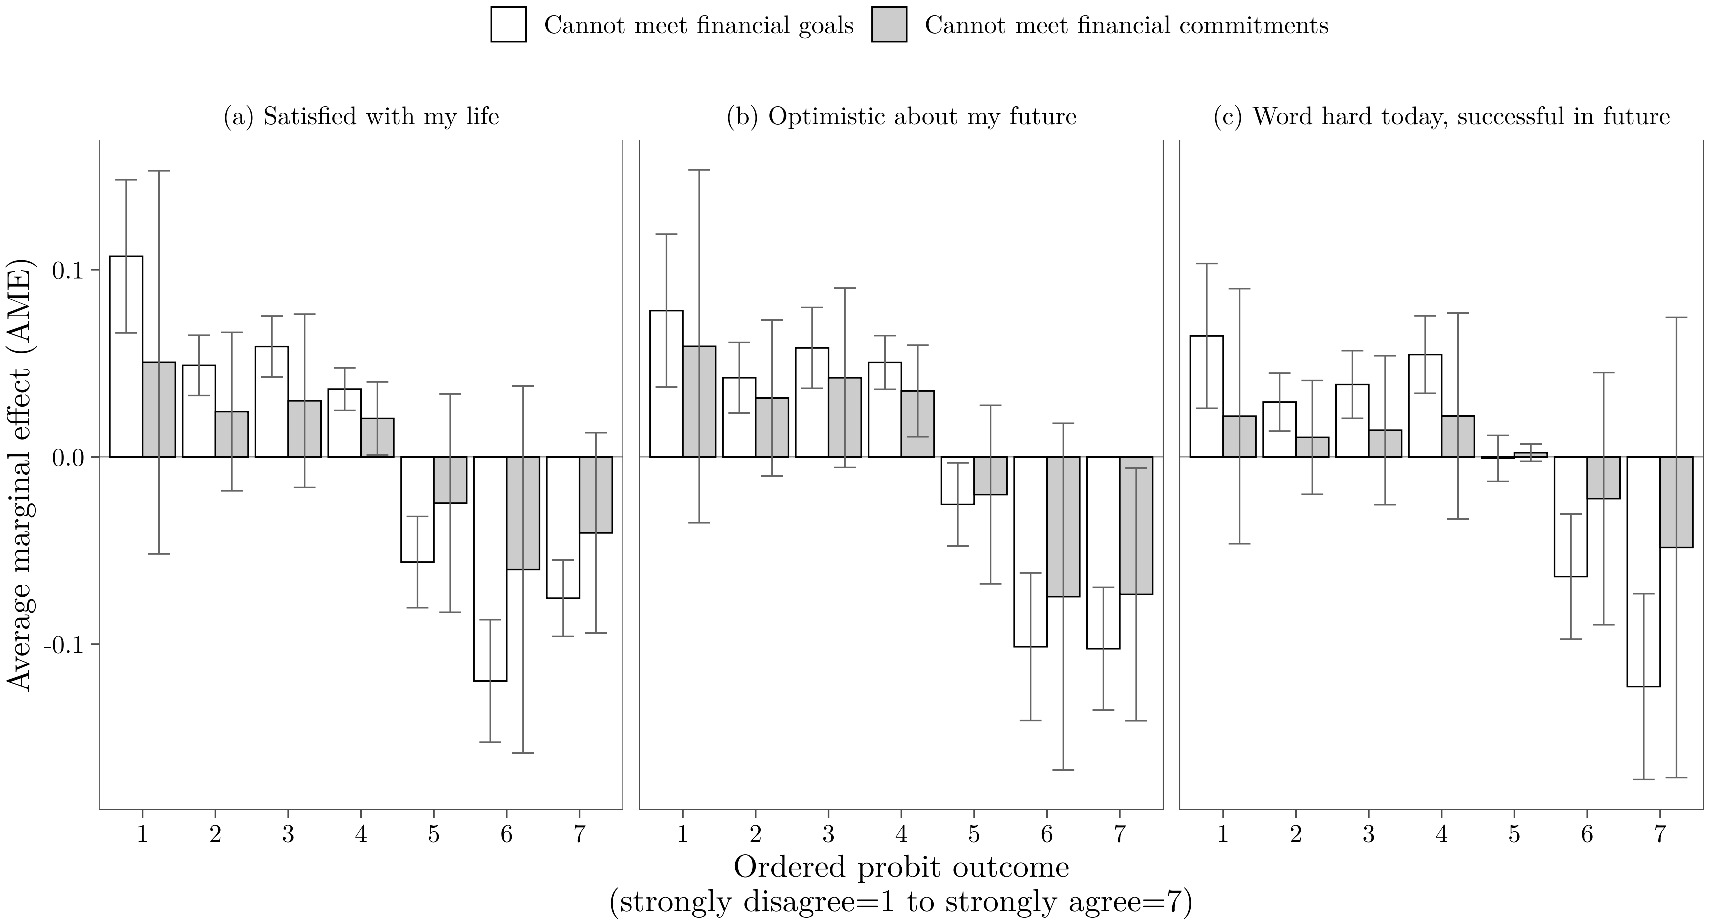

Table 1.1

Summary statistics

This table reports the summary statistics for our baseline data sample. The data is sourced from the 2016 National Financial well-being survey. The Appendix provides exact definitions of all variables.

	Mean	Min.	Max.	SD	Obs.
Aggregate subjective wellbeing	13.30	0.00	18.00	3.65	6314
Satisfied with my life	4.39	0.00	6.00	1.43	6363
Optimistic about my future	4.42	0.00	6.00	1.43	6337
Work hard today, successful in future	4.50	0.00	6.00	1.48	6329
Financial knowledge	-0.06	-2.05	1.27	0.81	6394
Financial distress	0.29	0.00	1.00	0.45	6394
Married or cohabiting	0.66	0.00	1.00	0.48	6394
Good health	0.85	0.00	1.00	0.36	6394
University education	0.38	0.00	1.00	0.48	6394
No dependent children	0.34	0.00	1.00	0.47	6394
Age	51.26	21.00	75.00	17.22	6394
Female	0.48	0.00	1.00	0.50	6394
Household income (1000s)	75.31	10.00	150.00	46.18	6394

Table 1.2

Financial knowledge and subjective well-being by financial distress

This table reports an analysis of the differences in subjective well-being means by financial knowledge levels for respondents with recent experiences of financial distress. The mean values capture the proportions of respondents with subjective well-being greater than the sample average along the respective dimensions. The columns report the counts and means for financial knowledge levels, together with the difference in proportions (Δ) and the results of t-tests for differences in means. The exact variable definitions can be found in the Appendix. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Financial knowledge					
	Low		High			
	Count	Mean	Count	Mean	Δ	t
Aggregate subjective well-being > average	1304	0.30	504	0.24	-0.06	-2.58***
Satisfied with my life > average	1310	0.17	505	0.09	-0.08	-5.16***
Optimistic about my future > average	1311	0.24	506	0.15	-0.10	-4.89***
Work hard today, successful in future > average	1311	0.32	505	0.22	-0.09	-4.17***

Table 1.3

Financial distress and subjective well-being

This table reports the results of OLS and ordered probit regressions of Equations (1.1) and (1.2) . The dependent variables in Columns (1) to (4) are the indicated dimensions of subjective well-being. The independent variable of interest is the financial distress indicator. The exact variable definitions can be found in the Appendix. The point estimates are OLS coefficients in Column (1) and ordered probit coefficients in Columns (2) to (4). Robust standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
Financial distress	-1.735*** (0.12)	-0.560*** (0.03)	-0.403*** (0.03)	-0.297*** (0.03)
Married or cohabiting	0.597*** (0.10)	0.290*** (0.03)	0.148*** (0.03)	0.068** (0.03)
Good health	2.336*** (0.14)	0.607*** (0.04)	0.628*** (0.04)	0.428*** (0.04)
University education	-0.052 (0.09)	-0.047 (0.03)	0.028 (0.03)	-0.066** (0.03)
No dependent children	0.242** (0.09)	-0.001 (0.03)	0.080*** (0.03)	0.082*** (0.03)
Age	-0.012*** (0.00)	0.007*** (0.00)	-0.005*** (0.00)	-0.011*** (0.00)
Female	0.087 (0.09)	0.095*** (0.03)	0.057** (0.03)	-0.049* (0.03)
Household income	0.003** (0.00)	0.001 (0.00)	0.000 (0.00)	0.001** (0.00)
Constant	11.738*** (0.23)			
Baseline predicted subjective wellbeing	13.298			
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Observations	6314	6363	6337	6329

Table 1.4

Financial distress, financial knowledge and subjective well-being

This table reports the results of OLS and ordered probit regressions. The dependent variables in Columns (1) to (4) are the indicated dimensions of subjective well-being. The independent variables of interest are interactions between the financial distress and high financial knowledge indicators. Financial knowledge greater than the sample average is defined as high. The exact variable definitions can be found in the Appendix. The point estimates are OLS coefficients in Column (1) and ordered probit coefficients in Columns (2) to (4). Robust standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
Financial distress=0 × High financial knowledge=1	-0.178*	-0.049	-0.061*	-0.081**
	(0.10)	(0.03)	(0.03)	(0.03)
Financial distress=1 × High financial knowledge=0 (γ_2)	-1.548***	-0.504***	-0.364***	-0.274***
	(0.14)	(0.04)	(0.04)	(0.04)
Financial distress=1 × High financial knowledge=1 (γ_3)	-2.365***	-0.744***	-0.566***	-0.454***
	(0.19)	(0.05)	(0.05)	(0.05)
Married or cohabiting	0.601***	0.291***	0.149***	0.068**
	(0.10)	(0.03)	(0.03)	(0.03)
Good health	2.350***	0.612***	0.633***	0.433***
	(0.14)	(0.04)	(0.04)	(0.04)
University education	0.009	-0.029	0.046	-0.045
	(0.10)	(0.03)	(0.03)	(0.03)
No dependent children	0.247***	0.000	0.081***	0.083***
	(0.09)	(0.03)	(0.03)	(0.03)
Age	-0.010***	0.008***	-0.004***	-0.011***
	(0.00)	(0.00)	(0.00)	(0.00)
Female	0.043	0.082***	0.043	-0.064**
	(0.09)	(0.03)	(0.03)	(0.03)
Household income	0.003***	0.001**	0.000	0.001***
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	11.652***			
	(0.23)			
Prob > F ($H_0 : \gamma_2 = \gamma_3$)	0.000	0.000	0.001	0.002
Baseline predicted subjective wellbeing	13.298			
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Observations	6314	6363	6337	6329

Table 1.5

Mechanism: financial goals and commitments

This table reports the results of OLS and ordered probit regressions. The dependent variables in Columns (1) to (4) are the indicated dimensions of subjective well-being. The independent variables of interest are triple interactions between the financial distress (FD) and high financial knowledge (FK) indicators, as well as indicators capturing whether individuals do not follow through with their own financial goals (Panel A) or their financial commitments towards others (Panel B), respectively. The exact variable definitions can be found in the Appendix. The point estimates are OLS coefficients in Column (1) and ordered probit coefficients in Columns (2) to (4). Robust standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
	Panel A: Financial goals			
FD=0 × FK=0 × Cannot meet financial goals=1	-1.875***	-0.724***	-0.461***	-0.342***
	(0.30)	(0.09)	(0.10)	(0.09)
FD=0 × FK=1 × Cannot meet financial goals=0	-0.208**	-0.070**	-0.066*	-0.091**
	(0.11)	(0.03)	(0.03)	(0.04)
FD=0 × FK=1 × Cannot meet financial goals=1	-2.184***	-0.691***	-0.621***	-0.379***
	(0.39)	(0.11)	(0.11)	(0.11)
FD=1 × FK=0 × Cannot meet financial goals=0	-1.398***	-0.477***	-0.340***	-0.253***
	(0.15)	(0.05)	(0.05)	(0.05)
FD=1 × FK=0 × Cannot meet financial goals=1	-3.133***	-1.008***	-0.717***	-0.538***
	(0.28)	(0.08)	(0.08)	(0.08)
FD=1 × FK=1 × Cannot meet financial goals=0 (λ_6)	-1.965***	-0.676***	-0.484***	-0.380***
	(0.20)	(0.06)	(0.06)	(0.06)
FD=1 × FK=1 × Cannot meet financial goals=1 (λ_7)	-4.710***	-1.349***	-1.100***	-0.892***
	(0.43)	(0.09)	(0.11)	(0.11)
Married or cohabiting	0.574***	0.287***	0.143***	0.064**
	(0.10)	(0.03)	(0.03)	(0.03)
Good health	2.211***	0.581***	0.607***	0.412***
	(0.14)	(0.04)	(0.04)	(0.04)
University education	-0.056	-0.049	0.031	-0.056*
	(0.10)	(0.03)	(0.03)	(0.03)
No dependent children	0.265***	0.008	0.087***	0.086***
	(0.09)	(0.03)	(0.03)	(0.03)

(Continued)

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
	Panel A: Financial goals			
Age	-0.011***	0.008***	-0.005***	-0.011***
	(0.00)	(0.00)	(0.00)	(0.00)
Female	0.027	0.079***	0.040	-0.068**
	(0.09)	(0.03)	(0.03)	(0.03)
Household income	0.003***	0.001**	0.000	0.001***
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	12.012***			
Observations	6314	6363	6337	6329
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Baseline predicted subjective wellbeing	13.298			
Prob > F ($H_0 : \gamma_6 = \gamma_7$)	0.000	0.000	0.000	0.000
	Panel B: Financial commitments			
FD=0 × FK=0 × Cannot meet financial commitments=1	-0.607	-0.283*	-0.096	-0.096
	(0.55)	(0.15)	(0.15)	(0.16)
FD=0 × FK=1 × Cannot meet financial commitments=0	-0.164	-0.048	-0.054	-0.080**
	(0.10)	(0.03)	(0.03)	(0.03)
FD=0 × FK=1 × Cannot meet financial commitments=1	-3.075***	-0.883***	-0.884***	-0.453*
	(0.85)	(0.26)	(0.23)	(0.24)
FD=1 × FK=0 × Cannot meet financial commitments=0	-1.456***	-0.483***	-0.341***	-0.271***
	(0.14)	(0.04)	(0.04)	(0.04)
FD=1 × FK=0 × Cannot meet financial commitments=1	-2.886***	-0.892***	-0.692***	-0.356***
	(0.41)	(0.12)	(0.12)	(0.13)
FD=1 × FK=1 × Cannot meet financial commitments=0 (λ_6)	-2.324***	-0.744***	-0.553***	-0.450***
	(0.19)	(0.05)	(0.05)	(0.06)
FD=1 × FK=1 × Cannot meet financial commitments=1 (λ_7)	-3.937***	-1.076***	-1.001***	-0.639**
	(1.07)	(0.28)	(0.28)	(0.26)
Married or cohabiting	0.586***	0.288***	0.145***	0.067**
	(0.10)	(0.03)	(0.03)	(0.03)
Good health	2.295***	0.598***	0.622***	0.429***
	(0.14)	(0.04)	(0.04)	(0.04)

(Continued)

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
	Panel B: Financial commitments			
University education	-0.011	-0.035	0.040	-0.048
	(0.10)	(0.03)	(0.03)	(0.03)
No dependent children	0.261***	0.005	0.085***	0.084***
	(0.09)	(0.03)	(0.03)	(0.03)
Age	-0.010***	0.008***	-0.005***	-0.011***
	(0.00)	(0.00)	(0.00)	(0.00)
Female	0.014	0.073***	0.036	-0.068**
	(0.09)	(0.03)	(0.03)	(0.03)
Household income	0.003***	0.001*	0.000	0.001***
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	11.776***			
	(0.23)			
Observations	6314	6363	6337	6329
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Baseline predicted subjective wellbeing	13.298			
Prob > F ($H_0 : \gamma_6 = \gamma_7$)	0.136	0.243	0.107	0.476

Table 1.6

Instrumental variable analysis

This table reports the Generalised Method of Moments (GMM) instrumental variable regression results for the second stage using Lewbel's (2012) method, where all included interactions between financial distress and financial knowledge are instrumented with various combinations of heteroskedasticity-based instruments and an external instrument capturing whether the respondents' parent completed university education. In addition to the baseline set of controls (i), variables are included that capture the respondents' recollections on interactions with their parents regarding financial education (ii). The reported coefficients are GMM point estimates. Robust standard errors are reported in parentheses. The exact variable definitions are reported in the Appendix. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Aggregate subjective well-being			
	(1)	(2)	(3)	(4)
Financial distress=0 × High financial knowledge=1	-0.246*	-0.231*	-0.278**	-0.263**
	(0.13)	(0.13)	(0.12)	(0.12)
Financial distress=1 × High financial knowledge=0	-0.721***	-0.718***	-0.707***	-0.704***
	(0.10)	(0.10)	(0.10)	(0.10)
Financial distress=1 × High financial knowledge=1	-1.438***	-1.423***	-1.363***	-1.346***
	(0.22)	(0.22)	(0.21)	(0.21)
(i) Baseline controls:				
Married or cohabiting	0.378***	0.378***	0.378***	0.378***
	(0.04)	(0.04)	(0.04)	(0.04)
Good health	0.761***	0.762***	0.768***	0.770***
	(0.06)	(0.06)	(0.06)	(0.06)
University education	-0.011	-0.014	-0.008	-0.011
	(0.04)	(0.04)	(0.04)	(0.04)
No dependent children	0.045	0.046	0.038	0.039
	(0.04)	(0.04)	(0.04)	(0.04)
Age	0.011***	0.011***	0.011***	0.011***
	(0.00)	(0.00)	(0.00)	(0.00)
Female	0.065*	0.067*	0.053	0.055
	(0.04)	(0.04)	(0.04)	(0.04)
Household income	0.001*	0.001*	0.001**	0.001**
	(0.00)	(0.00)	(0.00)	(0.00)
(ii) Childhood parent financial education controls:				
Discussed family financial matters with me	0.085**	0.084**	0.089**	0.088**
	(0.04)	(0.04)	(0.04)	(0.04)
Spoke to me about the importance of saving	0.060	0.060	0.055	0.055
	(0.05)	(0.05)	(0.04)	(0.04)
Discussed how to establish a good credit rating	0.128***	0.129***	0.114***	0.115***
	(0.04)	(0.04)	(0.04)	(0.04)
Taught me how to be a smart shopper	0.087**	0.087**	0.103**	0.102**
	(0.04)	(0.04)	(0.04)	(0.04)
Taught me that my actions determine my success in life	0.200***	0.198***	0.188***	0.186***
	(0.05)	(0.05)	(0.05)	(0.05)

(Continued)

	Aggregate subjective well-being			
	(1)	(2)	(3)	(4)
Provided me with a regular allowance	0.010 (0.04)	0.011 (0.04)	0.007 (0.03)	0.008 (0.03)
Provided me with a savings account	-0.028 (0.04)	-0.027 (0.04)	-0.029 (0.04)	-0.028 (0.04)
Constant	2.871*** (0.13)	2.869*** (0.13)	2.850*** (0.12)	2.848*** (0.12)
Observations	6363	6363	6363	6363
External instrument (University-educated parents)	No	Yes	No	Yes
Controls for generated instruments	(i)	(i)	(i), (ii)	(i), (ii)
R^2	0.19	0.19	0.19	0.19
Cragg-Donald F statistic	30.96	29.95	17.22	16.99
Stock-Yogo maximal IV size critical value:				
5%	19.67	19.77	20.59	20.61
10%	10.63	10.65	10.81	10.81
Hansen J statistic	13.64	14.46	36.53	37.76
Prob > J	0.75	0.76	0.58	0.57
J degrees of freedom	18	19	39	40

Table 1.7

Propensity score matching

This table reports the coefficient estimates of probit regressions. The results are based on the subsample of respondents with recent experiences of financial distress. The dependent variables are indicators taking the value one when the respondents' financial knowledge is greater than the full sample average (including respondents outside of financial distress), and zero otherwise. Column (1) reports the results for all repsondents in financial distress, which is the model to generate the individual propensity scores, while Column (2) reports the results for the sample obtained after matching on the propensity scores. The variables used in propensity score matching are the those shown in this table. The Appendix provides exact definitions of all variables. Standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	High financial knowledge	
	Propensity score model	Matched sample model
	(1)	(2)
Married or cohabiting	0.053	-0.033
	(0.07)	(0.09)
Good health	0.032	-0.035
	(0.08)	(0.10)
University education	0.465***	0.126
	(0.08)	(0.09)
No dependent children	0.054	0.079
	(0.07)	(0.09)
Age	0.016***	0.004
	(0.00)	(0.00)
Female	-0.325***	0.016
	(0.07)	(0.08)
Household income	0.007***	0.000
	(0.00)	(0.00)
Constant	-1.701***	-0.240
	(0.13)	(0.17)
Baseline predicted probability	.277	.500
χ^2	221.270	5.284
Prob > χ^2	0.000	0.625
Observations	1833	1016

Table 1.8

Covariate balancing

This table reports an analysis of the differences in means for the baseline controls by financial knowledge levels (FK) before and after propensity score matching. Financial knowledge is defined as low (high) when it is below (above) the average of the full sample. The columns report the mean for the respective variable, together with the difference in means between financial knowledge levels (Δ) and the results of t-tests for differences in means. The exact variable definitions can be found in the Appendix. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Full sample				Matched sample			
	FK				FK			
	Low	High	Δ	t	Low	High	Δ	t
Married or cohabitating	0.51	0.63	-0.12	-4.47***	0.62	0.63	-0.01	-0.26
Good health	0.71	0.76	-0.05	-2.05**	0.76	0.76	0.01	0.22
University education	0.16	0.38	-0.23	-10.73***	0.34	0.38	-0.04	-1.44
No dependent children	0.39	0.41	-0.02	-0.77	0.39	0.41	-0.02	-0.70
Age	43.12	49.13	-6.00	-7.21***	47.72	49.13	-1.40	-1.40
Female	0.56	0.42	0.14	5.24***	0.41	0.42	-0.01	-0.38
Household income	43.65	66.48	-22.82	-11.44***	65.11	66.48	-1.37	-0.52
Observations	1325	508			508	508		

Table 1.9

Baseline results on PSM sample

This table reports the results of OLS and probit regressions. The sample includes all respondents with high financial knowledge together with the matched low financial knowledge group of respondents obtained through propensity score matching. The dependent variables in Columns (1) to (4) are the indicated dimensions of subjective well-being. The independent variables of interest is high financial knowledge. The point estimates are OLS coefficients in Column (1) and ordered probit coefficients in Columns (2) to (4). Robust standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Aggregate subjective well-being	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
High financial knowledge	-0.484*	-0.172***	-0.074	-0.105
	(0.25)	(0.07)	(0.07)	(0.07)
Married or cohabiting	0.700**	0.328***	0.180**	-0.029
	(0.28)	(0.07)	(0.07)	(0.07)
Good health	2.372***	0.550***	0.577***	0.422***
	(0.30)	(0.08)	(0.08)	(0.08)
University education	-0.031	-0.026	0.017	-0.073
	(0.27)	(0.07)	(0.07)	(0.07)
No dependent children	0.375	-0.036	0.122*	0.168**
	(0.27)	(0.07)	(0.07)	(0.07)
Age	-0.014*	0.007***	-0.006**	-0.010***
	(0.01)	(0.00)	(0.00)	(0.00)
Female	-0.217	-0.019	-0.022	-0.057
	(0.25)	(0.07)	(0.07)	(0.07)
Household income	0.004	0.001	-0.000	0.001
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	9.921***			
	(0.55)			
Baseline predicted subjective wellbeing	11.537			
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Observations	1005	1007	1009	1009

Table 1.10

Mechanism on PSM sample

This table reports the results of OLS and ordered probit regressions. The sample includes all respondents with high financial knowledge together with the matched low financial knowledge group of respondents obtained through propensity score matching. The dependent variables in Columns (1) to (4) are the indicated dimensions of subjective well-being. The independent variables of interest are triple interactions between the financial distress (FD) and high financial knowledge (FK) indicators, as well as indicators capturing whether individuals do not follow through with their own financial goals (Panel A) or their financial commitments towards others (Panel B), respectively. The exact variable definitions can be found in the Appendix. The point estimates are OLS coefficients in Column (1) and ordered probit coefficients in Columns (2) to (4). Robust standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
	Panel A: Financial goals			
FK=0 × Cannot meet financial goals=1	-1.640***	-0.487***	-0.325***	-0.301**
	(0.44)	(0.12)	(0.12)	(0.12)
FK=1 × Cannot meet financial goals=0	-0.233	-0.128*	-0.010	-0.057
	(0.27)	(0.07)	(0.07)	(0.07)
FK=1 × Cannot meet financial goals=1	-3.032***	-0.817***	-0.634***	-0.577***
	(0.43)	(0.12)	(0.12)	(0.12)
Married or cohabiting	0.674**	0.331***	0.176**	-0.036
	(0.27)	(0.07)	(0.07)	(0.07)
Good health	2.124***	0.500***	0.536***	0.382***
	(0.30)	(0.08)	(0.08)	(0.08)
University education	-0.217	-0.075	-0.022	-0.108
	(0.27)	(0.07)	(0.07)	(0.07)
No dependent children	0.323	-0.051	0.113	0.161**
	(0.26)	(0.07)	(0.07)	(0.07)
Age	-0.015*	0.007***	-0.006***	-0.010***
	(0.01)	(0.00)	(0.00)	(0.00)
Female	-0.287	-0.039	-0.040	-0.072
	(0.25)	(0.07)	(0.07)	(0.07)
Household income	0.004	0.001*	-0.000	0.001*
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	10.584***			
	(0.54)			

	Aggregate subjective wellbeing	Satisfied with my life	Optimistic about my future	Work hard today, successful in future
	(1)	(2)	(3)	(4)
	Panel A: Financial goals			
Observations	1005	1007	1009	1009
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Baseline predicted subjective wellbeing	11.537			
	Panel B: Financial commitments			
FK=0 × Cannot meet financial commitments=1	-0.827	-0.204	-0.282*	0.014
	(0.64)	(0.17)	(0.17)	(0.17)
FK=1 × Cannot meet financial commitments=0	-0.474*	-0.171**	-0.075	-0.095
	(0.26)	(0.07)	(0.07)	(0.07)
FK=1 × Cannot meet financial commitments=1	-2.169**	-0.550**	-0.552**	-0.290
	(0.85)	(0.23)	(0.23)	(0.23)
Married or cohabiting	0.693**	0.327***	0.178**	-0.029
	(0.27)	(0.07)	(0.07)	(0.07)
Good health	2.324***	0.540***	0.564***	0.422***
	(0.30)	(0.08)	(0.08)	(0.08)
University education	-0.070	-0.035	0.006	-0.077
	(0.27)	(0.07)	(0.07)	(0.07)
No dependent children	0.368	-0.038	0.121*	0.168**
	(0.27)	(0.07)	(0.07)	(0.07)
Age	-0.015*	0.006***	-0.006***	-0.010***
	(0.01)	(0.00)	(0.00)	(0.00)
Female	-0.256	-0.029	-0.034	-0.060
	(0.25)	(0.07)	(0.07)	(0.07)
Household income	0.004	0.001	-0.000	0.001
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	10.126***			
	(0.56)			
Observations	1005	1007	1009	1009
Model	OLS	Ordered probit	Ordered probit	Ordered probit
Baseline predicted subjective wellbeing	11.537			

1.A Variable definitions

Variable	Definition
Panel A: Baseline analysis	
Aggregate subjective wellbeing	This variable is the sum of the subjective well-being dimensions “Satisfied with my life”, “Optimistic about my future” and “Work hard today, successful in future”. The variable takes integer values in the range from 3 to 21. Higher values indicate greater subjective wellbeing.
Satisfied with my life	This variable captures the respondens agreement or disagreement with the statement “I am satisfied with my life”. Possible answers are on the 7-point Likert scale from strongly agree to strongly disagree. Answers are mapped to integer values ranging from 1 to 7, where greater values indicate stronger agreement.
Optimistic about my future	This variable captures the respondens agreement or disagreement with the statement “I am optimistic about my future”. Possible answers are on the 7-point Likert scale from strongly agree to strongly disagree. Answers are mapped to integer values ranging from 1 to 7, where greater values indicate stronger agreement.
Work hard today, successful in future	This variable captures the respondens agreement or disagreement with the statement “If I work hard today, I will be more successful in the future”. Possible answers are on the 7-point Likert scale from strongly agree to strongly disagree. Answers are mapped to integer values ranging from 1 to 7, where greater values indicate stronger agreement.
Financial distress	This variable captures whether respondents experienced any financial distress “often” or “sometimes” in the past 12 months. Six dimensions of material hardship are captured. Specifically, respondents are asked whether each of the following statements applied “often”, “sometimes” or “never” to them in the past 12 months: 1. I worried whether our food would run out before I got money to buy more. 2. The food that I bought just didnt last and I didnt have money to get more. 3. I couldn’t afford a place to live. 4. I or someone in my household needed to see a doctor or go to the hospital but did not go because we couldnt afford it. 5. I or someone in my household stopped taking a medication or took less than directed due to the costs. 6. One or more of my utilities was shut off due to non-payment. This financial distress variable takes the value one if respondents answered “often” or “sometimes” to any of these six questions, and zero otherwise.
Financial knowledge	This variable captures the 10-item Knoll and Houts (2012) Financial Knowledge Scale. The variables ranges from -2.053 to 1.267. Higher values indicate greater financial knowledge.
Cannot meet financial goals	This variable takes the value one if the respondent chooses the response “very little” or “Not at all” regarding the satement “I follow-through on financial goals I set for myself”, and zero otherwise.
Cannot meet financial commitments	This variable takes the value one if the respondent chooses the response “very little” or “Not at all” regarding the satement “I follow-through on my financial commitments to others”, and zero otherwise.
Married or cohabiting	This variable takes the value one if respondents are married live with their partner, and zero otherwise.

(Continued)

Variable	Definition
Good health	This variable takes the value one if respondents indicate that their health in general is good, very good or excellent, and zero otherwise.
University education	This variable takes the value one if respondents obtained a bachelor's degree or a graduate/professional degree, and zero otherwise.
No dependent children	This variable takes the value one if respondents indicate that they have no children they financially support, and zero otherwise.
Age	The survey captures age of respondents in seven non-overlapping age brackets, between 18 and 74, and the eighth age bracket captures respondents older than 75. The variable Age for a respondent is equal to the midpoint age of the age bracket the respondents belong to. For respondents in the (first) eighth age bracket, the variable takes values equal to the (upper) lower limit of the age bracket.
Female	This variable takes the value one if the respondent is female, and zero otherwise.
Household income	The survey captures the income level of respondents (in 1000s), classified into nine non-overlapping income brackets and the ninth income bracket captures income of $150,000 or above. The variable Income for a respondent is equal to the midpoint income of the income bracket the respondents belong to. For respondents in the lowermost income bracket, the variable takes values equal to the upper limit of the income bracket. Similarly, for the uppermost income bracket, the variable takes values equal to the lower limit of the income bracket.
Panel B: Instrumental variable analysis	
Parent university education	This variable takes the value one if respondents' parent obtained a bachelor's degree or a graduate/professional degree, and zero otherwise.
Discussed family financial matters with me	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.
Spoke to me about the importance of saving	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.
Discussed how to establish good credit rating	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.
Taught me how to be a smart shopper	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.
Taught me that my actions determine my success in life	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.
Provided me with a regular allowance	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.
Provided me with a savings account	This variable takes the value one if the respondents' parent interacted with the respondent during childhood in the indicated manner, and zero otherwise.

2

Social Norm Enforcement and Mental Health

Highlights
∗ High welfarism – favourable attitudes towards the welfare state and benefit recipients – is associated with an increased prevalence of mental health problem of 13 percentage points, a 39% increase against predicted population-wide levels.
∗ To test the hypothesis that expressing high welfarism constitutes a deviation from work norms, resulting in social sanctioning and associated decreases in mental health, the changes in magnitude of the relationship under conservative governments is tested, who are known to be tougher on welfare. This reveals an intensifying of the relationship relative to labour governments.
∗ Under conservative governments, the relationship is stronger for employed individuals, which exhibit high conformity with work norms, relative to those not in employment, suggesting that in-group social norm deviations are sanctioned more severely.
∗ Additional analysis reveals that the observed relationships are stronger for females than for males. Further, the combination of high welfarism and mental health problems is associated with more favourable attitudes towards voting, suggesting increased motivations to induce social change.

2.1 Introduction

Policy makers are concerned about the significant costs that mental health problems inflict, rising to 4% of GDP (600 billion €) across the European Union (OECD/EU, 2018). Research finds that an array of factors, from individual characteristics and circumstances to wider economic conditions, influences mental health (Bridges & Disney, 2010; Currie et al., 2015; Gathergood, 2012; Grip et al., 2011). In this thesis chapter, we add to this body of research by investigating the effect of social norm enforcement (Bernhard et al., 2006; Goette et al., 2006) on mental health. As Posner (1997) notes: "Norms are also enforced by expressions of disapproval, by ridicule, and in extreme cases, by ostracism. The efficacy of the milder 'sanctions' lies in their implicit threat of ostracism, that is, of refusal of advantageous transactions." (p. 366) In light of the potential psychological repercussions these enforcement means can have, we thus ask: can the enforcement of social norms harm mental health?

Specifically, we use representative data on welfare attitudes in the UK to study the link between social norms enforcement and mental health. We propose that in societies that assign a normative value to work (Lindbeck, Nyberg & Weibull, 1999, 2003), expressing positive welfare attitudes can imply a deviation from social norms, resulting in social sanctioning and corresponding deteriorations in mental health. The UK provides an ideal setting to study this notion, as UK government welfare programmes constitute a significant mechanism of poverty alleviation; but, at the same time, the global stigmatization of welfare (see for example Field, 2002; Niskanen, 1996) has been particularly manifest in the UK society:

> 'Fairness', declared George Osborne, the then chancellor, in 2012, 'is about being fair to the person who leaves home every morning to go out to work and sees their neighbour still asleep, living a life on benefits.' Newspapers printed story after story about welfare fraudsters pinching

> from the public purse, from the woman with two Samoyed dogs who collected thousands of pounds a month and claimed, 'It's not worth my while working,' to the man who collected disability benefit while competing in bodybuilding contests. In 2014 one in ten Britons tuned in to 'Benefits Street', a documentary which featured welfare recipients drinking and fighting on rubbish-strewn streets. [...] Back then it felt impossible to be too mean to benefit claimants. They were a political piñata: whack them and votes fell out. (The Economist, 2019)

Thus, as public understanding, goes "benefit scroungers" claim welfare and live off the work of others (Geiger, 2017).[1] The variations in the UK political environment induced by the changes from labour to conservative governments in our sample period, and the concomitant changes in policy stances and public rhetoric on welfare, provide a plausible source of variation in the cost of sanctioning positive welfare attitudes and thus the consequent enforcements of the social norm that engaging in paid work is a normative good.

In our analysis, we use data from the British Social Attitudes (BSA) surveys from the years 2000, 2003, 2006, 2007 and 2013. These years capture repetitions of questions pertinent to our study. The surveys do not track individuals over time but rather sample a representative set of respondents in Britain with every new survey iteration; therefore, we pool all responses and obtain a cross-sectional sample of 3031 individuals, after removing respondent observations with missing values. The data elicit attitudes towards a wide range of issues, and our key explanatory variable is derived from respondents' attitudes towards the welfare state and welfare recipients, which is labelled welfarism. The individual questions that enter the construction of this aggregate welfarism score capture the respondents' views on whether 1) benefits discourage independence, 2) cutting benefits would damage too many lives, 3) the government should spend more money on benefits,

[1] In the UK, in light of the considerable public spending on welfare, "throughout his premiership, David Cameron, along with his chancellor, George Osborne, kept the opposition between 'hardworking people' and lazy benefit claimants right at the centre of their messaging on spending cuts." (de Vries & Reeves, 2017, in The Guardian).

4) people on the dole fiddle or way or another, 5) the unemployed could find a job if they really tried, 6) welfare discourages mutual support, 7) welfare recipients do not really deserve any help and 8) the welfare state is one of Britain's proudest achievements. The possible answers lie on a 5-point Likert scale ranging from "strongly disagree" to "strongly agree", which we map to the integers 1 to 5, where higher values indicate more favourable attitudes. Summing all answers yield the respondents' welfarism scores, ranging from 8 to 40; a composite variable shown to be of high reliability (Curtice, Clery, Perry, Phillips & Rahim, 2019). Our main outcome variable derives from a direct survey question capturing whether medical advice on a mental health problem has ever been sought, and we map "yes" answers to the integer one, and "no" answers to zero. From the BSA surveys, we further obtain key individual- and household-level characteristics, such as holding a University degree, the household income (in logs), being married or cohabitating, being female and living by oneself in a single household.

The relationship between high welfarism and mental health can be interpreted in the light of the literature on social norm enforcement and social identity. Identity is understood as identification with a social category and acceptance of the accompanying behaviours and attributions deemed appropriate. Accordingly, we provide a theoretical framework to guide the investigation of the basic relationship we posit in a simple extension of the model by Akerlof and Kranton (2000) to the context of welfarism and mental health problems. We show that work-related "identity norms" can cause mental health problems as a result of social sanctioning following from the expressions of positive welfare attitudes. Specifically, mental health problems can arise as individuals that express positive welfare attitudes threaten other individuals in their social identities and, accordingly, become sanctioned for their expression of welfare attitudes. The framework entails the testable implication that social sanctioning increases as its cost decreases.

Our empirical analysis begins by establishing the relationship between welfarism and mental health problems. As descriptive evidence reveals evidence of non-linearities in the relationship between welfarism and mental health problems, we construct a high welfarism indicator variable that takes the value one for individuals in the top 30% of the welfarism score distribution, and zero otherwise. The average marginal effects (AMEs) of 0.130 estimated from the baseline probit model suggests that high welfarism is associated with an increase in the mental health problem probability of 13 percentage points. In order to attenuate endogeneity concerns and establish the direction of the potential bias in the baseline regression of mental health problems on high welfarism, we further test the relationship in a Maximum Likelihood framework in which we explicitly model the potentially endogenous regressor, high welfarism. Specifically, we specify a two-stage recursive probit framework and identify the association between high welfarism and mental health problems in two ways: 1) through the frameworks non-linearities and 2) by instrumenting high welfarism with a set of dummy variables capturing the respondents' identification with political parties. Given the inclusion of a wide range of control variables, the instruments in approach 2) likely meet the instrumental variable relevance and exclusion restrictions. The obtained AME are equal to 0.425 and 0.364, suggesting that endogeneity introduces a downward bias in the baseline estimates, which thereby establish a lower bound on the effects. As the non-linear and instrumental variable identification approaches yield the same conclusions, our results do not rest on the validity of the instrumental variables.

Next, we exploit variation in the ruling governing party in the UK over time to capture changing political environments for welfare and test the implication that a reduction in the cost of sanctioning yields higher rates of mental health problems.[2] By testing whether the link between high welfarism and mental health

[2]Insofar as changes in majorities of political parties reflect evolving voter preferences due to economic conditions, our study is related to the economics literature on endogenous preferences

problems intensifies under conservative rule, we provide evidence in favour of the notion that expressing positive welfare attitudes can be sanctioned by other individuals as a result of them feeling threatened in their identities. Intuitively, when the UK government is drawn from the conservative party, the typical result is a political environment in which welfare spending and benefit recipients are viewed less favourably - in its extreme, resulting in political scapegoating.[3] We confirm the testable implication of our theoretical framework by finding that when the government is drawn from the conservative party, the relationship between high welfarism and mental health problems intensifies.

To identify the role of identity, we further split individuals by whether their labour market characteristics conform to work norms. Respondents that are employed or receive no welfare benefits conform to the societal ideal that working is a normative good, thereby adhering to the social identities associated with the working population. For these respondents, expressing positive welfare attitudes, constituting a deviation from their social identities, and can be sanctioned more strongly by other individuals subscribing to the same identity. In harsher political environments towards welfare, this sanctioning can occur more frequently, given the reduction in its cost. The regression results are consistent with these notions. Under Tory governments, high conformity, high welfare individuals report a higher prevalence of mental health problems (11.5 to 17.3 percentage point increase) than low conformity, high welfare individuals (8.5 to 11.5 percentage points).

We further provide evidence on the role of gender in an additional set of analyses. We find that the observed relationships are more pronounced for females than

(for example Alesina & Giuliano, 2015; Bowles, 1998; Bowles & Polania-Reyes, 2012; Callen, Isaqzadeh, Long & Sprenger, 2014; Giuliano & Spilimbergo, 2014; Malmendier & Nagel, 2011; Tanaka, Camerer & Nguyen, 2010). However, we do not pursue this point in this study.

[3] For some welfare claimants, counter-intuitively, this can result in more negative views of other welfare recipients, rather than an increased sense of solidarity with those in similar positions. The psychological literature labels this "cognitive distancing": a psychological coping mechanism to reconcile one's welfare receipt with the negative public opinion of welfare (Lott, 2002; Reutter et al., 2009), such as the image of benefit scroungers.

for males and that the gap in negative mental health outcomes increases between genders for survey responses elicited under conservative governments versus labour governments. Moreover, we document significant associations for high welfarism individuals with reports of mental health problems and the notion that voting is a duty, whereas these associations are not statistically significant, given no mental health problems. The results of this analysis suggest that respondents finding themselves worse off under conservative rule recognise voting as a mean to induce societal change.

Our analysis proceeds as follows. Section 2.2 describes the related literature; Section 2.3 presents our theoretical framework; Section 2.4 presents our data; Section 2.5 describes the results of our empirical analysis; Section 2.6 provides additional analysis on the role of gender and implications for attitudes towards voting; Section 2.7 discusses limitations and future research; while Section 2.8 concludes.

2.2 Related literature

Mental health influences financial outcomes and is also affected by economic circumstances and choices. As cited in Balloch, Engels and Philip (2020), a wide range of links exist between mental health and individual as well as macroeconomic circumstances. First, mental health influences financial outcomes. In an influential study, Gathergood (2012) shows that problematic mortgage debt significantly deteriorates mental health and is moderated by local house price movements that buffer the severity of mortgage payment arrears. The dependence of mental health outcomes on economic developments outside individual control are also developed in (Grip et al., 2011), who show that unexpected pension reform with negative implications on the individual income replacement rates in retirement affects the mental health of those close to the retirement age. More broadly, widening access to credit has net positive effects on mental health outcomes (Karlan & Zinman, 2010); while recessions impact mental health negatively, with differential impacts felt along the socio-economic gradient (Currie et al., 2015).

Second, financial outcomes are affected by mental health. Bogan and Fertig (2013) show that asset allocation is significantly determined by mental health. For instance, the share of household investments in risky assets, such as stock, mutual funds and investment trusts, decreases as mental health deteriorates; whereas single women increase their holdings of safe assets, such as, for example, savings and money market accounts. Further, Bogan and Fertig (2018) find that mental health problems decrease the individual likelihood to invest in retirement accounts and build retirement savings. Berkowitz and Qiu (2006) find that health status affects the allocation of household financial assets and is an important factor determining households' financial wealth. Further, Bogan, Fertig and Just (2019) show that psychological distress makes it more difficult for individuals to retain salaried jobs, thereby increasing the likelihood of self-employment. Balloch et al. (2020) suggest that such distress negatively affects the ability to accumulate wealth.

This thesis chapter is further related to the literature in identity economics that originates from the seminal contribution by Akerlof and Kranton (2000). Studies therein highlight the link between social identity and role of social norm enforcement (Bernhard et al., 2006; Goette et al., 2006); the effect of ethnic, racial, and gender category norms on time and risk preferences (Benjamin, Choi & Strickland, 2010); the link between personal identity and moral choices (Bénabou & Tirole, 2011); the effects of gender identity on the relative distribution of household incomes and marriage quality (Bertrand, Kamenica & Pan, 2015); and how identity influences consumption decisions (Bursztyn, Ferman, Fiorin, Kanz & Rao, 2017).

The social categories of workers and benefit scroungers are consistent with the pronounced and universal emphasis of the normative value of work in developed countries. In a theoretical analysis, (Lindbeck et al., 1999) show that both social norms and economic incentives determine the decision whether to take up welfare. Specifically, the social norm under consideration is that to be economically self-sufficient and use own earnings to sustain oneself. The larger the share of welfare recipient, the less pronounced this social norm becomes. Lindbeck et al. (2003) continue to how economic incentives and such social norms to work affect the endogenous determination of social insurance. Their findings suggest that if the norm to work is weakened through a higher share of welfare recipient, voters will decrease the generosity of welfare payment and the welfare stigma increases (Moffitt, 1983; Pescosolido & Martin, 2015).

2.3 Theoretical framework

In this section, we provide a simple theoretical framework, extending the model by Akerlof and Kranton (2000): we define the social categories of "worker" and "benefit scrounger" and analyse the game-theoretic interaction between two individuals, showing how mental health problems can arise as a result of the expression of welfare attitudes, which can attract sanctioning by others. We further derive a simple theoretical implication that we test in the empirical analysis of Section 2.5.

2.3.1 Social sanctions and mental health

In the following, we follow the notation by Akerlof and Kranton (2000). In our analysis, two social categories exist: worker and benefit scrounger, which we collect in the vector $\mathrm{C} = (\text{worker, benefit scrounger})'$. Workers have a higher social status than benefit scroungers. Each social category entails prescriptions, P, that indicate what behaviours or characteristics are appropriate for members of the given social category. In our context, these include, for instance, that workers stand on their own two feet and do not claim benefits to support their livelihoods. In contrast, the attributes of benefit scroungers include claiming benefits and living off the work of others.[4] Social identities then refer to which social categories individuals assign themselves and others.

Given these social identities, the expression of positive welfare attitudes, which we label high welfarism, by a given individual, HW_j, and by others, HW_{-j}, influence utility directly, but also indirectly through their effects in the context of j's own social identity, I_j:

$$U_j = U_j(HW_j, HW_{-j}, I_j) \tag{2.1}$$

$$I_j = I_j(HW_j, HW_{-j} \mid c_j, \epsilon_j, \mathrm{P}) \tag{2.2}$$

In particular, Equation 2.2 can be interpreted as utility externalities that arise from choices to express or not to express high welfarism, conditional on the assignments by j of herself and others to social categories, c_j, and the extent to which j conforms to her own social identity, ϵ_j. Depending on the particular social identity assignments, the effects on utility differ given whether a worker or a benefit scrounger express their attitudes to welfare.

Given these effects of others' choices in light of social identity, a simple game-theoretic interaction serves to show how workers can face incentives to sanction other workers that express high welfarism, as a result of feeling threatened in their identity. Sanctioning offending workers then restores the loss in self-image. The

[4]It is also reflected in the design of welfare programmes. For example, the UK's Department for Work and Pensions (2020) states that its new welfare system, "[Universal Credit,] has been introduced to give you the support you need to find and progress in work. We want you to be able to benefit from all the positives that work brings." (p. 1) This is suggestive of the view that benefit claimants have an innate aversion to work.

sanctioned workers consequently suffer a resulting utility loss due to an increase in mental health problems.

Specifically, let there be two workers, i and j. Worker i does not hold positive views on welfare and earns zero utility from expressing positive views. Worker j does hold positive welfare views, or high welfarism HW, and earns utility $V(HW)$ from expressing her attitudes; otherwise, she earns a utility of zero. Further, let both i and j correctly think of themselves and the other as workers, where worker identities entail the subscription, P, that "workers should earn their income" and "only benefit scroungers receive welfare and live off the work of others". So anyone who expresses views contrary to these subscriptions is in violation of the worker identity. In j's case, this would induce a loss in identity of I_s, where the subscript "s" stands for self.

If i and j are paired in social interaction, expressing positive welfare attitudes on the part of j diminishes i's worker identity and i has a loss in utility I_o where the subscript "o" stands for "others." This identity externality equals zero if j does not express any positive welfare attitudes. After observing that j expressed positive welfare attitudes, i may respond by sanctioning j at an effort, e. This causes j to incur utility losses due to deterioration in her mental health, MH.

Figure 2.1 represents this interaction between individual i and j. j can choose to express her high welfarism, HW, while i can decide to engage in sanctioning this expression at an effort, e. This game tree has four important subgame perfect equilibria (Akerlof & Kranton, 2000). First, consider the case in which the identity cost of expressing high welfarism is greater than the utility earned from such expression, $V(HW) < I_s$. In this case, j does not express her views and i is not faced with the decision of whether to sanction j. She then earns zero utility, while i earns her baseline utility, V. Second, assume that for j the utility earned from expressing high welfarism is greater than the identity costs incurred, $V(HW) > I_s$, and that for i the effort to sanction j is greater than identity externalities incurred, $I_o < e$. In this case, j expresses her views and earns utility $V(HW) - I_s$ while i does not sanction j and earns utility $V - I_o$.

Now assume that for j the utility earned from expressing high welfarism is greater than the identity costs incurred, $V(HW) > I_s$, and that for i it is more costly to suffer identity externalities than to engage in the effort to sanction i, $I_o > e$. Two further equilibrium outcomes are possible depending on the utility cost that sanctioning imposes on j. On the one hand, in case of the third outcome, if the utility net identity cost is greater than the mental health cost of being sanctioned, $V(HW) - I_s > MH$, j will express her views and earn utility $V(HW) - I_s - MH > 0$. On the other hand, in case of the fourth outcome, if utility after accounting for sanctioning is less than zero, she will refuse to express her views and earn zero utility. Accordingly, i will earn utility $V - e$ or V, respectively.

2.3.2 Testable implication

The model outlined in the previous section yields an implication that we will test empirically in the next sections. Specifically, consider that the model posits that sanctioning, in response to expressions of high welfarism, HW, by other workers, does not occur if its cost, e, is greater than the identity externalities incurred otherwise: $e > I_o$.

Now, consider a reduction in the costs of social sanctions. If these costs decrease, and it remains an optimal decision to express high welfarism after accounting for social sanctions, some workers that ignored expressions of high welfarism before will now find it optimal to sanction others. Consequently, an increase in mental health problems associated with this increase in sanctioning will become observable.

The empirical test in the later sections amounts to confirming an increase in mental health problems for high welfarism individuals as the cost of sanctioning reduces. The empirical challenge lies in identifying factors that plausibly reduce the cost of engaging in sanctioning for those workers that suffer identity externalities. One such factor we will consider is the variation in the political environment with respect to welfare friendliness due to changes in the governing parties.

2.4 Data

2.4.1 Data sample

We use data from the British Social Attitudes (BSA) surveys: a repeated cross-sectional survey, where a new set of respondents is obtained with every survey, representative of the attitudes of the British population towards a variety of societal issues. The majority of questions that are asked change year-to-year, yet some are periodically repeated to enable comparisons of attitudes over time. The years 2000, 2003, 2006, 2007 and 2013 contain repetitions of the questions that collect the information we are interested in: respondents' answers on questions concerning welfare recipients and the welfare state, mental health problems and socio-economic information, including detailed information on employment status and welfare receipt. As the BSA series of surveys do not have a panel dimension, we pool these years. After excluding observations with missing data in our baseline set of variables, we are left with a cross-sectional data sample of 3031 respondent observations for analysis.

2.4.2 Variable descriptions and summary statistics

Table 2.1 shows the descriptions and summary statistics (mean, minimum, maximum and standard deviation) for the variables in our data sample. The following subsections provide descriptions of our key explanatory variable, welfarism, our key outcome variable, mental health problems, as well as other individual- and household-level characteristics.

2.4.2.1 Welfarism

The key explanatory variable is labelled welfarism; a scale commonly used in the BSA data. Curtice et al. (2019) test the scale for reliability using Cronbach's alpha and find it to be highly reliable. It captures respondents' overall welfare attitudes and is derived from eight questions which that the respondents' degree of agreement or disagreement regarding various aspects of the welfare state and welfare recipients. Specifically, these questions capture the respondents' views on whether 1) benefits discourage independence, 2) cutting benefits would damage too many lives, 3) the government should spend more money on benefits, 4) people on the dole fiddle or way or another, 5) the unemployed could find a job if they really tried, 6) welfare discourages mutual support, 7) welfare recipients do not really deserve any help and 8) the welfare state is one of Britan's proudest achievements. The possible answers lie on a 5-point Likert scale ranging from "strongly disagree" to "strongly agree", which we map to the integers 1 to 5, where higher values indicate more favourable attitudes. Summing all answers yield the respondents' welfarism scores, ranging from 8 to 40, which, unlike the scale used in Curtice et al. (2019), we do not divide by eight to retain variability in the data. The mean score in our data equals 23.84, which amounts to 49.5% ($\approx (23.84 - 8)/(40 - 8)$) of the maximum positive overall welfare attitude.

2.4.2.2 Mental health problems

The key outcome in our analysis is the mental health problems variable. The BSA asks a direct question as to whether medical advice on a mental health problem has ever been sought, and we map "yes" answers to the integer one, and "no" answers to zero. This binary measure is easy to administer and therefore available in six BSA surveys. However, Bogan and Fertig (2018) use a similar question in their study of how mental health affects retirement savings behaviour and point out a number of limitations of this measurement approach, many of which apply in this context. In particular, the nature, timing and intensity of the mental health issue in question is not elicited and, because the question requires engagement with a medical practitioner, it captures socio-economic along with mental health status. Attenuating the timing issue – having *ever* sought medical advice – Aneshensel, Estrada, Hansell and Clark (1987) find that self-reports of lifetime diagnosis are often inconsistent over time and are more likely to capture a current diagnosis than a past diagnosis, which facilitates the analysis of trends in the changes of the answer distribution over time. Due to data limitations, we make use of this variable in our analysis; and recognizing its limitations, Table 2.1 indicates a prevalence of mental health problems in our sample of 33%.

2.4.2.3 Other characteristics

Table 2.1 reports a further set of individual- and household-level characteristics used in our analysis. It indicates that 19% of respondents obtained University education. Income is captured at the household-level and converted to lie on the logarithmic scale. 57% of respondents are either married or living with their partner, and 56% of respondents are female. The average age of respondents in our sample is 47.47 years; the youngest respondent is 21 years old, and the oldest is aged 65. Approximately 29% of respondents live in a household by themselves.

We further include a set of indicator variables that capture the respondents' identifications with UK political parties. One-fifth (20%) of respondents report that they do not identify with any political party, under one-third (28%) report identification with the conservative party, while approximately one-third (35%) report identification with the labour party, and under one-fifth (17%) with another party.

2.4.3 Descriptive evidence

In Figure 2.2, we begin by investigating the relationship between welfarism and mental health problems visually. The horizontal axis shows the full range of possible welfarism scores, while the vertical axis shows the average of the mental health problem variable in percentage points. Therefore, each dot in Figure 2.2 represents the prevalence of mental health problems at the given welfarism score estimated in our data sample. The dashed line indicates the fitted values obtained from the LOESS smoother, together with the 95% confidence intervals.

Three observations emerge from Figure 2.2. First, there appears to be a positive association between welfarism scores and mental health problems. Second, while the slope appears flat for the lower range of welfarism scores, it turns decidedly positive for higher scores. Third, the confidence intervals slightly widen for very high or very low welfarism scores, reflecting the lower number of respondents with welfarism scores in these regions of the distribution. Overall, this descriptive evidence suggests that welfarism scores and mental health problems are positively associated; specifically, when comparing the effects across the low and high region of welfarism scores.

In Figure 2.3, we provide the first evidence in support of the proposed mechanism that sanctioning of individuals can provide an explanation for the link underlying the relationship between welfarism and mental health problems. Intuitively, a political environment less favourable to welfare recipients and positive welfare views should increase the degree of sanctioning as the associated sanctioning costs reduce, thereby resulting in higher prevalence of mental health problems. Accordingly, Figure 2.3 plots the relationship between welfarism and mental health problems, where both dots and fitted values our split accordingly to the type of political party providing the government.

Two observations summarise Figure 2.3: first, we observe that the prevalence of mental health problems is significantly higher when the conservative party provides the government than when the labour party does so. Second, while the relationship between welfarism and mental health problems appears stable in low regions of the welfarism distribution, the trends diverge for higher scores. In the case of a conservative government, the positive slope is increasing with higher welfarism scores; in contrast, the slope is flattening in the same region of welfarism scores under labour governments. This evidence suggests the possibility that social sanctioning increases under conservative rule.

2.5 Empirical analysis

2.5.1 Welfare attitudes and mental health problems

We begin by establishing the baseline association between reports of high welfarism and mental health problems in a multivariate regression framework that controls for important confounding influences at the individual- and household-level. Moreover, as the relationship between welfarism and mental health problems is likely subject to endogeneity concerns, such as reverse causality or omitted variable bias, we attempt to identify the direction of potential bias introduced in the baseline regression estimates by employing a Limited Information Maximum Likelihood (LIML) framework.

To do so, we specify the following set of recursive probit equations:

$$MH_i = \mathbb{1}(\alpha + \gamma HW_i + X_i'\theta + \varepsilon_i > 0) \tag{2.3}$$

$$HW_i = \mathbb{1}(\beta + Z_i'\delta + X_i'\phi + \eta_i > 0), \tag{2.4}$$

with the following error term structure:

$$(\varepsilon_i, \eta_i)' \sim \mathcal{N}(0, \Sigma), \text{ where } \Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \tag{2.5}$$

where i is the respondent-level identifier and $\mathbb{1}(\cdot)$ is the indicator function, taking the value one if its argument is true and zero otherwise. The dependent variable in Equation (2.3) is the indicator variable capturing mental health problems (MH), and the explanatory variable of interest is the high welfarism indicator (HW, defined as scores in the top 30% of the welfarism score distribution). The dependent variable in Equation (2.4) is the high welfarism indicator, and the explanatory variables of interest are a set of instrumental variables, Z, to be specified. The error terms in Equations (2.3) and (2.4) are related to each other in a Seemingly Unrelated Regression (SUR) specification, parametrised as a multivariate normal distribution with covariance matrix Σ, where the variances are constrained to unity as a result of the probit normalization assumption. The parameter ρ captures the correlation between the error terms. The vector X captures important individual- and household-level characteristics, such as holding a University degree, the household income (in logs), being married or cohabitating, being female, the respondents age in levels and squares, as well as living by oneself in a single household, together with year dummies capturing time fixed effects.

This framework amounts to an instrumental variable approach for binary limited dependent variables in the first and second stage, Equation (2.3) and (2.4), respectively. It is attractive for several reasons. First, if the requirements for valid instrumental variables are met, the parameter of interest, γ, quantifies the relationship between HW and MH free of bias. Second, the joint estimation of the parameter ρ returns an estimate capturing the presence of endogeneity in the relationship between HW and MH. Specifically, conditional on HW, X and Z, an omitted driver of both HW and MH will introduce a non-zero correlation in the error terms of equation one and two, $\rho \neq 0$. We will obtain an estimate of ρ, thereby allowing for statistical testing of endogeneity. Third, Equation (2.3) is identified even if the elements of the vector δ are restricted to zero, amounting to the removal of the instrumental variables from the first stage equation. In this special case, the source of identification can originate from the specific form of the model's non-linearity (Mourifié & Méango, 2014; Roodman, 2011; Wilde, 2000). Importantly, identification of the parameter γ can thus be obtained when no appropriate instrumental variables suggest themselves, or the ones that are available do not unambiguously satisfy the instrumental variable requirements.

As instruments for high welfarism, we consider three binary variables, which take the value one if political identification with the conservative party, the labour party or another party, respectively, is reported, and zero otherwise. These three variables are likely to satisfy the requirements for valid instruments. In the UK, political party identification is strongly associated with varying degrees of welfare attitudes. For instance, members of the conservative party are known to hold less favourable views on welfare and benefit receipt, while supporters of the labour party hold more favourable views. Other parties exhibit heterogeneous stances on welfare. Therefore, we expect a strong first stage association – in other words, meeting of the instrumental variable relevance conditions – between the political party identifiers (Z) and the indicator variable capturing high welfarism (HW). We test for this strong first-stage empirically. Moreover, it appears unlikely that after controlling for the age, income and university education, political party identification infers a relationship with reports of mental health problems (MH) separate from its effect through (HW). Therefore, these political party identifiers are likely to satisfy the exclusion restriction for valid instrumental variables. Together the meeting of the relevance and exclusion restrictions constitute an instrumental variable approach to estimating the relationship between HW and MH in which endogeneity concerns are attenuated.

We estimate the association between HW and MH in three different ways. First, we obtain simple probit estimates by individually estimating Equation (1); second, we estimate Equations (1) to (3) without the inclusion of the instrumental variables; and third, we estimate Equations (1) to (3) and include the instrumental variables in the first stage regression. Table 2.2 reports the results, where Columns (1) to (3) show those for these different approaches, reflecting a successive increase in model complexity. The reported estimates are average marginal effects (AMEs).

In Column (1), the estimate of the relationship between HW and MH suggests that a unit increase in HW is associated with a 0.130 percentage point increase of $P(MH = 1)$. The baseline predicted probability of 0.335 indicates that this amounts to an approximately 30% increase versus the baseline. In Column (2), in which the first-stage regression of controls only is jointly estimated, the estimated association between W and MH increases to 0.425 percentage points. The estimates of ρ and their corresponding 95% confidence intervals indicates that the null of no endogeneity between HW and MH is rejected at the 10% level. This indicates that the estimates obtained from the simple probit regression in Column (1) are subject to bias. However, as the magnitude of the association is smaller in Column (1), we conclude that endogeneity introduces a downward bias. In

Column (3), in which HW is instrumented with the political party identifiers in Z, we observe that the association between HW and MH equals 0.364, approximately in line with the results from Column (2) but greater than those of Column (1). The first-stage F-statistic of 31.994 indicates a strong first-stage regression. Therefore, weak identification appears to not be a concern. The unreported first-stage coefficients of the instrumental variables are all strongly significant. Overall, the estimates in Table 2.2 reveal a strong association between HW and MH. Further, as endogeneity introduces a downward bias in the baseline estimates of this relationship, the estimations reported in the following sections constitute lower bounds of any the estimated associations.

2.5.2 Variation in sanctioning costs

We now turn to testing the implication of the theoretical framework described in Section 2.3. When the UK government is drawn from the conservative party, the typical result is a political environment in which welfare spending and benefit recipients are viewed less favourably - in its extreme, resulting in political scapegoating. In such a context, the cost of social sanctioning reduces. Individuals that are employed or receive no welfare benefits conform to the societal ideal that working is a normative good, thereby adhering to the social identities associated with the working population. Expressing positive welfare attitudes, constituting a deviation from the social identity that purports the normative value of work, will be sanctioned by other individuals subscribing to the same identity. In harsher political environments towards welfare that reduce the cost of sanctioning, the theoretical framework implies that it will occur at higher frequencies, manifesting in higher rates of mental health problems.

We exploit variation in the ruling governing party in the UK over time to capture changing political environments for welfare. By testing whether the link between high welfarism and mental health problems intensifies under conservative rule, we provide evidence in favour of the notion that expressing positive welfare attitudes can be sanctioned by other individuals as a result of them feeling threatened in their identities.

We construct an indicator variable (C_t) that takes the value 1 if the UK government at time t is drawn from the conservative party, and 0 otherwise. We interact this Tory government variable with the high welfarism indicator (HW) and estimate the relationship of these interactions with the mental health problem variable (MH). To capture the differing intensities of sanctioning for high and low conformity individuals, we estimate the relationships for G different subgroup of respondents, described in more detail below, where g is the group-level identifier. The following probit model is estimated separately for each group:

$$\begin{aligned} MH_i = \mathbb{1}(\alpha_g + \gamma_1(C_t) + \gamma_2(HW_i) + \gamma_3(HW_i \times C_t) + X_i'\theta + \varepsilon_i > 0), \\ \forall i \in \{i \mid s_g(i) = 1\}, \end{aligned} \tag{2.6}$$

where $s_g(i)$ takes the value 1 if the respondent i belongs to group g, and zero otherwise. The indicator function $\mathbb{1}(\cdot)$ takes the value one if its argument is true and zero otherwise. The dependent variable in Equation (2.6) is the indicator variable capturing mental health problems (MH), and the explanatory variable of interest is the high welfarism indicator (HW). The error term ε_i follows a standard normal distribution. The vector X captures important individual- and household-level characteristics, such as holding a University degree, the household income (in logs), being married or cohabitating, being female, the respondents' ages in levels and squares and living by oneself in a single household. The year fixed effects are collinear with the interactions of interest and therefore excluded from estimation.

We consider six respondent subgroups ($G = 6$) that are not mutually exclusive. The first three groups comprise respondents of high conformity with the normative value of work. The first of these groups includes respondents that are in employment, while the second group includes those that are not in receipt of welfare benefits, whereas respondents in the third group are both in employment and receive no benefits. The second three groups include respondents of low conformity with societal work norms. Respondents in these groups include, respectively, those that are not in employment, are in receipt of welfare benefits and are both not in employment and receive benefits. Accordingly, Table 2.3 reports the estimation results, where Columns (1) to (6) report the estimates corresponding to the respective subgroups.

Several important observations emerged from Table 2.3. It can be seen that for high welfarism individuals, conservative political party rule is associated with higher coefficient estimates. Tests for the differences between the coefficients under labour and conservative governments indicate a strong rejection of the null hypothesis that of equality of effects. This suggests that the relationship between welfarism and mental health problems intensifies under conservative rule and, importantly, the differences across columns indicate that this occurs disproportionately so for high conformity individuals. Overall, these estimates provide evidence in favour of the possibility that high welfarism is sanctioned more strongly in environments less friendly to the welfare state and welfare recipients.This confirms the implication from the theoretical framework from Section 2.3 that reductions in sanctioning costs increase mental health problems, and contributes social norm enforcement as a factor influencing mental health to the pertinent literature (Bridges & Disney, 2010; Currie et al., 2015; Gathergood, 2012; Grip et al., 2011).

2.6 Additional analysis

2.6.1 The role of gender

Research has documented significant gender differences in the propensity to experience mental health issues (Afifi, 2007). Given this disproportionate susceptibility to developing mental health issues, we test the notion that the relationship between welfarism and mental health problems manifests unequally between male and females.

We begin by estimating the following probit model in which we interact the indicator variable taking the value one for female respondents, and zero for male ones (F) with the respondents' welfarism scores (W):

$$MH_i = \mathbb{1}(\alpha + \gamma_1(F_i \times W_i) + \gamma_2((1 - F_i) \times W_i) + X_i'\theta + \varepsilon_i > 0) \tag{2.7}$$

where i is the respondent-level identifier. The indicator function $\mathbb{1}(\cdot)$ takes the value one if its argument is true and zero otherwise. Again, the dependent variable in Equation (2.7) is the indicator variable capturing mental health problems (MH). The error term ε_i follows a standard normal distribution. The vector X includes the control variables capturing an indicator for an obtained University degree, the household income (in logs), being married or cohabitating, the respondents' ages in levels and squares and living by oneself in a single household, as well as time dummies to capture time fixed effects.

Table 2.4 reports the estimation results. The coefficient estimates suggest that mental health problems increase as welfarism increases, for both males and females; however, though both estimates are highly significant, the magnitude of increase is greater for females. This greater gradient in the relationship between high welfarism and mental health problems for females is in line with research suggesting that women are more likely to suffer from mental health issues such as anxiety (Angst & Dobler-Mikola, 1985; Bruce et al., 2005; McLean, Asnaani, Litz & Hofmann, 2011) and depression (Nolen-Hoeksema & Girgus, 1994; Piccinelli & Wilkinson, 2000). As such, these results indicate one possible explanation regarding the differential likelihood in mental health problems for females versus males: females may be at higher risk of social sanctioning when they express positive welfare attitudes.

In order to provide economic magnitudes for the estimation results, we obtain the predicted probabilities across the full range of welfarism scores for both males and females. Figure 2.4 presents the results, where the welfarism scores are depicted on the horizontal axis, while the predicted probabilities are reported on the vertical axis. The black triangles show the predicted probabilities for females, while the white triangles show those for males. The error bars indicate 95% confidence intervals. It can be seen from Figure 2.4 that, while the predicted probabilities for males and females do not differ for very low welfarism scores, continuously increasing differences manifest as welfarism scores increase. Figure 2.4 thus implies that 1) the predicted probabilities for females are higher than for males across the majority of the welfarism score distribution and 2) the association between welfarism scores and mental health problems is stronger for females than it is for males. Overall, these results suggest support for the notion that the observed relationship between welfarism and mental health problems manifest unequally for females and males.

In order to shed further light on the notion that females may be more susceptible to social sanctioning, we further investigate the differential associations across gender under the different types of political rule. To do so, we make use of the indicator variable C_t that takes the value 1 if the UK government at time t is drawn from the conservative party, and 0 otherwise, and estimate the following probit model:

$$\begin{aligned} MH_i = \mathbb{1}(\alpha &+ \gamma_1((1-F_i) \times (1-C_t) \times W_i) \\ &+ \gamma_2((1-F_i) \times C_t \times W_i) \\ &+ \gamma_3(F_i \times (1-C_t) \times W_i) \\ &+ \gamma_4(F_i \times C_t \times W_i) + X_i'\theta + \varepsilon_i > 0) \end{aligned} \tag{2.8}$$

where i is again the individual identifier, F is the indicator capturing female respondents and W denotes welfarism scores. The indicator function $\mathbb{1}(\cdot)$ takes the value one if its argument is true and zero otherwise. The error term ε_i follows a standard normal distribution. The vector X includes the control variables capturing an indicator for an obtained University degree, the household income (in logs), the respondents' ages in levels and squares, being married or cohabitating and living by oneself in a single household. The time dummies are collinear with the interactions of interest and are therefore excluded from the estimation.

Table 2.5 reports the estimation results. The triple interaction of the female indicator, Tory government indicator and welfarism scores is highly significant in all cases and shows that the greatest susceptibility to mental health problems as welfarism scores increase is faced under Tory government rule, for both males and females. However, the gradient for males under Tory government is approximately equal those for females when non-Tory parties rule. To aid the economic interpretation of these estimation results, we again obtain the predicted probabilities for the full range of welfarism scores for both males, females and by whether the ruling party is drawn from the conservative or labour party, respectively. Figure 2.5 presents the results. Panel (a) shows the results for labour governments, while Panel (b) shows those for conservative governments. The welfarism scores are depicted on the horizontal axis, while the predicted probabilities are reported on the vertical axis. The black triangles show the predicted probabilities for females, while the white triangles show those for males. The error bars indicate 95% confidence intervals.

The results mirror those for Figure 2.4: while the predicted probabilities for males and females do not differ for very low welfarism scores under both ruling parties, continuously increasing differences manifest as welfarism scores increase. However, visual inspection of the differences between Panel (a) and (b) reveal that the gap between the predicted probabilities for females and males is notably greater under conservative than under labour party rule. This suggests that the relationship between welfarism and mental health problems disproportionally intensifies for females when the conservative party governs.

2.6.2 Voting for change

Having provided evidence on a significant link between welfarism and mental health problems, and that this relationship intensifies under conservative political governments, we turn our investigation to one possible response open to individuals: using their votes in elections to induce political change. If individual circumstances worsen under the policies and rhetoric of the incumbent political party, it appears likely that individuals recognise the importance of voting in order to bring on changes in the political system. Therefore, the possibility exists that individuals with high welfarism who report a higher prevalence of mental health problems will exhibit more favourable attitudes to voting participation.

The 2000 and 2013 BSA surveys provide a measure of the respondents' attitudes toward voting. Respondents indicate whether they believe that either 1) it is not worth voting, 2) that voting only matters if they care about who wins or 3) whether voting is a duty. We capture responses in the ordinal variable V where the lowest outcome refers to voting not being worth it, while the highest to the belief that voting is a duty. In order to understand how the combinations of high welfarism and mental health problems relate to voting attitudes, we estimate the following ordered probit model:

$$V_i = g(\alpha + \gamma_1(MH_i) + \gamma_2(HW_i) + \gamma_3(HW_i \times MH_i) + X_i'\theta + \varepsilon_i) \tag{2.9}$$

$$g(v^*) = \begin{cases} \text{“Not worth voting”} & \text{if} \quad -\infty < v^* \leq c_1 \\ \text{“Vote if care who wins”} & \text{if} \quad c_1 < v^* \leq c_2 \\ \text{“Duty to vote”} & \text{if} \quad c_2 < v^* < \infty \end{cases} \tag{2.10}$$

where i denotes individuals and $g(\cdot)$ is the ordered probit link function (Roodman, 2011). The error term in Equation (2.9), ε_i, follows a standard normal distribution. The vector X includes the control variables capturing an indicator for an obtained University degree, the household income (in logs), being married or cohabitating, being female, the respondents ages' in levels and squares and living by oneself in a single household, as well as time dummies to capture time fixed effects.

After estimation of the model in Equation (2.9), we obtain the average marginal effects (AMEs) for high welfarism, separated by reports of mental health problems. Table 2.6 presents the results. The columns show the AMEs for relating to the three possible outcomes for the voting attitudes variable; Columns (1) to (3) show the AMEs for positive reports of mental health problems and Columns (4) to (6) for negative ones. Two important trends become apparent in Table 2.6. First, for no reports of mental health problems, all coefficient estimates for high welfarism are not significant at the 10% level. In contrast, in the case of reports of mental health problems, all high welfarism estimates are strongly significant. This suggests that the combination of high welfarism and mental health problems is associated with changing attitudes towards voting. Second, turning to the specific estimates of these significant associations, Column (1) shows that the probability of reporting that voting is not worth it reduces by 8.88 percentage points for high welfarism, while Column (2) suggests an associated reduction of the probability that voting only matters if they care about who wins by 4.49 percentage points. Accordingly, Column (3) shows that high welfarism is associated with an increased probability of reporting a duty to vote by 13.7 percentage points.

Thus, the results in Table 2.6 provide evidence of more favourable attitudes towards voting for high welfarism individuals reporting mental health problems, thereby suggesting the existence of respondent tendencies to vote for political change in response to adverse circumstances.

2.7 Limitations and future research

Our findings identify a need for further research. Though we provide evidence on the possible existence of the sanctioning of high welfarism individuals for deviating from societal norms that posit the normative value of work, sanctioning is fundamentally unobservable in this analysis and thus remains latent. Future research that explicitly measures whether respondents are subject to sanctioning through survey questions can address this concern, thereby building on the empirical results that this work contributes to the literature, in which the sanctioning mechanism is identified through variation in incumbent governing parties and the associated differences in the ideological majority these different periods of political rule imply.

In addition, a limitation of this study is the lack of observing respondents across different survey iterations. Though our empirical analysis controls for important confounding factors – and takes steps to mitigate endogeneity concerns in an instrumental variable approach and by controlling for endogenous relationships parametrically – caution should be taken when assigning a causal interpretation to these findings. The results do not provide conclusive evidence that fostering more positive individual welfare attitudes will attract sanctioning and subsequent deteriorations in mental health. Future work could result from an effort to include measures of welfare attitudes in major longitudinal surveys, such as the Understanding Society and British Household Panel Survey. Time-ordered data would enable the measuring of changes in welfare attitudes and subsequent changes in mental health outcomes.

With respect to gender, further work could investigate the relationship between welfarism and mental health problems for non-binary gender identities and identify the exact channels that result in the differential effects across gender uncovered in this study. Finally, future work marrying the measurement of welfare attitudes to an existing longitudinal survey could provide important granular data on individual mental health; information that is already routinely elicited. For instance, together with welfare attitudes, respondents' perceptions of their mental health at the time of the survey could be queried, capturing specific mental health dimensions together with their experienced intensity.

2.8 Conclusion

In this study, we study the link between positive attitudes to the welfare state and benefit recipients (welfarism) and mental health problems. A set of empirical characteristics transpire from the data. First, we establish a strong association between high welfarism (defined as scores in the top 30% of the welfarism score distribution) and reports of mental health problems. The results from the empirical analysis indicate that the likely presence of endogeneity in this relationship introduces a downward bias in the baseline results; as such, the results of this study can be interpreted as measurements of lower bounds on the proposed relationships. Second, we show that the relationship between high welfarism and mental health problems intensifies under conservative governments, disproportionally so for individuals exhibiting high conformity with work-related social norms, displayed through being in employment or receiving no benefits. Third, we find that the observed relationships are more pronounced for females than for males and that the gap in negative mental health outcomes increases between genders for survey responses elicited under conservative governments versus labour governments. Fourth, we document significant associations for high welfarism individuals with reports of mental health problems and the notion that voting is a duty, whereas these associations are not statistically significant given no mental health problems.

The results have direct implications for policy makers. The evidence that the relationship between welfarism and mental health problems intensifies when the government is drawn from the Conversative party, who is known to be harsher on welfare, indicates that the political majority can influence how acceptable it is to reprimand individuals for perceived deviations from everyday social norms. The easier social sanctioning, the greater the ill effects on mental health are expected to become. Refraining from political scapegoating, as in the case of welfare recipients or other groups of society, and thereby increasing the cost of social sanctioning suggests itself as a way that policy makers can have direct influence on the mental health of its citizenry.

Figure 2.1

Game tree interaction between individuals i and j

This game tree shows the interactions between workers i and j in an adaption of the model by Akerlof and Kranton (2000). Individual j derives utility $V(HW)$ from expressing high welfarism, HW, and utility 0 if she does not express these positive welfare attitudes. In contrast, individual i does not hold positive welfare attitudes and gains no utility from expressing positive views; V is her baseline utility. I_s denotes the costs to the worker identity as a result of expressing positive welfare attitudes, while I_o denotes the resulting identity-externalities born by others. The effort involved in sanctioning is denoted e, while MH is the mental health utility loss as a result of being sanctioned.

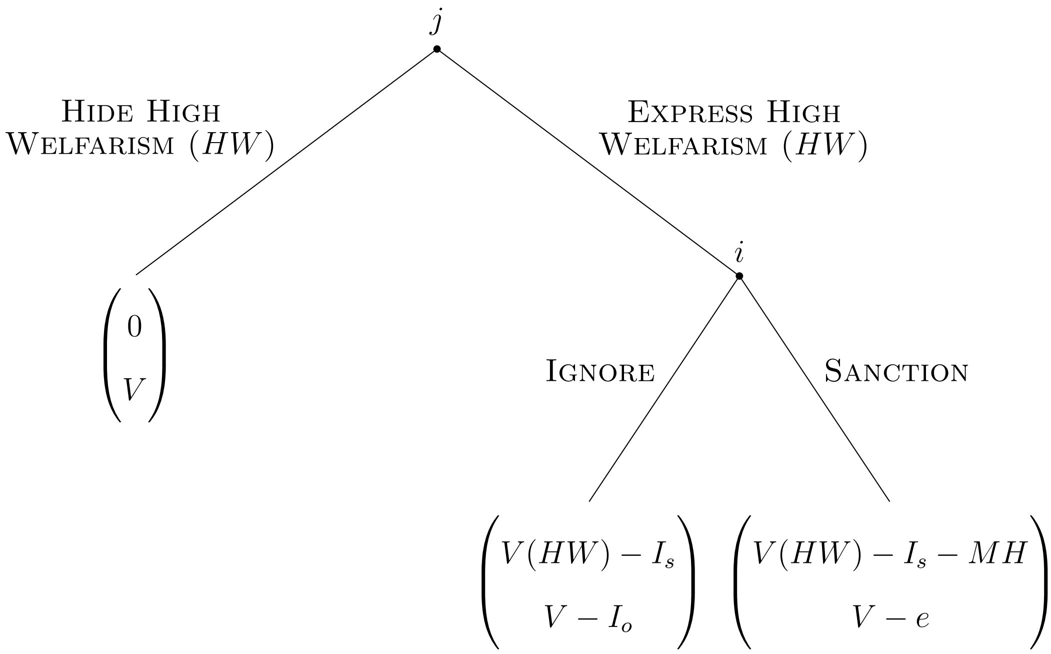

Figure 2.2

Welfarism and mental health problems

This figure shows the relationship between welfarism and mental health problems. The horizontal axis reports welfarism on a scale from 8 to 40, where higher values indicate more positive welfare attitudes. Each dot represents the prevalence of mental health problems in percentage points at the respective level of welfarism. The dashed line shows the fitted values obtained from the LOESS smoother, together with the 95% confidence intervals.

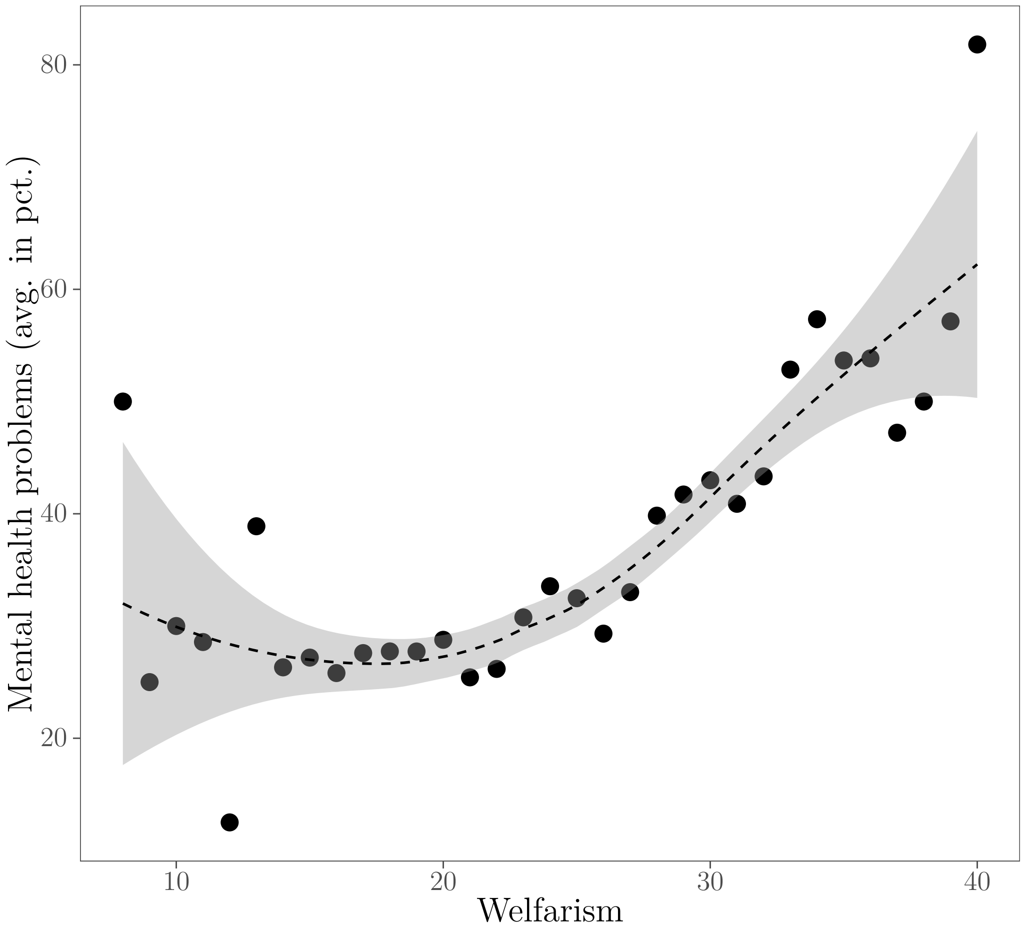

Figure 2.3

Welfarism and mental health problems by ruling political party

This figure shows the relationship between welfarism and mental health problems, split by the ruling political party in the given survey year. The horizontal axis reports welfarism on a scale from 8 to 40, where higher values indicate more positive welfare attitudes. Each square represents the prevalence of mental health problems in percentage points at the respective level of welfarism. The dashed line shows the fitted values obtained from the LOESS smoother, together with the 95% confidence intervals.

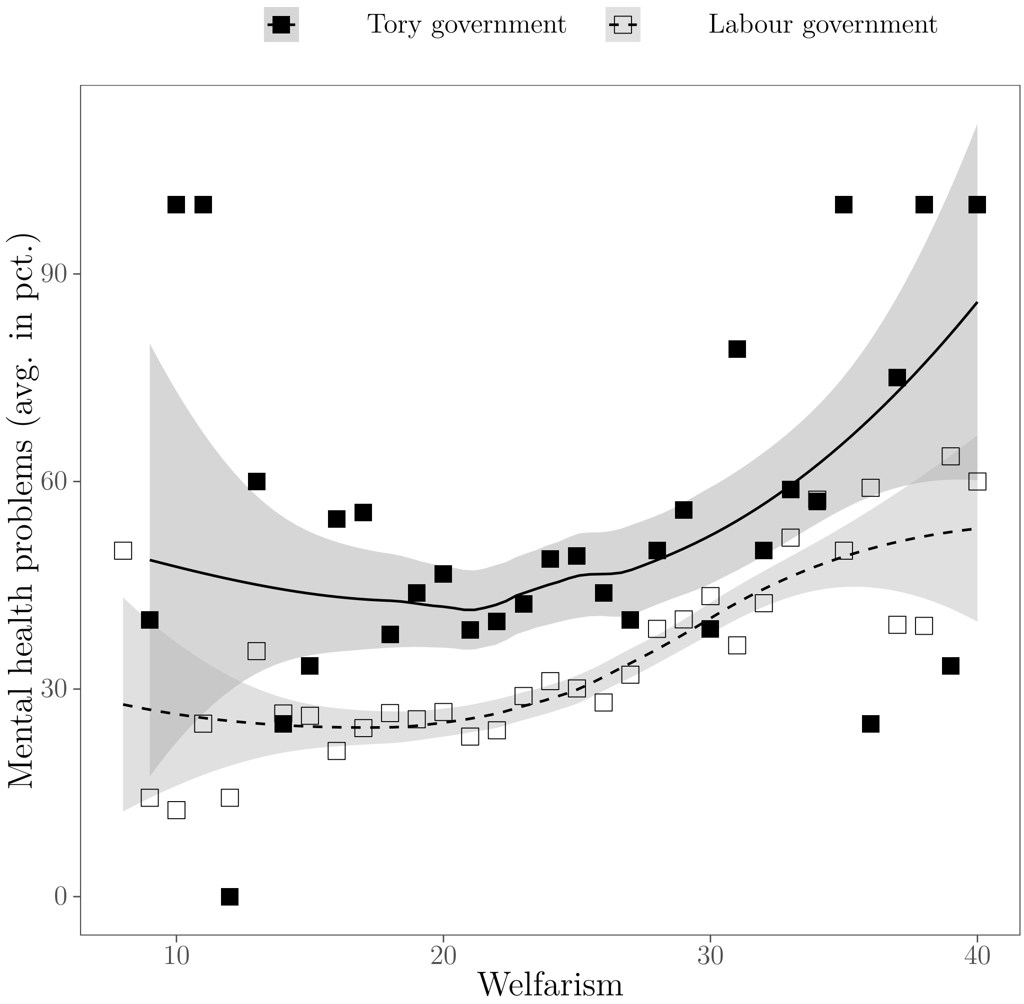

Figure 2.4

Welfarism, mental health problems and gender

This figure visualises average predicted probabilities. The horizontal axis shows welfarism on a scale from 8 to 40, where higher values indicate more positive welfare attitudes. Each dot represents the average predicted probability obtained at the respective level of welfarism. The solid line reports the predicted probabilities for females, while the dashed line reports those for males. The error bars indicate 95% confidence intervals.

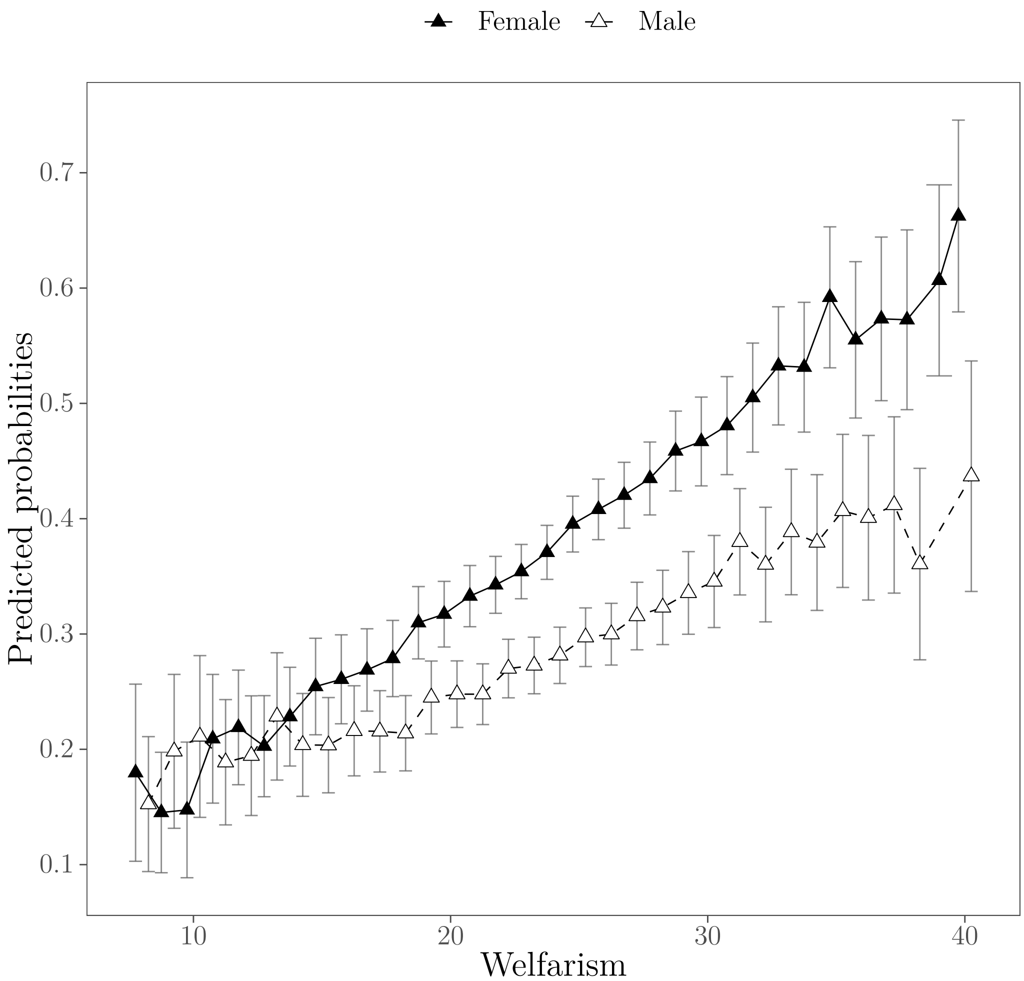

Figure 2.5

Gender and the political environment

This figure visualises average predicted probabilities. The horizontal axis shows welfarism on a scale from 8 to 40, where higher values indicate more positive welfare attitudes. Each dot represents the average predicted probability obtained at the respective level of welfarism. The solid line reports the predicted probabilities for females, while the dashed line reports those for males. Panel (a) shows the relationship under labour governments; Panel (b) under conservative governments. The error bars indicate 95% confidence intervals.

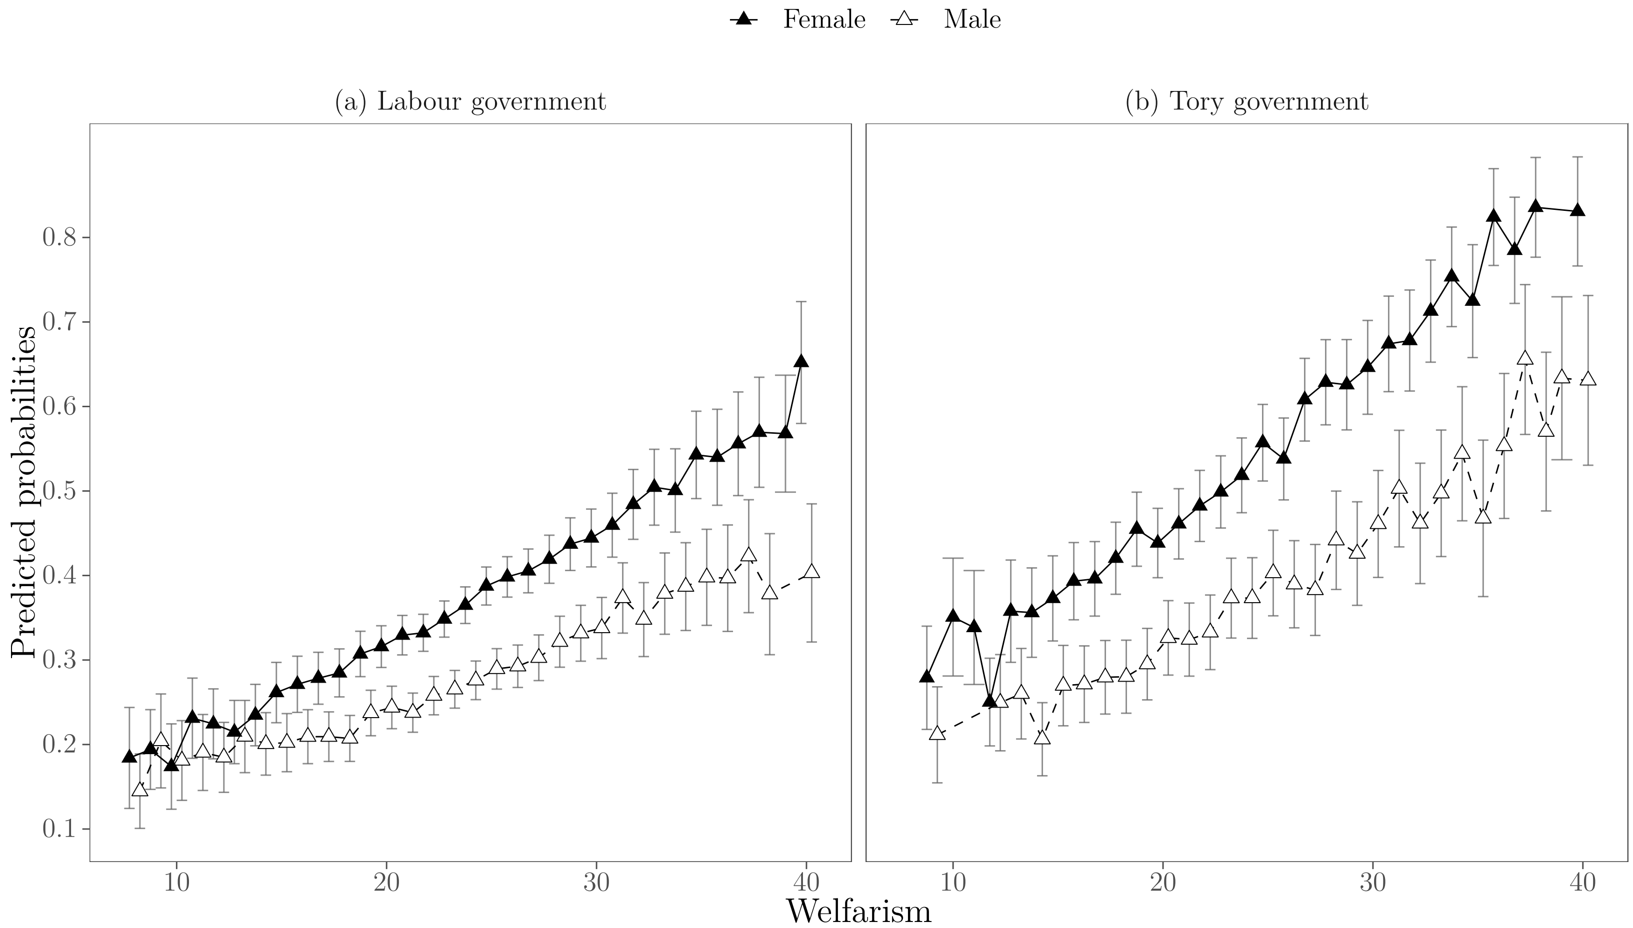

Table 2.1

Variable descriptions and summary statistics

This table reports the variable descriptions and summary statistics for the data sample. The data is sourced from the British Social Attitudes surveys in the years 2000, 2003, 2006, 2007 and 2013.

Variable	Description	Mean	Min.	Max.	SD
Welfarism	Captures welfare attitudes, higher answers indicate more positive views	23.84	8.00	40.00	5.03
Mental health problem	=1 if ever sought medical advice on mental health problem, =0 otherwise	0.33	0.00	1.00	0.47
University education	=1 if university education obtained, =0 otherwise	0.19	0.00	1.00	0.39
Household income	Household income converted to log scale	9.91	8.29	10.93	0.77
Married or cohabitating	=1 if married or cohabitating, =0 otherwise	0.57	0.00	1.00	0.50
Female	=1 if female, =0 if male	0.56	0.00	1.00	0.50
Age	Age in years	47.47	21.00	65.00	14.23
Single household	=1 if one person household, =0 otherwise	0.29	0.00	1.00	0.45
No party identification	=1 if no party identification, =0 otherwise	0.20	0.00	1.00	0.40
Conservative party identification	=1 if Conservative party identification, =0 otherwise	0.28	0.00	1.00	0.45
Labour party identification	=1 if Labour party identification, =0 otherwise	0.35	0.00	1.00	0.48
Other party identification	=1 if other party identification, =0 otherwise	0.17	0.00	1.00	0.38
Observations	3031				

Table 2.2

The relation between welfarism and mental health problems

This table reports average marginal effects (AMEs). The dependent variable is an indicator capturing the presence of mental health problems. High welfarism takes the value one for scores in the top 30% of the distribution, and zero otherwise. Column (1) shows the AMEs for the baseline probit model; Column (2) shows those for the two-stage probit model achieving identification through the model's non-linearity; and Column (3) shows those for the two stage probit model using an instrumental variable approach. In Column (2), the first stage involves a probit regression of high welfarism on the model controls, while in Column (3), three instrumental variables are included that capture political affiliation with the conservative, labour or other parties. Standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Mental health problem		
	Probit	Probit non-linear identification	Probit instrumental variable identification
	(1)	(2)	(3)
High welfarism	0.130***	0.425***	0.364***
	(0.02)	(0.06)	(0.06)
University degree	0.045**	-0.029	-0.010
	(0.02)	(0.03)	(0.03)
Household income	0.012	0.035***	0.030**
	(0.01)	(0.01)	(0.01)
Married or cohabitating	0.010	0.039*	0.033
	(0.03)	(0.02)	(0.02)
Female	0.096***	0.088***	0.093***
	(0.02)	(0.02)	(0.02)
Age	-0.003***	-0.002***	-0.003***
	(0.00)	(0.00)	(0.00)
Single household	0.061**	0.072***	0.072***
	(0.03)	(0.02)	(0.02)
Year f.e.	Yes	Yes	Yes
Baseline predicted probability	0.335	0.364	0.353
$\hat{\rho}$		-0.646	-0.487
		[-0.885,-0.137]	[-0.731,-0.135]
No. of instruments	0	0	3
First-stage F-statistic			31.994
Observations	3031	3031	3031

Table 2.3

Variation in sanctioning costs

This table reports regression results where the dependent variable is mental health problems and the independent variable of interest is the interaction between high welfarism and an indicator variable capturing whether the Conservative party is in government. The reported values are probit coefficient estimates. Columns (1) to (3) show the results for individuals with high conformity with societal norms that place value on strong labour market involvements, whereas Columns (4) to (6) show the results for low conformity individuals. Standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Mental health problem					
	High conformity individuals			Low conformity individuals		
	Employed	No benefits	Employed and no benefits	Not employed	Has benefits	Not employed and has benefits
	(1)	(2)	(3)	(4)	(5)	(6)
High welfarism=0 × Tory government=1	0.454*** (0.08)	0.376*** (0.09)	0.392*** (0.10)	0.493*** (0.08)	0.415*** (0.09)	0.442*** (0.10)
High welfarism=1 × Tory government=0 (λ_2)	0.398*** (0.07)	0.424*** (0.08)	0.462*** (0.09)	0.379*** (0.06)	0.348*** (0.07)	0.353*** (0.08)
High welfarism=1 × Tory government=1 (λ_3)	0.692*** (0.13)	0.873*** (0.15)	0.892*** (0.17)	0.599*** (0.10)	0.651*** (0.12)	0.601*** (0.13)
University education	0.124* (0.06)	0.077 (0.08)	0.098 (0.08)	0.157** (0.07)	0.177* (0.09)	0.263** (0.11)
Married or cohabitating	-0.056 (0.08)	0.058 (0.10)	-0.042 (0.11)	-0.053 (0.08)	0.070 (0.09)	0.008 (0.10)
Female	0.285*** (0.06)	0.306*** (0.07)	0.295*** (0.07)	0.279*** (0.05)	0.270*** (0.06)	0.269*** (0.07)
Age	0.026 (0.02)	0.037** (0.02)	0.029 (0.02)	0.070*** (0.02)	0.101*** (0.02)	0.119*** (0.02)
Age^2	-0.000* (0.00)	-0.000** (0.00)	-0.000 (0.00)	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)
Single household	0.108 (0.09)	0.255** (0.10)	0.167 (0.11)	0.066 (0.09)	0.199** (0.10)	0.119 (0.11)
Constant	-1.154*** (0.34)	-1.425*** (0.35)	-1.221*** (0.43)	-2.052*** (0.36)	-2.595*** (0.36)	-2.981*** (0.49)
Prob > χ^2 ($H_0 : \lambda_2 = \lambda_3$)	0.000	0.000	0.000	0.000	0.000	0.000
Observations	2242	1607	1270	2662	2027	1690

Table 2.4

Gender, welfarism and mental health

This table reports the coefficient estimates from a probit regression of mental health problems on the interaction of the female indicator and welfarism scores variables. Standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Mental health problem
Female=0 × Welfarism score	0.025***
	(0.01)
Female=1 × Welfarism score	0.038***
	(0.01)
University education	0.120*
	(0.06)
Household income	0.044
	(0.04)
Married or cohabitating	0.034
	(0.07)
Female	-0.048
	(0.24)
Age	0.049***
	(0.01)
Age^2	-0.001***
	(0.00)
Single household	0.182**
	(0.08)
Constant	-2.585***
	(0.47)
Year f.e.	Yes
Observations	3031

Table 2.5

Gender and variations in sanctioning costs

This table reports the coefficient estimates from a probit regression of mental health problems on the triple interaction of the female indicator, Tory government indicator and welfarism score variables. Standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Mental health problem
Female=0 × Tory government=0 × Welfarism score	0.026***
	(0.01)
Female=0 × Tory government=1 × Welfarism score	0.037***
	(0.01)
Female=1 × Tory government=0 × Welfarism score	0.033***
	(0.01)
Female=1 × Tory government=1 × Welfarism score	0.051***
	(0.01)
University education	0.121**
	(0.05)
Married or cohabitating	-0.003
	(0.06)
Female	0.085
	(0.20)
Age	0.055***
	(0.01)
Age^2	-0.001***
	(0.00)
Single household	0.141**
	(0.07)
Constant	-2.199***
	(0.26)
Observations	4269

Table 2.6

Voting for change

This table reports the average marginal effects (AMEs) from an ordered probit regression of attitudes to voting on the triple interaction of mental health problems, high welfarism and gender, as shown in Equation (2.9). The respective outcome to which the AMEs refer is indicated in the column heading. Columns (1) to (3) show the results for the subgroup with reports of mental health problems; Columns (4) to (6) show the results for that with no reports of mental health problems. 846 respondent observations from the 2000 and 2013 British Social Attitudes surveys are used in the estimation. Standard errors are reported in parentheses. The stars ***, ** and * denote levels of significance at 1, 5 and 10 percent, respectively.

	Mental health problems					
	Yes			No		
	Not worth voting	Vote if care who wins	Duty to vote	Not worth voting	Vote if care who wins	Duty to vote
	(1)	(2)	(3)	(4)	(5)	(6)
High welfarism	-0.088***	-0.049**	0.137***	-0.023	-0.014	0.037
	(0.03)	(0.02)	(0.05)	(0.03)	(0.02)	(0.05)
University education	-0.130***	-0.062***	0.192***	-0.124***	-0.070***	0.194***
	(0.03)	(0.01)	(0.04)	(0.03)	(0.01)	(0.04)
Married or cohabitating	-0.054*	-0.026*	0.080*	-0.052*	-0.029*	0.081*
	(0.03)	(0.01)	(0.04)	(0.03)	(0.02)	(0.04)
Female	-0.034	-0.016	0.051	-0.033	-0.018	0.051
	(0.02)	(0.01)	(0.03)	(0.02)	(0.01)	(0.03)
Age	-0.003***	-0.002***	0.005***	-0.003***	-0.002***	0.006***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Single household	-0.051	-0.024	0.075	-0.048	-0.027	0.075
	(0.03)	(0.02)	(0.05)	(0.03)	(0.02)	(0.05)
Observations	846	846	846	846	846	846

3

Endogenous Peer Choice and Investment Participation

Highlights

∗ A peer effects model for binary outcome variables is proposed that explicitly accounts for endogenous relationship formations in household networks. Given the high dimensionality of the parameter space, estimation takes place in a Bayesian econometric framework using Hamiltonian Monte Carlo sampling methods.

∗ The model is brought to the data for two outcome variables: a general investment income indicator variable and an interest + dividend income indicator variable. It is estimated both with exogenous and endogenous relationship formations. The model comparison reveals that endogenous peer choice accounts for approximately 25% of the peer effect in investment participation.

∗ To investigate the economic effect of low participation rates, a simulation exercise restricts the peer effect to zero and produces counter-factual individual investment participation. The exercise suggests that participation rates would be 7% to 10% higher in the absence of peer effects.

3.1 Introduction

Societal differences in financial inclusion have significant implications for wealth inequality (Favilukis, 2013), as well as labour market and economy-wide stability (Epstein & Shapiro, 2020). As individuals with greater incomes and education are more likely to marry one another (Nakosteen et al., 2004) and, at the same time, are more optimistic about future macroeconomic conditions (Das et al., 2020), the possibility exists that prudent financial behaviours, such as the use of investment products, perpetuate differently in groups of societies with homogeneous intra-group characteristics, with subsequent differential impacts on financial inclusion.

Prior research suggests that individuals influence each other in their decisions to participate in retirement schemes or the stock market (Arrondel, Calvo Pardo, Giannitsarou & Haliassos, 2020; Duflo & Saez, 2003; Haliassos, Jansson & Karabulut, 2019; Kaustia & Knüpfer, 2012). This literature however does not answer to what extent the tendency of individuals similar to each other to form relationships affects individual investment propensities via the peer effects channel. In this thesis chapter, we seek to fill this gap.

Specifically, this thesis chapter constitutes the first effort, to the best of our knowledge, to investigate whether such homophily in relationship formation affects investment participation while jointly testing for the presence of peer effects. To do so, we extend the selection-corrected spatial autoregressive model by Hsieh and Lee (2016), explicitly modelling latent variables driving both relationship formation and outcomes, to account for the limited-nature of binary dependent variables.

In order to solve the reflection problem (Manski, 1993), the difficulty of separating the peer effect from contextual effects, our model exploits the insights from the introduction of the spatial auto-regressive (SAR) model to the study of social networks (Lee, 2007; Lee, Liu & Lin, 2010; X. Lin, 2010). Specifically, non-linearities in the SAR model, generated by non-overlapping peer groups, identify the peer effect (Hsieh & Lee, 2016).

In this study, household networks that arise from split-off households in the longitudinal data obtained from Understanding Society and the British Household Panel Survey (BHPS) constitute these peer groups. Split-off households arise from existing households dissolving and their current members moving into new households. Members of these new households that are not part of the survey are invited to join, and their joining makes household networks observable in the data[1].

We focus on two measures of investment participation. First, respondents are asked whether they have received any form of investment income, namely any income from private pension/annuities, rent from boarders or lodgers, rent from any other properties and income from savings and investments in the last 12 months. This variable is available from 2011 onward. Second, we focus on a variable that captures whether respondents have received any income from interest or dividends. This variable is available from the years 2001 onward.

In the empirical analysis, we first estimate our binary network model without modelling the influence of latent variables, controlling for a host of demographic, socio-economic and spatial information. The estimation results suggest a positive and statistically significant peer effect with respect to both the general investment income and the interest and dividends variable. Further, we observe that the peers influence extends beyond that of their individual investment behaviour; peers educational and income levels infer a strong relation to investment participation. Unsurprisingly, individual age, personal income, retirement status and educational levels also influence individual investment propensities.

We proceed to investigate the effect of latent variables driving both investment

[1]Our definition of household networks relates to, but is broader than, the definition of extended family networks by Attanasio, Meghir and Mommaerts (2015), which are a strict subset of household networks.

participation and relationship formation. The optimal models for both the general investment as well as the interest and dividends variable suggest the existence of two latent variables that drive both investment participation and link formation. As a result, the peer effect for the investment income variable drops by 5% versus the baseline estimate; while that for the interest and dividends income variable drops by 25%. However, in both cases, the estimates remain strongly statistically significant. This corroborates the importance of homophily with respect to investment participation through the peer effect channel.

To test the effect peer effects have on overall investment participation levels in our data, we sample from the posterior predictive distributions of our selection-corrected model estimated with respect to both of our outcome variables of interest. Specifically, we first sample from the posterior distribution keeping the peer effects at their estimated values; second, we constrain the peer effects to zero. We find that peer effects depress overall investment participation in our data: we find that participation would be 7%-10% higher if peer effects were absent. This suggests the interpretation that individuals who are not participating in investments appear to effectively deter other individuals from investing.

This thesis chapter proceeds as follows. Section 3.2 lays out related literature. In Section 3.3, we describe the identification of the peer effect using household networks and the construction of our sample, as well descriptive statistics for our data. In Section 3.4, we introduce our Bayesian random effects probit model with self-selection of peer relationships. In Section 3.5, we report our estimation results; in Section 3.6 we show the effect of peers on overall investment participation; while Section 3.7 concludes.

3.2 Related literature

This thesis chapter relates to the literature on peer effects in financial behaviours. In their seminal contribution, Duflo and Saez (2003) show that informing a random subset of employees in a large university of the benefits of enrolling in a tax-deferred retirement plan increases not only the enrolment rates of those in attendance but also of those that were absent. Kaustia and Knüpfer (2012) document that individuals living in areas with better opportunities to learn from others are more likely to participate in the stock market. Using individuals expectations of their peers behaviours, Arrondel et al. (2020) find that both imitation of and information from peers are significant drivers of stock market participation. Making use of an exogenous allocation of refugees to different neighbourhoods in Sweden, Haliassos et al. (2019) show that exposure to financially literate increases participation intensity in private retirement schemes and the stock market, but only for educated refugee households and when interaction with their environment are possible.

The work in this chapter further contributes to the literature that shows how the self-selection of relationships can amplify peer effects due to unobserved characteristics that influence both the formation of peer relationships and economic outcomes. Weinberg (2007) appears to be the first to introduce a model of social interactions in which relationships between individuals form endogenously. Goldsmith-Pinkham and Imbens (2013) introduce a spatial autoregressive model that jointly models the self-selection of peers, including a binary latent variable that drives both the link formation between individuals and the continuous dependent variable of interest. For the case of one unobserved dimension, Badev (2013) extend this to binary choices in a Maximum Likelihood framework. Hsieh and Lee (2016) extend the latent variables model to include one or more continuous latent variables in a logit parametrisation of the self-selection equation; while Johnsson and Moon (2019) study unobserved drivers in a semi-parametric approach, leaving the form of the selection equation unspecified. Hsieh, Lee and Boucher (2019) extend the model by Hsieh and Lee (2016) to include incentives in selection, while Hsieh and Van Kippersluis (2018) allow for heterogeneity in peer effects.

Moreover, emerging developments in this literature estimate the peer effect jointly with unobservable social ties. In their seminal paper, de Paula, Rasul and Souza (2020) propose a method that uncovers social ties in panel data with no measurements on these relationships using the Adaptive Elastic Net GMM method. As cited by de Paula et al. (2020, p. 5), applications are already underway: for example, Fetzer, Souza, Eynde and Wright (2020) use this method to study the displacement of insurgency groups in Afghanistan as international security forces withdraw. Zhou (2019) contribute a model extension to account for link heterogeneity; while Z. Lin, Tang and Yu (2020) suggest the extension of link discovery and estimation of the peer effect for binary outcomes.

3.3 Data and household networks

3.3.1 Identification and household networks

In this section, we describe how split-off households in household longitudinal surveys can be employed to reveal household networks and how these, in turn, can be used to solve the reflection problem (Manski, 1993), the difficulty of separating the peer effect from contextual effects, which arises when studying investment participation under peer influence. To do so, we exploit the designs of longitudinal surveys.

Household longitudinal surveys, tracking the same individuals of a given household over time, suffer from attrition, i.e. respondents dropping out of the survey over time. Attrition threatens the surveys' representativeness, and in order to ensure that the surveys remain representative of their target population, the occurrence of split-off households is exploited: households that newly form when, for example, grown-up children move out of their parents' home and find shared flats, cohabitating couples break up and move into new homes or flatmates move in with their partners or other people. Members of these newly created households that were not part of the respective survey are invited to join, and by adding respondents to the survey in this way, the drop in respondent numbers is counteracted, ensuring the surveys' representativeness.

To illustrate, assume that every respondent is represented in a network by a node. Two nodes are connected by an edge if and only if these two respondents have lived together at some point in time. A collection of directly or indirectly connected nodes constitutes a household network. Household networks then grow when existing households dissolve and split-off households form: new nodes and edges appear when individuals begin living together and new household members join the survey.

Figure 3.7 illustrates how these spit-off households form and, consequently, household networks surface. The left-hand panel indicates an existing household at time t. Each node represents one of the household members a, b and c, and their nodes are connected as they live in the same household. At time $t + 1$, shown in the right-hand panel of Figure 3.7, their household has dissolved. Respondent a has moved in with individuals d, e and f; together they form the new household A. Respondent b has moved in with individual g, while respondent c is now living alone. Households A, B and C are split-off households, as they have split-off from the household at time t, and the individuals d, e, f and g have newly joined the survey. Together these new households form a household network. Household network continue to grow as time passes and more split-off households form.

Household networks vary in size, depending on the frequency of households in the respective network splitting, thereby providing econometric identification for peer effect parameters through the mechanism of non-overlapping peer groups (starting with Blume, Brock, Durlauf & Jayaraman, 2015; Bramoullé, Djebbari & Fortin, 2009; Lee, 2007). From an economic perspective, we argue that household networks provide a natural reference group to capture peer effects with respect to potentially many outcomes. Though friendships to individuals outside the household network are not observed in the data, it appears intuitive that the people respondents live with provide an important context for social learning, imitation or exchange of information.

3.3.2 Data sources and sample construction

We use data from the UK's Understanding Society and the British Household Panel Survey (BHPS) covering the years 2001 to 2016, fielded by the University of Essex and funded by the Economic and Social Sciences Research Council (ESRC), to conduct our analysis of investment participation under peer influence. Both surveys are closely related to each other and can be combined for analysis: in 1991, fieldwork started for the BHPS and in its 18th year, respondents were asked to join the larger Understanding Society survey, to which approximately 84% of respondents agreed. Aside from the BHPS sample of respondents, Understanding Society includes two further samples to ensure the its representativeness for the immigrant and ethnic subpopulations. These surveys are representative of the UK population and provide detailed insights into the respondents' lives over time. In particular, the data contain detailed information on the respondents financial and demographic characteristics. The full list of variables and details of their construction is reported in Appendix 3.A, and descriptive statistics are provided in the next section.

In order to construct our samples of interest, we separately require non-missing observations on the two key explanatory variables of interest: binary indicators capturing 1) whether the respondents receive any income from investments generally (private pension/annuities, rent from boarders or lodgers, rent from any other properties and income from savings and investments) captured from the year 2011 onwards and 2) whether the respondents receive any income from interest or dividends, elicited from 2001 onwards. Moreover, we also require for valid respondent observations to have non-missing information on a rich set of demographic characteristics, capturing in detail the respondents' age, education, employment status, number of children and gender. We additionally include the Government Office Region of where the respondents are domiciled in the UK and the year in which they are surveyed.

Next, we identify household networks using information on with whom respondents live in a given survey year, t. Specifically, both Understanding Society and the BHPS supply the table *egoalt* in which household membership of each respondent (the ego or "the self"; identified uniquely by the integer variable *pidp*) is reported jointly with their members of household (the alter or "the other"; identified uniquely by the integer variable *apidp*). Each row contains one *pidp*-*apidp* relationship, and as every respondent's identifier in *pidp* in one row appears as the identifier *apidp* in another, we remove redundant information by imposing that $pidp < apidp$.

In the next step, we iterate through the combined information on *pidp*-*apidp* pairs of the consecutive years t and $t+1$. In each step, we pool both year's information and remove redundant observations that arise when respondents live together at both time t and $t+1$. We then identify non-overlapping networks in the data by assigning directly or indirectly connected respondents into distinct groups. At this stage, the set of returned groups will contain households that have and have not split. In order to retain only those groups that result in household networks arising from split-off households, we require of valid groups to satisfy the network property that not all possible edges are present, i.e. some respondents never lived with each other. Put differently, in valid networks the number of actual edges is less than the number of possible ones; or, equivalently, the network density is less than unity.

This procedure results in a total sample of 5079 observations for our analysis with respect to the dividend and interest income variable, and 3200 observations for the investment income variable.

3.3.3 Descriptive statistics

In Panel A of Table 3.1, we show the counts, percentages and cumulative percentages of the relationship types of the respondents in the split-off households in our samples with respect to the investment income variable. Those for the interest + dividends income variable are unavailable due to data limitations. The table shows that a third (34.69%) of the relationship types in these households are that of a natural parent and child, while the next biggest groups constitute partner/cohabitee relationships (13.83%), two unrelated individuals living together (13.29%), natural siblings (12.07%) and husband and wife (10.48%). The large variability in relationship types is apparent in those types with less than 5% representation in the sample: it includes wider family relationships types, relationships between in-laws, foster relationships, in addition to relationship types that are missing in the data. Relationships of individuals that have never lived together are unrecorded and not reflected in the table.

In Panel B of Table 3.1, we show the absolute, relative and cumulative frequencies of household network sizes. With respect to the investment income variable, we observe variability in network sizes: most household networks have 3, 4 or 5 members, amounting to 86.15% of all networks. However, the distribution of household networks sizes exhibits a long tail, with the remaining networks being constituted of 6 to 15 members.

In Table 3.2, we report the newly-formed number of household networks, number of respondents and the degree of investment participation by year. With respect to the investment income variable in Panel (a), we observe data from year 2011 onward. The number of household networks in the years 2011 to 2016 lies between 76 and 175, with 329 to 820 network members. The average investment income participation lies between 21% to 29%. As missing values for the investment income variable are imputed, we observe a lot more respondents than for the interest and dividend income variable in Panel (b); for this, we observe data from the year 2001 to 2016, with the number of household networks identified in each year varying between 51 and 130, amounting to 207 to 593 respondents. Average interest and dividend income participation lies between 14% and 24%. Overall, we observe that investment participation across both measures and all years in our sample is low.

Table 3.3 indicates that the sample average investment participation as measured by the investment income and interest and dividend income variable is 26% and 18%, respectively. Additionally, Table 3.3 includes summary statistics for the individual- and household-level characteristics in our samples. The average age of respondents in our samples is 34.07 and 32.93. This is lower than the UK average of 39 according to the UK Census in 2011 (ONS, 2011). However, this is to be expected given our measurement strategy of household networks: as seen above, approximately a third (34.69%) of the observed relationships in our sample are those of a natural parent to one of its children. Children typically move out of the family home at in their early twenties, at which point the household has split, resulting in the parents being included in our sample at a younger age, depressing the average sample age.

We observe a lot of heterogeneity in the respondents' highest educational attainments. The majority of respondents hold A-levels (28%/29%), a university degree (25%/17%) or the GCSE (23%/25%). Some respondents have no educational qualifications (6%/10%), while others hold some other qualification (8%/10%) or some other higher qualification (10%/9%). At the same time, about two-thirds (66%/65%) of respondents are in full- or part-time employment as an employee or in self-employment. A small number of respondents in our sample look after their family or home (5%/6%), are retired (4%/3%) or are sick or disabled (2%/4%). The remaining respondents (23%/23%) indicated that they have another occupation. The average monthly gross labour income is equal to £1,154.01 and £1,002.99 in the different samples, respectively. On average, every fourth respondent in the sample is responsible for a child (0.23/0.29 number of children, on average). Further, the sample is made up of approximately equal numbers of females (51%/52%) and males (49%/48%), and it captures substantial spatial heterogeneity across the twelve Government Office Regions, varying from 3% (Northern Ireland) to 12% (South East).

3.4 Methodology

In this section, we describe our binary random effects multivariate probit model, its estimation with Bayesian Hamiltonian Monte Carlo sampling algorithms and its extension to account for endogeneity in the adjacency matrix of peers borrowing from the insights by Hsieh and Lee (2016) on how to model latent variables driving both link formation and outcomes.

3.4.1 Motivation

3.4.1.1 A simple example

The methodology proposed in this thesis chapter constitutes an extension of that by Hsieh and Lee (2016). The existing approach quantifies the peer effect in models with endogenous networks non-overlapping peer groups for continuous outcome variables. This chapter proposes the implementation for binary outcome variables. To see the importance of such an approach, consider the following example.

Let there be three individuals that influence each other in their investment propensity decisions, y_i^*, such that:

$$\begin{pmatrix} y_1^* \\ y_2^* \\ y_3^* \end{pmatrix} = \lambda \begin{pmatrix} 0 & w_{12} & w_{13} \\ w_{21} & 0 & w_{23} \\ w_{31} & w_{32} & 0 \end{pmatrix} \begin{pmatrix} y_1^* \\ y_2^* \\ y_3^* \end{pmatrix} + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \end{pmatrix} \tag{3.1}$$

where ε_i, for $i = 1, 2, 3$, are standard normal error terms. The latent propensities, y_i^*, are linked to the observable binary outcomes, y_i, in the conventional way through the indicator function $\mathbb{1}(\cdot)$: $y_i = \mathbb{1}(y_i^* > 0)$, where $\mathbb{1}(\cdot)$ takes the value one if its argument is true, and zero otherwise. $y_i = 1$ indicates that the individual i invests, while $y_i = 0$ means that she does not. The elements $w_{i,j}$ of the adjacency matrix take the value one if individual i has lived with individual j, and zero otherwise. For the case in which $i = j$, it also takes the value zero.

Simplifying the right-hand side of Equation (3.1) leads to:

$$\begin{pmatrix} y_1^* \\ y_2^* \\ y_3^* \end{pmatrix} = \begin{pmatrix} \lambda(w_{12}y_2^* + w_{13}y_3^*) + \varepsilon_1 \\ \lambda(w_{21}y_1^* + w_{23}y_3^*) + \varepsilon_2 \\ \lambda(w_{31}y_1^* + w_{32}y_2^*) + \varepsilon_3 \end{pmatrix} \tag{3.2}$$

From Equation (3.2) it is more easily seen that if there is a common driver that determines both w_{ij} and ε_i, where $i \neq j$, there will be bias in λ due to this omitted factor, which manifests in violation of the respective moment condition; specifically, this common driver will yield $E[(\sum_{\forall j, i \neq j} w_{ij} y_j^*)\varepsilon_i] \neq 0$.

To express this notion in more detail, assume that from the perspective of individual i, whether to live with j is determined by an unobserved factor z_i, so that $w_{ij} = F(z_i)$. And this factor also influences the propensity of i to invest, so that $\varepsilon_i = G(z_i) + \eta_i$ where η_i is white noise. Then $E[w_{ij}\varepsilon_i] = E[F(z_i)G(z_i)] \neq 0$, rendering w_{ij} endogenous and thereby introducing bias in an estimate of λ unless this endogeneity is explicitly accounted for.

This point can also be seen more generally. Consider Equation (3.1) expressed with appropriately defined matrices: $y^* = \lambda W(z)y^* + \varepsilon$, where the vector z collects the unobservable factors for the individuals. Simple rearranging yields $y^* = (I - \lambda W(z))^{-1}\varepsilon$, where I is the identity matrix and assuming the matrix $I - \lambda W(z)$ is invertible. Given the above relationships, we can see that $E[(I - \lambda W(z))^{-1}\varepsilon] \neq 0$.

3.4.1.2 Potential sources of endogeneity

What factors might these latent variables in the vector z capture? In order to be relevant in terms of endogeneity, they need to be (i) related to the likelihood that i has shared a household with j and (ii) to i's propensity to hold investments. In the context of the household networks introduced in Section 3.3.1, two types of factors are of significant import, reflecting the aforementioned distribution of relationship types visible in Table 3.2. These capture relationships between related individuals (for instance, mother/daughter) and those that are otherwise connected (for example, flatmates).

First, consider a relationship such as that between mother and daughter. At first glance, no endogeneity may seem to arise from these types of relationships; after all, mother and daughter share the same household almost by default. However, this shorthand view belies the role that genetics play in this relationship. More precisely, it is genetic similarity that determines their sharing of the same household. Conventionally, this is treated as unobservable. Accordingly, assume that the latent z_i captures the mother's genetic makeup and z_j that of her daughter. The genetic distance between the two can then be conceptualised as the difference $z_i - z_j = z_i - 0.5z_i = 0.5z_i$. This genetic difference then plays an important role in whether a household is shared; phrased in the notation of the preceding section: $w_{ij} = F(z_i) = 0.5z_i$. Further, genetics have been shown to exert significant influence over individual financial behaviours and outcomes (Barth, Papageorge & Thom, 2020; Cesarini, Johannesson, Lichtenstein, Sandewall & Wallace, 2010) and investment decisions (Barnea, Cronqvist & Siegel, 2010; Cronqvist & Siegel, 2014). In this context, genetics therefore constitute a potential source of endogeneity.

Second, in the case of relationships between individuals that are not related, demographic factors can play an important role in both household formation decisions as well as in investment participation. Research shows that individuals form relationships, and consequently households, on the basis of similarity in demographic and economic characteristics; for example, high-earning and highly educated individuals are more likely to marry one another (Nakosteen et al., 2004). At the same time, more specialist human capital such as financial literacy is known to significantly influence financial behaviours such as stock market participation (Van Rooij et al., 2011). While years of schooling and type of educational attainment are typically captured in surveys, more granular educational dimensions, such as financial literacy, are not. Therefore, in the context of this chapter, these characteristics are considered latent and therefore have the potential to generate a confounded measurement of the peer effect.

3.4.2 Data structures

We assume our data is made up of N individual-level observations that can be characterized by the $N \times 1$ vector of binary outcomes y, the $N \times (1+K+L)$ matrix of socio-economic and demographic individual- or household-level characteristics X, including the constant term, (where K and L are the number of continuous and dummy variables, respectively) and the $N \times D$ matrix Z. The data can then be jointly written as

$$\begin{pmatrix} y & X & Z \end{pmatrix}_{N \times (1+K+L+D)} \tag{3.3}$$

We assume that y and X are observable. In contrast, Z is assumed to contain unobservable individual characteristics of interest, possibly as a result of missing data or inherent measurement problems.

Further, we assume that each of the N observations in the matrix in (3.3) can be mapped to one and only one of G non-overlapping household networks. We assume that each network has m_g members, so that the total number of observations is equal to the sum of the G network sizes, $N = \sum_{g=1}^{G} m_g$.

Given these assumptions on the network structures, we can make partition the data in the matrix in (3.3) into the G component networks:

$$\begin{pmatrix} y_1 & X_1 & Z_1 \\ \vdots & \vdots & \vdots \\ y_g & X_g & Z_g \\ \vdots & \vdots & \vdots \\ y_G & X_G & Z_G \end{pmatrix} \quad \text{where} \quad y_g = \begin{pmatrix} y_{1,g} \\ \vdots \\ y_{i,g} \\ \vdots \\ y_{m_g,g} \end{pmatrix} X_g = \begin{pmatrix} x'_{1,g} \\ \vdots \\ x'_{i,g} \\ \vdots \\ x'_{m_g,g} \end{pmatrix} Z_g = \begin{pmatrix} z'_{1,g} \\ \vdots \\ z'_{i,g} \\ \vdots \\ z'_{m_g,g} \end{pmatrix}$$

$$\text{for } \; i = 1, \dots, m_g \text{ individuals}$$

$$\text{and } g = 1, \dots, G \text{ networks} \tag{3.4}$$

where y_g, X_g and Z_g are the outcomes, observable and unobservable characteristics, respectively, for the network g. The scalar $y_{i,g}$ denotes the outcome for individual i in network g, while $x'_{i,g}$ and $z'_{i,g}$ are row vectors of size $1 \times (1 + K + L)$ and $1 \times D$ of the observable and unobservable individual-level characteristics, respectively.

Each household network is associated with a $m_g \times m_g$ matrix capturing which individuals live together and which do not. This adjacency matrix is denoted W_g for $g = 1, \dots, G$. Its elements, $w_{ij,g}$, take values of one if the individuals i and j have lived together at any point in time, and zero otherwise, for $i = 1, \dots, m_g$ and $i = j, \dots, m_g$. For the purpose of this study, we assume that an individual does not live with itself, so that the $w_{ij,g} = 0$ if $i = j$, resulting in a diagonal of zeros. Two individuals living together intuitively implies an undirected relationship between i and j, resulting in symmetry of W_g: $w_{ij} = w_{ji}$ for all i and j:

$$W_g = (w_{ij,g}) = \begin{pmatrix} 0 & w_{12,g} & \cdots & w_{1m_g,g} \\ w_{21,g} & 0 & \cdots & w_{2m_g,g} \\ \vdots & \vdots & \ddots & \vdots \\ w_{m_g1,g} & w_{m_g2,g} & \cdots & 0 \end{pmatrix} \tag{3.5}$$

Given these assumptions, we are able to reconstruct any of these adjacency matrices from only their upper or lower triangular entries (excluding the diagonal entries), which facilitates efficient memory usage in computations. Importantly, we do not row-normalize W_g.

3.4.3 The likelihood function

Our outcome variable of interest, $y_{i,g}$ for individual i in network g, is an indicator variable that takes the value one if the individual receives a positive amount of investment income and zero otherwise. We ask whether the propensity of individual i to receive investment income is influenced by the respective propensities of those individuals in individual i's network g. For this purpose, we specify a linear-in-means peer effect model with group-level random effects (following Lee, Li and Lin (2014) in its specification of the random effects),

$$
\begin{aligned}
y_g^* &= \lambda W_g y_g^* + X_g \beta_1 + W_g X_g \beta_2 + \sigma_\alpha \alpha_g l_g + \varepsilon_g \\
\text{where } y_{i,g} &= \begin{cases} 1 \text{ if } y_{i,g}^* > 0 \\ 0 \text{ if } y_{i,g}^* \leq 0 \end{cases}
\end{aligned} \tag{3.6}
$$

where y^* is the vector of latent investment participation propensities, related to observed outcomes in the standard fashion. λ is the peer effect parameter of interest, while β_1 and β_2 are the coefficients of the respondents own covariates, X_g, and those of the respondents' peers group, $W_g X_g$, respectively. l_g is a vector of ones with length m_g. The terms $\sigma_\alpha \alpha_g l_g$, for $g = 1, \ldots, G$, capture the group random effects, where $\alpha_g \sim \mathcal{N}(0, 1)$ and $\sigma_\alpha > 0$, so that the products $\sigma_\alpha \alpha_g$ are i.i.d. $\mathcal{N}(0, \sigma_\alpha)$. The vector of errors follows a multivariate distribution with the identity matrix as its covariance matrix, $\varepsilon_g \sim \mathcal{N}(0, I_{m_g})$, reflecting the standard probit assumption of normalizing the standard deviation of the latent propensities to unity.

In order to capture the effects of unobserved variables on both the propensity to participate in financial investments and the likelihood of individuals living together when having similar latent characteristics, thus potentially amplifying the peer effect λ, we further adopt the formulation by Hsieh and Lee (2016) and assume that the error term $\varepsilon_{i,g}$ of individual i is jointly distributed with the vector of vector of unobserved variables, $z'_{i,g} = (z_{i1,g} \quad \cdots \quad z_{iD,g})$. Together they follow a multivariate distribution with a mean vector of zeros and covariance matrix Σ,

$$(\varepsilon_{i,g}, z_{i1,g}, \ldots, z_{iD,g})' \sim \mathcal{N}(0, \Sigma) \quad \text{where} \quad \Sigma = \begin{pmatrix} 1 & \rho_{\varepsilon z} \\ \rho_{z\varepsilon} & \Sigma_Z \end{pmatrix} \tag{3.7}$$

The covariance matrix of the vector $z_{i,g}$ is denoted by Σ_Z, while the vector $\rho_{z\varepsilon}$ captures the correlation of the latent variables in $z_{i,g}$ with the error term $\varepsilon_{i,g}$, thereby constituting the channel through which the individual unobserved characteristics can affect the investment participation propensities. Z, $\rho_{z\varepsilon}$ and Σ_Z cannot be identified so Σ_Z is normalized to equal the identity matrix of same size, I_D, and the correlation coefficients in $p_{\varepsilon z} = (p_{1,\varepsilon z} \cdots p_{D,\varepsilon z})$ are normalized to be positive Hsieh and Lee (2016).

To capture the effects of the latent variables in Z_g on the likelihood that two individuals live together, we focus on the undirected relationships between individuals i and j where $i < j$. This is equivalent to modelling the upper triangular potion (excluding the diagonal entries) of the adjacency matrix W_g, which fully characterizes the whole matrix due to the symmetry and diagonal of zeros of W_g, as discussed above . Specifically, we assume that the probability that i and j live together is a function of the absolute differences of continuous observables, indicators of shared binary characteristic and the absolute differences between their unobservable characteristics (Fafchamps & Gubert, 2007a, 2007b). We further assume that this probability can be approximated by the standard logistic function:

$$
\begin{aligned}
P(w_{ij,g} = 1 \mid c, \gamma, \theta, \delta, x'_{i,g}, z'_{i,g}) &= \frac{\exp(\psi_{ij,g})}{1+\exp(\psi_{ij,g})} \\
\text{where } \psi_{ij,g} &= c + \sum_{\bar{k}=1}^{\bar{K}} \gamma_{\bar{k}} |x_{i\bar{k},g} - x_{j\bar{k},g}| + \sum_{\bar{l}=1}^{\bar{L}} \theta_{\bar{l}} \mathbb{1}(x_{i\bar{l},g} = 1 \wedge x_{j\bar{l},g} = 1) \\
&\quad + \sum_{d=1}^{D} \delta_d |z_{id,g} - z_{jd,g}| \qquad (3.8)
\end{aligned}
$$

where $\bar{K}$ denotes the number of absolute differences of the continuous observable variables between i and j, $\bar{L}$ the number of indicators of shared binary characteristics and D the number of absolute differences in latent variables. $\mathbb{1}(\cdot)$ again denotes the indicator function that takes the value one if its argument is true, and zero otherwise. The regression constant is denoted by c. In order to identify the vector $\delta' = (\delta_1 \cdots \delta_D)$, we assume that its elements are sorted in ascending order, $\delta_1 \leq \ldots \leq \delta_D$.

Equations (3.6) to (3.8) together specify how the latent variables in Z_g affect both the propensity to participate in financial investments and the likelihood of living with individuals who have similar latent characteristics, thus introducing potential amplification of the peer effect λ. How such latent variables can act as confounders in the estimation of peer effects is easily seen from the conditional distribution of the latent propensities y_g given the unobserved covariates Z_g for $g = 1, \ldots, G$,

$$y_g^* = \lambda W_g y_g^* + X_g \beta_1 + W_g X_g \beta_2 + Z_g \Sigma_Z^{-1} \rho_{z\varepsilon} + \sigma_\alpha \alpha_g l_g + \nu_g \qquad (3.9)$$

The error term distribution equals $\nu_g \sim \mathcal{N}(0, \sigma_\nu^2 I_{m_g})$ with $\sigma_\nu^2 = 1 - \rho_{z\epsilon} \Sigma_Z^{-1} \rho_{\epsilon z}$. This conditional distribution in Equation (3.9) can be derived from Equation (3.7) given the standard properties of the multivariate normal distribution. The term $Z_g \Sigma_Z^{-1} \rho_{z\varepsilon}$ captures the effect of the unobserved variables on y_g^*, which would be erroneously attributed to λ if it were omitted.

In order to derive the joint distribution function of the latent propensities, the conditional distribution of y_g^* in Equation (3.9) can be manipulated by subtracting $\lambda W_g y_g^*$ from both sides, factoring out y_g^* on the left-hand side and pre-multiplying both sides by the Leontief-inverse $(I_{m_g} - \lambda W_g)^{-1}$. This yields:

$$y_g^* = B_{\lambda,g}^{-1}\mu_g + u_g \tag{3.10}$$

$$B_{\lambda,g} = I_{m_g} - \lambda W_g \tag{3.11}$$

$$\mu_g = X_g\beta_1 + W_g X_g \beta_2 + Z_g \Sigma_Z^{-1}\rho_{z\epsilon} + \sigma_\alpha \alpha_g l_g \tag{3.12}$$

$$u_g = B_{\lambda,g}^{-1}\nu_g \tag{3.13}$$

where $u_g \sim \mathcal{N}\big(0, \sigma_\nu^2\big[B'_{\lambda,g}B_{\lambda,g}\big]^{-1}\big)$. Equations (3.10) to (3.13) constitute the main model equations. Together they imply that the vectors y_g^*, for $g = 1, \ldots, G$, are distributed with heteroskedasticity, following a multivariate normal distribution with mean vector $B_{\lambda,g}^{-1}\mu_g$ and covariance matrix $\sigma_\nu^2\big[B'_{\lambda,g}B_{\lambda,g}\big]^{-1}$.

The product $B_{\lambda,g}^{-1}\mu_g$ can be obtained efficiently by system-of-equations solver, without requiring to compute the inverse $B_{\lambda,g}^{-1}$ explicitly; a process which would be both comparably slow and numerically unstable (Stan Development Team, 2018). Additionally, in order to avoid computation of the matrix inverse in the covariance matrix, we employ the parametrisation of the multivariate normal distribution in terms of the precision matrix, defined as the inverse of the covariance matrix. Therefore, the latent propensities for the networks g, for $g = 1, \ldots, G$, can be expressed as a multivariate normal distribution with mean vector $B_{\lambda,g}^{-1}\mu_g$

and precision matrix $(1/\sigma_\nu^2)B'_{\lambda,g}B_{\lambda,g}$,

$$y_g^* \sim \mathcal{N}(B_{\lambda,g}^{-1}\mu_g, (1/\sigma_\nu^2)B'_{\lambda,g}B_{\lambda,g}) \tag{3.14}$$

$$y_{i,g} = \begin{cases} 1 \text{ if } y_{i,g}^* > 0 \\ 0 \text{ if } y_{i,g}^* \leq 0 \end{cases} \tag{3.15}$$

Chib and Greenberg (1998) show that a multivariate distributions such that as described by Equation (3.14), for $g = 1, \ldots, G$, can be thought of as truncated at zero as a result of the bounds implied by Equation (3.15). In the below, we further draw on the insights by Chib and Greenberg (1998) to draw from the resulting truncated multivariate normal distribution.

Consequently, the probability of observing the vector of ones and zeros y_g, for $g = 1, \ldots, G$, is given by integrating over the density function of latent variables,

$$P(y_g \mid Z_g, \Omega) = \int_{A_{m_g,g}} \cdots \int_{A_{1,g}} P(y_g^* \mid Z_g, \Omega)\, dy_g^* \tag{3.16}$$

where $\Omega = \{\lambda, \beta_1, \beta_2, \sigma_\alpha, \{\alpha_g\}, \rho_{\varepsilon z}\}$ denotes the set of parameter of the peer effects equation and $P(\cdot)$ denotes probability mass or density functions. The bound of integration equals $A_{i,g} = (-\infty, 0]$ if $y_{i,g} = 0$ and $A_{i,g} = [0, \infty)$ if $y_{i,g} = 1$ for $i = 1, \ldots, m_g$. We exclude the remaining exogenous variables for notational convenience. The likelihood function for $y = (y'_1 \cdots y'_g)'$ is then given by

$$P(y \mid \{Z_g\}, \Omega) = \prod_{g=1}^{G} P(y_g \mid \{Z_g\}, \Omega) \tag{3.17}$$

The expression of the posterior density of the parameter values then follows from Bayes' theorem:

$$P(\{Z_g\}, \Omega \mid y) \propto P(\{Z_g\}, \Omega) \cdot P(y \mid \{Z_g\}, \Omega) \tag{3.18}$$

However, in our estimation, we will focus on the joint posterior distribution of the parameters and the latent $\{y_g^*\}$, which follows from Bayes' theorem and the relationship $P(Y, X) = P(Y \mid X)P(X)$ for two random variables Y and X,

$$\begin{aligned} P(\{y_g^*\}, \{Z_g\}, \Omega \mid y) \propto P(\{Z_g\}, \Omega) \cdot f(\{y_g^*\} \mid \{Z_g\}, \Omega) \\ \cdot P(y \mid \{y_g^*\}, \{Z_g\}, \Omega) \end{aligned} \tag{3.19}$$

Noting that the conditional probability of observing the vector y_g depends on whether the latent variables y_g^* fall in the set $A_g = A_{1,g} \times \cdots \times A_{m_g,g}$, where $\times$ denotes the Cartesian product, it holds that $P(y \mid \{y_g^*\}, \{Z_g\}, \Omega) = \mathbb{1}(\{y_g^* \in A_g\})$, where $\mathbb{1}$ denotes the indicator function as introduced above, which simplifies the posterior density to:

$$P(\{y_g^*\}, \{Z_g\}, \Omega \mid y) \propto P(\{Z_g\}, \Omega) \cdot f(\{y_g^*\} \mid \{Z_g\}, \Omega) \cdot I(\{y_g^* \in A_g\}) \tag{3.20}$$

By analogous reasoning, the posterior density, including the relevant terms for the selection equation, is

$$\begin{aligned} P(\{y_g^*\}, \{Z_g\}, \Omega, \Psi \mid y) \propto P(\{Z_g\}, \Omega, \Psi) \cdot P(\{y_g^*\} \mid \{Z_g\}, \Omega) \\ \cdot I(\{y_g^* \in A_g\}) \cdot P(\{W_g\} \mid \{Z_g\}, \Psi) \end{aligned} \tag{3.21}$$

where $\Psi = \{c, \gamma, \theta, \delta\}$ denotes the set of parameters of the selection equation in (3.8). Next, we turn to specifying the functional forms of the (mutually independent) prior distributions, $P(\{Z_g\}, \Omega, \Psi)$.

3.4.4 Prior densities

In order to ensure that the matrix $B_{\lambda,g}$ is invertible for all values of λ, we follow the standard approach of restricting λ to lie strictly between $-1/\tau_G$ to $1/\tau_G$ where

τ_G is equal to the number of edges of the best-connected node across all of the G networks (Hsieh & Lee, 2016; Hsieh & Van Kippersluis, 2018; Kelejian & Prucha, 2010). On this interval, we specify a uniform prior distribution over the admissible values of λ, $\lambda \sim U(-1/\tau_G, 1/\tau_G)$.

In Equation (3.6) and (3.8), we scale continuous regressors to have zero mean and a standard deviation of 0.5 and the binary regressors to have mean zero, as suggested by Gelman, Jakulin, Pittau, Su et al. (2008), and therefore follow Ghosh, Li, Mitra et al. (2018) in placing weakly informative Student-t priors on the elements in the coefficient vectors γ, θ, β_1 and β_2. Specifically, these priors are $\beta_{k,1} \sim t_6(0, 2.5)$ and $\beta_{l,2} \sim t_6(0, 2.5)$ for $k = 1, \ldots, K$ and $l = 1, \ldots, L$; $\gamma_{\bar{k},1} \sim t_6(0, 2.5)$ and $\theta_{\bar{l},1} \sim t_6(0, 2.5)$ $\bar{k} = 1, \ldots, \bar{K}$ and $\bar{l} = 1, \ldots, \bar{L}$, For the constant in the logistic regression, we choose $c \sim t_6(0, 10)$.

The standard deviation of the group-level random effects follows a weakly informative, Half-t prior distribution, restricting admissible values of σ_α to be positive, $\sigma_\alpha \sim Half\text{-}t_6(0, 10)$, while α_g, for $g = 1, \ldots, G$ have i.i.d. standard normal priors, $\alpha_g \sim \mathcal{N}(0, 1)$.

In accordance with the normalization assumption $\Sigma_Z = I_D$, we place informative i.i.d. standard normal priors on the elements of matrix of latent variables Z, $z_{i,d} \sim \mathcal{N}(0, 1)$ for $i = 1, \ldots, N$ and $d = 1, \ldots, D$. Additionally, we place informative beta priors on the D correlation coefficients, which are normalized to lie in the interval $[0, 1]$, $\rho_{d,z\epsilon} \sim \beta(1, 10)$, and i.i.d. standard normal priors on δ_d, $\delta_d \sim \mathcal{N}(0, 1)$, for $d = 1, \ldots, D$, in a way that the normalization constraint $\delta_1 \leq \ldots \leq \delta_d \leq \ldots \delta_D$ is satisfied.

3.4.5 Estimation

We use RStan (Stan Development Team, 2018) to draw from the posterior distribution with Hamiltonian Monte Carlo (HMC). HMC is a Markov chain Monte Carlo (MCMC) method that uses the analytical gradients of the density distributions to sample efficiently from the target posterior distribution (Betancourt, 2017; Betancourt & Girolami, 2015). The implementation details are complex and involve advanced differential geometry; however, from a user-perspective, estimating the parameters of the distribution amounts to writing a Stan program that specifies the prior distributions and likelihood functions that make up the posterior distribution. Stan then generates the draws from the posterior distribution.

3.5 Estimation results

In this section, we report the estimation results for both our binary spatial autoregressive (B-SAR, $D = 0$) model, assuming the adjacency matrix is exogenously given, and the binary selection-corrected spatial autoregressive model (B-SCSAR, $D > 0$), jointly modelling the peer relationships, using the methodology described in the previous section. For each estimation, we specify four Markov Chain Monte Carlo (MCMC) chains to run in parallel, using a different set of diffuse initial values for each. Each chain is set to run for 10000 iterations, where the first 4000 are used to arrive at the high density region of the stationary distribution. In order to determine the number of relevant latent variables, we estimate various B-SCSAR models, increasing the number of latent variables from $D = 1$ to $D = 4$, and check for convergence of all model parameters using the $\hat{R}$ convergence statistic by Gelman and Rubin (1992): a statistic indicative of whether all MCMC chains identify the same target distribution. In addition to the estimation results, we include plots of the MCMC chains and their respective correlograms for the peer effect parameter, λ. Moreover, we report results of the constrained and unconstrained posterior predictive distributions for the B-SAR model.

3.5.1 General investment income

In this subsection, the dependent variable takes the value one if the respondent has received any income from a private pension/annuities, rent from boarders or lodgers, rent from any other properties and income from savings and investments, and zero otherwise. To control for heterogeneity in the respondents, we include the respondents' own age, number of children, gross monthly income, occupational status (employed, retired, sick or disabled, other job status), education (degree, GCSE, A-level, higher degree, other education) and gender. For each respondent, we capture the peer group characteristics along these same variables. Moreover, we include year effects, region effects and group-level random effects. In the selection equation, the covariates include the difference of the respondents' ages, no of children and personal incomes, as well as binary variables that indicate whether the respondents share the same occupational status, education or gender. Moreover, we include year effects to account for possibly changes in relationship formation propensities over time. In both the outcome and selection equation, all continuous covariates are scaled to have mean zero and standard deviation of 0.5, while all binary covariates are centred at mean zero. The Appendix reports all variable definitions.

Figure 3.2 shows the sets of consecutive draws of the peer effects parameter, λ, from the posterior distribution. Panel (a) shows the MCMC draws for the B-SAR model, in which no latent variables are specified ($D = 0$), while Panel (b) shows the MCMC draws for the B-SCSAR model with two latent variables ($D = 2$). In both graphs we observe that all four MCMC chains appear to be mixing well. Figure 3.3 shows the corresponding correlograms for each of the four MCMC chains. Panel (a), again, shows the figures for the chains of the B-SAR model, while Panel (b) displays those of the B-SCSAR model. In all correlograms, the autocorrelations appear to die out quickly. In Panel (b), the fourth chain is the only one that exhibits some noteworthy autocorrelation beyond lag 15, indicating minor difficulty of the Hamiltonian Monte Carlo sampler to sample the posterior distribution given the initial values of this respective chain.

Table 3.4 reports the estimation results for the B-SAR ($D = 0$) and B-SCSAR($D = 2$) models. Panel A shows the results for the outcome equation. The point estimates correspond to the posterior means of the marginal distributions. For the BSAR model, the peer effects equals 0.058 and is strongly significant at the 1% level. This result indicates that, assuming that cohabiting relationships in the household networks are exogenously given, the behaviour of respondents' peers constitute a strong influence for their own investment participation. Among the demographic attributes, we observe that the respondents' age, personal income and educational levels (degree, A-level and higher degree) infer a strong positive relation to the respondents' investment participation propensities, significant at the 1% level. The number of children of the respondent are strongly negatively associated to investment, also significant at the 1% level. With respect to the characteristics of the respondents' peers, education (degree, A-level, higher degree) is positively associated to the respondents' individual investment propensity (at the 5%, 5% and 10% level, respectively).

With respect to the B-SCSAR model, we observe a peer effect of 0.055, highly significant at the 1% level. This estimate indicates that even after accounting for latent characteristics that drive both relationship formation and individual propensities, peers exert a significant influence on individual propensities; however, contrasting this result with the estimate obtained from the B-SAR model, we note that the peer effect drops by approximately 5%. It thus appears that if peers select each other based on unobserved characteristics that also drive investment participation propensities, the observable peer effect is amplified. The remaining estimates for the covariates (own and peers) in the B-SCSAR model are numerically similar in magnitude and significance to the ones in the B-SAR model.

Panel B of Table 3.4 reports the estimates for the selection equation of the B-SCSAR model. The propensity for individuals to live together is lower, the greater their age difference. Differences in number of children and personal income also decrease the likelihood of two individuals living together. All these estimates are significant at the 1% level. In contrast, individuals are more likely to have a cohabiting relationship if they share the same level educational attainment, significant at the 5% level. With respect to differences in the latent, the first and largest effect of these differences is significant at the 1% level, while the difference relating to the second latent variable is insignificant. Overall, the estimates of the selection equation confirm the existence of significant drivers of relationship formation, both observable and unobservable.

3.5.2 Interest and dividend income

In this subsection, the dependent variable takes the value one if the respondent has received any interest or dividend income during the last 12 months, and zero otherwise. As in the previous section, to control for respondent-level heterogeneity, we include the respondents' own age, number of children, gross monthly income, occupational status (employed, retired, sick or disabled, other job status), education (degree, GCSE, A-level, higher degree, other education) and gender. And again, for each respondent, we capture the characteristics of the relevant peer group along these same variables. We also include year effects, region effects and group-level random effects. In the selection equation, the covariates include the difference of the respondents' ages, no of children and personal incomes, as well as binary variables that indicate whether the respondents share the same occupational status, education or gender. We also include year effects to account for possible changes in relationship formation propensities over time. In both the outcome and selection equation, all continuous covariates are scaled to have mean zero and standard deviation of 0.5, while all binary covariates are centred at mean zero.

Figure 3.4 shows the sets of consecutive draws of the peer effects parameter, λ, from the posterior distribution. Panel (a) shows the MCMC draws for the B-SAR model, in which no latent variables are specified ($D = 0$), while Panel (b) shows the MCMC draws for the B-SCSAR model with two latent variables ($D = 2$). In both graphs we observe that all four MCMC chains appear to be mixing well. Figure 3.5 shows the corresponding correlograms for each of the four MCMC chains. Panel (a), again, shows the figures for the chains of the B-SAR model, while Panel (b) displays those of the B-SCSAR model. In the correlograms for the B-SAR models, we observe that the autocorrelations appear to die out quickly. In Panel b), exhibit some noteworthy autocorrelation, indicating more difficulties of the Hamiltonian Monte Carlo sampler to sample the posterior distribution given the initial values of this respective chain.

Table 3.5 reports the estimation results for the B-SAR ($D = 0$) and B-SCSAR ($D = 2$) models. Panel A shows the results for the outcome equation. The point estimates correspond to the posterior means of the marginal distributions. For the B-SAR model, the peer effects equals 0.085 and is significant at the 1% level. Among the demographic attributes, we observe that the respondents' age, personal income and educational levels (degree, GCSE, A-level and higher degree) infer a strong positive relation to the respondents' investment participation propensities, while the number of children relates negatively; all estimates being significant at the 1% level. Being male is also associated with higher investment propensities, albeit at the 5% level. With respect to the characteristics of the respondents' peers, the only significant predictor of the respondents' investment participation propensities is the peers' number of children: they infer a negative relationship to the respondents' investment propensities, significant at the 5% level. The remaining estimates for the covariates (own and peers) in the B-SCSAR model are numerically similar in magnitude and significance to the ones in the B-SAR model.

With respect to the B-SCSAR model, we observe a peer effect of 0.066, significant at the 5% level. The comparison in estimates of the peer effects between the B-SAR and B-SCSAR models indicate that latent characteristics that drive both relationship formation and individual propensities significantly amplify the peer effect in investment participation in interest- or dividend-yielding products. The peer effect drops by approximately 23% in magnitude, once endogenous relationship formations are accounted for.

Panel B of Table 3.5 reports the estimates for the selection equation of the B-SCSAR model. The propensity for individuals to live together is lower, the greater their age difference. Differences in number of children and personal income also decrease the likelihood of two individuals living together. All these estimates are significant at the 1% level. In contrast, individuals are more likely to form a cohabiting relationship if they share the same level educational attainment or occupational status, both significant at the 1% level. With respect to differences in the latent, the both of these are significant at the 1% level. Overall, the estimates of the selection equation confirm the existence of significant drivers of relationship formation, both observable and unobservable.

3.6 Peer effects and overall investment participation

In this section, we ask the question: what changes to the overall investment participation could we expect in the absence of peer effects? To do so, we inspect the posterior predictive distributions (PPD) of the B-SCSAR models of the previous two sections, with and without peer effects. Specifically, this amounts to simulating the individual investment participation from the fitted model without parameter restrictions and comparing it to simulations where the peer effect parameter is restricted to zero, $\lambda = 0$.

The relevant PPD takes the following form,

$$P\left(\{\widetilde{y}_g^*\} \mid \{Z_g\}, \Omega, \{X_g\}, \{W_g\}\right) \tag{3.22}$$

where $\tilde{y}_g^*$ denotes simulated values for the latent propensities for $g = 1, \ldots, G$. $P(\cdot)$ again denotes the probability density function. The parameters for the peer effect and selection equation are collected in the sets $\Omega = \{\lambda, \beta_1, \beta_2, \sigma_\alpha, \{\alpha_g\}, \rho_{\varepsilon z}\}$. The elements of the vector of simulated individual investment participation, $\tilde{y}_g$, take the value one if the corresponding simulated latent propensities are greater than zero, and zero otherwise.

For the general investment variable, the B-SCSAR model without restricting the peer effects parameter generates a simulated investment participation of 33.62%, while restricting the peer effect parameter to equal zero ($\lambda = 0$), yields a simulated participation level of 35.09%. The estimated probability that the investment participation simulated from the unrestricted model is greater than that of the restricted model equals 0.9018, and, therefore, is significant at the 10% level.

In the case of investment participation in interest or dividend yielding products, the unrestricted B-SCSAR model generates a simulated investment participation of 24.00%, while restricting the peer effect parameter to equal zero ($\lambda = 0$), yields a simulated participation level of 26.53%. The estimated probability that the investment participation simulated from the unrestricted model is greater than that of the restricted model equals 0.9985, and, therefore, is significant at the 1% level. Overall, the PPDs from the B-SCSAR models suggest that peers exert a negative influence on levels of investment participation in the UK.

3.7 Conclusion

This thesis chapter provides evidence that unobservable characteristics significantly influence peer choice jointly with the decision whether to participate in holding investments, using UK data from Understanding Society and the British Household Panel Survey for the years 2001 to 2016. In order to identify the peer effect in our empirical analysis, we use non-overlapping household networks generated through split-off households. Employing Bayesian binary choice models that uncover the peer effect and account for endogenous relationship formation, two important results transpire. First, the endogenous choice of peers accounts for approximately 25% of the peer effect in investment participation. This implies that approximately a quarter of the investment participation rate attributed to the effects of social learning or behavioural imitation are in fact a result of with whom individuals decide to associate. Second, a simulation exercise in which the peer effect, net of the effect of endogenous peer choice, is restricted to zero generates rates of counter-factual investment participation 7% to %10 higher than that when no restrictions are imposed.

Together, these findings have important implications for policy makers. As the evidence suggests that factors exist that make it more likely for individuals in social proximity to be self-similar as well as influence their propensities to invest, policy makers ought to implement measures that prevent imprudent financial behaviours to self-perpetuate in distinct groups in society. Ensuring that individuals from different socio-economic backgrounds have chances for interaction can be one approach policy makers can take. Such an approach is necessarily structural; it can be made operational, for instance, by ensuring that social housing is not relegated to remote geographic areas that make it difficult for its residences to interact with different socio-economic strata of society, thereby decreasing chances to strike up relationships with individuals different to themselves and missing out on observing examples of prudent financial behaviours. Further, policy makers can ensure that affordable office space is available in urban areas – for public and third sector organisations, for example – in which lucrative private sector businesses reside, so that social interactions and informal knowledge exchange between individuals of these sectors can be facilitated.

Figure 3.1

Household network growth process

This figure visualizes how "household networks" arise from "split-off households". Nodes indicate respondents and connected edges indicate that two respondents have lived together at some point in time, or currently do so. The left-hand panel shows an existing household at time t. Each node represents one of the household members a, b and c. The right-hand panel shows the new household configuration at time $t+1$. Respondent a has moved in with individuals d, e and f; together they form the new household A. Respondent b has moved in with individual g, while respondent c is now living alone. Their new households are denoted B and C. Households A, B and C are called split-off households, as they have split-off from the household at time t. The individuals d, e, f and g have newly joined the survey. Together, these new households form a household network, which continues to grow as time passes and more split-off households form.

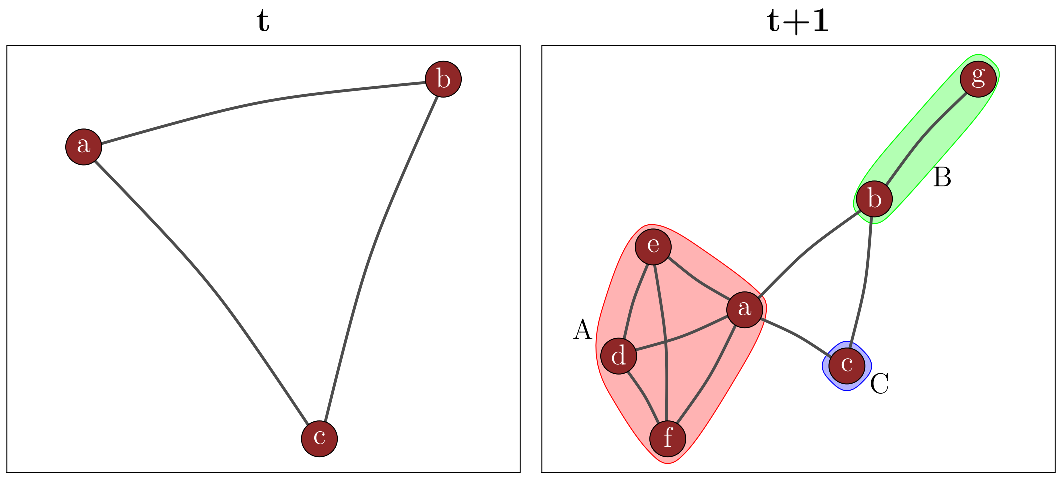

Figure 3.2

Traceplots for investment income models

This figure shows the traceplots of the peer effect (λ) for models in which the investment income indicator is the dependent variable. Panel (a) shows the traceplots for the model with no latent variables ($D = 0$), while Panel (b) shows those for the model with two latent variables ($D = 2$). The draws for four MCMC chains are overlaid, were each run is comprised of 10000 draws, where the first 4000 warmup draws of the Hamiltonian Monte Carlo sampler are discarded. Variable definitions are provided in Appendix A.

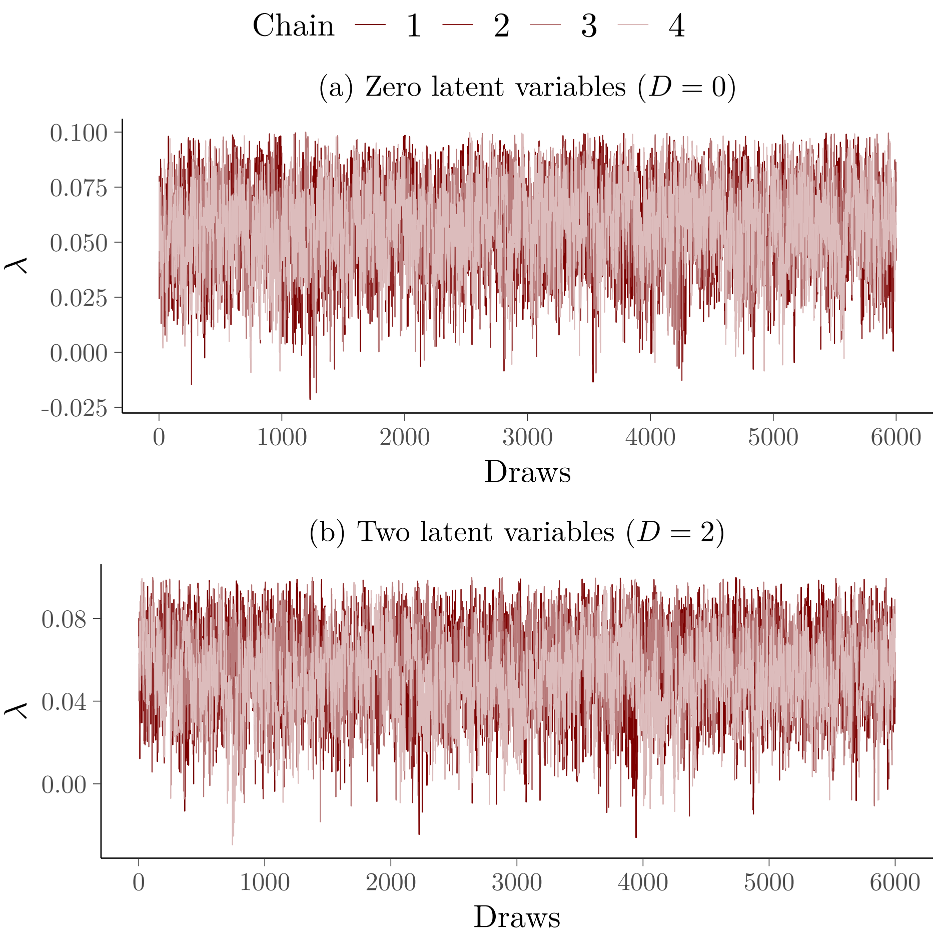

Figure 3.3

Correlograms for investment income models

This figure shows the correlograms of the peer effect (λ) for models in which the investment income indicator is the dependent variable. $\rho_\lambda(L)$ denotes the autocorrelation for λ relative to lag L. Panel (a) shows the correlograms for the model with no latent variables ($D = 0$), while Panel (b) shows those for the model with two latent variables ($D = 2$). The correlograms for four MCMC chains are overlaid, were each run is comprised of 10000 draws, where the first 4000 warmup draws of the Hamiltonian Monte Carlo sampler are discarded. Variable definitions are provided in the Appendix.

(a) Zero latent variables ($D = 0$)

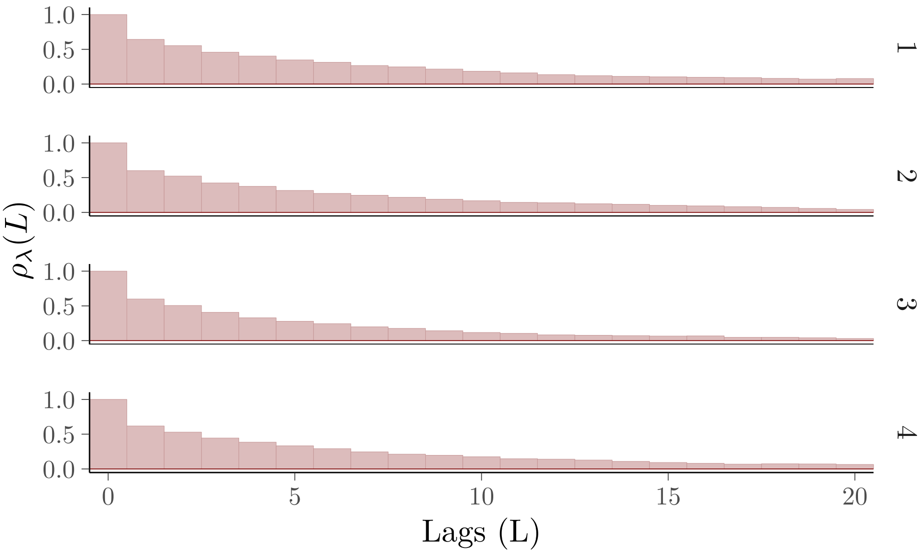

(b) Two latent variables ($D = 2$)

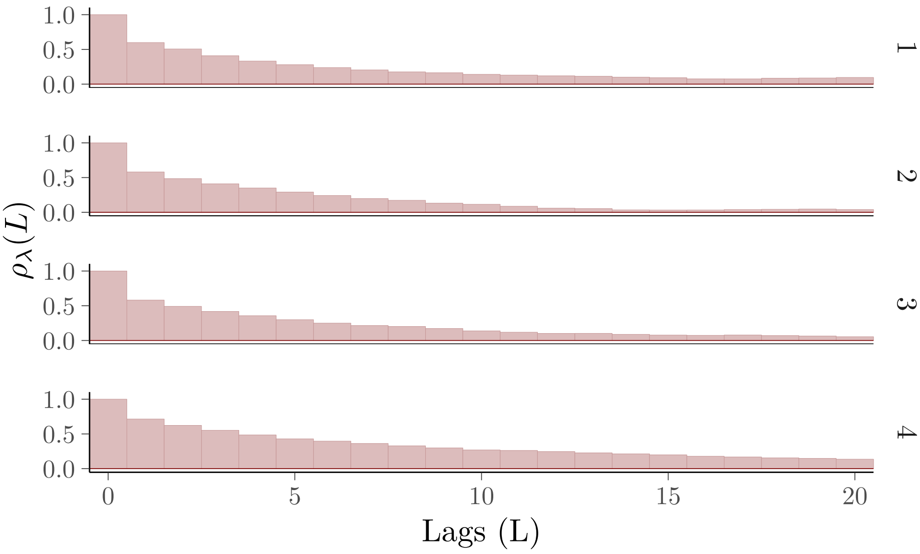

Figure 3.4

Traceplots for investment + dividend income models

This figure shows the traceplots of the peer effect (λ) for models in which the interest + dividend income indicator is the dependent variable. Panel (a) shows the traceplots for the model with no latent variables ($D = 0$), while Panel (b) shows those for the model with two latent variables ($D = 2$). The draws for four MCMC chains are overlaid, were each run is comprised of 10000 draws, where the first 4000 warmup draws of the Hamiltonian Monte Carlo sampler are discarded. Variable definitions are provided in the Appendix.

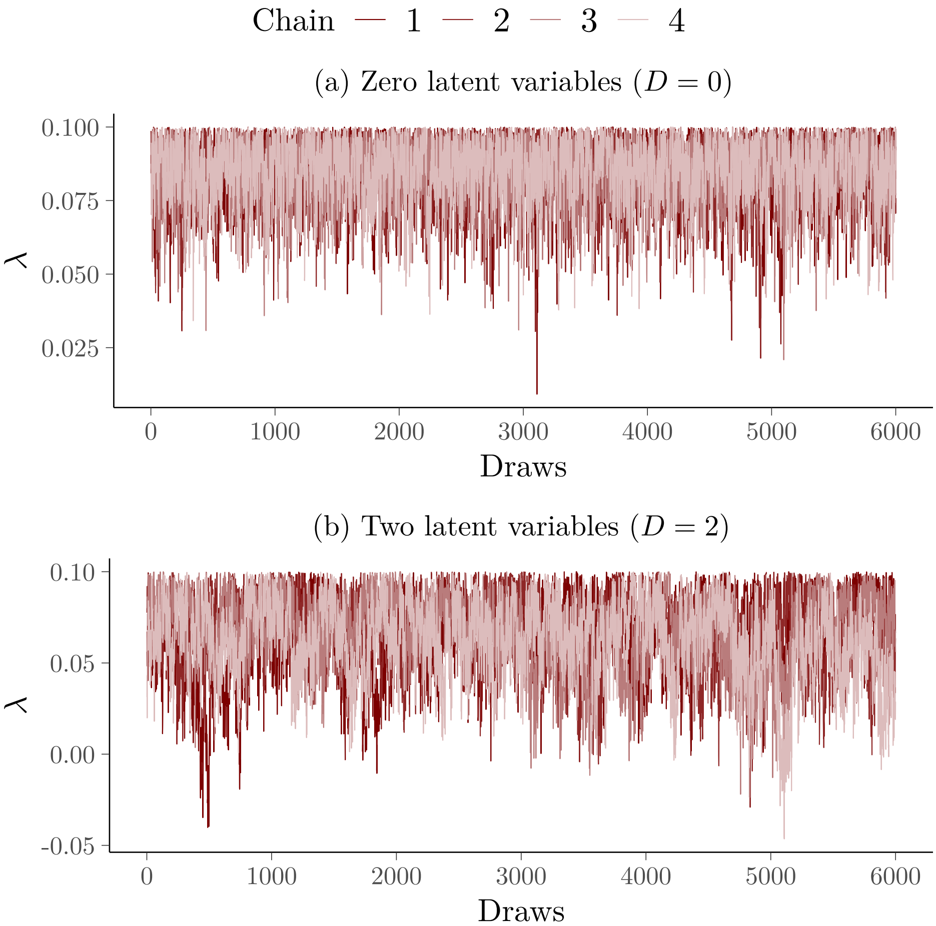

Figure 3.5

Correlograms for interest + dividend income models

This figure shows the correlograms of the peer effect (λ) for models in which the interest + dividend income indicator is the dependent variable. $\rho_\lambda(L)$ denotes the autocorrelation for λ relative to lag L. Panel (a) shows the correlograms for the model with no latent variables ($D = 0$), while Panel (b) shows those for the model with two latent variables ($D = 2$). The correlograms for four MCMC chains are overlaid, were each run is comprised of 10000 draws, where the first 4000 warmup draws of the Hamiltonian Monte Carlo sampler are discarded. Variable definitions are provided in the Appendix.

(a) Zero latent variables ($D = 0$)

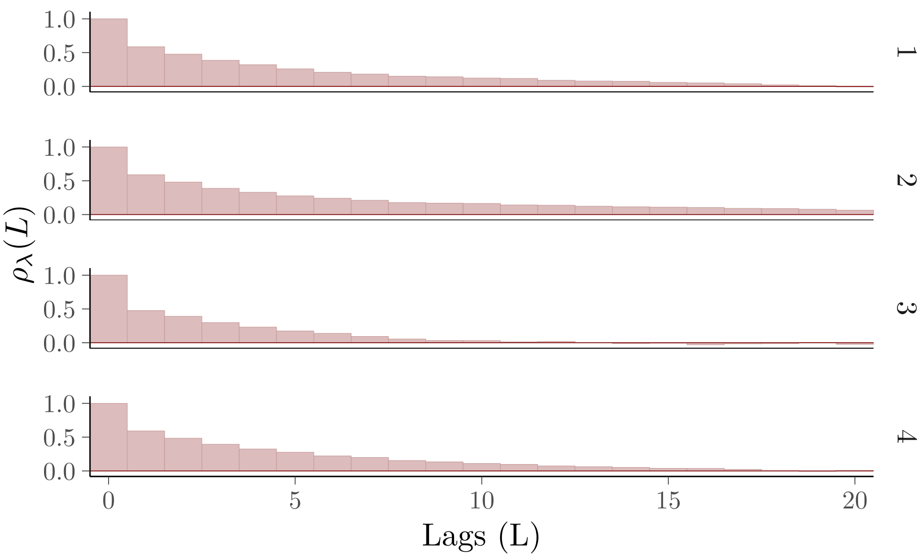

(b) Two latent variables ($D = 2$)

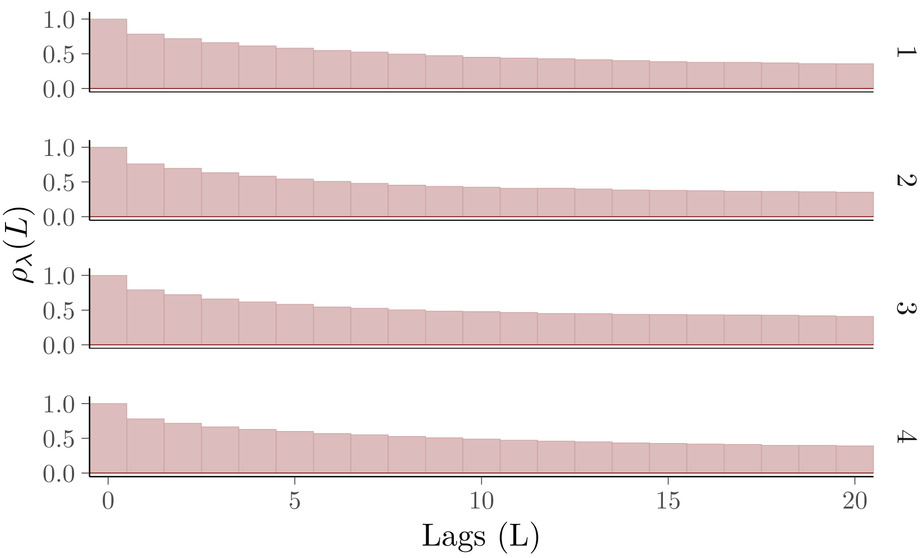

Table 3.1

Relationship type and network size summary statistics

This table reports sample summary statistics (counts, percentages and cumulative percentages) for observable relationship types across household networks (Panel A) and for network sizes (Panel B), respectively. The values are based on the investment income indicator due to data limitations.

	Count	Pct.	Cum. Pct.
Panel A: Relationship types			
Natural parent/child	1417	34.69	34.69
Partner/cohabitee	565	13.83	48.52
Non-relative	543	13.29	61.81
Natural brother/sister	493	12.07	73.88
Husband/wife	428	10.48	84.36
Missing	199	4.87	89.23
Step-parent/step-child	145	3.55	92.78
Landlord/landlady/tenant	60	1.47	94.25
Parent-in-law/son-in-law/sister-in-law	52	1.27	95.52
Brother-in-law/sister-in-law	49	1.20	96.72
Grandmother/grandfather and grandchild	41	1.00	97.72
Aunt/uncle and niece/nephew	24	0.59	98.31
Adoptive parent/child	19	0.47	98.78
Half-brother/half-sister	18	0.44	99.22
Cousin	16	0.39	99.61
Foster parent/child	6	0.15	99.76
Foster brother/sister	5	0.12	99.88
Other relative	5	0.12	100.00
Panel B: Network sizes			
3	174	23.87	23.87
4	270	37.04	60.91
5	184	25.24	86.15
6	56	7.68	93.83
7	19	2.61	96.43
8	16	2.19	98.63
9	6	0.82	99.45
10	1	0.14	99.59
11	2	0.27	99.86
15	1	0.14	100.00

Table 3.2

Network formation summary statistics across years

This table shows the number of newly formed household networks, the number of respondents therein and the levels of financial participation across these networks by year. The available data for the investment income variable covers the years 2011 to 2016, while for the interest + dividends variable is available from 2001 to 2016. The exact definitions of these variables are given in the Appendix.

	(a) Investment income			(b) Interest + dividends		
	Networks	Respondents		Networks	Respondents	
Year	Count	Count	Mean	Count	Count	Mean
2001	-	-	-	112	475	0.23
2002	-	-	-	72	315	0.17
2003	-	-	-	79	334	0.22
2004	-	-	-	75	327	0.15
2005	-	-	-	68	273	0.15
2006	-	-	-	85	369	0.19
2007	-	-	-	75	318	0.14
2008	-	-	-	68	289	0.18
2009	-	-	-	54	242	0.14
2011	175	820	0.27	130	593	0.19
2012	145	619	0.21	97	394	0.15
2013	123	547	0.29	86	360	0.25
2014	115	495	0.23	77	322	0.16
2015	95	410	0.29	64	261	0.16
2016	76	329	0.28	51	207	0.24

Table 3.3

Individual and household-level summary statistics

This table reports the sample summary statistics. Panel (a) shows the statistics for the sample derived from the investment income indicator, while Panel (b) shows those derived from the interest + dividend indicator. Variable definitions are reported in the Appendix.

	(a) Investment income				(b) Interest + dividends			
	Mean	SD	Min.	Max.	Mean	SD	Min.	Max.
Investment income	0.26	0.44	0	1	-	-	-	-
Interest + dividends	-	-	-	-	0.18	0.39	0	1
Age	34.07	15.14	16	94	32.93	14.51	15	96
A-level	0.28	0.45	0	1	0.29	0.45	0	1
Degree	0.25	0.43	0	1	0.17	0.38	0	1
GCSE	0.23	0.42	0	1	0.25	0.43	0	1
No qual	0.06	0.24	0	1	0.10	0.31	0	1
Other qual	0.08	0.26	0	1	0.10	0.30	0	1
Other higher qual	0.10	0.31	0	1	0.09	0.28	0	1
Labour income (1000s)	1.15	1.37	0	15	1.00	1.10	0	15
Employment	0.66	0.47	0	1	0.65	0.48	0	1
Family or home	0.05	0.22	0	1	0.06	0.23	0	1
Other occupation	0.23	0.42	0	1	0.23	0.42	0	1
Retired	0.04	0.19	0	1	0.03	0.17	0	1
Sick or disabled	0.02	0.15	0	1	0.04	0.18	0	1
Children	0.23	0.58	0	5	0.29	0.65	0	5
Female	0.51	0.50	0	1	0.52	0.50	0	1
Male	0.49	0.50	0	1	0.48	0.50	0	1
Year	22.01	1.67	20	25	17.15	4.67	10	25
East Midlands	0.11	0.09	0	1	0.09	0.08	0	1
East of England	0.09	0.08	0	1	0.07	0.07	0	1
London	0.07	0.07	0	1	0.06	0.05	0	1
North East	0.05	0.05	0	1	0.04	0.04	0	1
North West	0.10	0.09	0	1	0.09	0.08	0	1
Northern Ireland	0.03	0.03	0	1	0.06	0.06	0	1
Scotland	0.05	0.05	0	1	0.12	0.10	0	1
South East	0.12	0.11	0	1	0.11	0.10	0	1
South West	0.09	0.08	0	1	0.07	0.07	0	1
Wales	0.07	0.06	0	1	0.12	0.11	0	1
West Midlands	0.10	0.09	0	1	0.08	0.07	0	1
Yorkshire	0.13	0.11	0	1	0.08	0.08	0	1

Table 3.4

Estimation results for investment income models

This table reports the estimation results for the outcome equation (Panel A) and selection equation (Panel B) for the model in which the investment income indicator is the dependent variable. Definitions of all the variables are reported in Appendix A. The point estimates denote the mean of the respective marginal posterior distribution, while the corresponding standard deviations are reported in brackets. The stars ***, ** and * denote level of significance at 1, 5 and 10 percent, respectively.

	Panel A: Outcome equation			
	Zero latent variables ($D = 0$)		Two latent variables ($D = 2$)	
	Indivual	Peers	Individual	Peers
	(1)	(2)	(3)	(4)
Peer effect (λ)	0.058*** (0.019)		0.055*** (0.02)	
Age	0.717*** (0.08)	0.056 (0.19)	0.717*** (0.08)	0.055 (0.19)
No. of children	-0.308*** (0.08)	-0.167 (0.11)	-0.310*** (0.09)	-0.168 (0.11)
Labour income	0.487*** (0.09)	0.093 (0.11)	0.490*** (0.09)	0.097 (0.11)
Employed	0.041 (0.19)	-0.153 (0.25)	0.040 (0.19)	-0.153 (0.25)
Retired	0.500** (0.25)	0.021 (0.11)	0.505** (0.25)	0.024 (0.11)
Sick or disabled	-0.363 (0.31)	-0.157 (0.12)	-0.368 (0.32)	-0.156 (0.12)
Other job status	0.039 (0.20)	-0.186 (0.20)	0.041 (0.20)	-0.186 (0.21)
Degree	0.664*** (0.17)	0.400** (0.17)	0.661*** (0.17)	0.400** (0.17)
GCSE	0.285* (0.16)	0.101 (0.16)	0.282* (0.16)	0.099 (0.16)
A-level	0.437*** (0.16)	0.397** (0.17)	0.438*** (0.16)	0.398** (0.18)
Higher degree	0.639*** (0.18)	0.214* (0.12)	0.635*** (0.18)	0.210* (0.12)
Other education	0.187 (0.18)	0.090 (0.11)	0.185 (0.18)	0.088 (0.11)
Male	0.114 (0.07)	-0.099 (0.11)	0.114 (0.07)	-0.100 (0.11)
Selection equation	No		Yes	
Year effects	Yes		Yes	
Region effects	Yes		Yes	
Group effects	Yes		Yes	
Observations	3220		3220	
Networks	729		729	

(Continued)

	Panel B: Selection equation	
	Zero latent variables ($D = 0$)	Two latent variables ($D = 2$)
	(1)	(2)
Constant	-	10.324***
		(1.31)
Difference in age	-	-2.006***
		(0.25)
Difference in no. of children	-	-1.933***
		(0.24)
Difference in personal income	-	-0.736***
		(0.20)
Same occupational status	-	0.089
		(0.22)
Same education	-	0.547**
		(0.22)
Same gender	-	-1.003***
		(0.18)
δ_1	-	-6.461***
		(0.51)
δ_2	-	-0.312
		(0.83)
Year effects	-	Yes
Observations	-	6145

Table 3.5

Estimation results for interest + dividends models

This table reports the estimation results for the outcome equation (Panel A) and selection equation (Panel B) for the model in which the interest + dividend income indicator is the dependent variable. Definitions of all the variables are reported in Appendix A. The point estimates denote the mean of the respective marginal posterior distribution, while the corresponding standard deviations are reported in brackets. The stars ***, ** and * denote level of significance at 1, 5 and 10 percent, respectively.

	Panel A: Outcome equation			
	Zero latent variables ($D = 0$)		Two latent variables ($D = 2$)	
	Indivual	Peers	Individual	Peers
	(1)	(2)	(3)	(4)
Peer effect (λ)	0.085*** (0.01)		0.066** (0.02)	
Age	0.661*** (0.07)	0.113 (0.16)	0.669*** (0.07)	0.150 (0.17)
No. of children	-0.261*** (0.08)	-0.216** (0.10)	-0.276*** (0.08)	-0.232** (0.10)
Labour income	0.288*** (0.07)	0.106 (0.11)	0.297*** (0.08)	0.137 (0.11)
Employed	0.087 (0.17)	-0.125 (0.21)	0.073 (0.17)	-0.164 (0.23)
Retired	0.159 (0.23)	0.030 (0.09)	0.150 (0.24)	0.023 (0.10)
Sick or disabled	-0.446* (0.25)	-0.086 (0.10)	-0.476* (0.26)	-0.094 (0.11)
Other job status	0.151 (0.18)	-0.032 (0.18)	0.140 (0.18)	-0.067 (0.20)
Degree	0.963*** (0.14)	0.154 (0.13)	0.971*** (0.14)	0.152 (0.13)
GCSE	0.403*** (0.13)	0.053 (0.13)	0.408*** (0.13)	0.040 (0.13)
A-level	0.771*** (0.13)	0.157 (0.15)	0.780*** (0.13)	0.148 (0.16)
Higher degree	0.855*** (0.15)	0.050 (0.10)	0.876*** (0.15)	0.061 (0.10)
Other education	0.341** (0.14)	-0.054 (0.11)	0.340** (0.15)	-0.064 (0.11)
Male	0.136** (0.06)	0.051 (0.10)	0.136** (0.07)	0.057 (0.10)
Selection equation	No		Yes	
Year effects	Yes		Yes	
Region effects	Yes		Yes	
Group effects	Yes		Yes	
Observations	5079		5079	
Networks	1193		1193	

(Continued)

	Panel B: Selection equation	
	Zero latent variables ($D = 0$)	Two latent variables ($D = 2$)
	(1)	(2)
Difference in age	-	-2.240*** (0.24)
Difference in no. of children	-	-1.906*** (0.22)
Difference in personal income	-	-0.636*** (0.18)
Same occupational status	-	0.525*** (0.21)
Same education	-	0.615*** (0.21)
Same gender	-	-1.279*** (0.18)
δ_1	-	-6.929*** (0.50)
δ_2	-	-2.552*** (0.41)
Constant	-	13.568*** (1.07)
Year effects	-	Yes
Observations	-	9256

3.A Variable definitions

Variable	Definition
Investment income	Indicator variable: value one if the respondent receives investment income, and zero otherwise. Investment income includes income from private pensions/annuities, rent from non-family boarders or lodgers, rent from any other property, excluding tax plus income from savings and investment.
Interest plus dividends	Indicator variable: value one if the respondent receives interest or dividend income, and zero otherwise.
Age	Age of the respondent in years.
A-level	Indicator variable: value one if the highest educational attainment of the respondent are A-levels, and zero otherwise.
Degree	Indicator variable: value one if the highest educational attainment of the respondent is a university degree, and zero otherwise.
GCSE	Indicator variable: value one if the highest educational attainment of the respondent is the GCSE, and zero otherwise.
No qual	Indicator variable: value one if the respondent holds no educational qualifications, and zero otherwise.
Other qual	Indicator variable: value one if the respondent holds other educational qualifications, and zero otherwise.
Other higher qual	Indicator variable: value one if the respondent holds other higher educational qualifications, and zero otherwise.
Labour income	Gross monthly labour income of the respondent.
Employment	Indicator variable: value one if the respondent is an employee or self-employed, working full-time or part-time, and zero otherwise.
Family or home	Indicator variable: value one if the respondent looks after his/her family or home as a main occupation, and zero otherwise.
Other occupation	Indicator variable: value one if the respondent is not in employment, retired, sick or disabled or looks after family/home, and zero otherwise.
Retired	Indicator variable: value one if the respondent looks after his/her family or home as a main occupation, and zero otherwise.
Sick or disabled	Indicator variable: value one if the respondent is sick or disabled, and zero otherwise.
Children	Number of children of the respondent.
Female	Indicator variable: value one if the respondent is female, and zero otherwise.
Male	Indicator variable: value one if the respondent is male, and zero otherwise.
East Midlands	Indicator variable: value one if the respondent is domiciled in the Governmental Office Region East Midlands, and zero otherwise.
East of England	Indicator variable: value one if the respondent is domiciled in Governmental Office Region East of England, and zero otherwise.
London	Indicator variable: value one if the respondent is domiciled in Governmental Office Region London, and zero otherwise.

(Continued)

Variable	Definition
North East	Indicator variable: value one if the respondent is domiciled in Governmental Office Region North East, and zero otherwise.
North West	Indicator variable: value one if the respondent is domiciled in Governmental Office Region North West, and zero otherwise.
Northern Ireland	Indicator variable: value one if the respondent is domiciled in Governmental Office Region Northern Ireland, and zero otherwise.
Scotland	Indicator variable: value one if the respondent is domiciled in Governmental Office Region Scotland, and zero otherwise.
South East	Indicator variable: value one if the respondent is domiciled in Governmental Office Region South East, and zero otherwise.
South West	Indicator variable: value one if the respondent is domiciled in Governmental Office Region South West, and zero otherwise.
Wales	Indicator variable: value one if the respondent is domiciled in Governmental Office Region Wales, and zero otherwise.
West Midlands	Indicator variable: value one if the respondent is domiciled in Governmental Office Region West Midlands, and zero otherwise.
Yorkshire and the Humber	Indicator variable: value one if the respondent is domiciled in Governmental Office Region Yorkshire and the Humber, and zero otherwise.

Conclusion

This thesis set out to empirically analyse the relationship between financial literacy and subjective well-being (first essay), the relation between social norm enforcement and mental health (second essay) and the endogenous choice of peers with respect to investment participation (third essay). A multitude of observations emerges from these analyses.

In the first essay, we began by establishing that financial distress is associated with an approximately 13% drop in aggregate subjective well-being against predicted population-wide levels. Crucially, the analysis revealed that the relationship between financial distress and subjective well-being intensifies those individuals with high financial knowledge levels. For these individuals, financial distress is associated with a drop of 18% in aggregate subjective well-being, against relative to predicted overall levels. The stability of the relationship was confirmed for the individual subjective well-being dimensions.

Speaking to the driver of the observed relationship, the analysis revealed that when high financial knowledge individuals report that they fail to follow through with their own financial goals, the relationship intensifies disproportionately intensifies, yielding an estimated drop of 35% in aggregate subjective well-being. As we controlled for a wide range of individual- and household-level characteristics, this association suggests that high financial knowledge individuals have higher financial aspirations and again, the associations were confirmed for the individual subjective well-being dimensions.

In a set of additional analyses, we confirmed the conclusions drawn in an instrumental variable approach, and also found that relationships are stable when only the sample of financially distressed individuals is considered. These individuals are more likely to engage with financial education programmes aimed at individuals in difficult financial situations, which indicates that financial educators might find value in considering what effects erosions in subjective well-being might have on the extent of participants motivations to improve their circumstances.

In the second essay, we found that high welfarism, defined as favourable attitudes towards the welfare state and benefit recipients, is associated with an increased prevalence of mental health problem of 13 percentage points, a 39% increase against predicted population-wide levels. We formulated the hypothesis that expressing high welfarism constitutes a deviation of work norms, resulting in sanctioning and associated decreases in mental health.

We tested the notion by investigating the changes to the magnitudes of the relationship under Conservative governments, known to be tougher on welfare. The analysis revealed an intensifying of the relationship relative to Labour governments. Further, we found that the relationship is stronger for employed individuals, which exhibit high conformity with work norms, relative to those not in employment, with low conformity, suggesting that in-group social norms deviations are sanctioned more severely.

Additional analysis revealed that the observed relationships are stronger for females than for males. Further, we found that the combination of high welfarism and mental health problems is associated with more favourable attitudes towards voting, suggesting increased motivations to induce social change.

In the third essay, we proposed a peer effects model for binary outcome variables that explicitly models the empirical regularity that relationship formations occur on the basis of individuals similarities with one another. We use household networks that we obtain from the UK's major longitudinal surveys, Understanding Society and the British Household Panel Surveys (BHPS). Estimation took place in a Bayesian econometric framework using Hamiltonian Monte Carlo sampling methods.

We estimated the model on data for two outcomes variables: a general investment income indicator variable and an interest + dividend income indicator variable. We fit models both with exogenous and endogenous relationship formations. The resulting model comparison revealed that endogenous peer choice accounts for approximately 25% of the peer effect in investment participation. We further investigated the economic implications of low participation rates in the UK population in a simulation exercise, restricting the peer effect to zero. We produced counter-factual individual investment participation outcomes. The exercise suggested that participation rates would be 7% to 10% higher in the absence of peer effects.

A limitation that all three essays share as a result of data limitations is that identification of the results is derived from the cross-section of respondents in the samples. While all empirical analyses carefully control for factors that can constitute confounding influences to the observed relationships and employ measures to overcome causal identification challenges, future work building on these essays could exploit time-ordered data to examine changes that play out over time, if they becomes available, or test for the presence of causal effects bespoke randomised controls, where appropriate.

Overall, this thesis highlights three implications for the consideration of policymakers. First, policymakers should consider that financial aspirations can play an important role in the success of individuals enrolling in public financial education programmes. The higher financial aspirations, the more strongly financial distress can be experienced; and the resulting possible deterioration in subjective well-being may put educational success on the road to financial recovery at risk. Second, our results suggest the need for public discourse less loaded with stigmatization of the welfare state and benefit recipients due to its arguably identity-shaping effects and consequent social sanctioning. As the sanctioning of perceived deviations from work norms, through expressions of positive welfare attitudes, can put mental health at risk, changes to political rhetoric may contribute to population-wide improvements in mental health outcomes. Third, our results suggest a role for policy interventions that promote the known benefits of investment participation for the long-term economic well-being of household in the cross-section of society that holds no investments. Such policy efforts may facilitate the breaking up of possible self-sustaining tendencies of not participating in financial markets in various subgroups of society.

References

Chapter 1

Balloch, A., Nicolae, A. & Philip, D. (2015). Stock market literacy, trust, and participation. *Review of Finance*, *19*(5), 1925–1963.

Baum, C. F. & Lewbel, A. (2019). Advice on using heteroskedasticity-based identification. *The Stata Journal*, *19*(4), 757–767.

Brown, C. & Robinson, L. (2016). Breaking the cycle: From poverty to financial security for all. *Oakland, CA: PolicyLink*.

Caliendo, M. & Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, *22*(1), 31–72.

Carpena, F., Cole, S., Shapiro, J. & Zia, B. (2019). The ABCs of financial education: Experimental evidence on attitudes, behavior, and cognitive biases. *Management Science*, *65*(1), 346–369.

Chetty, R., Grusky, D., Hell, M., Hendren, N., Manduca, R. & Narang, J. (2017). The fading american dream: Trends in absolute income mobility since 1940. *Science*, *356*(6336), 398–406.

Consumer Financial Protection Bureau. (2017). Financial well-being in America. https://www.consumerfinance.gov/data-research/research-reports/financial-well-being-america/.

Deaton, A. (2008). Income, health, and well-being around the world: Evidence from the Gallup World Poll. *Journal of Economic Perspectives*, *22*(2), 53–72.

Dehejia, R. H. & Wahba, S. (2002). Propensity score-matching methods for non-experimental causal studies. *The Review of Economics and Statistics*, *84*(1), 151–161.

Diener, E., Oishi, S. & Tay, L. (2018). Advances in subjective well-being research. *Nature Human Behaviour*, *2*(4), 253.

Engels, C., Kumar, K. & Philip, D. (2020). Financial literacy and fraud detection. *The European Journal of Finance*, *26*(4-5), 420–442.

Grace, B. Y., Lee, D.-J., Sirgy, M. J. & Bosnjak, M. (2019). Household income, satisfaction with standard of living, and subjective well-being. the moderating role of happiness materialism. *Journal of Happiness Studies*. https://doi.org/10.1007/s10902-019-00202-x.

Graham, C. & Pettinato, S. (2002). Frustrated achievers: Winners, losers and subjective well-being in new market economies. *Journal of Development Studies*, *38*(4), 100–140.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 1029–1054.

Helliwell, J. F., Layard, R. & Sachs, J. (Eds.). (2019). World happiness report 2019. https://worldhappiness.report/ed/2019/.

Kahneman, D. & Deaton, A. (2010). High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, *107*(38), 16489–16493.

Karpman, M., Zuckerman, S. & Gonzalez, D. (2018). *Material Hardship among nonelderly adults and their families in 2017*. Urban Institute. https://www.urban.org/research/publication/material-hardship-among-nonelderly-adults-and-their-families-2017.

Klapper, L. & Panos, G. A. (2011). Financial literacy and retirement planning: The Russian case. *Journal of Pension Economics & Finance*, *10*(4), 599–618.

Klapper, L. & Panos, G. A. (2013). Financial literacy and its consequences: Evidence from Russia during the financial crisis. *Journal of Banking and Finance*, *37*(10), 3904–3923.

Knoll, M. A. & Houts, C. R. (2012). The financial knowledge scale: An application of item response theory to the assessment of financial literacy. *Journal of Consumer Affairs*, *46*(3), 381–410.

Lewbel, A. (2012). Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *Journal of Business & Economic Statistics*, *30*(1), 67–80.

Lindqvist, E., Östling, R. & Cesarini, D. (2020). Long-run effects of lottery wealth on psychological well-being. *The Review of Economic Studies*.

Lusardi, A., Michaud, P.-C. & Mitchell, O. S. (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy*, *125*(2), 431–477.

Lusardi, A. & Mitchell, O. S. (2007a). Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics*, *54*(1), 205–224.

Lusardi, A. & Mitchell, O. S. (2007b). Financial literacy and retirement preparedness: Evidence and implications for financial education. *Business Economics*, *42*(1), 35–44.

Lusardi, A. & Mitchell, O. S. (2011). *Financial literacy and planning: Implications for retirement wellbeing* (NBER Working Paper Series No. 17078). https://www.nber.org/papers/w17078.

McKernan, S.-M., Ratcliffe, C. & Iceland, J. (2018). *Policy Efforts to Reduce Material Hardship for Low-Income Families*. Urban Institute. https://www.urban.org/research/publication/policy-efforts-reduce-material-hardship-low-income-families.

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

Reeskens, T. & Vandecasteele, L. (2017). Economic hardship and well-being: Examining the relative role of individual resources and welfare state effort in resilience against economic hardship. *Journal of Happiness Studies*, *18*(1), 41–62.

Rodnyansky, A. & Darmouni, O. M. (2017). The effects of quantitative easing on bank lending behavior. *The Review of Financial Studies*, *30*(11), 3858–3887.

Roodman, D. (2011). Fitting fully observed recursive mixed-process models with cmp. *The Stata Journal*, *11*(2), 159–206.

Sacks, D. W., Stevenson, B. & Wolfers, J. (2012). The new stylized facts about income and subjective well-being. *Emotion*, *12*(6), 1181.

Schaffer, M. & Baum, C. (2012). IVREG2H: Stata module to perform instrumental variables estimation using heteroskedasticity-based instruments. Statistical Software Components S457555, Boston College Department of Economics, revised 26 Jun 2020.

Stock, J. H. & Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. Andrews & J. Stock (Eds.), *Identification and inference for econometric models: Essays in honor of Thomas J. Rothenberg* (pp. 17–80). Cambridge: Cambridge University Press.

Van Rooij, M., Lusardi, A. & Alessie, R. (2011). Financial literacy and stock market participation. *Journal of Financial Economics*, *101*(2), 449–472.

Yoong, J. (2011). Financial illiteracy and stock market participation: Evidence from the RAND American Life Panel. In O. S. Mitchell & A. Lusardi (Eds.), *Financial literacy: Implications for retirement security and the financial marketplace* (pp. 76–97). Oxford: Oxford University Press.

Chapter 2

Afifi, M. (2007). Gender differences in mental health. *Singapore Medical Journal*, *48*(5), 385.

Akerlof, G. A. & Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, *115*(3), 715–753.

Alesina, A. & Giuliano, P. (2015). Culture and institutions. *Journal of Economic Literature*, *53*(4), 898–944.

Aneshensel, C. S., Estrada, A. L., Hansell, M. J. & Clark, V. A. (1987). Social psychological aspects of reporting behavior: Lifetime depressive episode reports. *Journal of Health and Social Behavior*, 232–246.

Angst, J. & Dobler-Mikola, A. (1985). The Zurich study. V. Anxiety and phobia in young adults. *European Archives of Psychiatry and Neurological Sciences*, *235*(3), 171–178.

Balloch, A., Engels, C. & Philip, D. (2020). When it rains it drains: Psychological distress and household net worth. SSRN. https://ssrn.com/abstract=3521323.

Bénabou, R. & Tirole, J. (2011). Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics*, *126*(2), 805–855.

Benjamin, D. J., Choi, J. J. & Strickland, A. J. (2010). Social identity and preferences. *American Economic Review*, *100*(4), 1913–28.

Berkowitz, M. K. & Qiu, J. (2006). A further look at household portfolio choice and health status. *Journal of Banking & Finance*, *30*(4), 1201–1217.

Bernhard, H., Fehr, E. & Fischbacher, U. (2006). Group affiliation and altruistic norm enforcement. *American Economic Review*, *96*(2), 217–221.

Bertrand, M., Kamenica, E. & Pan, J. (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics*, *130*(2), 571–614.

Bogan, V. L. & Fertig, A. R. (2013). Portfolio choice and mental health. *Review of Finance*, *17*(3), 955–992.

Bogan, V. L. & Fertig, A. R. (2018). Mental health and retirement savings: Confounding issues with compounding interest. *Health Economics*, *27*(2), 404–425.

Bogan, V. L., Fertig, A. R. & Just, D. R. (2019). Self-employment and mental health. SSRN. http://dx.doi.org/10.2139/ssrn.2259765.

Bowles, S. (1998). Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature*, *36*(1), 75–111.

Bowles, S. & Polania-Reyes, S. (2012). Economic incentives and social preferences: Substitutes or complements? *Journal of Economic Literature*, *50*(2), 368–425.

Bridges, S. & Disney, R. (2010). Debt and depression. *Journal of Health Economics*, *29*(3), 388–403.

Bruce, S. E., Yonkers, K. A., Otto, M. W., Eisen, J. L., Weisberg, R. B., Pagano, M., . . . Keller, M. B. (2005). Influence of psychiatric comorbidity on recovery and recurrence in generalized anxiety disorder, social phobia, and panic disorder: A 12-year prospective study. *American Journal of Psychiatry*, *162*(6), 1179–1187.

Bursztyn, L., Ferman, B., Fiorin, S., Kanz, M. & Rao, G. (2017). Status goods: Experimental evidence from platinum credit cards. *The Quarterly Journal of Economics*.

Callen, M., Isaqzadeh, M., Long, J. D. & Sprenger, C. (2014). Violence and risk preference: Experimental evidence from Afghanistan. *American Economic Review*, *104*(1), 123–48.

Currie, J., Duque, V. & Garfinkel, I. (2015). The great recession and mothers' health. *The Economic Journal*, *125*(588), 311–346.

Curtice, J., Clery, E., Perry, J., Phillips, M. & Rahim, N. (Eds.). (2019). *British Social Attitudes: The 36th Report. Technical details*. London: The National Centre for Social Research.

de Vries, R. & Reeves, A. (2017). Why do people care more about benefit scroungers than billions lost to the rich? The Guardian. Last accessed 25 Oct 2020. https://www.theguardian.com/commentisfree/2017/nov/15/benefit-scroungers-billions-rich-paradise-papers-tax-avoidance.

Department for Work and Pensions. (2020). Universal Credit and you. Last accessed 25 Oct 2020. https://www.gov.uk/government/publications/universal-credit-and-you/draft-uc-and-you.

Field, F. (2002). Making welfare work: The politics of reform. *Scottish Journal of Political Economy*, *13*, 91–103.

Gathergood, J. (2012). Debt and depression: Causal links and social norm effects. *The Economic Journal*, *122*(563), 1094–1114.

Geiger, B. B. (2017). Benefit 'myths'? The accuracy and inaccuracy of public beliefs about the benefits system. *Social Policy & Administration*.

Giuliano, P. & Spilimbergo, A. (2014). Growing up in a recession. *Review of Economic Studies*, *81*(2), 787–817.

Goette, L., Huffman, D. & Meier, S. (2006). The impact of group membership on cooperation and norm enforcement: Evidence using random assignment to real social groups. *American Economic Review*, *96*(2), 212–216.

Grip, A. d., Lindeboom, M. & Montizaan, R. (2011). Shattered dreams: The effects of changing the pension system late in the game. *The Economic Journal*, *122*(559), 1–25.

Karlan, D. & Zinman, J. (2010). Expanding credit access: Using randomized supply decisions to estimate the impacts. *The Review of Financial Studies*, *23*(1), 433–464.

Lindbeck, A., Nyberg, S. & Weibull, J. W. (1999). Social norms and economic incentives in the welfare state. *The Quarterly Journal of Economics*, *114*(1), 1–35.

Lindbeck, A., Nyberg, S. & Weibull, J. W. (2003). Social norms and welfare state dynamics. *Journal of the European Economic Association*, *1*(2-3), 533–542.

Lott, B. (2002). Cognitive and behavioral distancing from the poor. *American Psychologist*, *57*(2), 100–110.

Malmendier, U. & Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, *126*(1), 373–416.

McLean, C. P., Asnaani, A., Litz, B. T. & Hofmann, S. G. (2011). Gender differences in anxiety disorders: Prevalence, course of illness, comorbidity and burden of illness. *Journal of Psychiatric Research*, *45*(8), 1027–1035.

Moffitt, R. (1983). An economic model of welfare stigma. *American Economic Review*, *73*(5), 1023–1035.

Mourifié, I. & Méango, R. (2014). A note on the identification in two equations probit model with dummy endogenous regressor. *Economics Letters*, *125*(3), 360–363.

Niskanen, W. A. (1996). Welfare and the culture of poverty. *The Cato Journal*, *16*, 1–15.

Nolen-Hoeksema, S. & Girgus, J. S. (1994). The emergence of gender differences in depression during adolescence. *Psychological Bulletin*, *115*(3), 424.

OECD/EU. (2018). *Health at a glance: Europe 2018: State of health in the eu cycle*. Paris: OECD Publishing. Retrieved from https://doi.org/10.1787/health_glance_eur-2018-en

Pescosolido, B. A. & Martin, J. K. (2015). The stigma complex. *Annual Review of Sociology*, *41*, 87–116.

Piccinelli, M. & Wilkinson, G. (2000). Gender differences in depression: Critical review. *The British Journal of Psychiatry*, *177*(6), 486–492.

Posner, R. A. (1997). Social norms and the law: An economic approach. *American Economic Review*, *87*(2), 365–369.

Reutter, L. I., Stewart, M. J., Veenstra, G., Love, R., Raphael, D. & Makwarimba, E. (2009). "Who do they think we are, anyway?": Perceptions of and responses to poverty stigma. *Qualitative Health Research*, *19*(3), 297–311.

Roodman, D. (2011). Fitting fully observed recursive mixed-process models with cmp. *The Stata Journal*, *11*(2), 159–206.

Tanaka, T., Camerer, C. F. & Nguyen, Q. (2010). Risk and time preferences: Linking experimental and household survey data from Vietnam. *American Economic Review*, *100*(1), 557–71.

The Economist. (2019). British attitudes to welfare have undergone a quiet revolution. Last accessed 25 Oct 2020. https://www.economist.com/britain/2019/07/11/british-attitudes-to-welfare-have-undergone-a-quiet-revolution.

Wilde, J. (2000). Identification of multiple equation probit models with endogenous dummy regressors. *Economics Letters*, *69*(3), 309–312.

Chapter 3

Arrondel, L., Calvo Pardo, H. F., Giannitsarou, C. & Haliassos, M. (2020). *Informative social interactions*. SSRN. https://ssrn.com/abstract=3171564.

Attanasio, O., Meghir, C. & Mommaerts, C. (2015). *Insurance in extended family networks* (Cowles Foundation Discussion Paper No. 1996). https://ssrn.com/abstract=3206754.

Badev, A. (2013). *Discrete games in endogenous networks: Theory and policy* (PSC Working Paper Series No. 45). https://repository.upenn.edu/psc_working_papers/45.

Barnea, A., Cronqvist, H. & Siegel, S. (2010). Nature or nurture: What determines investor behavior? *Journal of Financial Economics*, *98*(3), 583–604.

Barth, D., Papageorge, N. W. & Thom, K. (2020). Genetic endowments and wealth inequality. *Journal of Political Economy*, *128*(4), 1474–1522.

Betancourt, M. (2017). *A conceptual introduction to Hamiltonian Monte Carlo*. arXiv. https://arxiv.org/abs/1701.02434.

Betancourt, M. & Girolami, M. (2015). Hamiltonian Monte Carlo for hierarchical models. *Current trends in Bayesian methodology with applications*, *79*, 30.

Blume, L. E., Brock, W. A., Durlauf, S. N. & Jayaraman, R. (2015). Linear social interactions models. *Journal of Political Economy*, *123*(2), 444–496.

Bramoullé, Y., Djebbari, H. & Fortin, B. (2009). Identification of peer effects through social networks. *Journal of Econometrics*, *150*(1), 41–55.

Cesarini, D., Johannesson, M., Lichtenstein, P., Sandewall, Ö. & Wallace, B. (2010). Genetic variation in financial decision-making. *The Journal of Finance*, *65*(5), 1725–1754.

Chib, S. & Greenberg, E. (1998). Analysis of multivariate probit models. *Biometrika*, *85*(2), 347–361.

Cronqvist, H. & Siegel, S. (2014). The genetics of investment biases. *Journal of Financial Economics*, *113*(2), 215–234.

Das, S., Kuhnen, C. M. & Nagel, S. (2020). Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies*, *33*(1), 395–432.

de Paula, A., Rasul, I. & Souza, P. (2020). *Identifying network ties from panel data: Theory and an application to tax competition*. arXiv. https://arxiv.org/abs/1910.07452.

Duflo, E. & Saez, E. (2003). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, *118*(3), 815–842.

Epstein, B. & Shapiro, A. F. (2020). Increasing domestic financial participation: Implications for business cycles and labor markets. *Review of Economic Dynamics*.

Fafchamps, M. & Gubert, F. (2007a). Risk sharing and network formation. *American Economic Review*, *97*(2), 75–79.

Fafchamps, M. & Gubert, F. (2007b). The formation of risk sharing networks. *Journal of Development Economics*, *83*(2), 326–350.

Favilukis, J. (2013). Inequality, stock market participation, and the equity premium. *Journal of Financial Economics*, *107*(3), 740–759.

Fetzer, T., Souza, P., Eynde, O. & Wright, A. (2020). *Security Transitions*. HAL. https://halshs.archives-ouvertes.fr/halshs-02518234.

Gelman, A., Jakulin, A., Pittau, M. G., Su, Y.-S. et al. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, *2*(4), 1360–1383.

Gelman, A. & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences (with discussion). *Statistical Science*, *7*(4), 457–472.

Ghosh, J., Li, Y., Mitra, R. et al. (2018). On the use of Cauchy prior distributions for Bayesian logistic regression. *Bayesian Analysis*, *13*(2), 359–383.

Goldsmith-Pinkham, P. & Imbens, G. W. (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics*, *31*(3), 253–264.

Haliassos, M., Jansson, T. & Karabulut, Y. (2019). Financial literacy externalities. SSRN. https://ssrn.com/abstract=2870796.

Hsieh, C.-S. & Lee, L. F. (2016). A social interactions model with endogenous friendship formation and selectivity. *Journal of Applied Econometrics*, *31*(2), 301–319.

Hsieh, C.-S., Lee, L.-F. & Boucher, V. (2019). Specification and estimation of network formation and network interaction models with the exponential probability distribution. SSRN. https://ssrn.com/abstract=2967867.

Hsieh, C.-S. & Van Kippersluis, H. (2018). Smoking initiation: Peers and personality. *Quantitative Economics*, *9*(2), 825–863.

Johnsson, I. & Moon, H. R. (2019). Estimation of peer effects in endogenous social networks: Control function approach. *The Review of Economics and Statistics*. https://doi.org/10.1162/rest_a_00870.

Kaustia, M. & Knüpfer, S. (2012). Peer performance and stock market entry. *Journal of Financial Economics*, *104*(2), 321–338.

Kelejian, H. H. & Prucha, I. R. (2010). Specification and estimation of spatial autoregressive models with autoregressive and heteroskedastic disturbances. *Journal of Econometrics*, *157*(1), 53–67.

Lee. (2007). Identification and estimation of econometric models with group interactions, contextual factors and fixed effects. *Journal of Econometrics*, *140*(2), 333–374.

Lee, Li & Lin. (2014). Binary choice models with social network under heterogeneous rational expectations. *The Review of Economics and Statistics*, *96*(3), 402–417.

Lee, Liu & Lin. (2010). Specification and estimation of social interaction models with network structures. *The Econometrics Journal*, *13*(2), 145–176.

Lin, X. (2010). Identifying peer effects in student academic achievement by spatial autoregressive models with group unobservables. *Journal of Labor Economics*, *28*(4), 825–860.

Lin, Z., Tang, X. & Yu, N. (2020). Uncovering heterogeneous social effects in binary choices. *Journal of Econometrics*. https://doi.org/10.1016/j.jeconom.2020.08.005.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, *60*(3), 531–542.

Nakosteen, R. A., Westerlund, O. & Zimmer, M. A. (2004). Marital matching and earnings evidence from the unmarried population in Sweden. *Journal of Human Resources*, *39*(4), 1033–1044.

ONS. (2011). 2011 UK censuses. Retrieved October 24, 2020, from https://www.ons.gov.uk/census/2011census/2011ukcensuses

Stan Development Team. (2018). RStan: the R interface to Stan. R package version 2.17.3. http://mc-stan.org.

Van Rooij, M., Lusardi, A. & Alessie, R. (2011). Financial literacy and stock market participation. *Journal of Financial Economics*, *101*(2), 449–472.

Weinberg, B. A. (2007). *Social interactions with endogenous associations* (NBER Working Paper Series No. 13038). https://www.nber.org/papers/w13038.

Zhou, W. (2019). A network social interaction model with heterogeneous links. *Economics Letters*, *180*, 50–53.