

# Understanding Health, Labour, and Consumption Dynamics



Alexandra Elizabeth Brown

Lincoln College

University of Oxford

A thesis submitted for the degree of

*Doctor of Philosophy in Economics*

Trinity 2024

# Statement of Authorship

I certify that this thesis, submitted for the degree of Doctor of Philosophy in Economics at the University of Oxford, is my own work except where otherwise clearly stated.

- **Chapter 1: How should we model health as a dynamic process?** is solely my own work.
- **Chapter 2: Health-driven occupational changes** is co-authored with Nick Ridpath, a DPhil student in Economics at the University of Oxford. I was responsible for the data gathering, data analysis, literature review, and overall write up of the chapter. Nick was responsible for building a dynamic discrete choice model based on my data analysis. It was later decided to remove the structural modelling component of this piece, so his work became the theoretical framework, which I have since heavily revised. Nick is not also submitting this chapter as part of his DPhil.
- **Chapter 3: The consumption choices of ‘Generation Rent’** is solely my own work. An earlier version was submitted as part of the degree of MPhil in Economics at the University of Oxford but has since been heavily revised.

# Acknowledgements

*This thesis took me four years to write. During that period, I was supported by several wonderful friends, colleagues and communities in Oxford that I want to acknowledge and to whom I will be forever grateful.*

*First, I want to express a huge amount of gratitude to my supervisor for the past six(!) years, Professor Hamish Low. I have become a so much better economist under his supervision. Whether it was getting out the white board marker to patiently re-teach me some undergraduate economics concept I had long forgotten, inventing catch-phrases to make my presentations more exciting (“downward spirals of despair”), always having helpful suggestions whenever I was stuck no matter how far away my problem was from his research expertise, always somehow finding the time (with some frequent last-minute rescheduling) for his supervisees despite his horrifically busy schedule, or always being exceptionally encouraging and supportive; I am extremely thankful. I would also like to thank the Department of Economics for doctoral bursary funding and conference funding, and my parents for their support over the years from the other side of the world.*

*Residing with me in the DPhil space on the second floor of the Manor Road Building have been many smart, interesting and engaging individuals who have also made me a better economist. I want to particularly mention Stan and Nick who were stuck sitting next to me for years, and my former Reserve Bank of Australia colleagues; Pete and Martin. Special shout out to Juliette and JG for feeding me chocolate, choux buns, saucisson, and excellent gossip. I also want to acknowledge Merton College, which was such a wonderful community to be a member of and reside within, and probably was the key thing that kept me on track and working, especially during the first year of the DPhil. Special mention to the many wonderful people I worked with in the Welfare team at Merton: Simon, Jenny and the junior deans, and the ‘relentless dining’ fellows. I also want to thank Natalie and Frank for being so wonderfully supportive of my work as a lecturer at Lady Margaret Hall and ensuring my teaching workload was manageable so I could still do my degree.*

*Finally, George Orwell once observed that there is an ‘English characteristic which is so much a part of us that we barely notice it, and that is the addiction to hobbies and spare-time occupations.’ Moving to England in 2018 I found this to be true, and in particular I’m exceptionally grateful to the many friends I made in the sports clubs I joined - they also played an important role in getting me through my degree. I want to especially mention Heather, Vio, Kate S, Ali and Vania, alongside many other members of the Lincoln College Boat Club, Oxford University Triathlon Club, and LMH Women’s Lifting Club. I’m not sure whether church also counts as a hobby, but I’m also grateful to the communities at Merton College Chapel, St Mary Magdalen in Oxford, and the Balliol College Chapel Choir.*

*Thank you to you all.*

# Abstract

This thesis consists of three stand-alone papers that use British panel data and applied microeconomics techniques to further our understanding of health, labour and consumption dynamics. All papers are motivated by key public policy issues facing the United Kingdom at present. The first paper, *‘How should we model health as a dynamic process’*, reviews how health has been modelled in the literature as a statistical process. It identifies improvements to better capture persistence heterogeneity and fixed effects by using sophisticated non-linear panel data techniques and additional medical data. The second paper, *‘Health-driven occupational changes’*, shows that individuals are more likely to change occupation if their health worsens, and these occupation transitions are different to the occupation transitions of comparable individuals who remain healthy. The final paper, *‘The consumption choices of ‘Generation Rent’*, evaluates how the large increase in UK housing prices impacted the housing and non-housing budget shares of young renters by estimating housing price elasticities from Quadratic Almost Ideal Demand System estimation.

# CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>How should we model health as a dynamic process?</b>	<b>6</b>
2.1	Introduction . . . . .	7
2.2	Literature Review . . . . .	11
2.3	Data . . . . .	15
2.3.1	Health index construction . . . . .	15
2.3.2	Allostatic scores from biomarker data . . . . .	23
2.4	Modelling health as a dynamic process . . . . .	27
2.4.1	ARMA(p,q) baseline model . . . . .	27
2.4.2	Linear additive shock model . . . . .	35
2.5	Capturing more complex dynamics . . . . .	38
2.5.1	Recent health history . . . . .	38
2.5.2	Age and model stationarity . . . . .	43
2.5.3	Underlying health . . . . .	45
2.6	Non-linear health dynamics . . . . .	46
2.6.1	Non-linear persistence estimates of overall health . . . . .	47
2.6.2	Fixed effects . . . . .	51
2.6.3	Mental health persistence . . . . .	55
2.7	Conclusion . . . . .	58
<b>3</b>	<b>Health-driven occupational changes</b>	<b>60</b>
3.1	Introduction . . . . .	61

3.2	Literature Review . . . . .	64
3.3	Framework . . . . .	68
3.4	Data . . . . .	72
	3.4.1 Health data . . . . .	72
	3.4.2 Labour market transitions data . . . . .	77
3.5	Empirical Results . . . . .	82
	3.5.1 Occupation and employer transition probabilities . . . . .	82
	3.5.2 Changes to hours worked . . . . .	84
	3.5.3 Occupation changes . . . . .	88
	3.5.3.1 Empirical strategy . . . . .	88
	3.5.3.2 Key results . . . . .	90
	3.5.3.3 Heterogeneity analysis . . . . .	95
3.6	Conclusion . . . . .	97
<b>4</b>	<b>The consumption choices of ‘Generation Rent’</b>	<b>99</b>
4.1	Introduction . . . . .	100
4.2	Literature Review . . . . .	104
	4.2.1 Consumption and house price co-movement . . . . .	104
	4.2.2 Housing demand . . . . .	105
4.3	Data and Summary Statistics . . . . .	107
	4.3.1 Expenditure budget shares . . . . .	108
	4.3.1.1 Censored observations . . . . .	108
	4.3.1.2 Housing budget shares . . . . .	110
	4.3.1.3 Non-housing budget shares . . . . .	111
	4.3.2 Housing prices . . . . .	114
	4.3.3 House price and tenure changes . . . . .	116
4.4	Demand System Estimation . . . . .	118
	4.4.1 Demand system theory . . . . .	118
	4.4.2 Estimation details . . . . .	120
	4.4.3 QUAIDS coefficients & elasticity estimates . . . . .	122
	4.4.4 Robustness checks . . . . .	123
	4.4.5 Interpretation of elasticity estimates . . . . .	126
	4.4.5.1 Housing . . . . .	127

4.4.5.2	Food and fuel/light . . . . .	128
4.4.5.3	Other non-housing goods and services . . . . .	128
4.5	Welfare Analysis . . . . .	131
4.5.1	Estimation . . . . .	131
4.5.2	Compensating variation estimates . . . . .	133
4.6	Conclusion . . . . .	135
<b>5</b>	<b>Conclusion</b>	<b>137</b>
<b>A</b>		<b>140</b>
A.0.1	Attrition . . . . .	140
A.0.2	Regression output from health index construction . . . . .	141
A.0.3	Genetic data . . . . .	144
A.0.4	Replication exercise using non-detrended health data . . . . .	146
A.0.5	ARMA(p,q) groups from clustering analysis . . . . .	147
A.0.6	ARCH model estimates . . . . .	149
A.0.7	ARMA models with additional allostatic regressors . . . . .	151
A.0.8	Additional non-linear health persistence results . . . . .	152
A.0.9	Additional mental health persistence results . . . . .	155
<b>B</b>		<b>157</b>
B.0.1	PCA analysis for health category classification . . . . .	157
B.0.2	Additional summary statistics . . . . .	158
B.0.3	Additional analysis of hours changes . . . . .	160
B.0.4	Additional robustness checks . . . . .	162
B.0.5	Additional specifications . . . . .	163
B.0.6	Selection effects . . . . .	163
<b>C</b>		<b>167</b>
C.0.1	Additional summary statistics . . . . .	167
C.0.1.1	Price variation . . . . .	167
C.0.1.2	Trends in non-housing consumption & tenure . . . . .	167
C.0.1.3	Summary statistics table . . . . .	168
C.0.1.4	Budget shares . . . . .	170

C.0.2	Additional demand estimation derivations & results . . . . .	174
C.0.2.1	AIDS derivation . . . . .	174
C.0.2.2	Elasticity formula derivation . . . . .	175
C.0.2.3	Additional QUAIDS estimates . . . . .	176
C.0.2.4	EASI estimates . . . . .	177
C.0.2.5	Endogeneity . . . . .	178
C.0.3	Deriving Compensating Variation . . . . .	180

<b>Bibliography</b>		<b>181</b>
---------------------	--	------------

## **INTRODUCTION**

This thesis consists of three stand-alone papers that seek to further our understanding of health, labour and consumption dynamics using applied microeconomics techniques. Each paper makes up a chapter. My first chapter is about modelling the dynamics of health, as measured using health survey questions aggregated into a single index. My second chapter examines occupational mobility as an under-explored channel by which health impacts labour supply. In my third chapter, I estimate a demand system to assess how the expenditure of young renters was affected by the large and multi-decade increase in house and rental prices in the UK. My first two chapters were inspired by my work as a research assistant for a paper on the heterogeneous labour market impacts of covid-19 in the UK (Crossley, Fisher and Low, 2021). Using Understanding Society data, I noticed that those with pre-existing health conditions such as diabetes had much worse labour market outcomes in 2020–21 relative to healthy individuals in similar occupations and with otherwise similar observable traits. This led me to ask two sets of questions. First, how should we model the health of individuals over time in a way that is parsimonious but still captures the crucial statistical properties that impact economic decision making? And second, for those who were in poor health prior to 2020, how had their prior labour market history been affected by their poor health? Did their current occupation reflect adjustments they had made to manage their health condition? And did this make them more or less resilient to the current shock that they faced? The first two chapters are my

attempts to answer most of these questions. The third chapter is a heavily revised version of the thesis I submitted as part of my MPhil in Economics degree at the University of Oxford in 2020.

The first chapter, ‘How should we model health as a dynamic process?’, comprehensively reviews how health has been modelled in the literature as a statistical process, and identifies improvements. Health is an important determinant of an individual’s economic decision making, and accurately capturing key statistical features such as persistence and individual heterogeneity is important for answering several important and open questions in the literature. This chapter makes three main contributions to the literature. It provides an assessment of the different ways health dynamics have been modelled in the literature to date. Researchers have typically borrowed from the earnings dynamics literature and modelled health as a simple linear process, often discretised as a first-order Markov process, or the sum of a permanent and a transitory shock. To the author’s knowledge, there have been no prior attempts to systematically evaluate these modelling approaches and their underlying assumptions. I then suggest two improvements. I adapt Arellano, Blundell and Bonhomme (2017)’s quantile-based method from the earnings dynamics literature to produce non-linear persistence estimates that allow for a large amount of heterogeneity. This framework is able to capture several important features of health dynamics that other models cannot. These include accurately capturing that the persistence of the health process is higher among individuals in poor prior health, and that positive health shocks are less persistent than negative health shocks of comparable magnitudes. I also show that the persistence dynamics of mental health varies from overall health in meaningful ways. My third contribution is that I evaluate how we can best incorporate the increasing availability of medical data such as genetic data and biological marker ‘biomarker’ data from blood tests into modelling health dynamics, as these data have been shown to predict health outcomes of ostensibly healthy people. I show that using biomarker data can help us understand and model individual heterogeneity in health dynamics.

The second chapter, ‘Health-driven occupational changes’, evaluates the impact of worsening health on occupational mobility in the UK. It is a joint work with Nick Ridpath, a fellow DPhil in Economics student at the University of Oxford. While

there is a rich literature on the different ways that poor health impacts labour supply, particularly focussing on extensive margin adjustments and wages, we believe that we are the first to examine the impact of health shocks on the likelihood and nature of occupational changes. We show that individuals who suffer a health shock are 10-15 per cent more likely to change occupation or employer in the subsequent twelve months compared to employees who do not suffer a health shock. We also find that the occupations selected by those who have recently suffered a health shock differ from the occupations selected by comparable individuals who remain healthy. We model occupations as consisting of three tasks: cognitive, manual and interpersonal, that vary in intensity. We allow for health shocks to have different impacts on occupation choice depending on the type of health shock (physical, chronic/internal or mental), and occupation task requirements. The existence of alternate occupations that individuals can switch to functions as a form of partial insurance from the wage losses caused by poor health. We find that individuals who suffered a physical disability or mental health shock switched to occupations with lower cognitive intensity relative to the healthy. We also find that those who suffered a physical disability reduced the manual intensity of their occupation relative to the healthy. Less cognitively-intense jobs are typically jobs with lower overall task complexity, while less manual jobs can be more suitable for those with certain health conditions. Individuals who do not hold a degree and suffer a worsening of their mental health appear to be particularly vulnerable; we observe the largest declines in task intensity for this group relative to the healthy across multiple task domains. We find no relationship between health shocks and subsequent changes to occupation interpersonal intensity, a result which surprised us.

The final chapter, The consumption choices of ‘Generation Rent’, is motivated by observing in the data that the amount that young renters spend on rent relative to other consumption goods and services has sharply increased in the UK. The housing budget share of renters aged 18-35 increased by almost 15 percentage points between 1987 and 2018. This occurred alongside house prices increasing by more than twice the rate of overall inflation, rental prices increasing 1.5 times the rate of inflation, and a sharp fall in the homeownership rate of this age group. The relationship between young people’s consumption, housing prices and tenure choices is complex.

This chapter focusses on an aspect of this relationship, and evaluates how the large increase in housing prices impacted the housing and non-housing budget shares of young renters. I estimate housing own and cross-price elasticities of demand for six major expenditure categories by first estimating a Quadratic Almost Ideal Demand System (QUAIDS) with UK household expenditure survey data over 1987–2018. I find that the housing consumption of low-expenditure (a proxy for low-income) young renters is very responsive to house prices changes. However, the magnitude of this response decreased over the sample period, which coincided with large falls in social housing supply, limiting the ability of these households to respond to large price increases by obtaining cheaper rental accommodation. The housing price sensitivity of the housing consumption of median-expenditure young households was surprisingly stable over thirty years. Similarly, despite the large housing price increases, I observe quite stable preferences in the demand for food, fuel and light, leisure goods and services, and transport and other services over this time. In addition, my estimates of compensating variation show that despite the large housing price increases, the prices of most other goods and services declined over the sample period relative to income growth. Therefore, young renter households were better off in the late 2010s relative to the late 1980s, although not better off relative to the mid-2000s following stagnating wages and young renters allocating an increasing share of their budget to housing but also non-housing services where price inflation was greater. The welfare of low-expenditure renter households improved the most over this period, as social housing shielded many of them from the full impact of private rental increases, and they also particularly benefited from large falls in the price of food.

Both health and housing have spawned vast economic literatures, to which I contribute by focussing on previously under-explored topics that are relevant to public policy. All three papers are motivated by some of the key public policy issues facing the United Kingdom at present. On health, the UK government faces an increasing need to manage the costs of an ageing population, tackle the sharp rise in the number of working-age individuals out of the labour force on sickness and disability benefits, and manage the increasing resources required for public health provision. Improving our modelling of health as a dynamic process can help us better understand some of the complex ways that health impacts economic decision making, and can help poli-

cymakers better predict the impacts of new policies in this space. Understanding the occupation mobility patterns of those who fall ill but keep working can particularly help policymakers better support those with disabilities or chronic health conditions, as well as older workers, who are more likely to be in poorer health, remain in the workforce. Indeed, a UK government report by the Department of Work and Pensions (Salis et al., 2021) highlighted the lack of literature on this topic and called for further research into better understanding labour market transitions of those in poor health to better support policymakers. Housing affordability and the welfare of young renters is another important policy challenge in the UK, and is linked to rising productivity and inter-generational fairness concerns. Young renters make up 11 per cent of the working age sample, and understanding the biggest change to their consumption bundle in recent decades is also helpful for designing tax regimes and housing and other government support programmes.

## HOW SHOULD WE MODEL HEALTH AS A DYNAMIC PROCESS?

### Chapter Abstract

Health is a complex dynamic process that impacts many economic decisions in ways that remain poorly understood. This paper comprehensively reviews how health is modelled in the literature, showing that baseline models typically fail to take into account how persistence and frequency of health shocks vary by past health history and magnitude and direction of past shocks. Methods from the earnings dynamics literature are adapted to produce improved health persistence estimates. This paper also investigates how medical biomarker data can be incorporated in dynamic models of health as a proxy for underlying health. There is significant scope for further work in this area as more medical data becomes available to researchers.<sup>1</sup>

*JEL classifications:* I10, I31, C5

---

<sup>1</sup>This chapter has previously been published in August 2023 a working paper as part of the Department of Economics, University of Oxford, Discussion Paper Series (no. 1023). I would like to thank the many people who provided advice and feedback with this paper, especially Hamish Low, Francis DiTraglia, Steve Bond, participants at the RES 2024 Annual Conference, RES 2023 PhD Conference, Essex PhD Conference in Applied Economics, and Applied Microeconomics internal seminars at the University of Oxford

## 2.1 Introduction

Health is an important determinant of an individual's economic decision making, affecting labour supply, consumption, family composition and access to government-provided insurance. Accurately modelling health as a dynamic process is needed to answer several important and open questions in the literature, including what is the relationship between health and earnings inequality, how effective are the current government-provided safety nets for those who fall ill, and how best should governments respond to the increasing economic burdens of chronic disease, rising disability rates and an ageing population. There has been a significant amount of reduced-form work on these questions (Prinz et al. (2018) provides a good summary). More recently, structural approaches have been used to better understand some of the complex endogeneity between health and economic decision making. These models typically capture health dynamics in a highly simplified way to minimise computational burden. The contributions made by this chapter are relevant to both these reduced form and structural approaches.

This chapter makes three key contributions to the literature. Firstly, it provides a comprehensive assessment of the different ways health dynamics have been modelled in the literature. Researchers have typically borrowed from the earnings dynamics literature and modelled health as a simple linear process such as an autoregressive moving average process, often discretised as a first-order Markov process, or the sum of a permanent and transitory shock. To the author's knowledge, there have been no prior attempts to systematically evaluate these modelling approaches and their underlying assumptions. I use Understanding Society data, a commonly-used UK household panel dataset, to replicate the most common models of health dynamics using standard panel data techniques. I then evaluate how well they capture key features of the health process, focussing on estimates of persistence and cross-sectional heterogeneity caused by different health shock realisations. I show that an ARMA(1,1) model with a large AR(1) coefficient close to one and a moderately-sized negative MA(1) coefficient fit the data reasonably well. An alternative linear model that combines a permanent process and a transitory process can be a desirable alternative as it allows for two types of shocks with different properties, but at the expense of some very strong and potentially incorrect persistence assumptions.

One of the most important components of health dynamics to accurately model is persistence. An individual is likely to respond very differently to a highly persistent health shock compared to a moderately persistent one. In the ARMA(1,1) model, the AR(1) term captures the average persistence of health from last period, which is then modified by the MA(1) term depending on the magnitude of the prior period's error term. I show that persistence heterogeneity is much greater than captured by this model, and varies systematically by past health and the features of the health shock. On average, health shocks are more persistent if they are negative (a decline in health rather than an improvement), if the individual was in poor health prior to the shock, and if the health shock is large. The standard linear models of health dynamics do not capture this heterogeneity, and therefore tend to be overly-optimistic in modelling the health dynamics of those with a history of poor health who experience additional negative health shocks.

A related limitation is that these models are not particularly effective in capturing the different distributions of health shock risks that individuals face. While an ARMA(1,1) model can be estimated using GMM techniques that are fairly robust to various error distribution assumptions, the most obvious application of the model would impose a mean-zero independent and identically distributed (i.i.d.) normal error distribution, while the model that is a sum of a random walk and a moving average transitory process has a normally-distributed error term. I document several ways that the error terms, which I interpret as health shocks, deviate from an i.i.d. normal distribution. First, there is a strong relationship between past health and the expected distribution of future health shocks. Those in poor health face an increased risk of both large negative and large positive health shocks, while the variance of health shocks faced by those in good health is much lower. The variance of health shocks is higher for negative shocks than positive shocks, even when controlling for past health. Finally, the baseline models do not accurately replicate the higher order moments of the data.

The second contribution of this chapter is to address many of the limitations of these standard linear models of health by adapting a recent panel data technique from the earnings dynamics literature. I use Arellano, Blundell and Bonhomme (2017)'s quantile-based method to produce non-linear persistence estimates that allow for a

large amount of heterogeneity. One attraction of this framework is that it allows for persistence to vary depending on the size and sign of the health shock that occurs in period  $t$ , which cannot be done using the standard linear models due to the endogeneity between the shock and persistence estimates that relate health in period  $t - 1$  to health in period  $t$ . Applying this framework to my health data produces persistence estimates that range from 0.6 to 1.2, depending on prior health and characteristics of the shock in period  $t$ . This framework is able to capture that the persistence of the health process is higher among individuals in poor prior health, and that positive health shocks are typically less persistent than negative health shocks. These improved persistence estimates better capture the health risks faced by individuals, with implications for our understanding of the impact of health on economic decisions such as labour supply and consumption. I also estimate an extended version of this framework that is able to strip out time-invariant unobserved heterogeneity from the persistence estimates, and consider the wider applicability of the framework by applying it to produce non-linear persistence estimates of an index of mental health.

Finally, this chapter investigates how best to use increasingly-available medical data to improve health modelling. These data are available for a subset of individuals in the Understanding Society dataset. Previous studies have shown that biomarker data such as inflammation markers and steroid hormones in the blood can predict future adverse health outcomes in ostensibly healthy people (Davillas and Pudney, 2020a). To the author's knowledge, this data has never been used to better model health dynamics. I find that incorporating the biomarker data into my models of health dynamics does not improve their persistence estimates. However, the data can be used to better model the different health risks individuals face. I show that the ARMA(1,1) model performs less well in cases where the biomarker data suggest that an individual's underlying health is very poor. These are typically cases where an individual does not report any serious health conditions, but they face a significantly elevated risk of negative health shocks. This is an important source of risk to capture. I also find that variation in biomarker data is strongly correlated with the variation captured by the fixed effect component of the persistence estimates produced using the Arellano, Blundell and Bonhomme (2017) framework. This suggests that biomarker data can be used to better understand and model individual heterogeneity in health

outcomes, a topic that remains poorly understood.

The remainder of this chapter is structured as follows. Section 2 reviews the relevant literature and Section 3 describes the data, focussing on the construction of indices to capture observed and underlying health. Section 4 reviews the baseline dynamic health models in the literature and section 5 identifies their limitations. Section 6 applies methods from the earnings literature to produce non-linear estimates of persistence. Section 7 concludes.

## 2.2 Literature Review

There is a body of literature that develops methods of aggregating survey health data into an index of overall health, which I summarise in the data section of this chapter. However, answering questions on the relationship between health and economic outcomes often requires us to take a stance on how health evolves over time. There is some reduced form work on this question (O’Donnell, Van Doorslaer and Van Ourti, 2015), but the most common approach in the literature has been to apply the vast literature on modelling earnings dynamics to model health as a simple linear process, most commonly as an ARMA(p,q) process or the sum of a persistent and a transitory component. In the structural literature, a discrete version of this approach; a first-order Markov process with a small number of discrete health states, has commonly been used. However, the implications and limitations of these models has only very recently begun to be examined in the literature. I review the modelling health as a dynamic process literature, highlighting the gaps that this chapter seeks to fill.

The canonical papers that model the time series properties of the mean of earnings, such as Lillard and Willis (1978), MaCurdy (1982), and Abowd and Card (1989) use panel data to fit ARMA-type processes to earnings data. A recent example of this approach applied to health data is Blundell et al. (2020a), who represent health ( $\tilde{h}_t$ ) using the error correction model:  $\tilde{h}_t = \pi_t + \varepsilon_t$ . The persistent component ( $\pi_t$ ) evolves as a random walk:  $\pi_t = \pi_{t-1} + \eta_t$ , and  $\varepsilon_t$  is a MA(0) transitory component. Alternative specifications in the literature include modelling the persistent component as an AR(1) or higher order process so the effect of a shock to the persistent component decays over time, and adding more structure to the transitory component, such as by incorporating moving-average terms (Blundell et al., 2016), or by modelling health as a stock that decays (Wallenius, 2020). To reduce dimensionality, health processes that are included in structural models are typically discretised into a first-order Markov process that models transitions between discrete health states. There are many examples: Palumbo (1999), French (2005), De Nardi, French and Jones (2010), Attanasio, Kitao and Violante (2010), French and Jones (2011), Capatina (2015), Jung and Tran (2016), Braun, Kopecky and Koreshkova (2017), Imrohorglu and Zhao (2018), Jolivet and Postel-Vinay (2020), Nygaard (2021) and Amengual, Bueren and Crego

(2021). Earlier structural papers typically only modelled two health states, good and bad health, while more recent papers tend to include additional states, for example Jolivet and Postel-Vinay (2020) model four states of mental health: good, average, poor and severe. Some of these papers endogenise the health process by incorporating the impact of choices such as unhealthy consumption of cigarettes (Nygaard, 2021) or medical expenditure choices (Prados, 2018). Zweifel, Breyer and Kifmann (2009) model people choosing the level of health investment to marginally alter their transition probabilities between different health states. An important distinction between these Markov models and ARMA models is that the latter imposes a symmetry between positive and negative health shocks. Markov models do not have this feature, and the data suggest that the transition probability from good to poor health differs from the transition probability from poor health to good health.

Both ARMA and Markov models emphasise the state dependency of the health process. This can downplay the importance of individual heterogeneity in explaining the large cross-sectional variance in health observed in the data. Halliday (2008) finds that individual characteristics that trace back to childhood and early adulthood play an important role in determining how long health shocks persist, while the importance of state dependence varies significantly. However, he acknowledges that his first-order Markov model with only two health states limits his ability to pin down state dependence. Hauck and Rice (2004) similarly emphasise the importance of individual heterogeneity relative to state dependence in modelling mental health transitions. Pashchenko, Porapakarm and Nardi (2017) find that variation in health transitions due to ‘health types’ is much larger than variation due to state-dependence for men with a high-school education. Of particular interest is some ongoing work recently presented by De Nardi (2024), which identifies health types with different expected health trajectories. Performing clustering analysis on frailty measures, they identify five health types that they label as vigorous resilient, fair-health resilient, fair-health vulnerable, frail-resilient, and frail-vulnerable. They find that these types explain a large share of subsequent health trajectories of older adults, and significantly outperform forecasts of health trajectories based on initial health and a rich set of observables. These classifications are based on a clustering algorithm and the authors do not attempt to explain what causes these different health types. However,

the authors do highlight the recent empirical literature that emphasises the life-long economic consequences of genetics and early childhood experiences such as Barth, Papageorge and Thom (2020), Conti and Heckman (2010), Case, Fertig and Paxson (2004), Harris et al. (2016) and Cronqvist and Siegel (2015). Understanding the nature of this individual heterogeneity is of central importance to answering questions such as what causes the relationship between health and education, or health and earnings inequality, which currently remains poorly understood.

Some of the recent papers containing structural models have made progress in capturing additional complexity of health dynamics, most commonly by adding an extra variable that varies health shock risk such as ‘health type’ or ‘underlying health’. Pashchenko, Porapakarm and Nardi (2017) augment a standard first-order Markov model of health with transition probabilities that also depend on the duration of the current health spell and ‘health type’, which is a proxy for individual heterogeneity and affects transition likelihood. They find evidence of ‘duration dependence’ where the longer that someone has stayed in a particular state of health the less likely they are to transition states next period. This is not consistent with a low-order Markov process of health dynamics. Salvati (2021) incorporates a similar fixed-effect variable which is described as a proxy for high or low health into her model of health. She embeds a health equation into her life-cycle model that consists of an AR(1) process, a binary fixed effect term, a labour-market health interaction term, and various independent variables. Ozkan (2017) models two types of health capital: physical health capital that determines survival probability and preventative health capital that is subject to health shocks and can be modified by health investment. Keane, Capatina and Maruyama (2020) make progress in modelling individual heterogeneity by incorporating an asymptomatic health risk variable estimated with medical data. In this model, individuals have functional health that is subject to three types of shocks: predictable and persistent shocks, unpredictable and persistent shocks, and unpredictable and transitory shocks. Asymptomatic health risk captures conditions such as high cholesterol, high blood pressure and high BMI that do not directly affect daily life but increase the probability for future predictable adverse shocks to functional health. While these models have made progress in capturing health dynamics, these equations tend to be a small component of large and complex structural models

with computational demands that limit what these models can capture. The ‘black-box’ nature of these models can make it difficult to understand the mechanics of the interactions between health and other variables. This chapter identifies some of the limitations of modelling health in this way.

## 2.3 Data

The main dataset used in this chapter is Understanding Society - the UK Household Longitudinal Study. This is a longitudinal, nationally representative dataset with good coverage of health, education, employment, family life and income variables. I build an unbalanced panel using waves 1-12 of the study, which include observations from 2009–2021. Excluding a small number of individuals with insufficient health information results in a sample of 265,830 observations from 29,886 unique individuals. Table 2.1 reports the summary statistics of this sample and indicates good coverage over age, education, family type and employment.<sup>2</sup>

### 2.3.1 Health index construction

In many settings, the theoretically-ideal health index would be an overall stock of health measure, or a related concept such as a work-capacity index. Such an index would be continuous and bounded from below (death). Since these are unobservable concepts, we can instead construct a proxy index by aggregating various health data from household panel surveys. The available data can be grouped into three main categories. Objective health data are data on specific diagnoses and disabilities. Subjective health data are based on survey respondents' assessment of their own health. A third category of data is medical data such as pulse, blood pressure readings, blood tests or genomic data that can be used to predict health outcomes. Some of these medical data, such as genetic information, may be plausibly exogenous to any experiences or choices of the individual, which can be valuable for statistical analysis.

The limitations of each of these categories of health data as proxies for overall health has been thoroughly evaluated in several papers (Blundell et al. (2021) provides a good summary of this literature). To briefly summarise, objective measures are vulnerable to omitted variable bias, they can only capture a subset of relevant conditions, and often lack disease severity information. The rate of omission of life-changing medical diagnoses such as heart attacks and strokes reported by survey

---

<sup>2</sup>There is some gender imbalance in the sample (57% female, 43% male). This mostly reflects the raw Understanding Society data, which is split 55% female and 45% male. My sample is then further female skewed by men being more likely to enter the sample as proxies where partial information is provided about them by another household member, but they are unable or unwilling to respond themselves. Therefore, I do not have their subjective health scores and drop them from the sample.

Table 2.1: Summary statistics

	men	women
<i>age</i>		
<30	17,205	24,567
30-39	17,026	24,991
40-49	21,968	30,142
50-59	21,938	28,848
60-69	19,815	24,204
70-79	12,489	13,950
80-89	3,553	4,448
90+	282	404
<i>education</i>		
below GSCEs	25,743	31,135
GSCEs	31,755	41,889
A-level	13,476	16,622
degree	43,302	61,908
<i>family type</i>		
cohabitating/married	80,910	95,630
widowed	3,315	10,427
separated/divorced	7,111	16,091
single	22,772	29,068
<i>number of children</i>		
0	86,026	106,463
1	11,818	19,822
2	12,271	18,718
3	3,299	5,088
4+	862	1,463
<i>currently employed</i>		
yes	72,589	88,493
no	41,500	62,779
<i>occupation class</i>		
professional	6,324	4,820
managerial & technical	28,854	34,801
skilled non-manual	9,532	24,154
skilled manual	18,719	9,105
partially skilled	7,269	14,316
unskilled	2,651	2,005
<i>N (observations)</i>	<i>114,276</i>	<i>151,554</i>

respondents has been found to be surprisingly high when compared to linked hospital admission data, suggesting measurement error could be large (Caraballo et al., 2020). Subjective health measures are fairly crude and vulnerable to reporting error and justification bias. For a given disease presentation, people will vary hugely in how poorly they rate their health and to what degree they report that the disease has a negative impact on their life (French and Jones, 2017). Medical data are less commonly collected in household surveys and there is limited research on how best to use them to model health.

A challenge in the literature has been how best to use these data to construct an overall health index that minimises these biases and approximates the ideal theoretical health concept. Lack of consensus on this question has contributed to the ongoing uncertainty of the relationship between health and employment (Blundell et al., 2021). For example, large differences have been found when estimating the impact of poor health on labour supply using objective or subjective health data (Anderson and Burkhauser, 1984). To reduce these biases, a common approach has been to instrument subjective health data with objective data. This approach is still regularly used, with Blundell et al. (2020a) being a recent example, although the approach is not without criticism. Bound (1991) argues that the different types of biases affecting subjective health measures roughly offset, so that incorporating objective health data adds little value and may increase bias. Alternative approaches to aggregating health data have included taking the first principal component over a large number of objective measures (Poterba, Venti and Wise, 2017), constructing multiple indices (Blau and Gilleskie, 2001), and converting medical conditions into World Health Organisation disability weights that represent the magnitude of health loss associated with specific health outcomes, which can then be aggregated (Prados, 2018). A helpful contribution was made by Blundell et al. (2021) who comprehensively evaluated the different approaches in the literature to identify how should health data be combined to best represent overall health. They conclude that objective measures, provided that a large enough set of them are used, subjective measures, and subjective measures instrumented with objective measures can produce similar estimates of the impact of health on employment, and any of these modelling approach can reasonably be used. This finding was broadly supported by Hosseini, Kopecky and Zhao (2022),

who compares the performance of a ‘frailty index’ that aggregates objective indicators with a subjective health index, and an index constructed using principal component analysis, and similarly finds that the predictive power of the different approaches to be broadly comparable.

I follow the literature and use both subjective and objective health data to construct a single health index that functions as a proxy for an individual’s overall stock of health. The subjective data comes from the survey question ‘*In general, would you say your health is: poor, fair, good, very good or excellent?*’. The objective health data used is reported in Table 2.2.

To construct a single health index, I follow the approach of Blundell et al. (2020a) and estimate an ordered probit of an individual’s subjective reported health on a rich dataset of objective health measures, and then take the predicted values from this regression to be the individual’s health index. I run the following ordered probit regression where  $H_{it}^*$  is the unobserved continuous latent general health variable and  $H_{it}$  is the observed ordinal general health score assessed by the individual in period  $t$ .  $H_{it} = \{1, 2, 3, 4, 5\}$  where 1 = poor, 2 = fair, 3 = good, 4 = very good, and 5 = excellent.  $X_{it}$  is a vector of objective measures and some additional controls, and  $\epsilon_{it}$  is the individual error term. The included controls are age, sex, an employment dummy, occupation class, and month and year dummies. The Pseudo-R squared from these ordered probit regressions is around 0.2. Each wave is estimated separately, and a sample of the regression output is reported in Appendix A.0.2.

$$H_{it}^* = X_{it}'\beta_t + \epsilon_{it}, \quad \epsilon_{it} \sim \mathcal{N}(0, 1) \quad \forall i = 1 \dots N, t = 1 \dots T$$

$$H_{it} = j \quad \text{if } \mu_{jt} < H_{it}^* < \mu_{j-1,t} \quad j = \{1, 2, 3, 4, 5\}$$

The probability that individual  $i$  selects general health value  $j$  in period  $t$  is:


$$Pr(H_{it} = j) = \Phi(\mu_{jt} - X_{it}'\beta_t) - \Phi(\mu_{j-1,t} - X_{it}'\beta_t)$$

I then map the ordered probit fitted values onto the general health scores using a linear regression of the subjective scores onto the predicted values, and re-calculating the fitted values. The distribution of the original subjective health scores and constructed health index is shown in Figure 1.

Table 2.2: Objective health indicators

Objective measure	Data
Disabilities (specified as causing 'some difficulty' or 'much difficulty')	12 indicators: manual dexterity, mobility, lifting/moving objects, continence, hearing, sight, communication/speech, memory/ability to concentrate and learn, recognising danger, physical co-ordination, personal care, other
Mental wellbeing	General Health Questionnaire (GHQ) Caseness measure. Measures common mental health problems e.g. depression, anxiety, somatic symptoms, social withdrawal to detect those at risk of developing psychiatric disorders.
Ever diagnosed with condition	asthma, congestive heart failure, coronary heart disease, angina, heart attack, stroke, emphysema, hypothyroidism, chronic bronchitis, liver condition, epilepsy, hypertension, multiple sclerosis, COPD, osteoarthritis, rheumatoid arthritis, other arthritis, cancers: bowel/colorectal, lung, breast, prostate, liver, skin, other, diabetes: type 1, type 2, gestational and other, anxiety, depression, psychosis/schizophrenia, bipolar/manic depression, eating disorders, PTSD, other emotional/nervous/psychiatric condition, other chronic condition
Still have previously diagnosed condition	Conditions: asthma, congestive heart failure, coronary heart disease, angina, hypothyroidism, chronic bronchitis, liver condition, epilepsy, hypertension, COPD, osteoarthritis, rheumatoid arthritis, cancers: bowel/colorectal, breast, prostate and skin, type 2 diabetes, anxiety, depression, eating disorders, PTSD
Hospital out-patient	1-2 days, 3-5 days, 6-10 days, >10 days in the past year
Hospital in-patient	1-2 days, 3-5 days, 6-10 days, >10 days in the past year

Figure 2.1



The left hand side figure shows the raw health data; the right hand side shows the constructed health index data distribution

The constructed health index can be interpreted as the average subjective health score reported by all individuals with the same medical diagnoses and disabilities, controlling for individual characteristics such as age and sex. The distribution of these scores is left-skewed due to a large tail of individuals in poor health, and its kurtosis is around double that of a normal distribution, with many individuals bunching around the modal score.

Figure 2.2 indicates that differences in health index values between men and women are small. I include men and women in the same regression to calculate the health indices but include a gender dummy variable to allow for variation by gender. Average health index scores gradually decline with age, although they are fairly stable between the ages of 55-65. Variance in health scores increases with age, especially from the age of around 50. To strip out this decline in health by age, I demean the health index by regressing the health index against age, higher powers of age up to order four, sex, and month/year dummies. The residuals from this regression become the ‘demeaned health index’ that I use to model health as a dynamic process in subsequent chapters of this thesis

Figure 3 shows the distribution of changes to an individual’s demeaned health index over one year and ten years. In both cases the distribution has a slight negative skew, of -0.5 and -0.8 respectively. The approximate symmetry of shocks supports the use of simple linear ARMA models that impose symmetry of shocks.

Figure 2.2: Distribution of health index values by age and gender





Figure 2.3



A potential concern is that attrition rates vary systematically by health. In my dataset, around 18 per cent of observations do not have an observation next period, either due to attrition or missing data. Estimating a linear probability model of attrition indicates that those in the lowest health quintile are two percentage points more likely to not report health data next period relative to those in better health.

However, the literature is fairly sanguine about the risks of using health indices for economic research when there is differential attrition risk by health (Jones, Koolman and Rice, 2006; Pudney and Watson, 2013). I choose to follow this literature and do not attempt to adjust for attrition rates in my subsequent modelling of health dynamics. Further analysis of attrition in my dataset is reported in Appendix A.0.1.

### 2.3.2 Allostatic scores from biomarker data

Between 2010–12, a subset of 8,465 individuals from waves 2 and 3 of the main Understanding Society survey were visited by a nurse for a physical health check and gave a blood sample. I use this biomarker (biological marker) data to construct a second index that approximates a component of underlying health called ‘allostatic load’. This is a medical concept that reflects the risk from the cumulative effects of exposure to physical, psychosocial and environmental stressors that increase the risk of developing chronic diseases (Group, 2001). As a measure of cumulative wear and tear to the body, allostatic load is theoretically quite close to overall health stock or working capacity, although it cannot capture mental health or physical injury or disability.

Biomarker data can be used to improve health dynamics modelling for several reasons. They can be measured with less error than other health data that rely on an individual accurately describing their health. The availability of biomarker data is likely to grow rapidly following the increasing popularity of wearable health technology such as smart watches. They can help predict future health and mortality risk in ostensibly healthy individuals (Davillas and Pudney, 2020*b,c*). Davillas and Pudney (2020*a*) find that combining subjective health data with biomarker data significantly improves their predictions of future disability risk. This is because biomarker data incorporate health information such as kidney function and hormonal balance that may not be known by the individual, and because it offsets people’s bias towards over-weighting certain health information such as obesity and blood pressure, and under-weighting other information such as strength and lung function. Biomarker data can also help disentangle the endogeneity between health and economic outcomes, and have been used to better understand the income-health gradient (Davillas, Jones and Benzeval, 2019), the impact of economic insecurity and childhood economic circumstances on health (Niedzwiedz et al., 2017; Davillas and Jones, 2020), and comparing the health impact of becoming re-employed in poor-quality work compared to remaining unemployed (Chandola and Zhang, 2017).

To construct the allostatic score index, I normalise and then aggregate the biomarker data. I follow the approach of Davillas and Pudney (2020*a*) and take the simple average of the z-scores of 12 biomarkers and physical indicators reported in Table 2.3. I


then demean the allostatic scores by age and gender to match how I constructed my health index. The subsequent distribution is approximately normal (Figure 2.4).

Table 2.3: Biomarkers used in allostatic load index construction

Indicator	Data	Description
waist-to-height ratio	waist circumference, body mass index	obesity indicator
pulse	resting heart rate	lower heart rate associated with more efficient heart function
blood pressure	systolic, diastolic	two readings treated as separate indicators
lung function	forced vital capacity (FVC)	total amount of air forcibly blown out after a full inspiration using a spirometer
blood sugar	glycated haemoglobin levels (HbA1c)	measures glucose intolerance, a good indicator of diabetes risk
inflammation	C-reactive protein (CRP)	is a protein in the blood that rises in response to general chronic or systemic inflammation. High levels are risk factor for cardiovascular disease and mortality.
kidney function	creatinine	Creatinine is a waste product of muscle function, which is passed through the kidneys and excreted in urine. Glomerular filtration rate (eGFR) calculated using creatinine data according to calculation cited in Benzeval et al. (2014). Indicates how effectively the kidneys are ‘cleaning’ the blood.
liver function	albumin levels	albumin is main protein made by the liver. Low levels may be indicative of a loss of liver function
steroid hormone	dehydroepiandrosterone sulphate (DHEAS)	one of the primary mechanisms through which psychosocial stressors may affect health. Low levels associated with cardiovascular risk and all-cause mortality
cholesterol	high-density lipoprotein cholesterol (HDL)	‘good’ cholesterol that helps remove other forms of cholesterol from the bloodstream. High levels lower risk of cardiovascular disease.
grip strength	maximum grip strength	correlated with overall body strength, lower scores associated with decreased physical function, disability and mortality

A limitation of these data is that I only have one set of biomarker observations per individual. However, the predictive content of allostatic scores is fairly stable over time. I show this by regressing my health index against the allostatic score index, varying the time gap between the data used for the health index and the

Figure 2.4



allostatic score (Table 2.4). An allostatic score has similar predictive power for a health index based on survey data collected one year later, to a health index based on data collected ten years later. This suggests that allostatic scores capture a stable, long-term measure of health. There is a planned second round of biomarker data collection during wave 16 of Understanding Society in 2024–26, which can be used to check the stability of biomarker data over time more formally (Kumari, Al Baghal and Benzeval, 2022).

Table 2.4: Health index predictive content of allostatic scores

	Number of waves between collection of allostatic score and health index data								
	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10
allo.	0.527*** (0.0208)	0.533*** (0.0207)	0.526*** (0.0218)	0.517*** (0.0217)	0.522*** (0.0219)	0.562*** (0.0231)	0.583*** (0.0240)	0.468*** (0.0240)	0.439*** (0.0237)
R-sq	0.080	0.082	0.073	0.073	0.075	0.083	0.086	0.061	0.058
Obs	7434	7456	7400	7269	7024	6519	6249	5895	5536

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Alongside biomarker data, genetic data was also collected. However, the Understanding Society genetic data is safeguarded special license data. Therefore, I perform

some preliminary analysis using an alternate dataset, The English Longitudinal Study of Ageing (ELSA), on whether genetic data can be used as an additional health risk indicator when modelling health dynamics. The ELSA dataset reports polygenic scores for a variety of behavioural, emotional and health-related phenotypes, which estimate an individual's genetic propensity to develop various physical and mental health conditions. However, I find that while the polygenic scores do contain additional information on future health outcomes not captured by the the health or allostatic indices, the size of the additional information is too small to significantly improve my modelling of the overall health process. Further details of this analysis is reported in Appendix A.0.3.

## 2.4 Modelling health as a dynamic process

Health is a complex dynamic process that is subject to shocks that vary in magnitude and persistence. Heterogeneity between individuals is also large. In this section, I use panel data techniques to identify how best to model health as a simple linear process. I estimate two baseline models that replicate the two most commonly used approaches to modelling health dynamics: an ARMA(p,q) model, and a linear additive shock model that is the sum of a permanent process and a transitory MA(1) process. I find that an ARMA(1,1) model with a large AR coefficient and a moderately-sized negative MA coefficient best fits the data, although there are circumstances where the extra flexibility of the linear additive shock model in capturing two different shocks may be desirable. I then evaluate how effective these models are in capturing health dynamics accurately. I show that while these models can be appealing due to their simplicity and intuitive interpretation, they have some important limitations that I discuss in detail in the next section.

### 2.4.1 ARMA(p,q) baseline model

The two data attributes that I wish to capture in any baseline model of health are the persistence of innovations, and cross-sectional heterogeneity between individuals. My starting point is the simplest linear models that incorporate persistence; the autoregressive moving average (ARMA) class of models. I model health of individual  $i$  in period  $t$  ( $h_{it}$ ) as an ARMA(p,q) process that includes a fixed effect  $\mu_i$ :<sup>3</sup>

$$h_{it} = \sum_{k=1}^p \rho_k h_{i,t-k} + \sum_{j=1}^q \theta_j \varepsilon_{i,t-j} + \mu_i + \varepsilon_{it} \quad (2.1)$$

$$i = 1 \dots N, \quad t = 1 \dots T$$

The  $p$  lags of the  $\rho$  term make up the autoregressive AR(p) components, and the  $q$  lags of the  $\theta$  term make up the moving average MA(q) components. It is important to note that the health process I estimate is based on data that has been

---

<sup>3</sup>The ARMA models I describe in this section all allow for individual-specific fixed effects unless otherwise stated

detrended by age and gender. This was done by regressing the raw health index against these observable variables and taking the residuals as the detrended health index. This detrending is quite common in the literature, perhaps due to familiarity with modelling the component of earnings growth that is unexplained by observables such as experience and education. Furthermore, detrending by age removes the time trend as health declines over time, reducing the risk that the process is non-stationary. Small changes in survey design between waves is controlled for by including time dummies. Nonetheless, it may be attractive for the researcher to explicitly model the decline in health as people age. I replicate the key empirical work in this section with the original non-detrended health index, and report the results in Appendix A.0.4. I find that my results are robust to using a non-detrended index.

I test for stationarity, and find that at least a significant proportion of the series is stationary. I use the Born and Breitung (2016) test for panel series correlation, as it is designed to be robust to fixed effects and heteroskedasticity. A non-stationary pure random walk model would result in the autocorrelation of differenced health with its second (and higher) lag to be zero, which is not what we observe. Instead, this pattern of gradually decreasing autocorrelation in first differences is consistent with a persistent autoregressive process or a MA( $q$ ) process with a large  $q$ .<sup>4</sup>

In general, the literature finds mixed evidence of health following a random walk process as opposed to a highly persistent one. Blundell et al. (2020*a*) do find evidence of a random walk, while Blundell et al. (2016) estimate the coefficient on the first lag of health to be 0.9-1.1 depending on the sub-sample used, and Heiss, Venti and Wise (2014) estimate an overall coefficient of 0.9. Whether papers model health as a highly persistent or permanent process likely reflects sample selection or health index construction. For example, the use of a dataset such as ELSA or Health and Retirement Study (HRS) that only includes older individuals will have a higher proportion of highly-persistent health shocks compared to a more representative sample by age, which will contain a higher proportion of less-persistent health shocks such

---

<sup>4</sup>It is also interesting to note that the sign of the LM( $k$ ) test statistic in levels swaps from positive to negative from lag 4, indicating that health indices are positively correlated over short periods but negatively correlated over long periods. This is inconsistent with an ARMA(1,1) process, and could be explained by a combination of mean reversion and sample attrition. For example, an individual in poor health in period  $t$  is likely to also be in poor health in period  $t+1$  or period  $t+2$  but recovers by period  $t+4$  or attrits from the sample.

Table 2.5: Born and Breitung test for panel series correlation

	levels		first difference	
	LM(k)-stat*	p-value	LM(k)-stat	p-value
lag 1	36.30	0.000	-44.24	0.000
lag 2	24.63	0.000	8.36	0.000
lag 3	8.09	0.000	9.20	0.000
lag 4	-14.10	0.000	7.69	0.000
lag 5	-25.87	0.000	3.16	0.002
lag 6	-23.97	0.000	8.52	0.000
lag 7	-25.21	0.000	2.87	0.004
lag 8	-24.45	0.000	1.82	0.068
lag 9	-20.21	0.000	1.90	0.057
lag 10	-15.72	0.000	-1.22	0.224

\*LM(k) test statistic is a modified t test of  $\zeta = -1/(T - 1)$ .  $\zeta$  from equation  $h_{it} - \bar{h}_i = \zeta(h_{i,t-k} - \bar{h}_i) + \epsilon_{it}$ . k is the lag order being tested

as changes in mental health index scores. There are also different ways to construct health indices, and some may place more weight on more permanent health indicators such as disability diagnoses compared to indicators of temporary health conditions such as infectious disease history or mental health indexes. Blundell et al. (2020a) used ELSA data and constructed a health index that emphasised disability indicators, therefore it is unsurprising that they find evidence of a random walk.

I begin by estimating an AR(p) model using OLS with various values of p and no accounting for fixed effects, reported in Table 2.6.<sup>5</sup> The OLS estimates indicate that health is highly persistent, with the sum of coefficients on the lagged health terms consistently around 0.9. A major concern with using OLS is that the coefficient estimates may be spuriously high due to the presence of fixed effects. I strip them out using first differencing and avoid the resultant Nickell bias by using GMM estimation techniques. I use the Arellano-Bond ‘Difference GMM’ estimator which mitigates Nickell bias by instrumenting the lagged dependent variable terms with further lagged terms in levels. I re-estimate the AR(p) model now accounting for fixed effects, and adopting the following specifications which are selected to be conservative and robust: two-step estimator, time dummies, robust standard errors clustered at the individual level and an ‘unadjusted’ initial weighting matrix. I include the Windmeijer correction

<sup>5</sup>The sample size for the OLS and GMM estimations differ as the latter typically requires a higher t (additional lags) to generate the moment conditions

to correct for the usually negative bias in finite samples when the two-step estimator is used (Windmeijer, 2005). To prevent over-proliferation of instruments, I ‘collapse’ the instrument set and only include instruments based on the first to fifth lag of the variable being instrumented. My results are robust to various alternate specifications such as forward orthogonal deviations and different weighting matrices. I report the results of this exercise in Table 2.7. The MA(0) and MA(1) specifications indicate whether I allow the first lag ( $h_{i,t-2}$ ) to be used as an instrument in the first-differenced equations, which is a valid instrument if the errors follow an MA(0) but not MA(1) process.

Table 2.6: OLS estimates of the health process as an AR(p) model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
L1.health	0.850*** (382.52)	0.602*** (165.76)	0.535*** (117.42)	0.503*** (96.94)	0.485*** (82.66)	0.470*** (70.34)	0.469*** (62.31)	0.454*** (49.62)
L2.health		0.300*** (81.97)	0.244*** (47.88)	0.221*** (38.65)	0.213*** (32.27)	0.197*** (27.06)	0.203*** (25.85)	0.208*** (22.96)
L3.health			0.135*** (32.11)	0.108*** (18.63)	0.0938*** (13.57)	0.0825*** (10.67)	0.0878*** (10.23)	0.0922*** (9.33)
L4.health				0.0801*** (16.43)	0.0646*** (9.61)	0.0613*** (7.75)	0.0646*** (7.15)	0.0643*** (6.23)
L5.health					0.0533*** (8.93)	0.0350*** (4.29)	0.0352*** (3.79)	0.0423*** (3.89)
L6.health						0.0592*** (8.45)	0.0335*** (3.55)	0.0351** (2.99)
L7.health							0.0176* (2.17)	-0.00891 (-0.73)
L8.health								0.0149 (1.48)
Observations*	228,886	182,016	146,353	117,013	89,994	69,512	52,415	37,934

Standard errors in parentheses; clustered standard errors, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table 2.7: Difference-GMM estimates of the health process as an AR(p) model

	AR(1)		AR(2)		AR(3)		AR(4)		AR(5)	
	MA(0)	MA(1)	MA(0)	MA(1)	MA(0)	MA(1)	MA(0)	MA(1)	MA(0)	MA(1)
L1.health	0.241*** (0.00949)	0.944*** (0.0342)	0.478*** (0.0152)	0.971*** (0.0528)	0.571*** (0.0196)	1.008*** (0.0991)	0.574*** (0.0235)	1.184*** (0.156)	0.614*** (0.0281)	0.979*** (0.134)
L2.health			0.147*** (0.00748)	-0.0106 (0.0186)	0.203*** (0.0101)	-0.0453 (0.0560)	0.223*** (0.0125)	-0.190 (0.103)	0.241*** (0.0154)	-0.00671 (0.0888)
L3.health					0.0681*** (0.00678)	-0.0101 (0.0195)	0.0871*** (0.00861)	-0.0605 (0.0375)	0.0916*** (0.0106)	-0.00324 (0.0352)
L4.health							0.0351*** (0.00688)	-0.0170 (0.0158)	0.0380*** (0.00894)	-0.00104 (0.0169)
L5.health									0.00398 (0.00850)	-0.00925 (0.0106)
AB test, order 1, z score	-60.29	-31.22	-51.10	-18.73	-45.02	-9.10	-38.71	-6.38	-34.47	-6.14
AB test, order 1, p value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
AB test, order 2, z score	11.19	16.04	-1.79	9.89	2.00	4.58	-1.28	3.68	-2.32	2.64
AB test, order 2, p value	0.000	0.000	0.075	0.000	0.045	0.000	0.199	0.000	0.020	0.008
AB test, order 3, z score	-1.92	0.06	4.42	-0.47	-1.91	-0.54	0.75	0.66	2.65	2.49
AB test, order 3, p value	0.056	0.955	0.000	0.637	0.056	0.589	0.456	0.508	0.008	0.013
Hansen J test stat	550.65	0.91	127.31	2.51	35.64	6.83	28.25	5.51	20.60	10.99
Hansen J test p value	0.000	0.823	0.000	0.473	0.000	0.078	0.000	0.138	0.000	0.012
Moment conditions	16	15	16	15	16	15	16	15	16	15
Observations	222,095	222,095	184,734	184,734	151,622	151,622	121,698	121,698	94,513	94,513

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

These results strongly suggest that an ARMA(1,1) model best suits the data. All MA(0) specifications that include instruments based on the immediately preceding lag result in a strong rejection of the null of the Hansen J test, indicating that the model is wrongly specified. However, excluding this instrument, which the MA(1) specifications do, is typically sufficient to change the result of this test and fail to reject the null. For example, excluding  $h_{i,t-2}$  as an instrument for  $\Delta h_{i,t-1}$  and only using  $h_{i,t-3}$  and earlier lags leads to the non-rejection of the null. This strongly suggests that the errors follow an MA(1) process. Crucially, the Arellano-Bond autocorrelation tests also identify autocorrelation only up to the second order in most specifications. For all MA(1) specifications (except AR(5)), I find no evidence for third-order autocorrelation in the error terms, which is the key requirement for validity of the instruments used if I allow for the error component to follow a MA(1) process. In addition, when I exclude the first lag as an instrument, the point estimate of the coefficient on the first lag of health is much higher at around unity while the coefficients on all the subsequent lags are small and not significant. This suggests that including only one lag of health is sufficient.

Incorporating some additional moment conditions by using Blundell-Bond ‘System GMM’ estimation leads to improved ARMA(1,1) estimates. It is well known that the Arellano-Bond estimator does not function well when persistence is high. At the limit, if health follows a random walk ( $\rho_1 = 1$ ) then the difference GMM instruments are uninformative. Blundell and Bond (1998) suggest that there is a risk of serious finite sample bias at  $\rho_1$  values of 0.8 and higher, although they show that the bias is smaller with very large samples. The System GMM estimator typically performs much better in these circumstances. The additional moment conditions can also contribute to more precise coefficient estimation. This is particularly helpful as having to only use further lags as instruments due to the MA(1) error structure increases the risk of weak instruments. The additional initial moment restriction of  $\mathbb{E}(\varepsilon_{it}h_{i1}) = 0$  that is required for System GMM estimation is not a particularly onerous restriction for my data. Blundell and Bond (2023) state that this restriction holds automatically if the same process has generated the series for long enough before the start of the sample period. Since my first observation occurs at least 18 years after the the start of the health process, at birth, this may not be an unreasonable assumption.

Table 2.8 reports the AR(p) model coefficients estimated using System-GMM and allowing for MA(1) errors. Differences between the Difference and System GMM coefficient estimates are small, although using System GMM leads to much more precisely estimated coefficients, especially for the first lag. The coefficient estimates of the first lag are mostly not significantly different for the AR(1) AR(2) and AR(3) specifications, and the coefficients on additional lags are typically not significant. Therefore, including only one lag is sufficient to capture the persistence dynamics of this model.

Table 2.8: System-GMM estimates of the health process as an AR(p) model

	AR(1)	AR(2)	AR(3)	(AR4)	(AR5)
	MA(1) assumption				
L1.health	0.872*** (0.0123)	0.901*** (0.0310)	1.032*** (0.0790)	1.113*** (0.0864)	1.149*** (0.0980)
L2.health		-0.00940 (0.0170)	-0.0775 (0.0430)	-0.126** (0.0448)	-0.135** (0.0481)
L3.health			-0.0194 (0.0166)	-0.0361 (0.0199)	-0.0432* (0.0217)
L4.health				-0.0160 (0.00844)	-0.0115 (0.0116)
L5.health					0.00608 (0.00720)
AB test, order 1 z score	-53.86	-26.25	-11.81	-11.8	-10.78
AB test, order 1 p value	0.000	0.000	0.000	0.000	0.000
AB test, order 2 z score	19.19	12.69	6.27	6.88	6.32
AB test, order 2 p value	0.000	0.000	0.000	0.000	0.000
AB test, order 3 z score	0.032	-0.54	-0.39	0.08	1.75
AB test, order 3 p value	0.974	0.588	0.693	0.935	0.080
AB test, order 4 z score	0.867	0.54	-0.16	-0.36	-1.32
AB test, order 4 p value	0.386	0.592	0.871	0.721	0.188
Hansen J test stat	6.42	8.259	8.109	11.27	35.24
Hansen test p value	0.170	0.143	0.23	0.127	0.000
Moment conditions	16	17	18	19	20
Observations	222,095	184,734	151,622	121,698	94,513

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

I conclude that the best single-equation linear specification to capture the health process is the following ARMA(1,1) that accounts for fixed effects:

$$h_{it} = 0.87h_{i,t-1} - 0.33\varepsilon_{i,t-1} + \eta_i + \varepsilon_{it} \quad (2.2)$$

Since the AR term has already been estimated as 0.87, I estimate that the coefficient on the MA(1) term ( $\theta$ ) is -0.33 by re-arranging the ARMA(1,1) model as:  $h_{it} - 0.87h_{i,t-1} = \tilde{h}_{it} = \eta_i + \theta\varepsilon_{i,t-1} + \varepsilon_{it}$ . This is now a simple MA(1) process that can be estimated using GMM. I use the following three variance and covariance moments and report the coefficient estimates in Table 2.9:

$$\text{Var}(\tilde{h}_{it}) = \mathbb{E}\eta_i^2 + (1 + \theta^2)\mathbb{E}\varepsilon^2$$

$$\text{Cov}(\tilde{h}_{it}, \tilde{h}_{i,t-1}) = \mathbb{E}\eta_i^2 + \theta\mathbb{E}\varepsilon^2 \quad \text{Cov}(\tilde{h}_{it}, \tilde{h}_{i,t-2}) = \mathbb{E}\eta_i^2$$

Table 2.9: GMM estimates of MA(1) process

	$\rho = 0.87$
$\eta_i$	0.0380*** (0.00285)
$\theta$	-0.334*** (0.00593)
$\varepsilon_{it}$	0.335*** (0.00125)
Observations	222,095

The reported estimates in this table refer to the variance of the error components  $\eta_i$  and  $\varepsilon_{it}$ . Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

As a robustness exercise, I check whether there are large subgroups with health dynamics that are better captured by a different linear model. If this were the case, describing the health dynamics of the entire sample using a single ARMA(1,1) model may be misleading. I use the Sarafidis and Weber (2015) K-means clustering algorithm to divide the sample into as many clusters as required for the estimated slope coefficients of an AR(1) model to be the same within each cluster, accounting for individual-specific fixed effects. The algorithm divides my sample into two groups, containing 40 and 60 per cent of the sample respectively. This suggests that only two groups are needed to capture any heterogeneity in model coefficients. I then estimate AR(p) models separately for each group using GMM, re-assessing whether including one lag is sufficient and whether the error structure follows an MA(1) process. The regression tables are reported in Appendix A.0.5, as well as some summary statistics

for each group. I determine that the models for the two groups that best fit the data are an AR(2) and ARMA(1,1) respectively.

$$\text{Group 1: } h_{it} = 0.79h_{i,t-1} + 0.09h_{i,t-2} + \eta_i + \varepsilon_{it}$$

$$\text{Group 2: } h_{it} = 0.83h_{i,t-1} - 0.55\varepsilon_{i,t-1} + \eta_i + \varepsilon_{it}$$

The two models are quite similar. Both capture that health is a highly persistent process, and have an additional term that helps distinguish between highly-persistent health shocks such as chronic health conditions, and transitory health shocks. I conclude that an ARMA(1,1) model with fixed effects is sufficient to describe the entire sample and slope heterogeneity is not a significant concern.

## 2.4.2 Linear additive shock model

I conclude this section with estimating a slightly different model that allows for more flexibility in capturing shock persistence, but at the expense of imposing other restrictions. A specification used very commonly in the earnings dynamics literature, and sometimes in the health dynamics literature, relaxes the restriction of individuals being subject to only one type of shock. Instead, the variable is modelled as the sum of two independent random processes: a permanent shock process which is typically a random walk, and a transitory process which is either an MA(0) or MA(1):

$$y_{it} = p_{it} + v_{it}$$

$$\text{permanent process: } p_{it} = p_{i,t-1} + \zeta_{it}$$

$$\text{transitory process: } v_{it} = \varepsilon_{it} - \theta\varepsilon_{i,t-1}$$

An additive classical measurement error  $r_{it} \sim N(0, \sigma_r^2)$  can also be included. This clear distinction between permanent and transitory shocks reflects the influence of Friedman's Permanent Income Hypothesis on earnings dynamics research, but it is also conceptually attractive as researchers can cleanly classify most income shocks as either temporary, such as overtime or one-off-bonuses, or permanent, such as a job change (Meghir and Pistaferri, 2004). A similar intuition for health shocks being

divided into permanent shocks such as a physical disability and temporary shocks such as some mental health episodes is compelling, and adopted in papers such as Blundell et al. (2020a) and Blundell et al. (2016).

The canonical moment conditions used to estimate these models require that the permanent process is a random walk to achieve identification. This is a strong assumption for my data. It is challenging to distinguish between highly persistent and random walk processes in small-T panel data with significant individual heterogeneity, and I cannot reject that the coefficient on the lagged health term is 1 in many of the ARMA(p,q) models I estimated using GMM. However, my serial correlation tests do indicate that a reasonable proportion of the data are best characterised as following a persistent process rather than a random walk. It is also unclear how robust the resulting coefficient estimates are to small violations of the random walk assumption implied by the moment conditions. Proceeding with caution, I use the following moment conditions to estimate health as the sum of a permanent walk and MA(1) transitory process. Letting  $g_{it}$  be a change in  $h_{it}$  (equivalent to  $h_{i,t} - h_{i,t-1}$ ), we can identify the variance of the permanent component using the following moment condition from Meghir and Pistaferri (2004):

$$\mathbb{E}(\zeta_{it}^2) = \mathbb{E} \left[ g_{it} \left( \sum_{j=-(1+q)}^{1+q} g_{i,t+j} \right) \right]$$

. Since we cannot separately identify the variance of any measurement error, the variance of the transitory shock, and  $\theta$ , we can only use the moment conditions to place bounds on these coefficients with the following moment conditions:

$$\sigma_r^2 = \mathbb{E}(g_{it}, g_{i,t-1}) - \frac{(1 + \theta)^2}{\theta} \mathbb{E}(g_{it}, g_{i,t-2})$$

$$\sigma_\varepsilon^2 = \frac{\mathbb{E}(g_{it}, g_{i,t-2})}{\theta}$$

By setting  $\sigma_r^2$  to zero we can estimate the lower or upper bound of  $\theta$ , which we assume is bounded between -1 and 1. The sign of  $\mathbb{E}(g_{it}, g_{i,t-2})$  defines the sign of  $\theta$ . In my case it is negative, therefore the maximum value of  $\theta$  is the case where  $\sigma_r^2 = 0$ . I use these moment conditions to estimate the variance of the two shocks, as well as the

coefficients of the MA(1) transitory process. These estimates are reported in Table 2.10. My estimates of the magnitude of the variances of the two shocks are quite similar to the findings of Blundell et al. (2016), although they do not find evidence of an MA(1) transitory process. Blundell et al. (2020a) obtain quite different results and argue that transitory and permanent shocks contribute fairly equally to health variance. However, they use a very different estimation strategy and do not use these canonical moments from the earning dynamics literature.

Table 2.10: Coefficient estimates of linear additive shock model

Variable	Estimate
$\mathbb{E}(\zeta_{it}^2)$	0.155*** (47.26)
$\sigma_\varepsilon^2$ if $\sigma_r^2 = 0$	0.050*** (40.89)
$\theta$ if $\sigma_r^2 = 0$ (upper bound)	-0.072*** (-6.42)

t statistics in parentheses, \*p<0.05, \*\*p<0.01, \*\*\*p<0.001

If transitory shocks do not explain much variation of the overall health process, then using an ARMA(p,q) model that only allows for one type of shock is sufficient. My results suggest that the permanent process is responsible for the majority of the variance in the health process over time, although the transitory process does make some contribution. In addition, the ARMA(p,q) model does not require the persistent shock to follow a random walk. I conclude that the ARMA(p,q) model is a superior fit for my data.

I further consider the limitations of these two models in the next chapter, and suggest some improvements.

## 2.5 Capturing more complex dynamics

The two linear health models estimated in the previous section are simple to use and incorporate into more complex structural models. However, there is a cost to their simplicity. Since the baseline ARMA(p,q) model attempts to capture the average persistence of a health shock, it imposes uniformity of persistence on shocks of different sizes, for positive and negative health shocks, and for individuals with very different levels of health and health histories pre-shock. I find evidence of significant heterogeneity in persistence once I allow persistence to vary by these characteristics. Simple extensions of the ARMA(p,q) baseline model can capture some of this variation, however we can make further progress with more sophisticated modelling approaches, which I discuss in the subsequent chapter. In addition, I do not need to assume stationarity or that the error terms follow a white noise process for my ARMA(p,q) coefficient estimates to be valid. However, I document some features of the error distribution that are important to capture when modelling the heterogeneity in health shock risk that individuals face. I show that biomarker data can be used to capture some of the elevated negative health shock risk faced by some individuals.

This section focusses on the ARMA(1,1) model as my preferred linear model, but most of the limitations I identify can also be applied to the linear additive shock model. De Nardi, Fella and Paz-Pardo (2019) provide a good summary of the key limitations of this model when applied to earnings data, which are equally valid when using health data. The key model assumptions they identify that do not match the data are: age independence of the second and higher moments of the conditional distribution of both the transitory and persistent components, normality of the shock distribution, and linearity of the process of the persistent component.

### 2.5.1 Recent health history

The average persistence of a health shock varies significantly depending on the health history of the individual prior to the shock taking place. This makes intuitive sense; someone's capacity to recover from an illness is a function of how healthy they were just prior to getting sick. The MA term in an ARMA(1,1) model takes into account the size of the shock last period, but there is significant additional persistence

information in the level of health. The simplest way to capture this would be to add an interaction term to the baseline ARMA(1,1) model that assigns individuals to a quintile of their health just before the shock, and interact it with the lagged health term, which I report in Table 2.11.

Table 2.11: ARMA(1,1) model with interaction dummy for lagged health level quintile

	Difference-GMM	System-GMM
Q1 - lagged health index	0.966*** (0.0697)	0.923*** (0.0163)
Q2 - lagged health index	1.102*** (0.159)	0.910*** (0.104)
Q3 - lagged health index	0.754*** (0.100)	0.665*** (0.0719)
Q4 - lagged health index	0.901*** (0.0569)	0.829*** (0.0433)
Q5 - lagged health index	0.871*** (0.0349)	0.868*** (0.0260)
AB test, order 1 z score	-24.76	-56.093
AB test, order 1 p value	0.000	0.000
AB test, order 2 z score	14.3	20.835
AB test, order 2 p value	0.000	0.000
AB test, order 3 z score	0.0567	0.0915
AB test, order 3 p value	0.9548	0.927
Hansen J test stat	25.925	67.035
Hansen J test p value	0.0388	0.000
Observations	222,095	222,095

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

GMM estimation specifications identical to those used in baseline model in section 4

This exercise of relaxing the restriction that the autoregressive parameter is common across quintiles of the lagged health index suggests that persistence is the highest for those with prior bad health, and lowest for those with prior average health. However, these results should be taken with extreme caution as the Hansen test strongly


rejects the validity of the over-identifying restrictions used by both the Difference-GMM and System-GMM estimators. This approach also is unable to allow coefficients to depend on the sign or magnitude of the health shock between periods  $t - 1$  and  $t$ . I adopt more complex econometric techniques in a later section of this chapter to capture this heterogeneity in the persistence of health shocks.

As well as the relationship between recent health history and persistence, there is also a relationship between recent health history and the expected distribution of future health shocks. This is difficult to capture in a simple linear model but is an important component of health risk to capture. To illustrate the relationship between past health and the expected distribution of health shocks, I graph the higher moments of the health index (variance, skewness, and kurtosis) as a function of that individual's health percentile in the previous period, where 1 is the bottom health percentile of all individuals and 99 is the highest health percentile in Figure 2.5. The variance depicted in the top panel is calculated separately for the subset of individuals who experienced a 'positive health shock', meaning their reported health improved between the current and immediately prior period, and those who suffered a negative health shock, which is defined as the opposite.

I find that variance, skewness, and kurtosis all systematically vary by health the previous period. Notably, those in poor health have more volatile health in subsequent periods, with increased risk of both large negative and positive changes to their health relative to those in good or average health. This elevated risk is difficult to capture using simple linear models. One plausible way of capturing this feature of the data, is estimating an autoregressive conditional heteroskedasticity (ARCH) model. ARCH models are able to capture differences in variance depending on the size of the error term in the previous period. For example, a large shock in period  $t - 1$  may mean a large shock is more likely in period  $t$ . Figure 2.6 compares the histograms of health changes between period  $t - 1$  and  $t$  for those who experienced a greater than one standard deviation change in health in the prior period (between  $t - 2$  and  $t - 1$ ) to those who did not. It shows that large changes in health is associated with more volatile health next period, and that, on average, a large negative health shock is associated with an improvement in health next period, and the opposite holds for those who experienced a positive health shock in the prior period. Interestingly, the


distribution of health changes for those who experienced large positive or negative shocks in the prior period is closer to a normal distribution than the distribution of health changes of those who experienced neither. For this group, there is very little mass in the tails as stable health in the past is correlated with stable health in the next period.

Figure 2.5



x-axis is individual's health percentile in the previous period, where 99 = best health percentile.  
y-axis is units of higher order moment being graphed

Figure 2.6



I assess this relationship between errors and variance more formally by estimating an ARCH(1) model with the following specification for health variance:  $\exp(\gamma_0 + \gamma_1 \varepsilon_{i,t-1}^2 + \gamma_2 \varepsilon_{i,t-1})$ . The  $\gamma_2$  term accounts for possible heterogeneity between positive and negative shocks. I describe my estimation procedure in Appendix A.0.6, but I do not find any evidence that  $\gamma_1$  or  $\gamma_2 \neq 0$ , and therefore do not find evidence of ARCH effects in my data. However, this specification only models the relationship between shock magnitude in two consecutive periods. I do find evidence that individuals who experience a large negative health shock are more likely to experience another large negative health shock in subsequent years. However, a majority of these later shocks occur several years afterwards, which cannot be captured in an ARCH(1) model and requires a more complex econometric approach. Table 2.12 reports the number of large negative shocks, defined as at least one standard deviation fall in the detrended health index, experienced by those of different ages in the sample. Conditional on experiencing one negative shock, individuals are more likely to experience a second. For example, those aged 20-29 at the beginning of the sample period have a 16 per

cent chance of experiencing a negative shock in the next decade, but 28 per cent of those who experienced one negative shock experienced a second negative shock, with an average gap between shocks of four years. The average gap between negative health shocks rises with age.

Table 2.12: Number of large negative shocks over 10 years, population share by age<sup>†</sup>

Age in first wave	0 shocks	1 shock	2+ shocks	average yrs b/tween shocks
20-29	0.801	0.155	0.044	3.5
30-39	0.784	0.169	0.047	3.7
40-49	0.744	0.198	0.058	4.2
50-59	0.722	0.222	0.056	4.4
60-69	0.704	0.242	0.054	4.6
70-79	0.610	0.300	0.090	4.5


<sup>†</sup>shocks of at least one standard deviation

## 2.5.2 Age and model stationarity

A different source of heterogeneity in persistence and shock distribution is age of the individual. Age is closely related to the statistical property of stationarity. Since the time dimension of my panel data is fairly short, stationarity is difficult to assess. However, the ARMA(1,1) process I estimated is stationary in the long run, provided that  $\rho + \theta \neq 0$ ,  $\rho < |1|$ , and some not particularly onerous restrictions are imposed on the distribution of  $\varepsilon_{it}$ . Stationarity implies that the moments of the data are age independent. For the first moment, this is mechanically achieved by detrending the health index by age and age polynomials. However, higher moments of the detrended health data are not age-independent. Figure 2.7 shows the second, third and fourth moments of the detrended health data by age. Older individuals are more likely to experience health shocks, and so the standard deviation of the health index increases with age. The distribution of the detrended health index of older people is less negatively skewed, reflecting their increased propensity to experience positive health shocks. Young people are much less likely to experience positive health shocks as

their health is typically good and so cannot be improved further. The health index distribution for young people is platykurtic and so extreme health changes are rarer, while the kurtosis for older people is close to a normal distribution.

Figure 2.7



Higher moments do systematically vary by age, which should be captured in life-cycle models or models that consider the long-term impacts of health shocks on economic outcomes. This can be achieved by imposing a shock sequence that is a function of age rather than assuming a normal distribution for the error term. The ARMA(1,1) coefficients I estimate using GMM are robust to conditional heteroskedasticity and


the patterns of kurtosis and skewness I identify. Assuming mean-zero errors and no serial correlation of the errors is sufficient for this to be the case (Arellano and Bond, 1991). However, these higher moments are an important component of capturing the health risk people face.

### 2.5.3 Underlying health

I conclude with considering how allostatic scores can be used as an additional data source to improve the performance of linear health dynamics models. I find that the main informational content of allostatic scores relates to the likelihood of a large negative health shock in the future. In addition, I find that allostatic scores do not help predict the persistence of already realised health shocks (see Appendix A.0.7 for further details).

The ARMA(1,1) model has the worst performance when predicting the health of individuals with poor allostatic scores. Figure 2.8 shows the average difference between the level of health predicted in period  $t$  using my preferred ARMA(1,1) model as estimated using System GMM (see section 2.4.1 for further details), and actual health in period  $t$  by allostatic score quintile. These are the residuals for the health equations in levels, averaged for each allostatic score quintile. The biggest forecast misses occurs for the population with the worst 10 per cent of allostatic scores, indicating very poor underlying health. If someone is in average health in period  $t - 1$  but has bad underlying health, they are much more likely to be hit by a large negative shock in period  $t$ . In these cases, the ARMA(1,1) model performs the most poorly and significantly overestimates their level of health. This result is in line with previous research that finds that biomarker data can predict future negative health outcomes among ostensibly healthy people (Davillas and Pudney, 2020c). This increased propensity to experience a large negative health shock is an important source of risk to capture in models of health dynamics.

Figure 2.8



## 2.6 Non-linear health dynamics

The complex dynamics of health, as described in the prior sections, can be better understood by adapting the latest panel data techniques. I estimate the health process using a non-linear panel data framework developed by Arellano, Blundell and Bonhomme (2017). This method is from the earnings literature, although it has been applied to a small number of non-earnings contexts, such as non-linear productivity and investment dynamics in firms (Melcangi and Sarpietro, 2024).<sup>6</sup> A major attraction of this method is that it allows for heterogeneity in persistence to depend on the size and direction of the health shock that occurs in period  $t$ . This is not possible to do using the methods used to estimate the ARMA(p,q) models, due to the fundamental endogeneity between the shock in period  $t$  and the persistence estimates that relate health in period  $t - 1$  to health in period  $t$ . This framework also allows persistence estimates to vary by the level of health in period  $t - 1$ , which I previously showed can have a large impact on persistence estimates.

Adapting the Arellano, Blundell and Bonhomme (2017) framework to a health

---

<sup>6</sup>Dal Bianco and Moro (2022) have written a working paper concurrent to this one that also applies this framework to a health context

context produces persistence estimates that range from 0.6 to 1.2. While the linear methods from the prior section produce persistence estimates around the midpoint of these estimates, this range is large enough to have meaningful implications for economic decision making. People faced with a health shock with persistence at the lower end of this range are likely to behave quite differently to those facing a much more persistent health shock. I also document some interesting patterns in how persistence estimates vary depending on whether the shock at period  $t$  is positive or negative, the magnitude of the shock, and the level of health immediately prior to the shock. I find that negative health shocks are more persistent than positive health shocks, and that negative health shocks are more persistent if someone was in poor health prior to the shock. I also estimate an additional model that includes fixed effects as an additional source of heterogeneity. Accounting for fixed effects does reduce the persistence estimates a little, especially for those in poor health who experience large negative shocks. I find some evidence that the size and sign of the fixed effect is correlated with allostatic scores, which helps us understand the variation captured by the fixed effect. I conclude this section by extending this method to better capturing the complex dynamics of other health indicators by estimating the non-linear persistence of an index of mental health.

### 2.6.1 Non-linear persistence estimates of overall health

The non-linear framework of Arellano, Blundell and Bonhomme (2017) models their variable of interest as the sum of a persistent component ( $\eta_{it}$ ) and a transitory innovation ( $\varepsilon_{it}$ ). The linear model estimated in the previous section as also the sum of a permanent component and transitory innovation can be considered a special, highly-restrictive case of Arellano, Blundell and Bonhomme (2017)'s framework. The persistent component is assumed to follow a general first-order Markov process, and so the  $\eta_{it}$  terms are dependent over time, although the nature of their dependence does not need to be specified, allowing for flexible temporal dynamics. The  $\tau$ th conditional quantile ( $\tau \in (0, 1)$ ) of this persistent component, given  $\eta_{it-1}$ , is  $Q_t(\eta_{it-1}, \tau)$ .  $v_{it}$  is then defined as a random process such that :

$$\eta_{it} = Q_t(\eta_{i,t-1}v_{it}), \text{ where } (v_{it}|\eta_{i,t-1}, \eta_{i,t-2} \dots) \sim \text{Uniform}(0,1)$$

. The quantile function maps draws of  $v_{it}$  from a uniform distribution into quantile draws for the persistent component. The transitory component  $\varepsilon_{it}$  is assumed to be mean-zero, independent over time, and independent of  $\eta_{i,t-s}$  for all  $s$ , and is assumed to also include any measurement error. This method allows for general forms of heteroskedasticity, conditional skewness and kurtosis in  $\eta_{it}$ . A caveat to this specification is that it excludes the possibility for the transitory component to follow an MA(1) process, which I do find some evidence for when estimating the baseline models. The  $t$  subscript refers to age. The permanent and transitory components are assumed to be mean-independent of age  $t$ , but the conditional quantile functions and marginal distributions of the transitory component may all depend on  $t$ . Non-linear persistence ( $\rho_t$ ) of the persistent component can then be defined as:

$$\rho_t(\eta_{i,t-1}, \tau) = \frac{\partial Q_t(\eta_{it-1}, \tau)}{\partial \eta}, \quad \rho_t(\tau) = \mathbb{E} \frac{\partial Q_t(\eta_{it-1}, \tau)}{\partial \eta}$$

$\delta Q_t / \delta \eta$  is the partial derivative of  $Q_t$  with respect to its first argument, and the expectation is taken with respect to the distribution of  $\eta_{t-1}$ . This approach estimates persistence as the derivative effect of how much the persistent component of earnings in period  $t$  varies with the persistent component of earnings in period  $t - 1$  when hit with a shock in period  $t$ . I estimate  $\rho_t(\eta_{i,t-1}, \tau)$  of the health process, which is the persistence of  $\eta_{i,t-1}$  when hit by shock  $v_{it}$  with rank  $\tau$ .

A key attraction of this method is that it allows for one shock (such as a very large or small realisation of  $v_{it}$ ) to wipe out the memory of past shocks. This incorporates an important additional source of heterogeneity in health shock persistence that is unavailable in the simple linear models. This allows, for example, a big negative shock in period  $t$ , such as a sudden permanent severe disability, to wipe out the persistence of past shocks. By contrast, the ARMA(p,q) and simple linear additive shock models cannot allow  $\rho$  to vary by any features of the shock that occurs in period  $t$ . Despite its computational complexity, the method is easy to use as Arellano, Blundell and Bonhomme (2017) make available full MATLAB replication files.<sup>7</sup> Furthermore, De Nardi, Fella and Paz-Pardo (2019) propose a simulation-based method

---

<sup>7</sup>All replication files and supplementary material can be downloaded from: <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13795>

to discretize nonlinear and non-normal stochastic processes, so that these estimates can be incorporated into a life-cycle model with minimal state-space cost.

To estimate the model, the quantile functions for  $\varepsilon_{it}$ ,  $\eta_{i1}$  and  $\eta_{it}$  are first parameterised as low order Hermite polynomials. Since the persistent and transitory components of the process are not separately observable, the estimation algorithm begins with an initial guess for the coefficients and then iterates sequentially between draws from the posterior distribution of the latent persistent component and quantile regression estimation until convergence is achieved. The algorithm used is closely related to the stochastic EM algorithm (Diebolt and Celeux, 1993), although the quantile specification of the model avoids the need for a likelihood-based approach to estimation.

I apply this method to estimating the persistence of health, and report the results in Figure 2.9 and Table 2.13 by deciles for the magnitude of the shock at period  $t$  and health decile in period  $t-1$ . Since the health index has been de-measured by age, the health shocks are approximately symmetric, so the lowest decile consists of large negative health shocks, the median decile consists of very small health shocks or unchanged health, and the highest decile consists of very large positive shocks. I report both the persistence estimates for the overall health process, and just the persistent component  $\eta_{it}$ , which strips out the transitory component from the overall estimates. The persistent-component-only estimates are on average higher, with two notable exceptions; large positive shocks experienced by those in poor prior health, and large negative shocks experienced by those in prior good health. Transitory shocks are likely to be more important in these cases.

I find that persistence of health shocks varies greatly, depending on past health, shock size and sign. While the average of my estimates is approximately the estimate of persistence from my baseline models, my non-linear persistence estimates range from 0.6 to 1.2. Such variation has significant implications for economic decision making. Furthermore, there is a large difference in the persistence of positive health shocks and negative health shocks. Large negative health shocks are almost twice as persistent as large positive health shocks. Another notable result is that those in poorer health pre-shock take much longer to recover from a negative shock relative to those in better health pre-shock. Individuals who are both in poor health in

Table 2.13: Non-linear health persistence estimates

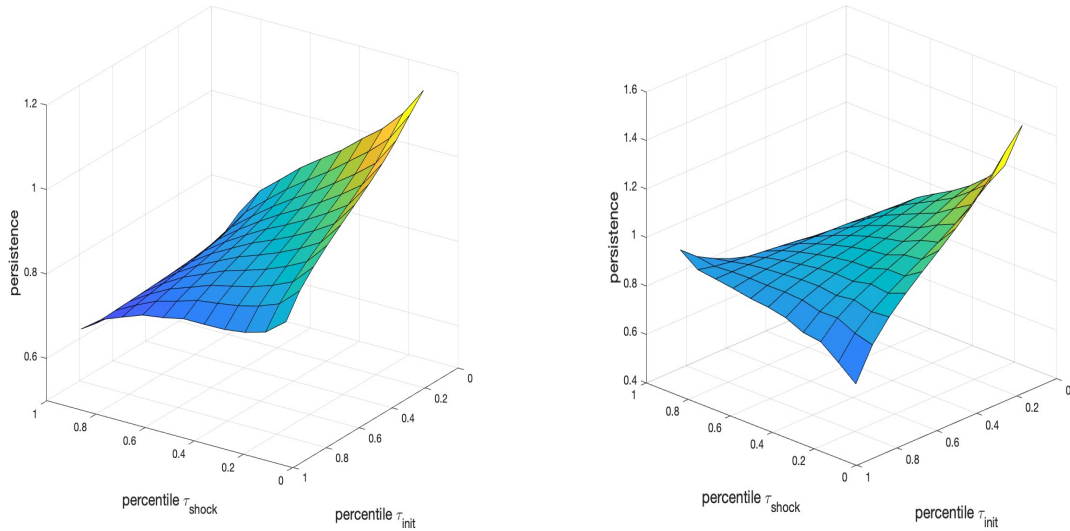
	Shock size percentiles*										
	1	2	3	4	5	6	7	8	9	10	11
$health_{t-1}^*$	Overall health persistence										
1	1.16	1.10	1.05	1.01	0.97	0.94	0.90	0.86	0.82	0.76	0.67
2	1.13	1.06	1.01	0.98	0.94	0.91	0.87	0.83	0.79	0.73	0.65
3	1.08	1.02	0.98	0.94	0.91	0.88	0.85	0.81	0.78	0.72	0.64
4	1.05	0.99	0.95	0.92	0.89	0.86	0.83	0.80	0.77	0.71	0.64
5	1.02	0.96	0.93	0.90	0.87	0.85	0.82	0.79	0.76	0.71	0.64
6	0.99	0.93	0.90	0.88	0.86	0.83	0.81	0.78	0.75	0.71	0.64
7	0.96	0.91	0.88	0.86	0.84	0.82	0.80	0.78	0.74	0.71	0.65
8	0.93	0.88	0.86	0.84	0.82	0.81	0.79	0.77	0.74	0.70	0.65
9	0.90	0.85	0.83	0.82	0.80	0.79	0.78	0.76	0.73	0.70	0.65
10	0.87	0.82	0.80	0.79	0.78	0.78	0.76	0.75	0.73	0.70	0.66
11	0.81	0.77	0.76	0.75	0.75	0.75	0.74	0.74	0.72	0.70	0.66
	Persistent component of health shocks										
1	1.45	1.25	1.17	1.11	1.06	1.02	0.98	0.92	0.86	0.77	0.56
2	1.37	1.21	1.15	1.09	1.05	1.01	0.98	0.93	0.87	0.80	0.62
3	1.28	1.15	1.11	1.06	1.03	1.00	0.96	0.92	0.88	0.81	0.67
4	1.21	1.10	1.07	1.03	1.01	0.98	0.95	0.92	0.88	0.83	0.71
5	1.14	1.05	1.03	1.00	0.99	0.97	0.94	0.91	0.88	0.84	0.74
6	1.08	1.01	1.00	0.98	0.97	0.95	0.93	0.91	0.88	0.84	0.77
7	1.02	0.97	0.97	0.95	0.95	0.94	0.92	0.90	0.88	0.85	0.79
8	0.96	0.93	0.94	0.93	0.93	0.92	0.91	0.90	0.88	0.86	0.82
9	0.89	0.89	0.91	0.90	0.91	0.91	0.90	0.89	0.88	0.87	0.85
10	0.81	0.83	0.86	0.86	0.88	0.89	0.88	0.89	0.88	0.88	0.89
11	0.68	0.74	0.79	0.80	0.84	0.85	0.86	0.87	0.88	0.89	0.95

\*1=most negative, 11=most positive

Figure 2.9: Non-linear persistence estimates

(a) Health ( $\eta_{it} + \varepsilon_{it}$ )

(b) Persistent component only ( $\eta_{it}$ )



period  $t - 1$  and then experience a large negative health shock in period  $t$  have an estimated persistence coefficient of 1 or more, suggesting that a negative shock is likely to be permanent for these individuals. By comparison, an ARMA(p,q) model will underestimate the persistence of a large negative health shock and overestimate the pace and magnitude of recovery, especially for those with poor past health.


## 2.6.2 Fixed effects

These non-linear persistence estimates demonstrate the crucial importance of allowing for heterogeneity in health shock features and health history when estimating persistence. Time-invariant, individual fixed effects are an additional important source of heterogeneity, and not accounting for them may bias the persistence estimates upwards. The literature also emphasises the importance of individual heterogeneity, such as initial conditions from childhood, education, or generic variation, as potentially more important than state dependence in determining health outcomes (Halliday, 2008). I re-estimate persistence allowing for fixed effects by using an extension to the Arellano, Blundell and Bonhomme (2017) framework included in their supplementary appendix. I find that accounting for fixed effects does reduce the persistence estimates, and the magnitude of the reduction varies by past health and shock

magnitude. The reductions are largest for those in prior poor health who experience a large negative health shock. Therefore, the extremely high persistence previously observed for this group partially reflects fixed effects, although the new persistence estimates remain high. Accounting for fixed effects also removes the asymmetry between positive and negative shock persistence.

To capture time-invariant fixed effects, the persistent component  $\eta_{it}$  is now defined as being equal to  $Q_t(\eta_{i,t-1}, \zeta_i, v_{it})$  where  $\zeta_i$  is the fixed effect. I report the new persistence estimates in Figure 2.10 and Table 2.14. The two graphs that make up Figure 2.10 illustrate the same data, but I rotate the plane around the persistence axis to better illustrate the range of the persistence estimates. I report the estimates for the persistent component rather than overall health as estimates of this component are most likely to be overstated by not accounting for fixed effects.

Figure 2.10: Persistent component of health, accounting for fixed effects



The two graphs show the same data, just rotated around the axis for persistence

Several of the key results from the original non-linear persistence estimates are unaffected by accounting for fixed effects. The range of the persistence estimates, depending on past health and characteristics of the shock at period  $t$ , remain large, ranging from 0.3 to 1.0 depending on the features of the shock at period  $t$  and past health. Persistence estimates in cases of negative health shocks continue to be much

Table 2.14: Persistent component of health, accounting for fixed effects

$health_{t-1}^*$	Shock size deciles*										
	1	2	3	4	5	6	7	8	9	10	11
1	1.01	0.99	0.99	0.97	0.92	0.86	0.81	0.72	0.62	0.50	0.34
2	0.96	0.98	0.99	0.99	0.97	0.93	0.90	0.84	0.75	0.66	0.52
3	0.91	0.94	0.96	0.98	0.97	0.95	0.93	0.89	0.82	0.74	0.63
4	0.86	0.90	0.93	0.96	0.96	0.95	0.94	0.91	0.86	0.79	0.70
5	0.81	0.86	0.90	0.93	0.94	0.95	0.95	0.92	0.88	0.83	0.76
6	0.76	0.82	0.86	0.90	0.92	0.93	0.94	0.93	0.90	0.86	0.80
7	0.71	0.78	0.83	0.87	0.89	0.92	0.93	0.93	0.91	0.88	0.84
8	0.66	0.74	0.79	0.83	0.86	0.89	0.91	0.93	0.92	0.90	0.88
9	0.61	0.69	0.74	0.79	0.83	0.87	0.89	0.92	0.92	0.92	0.92
10	0.54	0.62	0.68	0.74	0.78	0.83	0.87	0.90	0.92	0.93	0.97
11	0.41	0.50	0.57	0.63	0.68	0.75	0.80	0.86	0.91	0.95	1.03

\*1=most negative, 11=most positive

higher for individuals in poor health, and the opposite is true for positive health shocks.

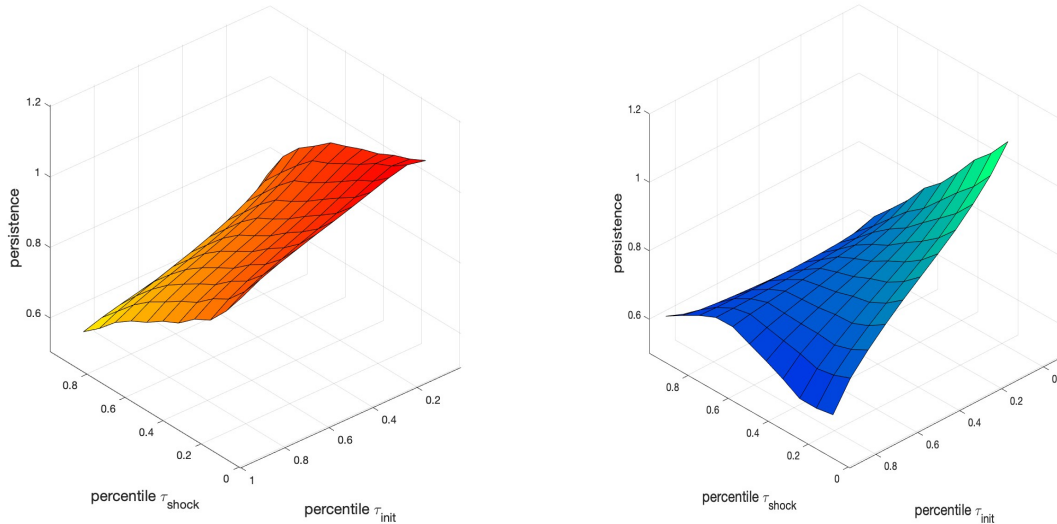
On average, the persistence estimates are smaller when fixed effects are taken into account. This is mostly driven by reductions to the persistence estimates in cases of negative shocks at period  $t$ . The largest reductions are observed for the largest decile of negative shocks, where persistence estimates fall by 0.3-0.4. Accounting for fixed effects has little impact on the persistence estimates when there are positive shocks. As a result, while the original non-linear estimates were much higher for negative shocks than positive shocks, this difference disappears when we take fixed effects into account. Accounting for fixed effects also reduces the persistence estimates for those in very poor health in  $t - 1$  who experience a positive shock in period  $t$ . These results suggest that those who experience large negative shocks, or have a history of poor health, are also more likely to have some unobserved time-invariant trait that subtracts from overall health, such as poor underlying health, and this partially explains the persistently very poor health we observe after large negative shocks and among those with poor health in the prior period. This result is not symmetrical for those who experience positive shocks.

While these fixed effects cannot be observed directly, I do find some evidence that they are related to allostatic scores. Allostatic scores attempt to measure underlying

health, which may be associated with vulnerability to suffer negative health shocks, and propensity and speed of recovery from them. I divide my sample into two groups based on whether allostatic scores are above or below the sample median allostatic score, and then re-estimate persistence for these two groups (Figure 2.11). High allostatic scores indicate poor underlying health while low allostatic scores indicate good underlying health. These estimates are for the entire index, rather than just the persistent sub-component.

Figure 2.11: Non-linear overall persistence estimates; by allostatic score

(a) High allostatic scores (poor underlying health)      (b) Low allostatic scores (good underlying health)



I observe significant differences in the persistence estimates of those with better and worse allostatic scores. Those with worse underlying health experience more persistent negative health shocks and less persistent positive health shocks. The biggest difference between them is that the persistence estimates for those who suffered a large negative health shock but were in good prior health are about 0.5 units lower than for the group with worse allostatic scores. There are two possible reasons why these persistence estimates vary by allostatic score. Allostatic scores may be correlated with the persistence of shocks that people experience. For example, those with worse underlying health may be more vulnerable to highly persistent chronic health conditions. Alternatively, there may be a high correlation between allostatic scores and fixed effects. I find that the persistence differences between the higher and lower

allostatic score groups can be significantly reduced if I use the estimation procedure that takes fixed effects into account. This suggests that these differences mostly reflect fixed effects. I show this in Table 2.15, which reports the difference between the non-linear persistence estimates that takes fixed effects into account for the groups with good and bad allostatic scores. I subtract the estimates for the group with bad (above average) allostatic scores from the group with good (below average) allostatic scores; the difference now only ranges from -0.2 and 0.1.

Table 2.15: Difference between persistent component estimates of poor and good allostatic health groups, accounting for fixed effects

$health_{t-1}^*$	Shock size deciles*										
	1	2	3	4	5	6	7	8	9	10	11
1	0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.0	0.0	0.0	-0.1
2	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1
3	0.1	0.0	0.0	0.0	0.0	0.0	-0.1	-0.1	-0.1	-0.1	-0.1
4	0.1	0.0	-0.1	0.0	0.0	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
5	0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
6	0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
7	0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.2	-0.2	-0.2	-0.1
8	0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.2	-0.2	-0.2	-0.2	-0.1
9	0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.2	-0.2	-0.2	-0.2	-0.1
10	0.1	-0.1	-0.1	-0.1	-0.1	-0.2	-0.2	-0.2	-0.2	-0.2	-0.1
11	0.2	-0.1	-0.1	-0.1	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.1

\*1=most negative, 11=most positive

This result strongly suggests that allostatic scores capture some aspect of fixed effects that are helpful to include when modelling health dynamics. There is significant scope for further research on modelling these fixed effects and identifying whether they relate to, for example, education, early childhood experiences, or genetics.


### 2.6.3 Mental health persistence

This chapter has focussed on modelling overall health. However, the methods used in this chapter can be easily applied to other health indices used in the economics literature, which may have very different persistence profiles. I calculate the persistence of GHQ scores, which are a sub-component of my health index and can be

considered a measure of overall mental health. GHQ (general health questionnaire) is a questionnaire designed to identify non-psychotic and minor psychiatric disorders such as anxiety and depression, and provides a mental health score ranging from 1 to 36. I de-trend the raw GHQ scores from age, gender, and time trends, and first estimate an ARMA(p,q) model as a linear baseline. I find that GHQ scores can be represented as an ARMA(1,1) model in a similar manner to an overall health index, although the level of persistence is lower, with the coefficient on the lagged health term estimated as 0.7. I report these ARMA estimation results in Appendix table A.15.

I then follow the same procedure as above and calculate the non-linear persistence of the mental health index using the Arellano, Blundell and Bonhomme (2017) framework. I report a table of my persistence estimates by past health shock and past health decile in Appendix table A.16 and illustrate the estimates in Figure 2.12. The two graphs show the same data, but I rotate the plane around the persistence axis to better illustrate the range in the persistence estimates. These estimates are for the complete mental health index, rather than just the persistent component.

Figure 2.12: Overall persistence of mental health index



The two graphs show the same data, just rotated around the axis for persistence

In some ways, the mental health persistence graphs resemble the overall health persistence graphs. In both cases, the persistence of shocks varies significantly de-

pending on past health history and magnitude and sign of the shock in period  $t$ , outcomes tend to be worse if the individual is in prior bad health, and negative shocks are more persistent than positive shocks. However, there are some significant differences, which may be important to capture when considering the impact of mental health shocks on economic decision making; the literature on this is very nascent (Jolivet and Postel-Vinay, 2020; Abramson, Boerma and Tsyvinski, 2024). For those with good prior mental health, persistence estimates are around 0.5-0.6, and do not vary much by shock sign or magnitude. However, for those with bad prior mental health, persistence estimates have a huge range. Of note, the persistence of very large negative shocks is over 1, while the persistence of very large positive shocks is around 0.1. This indicates that for individuals already struggling with their mental health, large improvements are highly transient but any further declines are permanent, suggesting a ‘downward spiral of despair’ mechanism and that the capacity for recovery is limited.

## 2.7 Conclusion

This chapter investigated how best to model health as a dynamic process. It evaluated the strengths and weaknesses of the most commonly used approaches in the literature, and adapted recent techniques from the earnings dynamics literature to better capture some of the complexities around modelling shock persistence, frequency and magnitude. It also explored how biomarker data can be used to improve our ability to model health dynamics, although further research in this area is recommended as the increasing availability of genetic and other medical data offers researchers the opportunity to model health in increasingly sophisticated ways.

I conclude with several suggestions to further develop this research. First, as the Arellano, Blundell and Bonhomme (2017) non-linear persistence framework becomes better known, other researchers are suggesting modifications and improvements, which could also be applied to a health context. For example, Almuzara (2020) developed a ‘heterogeneous transitory risk’ (HTR) model that offers a sophisticated way of separately identifying the permanent and transitory components of a shock while also permitting dependence between them. This cannot be done using the Arellano, Blundell and Bonhomme (2017) approach. Health could be an interesting application of this model, as many individuals suffer from multiple health conditions, and capturing interactions between different conditions with different persistence profiles, such as a long-term chronic health condition and shorter-term mental health shock, could further improve our health dynamics modelling.

Second, the focus of this paper has been to improve our ability to statistically predict health dynamics, because doing so helps us understand how health impacts economic decision making. However, I do not consider to what degree my predictions map onto how individuals understand and predict their own health trajectory. There are different ways to characterise this relationship. It could be a process of learning where a series of positive and negative health shocks helps people gradually learn their ‘health type’, or individuals could have stable, long-term biases to be overly pessimistic, optimistic, or broadly correct about their future health outcomes, which has little relationship with their actual health histories. An interesting avenue for future research is to better understand to what degree people modify their expectations of the frequency and persistence of future health shocks following a period

of poor health. Understanding this relationship is important, as the persistence and distribution of health shocks affects both ex-ante choices (how people prepare in advance for a potential health shock) and ex-post choices (how people respond to a realised health shock). For example, a bad health shock may directly impact someone's savings behaviour as they have to stop working for a while and that reduces their income, but it might also affect their savings behaviour by modifying their priors about their future health, which will continue to impact their savings behaviour even when fully recovered. A related improvement could be to separately consider the impact of negative and positive health shocks. Much of the literature focusses on modelling negative health shocks, with little attention paid to recoveries, perhaps because overall health indices heavily feature chronic health conditions and disabilities from which full recovery is unlikely.

A final suggestion is that there is significant scope to improve our understanding of the dynamics of sub-components of overall health, such as mental health. I showed that in some important ways, the statistical properties of mental health dynamics differs from the dynamics of overall health. There are lots of potential avenues for future work on this topic. For example, researchers could build up a richer picture of how mental health dynamics vary by observable characteristics such as age, gender or education level, or the interaction between mental health dynamics and economic events such as unemployment (De Vera, Garcia-Brazales and Lin (2023) is a current working paper on a closely-related topic). My method of modelling mental health is based on a score from a short questionnaire from the psychology literature, which is common practise in the health economics literature but could also be improved, such as by making adjustments for the fact that it is measured as a non-negative integer that is bounded from above and below (Mullahy, 2024), or by incorporating additional data sources, such as high-frequency health information from wearable health technology.

## **HEALTH-DRIVEN OCCUPATIONAL CHANGES**

### **Chapter Abstract**

Poor health impacts labour supply in varied and complex ways. This paper examines an under-explored aspect of this relationship: how suffering a health shock affects occupational mobility. Occupational change commonly occurs after health shocks. Individuals are 10-15 per cent more likely to change occupation or employer in subsequent months relative to those who remain healthy. We document how these newly chosen occupations differ from the occupation mobility patterns of the healthy. Those who newly report a physical disability switch to less cognitive and less manual occupations, those who report worsening mental health switch to less cognitive occupations, and those who report a new chronic health condition switch to less manual occupations, relative to their healthy counterparts. Lower cognitive intensity jobs are jobs with lower overall task complexity, while less manual jobs can be more suitable for those with certain health conditions. Individuals who do not hold a degree and report worsening mental health appear to be particularly vulnerable; we observe the largest declines in overall task intensity for this group.

*JEL classifications:* J24, J62, I10

## 3.1 Introduction

Poor health impacts an individual's labour market trajectory in many different ways. Some are 'demand side' changes imposed by an employer, others are 'supply side' changes where an individual modifies their labour supply to better manage their health condition. We focus on the latter. While there has been significant research on extensive margin labour supply adjustments to poor health, including early retirement (French, 2005) or stopping work to access disability benefits (Low and Pistaferri, 2015), there has been very little attention paid to some intensive margin adjustments, especially occupation mobility. Our paper seeks to fill this gap. We believe that we are the first to analyse the impact of different types of poor health shocks on the likelihood and nature of occupation change.

We report the following findings. First, we show that individuals who suffer a health shock are 10-15 per cent more likely to change occupation or employer in the subsequent twelve months compared to employees who do not suffer a health shock, and they are also more likely to modify hours worked. We define occupational change broadly, and include changes to tasks and responsibilities within an organisation, such as a promotion or a lateral move to a new team. A health shock is defined as a newly reported or newly worsened health condition in our survey data, which we categorise into physical disabilities, mental health conditions, and chronic health conditions.

The welfare implications of this increased propensity to change occupation are unclear, as the relationship between health and occupation choice is complex. To provide a framework to interpret our empirical analysis, we outline a two-stage labour supply model. In the first stage, an individual who is currently working and has suffered a recent health shock identifies which occupation offers them the highest wage. The individual then chooses between switching to this highest-wage occupation, or exiting the labour force and seeking sickness benefits. Occupations are modelled as bundles of cognitive, manual, and interpersonal tasks that vary in intensity, and an individual's task-specific productivity is reduced by health shocks. A crucial component of our framework is that different types of health shocks have different impacts on task-productivity, and therefore wages. For example, an individual who suffers a physical health shock such as becoming paralysed will no longer be able to perform

highly manual tasks, but their injury should have no impact on their ability to perform cognitive tasks. This health shock will more likely lead to occupation change if the individual previously worked in a highly manual job. The existence of alternate occupations that individuals can switch to therefore functions as a form of partial insurance from the wage losses caused by poor health.

The next part of our chapter examines the types of occupations selected by those who recently suffered a health shock. We use GMM estimation to compare the cognitive, manual, and interpersonal intensity of occupations newly selected by those who suffered a physical disability, chronic health shock or mental health shock over the previous six months, compared to the occupation transitions of the healthy. We find that individuals who suffered a physical disability or mental health shock switched to occupations with lower cognitive intensity relative to the healthy. These are typically occupations with fewer responsibilities and lower overall task complexity. Since cognitive intensity is strongly associated with pay levels, this result maps onto these individuals selecting occupations that offer lower wages. We also find that, relative to the healthy, those who suffered a physical disability reduced the manual intensity of their occupation. In theory, the relationship between health shocks and manual task intensity is ambiguous. While jobs with more intense manual requirements may be unsuitable for people with certain disabilities and health conditions, many low-skilled occupations with easy entry conditions also have medium-to-high manual content, and therefore may be chosen by individuals whose health forces them to leave their prior occupation. We do not find any evidence for this latter case. We also obtain some interesting results from our heterogeneity analysis. The decline in occupational cognitive task intensity is concentrated among those who do not hold a university degree. Individuals who do not hold a degree and suffer a worsening of their mental health appear to be particularly vulnerable; we observe the largest declines in cognitive and manual task intensity for this group relative to the healthy.

Finally, we find no relationship between health shocks and subsequent changes to occupation interpersonal intensity. Productivity in performing interpersonal tasks could be unaffected by most health shocks. While we think this is unlikely to be true, especially for mental health shocks, there is very little research on this topic. More highly interpersonal jobs, such as medical practitioners, restaurant owners, teachers,

and social workers, may also offer individuals more opportunities to advocate for themselves for support in the workplace, and individuals who work in highly interpersonal jobs may be more skilled in doing so. While we cannot test these hypotheses using our data, we do find some evidence that flexibility, a key job trait desired by those in poor health (Florisson et al., 2022), is positively correlated with interpersonal intensity, but negatively correlated with cognitive intensity. Better understanding this relationship would be a valuable avenue for further research, especially as there is increasing evidence of the value of high-interpersonal jobs (Deming, 2017), which are increasingly commanding higher wages and are less vulnerable to automation (Autor, 2015), especially for individuals who might otherwise struggle to access high-quality jobs (Aghion et al., 2023).

The public policy relevance of this work is clear. A current major UK public policy objective is to better support those with disabilities and chronic health conditions remain in or re-enter the workforce (HM Treasury, 2023). There is increasing concern over the sharp rise in the number of people in the UK out of work and receiving sickness benefits. As of mid-2024, over 10 per cent of the UK's working age population received at least one health-related benefit, and this share is projected to grow further, with significant government budgetary implications (Ray-Chaudhuri and Waters, 2024). Better supporting older workers, who are more likely to be in poorer health, remain in the workforce will also help reduce the fiscal burden of an ageing population. Despite this clear policy relevance, research on the labour market trajectories of those in poor health who remain in the workforce instead of stopping work is scarce. Indeed, a recent UK government report by the Department of Work and Pensions (Salis et al., 2021) highlighted this lack of literature and called for further research into better understanding labour market transitions of those in poor health. We hope this chapter can contribute to this research gap, although much more work is needed.

The rest of our chapter proceeds as follows. Section 2 summarises the relevant literature, section 3 provides a theoretical framework for understanding the impact of health shocks on occupational choices, section 4 describes our data, focussing on how we model health and occupations, section 5 reports our empirical results and section 6 concludes.

## 3.2 Literature Review

This chapter makes several contributions to the literature. We contribute to the body of reduced form work on the impact of health shocks on various labour market outcomes. Our focus is a significantly under-researched outcome: occupation and employer changes. Much of the existing literature focusses on estimating the extensive margin response (Garcia-Gomez, Jones and Rice, 2008), although papers that investigate the impact of health shocks on earnings (Jolly, 2013; Dobkin et al., 2018; Charles, 2003; García-Gómez et al., 2013) and hours (Bound and Burkhauser, 1999; Gannon and Roberts, 2011) are also common. These labour market outcomes are also increasingly examined in the structural literature, often with the purpose of evaluating health-related public policies, such as retirement ages and the retirement decision (French, 2005; Blau and Gilleskie, 2001; Blundell et al., 2021) disability insurance (Low and Pistaferri, 2015), and means-tested health insurance (Keane, Capatina and Maruyama, 2020). In a similar vein, Bound, Stinebrickner and Waidmann (2010) estimate a dynamic programming model of the retirement decision but also include the option of switching to a ‘bridge job’ that is typically worse paid but may be more flexible or require fewer hours or be lower stress.

To reduce concerns around unobserved heterogeneity, reverse causality and sample selection, the reduced form literature tends to focus on sudden, unpredictable and random shocks that happen to previously healthy people, such as being hit by a car (Dano, 2005; Halla and Zweimüller, 2013), a traumatic injury from playing professional football (Carrieri, Jones and Principe, 2018) or acute sudden health conditions such as strokes, cancer or heart attacks (Jones, Rice and Zantomio, 2020; Tanaka, 2021). The structural literature tends to use a simple index of overall health to reduce computational burden (Brown, 2023). We consider a broader range of health conditions, including mental health conditions, which constitute a significant share of the disease burden affecting labour supply, and also identify differences in the labour supply impact of different types of health shocks.

In general, the literature finds that following a health shock, the likelihood of leaving the labour force increases, and for those who remain working, there is a decline in hours worked and a highly-persistent decline in hourly earnings (Tanaka, 2021). Notably, changing hours and accepting a lower wage, potentially reflecting a decline

in a worker's productivity due to poor health (Grossman, 1972), are adjustments available to workers that may allow them to remain in employment rather than having to stop working. Flexibility in hours and work location, often achieved by switching to self-employment, is an additional adjustment available (Harris, Zhao and Zucchelli, 2021). There is some evidence that without these options, exiting the labour force is more likely following a health shock (Simonetti et al., 2022). This chapter examines to what extent changing occupation and/or employer are additional means of adjustment that can be used by workers who suffer a health shock to remain working.

Our interest in occupational choice is based on the idea that different types of health shocks may impact an individual's productivity in different ways depending on the tasks required for that occupation. For example, a physical disability may have a bigger impact on someone in an occupation requiring manual labour compared to an office job, and so switching to an office job could be a helpful adjustment. Hudomiet et al. (2018) investigate an aspect of this relationship. They examine how different types of health shocks that worsen 'large muscle' physical strength, fine motor skills or episodic memory (a cognitive task) affect workers differentially depending on the demands of their occupation, but obtain mixed results. They do find that those who suffer a worsening in large-muscle physical strength and work in a highly physical job are more likely to stop working, report depressive symptoms, and reduce their self-reported likelihood of working after the age of 65 compared to those working in less physical occupations. However, the authors obtain a null result for the other two traits they investigate: fine-motor skills and cognitive skills. Hudomiet et al. (2018) suggest this could be because jobs with high cognitive or fine-motor skill demands are more likely to be more flexible, or that workers in these jobs have better alternative jobs they could switch to.

One of the very few papers that does model the relationship between health shocks and subsequent job choices is Jolivet and Postel-Vinay (2020). They build a structural life-cycle model with mental health shocks, and show that mental health shocks have a bigger impact on subsequent employment and income if a job is high-stress or the worker faces other non-health adverse labour supply shocks. Much of the adjustment is via the extensive margin, where workers quit their job and enter a potentially lengthy period of unemployment. These workers are then more likely to accept a

new job that is lower stress, although potentially lower paid. There is also some recent research on the impact of poor mental health on economic decision making (Abramson, Boerma and Tsyvinski, 2024). Harris (2019) estimates a dynamic discrete choice model of occupation choice where body weight affects both the distribution of wage offers and the non-monetary costs of participating in each occupation. Both the wage offers and non-monetary costs are a function of occupation-specific job requirements, which include the intensity of mental, physical and social content, to capture that the cost of obesity is higher for some job tasks than others. They find that obesity imposes barriers to occupational mobility; in particular, it is harder to progress a career and switch to professional and managerial occupations with high social content. Biro et al. (2023) explore the wage penalty from not being able to receive wage offers (internal and external) while off sick due to an unexpected accident. Their paper identifies long-term wage losses from short-term absences of 3-12 months, but find that much of the wage penalty is due to missed opportunities to change to a better paying employer, not reduced chances to switch occupation. Finally, there are some papers that consider the related question of the impact of occupation choice on health (Michaud and Wiczer, 2018). While these papers have made some progress in understanding the relationship between task requirements of different jobs, health, and subsequent labour market outcomes, the relationship remains poorly understood.

Outside of the health change context, there is a literature on occupational and job change that this chapter builds on. There are various canonical job and/or occupation change models, which typically situate these labour supply changes in a general equilibrium framework. Such models include island economy models, which are based on the equilibrium search framework of Lucas and Prescott (1974), the search and matching model literature (Mortensen and Pissarides, 1994), extended to occupation choice (Carrillo-Tudela and Visschers, 2023), the Roy (1951) model of worker mobility in the presence of sectoral productivity shocks, and horizontal sorting due to match-specific shock models (McCall, 1990). For example, Busch (2020) extends an island economy framework to occupational switching in the presence of task-specific human capital. Individuals can only imperfectly transfer human capital between occupations, with the amount that can be transferred (a cost of switching) a function of distance in the task space between occupations. While Busch (2020) does not mention health

directly, a major cause of negative productivity shocks that this chapter explores are health shocks. His results also support the findings of Groes, Kircher and Manovskii (2014), who argue that occupational mobility is ‘U shaped and directional’ so that those at the bottom and top of a wage distribution are more likely to switch occupations. In addition, low-wage earners tend to switch to occupations with lower average wages, relative to their former occupation, while the reverse is true for high-wage earners. Since health shocks are highly correlated with low wage work (Benzeval, Davillas and Jones, 2017), this result is consistent with those in poor health swapping to lower-wage occupations. Occupational changes following health shocks could also be consistent with workers having to gather information about what skills they now have. Sanders (2012) models young workers choosing between transitioning to an occupation with similar multi-dimensional (cognitive, manual) skill requirements to their current occupation, or a more dissimilar occupation, as they weigh up furthering skill-specific capital accumulation, and gathering information about their ability. In a similar spirit, Guvenen et al. (2020) model multi-dimensional skill (mis-)match between workers and occupations which reduces as workers learn from occupation switching until their beliefs converge to their true skill portfolio.

### 3.3 Framework

To provide a framework for our empirical analysis, we first outline a model of health and occupation choice in which an individual's productivity and utility cost of work changes after a health shock. We set up a two-stage labour supply decision framework somewhat analogous to a two-stage budgeting model. First, individuals identify the highest-wage occupation available to them, taking into account their task-specific productivity, which is a function of their health. Then, they choose between working in this occupation, or exiting the labour force.<sup>1</sup>

#### Key variables:

**Health:** Individuals have a stock of health  $h$ , which is the sum of three types of health: physical health, mental health, and internal health:  $h = h^p + h^m + h^i$ . These three types map onto our health data that we describe in the next section. Individuals suffer shocks to  $h^p$ ,  $h^m$  or  $h^i$ . Shocks follow a normal distribution. We focus on scenarios where people respond to large negative health shocks, but our framework could also be applied to health improvements.

**Utility:** Individuals have a utility function where the utility cost of work is a function of health. We make the following simplifying assumptions: consumption utility is not a function of health, individuals consume all their wage income ( $w_i$ ) each period and receive no non-wage income so that:  $u(c_i, l(h_i)) = u(w_i, l(h_i))$

**Occupations & Wages:** We model occupations ( $o$ ) as a bundle of three tasks: cognitive, manual, and interpersonal (C,M,I) that we index to  $j$  such that  $j = \{C, M, I\}$ . The wage for each occupation is the sum of three task-specific wages, scaled by that tasks's intensity in that occupation. A crucial component of our wage equation is that different types of health shocks affect an individual's task-specific productivity in different ways. For example, a large fall in physical health may cause a reduction in an individual's manual task productivity, but have no impact on their ability to perform cognitive or interpersonal tasks. This health shock will only have a large impact on wages if the individual works in a highly-manual occupation. This

---

<sup>1</sup>Since we focus on labour supply decisions, we abstract from labour demand factors such as hiring discrimination or increased risk of non-voluntary job separation.

mechanism allows identical health shocks to have very different impacts on wages depending on task composition.

Adapting notation by Bachmann et al. (2022), wages of individual  $i$  in occupation  $o$  take the form:<sup>2</sup>

$$w_{i,o}(z_i, h_i) = \sum_{j=C,M,I} \alpha_o^j \lambda_o^j \theta^j(z_i, h_i^p, h_i^m, h_i^i) \quad (3.1)$$

$\alpha_o^j$  is the intensity of task  $j$  for occupation  $o$ , and ranges between 0-1 for each  $j$ . For example, if  $\alpha_o^C > 0.7$ , this would indicate an occupation where cognitive tasks are important. We source this task intensity data for each occupation from Lise and Postel-Vinay (2020).  $\lambda_o^j$  is the occupation-specific wage per efficiency unit of task  $j$ . On average,  $\lambda_o^C > \lambda_o^I > \lambda_o^M$ , but there is large variation between occupations. We can identify  $\lambda_o^j$  by comparing an individual's wage in different occupations with similar task intensities (as  $\theta^j(\cdot)$  would be held constant and  $\alpha_o^j$  is known).  $\theta^j(z_i, h_i)$  represents an individual's productivity in performing task  $j$ , which is a function of their health and other factors ( $z_i$ ), such as skill in performing occupation tasks, which is not a function of current health. Note that  $\theta^j(\cdot)$  is not occupation specific; it is an individual's productivity at  $j$  tasks in any occupation.

While wages are a crucial component of our framework, we are not able to estimate equation 3.1 in our subsequent empirical work. This is because the available wage data in our LFS dataset is too limited to do so. We discuss this issue further in the data section.

## Two-stage labour supply decision framework

**Stage one:** In the first stage, an individual who is currently in the labour force and has recently suffered a health shock identifies the best occupation for them, which we define as the one that offers the highest wage given their current health and productivity levels.<sup>3</sup> Each period, individuals receive a set of job offers from different

<sup>2</sup>We do not take the logarithm of wages as this would remove the separability of the remuneration from the three different tasks

<sup>3</sup>While we do not allow the utility cost of work to vary by occupation type in our framework, in reality the correlation between the impact of a health shock on the utility cost of work, and productivity in a particular occupation is likely to be high. Therefore, our framework will still be able to accurately predict most occupation changes.

occupations. We denote the alternate occupation offering the highest wage out of the set of offers as  $\hat{o}_i$  and  $\hat{w}_{i_o}$  respectively, and maintain  $o_i$  and  $w_{i_o}$  notation for the individual's current occupation and wage. For given values of an individual's task-specific productivities ( $\theta^j(z_i, h_i)$ ), individuals face a distribution of best alternate occupation wages that we assume follows a normal distribution:  $N(\mu_w, \sigma_w)$ .<sup>4</sup> We assume that the average wage offer at any level of health is lower than an individual's original occupation wage, otherwise individuals would change jobs too frequently. Increases in  $\sigma_w$  will increase the likelihood of individuals changing occupation; the higher the variance of job offers, the more likely that individuals will receive high wage offers.

Individuals will change occupation if  $u(w_{i_o}, l(h_i)) < u(\hat{w}_{i_o}, l(h_i))$ . In other words, given their current productivity level and skills,  $\hat{w}_{i_o}(z_i, h_i) > w_{i_o}(z_i, h_i)$ . Conversely, an individual will remain in their current occupation if the opposite inequality condition holds, and  $u(w_{i_o}, l(h_i)) \geq u(\hat{w}_{i_o}, l(h_i))$ . In this case, an individual may have suffered a reduction in their wage due to their health, but their current wage is still higher than or equal to the best alternate occupation offer. Once the individual has determined the occupation that offers them the highest wage, we designate it as  $o_i^*$  offering wage  $w_{i_o}^*$ .

**Stage two:** In the second stage, individuals decide whether to remain in the labour force following a health shock. If they stop working due to poor health, they receive sickness benefit  $s$ . We assume that sickness benefits are not very generous, and so  $s < w_{i_o}^* \forall i$ . We also make a simplifying assumption that the utility cost of work is a function of health but not occupation. Therefore, for a given best available wage  $w_{i_o}^*$  for individual  $i$ , there exists a 'reservation' level of health  $k$ , such that if  $h_i < k$ , due to some combination of shocks to  $h^p$ ,  $h^m$  or  $h^i$ , an individual's disutility cost of work is too high relative to their consumption utility from their wage income if they work. If  $k = h_i$ , an individual is indifferent between the two options:

$$u_i(w_i^*, l(k) = \text{work}) = u_i(s, l(k) = \text{not work})$$

---

<sup>4</sup>While the overall distribution of wage offers is unlikely to be normally distributed, once we control for an individual's productivity, imposing a normal distribution on the residual components of the wage function is a more reasonable assumption.

We are able to observe some of these choices in our data; we observe whether an individual remains in their original occupation, changes occupation or stops working following a health shock. However, we do not observe an individual's best alternate occupation if they choose to remain in the same occupation or stop working.

## 3.4 Data

Our analysis is performed using a UK longitudinal panel data set. We choose to focus on the UK because of the current heightened policy interest in reversing the recent large increase in the share of the British working age population out of work due to illness and disability (HM Treasury, 2023). Not using US data also allows us to abstract from the complex labour supply incentives around obtaining or preserving health insurance access following a health shock. Our data set is constructed using the Labour Force Survey (LFS), a quarterly data set that contains data on individuals for five consecutive quarters, where the fifth (final) survey is administered one year after the first. We supplement this LFS data with data from the Annual Population Survey (APS), which itself is based on LFS data but also includes additional boosts to achieve better geographic coverage. We use the LFS because it captures both rich labour market data, and data on longer-term health conditions that are most likely to impact labour supply. It is also the largest household study in the UK and therefore provides sufficient cross-sectional variation to study different labour market transitions. Unfortunately, the wages data available from the LFS is limited. While it does include an hourly pay variable, this is only reported for the first and fifth waves (if the individual remains in the sample for the full five waves). In addition, answering this question is voluntary and around 1/3 of the sample chooses not to do so. Non-response rates are higher among lower-skilled, lower-paid occupations. As a result, we restrict our analysis to focussing on the relationship between health changes and occupation changes, rather than the impact of health on wages directly.

We restrict our sample to 2010–2019 due to a change in some survey questions in 2010 and to avoid the impact of Covid-19. We focus on individuals who suffer a health shock while working, therefore we drop all individuals from the sample who report that they do not work for all periods they are observed. Our dataset includes 1,115,013 observations of 302,513 unique individuals who are surveyed 2-5 times.

### 3.4.1 Health data

The LFS reports two types of health data. Individuals can report a ‘health limit’ meaning that health problems affect the kind of paid work they could have done

that quarter. Individuals can also report whether they have any of 16 longer-term health conditions.<sup>5</sup> We make use of both types of health data in our analysis. Table 3.1 reports the 16 conditions, their prevalence by gender, as well as the share of individuals with each health condition who report that their health limits the type of work they can do.

Table 3.1: Summary of health data

Condition	Incidence		Health limits work* share of total sample (%)	
	men	women	men	women
Problems or disabilities connected with:				
(1) arms or hands	3.4	5.2	55.7	66.0
(2) legs or feet	5.6	6.4	51.6	54.2
(3) back or neck	5.2	6.6	56.2	57.2
difficulty in seeing	0.9	0.7	46.6	51.4
difficulty in hearing	1.5	1.1	36.8	40.8
a speech impediment	0.1	0.1	66.0	69.8
severe disfigurement, skin conds., allergies	1.9	2.2	29.7	31.8
chest/breathing problems, asthma, bronchitis	5.1	5.4	27.6	28.2
heart, blood pressure, circulation problems	8.3	5.5	23.9	23.5
stomach, liver, kidney, digestive problems	2.8	3.0	31.9	33.3
diabetes	3.4	2.0	23.0	24.1
depression, bad nerves, anxiety	2.5	4.3	47.0	40.9
epilepsy	0.4	0.4	49.2	45.5
severe or specific learning difficulties	0.6	0.3	64.9	55.4
mental illness, other nervous disorders	0.8	1.1	62.0	52.6
other progressive illness e.g. cancer, MS	0.9	1.2	44.0	52.4
other	2.9	5.5	30.8	27.5

\*Share of individuals with diagnosed condition who report that their health limits their work

The prevalence of health conditions are fairly similar between men and women, with the exception of problems or disabilities connected with arms or hands, legs or feet, back or neck, and depression, bad nerves and anxiety, which are more common among women, and heart conditions and diabetes, which are more common among men. The share of individuals who report that their health condition limits the type of work they can do varies significantly by condition, with individuals with mental

<sup>5</sup>See Labour Force Survey - Volume 3 - Details of LFS variables for further detail. From 2020, an 18th category was included for autism, which we do not include in our analysis.

illnesses, learning difficulties and some rarer conditions being most likely to report this. To make our analysis more tractable we aggregate these 16 health conditions (excluding ‘other’) into three categories which we label as physical disabilities, chronic health conditions, and mental health conditions as described in Table 3.2. These three categories map fairly well onto the first three components of a principal component analysis of the 16 conditions (Appendix B.0.1).

Table 3.2: Classification of health conditions into three categories

Category	Conditions
Physical disability	problems or disabilities connected with: <ul style="list-style-type: none"> <li>- arms or hands</li> <li>- legs or feet</li> <li>- back or neck</li> </ul> difficulty in seeing difficulty in hearing a speech impediment epilepsy
Chronic condition	severe disfigurement, skin cond., allergies, chest/breathing issues, asthma, bronchitis, heart, blood pressure, circulation problems, stomach, liver, kidney, digestive problems, diabetes
Mental health condition	depression, bad nerves, anxiety severe or specific learning difficulties mental illness, other nervous disorders

Our empirical approach is to compare the labour market responses of individuals who are working but then their health worsens, to those who remain healthy. We describe an incident of worsening health between two consecutive survey waves as a ‘health shock’. This approach means we do not consider the labour market behaviour of individuals whose health condition began prior to them entering the LFS survey, and is stable while they remain in the survey panel. This is because we do not observe their labour supply choices when healthy.<sup>6</sup> We identify a ‘health shock’ in the data in two different ways: as a health condition that is newly reported in a later survey wave, or the worsening of a pre-existing health condition. We define the latter case as when an individual with a pre-existing health condition newly reports that their health limits the work they can do. Using diabetes as an example, we would classify

<sup>6</sup>We do not consider the reverse situation of individuals recovering from a prior health shock as most suffer long-term conditions and we do not have a long enough sample size to consider recovery

an individual as having suffered a chronic health shock in period  $t$  if they report having diabetes in period  $t$  but not in period  $t - 1$ , or if they report having diabetes in both period  $t$  and period  $t - 1$  but report that their health limits the work they can do in period  $t$  but not in  $t - 1$ , and do not report any other new health condition in period  $t$ . Around 85-90 per cent of the health shocks in our sample are new health conditions. We report summary statistics for those who suffer a disability, mental and chronic health shock, as well as those who remain healthy in Table 3.3. The correlations between observable traits and health shocks are all as expected, and the sample appears to be well balanced across a wide range of observable variables.

Our chosen method of health shock identification may potentially be vulnerable to justification bias, where people inaccurately report their health to justify their labour market outcomes (Bound, 1991). Individuals who are unemployed or in low-status jobs may be more likely to overstate how bad their health is. We do not think this bias is a major threat to our empirical strategy. The prevalence and magnitude of justification bias remains contested in the literature (Kapteyn, Smith and van Soest, 2011). The strongest evidence for justification bias has been found in cases of unemployment or accessing disability benefits, which typically requires an individual to not work (Black, Johnston and Suziedelyte, 2017). There is much less evidence for justification bias in health reporting to justify occupation or employer changes. Our method of identifying health shocks may also be vulnerable to under-reporting, especially for mental health conditions where individuals may be experiencing symptoms but have not received a diagnosis, or they do not wish to disclose a diagnosis. If this issue is significant, then our estimates are likely to be a lower bound of the impact of mental health shocks on occupation transitions. A related concern is measurement error. Attempts by the literature to estimate the magnitude of measurement error in survey responses to medical questions by comparing them to linked data on hospital admissions has typically found the non-reporting rate of serious health conditions to be surprisingly high (Caraballo et al., 2020). We do not try to adjust for measurement error, which may also bias our results towards zero.

Table 3.3: Summary statistics: percentage shares by category for each health shock\*

	share of total sample			
	split by category and health shock type (%)			
	healthy	physical	mental	chronic
Age				
under 30	20.1	7.8	26.2	12.3
30-39	22.6	12.0	22.3	14.5
40-49	25.4	24.3	25.0	23.4
50-59	21.7	37.5	21.2	34.1
60+	10.3	18.4	5.3	15.7
Sex				
male	50.0	48.3	40.5	48.4
female	50.0	51.7	59.5	51.6
Degree				
degree	66.5	74.9	71.1	70.2
non-degree	33.6	25.1	28.9	29.8
Employment status				
employee	79.1	74.0	74.8	77.3
self-employed	13.3	15.0	9.6	13.1
not working/other	7.6	10.9	15.6	9.6
Hourly pay (2010 prices)				
25th percentile	8.0	7.9	7.2	7.9
median	11.7	11.2	9.8	11.5
75th percentile	14.6	13.9	12.2	14.5
Occupation group				
managers, directors, senior officials	10.8	9.9	7.0	10.4
professional	22.0	18.4	18.4	21.4
associate professional/technical	14.1	13.2	14.2	14.0
administrative/secretarial	11.2	11.2	12.2	11.8
skilled trades	10.1	11.9	7.2	9.1
caring, leisure, other services	8.8	10.7	12.8	9.9
sales/customer service	7.2	7.4	10.9	7.3
process, plant, machine ops.	5.9	6.7	4.3	6.4
elementary	9.9	10.8	13.2	9.8
Hours working				
<10	3.7	3.8	5.1	4.1
10-19	9.2	10.7	14.0	9.7
20-29	12.2	13.9	15.8	13.0
30-39	30.3	29.7	31.6	31.2
40-49	31.2	28.0	23.9	29.0
50-59	8.9	8.8	6.4	8.5
60+	4.6	5.0	3.3	4.5
<i>N</i>	<i>543,649</i>	<i>26,075</i>	<i>10,252</i>	<i>37,581</i>

\*Excludes wave 1 observations as we need two consecutive observations to classify health status

### 3.4.2 Labour market transitions data

We examine two types of labour market transitions, which we label as ‘occupational change’ and ‘employer change’. We define occupational change as whether an individual reports that their job has a different SOC (UK standard occupation classification code) relative to their job in the last sample wave. The LFS data reports 4-level SOC codes, identifying almost 400 separate occupations. The majority of occupation changes we observe constitute small changes in role tasks and responsibilities within an organisation, such as ‘hairdressers and barbers’ to ‘hairdressing and beauty salon management’, from ‘medical practitioner’ to ‘medical radiographer’, from ‘primary and nursery education teaching professionals’ to ‘teaching assistant’. We define employer change using a variable that reports the duration of time an individual has been with their employer. We identify an employer change if an individual is continuously employed in two consecutive waves and reports that his job in the latter wave commenced in the past six months, and at least six months later than the prior job’s reported commencement. This is a conservative approach and may not capture some employer changes, such as if an individual changes employers twice in two consecutive quarters.

Occupational change is much more frequent than employer change, as shown in Table 3.4. Less than 20 per cent of occupation changes observed in the data coincide with change in employer, while over 40 per cent of employer changes also involve some change in occupation. This result is in line with other research. Carrillo-Tudela et al. (2016) find that around half of UK individuals who changed employers also changed occupation or industry. Our analysis excludes individuals who experienced a period of unemployment between jobs, unless that period of unemployment is short so that they are able to report being employed in consecutive survey waves. While occupational change following a period of unemployment is common (Carrillo-Tudela et al., 2016), these labour market outcomes are likely to differ substantially from those who changed occupation or employer without suffering a lengthy period of unemployment (Huckfeldt, 2022). In addition, periods of unemployment may worsen health (Picchio and Ubaldi, 2022), complicating an examination of the relationship between health shocks and labour market transitions. In practise, this represents only around 2 per cent of those in our sample who suffer a health shock; the remaining 91

per cent remain working, and 6 per cent cease working in all subsequent surveys they participate in.

Table 3.4: Occupation and employer change frequency

	Employer change	Employer unchanged	<i>N</i>
Occupation change	14,012	61,934	<i>75,946</i>
Occupation unchanged	20,921	595,560	<i>616,481</i>
<i>N</i>	<i>34,933</i>	<i>657,494</i>	<i>692,427</i>

To make our analysis of almost 76,000 occupation changes more tractable, we model each occupation as a bundle of three key tasks. We adopt a method from Lise and Postel-Vinay (2020), who perform Principal Component Analysis on approximately 200 occupation descriptors from the O\*NET database; a popular US-government-funded database of occupation-specific descriptors. They identify three principal components, which they label as cognitive, interpersonal and manual.<sup>7</sup> Each occupation is assigned a score of between 0 and 1 for each component. Since there are differences in the standard occupation classifications used by US and UK statistical agencies, Blundell et al. (2020*b*) map Lise and Postel-Vinay (2020)’s scores onto UK occupation classifications. These are the data we use for our occupation task intensity scores. Table 3.4.2 provides examples of low, medium, and high cognitive, manual, and interpersonal content occupations. Low, medium, and high classifications correspond to the first, second, and third terciles of the cognitive, interpersonal, and manual score distributions. The distributions are quite similar for the three traits and are approximately normal with a standard deviation of around 0.2 units.<sup>8</sup>

Finally, we map the probability of suffering health shocks onto the distribution of occupations by cognitive, manual, and interpersonal intensity. We graph each occupation as a function of their cognitive, manual, and interpersonal intensity, with the dots shaded according to the share of individuals in that occupation who suffer a health shock. Each dot represents one occupation at the 3-digit SOC level, which equates to around 90 unique occupations. The clustering of dots at certain points of the cube reflects the correlation of cognitive, manual, and interpersonal task intensity in occupations. These plots also illustrate clear differences in the average cognitive,

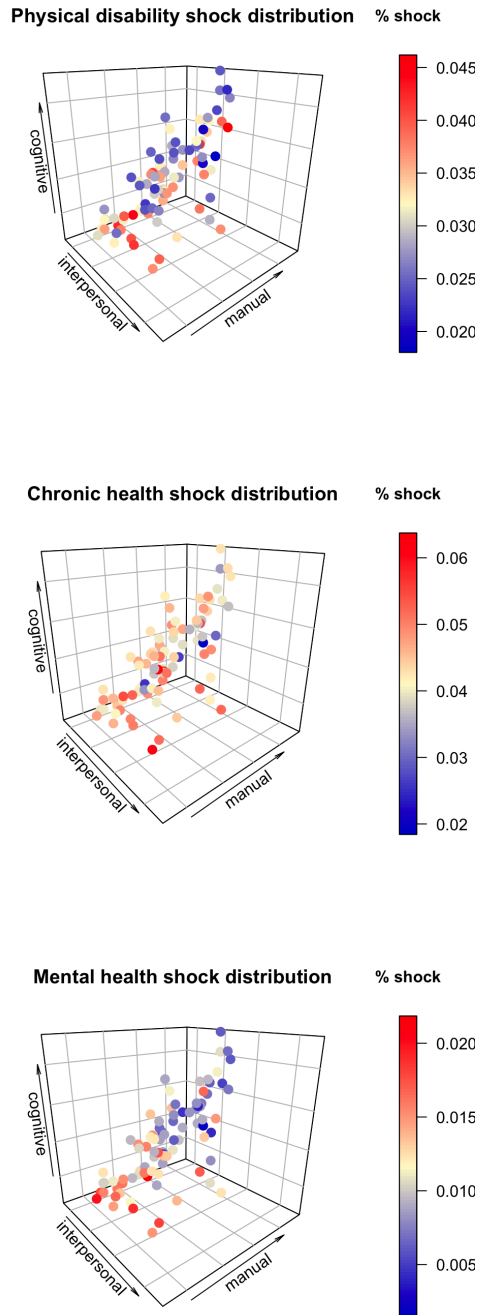
<sup>7</sup>‘Cognitive’ is the component containing the mathematics skills descriptor, ‘manual’ contains the mechanical knowledge descriptor, ‘interpersonal’ contains the social perceptiveness descriptor.

<sup>8</sup>Changes to task intensity are also approximately normal; see Appendix B.0.2 for details.

manual, and interpersonal content of occupations with higher and lower rates of health shocks. Those who suffer a physical disability or mental health shock are more likely to cluster in low-cognitive and low-interpersonal work. There does not appear to be a clear correlation between an occupation's manual content and health shock likelihood. The occupations where individuals are least likely to report a physical health shock are white collar professional jobs, while individuals working in 'elementary cleaning occupations' are the most likely to report a physical disability. Mental health shocks are least likely to be reported by pilots, chief executives and senior officials, senior police and military officers (there is potentially under-reporting by this group).

Mental health shocks are most likely to be reported by caring personal services jobs such as care workers and nursing assistants, and customer services roles such as call centre operators. Chronic health shocks are more similarly distributed across occupations.

Figure 3.1: Share of individuals who suffer a health shock, by occupation



Each dot represents one occupation group, graphed in the cognitive-manual-interpersonal (CMI) intensity space. Clustering patterns indicate common combinations of CMI intensities. Dots shaded according to the share of individuals who suffer a health shock

Table 3.5: Examples of occupations by task intensity

low cognitive occupations			
	low interpersonal	medium interpersonal	high interpersonal
low manual	Bank/post office clerks	Care workers, home carers	Social workers
	Receptionists	Teaching assistants	Welfare/ housing associate professionals
	Waiters and waitresses	Sales and retail assistants	Probation officers
	Retail cashiers/check-out operators	Musicians	youth/community workers
medium manual	Shelf fillers	Bar staff	Prison service officers
	Cleaners and domestics	Pharmacy/dispensing assistants	Undertakers, mortuary/crematorium assistants
	Cooks	Security guards	-
	Postal workers, mail sorters, couriers	beauticians	-
high manual	Farm workers	(none)	(none)
	Bus and coach drivers, van drivers	-	-
	Food, drink, tobacco process operatives	-	-
medium cognitive occupations			
	low interpersonal	medium interpersonal	high interpersonal
low manual	Book-keepers, payroll managers, wages clerks	Authors, writers and translators	Primary/secondary/higher education teachers
	Finance officers	Personal assistants and other secretaries	Clergy
	Financial accounts managers	Journalists, newspaper and periodical editors	Psychologists
	-	School secretaries	Insurance underwriters
medium manual	-	Chefs	Nurses
	-	Office managers/supervisors	National government administrative occupations
	-	Sales supervisors	Physiotherapist
	-	Hairdressers and barbers	Sports coaches, instructors, officials
high manual	carpenters and joiners	Farmers	Restaurant/catering managers/proprietors
	gardeners/landscape gardeners	Medical and dental technicians	Police officers (sergeant and below)
	large goods vehicle drivers	Cleaning and housekeeping managers/supervisors	Veterinarians
	Vehicle technicians, mechanics, electricians	Health and safety officers	Paramedics
high cognitive occupations			
	low interpersonal	medium interpersonal	high interpersonal
low manual	-	chartered and certified accountants	solicitors/barristers/judges
	-	tax experts	HR managers and directors
medium manual	-	business sales executives	medical practitioners
	-	-	finance and investment analysts and advisors
	-	IT project and program managers	property, housing and estate managers
	-	graphic designers	sales account and business development managers
high manual	electricians and electrical fitters	actuarial, economists, statisticians	management consultants and business analysts
	metal working production and maintenance	civil engineers	chief executives and senior officials
	programmers and software development	production managers and directors in manufacturing	biological scientists and biochemists
	laboratory technicians	R&D managers	architects
		chartered surveyors	publicans and managers of licensed premises

## 3.5 Empirical Results

This section describes our key empirical findings. Referring back to our theoretic framework, we proceed backwards and start with evaluating the data that maps onto the second stage of our decision framework: whether an individual stays in the workplace in their best available occupation (either by retaining the same occupation or switching to a different one), or drops out of the labour force. We establish that suffering a health shock increases an individual’s likelihood of both changing their occupation or employer, and stopping work. We then identify how an individual’s best available occupation may differ from their previous one. This is an approximation of the first stage of our decision framework as we cannot observe the best available occupation for individuals who choose to stop working in stage two.

We find that following a health shock, people switch to less complex occupations that have lower task intensity across multiple domains. In particular, we observe declines in cognitive task intensity, which can proxy for overall occupation task complexity. Second, we highlight the importance of modelling health as a multi-dimensional variable when analysing labour market mobility, as individuals suffering different health conditions display different occupation mobility patterns. Individuals who do not hold a degree and suffer a mental health shock appear to be particularly vulnerable; we observe the largest declines in task intensity across multiple domains in this group. Finally, we find no evidence that suffering a health shocks leads individuals to change the interpersonal task intensity of their occupation. This is a puzzling null result, and we consider various explanations.

### 3.5.1 Occupation and employer transition probabilities

We find that individuals who suffered a health shock are one-to-two percentage points more likely to change occupation or employer in the subsequent three-to-six months, depending on regression specification used. In our entire sample, around 11 per cent change occupation and 5 per cent change employer each quarter. Therefore, our findings represent a 10-20 per cent increase in the likelihood of occupation or employer transition among those who recently suffered a health shock. We report estimates from three different estimation strategies: OLS, fixed effect models and

fixed-effect multinomial logits. The latter specification allows us to control for selection bias from some individuals stopping work when they get sick. We estimate the  $\beta$  coefficients of the following equation:<sup>9</sup>

$$\text{transition likelihood}_{i,t} = \beta_1 h_{i,t}^p + \beta_2 h_{i,t}^m + \beta_3 h_{i,t}^i + X_{i,t} + \gamma \text{job traits}_{i,t-1} \quad (3.2)$$

The outcome variable for this regression is a binary indicator of whether an individual transitions job or occupation, and therefore the estimation sample is conditioned on remaining in employment or self-employment.<sup>10</sup> These  $\beta$  coefficients capture the difference in likelihood of occupation and employer change for those who suffered a health shock over the past six months, compared to the healthy. Our OLS and fixed effect (FE) estimates are reported in Table 3.6. Once we strip out time-invariant heterogeneity, those who suffered a physical disability shock are one percentage point more likely to change occupation, and those who suffered a chronic health shock are one percentage point more likely to change occupation or employer. We also observe interesting occupational change likelihood patterns by prior occupation and health shock type. Figure 3.2 illustrates the  $\beta$  coefficients estimated using OLS as reported in Table 3.6, but estimated separately for each pre-shock occupation group. We report the equivalent graph reporting differences in job change likelihood by prior occupation in Appendix figure B.2, and report the raw data of occupation and job change likelihood by prior occupation in Appendix table B.2

Our preferred specification is the fixed-effect multinomial logit, as reported in Table 3.7. We model individuals choosing between stopping working (which does not include going on sick leave) and reporting no longer being employed, changing employer and/or occupation, or remaining working in the same occupation with the same employer. Those who suffer a physical disability shock, mental health shock, or chronic health shock are 18, 12 and 16 per cent more likely respectively to change occupation relative to the healthy.<sup>11</sup> Those who suffer health shocks are similarly more likely to change employer, and both change occupation and employer in the

---

<sup>9</sup> $h_p$ ,  $h_m$  and  $h_i$  represent physical, mental and internal/chronic health shocks as described in our theoretic framework

<sup>10</sup>If an individual switches from working for an employer to being self-employed (or visa versa), we would count that an employer/job change.

<sup>11</sup>For example, if the log-odds coefficient is 0.1673, that equates to  $e^{0.1673} = 1.1821$ , which is an 18 per cent increase in likelihood relative to the baseline choice

Table 3.6: Probability of changing occupation or employer

	occupation change		job change	
	OLS	FE	OLS	FE
physical disability shock	0.0219*** (0.0027)	0.0148*** (0.0032)	0.0075*** (0.0020)	0.0047** (0.0024)
mental health shock	0.0266*** (0.0045)	0.0053 (0.0055)	0.0171*** (0.0036)	0.0057 (0.0044)
chronic health shock	0.0192*** (0.0022)	0.0128*** (0.0027)	0.0175*** (0.0017)	0.0138*** (0.0021)
$R^2$	0.0153	0.0051	0.0090	0.0027
$N$	668,474	668,474	665,296	665,296

Clustered standard errors in parentheses, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$   
health shocks over past 3 months, additional controls: lagged C,M,I intensity,  
lagged hours, age, sex, ethnicity, education, time dummies, pay, region

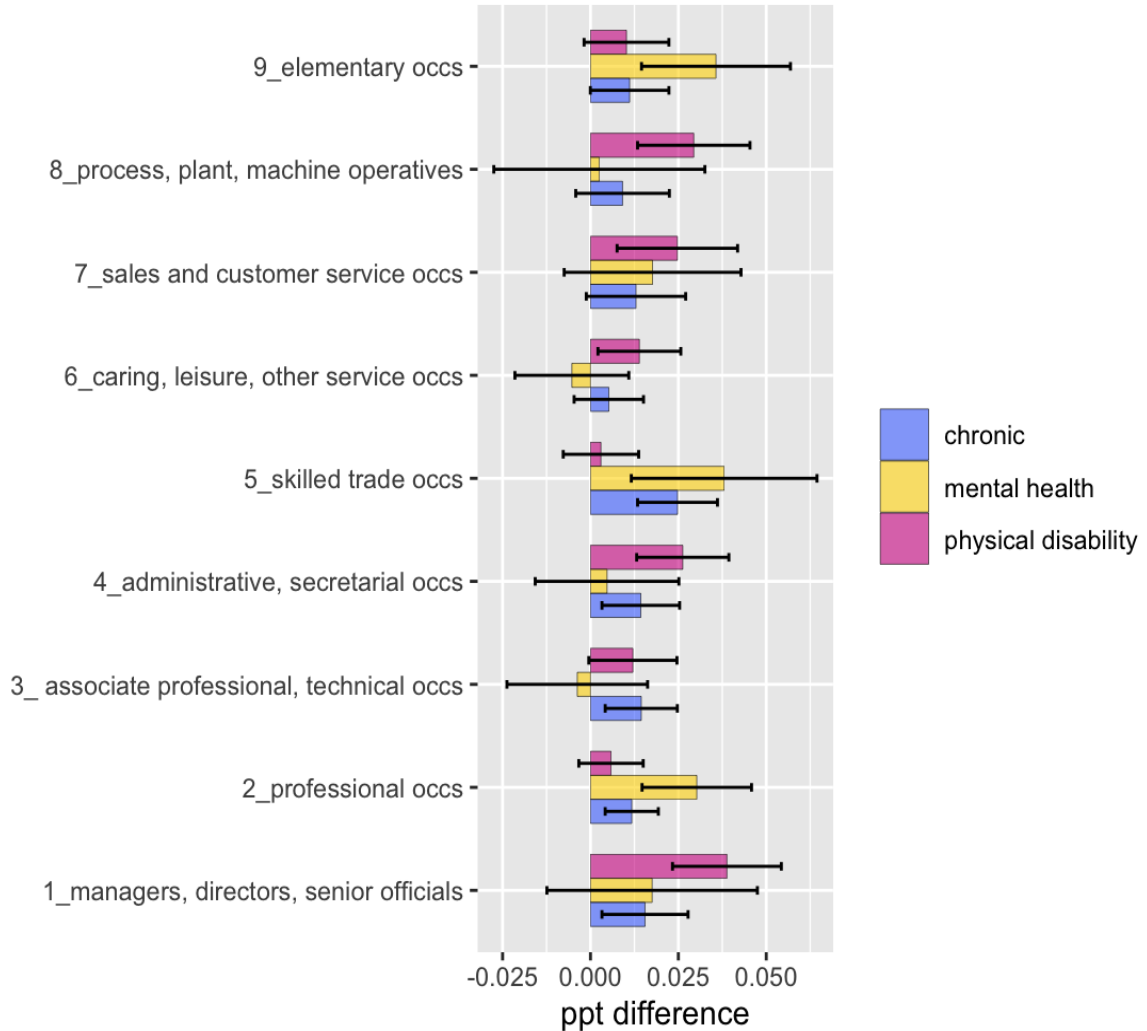
next six months. The most notable difference between these results and our prior OLS and fixed effect results is that the  $\beta$  coefficient on the mental health shock dummy is now significant. This is unsurprising, as those who suffer mental health shocks have the highest probability of stopping work, so not accounting for this likely biases the former estimates downwards.

### 3.5.2 Changes to hours worked

Changing hours have been documented in several papers as another common labour supply response to health shocks (Gannon and Roberts, 2011). While the literature typically finds a negative relationship between suffering a health shock and hours worked, the relationship is theoretically ambiguous as some individuals may increase their hours in an attempt to mitigate the decline in hourly wages associated with poor health. There also may be demand-side effects as employers reduce the hours of unwell employees. The share of individuals who report changes to their weekly working hours over a three month period is high; a little under half of our sample report different average weekly hours in consecutive survey waves. Both increases and decreases in weekly hours worked are more commonly observed among individuals who report a recent health shock relative to those who remain healthy.

We estimate a conditional fixed effect multinomial logit that models individuals choosing between: stopping work, a reduction of five hours worked per week or more,

Figure 3.2: Occupation change likelihood by prior occupation difference from healthy individuals†



† Estimated  $\beta$  coefficients from OLS regression estimates of equation 3.2, estimated separately by occupation group

small changes in hours worked that range from -4 to 4 hours, and an increase of five hours or more. A little under 50 per cent of those who report a change in weekly hours between survey waves report a change of at least five hours per week, so our specification captures most larger changes in hours. We include health shock dummies as regressors, as well as age and previous job traits, and estimate the model separately for men and women, as well as those who worked full-time and part-time in the prior quarter. We estimate the model separately for men and women because

Table 3.7: Fixed effects multinomial logit - log odds ratio

<b>0: baseline (work unchanged)</b>	
<b>1: stop working</b>	
physical disability shock	0.1345* (0.0729)
mental health shock	0.4272*** (0.1014)
chronic health shock	-0.0012 (0.0643)
<b>2: change occupation</b>	
physical disability shock	0.1673*** (0.0354)
mental health shock	0.1128* (0.0579)
chronic health shock	0.1518*** (0.0296)
<b>3: change employer</b>	
physical disability shock	0.0801 (0.0565)
mental health shock	0.2277*** (0.0856)
chronic health shock	0.3086*** (0.0446)
<b>4: change occupation and employer</b>	
physical disability shock	0.1772** (0.0727)
mental health shock	-0.0990 (0.1007)
chronic health shock	0.2506*** (0.0578)
<i>N</i>	198,969
Standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01	
Additional controls: age, age quadratics, hours worked pre-shock	

previous research has shown that hours worked by men and women can respond very differently to shocks (Attanasio et al., 2018). We estimate the model separately by prior full or part time work status to partially control for lagged hours<sup>12</sup> Our specification takes into account time-invariant heterogeneity as well as selection bias from individuals suffering a health shock being more likely to stop working.

Men and women who suffer a health shock are more likely to decrease their hours

<sup>12</sup>We can include lagged hours as an independent variable in the multinomial logit results on occupation transitions reported in Table 3.7 without risking dynamic panel bias, which is not the case here.

Table 3.8: Fixed effect multinomial logit: Changes in weekly hours - log odds ratio

	men		women	
	full time	part time	full time	part time
<b>0: stop work</b>				
physical disability	-0.0510 (0.1392)	0.3659 (0.2376)	0.0481 (0.1778)	0.1491 (0.1268)
mental health	0.7938*** (0.2130)	0.0379 (0.2926)	0.3981* (0.2374)	0.4083** (0.1795)
chronic	0.0521 (0.1260)	0.1102 (0.1868)	-0.1467 (0.1518)	-0.0392 (0.1156)
<b>1: <math>\geq 5</math>hr decline</b>				
physical disability	0.0834* (0.0484)	0.0923 (0.1296)	0.1134* (0.0672)	0.0871 (0.0770)
mental health	0.2858*** (0.0942)	0.0652 (0.2162)	0.0810 (0.1026)	0.0075 (0.1136)
chronic	0.1178*** (0.0397)	0.2067* (0.1103)	0.1222** (0.0550)	0.0427 (0.0667)
<b>2: baseline (hours stable)</b>				
<b>3: <math>\geq 5</math>hr increase</b>				
physical disability	0.2185*** (0.0581)	0.0906 (0.1267)	0.2390*** (0.0847)	0.1642** (0.0786)
mental health	0.1945* (0.1126)	-0.1118 (0.1888)	0.1259 (0.1281)	0.2285** (0.1120)
chronic	0.2379*** (0.0481)	0.0449 (0.1081)	0.2334*** (0.0698)	-0.0079 (0.0663)
<i>N</i>	99,250	17,648	47,112	45,803

Multinomial logit: choices (0,1,2,3) = health shock +  $X_{it}$  + age + pre-shock job traits  
standard errors in parentheses, \* p<0.1, \*\* p<0.05, \*\*\* p<0.01

by at least five, but also increase their hours by at least five, relative to those who remain healthy. These results are fairly broad-based by type of health shock, although men who suffer a mental health shock are most likely to stop working or decrease their hours by at least five. We conclude that while the average response of individuals who suffer a health shock but remain in the labour force is to work fewer hours than comparable individuals who do not suffer a health shock, this result obscures significant heterogeneity across multiple dimensions.<sup>13</sup>

### **3.5.3 Occupation changes**

As well as the increased propensity to change occupations, we also find that the new occupations selected into by those who suffered a health shock differ from those selected into by individuals who remained healthy. Our empirical approach compares the new and old occupations of those who suffered health shocks, using the occupation changes of healthy individuals as a comparison baseline. This strategy requires us to assume that if individuals working in a specific occupation and with specific observable traits who suffered a health shock had instead remained healthy, they would have followed the same occupational mobility patterns as those who remained healthy. We compare occupations by comparing their cognitive, manual, and interpersonal task intensity. We first detail our empirical strategy, before presenting and discussing our key results and performing heterogeneity analysis.

#### **3.5.3.1 Empirical strategy**

We use general method of moments (GMM) estimation to identify whether the cognitive, manual, and interpersonal intensity of occupations newly selected by those who suffered a health shock differ from occupations newly selected by those who did not suffer a health shock, controlling for pre-shock occupation and fixed effects. We use GMM estimation as it allows us to account for unobserved heterogeneity, as well as the cognitive, manual, and interpersonal content of occupations pre-shock without risking dynamic panel bias. We estimate equations that follow the below

---

<sup>13</sup>See Appendix B.0.3 for additional analysis of hours data

form (replacing cognitive with manual or interpersonal as required):

$$\begin{aligned} \text{cognitive}_{i,t} = & \beta_1^p h_{i,t}^p + \beta_1^m h_{i,t}^m + \beta_1^i h_{i,t}^i + \beta_2^p h_{i,t-1}^p + \beta_2^m h_{i,t-1}^m + \beta_2^i h_{i,t-1}^i + \gamma_1 \text{cognitive}_{i,t-1} + \\ & \gamma_2 \text{cognitive}_{i,t-2} + v_1 \text{manual}_{i,t-1} + v_2 \text{interpersonal}_{i,t-1} + X_{it} + \varepsilon_{it} \end{aligned} \quad (3.3)$$

We once again focus on estimating and interpreting the  $\beta$  coefficients, which capture the average impact of suffering a health shock on the cognitive, manual or interpersonal intensity of occupations held in period  $t$ . We include two lags of each type of health shock to account for both health shocks that occur in the same quarter as a potential occupation change and the quarter before a potential occupation change, as well as two lags of the dependent variable and one lag of the other two tasks. We report results from both Arellano-Bond ‘Difference-GMM’ and Blundell-Bond ‘System GMM’ estimation, and additionally report estimates using OLS as a robustness check in Appendix B.0.4.<sup>14</sup> We report two sets of results; one estimated only using the sub-sample who report changing occupations, the other estimated using our entire sample, many of whom do not change occupation. This allows us to directly compare the occupation transitions of the healthy to the sick, as well as observe broader compositional changes.

We use a conservative set of specifications for our GMM estimations: two-step estimator, time dummies, robust standard errors clustered at the individual level, and an unadjusted initial weighting matrix with the Windmeijer correction to correct for finite-sample bias (Windmeijer, 2005). To prevent over-proliferation of the instruments, we collapse the instrument set. All our specifications do not fail to reject the null of the Hansen J test for overidentifying restrictions (Hansen, 1982), a check of instrumental validity. In addition, we believe our results are reasonably robust to the threat of selection bias from those stopping work following a health shock and therefore not being captured in our estimates of occupation change. Our use of difference GMM (although not selection GMM) is robust to some forms of selection

---

<sup>14</sup>The additional initial moment restriction of  $\mathbb{E}(\varepsilon_{it} \text{task}_{i1}) = 0$  that is required for System GMM estimation is not a particularly onerous restriction for the majority of individuals in my data.  $\text{task}_{i1}$  refers to the first occupation held by the individual; for the majority this is many years before the health shock I observe in my dataset. Blundell and Bond (2023) note that this additional moment restriction holds automatically if the data generating process begun long enough before the start of the sample period.

bias (Baltagi et al., 2023), although the differences between our difference and system GMM results are small. The key form of selection bias to which our estimation approach would not be robust, is cases where the lagged dependent variable is part of the selection equation. We run a regression to check whether the lagged occupation intensity variable is correlated with the likelihood of stopping work, and do not find much evidence that this is the case (see Appendix B.0.6 for further details).

### **3.5.3.2 Key results**

We report our key results in Table 3.9. For each of the three tasks, we report four sets of regression results that estimate equation 3.3 by using the sub-sample of individuals who changed occupation or the full sample, and by using difference or system GMM.

Table 3.9: Changes in occupation task intensity

	cognitive				interpersonal				manual			
	occ change subsample		full sample		occ change subsample		full sample		occ change subsample		full sample	
	diff-GMM	sys-GMM	diff-GMM	sys-GMM	diff-GMM	sys-GMM	diff-GMM	sys-GMM	diff-GMM	sys-GMM	diff-GMM	sys-GMM
physical	-0.169**	-0.122***	-0.009*	-0.008*	-0.068	-0.043	-0.002	-0.002	-0.125	-0.165**	-0.007	-0.008
	(0.069)	(0.049)	(0.005)	(0.005)	(0.066)	(0.035)	(0.005)	(0.005)	(0.099)	(0.071)	(0.006)	(0.006)
L.physical	-0.0681	-0.016	-0.007**	-0.007**	-0.036	-0.021	-0.002	-0.002	-0.059	-0.087*	-0.006*	0.005
	0.044	(0.021)	(0.003)	(0.003)	(0.047)	(0.016)	(0.003)	(0.003)	(0.063)	(0.048)	(0.003)	(0.003)
mental	0.0068	-0.138	-0.018**	-0.019**	-0.038	-0.067	0.004	0.003	0.044	0.083	-0.019**	-0.020**
	(0.0833)	(0.129)	(0.009)	(0.009)	(0.085)	(0.075)	(0.008)	(0.008)	(0.078)	(0.065)	(0.009)	(0.009)
L.mental	-0.007	-0.063**	-0.007	-0.008	-0.017	-0.003	0.001	0.001	0.056	0.063	-0.005	-0.006
	(0.080)	(0.028)	(0.005)	(0.006)	(0.057)	(0.023)	(0.005)	(0.005)	(0.065)	(0.059)	(0.005)	(0.005)
chronic	-0.033	-0.027	0.008*	0.009**	-0.011	0.002	0.005	0.005	-0.077*	-0.092**	0.002	0.002
	(0.044)	(0.036)	(0.004)	(0.004)	(0.039)	(0.029)	(0.004)	(0.004)	(0.041)	(0.040)	(0.004)	(0.004)
L.chronic	-0.004	0.006	0.006*	0.006**	0.006	0.014	0.003	0.003	0.028	-0.030	0.002	0.001
	(0.026)	(0.014)	(0.003)	(0.003)	(0.025)	(0.012)	(0.003)	(0.003)	(0.027)	(0.023)	(0.002)	(0.003)
L.cognitive	-0.405***	-0.390***	0.596***	0.696***	0.008	0.019	0.023	0.011	0.053	0.059	-0.012	-0.019
	(0.080)	(0.063)	(0.061)	(0.016)	(0.056)	(0.046)	(0.049)	(0.013)	(0.057)	(0.050)	(0.047)	(0.028)
L2.cognitive	0.175**	0.114***	0.047***	0.067***								
	(0.070)	(0.034)	(0.016)	(0.010)								
L.interpersonal	0.041	-0.052	-0.005	-0.020	-0.446***	-0.521***	0.595***	0.555***	0.017	-0.001	0.007	0.008
	(0.089)	(0.073)	(0.065)	(0.016)	(0.099)	(0.060)	(0.084)	(0.018)	(0.071)	(0.063)	(0.056)	(0.027)
L2.interpersonal					0.136*	0.067**	0.048***	0.060***				
					(0.078)	(0.033)	(0.017)	(0.011)				
L.manual	0.005	-0.053	0.018	-0.013	-0.020	-0.041	-0.054	0.011	-0.292***	-0.483***	0.611***	0.733***
	(0.072)	(0.061)	(0.056)	(0.013)	(0.058)	(0.047)	(0.049)	(0.013)	(0.103)	(0.057)	(0.080)	(0.035)
L2.manual									0.317***	0.091***	0.020	0.048***
									(0.104)	(0.035)	(0.018)	(0.011)
L3.manual									0.236***		-0.004	
									(0.077)		(0.014)	
Hansen J test	0.529	0.399	0.813	0.752	0.605	0.781	0.8985	0.9798	0.728	0.134	0.3601	0.6648
N	13,850	13,850	148,803	148,803	13,850	13,850	148,803	148,803	13,360	13,850	145,107	148,803

Standard errors in parentheses, \* p<0.1, \*\* p<0.05, \*\*\* p<0.01, Include 3rd lag for (9) and (11) to ensure Hansen J test valid

We find that individuals who suffered a physical disability shock switched to occupations with an average cognitive and manual intensity of 0.1-0.2 units lower than the healthy. A one standard deviation reduction in task intensity is equivalent to around 0.2 units, therefore the effect size we observe is substantial. We also find some evidence that those who suffered a mental health shock switched into occupations with less cognitive and manual intensity. These results vary a little by specification, possibly because our mental health shock sub-sample is much smaller than our chronic or physical health shock sub-sample. Our results for chronic health shocks are more mixed; our estimates using the occupation change sub-sample indicate that they switched into occupations with lower manual intensity than the healthy, while our estimates using the full sample suggest that chronic health shocks are associated with switching to more cognitively-intense occupations. We interpret the latter result as a compositional effect. On average, previous research has found that individuals switch to occupations with increased cognitive intensity (Lise and Postel-Vinay, 2020), and we do not find any evidence that those who suffered a chronic health shock behaved differently from the healthy in this respect. Since those who suffer a chronic health shock are more likely to change occupation than the healthy, a combination of these two effects would result in the coefficient on chronic health shocks for the cognitive regression being insignificant when the sample just consists of those who change occupation, but positive and significant when the whole sample is used.

Our most consistent result across our various specifications, is that individuals select into occupations with lower cognitive intensity following a physical or mental health shock, relative to the healthy. Highly cognitive jobs are typically more complex jobs with higher levels of responsibility that command higher salaries. ‘Cognitive’ is also the first principal component identified by the principal component analysis run by Lise and Postel-Vinay (2020) on all job tasks, and can be considered a proxy for overall occupation task intensity. Individuals may switch to less cognitive jobs to seek less-demanding jobs that they can better manage while in poor health. However, the gap between the average cognitive intensity of the new occupations of the healthy and the sick may also reflect the latter group missing out on the ‘upside’ of labour market transitions such as promotions. The reduction in cognitive intensity maps onto individuals selecting occupations that offer lower pay on average. Cognitive intensity is

the task that has the strongest relationship with pay (see Appendix B.0.5 for details). We re-estimate our preferred System-GMM specification using the average wage and standard deviation of an occupation as the dependent variable, but otherwise following the structure of equation 3.3.<sup>15</sup> We find that if an individual switches occupations following a new physical disability or worsening in mental health, the new occupation chosen has lower average pay, as well as lower standard deviation of pay, relative to the new occupations of those who remain healthy.<sup>16</sup>

Ex ante, the relationship between health shocks and manual task intensity is theoretically ambiguous. The onset of a physical disability may have a large impact on an individual's productivity in performing manual tasks ( $\theta^M$  in our model) but not cognitive or interpersonal tasks, so individuals will switch to a less manually-intense job (lower  $\alpha^M$ ) to seek higher wages. On the other hand, individuals who suffer a health shock and are forced to leave their occupation may have to switch to a low-skilled occupation with low entry conditions. These types of occupations often have medium-to-high manual content, such as bus drivers, gardeners, warehouse workers, shelf-fillers, cleaners, and cooks. We do not model the occupation offer distribution, but this outcome could be more likely if a health shock also has a large impact on cognitive or interpersonal productivity ( $\theta^C$  and  $\theta^I$ ). We find strong evidence that those who suffer a physical disability reduce the manual intensity of their occupation relative to the healthy, while our evidence is more mixed for mental and chronic health shocks. This suggests that physical health shocks have the biggest impact on  $\theta^M$ .

Our null result for any relationship between health shocks and subsequent interpersonal task intensity surprised us, as we expected mental health shocks in particular to have a negative impact on an individual's productivity in performing highly-interpersonal tasks. There are several potential explanations for our null result. There may be no relationship between health shocks and an individual's ability to perform in highly interpersonal occupations, perhaps because interpersonal skills are broadly fixed over an individual's working life (Lise and Postel-Vinay, 2020). While we would expect that some mental health conditions such as depression or social anxiety would worsen an individual's interpersonal skills, we cannot find any research specifically

---

<sup>15</sup>Since our wage data is limited, we restrict our wage analysis to considering occupation averages rather than individual wage trajectories.

<sup>16</sup>See Appendix Table B.5 for regression table output

on the relationship between health shocks and interpersonal/social skills. A second possibility is that a more interpersonal job may give individuals more opportunities to advocate for themselves for support in the workplace (Szerman, 2024), and that individuals employed in more interpersonal jobs are also likely to have better interpersonal skills that enable this advocacy to be successful. There is increasing evidence that, relative to other skills, social skills are particularly crucial for labour market success (Noray, 2020).

While we cannot test these hypothesis with our data, we can test a third possibility; that occupations with higher interpersonal task intensity are more likely to have other features desired by those who are in poor health. We find evidence of a positive correlation between interpersonal task intensity and job flexibility, which is a key job trait desired by those in poor health (Florisson et al., 2022).

Table 3.10: Occupation flexibility and task intensity

	OLS			fixed effects		
	hours vary	part time	region match	hours vary	part time	region match
	(1)	(2)	(3)	(4)	(5)	(6)
interpersonal	0.3175*** (0.0050)	0.1103*** (0.0023)	0.0566*** (0.0033)	0.1143*** (0.0207)	0.0222*** (0.0082)	-0.0038 (0.0100)
cognitive	-0.2110*** (0.0048)	-0.1400*** (0.0021)	-0.0375*** (0.0032)	-0.1042*** (0.0200)	-0.0527*** (0.0076)	-0.0126 (0.0095)
manual	0.2365*** (0.0042)	0.0226*** (0.0018)	0.0271*** (0.0028)	0.0742*** (0.0177)	-0.0070 (0.0067)	0.0026 (0.0081)
hours	-0.0087*** (0.0002)	-0.0536*** (0.0001)	-0.0025*** (0.0001)	0.0031*** (0.0005)	-0.0450*** (0.0003)	-0.0013*** (0.0002)
hours <sup>2</sup>	0.0002*** (0.0000)	0.0004*** (0.0000)	0.0000*** (0.0000)	0.0000*** (0.0000)	0.0003*** (0.0000)	0.0000*** (0.0000)
<i>N</i>	667,830	835,667	835,173	667,830	835,667	835,173

Standard errors in parentheses, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Additional controls: age, sex, pay, education, ethnicity

Work ‘flexibility’ can refer to a large range of workplace practises, and we do not

have direct measures of flexibility in our database.<sup>17</sup> However, we do have several variables that could be associated with flexibility: whether hours worked vary week-to-week, which would capture both shift-based work where hours are likely to vary by employer demand as well as jobs where an individual has significant flexibility to set their own weekly working hours, whether an individual works part-time hours, defined as under 35 hours per week, and whether an individual lives in the same region as where they work. The latter variable is a crude attempt to capture commuting time as well as whether the individual works from home. We regress our proxies for flexibility against cognitive, manual, and interpersonal intensity and report our results in Table 3.10. We find that interpersonal intensity is strongly positively correlated with all our indicators of flexibility, while cognitive intensity is negatively correlated and manual intensity is more weakly positively correlated. Therefore, individuals who report that their work is flexible are more likely to be in occupations with high interpersonal content.

### 3.5.3.3 Heterogeneity analysis

We re-estimate equation 3.3 separately for those with a university degree, and those with a high school education or below, using system GMM and the full sample. We find that declines in cognitive intensity following a health shock are concentrated among those who do not hold a university degree. The effect size is particularly strong for individuals who suffered a mental health shock and do not hold a degree; we observe large declines in both average cognitive intensity as well as average manual intensity for this group. The magnitude of these changes suggest that low-educated individuals who suffer a mental health shock are particularly vulnerable to the labour market consequences of health shocks, and further research into their employment outcomes is recommended.

We also replicate this analysis by sex; we report the full results in Appendix table B.6. Men and women respond in a similar way to physical health shocks, but women seem to make larger changes to their occupation task intensity following a mental health shock. In particular, we observe large declines in the average cognitive content

---

<sup>17</sup>A possible extension to this work could be to source flexibility data from elsewhere, such as from recent work by Adams-Prassl et al. (2023)

of occupations held by women who suffer a mental health shock, relative to healthy women.

Table 3.11: System GMM: Changes in occupation content, by education level

	cognitive		interpersonal		manual	
	no degree	degree	no degree	degree	no degree	degree
	(1)	(2)	(3)	(4)	(5)	(6)
physical disability	-0.0095 (0.0059)	-0.0042 (0.0086)	-0.0049 (0.0058)	0.0061 (0.0072)	-0.0047 (0.0069)	-0.0144 (0.0109)
L.physical disability	-0.0064* (0.0035)	-0.0078 (0.0058)	-0.0028 (0.0037)	0.0006 (0.0055)	-0.0028 (0.0040)	-0.0111* (0.0064)
mental	-0.0316*** (0.0118)	0.0025 (0.0129)	0.0083 (0.0103)	-0.0086 (0.0105)	-0.0309*** (0.0109)	0.0009 (0.0153)
L.mental	-0.0146** (0.0073)	0.0051 (0.0076)	0.0033 (0.0066)	-0.0044 (0.0058)	-0.0108* (0.0065)	0.0041 (0.0094)
chronic	0.0096* (0.0058)	0.0081 (0.0066)	0.0067 (0.0046)	0.0020 (0.0066)	0.0015 (0.0050)	0.0052 (0.0063)
L.chronic	0.0054 (0.0035)	0.0068 (0.0044)	0.0026 (0.0031)	0.0025 (0.0044)	0.0007 (0.0031)	0.0032 (0.0043)
Hansen J stat p-value	0.9073	0.7575	0.9909	0.9405	0.8153	0.5880
<i>N</i>	96,025	52,778	96,025	52,778	96,025	52,778

Standard errors in parentheses, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

## 3.6 Conclusion

This chapter sought to address a gap in the literature and investigated how suffering a health shock affects subsequent occupational mobility. We found that those who report a new physical disability, worsening of mental health, or new chronic health condition are more likely to change occupation or employer over the subsequent few months, as well as stop working and drop out of the labour force. We also find that the occupations selected by these individuals differ from the occupation choices of comparable individuals who remain healthy. Those who reported a new physical disability switched to occupations with lower cognitive and manual task intensity, those who reported worsening mental health switched to occupations with lower cognitive intensity, and those who reported a new chronic health condition switched to occupations with lower manual intensity, relative to the occupation choices of those who remained healthy. These results are broadly consistent with individuals who suffer a health shock switching to jobs with fewer responsibilities that may be easier to manage with their health condition, as well as switching to jobs with lesser manual requirements if they suffered a physical or chronic health shock. The reduction in cognitive intensity may also reflect individuals with physical disabilities or poor mental health facing additional barriers to being promoted to better paid, more task-intensive work. Better understanding the occupation mobility patterns of individuals who suffer a health shock but remain in the labour force can contribute to policy work in supporting individuals in poor health remain in or return to the labour force, which is currently a UK government policy priority. In addition, our work highlights the importance of separately considering different health conditions rather than relying on a single health variable or index for analysis, as occupation change patterns differ by health condition.

There are many interesting potential extensions to our research. Our finding that individuals do not seek to reduce the interpersonal content of their job following a health shock surprised us. Further research could unpack this result, especially as prior literature has found a positive relationship between interpersonal task intensity and positive labour market outcomes for those with disabilities or low education levels (Aghion et al., 2023). Identifying how different types of health shocks erode

task-specific skills, for example determining whether mental health shocks erode interpersonal task productivity, could help explain the mechanisms behind the occupation transitions we describe in this paper, as could a further examination of the relationship between flexibility and task intensity for different occupations. A dataset with much better wage data, such as an administrative dataset with linked health data, could be used to identify the wage consequences of different types of occupation changes following health shocks. Finally, such a dataset could be combined with our theoretic framework and empirical results to build a structural model of occupation choices following health shocks that could be used to support government policy making.

## **THE CONSUMPTION CHOICES OF 'GENERATION RENT'**

### **Chapter Abstract**

The large increase in UK house and rental prices over the past thirty years has significantly outpaced income growth and has been accompanied by a large decline in the homeownership rate of young people and a large increase in the share of the budget young renters allocate to housing. It is unclear to what degree this large increase in housing budget shares reflects changing preferences. I estimate housing elasticities of demand for six major expenditure categories using a Quadratic Almost Ideal Demand System (QUAIDS) with UK household expenditure survey data over 1987–2017, and document how housing elasticities varied over time and between lower and higher expenditure households. Despite the large increase in housing and rent costs over this period, both the housing own-price and cross-price compensated elasticities for other goods and services are surprisingly stable for both higher and lower expenditure households. An exception is attributed to the decline in social housing, which is important determinant of spending patterns of lower-income households.

*JEL classifications:* D12, D15, R21

## 4.1 Introduction

In the United Kingdom, real house prices have almost quadrupled over the past 40 years, significantly outpacing real income growth (Miles and Monro, 2019). This has reduced the affordability of housing, particularly for first home buyers, and has also pushed up the cost of renting (Cribb and Simpson, 2018). As a result, there has been a broad-based decline in the homeownership rate of young people, concentrated among middle-income cohorts. The share of young people aged 25–34 with middle-quintile incomes who are owner occupiers fell from 65 per cent in 1995–96 to 27 per cent twenty years later (Muellbauer, 2018). Young people who in prior generations would have transitioned to being owner-occupiers are remaining in the private rental sector as ‘Generation Rent’. These households are also spending much more of their income on rent. From the late 1980s to the late 2010s, the budget share that young renters aged 18-35 allocated to housing increased by almost 15 percentage points. As a result, the budget share young renters allocated to most other consumption goods declined.

In the media, ‘Generation Rent’ is popularly depicted as being forced to spend an increasing amount of their income on rent due to rising rental costs, stagnating wages and a shortage of social housing (Joyce, Mitchell and Keiller, 2017), causing increased hardship for this cohort (Broome et al., 2023). However, the impact of house prices on the consumption of renters is complex. First, long-term increases in house prices will typically translate into higher rents. Between 1988–2017, house prices in the UK increased by 400 per cent, private sector housing rents by 300 per cent, and overall CPIH (a broad measure of the consumer price index that includes housing) by a little over 200 per cent.<sup>1</sup> While there was significant regional variation, house price increases outstripped CPIH increases throughout the UK. Changes to the relative price of one good will impact both total consumption and how an individual allocates their budget between different types of consumption goods. Second, there are dynamic effects from the interaction of house prices and tenure choices. Over 70 per cent of young renters in England aspire to become homeowners (Department for Levelling Up, 2023). These households could respond to an increase in house prices by

---

<sup>1</sup>House price data is from Nationwide House Price Index, private sector housing rents is a sub-component of the CPIH series produced by the UK’s Office for National Statistics (ONS)

reducing their consumption to afford the larger housing down-payment now required. Since housing expenditure is typically more difficult to reduce than other consumption items such as restaurants or leisure goods and services (Chetty and Szeidl, 2016), this would mechanically cause the housing budget share to increase. Conversely, some young renter households may give up or delay their homeownership aspirations in response to rising house prices, and therefore stop saving and increase their housing or non-housing consumption. Increasing house purchase down-payment requirements and credit constraints exacerbate these effects.

This chapter focusses on interpreting the impact of housing price increases on the housing and non-housing budget shares of young renters, which I define as households with a head aged under 35 years old, in England, Scotland and Wales between 1987–2018.<sup>2</sup> The large increase in the housing budget share is consistent with a range of different preferences. I am particularly interesting in determining whether the large increase in housing prices is correlated with the demand by young renters for non-housing goods and services becoming more sensitive to housing prices, and if there are differences between higher and lower expenditure households. In addition, there were large changes to the prices of other non-housing goods and services over the sample period, such as large falls in real food prices, which will also impact budget shares. To disentangle these various factors, I estimate a Quadratic Almost Ideal Demand System (QUAIDS) (Banks, Blundell and Lewbel, 1997), which is a variant of the Almost Ideal Demand System (Deaton and Muellbauer, 1980*a*). I estimate a six equation demand system for food, housing, fuel & light, leisure, other goods, and transport & other services, using household expenditure data from the UK’s Living Costs and Food Survey and its predecessor surveys, which I use to calculate compensated (Hick-sian) cross-price and income elasticities for housing. I use these housing elasticities to identify the impact of housing price changes on how young renter households allocate their spending between housing and non-housing goods, taking into account changes to total expenditure as well as price changes for all goods in my demand system. I focus on identifying how these elasticities varied over my three-decade sample period

---

<sup>2</sup>Throughout this paper, ‘housing prices’ refers to the buildings people live in, and typically means rented accommodation prices as my sample are renters, while ‘house prices’ is the purchase price of houses (not flats) and typically refers to a hedonic house price index.

as rental and house prices rose above CPIH, and how they differed between lower and higher expenditure households.

I obtain the following results. First, the median compensated housing own-price elasticity became less negative over my sample period as rental households reduced how much they spent on rent by a smaller amount as house prices rose. This reduction was driven by low-expenditure young renter households, and matches a sharp decline in the share of this group living in social housing (Waters and Wernham, 2023). The increasing lack of affordable housing would prevent these households from responding to housing price increases by seeking cheaper rent elsewhere. The compensated housing own-price elasticity for median-expenditure households was surprisingly stable over the sample period, despite the large increase in house prices and decline in the homeownership rate of this group. Second, I find little evidence that non-housing consumption of young renters, and especially consumption of low-expenditure young renters, has become more sensitive to housing price changes over the sample period, including during periods where the gap between housing price and income growth was the largest. Food cross housing-price elasticity was broadly unchanged, especially for lower-expenditure households. Leisure goods, leisure services and catering would be the easiest expenditure category to reduce if an individual was experiencing financial stress, but the leisure cross housing-price elasticity was positive and only declined a little during the first decade of the sample. I conclude that the consumption preferences of young renter households were surprisingly robust to the large increase in house prices, rent prices and decline in homeownership rates.

I then aggregate all these results into an overall measure of welfare. I calculate the compensating variation of all price changes experienced by young rental households in my sample. Compensating variation is the amount of additional money a household would need to return to their original utility following these price changes. I find that while housing has become much more expensive over my sample period, the price of most other goods and services declined relative to the disposable income growth of young households, and so young renter households in 2018 are better off than those in 1987. However, the purchasing power of young households has stagnated following the Great Recession due to weak income growth relative to price growth, made worse by households allocating an increasing share of their budget towards relatively high-

price-inflation items such as housing and services. Relative to the late 1980s, lower-expenditure young renter households in the late-2010s are particularly better off as social housing partially shielded many of them from the large rental price increases in the private rental market, and they also benefited from large declines in food prices. Food makes up a much larger share of their budget compared to higher-expenditure households.

There is a public policy interest in better understanding how the large increase in house prices has impacted young renters, as well as to what degree the large increase in housing budget shares of this cohort reflects financial stress. Young renters make up 11 per cent of the working age sample, and understanding the biggest change to their consumption bundle in recent decades is helpful for designing tax regimes and government support programmes. Over the sample period, government housing support for young renters in the form of housing benefits was narrowly targeted at those in the bottom expenditure quintile, and this has become much less generous, especially since 2011 (Waters and Wernham, 2023). There has also been a series of policies by the UK government aimed at increasing the rate of youth homeownership, including reductions in stamp duty, Help to Buy, and specific ISAs. Many of these were only introduced in the latter years of the sample, although a shared homeownership programme was introduced in 1980. A better understanding of the take-up of these programmes over the expenditure distribution of young renters is beyond the scope of this work, but would be a very helpful complement.

The rest of the chapter proceeds as follows. Section two summarises the relevant literature. I then introduce my dataset and describe the major trends in expenditure budget shares, housing prices, and housing tenure of young renters in section three. Section four reports my QUAIDS demand system elasticity estimates, followed by a section of welfare analysis, and section six concludes.

## 4.2 Literature Review

Housing is a vast topic in the microeconomic and macroeconomic literature. Within it, there are two main areas to which this chapter contributes. There is a literature trying to explain why house prices and consumption are so highly correlated in the macroeconomic data. Much of this literature focusses on the consumption of homeowners and mortgagors, and much of it focusses on total expenditure rather than identifying differences by type of consumption good or service. This chapter contributes to this literature by focussing on the consumption response of renters, and separately considering expenditure budget shares of different goods and services. This has previously received very little research attention despite worsening housing affordability and increasing cost of living concerns for young renters (Waters and Wernham, 2023). There is also a sizeable literature on modelling and estimating housing demand, including calculating price and income elasticities. Once again, much of this literature focusses on homeowners, and does not incorporate housing demand into a broader demand system framework, which this chapter does. This enables me to evaluate the welfare implications of house price increases on renters.

### 4.2.1 Consumption and house price co-movement

Several channels have been proposed to explain the strong positive relationship between house price movements and consumption observed in the aggregate data, which can inform us about how the consumption of renters is likely to respond to house price increases. These channels include the collateral channel, which is the increased borrowing capacity of credit-constrained homeowners when the value of their home rises (Mian, Rao and Sufi, 2013), substitution effects, income effects due to changes in future implicit rental costs, and endowment income effects (wealth channel) from the re-evaluation of the housing endowment (Berger et al., 2017). In addition, since house prices are equilibrium objects, shocks that affect house prices may also directly impact consumption, so some of the co-movement may simply reflect a common omitted factor such as expected future income or productivity growth (Attanasio et al., 2009). The relative importance of these channels in the UK during my period of interest remains contested. Campbell and Cocco (2007) and Attanasio et al. (2009) used

the same UK data but came to opposite conclusions on this question. These different channels predict different consumption and saving responses by young renters to house prices increases. If much of the house prices and consumption co-movement in the UK in my period was driven by wealth or collateral channels, the response by renters should be minimal. However, if the correlation is driven by a common factor, then an increase in house prices should have a larger (and positive) effect on all young households, both homeowners and renters, who consumption smooth by increasing their current consumption. The majority of empirical papers on this topic have found little to no consumption response by young renters to house price changes (Campbell and Cocco, 2007; Bijlsma and Mocking, 2017; Rouwendal and Alessie, 2002; Browning, Gørtz and Leth-Petersen, 2013; Zhang, 2019), but there are exceptions. Attanasio et al. (2009) and Disney, Gathergood and Henley (2010) both find a positive response by young renters to price increases, and Guiso, Paiella and Visco (2005) find a negative effect in Italy in the 1990s and early 2000s.

In addition, the literature documents several highly-relevant consumption response heterogeneities, including asymmetric responses to house price increases and falls (Lee, 2023), and large differences in responses by age. Responses may also be non-linear by size of house-price shock, as components of non-housing consumption such as food and leisure may be easier to adjust than consumption ‘commitments’ such as housing (Chetty and Szeidl, 2016). Khorunzhina (2020) finds some evidence of younger households exhibiting stronger non-separability between non-durable consumption and housing, but emphasised that large gaps in the literature remain and called for more research in this area.

### **4.2.2 Housing demand**

There have been many attempts over the decades to estimate price and income elasticities for housing; Polinsky and Ellwood (1979); Ermisch, Findlay and Gibb (1996); Harmon (1988); Zabel (2004); Røed Larsen (2014) are a few examples. Estimates vary hugely, depending on the empirical method used, and once again, the majority of these papers focus on homeowners and mortgagors. A paper similar in motivation to this one is Albouy, Ehrlich and Liu (2016), who produce temporal and spatial elasticity estimates of housing demand by renters to better understand the

large housing budget share increase by renters in the US. They find that housing demand is income and price inelastic, and also find some evidence of economies of scale in housing. More recently, attention has shifted to estimating housing supply elasticities, motivated by supply constraint issues, especially in parts of the US (Baum-Snow and Han, 2024; Saiz, 2010).

As well as changing the amount of housing consumed, house price changes will also affect consumption by modifying tenure decisions. While some young renters with homeownership aspirations may respond to a sustained increase in house prices by decreasing their consumption to save for the now larger house deposit required, others may delay or give up their homeownership ambitions. This ‘discouragement effect’ should increase consumption (Engelhardt, 1994). Alternatively, if credit constraints are non-binding, then people living in locations with higher (positive) house price growth may be incentivised to switch from renting to homeownership at a younger age to protect themselves from future house price increases (Banks et al., 2016), or switch to saving to buy a flat instead of a house (Attanasio et al., 2012). We can observe this consumption response heterogeneity in some older empirical literature that observes how house prices affect the use of savings products for a housing down-payment (Yoshikawa and Ohtaka, 1989; Engelhardt, 1994; Sheiner, 1995). These papers observe significant heterogeneity among households; while some households dropped out of using the savings product when prices increased, those who continued to save using the product increased their rate of saving. This heterogeneity is also clearly demonstrated in a life-cycle model by Li et al. (2016) that jointly identifies intra-temporal and inter-temporal preferences over non-housing consumption and housing consumption. When the intra-temporal elasticity of substitution (ES) is low, a large and persistent house price appreciation coupled with an increase in income leads to a significant increase in homeownership rates, alongside an increase in non-housing consumption. However, when ES is high, an increase in house price growth will encourage renters to substitute to non-housing consumption, and homeownership rates respond very little. There is very little empirical consensus on the magnitude of this ES coefficient, or how it has changed over time or differs between households.<sup>3</sup>

---

<sup>3</sup>ES is defined as substitution across goods within a given time period holding real income constant, and is closely related to Hicksian cross-price demand elasticity which is the focus of this chapter, although calculated differently.

### 4.3 Data and Summary Statistics

To estimate a demand system for young renters I use several cross-sectional datasets of household expenditure in the United Kingdom: the Living Costs and Food Survey (LCFS) which began in 2001, and its predecessor surveys, the Family Expenditure Survey (FES) and National Food Survey (NFS). To ensure consistency between pre and post-2001 data, I use a dataset produced by the Institute for Fiscal Studies (IFS) that cleans these cross-sectional datasets and constructs aggregated consumption categories using the LCFS data to make them consistent with the prior FES expenditure categories (Oldfield et al., 2019). The LCFS expenditure data are classified into 12 consumption categories based on COICOP guidelines, which is a standardised UN classification, while FES data reports 14 main expenditure categories. The 14 FES categories cannot be easily mapped onto the 12 LCFS categories; such mapping needs to be done with highly-disaggregated data, which is what the IFS dataset does. The IFS dataset also includes harmonized household demographic information such as age of household head, region, number of adults, and number of children. I merge these with some additional data directly from the LCFS and FES that is not included in the IFS dataset. My dataset consists of 171,499 individuals and covers 1987–2018 Q1. Of these, 14,747 are young (18–35) renter households, which are the focus of my analysis. The IFS dataset ends in 2018 Q1 and I do not extend it to 2024 to avoid the volatility of 2020–21, and because my dataset ends at approximately the trough of the UK youth homeownership rate, which has since increased a little (Cribb, 2024). I summarise the key features of the data by age and homeowner status in in Appendix table C.1.

My price index data is also from an IFS data release that aggregates retail price indices (RPI) into categories to match the 14 FES expenditure categories, and are also cleaned to be more consistent over time.<sup>4</sup> I construct my own index for housing prices using CPIH (Consumer Prices Index including owner occupiers’ housing costs) data, as the IFS-constructed housing RPI is based on actual housing expenditure by renters and homeowners as reported in the expenditure survey, and therefore includes changes in mortgage costs rather than imputed rent for owner-occupiers. This results

---

<sup>4</sup>See Price Data User Guide (Zoe Oldfield, 2022) for details: <https://ifs.org.uk/internal-datasets/price-data-accompany-ifs-derived-fesefslcfs-variables>

in a more volatile index that moves with interest rate changes, which is not relevant for my analysis of how renters respond to housing price changes. I describe in more detail how I construct my housing index in a following sub-section.

### 4.3.1 Expenditure budget shares

Table 4.1 reports the average expenditure budget shares for young households with a household head aged 18-35 who rent, are mortgagors, or own outright, as well for the full sample.

Table 4.1: Budget shares by age group\*  
1987–2018

FES categories	Young (18-34)			Full sample		
	rent	mortgagor	own	rent	mortgagor	own
food	0.162	0.114	0.144	0.196	0.134	0.185
catering	0.058	0.060	0.068	0.047	0.055	0.048
alcoholic drinks	0.041	0.044	0.045	0.039	0.042	0.040
tobacco	0.039	0.014	0.018	0.040	0.015	0.013
housing	0.229	0.265	0.152	0.222	0.230	0.145
fuel, light & power	0.074	0.041	0.054	0.083	0.046	0.074
household services	0.055	0.062	0.074	0.058	0.057	0.067
clothing & footwear	0.061	0.061	0.069	0.052	0.059	0.051
personal goods & services	0.038	0.038	0.045	0.037	0.040	0.050
motoring	0.094	0.151	0.160	0.082	0.150	0.138
fares and other travel	0.034	0.023	0.027	0.029	0.023	0.018
leisure goods	0.044	0.046	0.049	0.043	0.046	0.049
leisure services	0.070	0.082	0.095	0.071	0.103	0.122
<i>N</i>	14,771	18,129	832	49,666	72,257	52,333

\*Ages based on household head, excludes Northern Ireland, excludes household goods

#### 4.3.1.1 Censored observations

A significant fraction of households consume zero goods from one or more of the 14 expenditure categories. Since zero observations are implicitly non-positive consumption choices that are censored to be non-negative, not accounting for them may bias demand coefficient estimates. The literature proposes various solutions to this issue, including constructing virtual prices which would induce households to consume exactly zero of that commodity (Pitt and Lee, 1986), modified two-stage Heckman

selection models (Heien and Wessells, 1990) and simulated method of moments estimators (Amano-Patiño, 2019). Unlike much of this literature, my setting is quite straightforward as I only care about large expenditure categories rather than more granular categories where zero budget shares are much more common. In addition, since I focus on the housing choices of young renters, the budget share of greatest interest is never zero by construction. Therefore, a simple approach to this issue suffices for my purposes. I construct three amalgamated variables: leisure, other goods, and transport and other services, which significantly reduces the number of censored observations. I then remove households with zero or negative expenditure in any of the six budget shares from my sample. This removes seven per cent of households in the sample, the large majority of which are older households. I also exclude durable (household) goods such as furniture and electrical appliances from my subsequent analysis because they provide a consumption flow over many time periods. Estimating a demand system with non-durable consumption expenditure is consistent with assuming the utility function has weak separability between the total consumption of durable and non-durable goods (Deaton and Muellbauer, 1980*b*), and is common practise in the literature. I report these amalgamated consumption variables that I use for my demand estimation in Table 4.2.

Table 4.2: Amalgamated consumption categories

New categories	Components	Share zero observations (%)
Food	food	0.42
Housing	housing	1.07
Fuel & light	fuel & light	2.93
Leisure	leisure goods, leisure services, catering	0.39
Other goods	clothing & footwear, personal goods & services, tobacco, alcohol	2.73
Transport & other services	household services, fares, motoring	1.07

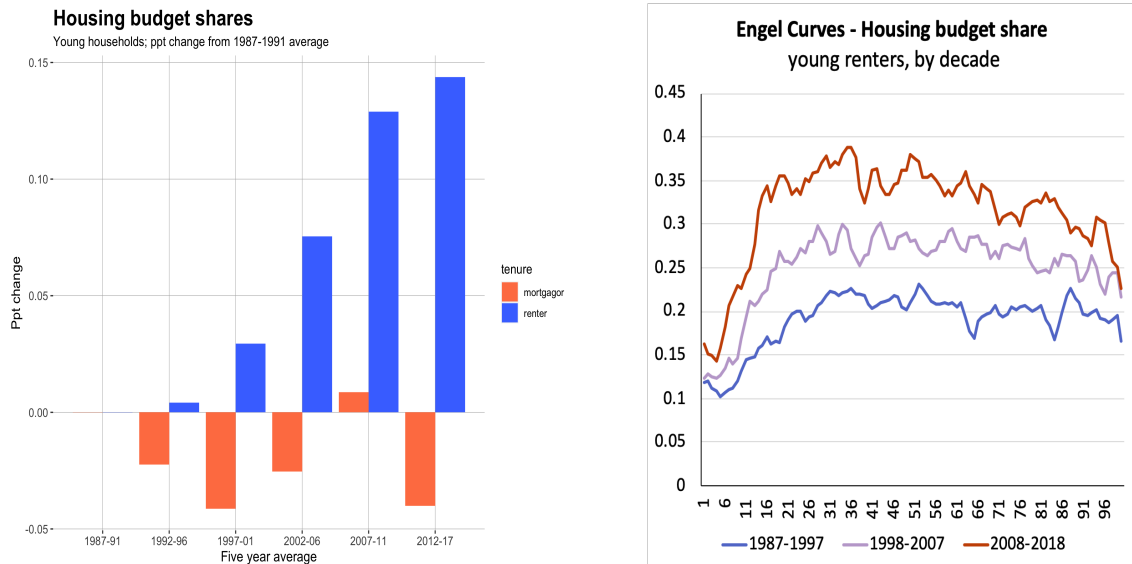
I now describe some trends in budget shares over my sample period. While changes in budget shares reflect a combination of relative price changes and preference changes that I am better able to disentangle in my subsequent demand system estimation, I document some interesting patterns in how the budget shares of young renters changed over time as house prices increased and renters allocated an increasing budget share

to rent. I also construct Engel curve graphs to illustrate how budget shares vary with total expenditure, which I use as a proxy for total income levels.

#### 4.3.1.2 Housing budget shares

On average, housing made up 23 per cent of the non-durable expenditure of young (18–35) renter households. These budget shares sharply increased over the sample period. As shown in the left panel of Figure 4.1, the budget share allocated to housing by young renters increased by almost 15 percentage points between 1987–91 and 2012–17, while it declined a little for mortgagors over this period. The increase was largest for lower-expenditure young renter households. I construct housing Engel curves with the housing budget share on the y-axis and the total expenditure percentile on the x-axis.<sup>5</sup> The Engel curve for housing shifted up across the expenditure distribution between 1987–2018 as people spent a higher share of their total budget on rent, with the largest increases occurring for lower expenditure renter households at around the 20-40<sup>th</sup> expenditure percentile. This also caused the Engel curve to change from being flat prior to 2007 to downward sloping above the 35<sup>th</sup> percentile of expenditure for the 2008–2018 period, as shown in the right panel of Figure 4.1.

Figure 4.1



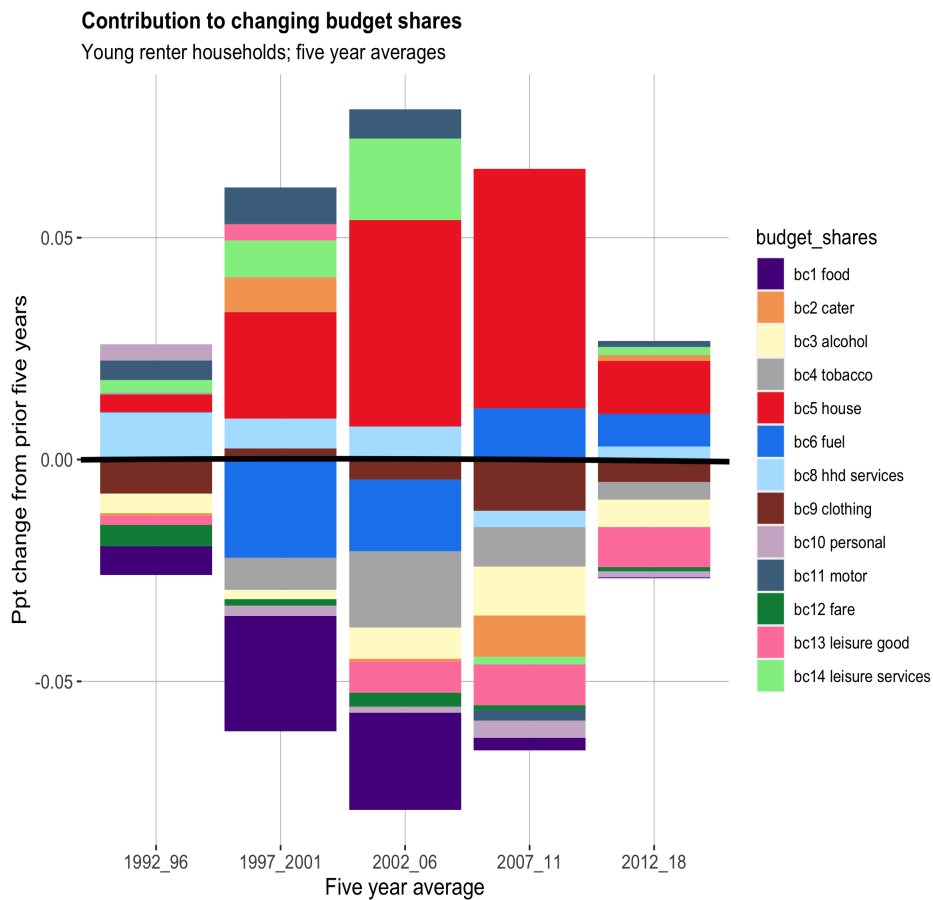
<sup>5</sup>I include both young renters and owner-occupier households when calculating expenditure percentiles to control for increases in the share of higher-income young households renting over time, and the percentiles are calculated by year.

Housing expenditure for those in the bottom 10 per cent of total expenditure remained fairly low; this cohort would consist of many individuals in subsidised housing, as well as individuals not paying market rents for other reasons.

#### 4.3.1.3 Non-housing budget shares

The large increase in housing budget shares by young renters was at the expense of most other expenditure categories, as illustrated in Figure 4.2. Relative to 1987–91, the budget shares of clothing, alcohol, tobacco, fuel, leisure goods, food, and fares were all smaller. The only categories where young renters increased their budget shares were housing services, motoring, and leisure services, and these increases mostly occurred during the pre-2006 part of the sample.

Figure 4.2



I report the equivalent figure for young mortgagors in Appendix figure C.3; their

budget shares were more stable over the sample period. I also graph changes to each budget share separately in Appendix C.0.1.4.

To determine how these budget share changes varied by household expenditure levels, I constructed Engel curves for the rest of my amalgamated consumption categories: food, fuel and light, housing, leisure, other goods, transport and other services. Food Engel curves for young renters shifted down fairly consistently over the expenditure distribution from the first to the second decade of my sample as individuals allocated a smaller share of their budget to food, and then were broadly unchanged between 1998–2007 and 2008–2017 despite much weaker wage growth in the UK following the financial crisis (Emmerson, Johnson and Ridpath, 2024).<sup>6</sup> Food prices fell in real terms over the sample period, which may explain the broad-based decline in the budget share of food. Leisure Engel curves, which include leisure goods, leisure services and catering, were fairly unchanged between the first and second decade of my sample, but then shifted down as individuals reduced their leisure budget shares between 1998–2007 and 2008–2018. The declines were largest for individuals around the middle half of the expenditure distribution. Many of the items in the leisure consumption category, such as hobbies, entertainment and holiday expenses, would be the first to be reduced if a household is trying to cut back on expenditure.<sup>7</sup> However, very high expenditure renters increased their leisure budget share over the sample period.


Energy prices were volatile over the sample period. Electricity prices were high in the early 1980s, fell significantly over the latter part of the 1990s and early-to-mid 2000s, and then rose again. The downward shift in the fuel and light Engel curve in the second decade of my sample matches this period of low energy prices. The ‘other goods’ category consists of alcohol, tobacco, clothing and footwear, and personal goods and services. The budget shares of alcohol and tobacco both declined sharply over the sample period, alongside steep price increases. There was also a decline in the budget share of clothing and footwear, but this category experienced

---

<sup>6</sup>My food variable does not include catering and restaurants


<sup>7</sup>The main categories of leisure goods are: TV/video/audio equipment, sports, camping and outdoor goods and equipment, newspapers, magazines, books, stationery, and toys, hobbies and photography, and the main categories of leisure services are: entertainments, social events, sport, TV and video licence, rental, subscriptions, education and training, hotels and holiday expenses, betting stakes and betting winnings.

Figure 4.3



large real price declines. The budget share of personal goods and services, which includes personal articles, chemist goods, and personal services only declined a little over the sample period.

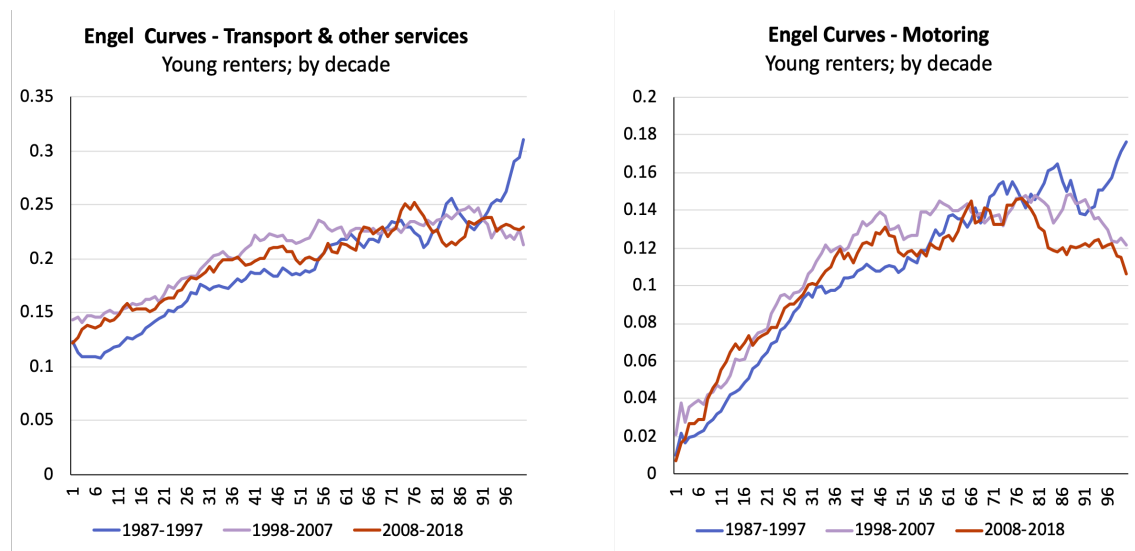
Figure 4.4



The Engel curve for transport and other services, which includes motoring, fares and household services, shifted up for individuals in the bottom two-thirds of the expenditure distribution over the sample period as they allocated a greater share of their budget to these services, but shifted down for high-expenditure individuals. This

mostly reflects large falls in motoring expenditure by high-expenditure individuals, as shown in the right panel of Figure 4.5. Around 60 per cent of this category is made up of motoring expenditure, with transport fares and household services each making up around 15 per cent.<sup>8</sup> This is consistent with high-expenditure rental households living in expensive inner-city areas where they do not need to spend much on transport, while lower-expenditure households may need to live further away to reduce their rent burden and so spend more on transport. Car ownership rates for young renter households remained fairly stable over the sample period.<sup>9</sup>

Figure 4.5



### 4.3.2 Housing prices

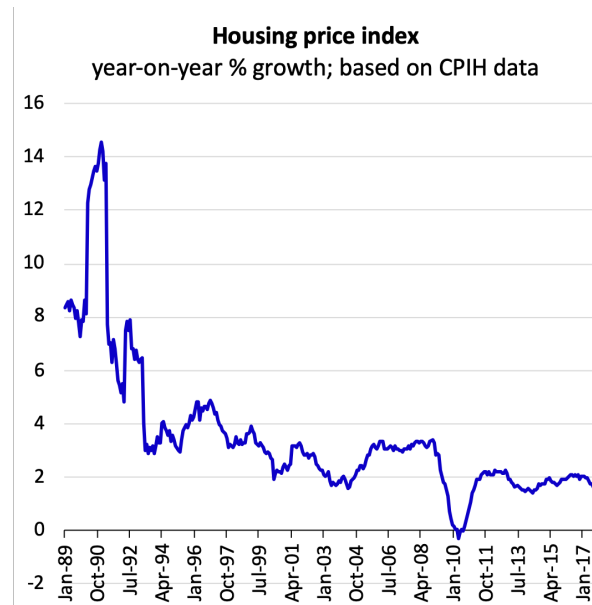
For my demand system estimation, I need price indices that match each consumption category. For all non-housing categories, I use RPI data that has been aggregated to match the FES consumption categories. However, I construct my own price index for housing using CPIH data from the UK’s Office of National Statistics, as the RPI housing index data is based on actual expenditure for owner-occupiers such as mortgage and interest payments, rather than imputed rent that captures the consumption flow from consuming housing services. I aggregate the following indices,

<sup>8</sup>Household services includes postage, telephones, domestic services, and fees and subscriptions

<sup>9</sup>See Appendix figure C.4 for graph of car ownership rates over time by expenditure quintile

which I weight using my budget share data: actual rentals for housing, owner occupiers housing costs (imputed rent), regular maintenance and repair of the dwelling, water supply and miscellaneous services relating to the dwelling, and council tax and rates. Figure 4.6 presents a graph of this index.

Figure 4.6



A limitation of my empirical approach, is that I only have national price indices for rents and the other components that make up my housing price index, therefore I cannot capture regional housing price variation in my demand system estimation. I do include regional dummies as a demographic shifter variable, but this can only shift estimated budget shares for all consumption goods equally. Unfortunately, time series of regional prices in the UK that could be applied to my six consumption categories do not exist. Regional price indices for some consumption variables do exist, but are still in an experimental stage and do not have a long enough back history for my analysis.<sup>10</sup> Using national price indices in my demand system estimation is equivalent

<sup>10</sup>In the 2003 budget, the then UK Chancellor announced plans to produce regional price indices for the UK, but this remains a future work plan for the ONS. The ONS has provided a few snapshots of regional price levels in 2000, 2003, 2004, 2010 and 2016, but they only consider London, Scotland, England (excl. London), Wales and Northern Ireland, and they do not including housing costs. In general, these indices show little price dispersion for food, clothing, footwear, alcohol, tobacco, and greater price dispersion where there is an element of service due to variability in labour and rental costs. See for details: <https://www.ons.gov.uk/economy/inflationandpriceindices/articles/relativeregionalconsumerpricelevelsuk/2016>

to assuming relative prices for each consumption good in my demand system are the same between regions. This will, for example, overstate Hicksian housing own-price elasticity for those in high-housing-price areas, and understate it for those in low-housing-price areas. However, while the magnitudes do vary, the increase in housing rents and house prices relative to regional income and regional non-housing prices has been broad-based by region, and has not just been a London and south-east England phenomenon (Figure C.0.1). Therefore, trends in housing own and cross-price elasticities constructed using national price data should still be informative on how young renters have responded to house price increases.


### 4.3.3 House price and tenure changes

I conclude the summary statistics by briefly describing two other trends in the data that are highly relevant for this work.<sup>11</sup> First, the left hand panel of Figure 4.7 depicts the change in price levels for three indices: the Nationwide UK house price index, an index of actual rentals paid for housing which is from the CPIH and is a key input into my housing index, and the overall CPIH index, which captures overall consumer price changes over my sample period and does include rents. These indices show that while house prices have been more volatile and increased more than rents, both have risen by far more than overall consumer prices. I discuss regional variation in these indices in Appendix C.0.1. Second, I observe a steady decline in the homeownership rate of young people in the second half of my sample, as shown in the right hand panel of Figure 4.7. The decline is largest for middle-income young households, although the homeownership rate for households in the top disposable income quintile, who presumably would have the least financial difficulty transitioning to homeownership, declined by almost 20 per cent. Households with disposable income in the bottom 20 per cent of young households have consistently low homeownership rates. The decline in homeownership is broad-based by region, and is not limited to regions with the most expensive houses or fastest house price growth (Cribb and Simpson, 2018)

---

<sup>11</sup>I report additional summary statistics and graphs on these topics in Appendix C.0.1

Figure 4.7



## 4.4 Demand System Estimation

This section presents my estimates for the housing cross-price elasticity of demand for six major expenditure categories: food, fuel and light, housing, leisure, other goods, transport and other services. Since I focus on renters, I can treat housing as a consumption good just like the other five in my demand system. To calculate these elasticities, I estimate an Almost Ideal Demand System (AIDS). I briefly describe the theoretical properties of AIDS, my estimation process, and interpret the resulting elasticity estimates.

### 4.4.1 Demand system theory

Demand system estimation allows us to identify the impact of house price changes on the allocation of a budget between differentiated products. Almost Ideal Demand System estimation is the most widely used method of estimating demand functions. Starting with an expenditure function, Deaton and Muellbauer (1980*a*) derive Marshallian demand equations expressed in terms of expenditure (budget) shares  $w_i = x_{it}p_{it}/m_t$ :

$$w_{it} = \alpha_i + \sum_j \gamma_{ij}^* + \ln p_{jt} + \beta_i \ln \left( \frac{m_t}{a(\mathbf{p})} \right) \quad (4.1)$$

where  $\gamma_{ij}^* = \frac{1}{2}(\gamma_{ij} + \gamma_{ji})$ , which equals  $\gamma_{ij}$  if symmetry is imposed,  $p_{it}$  is the price of good  $i$  at time  $t$ ,  $x_{it}$  is real household expenditure for good  $i$ ,  $m_t$  is total nominal expenditure,  $\alpha$ ,  $\beta$  and  $\gamma$  are coefficients I estimate, the consumption goods are indexed to  $j$ , and  $a(\mathbf{p})$  is a translog price index such that  $\ln a(\mathbf{p}) = \alpha_0 + \sum_i \alpha_i \ln p_{it} + \frac{1}{2} \sum_i \sum_j (\frac{1}{2}(\gamma_{ij} + \gamma_{ji})) \ln p_{it} \ln p_{jt}$ . This approach requires a separability assumption that preferences over different products within one group is independent of preferences over different products within another group. Then multi-stage budgeting can be assumed, so that households first allocate total expenditure between a small number of highly aggregated groups, and then allocate within each sub-group. I provide a summary of Deaton and Muellbauer (1980*a*)'s derivation of equation 4.1 in Appendix C.0.2.1.

Since the AIDS framework implies that Engel curves are linear in the log of total expenditure, I incorporate a common extension by Banks, Blundell and Lewbel (1997). Their Quadratic Almost Ideal Demand System (QUAIDS) model includes a quadratic log expenditure term, scaled by the Cobb-Douglas price aggregator  $b(\mathbf{p}) = \prod_i p_{it}^{\beta_i}$ . This allows for curvature in the Engel curves, which permits goods to be necessities at some income levels and luxuries at other income levels:

$$w_{it} = \alpha_i + \sum_j \gamma_{ij}^* \ln p_{jt} + \beta_i \ln \left( \frac{m_t}{a(\mathbf{p})} \right) + \frac{\lambda_i}{b(\mathbf{p})} \left[ \ln \left( \frac{m_t}{a(\mathbf{p})} \right) \right]^2 \quad (4.2)$$

There are six properties that an estimated AIDS model must hold to be consistent with demand theory: adding up, homogeneity, symmetry, positivity, monotonicity and curvature/negativity (Deaton and Muellbauer, 1980*b*). The adding-up condition is automatically satisfied as, by definition, observed budget shares always sum to one. To ensure this is satisfied by the estimated budget shares, one imposes the following constraints:  $\sum_i \alpha_i = 1$ ,  $\sum_i \beta_i = 0$  and  $\sum_i \gamma_{ij} = 0 \forall j$ .

To incorporate demographic variables ( $\mathbf{z}$ ), I adopt the method established by Ray (1983) that I summarise in Appendix C.0.2.1. I use my dataset to estimate the following expenditure share equation for each good (omitting time subscripts for readability):

$$w_i = \alpha_i + \sum_j \gamma_{ij}^* \ln p_j + (\beta_i + \eta_i \mathbf{z}) \ln \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} + \frac{\lambda_i}{b(\mathbf{p})c(\mathbf{p}, \mathbf{z})} \left[ \ln \left\{ \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} \right\} \right]^2 \quad (4.3)$$

I estimate the coefficients:  $\alpha$ ,  $\beta$ ,  $\gamma$ ,  $\eta$ ,  $\lambda$  and  $\rho$ .  $c(\mathbf{p}, \mathbf{z})$  is another Cobb-Douglas price aggregator:  $\prod_{j=1}^k p_j^{\eta_j \mathbf{z}}$ .  $\bar{m}_0$  is defined as the ratio of the expenditure required to obtain a certain utility level at a given price for a household ( $h$ ) with demographic features ( $\mathbf{z}$ , where  $\mathbf{z}$  is a vector of  $s$  characteristics, indexed to  $r$ ), and can be decomposed into two factors:  $\bar{m}_0(z)$ , which captures the change in household expenditure as a function of demographic variables only, and  $\phi(\mathbf{p}, \mathbf{z}, u)$ , which captures changes in relative prices and actual goods consumed. We can parameterise  $\bar{m}_0(z)$  as

$\bar{m}_0(\mathbf{z}) = 1 + \rho\mathbf{z}$ , where  $\rho$  is a vector of parameters to be estimated. The adding up condition now requires that  $\sum_{j=1}^k \eta_{rj} = 0$  for  $r=1\dots s$ .

I also impose parameter restrictions to ensure symmetry, which means that the cross-price derivatives of Hicksian demand functions are equal, and so  $\gamma_{ij} = \gamma_{ji} \forall i \neq j$ . If the adding up and symmetry properties are both met, this implies the homogeneity condition of degree zero in total expenditure and prices is also met, and  $\sum_i \gamma_{ij} = 0 \forall i$ . Positivity means no negative prices, which I do not have in my dataset. The final two properties; monotonicity and curvature, are not imposed by my estimation but I check my estimates for violations. Monotonicity requires that the first-order derivatives of the cost function must be non-negative, which is satisfied if the estimated budget shares are positive. Curvature, which is also referred to as negativity, requires that the Hessian matrix of the expenditure function is negative semi-definite, and therefore Hicksian own-price elasticities are not positive.

#### 4.4.2 Estimation details

I estimate a QUAIDS model using the Stata *demandsys* package that uses an iterated feasible generalised nonlinear least-squares estimator to solve the system.<sup>12</sup> I estimate the model on my full sample of renters, and then report elasticity estimates calculated only on the sub-sample of renters aged under 35. I use the full sample of renters for estimating the demand system coefficients because if I just use young renters, the sample becomes too small. I exclude owner-occupiers as housing is both an investment and consumption good for this group, and I do not have imputed rent data that only captures their consumption of housing services.<sup>13</sup>

I include five ‘shifter’ variables that enter the QUAIDS system linearly and shift the constant  $\alpha_i$  term up or down. Since they enter linearly, we must assume that the impact of these mostly demographic variables on the budget shares is independent of total expenditure (Ray, 1983). There are modelled as follows:

$$\alpha_i = \phi_{0,i} + \phi_{1,i}\text{region} + \phi_{2,i}\text{kids} + \phi_{3,i}\text{adults} + \phi_{4,i}\text{cohort} + \phi_{5,i}\text{tenure IMR}$$

<sup>12</sup>Further estimation details is available at: <https://www.stata.com/manuals/rdemandsys.pdf>

<sup>13</sup>As a robustness exercise I do re-estimate the QUAIDS model using this group and report the results in Appendix table C.2

$\phi_{1,i}$  is the coefficient on which of the 11 regions in England, Wales or Scotland the household resides in<sup>14</sup>,  $\phi_{2,i}$  is the coefficient on number of children in the household,  $\phi_{3,i}$  is the coefficient on number of adults in the household,  $\phi_{4,i}$  is the coefficient on five-year birth cohorts, which can be considered a crude time trend, and  $\phi_{5,i}$  is the coefficient on the inverse mills ratio from a tenure selection equation. This last term is included to account for the selection of individuals into renting or homeownership. Following the structure of a two-stage Heckman model (Heckman, 1979), I first estimate a probit model to capture selection into being a renter rather than owner-occupier:

$$\mathbb{P}(\text{tenure} = \text{rent} \mid X) = \mathbb{P}(X_i\beta + \varepsilon_i > 0) = \Phi(X\beta)$$

$X$  is a vector of all right hand side variables in the QUAIDS model, some selection variables as well as some additional demographic variables I do not include in the main QUAIDS model so that the model does not become too computationally demanding. These variables are: prices, total expenditure, number of adults, number of kids, cohort, region, age and sex of household head, marital status, education, age and education interactions. The inverse mills ratio is then calculated as:

$$\text{inverse mills ratio} = \frac{\phi(X\hat{\beta})}{\Phi(X\hat{\beta})}$$

I specify the value of  $\alpha_0$  to be equal to zero, which is common practise. Deaton and Muellbauer (1980a) note that the practical identification of  $\alpha_0$  is problematic. Since it can be interpreted as the consumption required for a minimal standard of living when prices are unity, a common approach is to set  $\alpha_0$  to be a little below the lowest value of the log of total expenditure observed in the data (Banks, Blundell and Lewbel, 1997), which is often zero.

Estimating this demand systems requires several assumptions to be made. Perhaps the strongest for my research question, is that I assume consumption choices are independent between time periods, and therefore there is no scope for inter-temporal

---

<sup>14</sup>The 11 regions are the nine regions of England as identified by the UK government (North East, North West, Yorkshire and the Humber, East Midlands, West Midlands, East of England, London, South East, South West) plus Scotland and Wales

substitution or changes to saving behaviour. An important channel available to young renters responding to large house price increases is to change their saving behaviour, either to save more for a housing deposit and push consumption into the future, or to give up saving and do the reverse. There may also be wealth effects from young renters who aspire to become future homeowners and expect house prices to continue appreciating after they transition into homeownership. However, I am constrained by my data as I only have cross-sectional data, and high-quality disaggregated consumption panel data for the UK does not exist. Since I do have cross-sectional data on how total expenditure changed over time as house prices rose I would be able to identify some of these inter-temporal channels using my dataset, but I choose to focus on intra-temporal substitution between expenditure categories and assume inter-temporal consumption behaviour is unchanged.

#### 4.4.3 QUAIDS coefficients & elasticity estimates

I report the QUAIDS coefficient estimates in Table 4.3.  $\alpha_i$  is the intercept for the budget share equation for good  $i$ ,  $\beta_i$  is the expenditure coefficient that captures the change in the budget share of good  $i$  with respect to a one per cent change in total deflated expenditure, and  $\lambda_i$  has a similar interpretation with the squared deflated expenditure term.  $\gamma_{ji}$  are price coefficients and represent the change in budget share of good  $i$  in response to a one per cent change in the price of good  $j$ . If  $\gamma_{ji} > 0$ , then goods  $i$  and  $j$  are substitutes, and if  $\gamma_{ji} < 0$ , then goods  $i$  and  $j$  are complements. For young renters, housing is a substitute with most other goods. All estimated budget shares are positive, indicating that the monotonicity demand condition is satisfied. I report the  $\rho$  estimates in the footnotes, and do not report the  $\eta$  terms as these are very small in magnitude; these two sets of coefficients originate from the incorporation of demographic variables.

I use my coefficient estimates to construct price and income elasticities. I report compensated (Hicksian) price and income elasticities in Table 4.4. Compensated price elasticities estimate the impact of the change in the price of a good on the demand for that good (own-price) or another good (cross-price) in the demand system, holding total utility constant, and income elasticities estimate the impact of a change in total

Table 4.3: QUAIDS coefficient estimates - renters

$Good_i$	$\alpha_i$	$\beta_i$	$\gamma_{1i}$ (food)	$\gamma_{2i}$ (house)	$\gamma_{3i}$ (fuel)	$\gamma_{4i}$ (leisure)	$\gamma_{5i}$ (goods)	$\gamma_{6i}$ (services)	$\lambda_i$
food	-0.306* (0.016)	-0.015* (0.006)	0.079* (0.008)	.	.	.	.	.	0.037* (0.002)
house	0.237* (0.045)	-0.446* (0.013)	-0.112* (0.010)	0.416* (0.034)	.	.	.	.	-0.140* (0.007)
fuel	-0.018 (0.011)	0.063* (0.004)	0.043* (0.003)	-0.052* (0.008)	0.074* (0.002)	.	.	.	0.036* (0.002)
leisure	0.634* (0.029)	0.219* (0.010)	-0.080* (0.009)	-0.111* (0.021)	-0.014* (0.005)	0.287* (0.019)	.	.	0.044* (0.003)
other goods	0.487* (0.024)	-0.064* (0.009)	0.067* (0.007)	0.200* (0.015)	-0.058* (0.004)	-0.246* (0.012)	-0.099* (0.013)	.	-0.020* (0.002)
other services	-0.033 (0.031)	0.213* (0.009)	0.005 (0.012)	-0.341* (0.019)	0.007 (0.005)	0.163* (0.016)	0.136* (0.013)	0.004* (0.024)	0.044* (0.003)

\*= $p < 0.01$ , standard errors in parentheses  
 $\rho$  coefficients: imr -0.384\* (0.025), adults 0.686\* (0.046), kids 0.647\* (0.040), cohort 0.044\* (0.005), region -0.005 (0.004).  $\eta$  coefficients not reported as all are very small (under 0.01)

income on the quantity demanded of a good.<sup>15</sup> My income elasticity estimates are all positive, indicating that all commodities in my demand system are normal goods. Food and fuel/light are necessities as their income elasticity is below 1, while the other goods are luxuries. Compensated own-price elasticities are negative for all goods except leisure (for which the elasticity estimate is not significant), therefore satisfying the curvature/negativity demand property. The elasticities of greatest relevance to my research question are the housing cross-price elasticities as they indicate how the non-housing budget shares of renters responded to house price increases. I describe these in further detail in Section 4.4.5.

#### 4.4.4 Robustness checks

I perform several several robustness checks on my QUAIDS estimates. First, I re-estimate my budget shares as a Exact Affine Stone Index (EASI) implicit Marshallian demand system (Lewbel and Pendakur, 2009). The EASI demand system can be considered a generalisation of the AIDS system. While more complex, its two

<sup>15</sup>Formulas for these elasticities are reported in Appendix C.0.2.2

Table 4.4: Average income and Hicksian elasticity estimates: Young renters

	Food	Housing	Fuel	Leisure	Other goods	Other services
income	0.232** (0.008)	1.412** (0.009)	0.181** (0.010)	1.346** (0.009)	1.017** (0.008)	1.159** (0.009)
food price	-0.481** (0.057)	-0.063 (0.032)	0.531** (0.037)	-0.337** (0.052)	0.575** (0.038)	0.142* (0.066)
housing price	-0.243** (0.051)	-1.439** (0.064)	0.808** (0.055)	1.089** (0.063)	0.882** (0.056)	-0.427** (0.066)
fuel price	0.240** (0.017)	0.324** (0.016)	-0.088** (0.023)	-0.219** (0.022)	-0.173** (0.017)	-0.078** (0.021)
leisure price	-0.410** (0.062)	0.826** (0.048)	-0.539** (0.057)	0.077 (0.092)	-0.983** (0.058)	0.430** (0.078)
other goods price	0.713** (0.049)	0.705** (0.046)	-0.533** (0.048)	-1.066** (0.063)	-1.455** (0.073)	1.169** (0.066)
other services price	0.180* (0.085)	-0.352** (0.054)	-0.180** (0.058)	0.457** (0.084)	1.154** (0.065)	-1.236** (0.127)

\* =  $p < 0.05$ , \*\* =  $p < 0.01$ , standard errors in parentheses

main benefits over an AIDS system is that it allows for much more flexibility in estimated Engel curves (income expansion paths), which can have any shape over real expenditures, unlike AIDS, which restricts Engel curves to be linear, and QUAIDS which restricts them to be quadratic. Second, it allows for unobserved preference heterogeneity, which is not possible with AIDS models as this would require additive errors to be heteroskedastic (Lewbel and Pendakur, 2009). Using the notation from their paper, budget share ( $w_j$ ) equations take the following form:

$$w_j = \sum_{r=0}^5 \mathbf{b}_{rj} y^r + \sum_{l=1}^L (\mathbf{C}_{lj} z_l + \mathbf{D}_{lj} z_l y) + \sum_{l=0}^L \sum_{k=1}^J \mathbf{A}_{lkj} z_l \mathbf{p}_k + \sum_{k=1}^J \mathbf{B}_{kj} \mathbf{p}_k y + \varepsilon_j \quad (4.4)$$


There are  $j$  goods indexed by  $k$ , and  $L$  demographic characteristics indexed by  $l$ . Equation 4.4 models the budget share for good  $j$  as a function of: an (up to) fifth order polynomial in total expenditure ( $y$ ), demographic characteristics ( $z$ ), log prices of each good ( $p_k$ ), and several sets of interaction terms:  $z_l y$ ,  $z_l p_k$  and  $p_k y$ . I estimate an EASI model using my renter sample that includes a full set of interaction terms, and otherwise has the same specifications as my QUAIDS model.<sup>16</sup> I report the

<sup>16</sup>I use a combination of replication Stata code published with the Lewbel and Pendakur (2009) paper, and the *easi* R package

coefficient estimates in Appendix table C.3.

The purpose of this robustness check is to determine how sensitive my elasticity estimates are to my demand estimation specifications. The housing Engel curve constructed using my data fits a quadratic shape fairly well, therefore the Engel curve flexibility offered by the EASI model is less helpful for my purposes. However, accounting for unobserved heterogeneity could have a large impact on my estimates. I report the full elasticity results estimated using the EASI demand system in Appendix table C.4 and graph the compensated housing cross price elasticities estimated using QUAIDS and EASI system in Figure 4.8, including error bars to represent the 95% confidence intervals. The only variable where the QUAIDS and EASI estimates significantly diverge is the other goods variable. Since this variable includes alcohol and tobacco where consumption patterns vary hugely, it is not surprising that a model that allows for preference heterogeneity has the biggest impact on the estimated elasticity for this variable. I conclude that my compensated housing cross-price elasticity estimates are fairly robust to demand estimation specifications, with the exception of estimates for the other goods variable. I therefore omit a discussion of the other goods elasticity estimates in the next section.

Figure 4.8



My second robustness test is to account for possible endogeneity caused by total expenditure being both an explanatory variable and the denominator of the budget share, which is my dependent variable. I follow Banks, Blundell and Lewbel (1997), and use the Two-Stage Least Squares (2SLS) estimator, with the expenditure and expenditure-squared terms being instrumented with disposable household income and its square. Household income is a commonly used instrumental variable for expenditure in the demand literature. If the potential endogeneity is due to heterogeneity in tastes, then income is a valid instrument if labour supply is weakly separable from consumption (Attanasio and Lechene, 2014). The validity of this assumption is unclear, although there is some empirical evidence that labour supply of young renters does not respond to house price changes (Disney and Gathergood, 2018). I report the results of this IV estimation in Appendix C.0.2.5. I find that the impact of instrumenting expenditure with household income on my coefficient estimates to be fairly small, and it also significantly reduces my sample size as I do not have disposable household income data for my entire sample. Therefore, I choose to retain my baseline QUAIDS elasticity estimates.


#### **4.4.5 Interpretation of elasticity estimates**

I determine how young renters responded to housing price increases by interpreting the compensated housing cross-price elasticities. I focus on how these elasticities varied over my sample period as rental and housing prices rose above CPI, and how they varied between lower and higher expenditure households. Since changes in elasticities indicate changing preferences (if we assume CES (Constant elasticity of substitution) preferences when moving up or down demand schedules), I interpret households becoming more sensitive to housing price changes as potential evidence of the impact of the large increase in house prices, especially if I observe this among lower-expenditure households. However, I find little evidence of rising housing cross-price elasticity in my sample period for most goods and services and for most households. I do find some evidence that lower-expenditure households have become less able or willing to reduce their housing expenditure when house prices rise, which may reflect the decline in social housing.

#### 4.4.5.1 Housing

The average compensated housing own-price elasticity became less negative over my sample period, especially between 1996–2007, as shown in Figure 4.9. This increase was driven by lower-expenditure households. While low-expenditure households still reduced their housing expenditure the most when housing prices rose in the mid-2010s, the magnitude of their response was much smaller relative to the late-1980s. The change in compensated housing own-price elasticity for this group is very large, and matches the decline in social housing (Waters and Wernham, 2023). Between 1987–1989 and 2017–2019, the share of young households in their 30s and in the bottom-third income tercile who lived in social rental accommodation fell from over 40 per cent to below 30 per cent, replaced almost entirely by private renting (Waters and Wernham, 2023). The increased inability to access social housing and therefore having to pay market rates can explain this large increase in compensated own-price housing elasticity, as individuals increasingly could not access cheaper accommodation to switch to when faced with housing price increases.

Figure 4.9



For the median expenditure cohort, the average compensated housing own-price elasticity was surprisingly stable over the sample period at around -1, despite the particularly large decline in homeownership rate of this group. For high expenditure cohorts, their compensated own-price housing elasticity was also fairly stable over

the first two decades of the sample but then declined a little to converge with the elasticity of the median households.

#### **4.4.5.2 Food and fuel/light**

Food and fuel/light are the two expenditure categories where if individuals demonstrated increasing demand sensitivity to housing price increases, this could be a sign of financial stress as individuals reduce their consumption of essential everyday goods to keep paying the rent. This is not what I observe, as illustrated in Figure 4.10. For the lowest expenditure households, compensated food cross housing price elasticity remained stable at just above zero, indicating that these individuals did not modify their food demand in response to housing price changes. For higher expenditure households, this elasticity was quite volatile over the sample period. These households allocated a higher share of their budget to catering and restaurants than lower-expenditure households, which are a substitute to (home) food. The budget share allocated to catering and restaurants increased and then decreased over the sample period, and changes to the food price elasticity for higher-expenditure households may reflect this as food-type preferences changed. I observe little difference in the compensated fuel-housing cross-price elasticity by expenditure quintile.

#### **4.4.5.3 Other non-housing goods and services**

Leisure goods, leisure services and catering would be the easiest expenditure category to cut down on if a household is experiencing financial stress, which would be consistent with the leisure-housing cross-price elasticity becoming more negative over time. However, I once again do not observe this. For high-expenditure households, the leisure-housing cross price elasticity remained just above zero, indicating that leisure consumption by high expenditure households was unaffected by changes to housing prices. The leisure demand of lower-expenditure households is much more responsive to house price changes, but the direction is positive. As housing prices increased, these individuals spent more on leisure goods, leisure services and catering. Because there is a strong positive correlation between housing prices and income (Attanasio et al., 2009), and the leisure income elasticity is large and above 1 for all expenditure quintiles, as shown in Figure 4.11, it suggests that this effect dominated the budget

pressures from rising rents. The leisure budget shares for lower-expenditure young renters do decline over this period, but this decline is driven by the leisure goods sub-component, for which real prices fell by almost 200 per cent over the sample period.

The housing cross-price elasticity of demand for transport and other services became less negative over the sample period, indicating that individuals reduced their spending on these items by less when house prices rose. The aggregate includes motor-ing, public transport fares and household services, particularly phone bills. Differences by expenditure quintile were small.

Figure 4.10






Figure 4.11



## 4.5 Welfare Analysis

The multi-decade rise in house and rent prices above CPI and the concomitant multi-decade decline in homeownership in the UK occurred alongside the decline in the relative prices of many other goods and services, as well as strong expenditure growth until the Great Recession, as illustrated in Figure 4.12. I aggregate the welfare consequences of all these price changes for young renter households between 1987–2017 by calculating the compensating variation (CV) of these price changes. I find that the overall welfare of young renter households improved from the beginning of the sample until the Great Recession as income growth and declining prices of non-housing goods and services offset rising housing prices. Since then, the welfare of young renter households declined a little due to weak wage growth and an ongoing re-allocation of expenditure towards housing and services consumption, where price growth has been stronger.

Figure 4.12



### 4.5.1 Estimation

Compensating variation is the amount of additional money a household would need to return to their original utility following a change in prices at  $t = 1$ , keeping

total expenditure  $w$  otherwise fixed between two periods ( $t = 0, 1$ )

$$CV = e(p_1, u_1) - e(p_1, u_0) = w - e(p_1, u_0) \quad (4.5)$$

I follow the approach of Friedman and Levinsohn (2001), and take a second-order Taylor expansion of the expenditure function  $e(p, u)$  with respect to price. A full derivation of this expansion is provided in Appendix C.0.3. This results in the following compensating variation expression:

$$CV \approx \sum_{i=1}^n w_i \Delta \ln p_i + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n w_i \varepsilon_{ij} \Delta \ln p_i \Delta \ln p_j \quad (4.6)$$

$w_i$  is the budget share of commodity  $i$ , there are  $n$  commodities,  $\Delta \ln p_i$  is the change in log price of good  $i$ , and  $\varepsilon_{ij}$  is the compensated price elasticity of good  $i$  with respect to the price change of good  $j$ , which I obtain from my QUAIDS estimates. The first term in equation 4.6 can be interpreted as the income effect of a price change, and the second term captures substitution effects. Price increases of goods that are more inelastic will generate larger compensating variation estimates, as households are less willing or able to substitute to other relatively cheaper goods. Since this calculation assumes total expenditure is constant between the two periods, I deflate prices by growth in the total expenditure of young households in my sample. Therefore, compensating variation will be positive if prices grow faster than expenditure growth.

I calculate the compensating variation for the price changes that occur each year of my sample relative to the base year of 1987. Since my dataset is cross-sectional, to obtain 1987 budget shares ( $w_i$ ), I match post-1987 households to the budget allocation choices of an average 1987 household based on region, age and tenure choice. For example, I match a 27 year old renter living in Yorkshire and the Humber with the average budget shares of all 27 year old renters living in Yorkshire and the Humber in 1987. To incorporate differences in expenditure growth by expenditure quintile, I deflate the prices faced by households in each quintile by the median household expenditure growth of renters or homeowners in that quintile. The expenditure quintiles are calculated based on all young people to control for the compositional shift

towards more higher-income renters.


### 4.5.2 Compensating variation estimates

I report the results of my compensating variation estimates in Figure 4.13. Compensating variation is measured as a proportion of 1987 household expenditure, which is the base year. Housing prices is the only expenditure category price that grew at a consistently faster pace than expenditure. As a result, compensating variation is negative for the majority of the sample period. It was positive during the late 1980s–1997 for the lower-expenditure quintiles, reflecting the recession of 1991–92 and subsequent years of high unemployment and low household income growth relative to retail price inflation. 2000–2007 was the period when compensating variation was most negative, reflecting strong expenditure growth relative to price changes. For the last decade of the sample, compensating variation remained negative but smaller, indicating that the average total welfare of young renter households was higher than at the beginning of the sample, but lower than in 2000–2007. The compensating variation for the bottom expenditure quintile was more negative than for other expenditure quintiles. This is mostly due to their smaller housing budget share as many live in social housing or receive housing benefits so they are less negatively affected by the rise in house prices. In addition, this quintile allocated a far larger share of their budget to food, which saw large price declines.<sup>17</sup>

---

<sup>17</sup>Over the sample period, the lowest expenditure quintile allocated an average of 21 per cent of their budget to food, compared to 13 per cent by the median young renter household.

Figure 4.13



Compensating variation is reported as a proportion of 1987 household expenditure

I conclude that the decline in homeownership rates of young people between the early 1990s and the Great Recession was accompanied by an increase in the overall welfare of young people as measured using compensating variation. This is because household income growth was robust over this period, so the purchasing power of households over non-housing goods and services increased, offsetting the negative welfare consequences of high house price growth. However, following the Great Recession, the average welfare of young renter households declined relative to the 2000s. This decline has been broad based by expenditure quintile, and can be attributed to various factors, including weak income/expenditure growth but also a re-allocation of expenditure towards housing and services consumption, such as leisure services, where price increases have been larger.

## 4.6 Conclusion

The budget share that young renters aged 18-35 allocated to housing increased by almost 15 percentage points between 1987 and 2018, at the expense of most other goods and services. Over this period, house prices in the UK increased by more than twice the rate of overall inflation, rental prices increased by 1.5 times the rate of inflation, the youth homeownership rate fell sharply, especially for middle-income households, and there was rising media and government concern around worsening housing affordability for first home buyers (Cribb and Simpson, 2018). This chapter evaluated how the large increase in housing prices impacted the housing and non-housing budget shares of young renters in the UK over this period. I estimated a QUAIDS model to disentangle preferences from relative price changes, and used the coefficients to calculate compensated (Hicksian) cross-price and income elasticities for housing.

I find that lower-expenditure households became less responsive to housing price increases over the sample period, which coincided with a large fall in social housing supply, limiting the ability of these households to respond to large price increases by obtaining cheaper rental accommodation. However, the price sensitivity of the housing consumption of median-expenditure young renter households to housing price changes was surprisingly stable over thirty years. I do observe some heterogeneity in how the demand by relatively higher and lower-spending households for different non-housing goods and services responded to housing price change over the sample period, but these can be explained by non-housing factors. In addition, my estimates of compensating variation suggest that despite the large housing price increases, the prices of most other goods and services declined over the sample period relative to income growth, and so young renter households are better off in the late 2010s relative to the late 1980s, although not better off relative to the mid-2000s. The welfare of low-expenditure renter households improved the most over this period, as social housing shielded many of them from the full impact of private rental increases, and they also particularly benefited from large falls in the price of food.

There are various potential extensions to this research. There have been several government programmes aimed at improving the rate of youth homeownership. These impact consumption of young households in complex ways. For example, van

Horen and Tracey (2022) showed that relaxing down-payment constraints via the UK's Help-to-buy scheme increased the rate of home purchases, but also increased non-durable and durable consumption, mostly due to local demand effect spillovers. Further research could examine the interaction of these policies with the housing and non-housing expenditure of young renter households, to better understand the welfare and distributional consequences of these policies. Further research could also incorporate additional margins of adjustment available to young renters in response to house price shocks, such as changes to labour supply or regional migration. More complex measures of welfare that better account for preference heterogeneity (Maes and Malhotra, 2024) could also be considered.

## CONCLUSION

This thesis presented a set of three applied microeconomics papers that contribute to the field's understanding of health, labour and consumption dynamics. In my first chapter, I comprehensively reviewed the most common ways that health dynamics are modelled in the literature and suggested two improvements that help us explain how people respond to health shocks. In my second chapter, my co-author and I filled a gap in understanding the many ways health impacts labour supply and showed that a sudden worsening in health increases the likelihood of occupation change and affects the types of occupations chosen, and that this varies by type of health shock. In my third chapter, I showed that despite the large increase in housing prices and resulting large increase in housing budget shares of young renters, the preferences of the majority of young renters for food, fuel, leisure, transport and other services, as well as housing, have remained surprisingly unchanged by this significant shock. All three topics are highly pertinent to current public policy debates in the United Kingdom.

The health dynamics chapter was designed to assist other researchers working on better understanding how health impacts economic outcomes, as I both review common approaches to modelling health, and suggest improvements. There are also several interesting ways one could extend the results of this chapter. Further cutting edge developments in the econometrics literature could be applied to modelling health dynamics, such as models that can account for interactions between permanent and

transitory shocks. This could potentially reveal interesting statistical properties of the health of individuals who suffer multiple health conditions that interact, such as a chronic and mental health condition. Second, I focus on predicting and describing health dynamics but do not consider how this maps onto an individual's own understanding of how their health will likely evolve in the future. Currently, we do not have a good understanding of how, and to what degree, people modify their expectations of the frequency and persistence of future health shocks following a period of poor health, from which they may have recovered. Researching this relationship would be a valuable addition to our understanding how an individual's health affects the economic decisions that they make, which is mediated by how an individual expects their health to change in the future. Finally, I show that the statistical properties of mental health dynamics such as shock persistence and frequency differ from the dynamics of overall health. The economics of mental health is a nascent but fast growing area in the literature, and further research could further investigate mental health dynamics, such as determining how these dynamics vary by observable characteristics, as well as considering alternate ways to construct mental health indices and incorporate new data, such as high-frequency health data from wearable health devices.

The second chapter documented a relationship between health and occupational change in that data that had previously received little attention in the literature, despite its relevance to UK public policy around supporting those in poorer health return to or remain in work. Further research could investigate the lack of relationship between suffering a health shock and switching to less interpersonal work, as our prior was that mental health shocks should reduce productivity in performing interpersonal tasks. This could include further analysis of the relationship between flexibility, task content and occupation mobility. Identifying how different types of health shocks erode task-specific skills and therefore productivity could help us better understand the mechanisms behind the occupation changes we document in this chapter. Finally, our analysis was limited by our sample size as even though we used a large UK panel dataset, we needed individuals who experienced two shocks within a short period: a sudden worsening of their health, and a change in their occupation, and we needed a large enough sample of these individuals to identify differences by type of health shock and occupation task intensity. A richer dataset, such as an administrative dataset

with linked health data could be used to discover further details on the impact of different types of health shocks on occupational change, and could also be used to develop a structural model based on our empirical results and theoretic framework.

The final chapter used methods from the consumer demand literature to better understand the impact of a huge price change on young renters, a group that has received significant media and government policy attention due to the sharp fall in homeownership of young households, and the rising rental burden they face. Further research could go in two directions. First, there are various ways to further develop the technical aspects of this paper. This includes using more complex measures of welfare that better account for preference heterogeneity, calculating imputed rent measures for young homeowners and incorporating them into the demand system estimation, and considering methods to better capture regional variation in price indices. Second, there are other margins of adjustment available to young renters to mitigate the impact of housing price increases, which will also affect their consumption. These include changing their labour supply or regional migration, and they may also benefit from one of the government programmes aimed at improving the rate of youth homeownership. The impact of these other margins of adjustment on the housing and non-housing consumption of young renters in the UK remains poorly understood.

### **A.0.1 Attrition**

I estimate a linear probability model of attrition to identify the magnitude of any relationship between health and attrition. Those in the poorest health quintile are around two percentage points more likely to drop out of the sample next period. In general, the literature is fairly sanguine about the risks of using health indices for economic research when there is differential attrition risk by health. Jones, Koolman and Rice (2006) find that response rates to the British Household Panel Survey (BHPS) vary by health, with elderly or low-income individuals who start the survey in poor health particularly likely to attrit, but finds that attempting to account for this using inverse probability weights is unnecessary for most research applications. Similarly, Pudney and Watson (2013) investigate the impact of reducing the effort made to chase up non-responders to BHPS and HILDA (an Australian panel dataset) surveys. While they find that the effort exerted to chase up non-respondents changed the sample prevalence of disability and ill health, their subsequent statistical modelling of the relationship between health and unemployment is unaffected. I conclude that the observed level of attrition in my dataset does not pose a significant threat to the robustness of my subsequent analysis.

Table A.1: Linear probability model of attrition

	Missing next period
health index quintile 1 (lowest)	0.0213*** (8.14)
health index quintile 2	-0.00200 (-0.80)
health index quintile 3 (baseline)	0
health index quintile 4	-0.00131 (-0.52)
health index quintile 5 (highest)	0.00166 (0.62)
age	-0.0205*** (-5.88)
age squared	0.0577*** (5.08)
age cubed	-0.0861*** (-5.61)
age quartic	0.0498*** (6.75)
sex	-0.00706*** (-4.34)
Observations	228,886

*t* statistics in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

## A.0.2 Regression output from health index construction

I estimate the health index separately for each data wave using an ordered probit. Below I report the output for the the second last wave of data (wave 11) as an example to show how different objective variables contribute to the final index. Most of the objective variables are dummy variables.

Table A.2: Estimation of health index - wave 11

variables	coefficient estimate	t-stat
mobility - some difficulty	-0.701***	(-18.97)
mobility - significant difficulty	-1.069***	(-13.40)
lifting, carrying, moving objects - some difficulty	-0.405***	(-10.77)
lifting, carrying, moving objects - significant difficulty	-0.489***	(-7.65)
manual dexterity - some difficulty	-0.173***	(-3.05)
manual dexterity - significant difficulty	-0.401***	(-3.90)

continence - some difficulty	-0.241***	(-4.81)
continence - significant difficulty	-0.339***	(-3.33)
hearing - some difficulty	-0.0565	(-0.76)
hearing - significant difficulty	-0.0403	(-0.34)
sight - some difficulty	-0.101	(-1.44)
sight - significant difficulty	-0.134	(-1.18)
communication, speech problems - some difficulty	-0.140	(-1.23)
communication, speech problems - significant difficulty	-0.455*	(-2.20)
memory, ability to concentrate, learn, understand - some difficulty	-0.333***	(-6.02)
memory, ability to concentrate, learn, understand - sig. difficulty	-0.398***	(-4.19)
recognise danger - some difficulty	0.116	(0.63)
recognise danger - significant difficulty	0.543*	(2.25)
physical coordination - some difficulty	-0.0759	(-1.29)
physical coordination - significant difficulty	0.111	(0.98)
personal care - some difficulty	-0.144	(-1.91)
personal care - significant difficulty	-0.321**	(-2.87)
other - some difficulty	-0.581***	(-14.95)
other - significant difficulty	-0.745***	(-10.50)
<hr/>		
age 20-24	-0.244***	(-4.32)
age 25-29	-0.415***	(-7.33)
age 30-34	-0.452***	(-7.73)
age 35-39	-0.531***	(-9.61)
age 40-44	-0.649***	(-12.00)
age 45-49	-0.653***	(-12.49)
age 50-54	-0.732***	(-14.15)
age 55-59	-0.692***	(-13.66)
age 60-64	-0.764***	(-14.60)
age 65-69	-0.625***	(-11.88)
age 70 and older	-0.616***	(-12.39)
female	0.117***	(6.43)
<hr/>		
asthma - ever had	0.0609	(0.29)
arthritis - ever had	0.0407	(0.26)
congestive heart failure - ever had	-0.285	(-0.58)
coronary heart disease - ever had	0.855**	(2.76)
angina - ever had	-0.175	(-0.45)
heart attack - ever had	-0.358	(-1.93)
angina - ever had	-0.0386	(-0.27)
emphysema - ever had	-0.480	(-1.15)
hypothyroidism - ever had	-0.0582	(-0.14)
chronic bronchitis - ever had	-0.532	(-1.15)
chronic liver condition - ever had	0.405	(1.33)
cancer - ever had	0.324*	(2.17)
diabetes - ever had	-0.277	(-1.25)
epilepsy - ever had	0.0659	(0.16)
high blood pressure - ever had	-0.0356	(-0.32)
other chronic condition - ever had	-0.209***	(-3.78)
multiple sclerosis - ever had	0.148	(0.51)
COPD - ever had	0.158	(0.43)
emotional, nervous, psychiatric problem - ever had	-0.252	(-1.08)
other cancer - ever had	-0.590**	(-2.90)
anxiety - ever had	-0.0187	(-0.07)
depression - ever had	-0.267	(-1.41)
asthma - still have	-0.194	(-0.89)
arthritis - still have	-0.128	(-0.78)
congestive heart failure - still have	-0.327	(-0.60)
coronary heart disease - still have	-1.019**	(-3.00)

angina - still have	-0.310	(-0.73)
hypothyroidism or underactive thyroid - still have	-0.0505	(-0.12)
chronic bronchitis - still have	0.144	(0.27)
liver condition - still have	-0.606	(-1.75)
cancer - still have	-0.884***	(-4.44)
diabetes - still have	-0.115	(-0.50)
epilepsy - still have	-0.105	(-0.20)
high blood pressure - still have	-0.134	(-1.14)
COPD - still have	-0.488	(-1.28)
emotional or nervous or psychiatric condition - still have	0.521	(1.91)
anxiety - still have	-0.0641	(-0.24)
depression - still have	0.000729	(0.00)
<hr/>		
1-2 visits to hospital outpatient in yr	-0.132***	(-6.50)
3-5 visits to hospital outpatient in yr	-0.384***	(-12.12)
6-10 visits to hospital outpatient in yr	-0.476***	(-9.93)
>10 visits to hospital outpatient in yr	-0.652***	(-9.31)
No job dummy	-0.269***	(-9.84)
professional occupation	0.0977*	(2.02)
skilled non-manual occupation	-0.154***	(-5.00)
skilled manual occupation	-0.1000**	(-2.95)
partly skilled occupation	-0.188***	(-5.45)
unskilled occupation	-0.181**	(-2.58)
<hr/>		
GHQ score - 1	0.192	(0.66)
GHQ score - 2	-0.309	(-1.30)
GHQ score - 3	-0.261	(-1.19)
GHQ score - 4	-0.424*	(-2.12)
GHQ score - 5	-0.476**	(-2.60)
GHQ score - 6	-0.531**	(-3.04)
GHQ score - 7	-0.639***	(-3.65)
GHQ score - 8	-0.720***	(-4.11)
GHQ score - 9	-0.851***	(-4.85)
GHQ score - 10	-0.962***	(-5.49)
GHQ score - 11	-1.051***	(-6.01)
GHQ score - 12	-1.163***	(-6.66)
GHQ score - 13	-1.226***	(-6.93)
GHQ score - 14	-1.221***	(-6.82)
GHQ score - 15	-1.268***	(-7.04)
GHQ score - 16	-1.354***	(-7.48)
GHQ score - 17	-1.292***	(-7.10)
GHQ score - 18	-1.433***	(-7.76)
GHQ score - 19	-1.389***	(-7.48)
GHQ score - 20	-1.402***	(-7.47)
GHQ score - 21	-1.447***	(-7.61)
GHQ score - 22	-1.634***	(-8.59)
GHQ score - 23	-1.698***	(-8.95)
GHQ score - 24	-1.497***	(-7.89)
GHQ score - 25	-1.788***	(-8.56)
GHQ score - 26	-1.864***	(-8.35)
GHQ score - 27	-1.823***	(-6.19)
GHQ score - 28	-2.013***	(-7.98)
GHQ score - 29	-1.703***	(-6.75)
GHQ score - 30	-2.027***	(-7.47)
GHQ score - 31	-1.980***	(-7.51)
GHQ score - 32	-1.922***	(-6.90)
GHQ score - 33	-2.272***	(-6.62)
GHQ score - 34	-2.033***	(-8.03)

GHQ score - 35	-1.946***	(-4.30)
GHQ score - 36	-1.527***	(-3.66)
Education - GSCEs	0.0511*	(1.96)
Education - Yr12	0.125***	(3.66)
Education - Degree	0.275***	(10.81)
pregnant	0.174	(1.70)
<i>N</i>	24,987	

*t* stats in parentheses, exclude reporting time dummies coefficients, \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\*  $p < 0.001$

### A.0.3 Genetic data

Polygenic scores (PGS) are scores constructed using genetic data that estimate an individual’s propensity to express a phenotype, which is an observable trait. They are calculated from genome-wide association studies (GWAS), which are systematic analyses of genetic variation across the entire human genome and their association with various phenotypes. I make use of the latest (2022) version of the ELSA polygenic scores (Ajnakin and Andrew Steptoe, 2022) and select some of them to capture health conditions with the biggest disease burdens, which I report in Table A.3. I use these polygenic scores to create two health aggregates; one capturing chronic physical health conditions, the other mental health. The correlation between the two indices is small.

Table A.3: Polygenic score aggregation

Physical index	Mental index
Coronary artery disease (2016)	Alzheimer’s disease (2019)
Type II diabetes (2018)	Depressive Symptoms
Rheumatoid arthritis	Major depressive disorder (2018)
Myocardial infarction	Anxiety (case-control)
Migrane (2016)	Schizophrenia (2020)
Chronic pain	Bipolar disorders (2021)
Waist-hip-ratio	Subjective wellbeing
	Loneliness

My method of aggregation is identical to how I aggregate the biomarker data. I normalise each PGS and then aggregate them. I consider alternate aggregation methods, including factor analysis and converting each PGS into a binary variable

with the highest 10-20 per cent of scores coded as '1', however the resulting indices are all quite similar. To assess the predictive value of my PGS indices, I regress a health index constructed using ELSA data, that is designed to be as similar as possible to my main health index constructed using Understanding Society data, against normalised polygenic scores for all the PGS that make up my indices for mental and physical health. I report the results of this exercise in Table A.4. The major depressive disorder PGS has the highest predictive power although the predictive power of many of the eight mental health PGS are quite similar. The predictive power of the physical PGS are much more varied, with chronic pain being by far the most important.

Table A.4: How well individual polygenic scores predict health index

z-score of PGS	components of	
	mental index	physical index
z score: depressive symptoms	-0.0302*** (0.00486)	
z score: major depressive disorder	-0.0521*** (0.00515)	
z score: anxiety	0.00763 (0.00490)	
z score: schizophrenia	0.0192*** (0.00575)	
z score: bipolar	-0.00397 (0.00468)	
z score: subjective well-being	0.0119** (0.00433)	
z score: Alzheimer's	-0.0170*** (0.00394)	
z score: loneliness	-0.0340*** (0.00490)	
z score: arthritis		0.00911* (0.00426)
z score: coronary heart disease		0.00193 (0.00401)
z score: diabetes		-0.0342*** (0.00409)
z score: chronic pain		-0.112*** (0.00416)
z score: myocardial infarction		-0.0192*** (0.00408)
z score: waist-hip ratio		-0.0121** (0.00394)
z score: migraines		-0.00235 (0.00378)
<i>N</i>	37,543	37,546

I then repeat this exercise, but with the PGS scores aggregated into two indices that capture mental and physical health, and report the results in Table A.5. I find that while both indices based on genetic data are significantly correlated with the level of health index and its variance over time, and contain additional information not captured by lagged health index terms or biomarkers, the size of the coefficients are very small. I conclude that the coefficient sizes are too small to be a useful addition to modelling the overall dynamics of the health index. However, a more granular approach may be more effective. This could include using only individual PGS with higher predictive power or for diseases with high incidence rates such as diabetes and depression rather than also including rarer conditions such as schizophrenia.

Table A.5: How well aggregated polygenic scores predict health index

	level of health index			variance of health index		
	(1)	(2)	(3)	(4)	(5)	(6)
mental index	-0.00746*** (-8.42)	-0.00125* (-2.10)	-0.00768*** (-6.31)	0.000810*** (4.48)	0.000392 (1.93)	0.000654* (2.52)
physical index	-0.0229*** (-18.28)	-0.00476*** (-5.42)	-0.0187*** (-10.71)	0.00165*** (6.24)	0.000809** (2.63)	0.00136*** (3.56)
mental index <sup>2</sup>	-0.000496*** (-4.73)	-0.000187* (-2.45)	-0.000463** (-3.19)	0.0000500* (2.37)	0.0000226 (0.93)	0.0000302 (1.01)
physical index <sup>2</sup>	-0.000519* (-2.05)	0.0000618 (0.34)	-0.000417 (-1.18)	0.0000270 (0.49)	0.0000366 (0.57)	0.00000785 (0.10)
lagged health index		0.840*** (201.75)			-0.0485*** (-29.36)	
allostatic index			-0.385*** (-18.59)			0.0330*** (10.38)
<i>N</i>	37,543	25,256	18,304	37,412	25,256	18,203

*t* statistics in parentheses; also control for age and gender, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

#### A.0.4 Replication exercise using non-detrended health data

I replicate the preferred baseline ARMA(1,1) model estimated using health data that has not been detrended by age and sex, and a difference-GMM specification. This exercise shows that the persistence estimates are robust to being detrended by age and sex; the slightly higher persistence estimates in this case are due to a gradual decline in health as people age being captured by the persistence estimates.

Table A.6: Estimation of ARMA(1,1) model with non-detrended health index

	health index
L.health.index	1.038*** (0.0227)
Hansen J test stat	2.756
Hansen J p value	0.431
$N$	222,095

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

### A.0.5 ARMA(p,q) groups from clustering analysis

The partition clustering algorithm makes an initial partition of individuals into clusters based on the number of desired clusters, and then reallocates individuals until the final partition minimises the residual sum of the squared objective function. The optimal number of clusters for a given model is selected as the one with the lowest model information criterion (MIC). I describe the main characteristics of the two clusters I obtain after performing clustering analysis in Table A.7. Relative to group 2, group 1 contains individuals with worse health on average. The individuals in the two groups are just as likely to experience a large negative health shock of at least one standard deviation (6.6 and 6.7 per cent of observations respectively), but the second group is more likely to experience a positive shock (5.3 versus 6.4 per cent).

Table A.7: Group characteristics\*

	Group 1	Group 2
mean health index	0.024	0.134
median health index	0.186	0.254
standard deviation of health indices	0.670	0.565
mean allostatic score	0.023	-0.046
$N$	28,644	44,231

\* Only include individuals where T=12

I then replicate my difference and system GMM specifications I used for estimating the ARMA(p,q) specification for my main health index for each group separately in Table A.8.

Table A.8: GMM estimates of the health process, by cluster

	Difference GMM								System GMM	
	group 1				group 2				group 1	group 2
	MA(0)	MA(0)	MA(1)	MA(1)	MA(0)	MA(0)	MA(1)	MA(1)	MA(0)	MA(1)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
L.health	0.709*** (0.0203)	0.791*** (0.0218)	0.898*** (0.0311)	0.922*** (0.100)	0.104*** (0.0119)	0.187*** (0.0172)	0.997*** (0.132)	0.893*** (0.164)	0.742*** (0.00819)	0.828*** (0.0461)
L2.health		0.0876*** (0.0124)		-0.0207 (0.0822)		0.0737*** (0.0127)		-0.00265 (0.0236)	0.0807*** (0.0106)	
AB test, order 1 z score	-26.87	-27.77	-21.04	-6.27	-35.4	-31.92	-9.46	-6.99	-33.00	-19.15
AB test, order 1 p value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
AB test, order 2 z score	3.44	-1.21	3.59	1.27	3.36	-1.46	6.02	4.2	-1.11	9.39
AB test, order 2 p value	0.001	0.226	0.000	0.203	0.001	0.145	0.000	0.000	0.266	0.000
Hansen J test statistic	63.74	3.39	1.47	1.37	78.77	43.01	1.71	7.19	12.74	4.09
Hansen J p value	0.000	0.495	0.689	0.712	0.000	0.000	0.634	0.066	0.047	0.394
Observations	28,644	26,040	28,644	26,040	44,231	40,210	44,231	40,210	26,040	44,231

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

This exercise shows that Group 1 is best modelled as an AR(2) process and group 2 as an ARMA(1,1) process. Table A.9 reports the results of estimating the MA(1) coefficient for group 2, which I estimate to be -0.54.

Table A.9: GMM estimation of MA(1) coefficient - group 2

$\rho = 0.83$	
$\eta_i$	0.0851*** (0.00280)
$\theta$	-0.544*** (0.0249)
$\varepsilon_{it}$	0.289*** (0.00411)
$N$	44,231
Standard errors in parentheses, * $p < 0.05$ , ** $p < 0.01$ , *** $p < 0.001$	

## A.0.6 ARCH model estimates

An autoregressive conditional heteroskedasticity (ARCH) model allows us to account for whether individuals who have recently experienced a health shock, captured in the model as a large  $\varepsilon_{it}$  term, are more or less likely to experience additional health shocks in the subsequent period. I estimate an ARCH(1) model with the following exponential variance specification:  $\sigma_{it}^2 = \exp(\gamma_0 + \gamma_1 \varepsilon_{i,t-1}^2 + \gamma_2 \varepsilon_{i,t-1})$ . The  $\gamma_2$  term allows for asymmetry between negative and positive shocks. To estimate ARCH effects using GMM I use the following moment condition derived by Arellano (1995). For robustness I also estimate the more general specification by Meghir and Windmeijer (1999) and obtain similar results.

$$\mathbb{E} \left[ h_{i,t-k} \left( \varepsilon_{i,t-1}^2 - \frac{\varepsilon_{i,t}^2 (1 + \sigma_{it-1}^2)}{(1 + \sigma_{i,t}^2)} \right) \right] = 0, k = 1 \dots t - 3$$

An ARCH effect exists if the  $\gamma_1$  and  $\gamma_2$  terms on the lagged error terms in the variance specification ( $\sigma_{it}^2$ ) are significantly different from zero. I report my results in Table A.10. I find that while the point estimates for the  $\gamma$  terms are reasonable, and suggest that a person who experiences a large health shock has more volatile health next period, especially if they suffer a negative health shock, the estimates are not

statistically significant. I conclude that there is no evidence for ARCH effects in the data.


Table A.10: GMM estimates of ARCH(1) effect

variables	estimates
$\gamma_0$	-2.5176 (5.7445)
$\gamma_1$	0.3301 (0.5848)
$\gamma_2$	-0.1648 (0.8703)

standard errors in parentheses, \* =  $p < 0.05$

I also consider shock heterogeneity in a different way. Figure A.1 graphs the relationship between the magnitude of a health shock in period  $t$  and the extent of recovery/mean-reversion five years later. I plot the initial health shock ( $h_{i,t} - h_{i,t-1}$ ) on the x axis and the difference between health at period  $t$  and five years later ( $h_{i,t+5} - h_{i,t}$ ) on the y axis. A steeper downward-sloping line suggests a faster rate of mean-reversion, while a flat line would indicate that there has been no change in health between period  $t$  and period  $t + 5$ . I graph this relationship for each sextile based on average health in the period prior to the shock ( $t + 1$ ). The design of this graph was adapted from an earnings dynamics graph by Guvenen et al. (2021). Figure A.1 shows that there is significantly less recovery from negative shocks than would be predicted by an ARMA(1,1) model with normally distributed white-noise shocks, which would achieve around a 60 per cent mean revision, especially among those in persistently poor health prior to the negative shock. The health dynamics of those whose health is in the bottom sixth of the sample are much less well captured by an ARMA(1,1) model relative to those in the upper two-thirds. This exercise suggests that simple ARMA models are adequate for approximating the health process for the healthier section of the population, but are much less accurate for those with a history of poor health.

Figure A.1



### A.0.7 ARMA models with additional allostatic regressors

Adding allostatic scores as an additional regressor, or as a dummy variable for whether individuals have ‘bad’ allostatic scores interacted with the lagged health term, has very little impact on my linear estimates of health. Table A.11 reports the regression output if I include these additional terms in my preferred specification.

Table A.11: Health index estimates with additional allostatic score regressors

	(1)	(2)
L.health index	0.890*** (0.0286)	0.876*** (0.0218)
allostatic scores	0.203 (0.302)	-0.348 (0.178)
L.health index $\times$ allostatic scores		-0.131 (0.0711)
Observations	70,039	70,039

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

### A.0.8 Additional non-linear health persistence results

I report the full results of re-estimating the non-linear persistence model when the sample is split into two sub-groups based on allostatic score in Tables A.12 and A.13. Table A.14 reports output from the same estimation process, where I separately estimate the health persistence for individuals with higher and lower allostatic scores. However, this time I only report the persistence results for the persistent component ( $\rho_t(\tau)$ ) rather than the entire index. I do not control for fixed effects in either set of results.

Table A.12: Non-linear persistence estimates: low allostatic scores sub-sample (healthy)

$health_{t-1}^*$	Shock size percentiles*										
	1	2	3	4	5	6	7	8	9	10	11
1	1.15	1.09	1.04	0.98	0.92	0.88	0.82	0.76	0.71	0.64	0.55
2	1.08	1.02	0.98	0.93	0.87	0.84	0.79	0.74	0.70	0.63	0.54
3	1.02	0.97	0.93	0.89	0.84	0.81	0.78	0.73	0.69	0.63	0.55
4	0.98	0.92	0.89	0.86	0.82	0.79	0.77	0.72	0.69	0.63	0.55
5	0.93	0.89	0.85	0.83	0.80	0.78	0.76	0.72	0.69	0.63	0.56
6	0.89	0.85	0.82	0.80	0.78	0.77	0.75	0.72	0.68	0.63	0.57
7	0.85	0.81	0.79	0.78	0.76	0.75	0.74	0.72	0.68	0.64	0.58
8	0.81	0.77	0.75	0.75	0.74	0.74	0.74	0.71	0.68	0.64	0.59
9	0.76	0.73	0.72	0.72	0.72	0.73	0.73	0.71	0.69	0.65	0.60
10	0.70	0.68	0.67	0.69	0.70	0.71	0.72	0.71	0.69	0.65	0.61
11	0.62	0.61	0.61	0.64	0.67	0.69	0.71	0.71	0.69	0.66	0.63

\*1=most negative, 11=most positive

Table A.13: Non-linear persistence estimates: high allostatic score sub-sample (un-healthy)

$health_{t-1}^*$	Shock size percentiles*										
	1	2	3	4	5	6	7	8	9	10	11
1	1.09	1.06	1.05	1.02	1.01	0.99	0.97	0.93	0.90	0.85	0.77
2	1.10	1.07	1.05	1.02	0.99	0.96	0.94	0.90	0.85	0.80	0.72
3	1.08	1.05	1.02	0.99	0.96	0.93	0.91	0.87	0.82	0.77	0.69
4	1.07	1.02	1.00	0.96	0.93	0.90	0.88	0.84	0.79	0.74	0.66
5	1.05	1.00	0.97	0.94	0.90	0.88	0.85	0.81	0.77	0.72	0.64
6	1.03	0.98	0.95	0.92	0.88	0.85	0.83	0.79	0.74	0.70	0.63
7	1.02	0.96	0.93	0.90	0.86	0.83	0.81	0.77	0.73	0.68	0.62
8	1.00	0.94	0.90	0.87	0.84	0.81	0.78	0.75	0.71	0.66	0.60
9	0.98	0.92	0.88	0.85	0.81	0.79	0.76	0.73	0.69	0.64	0.59
10	0.96	0.90	0.85	0.82	0.79	0.76	0.73	0.71	0.67	0.62	0.58
11	0.93	0.87	0.81	0.79	0.75	0.73	0.70	0.67	0.64	0.60	0.56

\*1=most negative, 11=most positive

Table A.14: Estimates of coefficient of persistent component  $\rho_t(\tau)$

	Shock size percentiles*										
	1	2	3	4	5	6	7	8	9	10	11
<i>health</i> <sub><i>t</i>-1</sub> *	Sub-sample with below-average (good) allostatic scores										
1	1.48	1.35	1.25	1.15	1.05	0.97	0.91	0.86	0.80	0.68	0.39
2	1.35	1.21	1.13	1.04	0.96	0.90	0.85	0.81	0.76	0.66	0.46
3	1.24	1.11	1.02	0.96	0.89	0.85	0.80	0.76	0.72	0.64	0.50
4	1.15	1.02	0.95	0.89	0.84	0.80	0.76	0.73	0.69	0.63	0.53
5	1.08	0.95	0.88	0.84	0.80	0.77	0.73	0.70	0.67	0.61	0.55
6	1.02	0.90	0.83	0.79	0.76	0.74	0.71	0.68	0.65	0.60	0.56
7	0.96	0.85	0.79	0.76	0.73	0.71	0.68	0.66	0.63	0.59	0.58
8	0.91	0.80	0.74	0.72	0.70	0.69	0.66	0.64	0.61	0.58	0.59
9	0.86	0.75	0.70	0.68	0.67	0.66	0.64	0.62	0.60	0.57	0.60
10	0.80	0.69	0.65	0.64	0.64	0.63	0.61	0.59	0.58	0.56	0.61
11	0.70	0.61	0.57	0.57	0.58	0.59	0.57	0.56	0.55	0.54	0.63
	Subsample with above-average (bad) allostatic scores										
1	1.35	1.25	1.16	1.03	0.96	0.92	0.89	0.83	0.78	0.68	0.52
2	1.26	1.19	1.11	1.00	0.94	0.91	0.89	0.83	0.79	0.71	0.58
3	1.17	1.13	1.05	0.96	0.91	0.89	0.87	0.83	0.79	0.73	0.61
4	1.08	1.07	1.00	0.93	0.89	0.88	0.86	0.82	0.79	0.73	0.64
5	1.01	1.01	0.96	0.90	0.87	0.86	0.85	0.82	0.79	0.74	0.66
6	0.94	0.96	0.92	0.87	0.85	0.85	0.84	0.81	0.79	0.74	0.67
7	0.88	0.92	0.88	0.85	0.84	0.84	0.83	0.81	0.79	0.75	0.69
8	0.82	0.87	0.84	0.82	0.82	0.83	0.82	0.81	0.79	0.75	0.70
9	0.75	0.82	0.80	0.80	0.80	0.81	0.81	0.80	0.79	0.75	0.72
10	0.67	0.76	0.75	0.76	0.78	0.80	0.80	0.79	0.78	0.75	0.73
11	0.53	0.66	0.67	0.71	0.74	0.77	0.77	0.78	0.78	0.76	0.76

\*1=most negative, 11=most positive

### A.0.9 Additional mental health persistence results

Table A.15 reports the coefficient estimates for the lagged GHQ term when estimating an ARMA(1,1) model for the mental health index using GMM.

Table A.15: ARMA(1,1) model of GHQ scores

	Diff-GMM	Sys-GMM
L.ghq	0.610*** (0.0508)	0.693*** (0.0388)
AB test, order 1, z score	-17.27	-22.72
AB test, order 1, p value	0.000	0.000
AB test, order 2, z score	8.960	11.42
AB test, order 2, p value	0.000	0.000
AB test, order 3, z score	0.158	0.262
AB test, order 3, p value	0.874	0.793
Hansen J test stat	2.223	7.41
Hansen J test p value	0.527	0.115
Observations	222,095	222,095

Standard errors in parentheses, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$   
Incorporating additional lagged terms found to be insignificant

I now report some additional results from the non-linear persistence estimates of mental health. First, I report the persistence estimates in Table A.16 that accompany the graphs I include in section 6 (Figure 2.12). These data are the overall mental health persistence estimates, rather than the estimates for the persistent component only. Second, I illustrate the impact of taking fixed effects into account when calculating mental health persistence in Figure A.2. These calculations consider the persistent component only.

Table A.16: Mental health persistence estimates



$health_{t-1}^*$	Shock size percentiles*										
	1	2	3	4	5	6	7	8	9	10	11
1	1.09	1.00	0.90	0.78	0.66	0.56	0.45	0.34	0.26	0.18	0.06
2	1.08	0.96	0.86	0.76	0.66	0.58	0.49	0.40	0.33	0.25	0.13
3	1.02	0.90	0.81	0.73	0.66	0.59	0.52	0.44	0.38	0.30	0.19
4	0.97	0.86	0.78	0.71	0.65	0.60	0.54	0.48	0.41	0.34	0.24
5	0.92	0.82	0.74	0.69	0.65	0.61	0.56	0.50	0.45	0.37	0.28
6	0.87	0.77	0.71	0.67	0.65	0.62	0.58	0.53	0.48	0.40	0.32
7	0.81	0.72	0.67	0.65	0.64	0.62	0.60	0.56	0.51	0.44	0.36
8	0.75	0.67	0.63	0.63	0.63	0.63	0.62	0.59	0.54	0.47	0.41
9	0.69	0.62	0.60	0.61	0.63	0.64	0.64	0.62	0.57	0.51	0.46
10	0.63	0.56	0.55	0.59	0.62	0.65	0.66	0.65	0.61	0.54	0.50
11	0.54	0.48	0.50	0.55	0.61	0.66	0.69	0.70	0.65	0.59	0.57

\*1=most negative, 11=most positive

Figure A.2: Persistence of GHQ index

(a) Persistent component only

(b) Persistent component + fixed effect



**B.0.1 PCA analysis for health category classification**

C (component) 1 maps onto physical disabilities, C2 maps onto mental health, C3 maps onto chronic health conditions.

Table B.1: Principal component analysis

	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10
Problems - arms or hands	0.550	-0.178	-0.117	-0.036	-0.005	-0.018	0.047	0.017	-0.006	0.048
Problems - legs or feet	0.565	-0.151	-0.069	-0.045	0.001	-0.025	0.041	0.002	0.001	0.029
Problems - back or neck	0.520	-0.117	-0.112	-0.002	-0.048	-0.009	0.019	0.040	-0.016	-0.022
Difficulty in seeing	0.097	0.102	0.161	0.031	0.263	0.136	-0.593	-0.146	0.373	0.319
Difficulty in hearing	0.145	0.140	0.183	0.096	0.239	0.093	-0.491	-0.039	0.020	-0.440
A speech impediment	0.054	0.106	-0.001	0.124	0.611	-0.112	0.106	-0.198	-0.310	-0.052
Skin conditions, allergies	0.138	0.275	0.197	0.505	-0.232	-0.092	0.077	-0.082	0.082	0.115
Chest/breathing problems	0.110	0.281	0.247	0.470	-0.238	-0.125	0.154	-0.113	0.083	0.129
Heart, blood pressure	0.141	0.205	0.482	-0.367	-0.015	-0.109	0.119	-0.022	-0.002	-0.064
Stomach, liver, kidney, digestive	0.098	0.259	0.175	0.009	-0.131	0.352	0.058	0.174	-0.255	-0.571
Diabetes	0.052	0.134	0.421	-0.500	0.019	-0.253	0.170	-0.141	0.108	0.121
Depression, bad nerves, anxiety	0.086	0.538	-0.364	-0.178	-0.122	0.007	-0.050	-0.038	0.038	0.001
Depression, bad nerves or anxiety	0.009	0.083	-0.019	0.059	0.317	0.123	0.346	0.524	0.668	-0.127
Learning difficulties	0.023	0.210	-0.098	0.153	0.488	-0.238	0.254	-0.018	-0.202	0.075
Mental illness, phobia, panics	0.033	0.479	-0.456	-0.220	-0.080	-0.044	-0.049	-0.113	0.093	0.024
Other progressive illness	0.036	0.051	0.066	-0.051	0.101	0.808	0.294	-0.316	-0.050	0.292
other	0.044	0.193	0.122	-0.041	0.016	0.092	-0.199	0.691	-0.416	0.465
Proportion explained	0.110	0.081	0.071	0.065	0.063	0.059	0.059	0.058	0.058	0.056
Cumulative proportion	0.110	0.191	0.262	0.327	0.390	0.449	0.508	0.566	0.624	0.681

Only report first 10 components, so omit C11-17

## B.0.2 Additional summary statistics

We provide some additional information on the distribution of occupation changes. We graph the distribution of occupation content changes in Figure B.1. The correlations between task changes is as follows: the correlation between cognitive and interpersonal is 0.60, cognitive and manual is 0.42 and manual and interpersonal is -0.17.

Figure B.1: Histogram of occupational content changes





Table B.2: Labour market transitions by occupation and health shock type, share of total

	change occupation				change job				stop working			
	Disability	Mental	Chronic	Healthy	D	M	C	H	D	M	C	H
corporate managers and directors	19.2	18.8	17.9	14.2	6.5	6.3	6.2	5.2	2.3	2.0	2.1	1.5
other managers and proprietors	15.1	15.5	13.3	13.3	4.7	5.4	4.8	4.4	3.4	4.3	3.0	2.2
science, research, engineering and tech prof.	16.7	19.6	16.1	14.2	4.4	6.1	6.3	4.9	3.0	4.0	2.1	1.4
health professionals	3.8	4.6	4.1	3.7	4.7	7.0	6.3	4.8	2.4	2.1	2.0	1.2
teaching and educational prof.	10.6	10.4	10.4	8.9	4.2	7.7	5.8	4.4	5.1	3.2	3.7	2.4
business, media and public service prof.	12.8	17.8	12.8	11.6	4.9	6.9	6.2	5.2	2.2	3.3	1.9	1.7
science, engineering and tech. associate prof.	18.8	15.0	18.4	16.3	4.6	6.6	5.4	4.4	2.1	7.5	2.6	1.7
health and social care associate prof.	15.9	11.3	16.5	13.1	4.3	4.4	5.7	4.1	3.2	3.3	3.2	2.2
protective service occupations	6.8	4.8	7.3	5.9	2.6	2.4	2.9	2.6	3.4	2.4	1.7	0.9
culture, media and sports occupations	7.7	5.1	8.2	7.5	4.2	4.4	4.0	4.4	4.6	7.0	4.7	3.4
business and public service associate prof.	14.7	16.7	14.9	13.8	5.1	6.6	4.9	5.2	2.9	3.4	2.7	2.1
administrative occupations	15.5	15.0	14.0	13.2	5.6	5.4	6.0	4.5	3.7	4.3	3.3	2.7
secretarial and related occupations	9.1	8.9	11.0	7.7	4.6	8.6	6.0	4.6	5.0	7.9	4.1	2.4
skilled agricultural and related trades	4.0	6.5	5.0	4.6	1.0	3.2	1.9	2.5	4.2	14.5	5.5	1.6
skilled metal, electrical and electronic trades	12.4	14.3	14.4	10.7	4.4	6.0	5.2	4.5	3.2	7.2	2.6	1.5
skilled construction and building trades	6.4	12.1	7.7	5.9	3.4	5.1	4.4	3.7	3.0	4.7	3.5	1.8
textiles, printing and other skilled trades	8.5	9.9	9.9	7.5	4.6	5.8	6.5	5.3	5.4	7.4	3.3	2.5
caring personal service occupations	9.5	9.4	8.9	8.1	6.2	8.2	6.5	5.3	4.7	5.5	4.3	3.0
leisure, travel and related personal service	9.5	5.6	8.9	7.5	4.0	5.2	5.0	4.6	3.2	4.8	4.0	3.1
sales occupations	11.9	13.2	12.6	11.0	6.6	8.5	6.9	6.1	5.3	8.2	5.6	4.3
customer service occupations	17.4	17.6	17.9	16.5	6.1	6.8	8.6	5.7	3.9	6.5	3.2	3.1
process, plant and machine operatives	17.2	17.3	14.9	14.6	4.4	7.1	4.8	4.7	3.7	7.1	3.2	2.4
transport/mobile machine drivers/operatives	8.0	8.1	7.1	6.7	4.9	6.5	5.1	5.2	2.8	4.2	3.0	1.8
elementary trades and related occupations	17.4	15.8	14.3	14.0	6.0	7.9	7.3	6.8	6.5	15.8	8.9	4.7
elementary administration and service	7.4	10.8	8.7	8.3	5.6	9.3	6.4	6.2	5.2	8.8	4.9	4.6

Figure B.2: Jobs change likelihood by prior occupation - difference from healthy individuals




### B.0.3 Additional analysis of hours changes

We firstly graph the distribution of the reported changes in hours between two quarters for those who suffer a health shock and for those who remain healthy in Figure B.3. Those who suffer a health shock are more likely to increase or decrease their hours relative to those who remain healthy.

We then estimate the mean change in hours worked following a health shock by replicating the system-GMM specification from a prior section of this paper. We estimate the following equation only using the individuals in our sample who change

Figure B.3



occupation, change jobs, as well as the full sample. We do not use the first lag of the hours variable as an instrument as including it leads to a strong rejection of the null of the Hansen J-test, while only including the second and later lags does not, suggesting that the errors follow an MA(1) process. We also include a second lag of the dependent variable when we use the full sample for our estimation, as hours are much less persistent for the sub-sample that changes occupation relative to the full sample.

$$\begin{aligned} \text{weekly hours}_{i,t} = & \beta_1^D \text{disability shock}_{i,t} + \beta_1^M \text{mental shock}_{i,t} + \beta_1^C \text{chronic shock}_{i,t} + \\ & \beta_2^D \text{disability shock}_{i,t-1} + \beta_2^M \text{mental shock}_{i,t-1} + \beta_2^C \text{chronic shock}_{i,t-1} + \\ & \gamma_1 \text{weekly hours}_{i,t-1} + \gamma_2 \text{weekly hours}_{i,t-2} + \varepsilon_{it}. \end{aligned} \tag{B.1}$$

We find that the sub-sample that changed occupation or job following a physical disability shock or mental health shock also, on average, worked fewer hours in their new occupation than those who changed occupation while remaining healthy. When we repeat our estimation using the full sample, only the coefficient on the lagged physical disability shock term continues to be negative and significant. However,

there is a risk that this specification is vulnerable to selection bias because our lagged dependent variable is correlated with the selection equation. Those who work fewer hours are more likely to stop working following a health shock. One way to potentially reduce this bias would be estimate full time and part time workers separately; we are not able to do this because our sample becomes too small.

Table B.3: Average weekly hours - system GMM

	(1) Occupation change sub-sample	(2) Employer change sub-sample	Full sample
physical disability shock	-1.0684* (0.5836)	-6.6858*** (2.5892)	0.0249 (0.1249)
L.physical disability shock	-0.6131*** (0.2284)	0.3725 (0.6602)	-0.1524** (0.0665)
mental health	0.0793 (1.2951)	2.5348 (3.5931)	-0.2490 (0.2388)
L.mental health	-1.0143** (0.4409)	-2.2415*** (0.8614)	-0.0075 (0.1567)
chronic	0.2941 (0.4956)	1.8911 (1.4073)	-0.0135 (0.1067)
L.chronic	0.3050 (0.2013)	-0.3354 (0.4599)	-0.0624 (0.0554)
L.hours	0.7302*** (0.1046)	0.1356 (0.1120)	1.1761*** (0.0958)
L2.hours		-0.2424*** (0.0596)	
Hansen J-stat p val	0.3162	0.4320	0.0646
<i>N</i>	73405	33097	366768

Standard errors in parentheses, \* p<0.1, \*\* p<0.05, \*\*\* p<0.01

#### B.0.4 Additional robustness checks

As a robustness check, we re-estimate our main occupation change specification using OLS rather than GMM. We regress cognitive, manual, and interpersonal content of jobs in period  $t$  against dummies for suffering a health shock in the past six months, controlling for job characteristics pre-shock. To reduce the risk of dynamic panel (Nickell) bias, we run several specifications that control for job characteristics pre-shock. The three specifications are: do not control for lagged cognitive, manual, and

interpersonal intensity (specification 1), include lagged occupation codes at the two digit level (25 categories) instead of lagged dependent variables (specification 2), and include lagged cognitive, manual, and interpersonal intensity as lagged dependent variables (specification 3). These are reported in Table B.4.

Table B.4: Occupation content OLS regressions - full sample

	Spec 1	2 cognitive	3	1	2 interpersonal	3	1	2 manual	3
disab	-0.0046 (0.0033)	-0.0010 (0.0020)	0.0001 (0.0014)	-0.0001 (0.0028)	0.0000 (0.0018)	0.0003 (0.0012)	0.0041 (0.0032)	0.0015 (0.0022)	0.0009 (0.0012)
mental	-0.0121** (0.0048)	-0.0104*** (0.0030)	-0.0034* (0.0020)	0.0010 (0.0042)	0.0004 (0.0028)	-0.0014 (0.0018)	-0.0163*** (0.0047)	-0.0070** (0.0031)	-0.0022 (0.0016)
chronic	0.0004 (0.0029)	0.0001 (0.0017)	0.0015 (0.0011)	0.0039 (0.0026)	0.0024 (0.0016)	0.0021** (0.0010)	0.0000 (0.0030)	0.0007 (0.0020)	0.0012 (0.0011)
<i>N</i>	218386	216350	216350	218386	216350	216350	218386	216350	216350

Standard errors in parentheses, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$   
 Controls: age, sex, educ, ethnicity, time, lagged job traits

## B.0.5 Additional specifications

We consider two additional specifications that relate to pay and gender. First, the various channels by which health shocks influence subsequent labour market transitions will also impact pay. Due to the limited pay data in the LFS, we focus on comparing the average salaries of occupations, rather than an individual's wage trajectory following a health shock, and estimate an occupation's average salary and its standard deviation in Table B.5. These regressions only includes individuals who changed occupations.

Second, we report our main specification re-estimated by sex in Table B.6.

## B.0.6 Selection effects

It is possible that those who stop working bias the estimates of cognitive, manual, content, especially if the lagged dependent variables are correlated with selection in our GMM specifications. To assess the significance of this risk, we regress a dummy variable that captures whether an individual stops working in period  $t$  on lags of interpersonal, manual, cognitive task intensity, lagged hours worked, as well as my standard demographic controls such as age, sex, and education level and report the

Table B.5: System GMM: average wage and standard deviation of new occupation

	(1) occupation ave. salary	(2) occupation st. dev. salary
disability shock	-0.0556** (0.0229)	-0.0213** (0.0107)
L.disability shock	-0.0289** (0.0147)	-0.0033 (0.0032)
mental health shock	-0.0205 (0.0436)	0.0072 (0.0152)
L.mental health shock	-0.0584** (0.0246)	-0.0108** (0.0050)
chronic health shock	-0.0170 (0.0243)	0.0123 (0.0098)
L.chronic health shock	0.0159 (0.0111)	-0.0002 (0.0029)
Hansen J stat p-value	0.5352	0.4316
<i>N</i>	38,294	40,158

Standard errors in parentheses, \* p<0.1, \*\* p<0.05, \*\*\* p<0.01  
Additional controls: first two lags of dependent variable

results in Table B.7. We find that there is no relationship between lagged occupation task intensity and stopping work. This is a good result for our main occupation content change specification. While this finding does not fully absolve us from the risk of selection bias, it does reduce it.

A related concern is attrition rates from the sample following a health shock. We report attrition rates between waves 1-5 of the survey by health shock. We do observe a small uptick in attrition rates for those who suffer a health shock, especially those with mental health conditions.

Table B.6: Changes in occupation content, by sex, difference GMM

	(1)	(2)	(3)	(4)	(5)	(6)
	cognitive		interpersonal		manual	
	men	women	men	women	men	women
physical disability	-0.0116 (0.0078)	-0.0058 (0.0059)	-0.0019 (0.0071)	-0.0023 (0.0062)	-0.0108 (0.0091)	-0.0044 (0.0070)
L.physical disability	-0.0085* (0.0046)	-0.0065* (0.0039)	-0.0017 (0.0045)	-0.0026 (0.0041)	-0.0037 (0.0057)	-0.0071* (0.0037)
mental	-0.0088 (0.0159)	-0.0246** (0.0108)	0.0166 (0.0150)	-0.0066 (0.0079)	-0.0180 (0.0188)	-0.0197** (0.0083)
L.mental	-0.0043 (0.0096)	-0.0099 (0.0065)	0.0104 (0.0093)	-0.0067 (0.0047)	-0.0038 (0.0106)	-0.0065 (0.0048)
chronic	0.0155** (0.0072)	0.0020 (0.0052)	0.0077 (0.0054)	0.0025 (0.0053)	0.0019 (0.0067)	0.0015 (0.0044)
L.chronic	0.0077* (0.0044)	0.0043 (0.0035)	0.0025 (0.0037)	0.0021 (0.0035)	-0.0007 (0.0043)	0.0031 (0.0029)
Hansen J stat p-value	0.9299	0.4068	0.4698	0.9528	0.9077	0.4045
<i>N</i>	74,620	74,183	74620	74,183	74,620	74,183

Standard errors in parentheses, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table B.7: Likelihood of stopping work

	(1)
L.cognitive	-0.0043 (0.0038)
L2.cognitive	0.0043 (0.0053)
L3.cognitive	0.0010 (0.0045)
L.manual	0.0039 (0.0035)
L2.manual	0.0007 (0.0044)
L3.manual	-0.0027 (0.0036)
L.interpersonal	0.0038 (0.0042)
L2.interpersonal	-0.0064 (0.0061)
L3.interpersonal	0.0015 (0.0051)
L.hours	0.0001** (0.0001)
L2.hours	0.0001 (0.0001)
L3.hours	-0.0001* (0.0000)
age	0.0079** (0.0031)
age2	-0.0265** (0.0111)
age3	0.0408** (0.0174)
age4	-0.0243** (0.0098)
sex	0.0002 (0.0005)
education	-0.0018*** (0.0005)
pay in period 1	0.0027*** (0.0007)
<i>N</i>	188,966

Standard errors in parentheses, \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table B.8: Attrition rate; by health shock type

	t-1	t (shock)	t+1	t+2	t+3
healthy	100	97.3	95.9	94.9	94.1
new disability at t	100	96.6	94.8	93.6	91.8
new mental health condition at t	100	95.2	93.1	91.4	91.0
new chronic health condition at t	100	96.8	95.2	94.4	93.0

## **C.0.1 Additional summary statistics**


### **C.0.1.1 Price variation**

I report additional information on house price increases and price changes in the RPI categories I use for my analysis in Figure C.1. The house price changes graph reports total change to the Nationwide housing index by region, which I deflated using the RPIX. This shows that house prices increased above inflation (measured using RPIX) across the UK. I also report total percentage changes to the RPI series I use over my sample period in the right hand panel of this figure. Note I report the full set of IFS-derived RPI indices, and so the housing index I use for my demand estimation is different to the RPI-housing index.

### **C.0.1.2 Trends in non-housing consumption & tenure**

I report additional statistics on non-housing consumption and homeownership by age cohorts in Figure C.2. Both indicate a substantial decline for younger cohorts relative to older cohorts. While real consumption excluding housing continues to be consistent with a life-cycle model, peaking around the age of 45–50 before declining, the real non-housing expenditure of younger cohorts born in the 1980s and 1990s is a little lower than older cohorts in levels. This must partially reflect the large increase in consumption expenditure by renters.

Figure C.1



This decline in homeownership likely reflects a combination of shorter-term cyclical factors and longer-term structural trends. The literature suggests that some of the decline in UK youth homeownership has been temporary, and historically, young cohorts with low homeownership rates often catch up when they age into middle-age (Bottazzi, Crossley and Wakefield, 2012). However, the graph of homeownership rates by age for five-year-birth-cohorts indicates that the sharp decline in the homeownership rate for several of the younger cohorts does not appear to be significantly reversing by the age of 40.


### C.0.1.3 Summary statistics table

Here is a table summarising my full dataset by age and tenure status.

Table C.1: Summary Statistics

		Age 18-34 household head			Full sample		
		rent	mortgagor	own	rent	mortgagor	own
Age hhd head							
	18-24	4071	1532	94	4071	1532	94
	25-29	5390	6448	255	5390	6448	255
	30-34	5,276	10,148	483	5,276	10,148	483
	35-49	.	.	.	12376	33995	4411
	50-64	.	.	.	9485	17541	16910
	65+	.	.	.	13034	2592	30180
Marital status							
	Married	3,743	10,287	335	16,863	49,382	30,682
	Cohabitate	3,113	3,598	161	5,058	7,347	1,394
	Single/divorced	7,863	4,221	336	27,683	15,512	20,040
Age gap with partner: spouse is:							
	yrs younger	111	109	2	2086	3444	2250
	5-9yrs younger	1020	1522	58	4420	9741	6027
	1-4yrs younger same age	2957	6255	195	8,532	24,902	14,246
	1-4yrs older	962	2,127	80	2,297	6,757	4,107
	5-9yrs older	1272	2,908	108	3,201	9,021	4,316
	10+ yrs older	411	775	32	995	2224	984
		121	191	20	363	585	283
Sex hhd head							
	Male	15,989	653	25,881	30,774	63,079	38,592
	Female	2,139	179	7,816	18,892	9,178	13,741
Number adults hhd							
	1	6,440	3,722	278	22,727	12,614	17,106
	2	7,344	13,977	500	21,897	48,832	29,570
	3	565	332	36	3,569	7,962	4,364
	4+	388	97	18	1,473	2,849	1,293
Number kids 0-4							
	0	8,843	11,279	548	58,230	41,176	51,353
	1	4,336	5,020	201	10,601	6,385	771
	2+	1,558	1,829	83	3,426	2,105	209
Number kids 5-10							
	0	10,517	14,657	646	41,135	56,207	50,756
	1	2,865	2,399	125	5,909	10,989	1,142
	2+	1,355	1,072	61	2,622	5,061	435
Number kids 10-16							
	0	13,388	17,259	764	42,860	56,978	49780
	1	1,017	708	54	4,818	10,587	1,879
	2+	332	161	14	1,988	4,692	674
Region							
	North East	816	971	35	3,159	3,588	2,409
	North West	1,598	2,195	95	5,339	8,302	5,969
	Yorkshire/Humber	1,445	1,727	81	4,724	6,583	4,703
	East Midlands	1,059	1,491	68	3,419	5,789	4,229
	West Midlands	1,333	1,583	79	4,532	6,498	4,933
	Eastern	862	1,130	41	3,045	4,937	4,144
	London	2,104	1,803	117	6,423	6,689	3,947
	South East	2,079	3,314	141	6,891	13,615	9,158
	South West	1,196	1,580	60	3,934	6,784	5,709
	Wales	710	825	45	2,389	3,451	3,136
	Scotland	1,535	1,509	70	5,811	6,021	3,996
<i>N</i>		14,747	18,128	832	49,666	72,257	52,333

Figure C.2



### C.0.1.4 Budget shares

Figure C.3

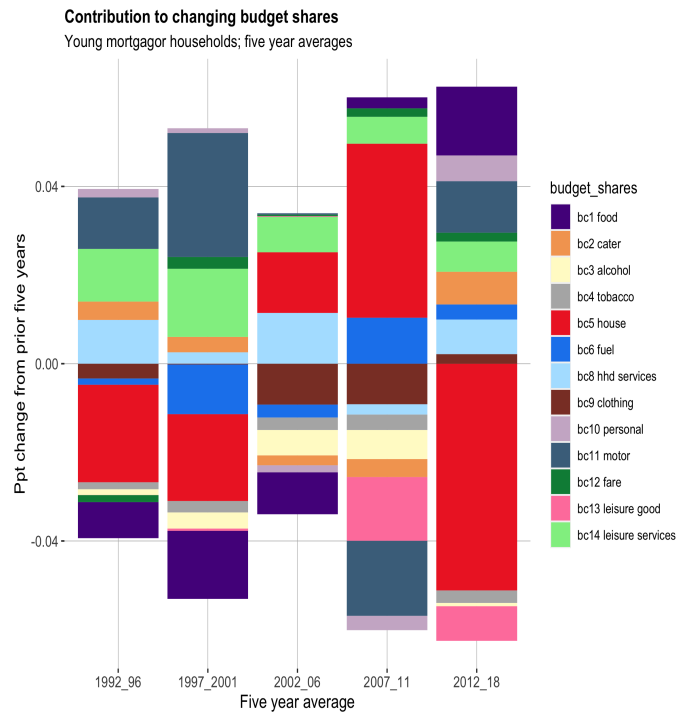



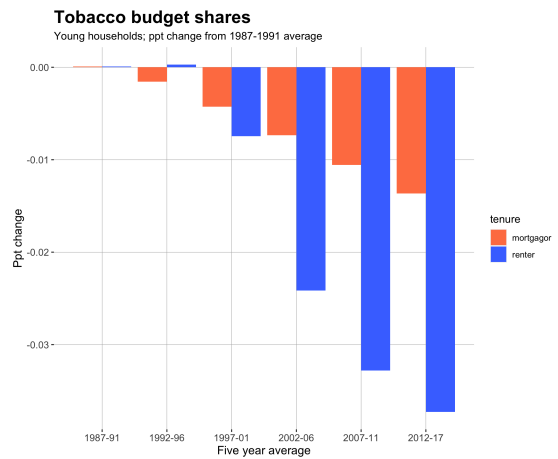
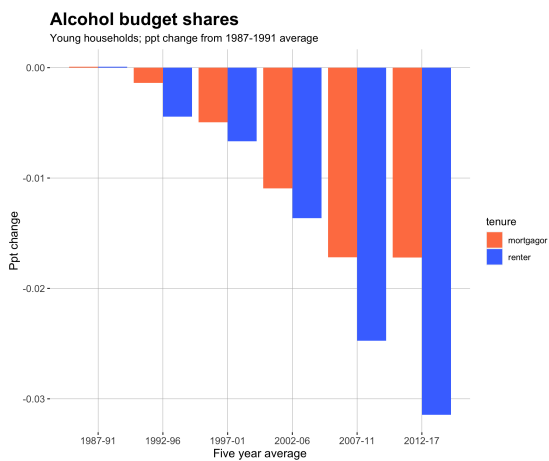

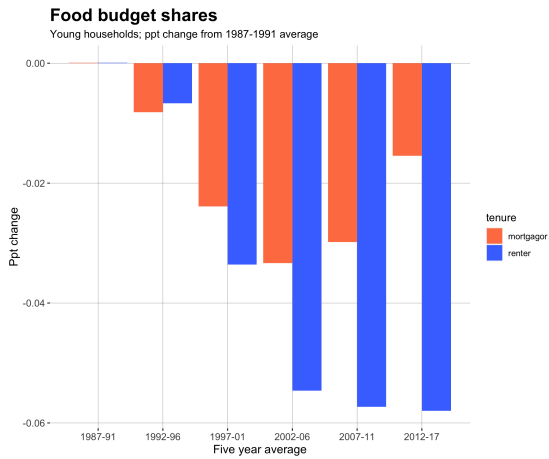

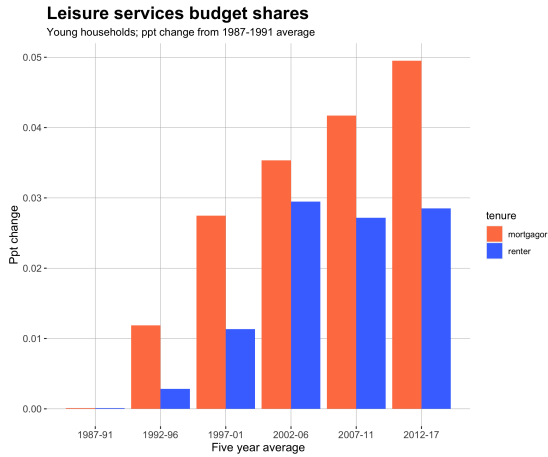






Figure C.4



I also report changes to each budget share for young owners and renters relative to a 1987–91 average.





## C.0.2 Additional demand estimation derivations & results

### C.0.2.1 AIDS derivation

The derivation of AIDS demand functions is a fairly simple application of consumer theory. Starting with the AIDS expenditure function:

$$\ln e(p_t, U_t) = \alpha_0 + \sum_i \alpha_i \ln p_{it} + \frac{1}{2} \sum_i \sum_j \gamma_{ij} \ln p_{it} \ln p_{jt} + U_t \beta_0 \prod_i p_{it}^{\beta_i}$$

where  $p_{it}$  is the price of good  $i$  at time  $t$ ,  $U_t$  is utility at time  $t$ , and  $\alpha$ ,  $\beta$  and  $\gamma$  are coefficients to estimate. By applying Shephard's Lemma and replacing the unobservable utility term with an indirect utility function that corresponds to the expenditure function, we can then derive Marshallian demand function in terms of  $x_{it}$  for the quantity consumed of good  $i$  at time  $t$  in terms of total expenditure  $m$  and prices:

$$x_{it}(p_t, m_t) = \frac{m_t}{p_t} \left( \alpha_i + \sum_j \left( \frac{1}{2}(\gamma_{ij} + \gamma_{ji}) \right) \ln p_{jt} + \beta_i \ln \left( \frac{m_t}{P_t} \right) \right)$$

$P_t$  is a translog price index:  $\ln P_t = \alpha_0 + \sum_i \alpha_i \ln p_{it} + \frac{1}{2} \sum_i \sum_j \left( \frac{1}{2}(\gamma_{ij} + \gamma_{ji}) \right) \ln p_{it} \ln p_{jt}$ , and  $\gamma_{ij}$  is  $\frac{1}{2}(\gamma_{ij} + \gamma_{ji})$  if symmetry is imposed. The final step is to simplify the Marshallian demand equations by expressing them as expenditure (budget) shares  $w_i = \frac{x_{it} p_{it}}{m_t}$ :

$$w_{it} = \alpha_i + \sum_j \left( \frac{1}{2}(\gamma_{ij} + \gamma_{ji}) \right) \ln p_{jt} + \beta_i \ln \left( \frac{m_t}{P_t} \right)$$

To incorporate demographic variables based on Ray (1983),  $m_{0h}$  is defined as the ratio of the expenditure required to obtain a certain utility level at given price  $\mathbf{p}$  for household  $h$  with  $\mathbf{z}$  demographic features, relative to the expenditure of a reference household  $R$ :

$$e(\mathbf{p}, \mathbf{z}, u) = m_{0h}(\mathbf{p}, \mathbf{z}, u) e^R(\mathbf{p}, u)$$

The key assumption of this approach is the impact of demographic variables on budget shares is independent of total expenditure. Therefore,  $m_{0h}$  can be decomposed into two factors:  $\bar{m}_0(z)$  which captures the change in household expenditure as a function

of demographic variables only, and  $\phi(\mathbf{p}, \mathbf{z}, u)$ , which captures changes in relative prices and actual goods consumed, so that:

$$e(\mathbf{p}, \mathbf{z}, u) = \bar{m}_0(\mathbf{z})\phi(\mathbf{p}, \mathbf{z}, u)e^R(\mathbf{p}, u)$$

Continuing to follow Ray (1983), we can parameterise the components of  $m_{0h}(\mathbf{p}, \mathbf{z}, u)$  as follows (where  $\rho$  is a vector of parameters to be estimated):

$$\bar{m}_0(\mathbf{z}) = 1 + \rho\mathbf{z}$$

$$\ln \phi(\mathbf{p}, \mathbf{z}, u) = \frac{\prod_{j=1}^k p_j^{\beta_j} (\prod_{j=1}^k p_j^{\eta'_j \mathbf{z}} - 1)}{\frac{1}{u} - \sum_{j=1}^k \lambda_j \ln p_j}$$

The expenditure share equations that take into account demographics are:

$$w_i = \alpha_i + \sum_{j=1}^k \gamma_{ij} \ln p_j + (\beta_i + \eta_i \mathbf{z}) \ln \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} + \frac{\lambda_i}{b(\mathbf{p})c(\mathbf{p}, \mathbf{z})} \left[ \ln \left\{ \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} \right\} \right]^2$$

where  $c(\mathbf{p}, \mathbf{z}) = \prod_{j=1}^k p_j^{\eta'_j \mathbf{z}}$ . The adding up condition now requires that  $\sum_{j=1}^k \eta_{rj} = 0$  for  $r=1\dots s$ . See Ray (1983) for further details of this derivation.

### C.0.2.2 Elasticity formula derivation

The equations for income, compensated price and uncompensated price elasticity as derived from QUAIDS models are listed below. They can be obtained by differentiating the budget share equation with respect to total income and prices. I omit the algebraic details of the derivation, but the elasticity formulas are listed below. The notation is from Poi (2012).

The income elasticity for good  $i$   $\mu_i$  is:

$$\mu_i = 1 + \frac{1}{w_i} \left[ \beta_i + \eta'_i \mathbf{z} + \frac{2\lambda_i}{b(\mathbf{p})c(\mathbf{p}, \mathbf{z})} \ln \left\{ \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} \right\} \right]$$

Where  $b(\mathbf{p})$  is the Cobb-Douglas price aggregator  $\prod_i p_{it}^{\beta_i}$ ,  $\ln a(\mathbf{p})$  is the transcendental logarithm function:  $\alpha_0 + \sum_i \alpha_i \ln p_{it} + \frac{1}{2} \sum_i \sum_j \gamma_{ij} \ln p_{it} \ln p_{jt}$ ,  $c(\mathbf{p}, \mathbf{z})$  is the Cobb-Douglas price aggregator scaled by the demographic variables  $\mathbf{z}$ :  $\prod_{j=1}^k p_j^{\eta'_j \mathbf{z}}$ ,  $w_i$  is

the budget share of good  $i$ ,  $m$  is total expenditure and  $\bar{m}_0(\mathbf{z}) = 1 + \rho' \mathbf{z}$  and is a component of the function that scales expenditure function to account for the household characteristics.

The uncompensated price elasticity for good  $i$  with respect to a price change of good  $j$ ,  $\epsilon_{ij}^U$ , is:

$$\epsilon_{ij}^U = -\delta_{ij} + \frac{1}{w_i} \left( \gamma_{ij} - \left[ \beta_i + \eta'_i \mathbf{z} + \frac{2\lambda_i}{b(\mathbf{p})c(\mathbf{p}, \mathbf{z})} \ln \left\{ \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} \right\} \right] \times \left( \alpha_j + \sum_i \gamma_{ji} \ln p_i \right) \right) - \frac{1}{w_i} \left( \frac{(\beta_j + \eta'_j \mathbf{z}) \gamma_i}{b(\mathbf{p})c(\mathbf{p}, \mathbf{z})} \left[ \ln \left\{ \frac{m}{\bar{m}_0(\mathbf{z})a(\mathbf{p})} \right\} \right]^2 \right)$$

$\delta_{ij}$  is the Kronecker delta, which takes the value of 1 if  $i = j$  and zero otherwise. The compensated price elasticity for good  $i$  with respect to a price change of good  $j$ ,  $\epsilon_{ij}^C$ , can be derived from the Slutsky equation:

$$\epsilon_{ij}^C = \epsilon_{ij}^U + \mu_i w_j$$

### C.0.2.3 Additional QUAIDS estimates

Table C.2: QUAIDS coefficient estimates - Owner occupiers

<i>Good<sub>i</sub></i>	$\alpha_i$	$\beta_i$	$\gamma_{1i}$ (food)	$\gamma_{2i}$ (house)	$\gamma_{3i}$ (fuel)	$\gamma_{4i}$ (leisure)	$\gamma_{5i}$ (goods)	$\gamma_{6i}$ (services)	$\lambda$
Food	0.2379* (0.0080)	-0.0241* (0.0022)	0.071* (0.004)	0.101* (0.003)	0.026* (0.001)	-0.146* (0.005)	0.021* (0.003)	-0.073* (0.006)	0.022* (0.001)
House	0.2730* (0.0140)	-0.0501* (0.0035)	.	0.071* (0.006)	-0.029* (0.001)	-0.165* (0.006)	0.030* (0.004)	-0.008* (0.007)	-0.006 (0.001)
Fuel	-0.0412* (0.0032)	-0.0026 (0.0016)	.	.	0.035* (0.001)	-0.006* (0.002)	-0.017* (0.001)	-0.010* (0.002)	0.014* (0.000)
Leisure	-0.0026 (0.0141)	0.0865* (0.0043)	.	.	.	0.294* (0.010)	-0.149* (0.005)	0.173* (0.010)	-0.001 (0.001)
Other goods	0.1340* (0.0098)	-0.1392* (0.0040)	.	.	.	.	0.086* (0.005)	0.029* (0.006)	-0.042* (0.001)
Other services	0.3990* (0.0179)	0.1294* (0.0045)	.	.	.	.	.	-0.111* (0.016)	0.012* (0.001)

\*=p<0.01, suppress coefficients on scaling factors and interaction terms

### C.0.2.4 EASI estimates

Table C.3: EASI coefficient estimates - renters

	Food		Housing		Fuel		Leisure		Other Goods	
	coeff	st error	coeff	st error	coeff	st error	coeff	st error	coeff	st error
$\alpha$	-0.094	0.053	0.305*	0.091	0.235*	0.022	0.604*	0.076	-0.043	0.059
$y^1$	-0.160*	0.022	0.091*	0.037	0.055*	0.009	0.098*	0.032	-0.032	0.024
$y^2$	0.011*	0.001	-0.042*	0.002	0.006*	0.001	0.025*	0.002	-0.004*	0.002
$y^3$	0.004*	0.000	-0.003*	0.000	0.001*	0.000	0.001*	0.000	-0.001*	0.000
$z1$	0.008	0.042	0.053	0.071	-0.076*	0.017	-0.114	0.060	0.037	0.046
$z2$	0.029*	0.001	-0.020*	0.001	0.005*	0.000	-0.016*	0.001	0.011*	0.001
$z3$	0.020*	0.000	-0.026*	0.001	0.003*	0.000	-0.001	0.001	0.002*	0.001
$z4$	-0.002*	0.000	0.009*	0.000	-0.001*	0.000	-0.005*	0.000	-0.006*	0.000
$z5$	0.002*	0.000	0.002*	0.000	0.000	0.000	-0.002*	0.000	-0.001*	0.000
$y^*z1$	-0.004*	0.002	-0.009*	0.003	0.005*	0.001	-0.002	0.002	0.000	0.002
$y^*z2$	-0.012*	0.001	0.014*	0.001	0.000	0.000	-0.004*	0.001	0.007*	0.001
$y^*z3$	-0.009*	0.000	0.012*	0.001	-0.003*	0.000	-0.002*	0.001	0.002*	0.001
$y^*z4$	0.003*	0.000	0.004*	0.000	0.000	0.000	-0.003*	0.000	-0.005*	0.000
$y^*z5$	0.000	0.000	0.002*	0.000	0.000*	0.000	-0.001*	0.000	-0.001*	0.000
$ps1$	0.003	0.029	-0.060*	0.023	0.031*	0.007	0.002	0.028	0.081*	0.021
$ps2$	-0.060*	0.023	0.062	0.040	0.095*	0.010	0.128*	0.034	-0.165*	0.026
$ps3$	0.031*	0.007	0.095*	0.010	0.036*	0.004	-0.048*	0.011	-0.103*	0.008
$ps4$	0.002	0.028	0.128*	0.034	-0.048*	0.011	-0.012	0.047	-0.129*	0.027
$ps5$	0.081*	0.021	-0.165*	0.026	-0.103*	0.008	-0.129*	0.027	0.275*	0.029
$y^*ps1$	-0.045*	0.012	-0.024*	0.009	0.009*	0.003	-0.026*	0.011	0.027*	0.008
$y^*ps2$	-0.024*	0.009	0.089*	0.016	0.043*	0.004	-0.011	0.014	-0.042*	0.010
$y^*ps3$	0.009*	0.003	0.043*	0.004	-0.015*	0.002	0.012*	0.004	-0.039*	0.003
$y^*ps4$	-0.026*	0.011	-0.011	0.014	0.012*	0.004	0.014	0.019	0.005	0.011
$y^*ps5$	0.027*	0.008	-0.042*	0.010	-0.039*	0.003	0.005	0.011	0.014	0.011
$z1^*ps1$	0.007	0.023	-0.005	0.018	-0.006	0.006	-0.047*	0.022	-0.027	0.017
$z1^*ps2$	-0.005	0.018	0.015	0.032	-0.038*	0.008	-0.032	0.027	0.050*	0.020
$z1^*ps3$	-0.006	0.006	-0.038*	0.008	-0.005	0.003	0.004	0.008	0.058*	0.007
$z1^*ps4$	-0.047*	0.022	-0.032	0.027	0.004	0.008	0.045	0.037	0.027	0.021
$z1^*ps5$	-0.027	0.017	0.050*	0.020	0.058*	0.007	0.027	0.021	-0.122*	0.023

No variable 6 coefficients, y is expenditure, z is demographic term, s are budget shares, ps is price of budget share \* =  $p < 0.05$ ,  $\alpha$  is constant

Table C.4: EASI elasticity estimates - Hicksian price & income

	Food	Housing	Fuel	Leisure	Other Goods	Other Services
food	-0.197* (0.072)	-0.432* (0.051)	0.044 (0.045)	-0.132 (0.085)	0.259* (0.062)	-0.505* (0.148)
housing	-0.384* (0.058)	-0.844* (0.084)	0.505* (0.056)	0.796* (0.102)	-0.663* (0.072)	-0.114 (0.179)
fuel	0.068* (0.019)	0.168* (0.021)	-0.501* (0.025)	-0.379* (0.032)	-0.085* (0.025)	-0.099 (0.053)
leisure	-0.037 (0.070)	0.557* (0.074)	-0.772* (0.062)	-1.154* (0.144)	-0.639* (0.076)	0.455* (0.203)
other goods	0.319* (0.053)	-0.495* (0.054)	-0.125* (0.050)	-0.634* (0.079)	-0.054 (0.084)	0.014 (0.148)
other services	-0.270 (0.148)	-0.020 (0.179)	-0.001 (0.053)	0.595* (0.203)	0.098 (0.148)	-0.738 (0.855)
income	0.486	1.206	0.437	1.288	1.074	1.243

\* =  $p < 0.05$ . Do not report income elasticity standard errors as unavailable with R package used for elasticity estimation

### C.0.2.5 Endogeneity

Table C.5: Mean Income & Compensated Price Elasticity Estimates - Baseline vs Instrumental Variables - Young Households

	Food		Housing		Fuel		Leisure		Other Goods		Other Services	
	BL	IV	BL	IV	BL	IV	BL	IV	BL	IV	BL	IV
Income	0.393*	0.413*	0.985*	1.140*	0.244*	0.259*	1.343*	1.296*	0.961*	0.815*	1.379*	1.350*
Food price	-0.910*	-0.834*	0.315*	0.339*	0.531*	0.553*	-0.019	-0.061	0.531*	0.457*	-0.303*	-0.289*
Housing price	0.523*	0.555*	-0.613*	-0.837*	1.016*	1.060*	0.09	0.116*	-0.025	0.148*	-0.019	0.023
Fuel price	0.228*	0.237*	0.262*	0.277*	-0.475*	-0.432*	-0.128*	-0.138*	-0.198*	-0.214*	-0.039	-0.048
Leisure price	-0.025	-0.079	0.070	0.091	-0.387*	-0.415*	-0.088	-0.099	-1.269*	-1.181*	1.152*	1.120*
Other goods price	0.631*	0.548*	-0.018	0.108	-0.551*	-0.600*	-1.171*	-1.098*	-0.342*	-0.460*	1.052*	1.018*
Other services price	-0.446*	-0.427*	-0.017	0.021	-0.136	-0.166	1.317*	1.279*	1.303*	1.251*	-1.842*	-1.824*

'BL' = 'baseline model', 'IV' = 'model with instrumental variables'; \*=significant at p=0.05; standard errors omitted. The baseline estimates reported in this table are different to those reported with my baseline model in Chapter 6.3 of this thesis. This is for several reasons. First, I am not able to match all households with disposable household income data, and so these results are estimated with a reduced sample. Second, I use all households (homeowners and renters) for the estimation to overcome small sample size issues caused by only being able to use a restricted sample with the disposable household income data

### C.0.3 Deriving Compensating Variation

This derivation is based on Friedman and Levinsohn (2001). We start with the expenditure function  $C(p, u)$ . Taking a second order Taylor expansion of the minimum expenditure function with respect to price:

$$\Delta C \approx q\Delta p + \frac{1}{2}\Delta p^T s \Delta p$$

Where  $p$  is a vector of price changes for  $n$  goods,  $q$  is a vector of consumption good quantities, and  $s$  is the  $n \times n$  matrix of compensated derivatives of demand. This expression can be re-written in terms of budget shares and logarithms:

$$\Delta \ln C \approx \sum_{i=1}^n w_i \Delta \ln p_i + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n c_{ij} \Delta \ln p_i \Delta \ln p_j$$

where  $c_{ij} = P_i s_{ij} p_j / C$ , and  $s_{ij}$  is the Slutsky derivatives. With some algebraic manipulation, it can be shown that  $c_{ij} = w_i \varepsilon_{ij}$ , resulting in the expression:

$$CV \approx \sum_{i=1}^n w_i \Delta \ln p_i + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n w_i \varepsilon_{ij} \Delta \ln p_i \Delta \ln p_j$$

where  $w_i$  is the budget share of commodity  $i$  (out of  $n$  different commodities) in the initial period before the price change,  $\Delta \ln p_i$  is the change in price of good  $i$  and  $\varepsilon_{ij}$  is the compensated price elasticity of commodity  $i$  with respect to the price change of good  $j$ ,

## BIBLIOGRAPHY

- Abowd, John, and David Card.** 1989. “On the Covariance Structure of Earnings and Hours Changes.” *Econometrica*, 57(2): 411–45.
- Abramson, Boaz, Job Boerma, and Aleh Tsyvinski.** 2024. “Macroeconomics of Mental Health.” National Bureau of Economic Research Working Paper 32354.
- Adams-Prassl, Abi, Maria Balgova, Matthias Qian, and Tom Waters.** 2023. “Firm concentration & job design: the case of schedule flexible work arrangements.” Institute for Fiscal Studies IFS Working Papers WP2023/14.
- Aghion, Philippe, Antonin Bergeaud, Richard Blundell, and Rachel Griffin.** 2023. “Social Skills and the Individual Wage Growth of Less Educated Workers.” Institute for Fiscal Studies Working paper 23/25.
- Ajnakin, Olesya, and University College London Andrew Steptoe.** 2022. “The English Longitudinal Study of Ageing (ELSA) Polygenic Scores.” University College London Documentation Report.
- Albouy, David, Gabriel Ehrlich, and Yingyi Liu.** 2016. “Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis.” National Bureau of Economic Research Working Paper 22816.
- Almuzara, Martin.** 2020. “Heterogeneity in transitory income risk.” Unpublished mimeo; available at [https://martinalmuzara.com/files/transitory\\_risk/transitory\\_risk.pdf](https://martinalmuzara.com/files/transitory_risk/transitory_risk.pdf).
- Amano-Patiño, Noriko.** 2019. “Nutritional Inequality: The Role of Prices, Income, and Preferences.” *Cambridge Working Papers in Economics*, 1909.
- Amengual, Dante, Jesús Bueren, and Julio A. Crego.** 2021. “Endogenous health groups and heterogeneous dynamics of the elderly.” *Journal of Applied Econometrics*, 36(7): 878–897.
- Anderson, Kathryn H., and Richard V. Burkhauser.** 1984. “The importance of the measure of health in empirical estimates of the labor supply of older men.” *Economics Letters*, 16(3): 375–380.

- Arellano, Manuel.** 1995. “Panel data models with predetermined instruments, Paper presented at Review of Economic Studies lecture, Royal Economic Society Conference, Kent.”
- Arellano, Manuel, and Stephen Bond.** 1991. “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations.” *Review of Economic Studies*, 58(2): 277–297.
- Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme.** 2017. “Earnings and Consumption Dynamics: A Non-linear Panel Data Framework.” *Econometrica*, 85(3): 693–734.
- Attanasio, Orazio, and Valérie Lechene.** 2014. “Efficient Responses to Targeted Cash Transfers.” *Journal of Political Economy*, 122(1): 178 – 222.
- Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos.** 2018. “Aggregating Elasticities: Intensive and extensive margins of women’s labor supply.” *Econometrica*, 86(6): 2049–2082.
- Attanasio, Orazio P., Laura Blow, Robert Hamilton, and Andrew Leices-ter.** 2009. “Booms and Busts: Consumption, House Prices and Expectations.” *Economica*, 76(301): 20–50.
- Attanasio, Orazio P., Renata Bottazzi, Hamish W. Low, Lars Nesheim, and Matthew Wakefield.** 2012. “Modelling the demand for housing over the life cycle.” *Review of Economic Dynamics*, 15(1): 1 – 18.
- Attanasio, Orazio, Sagiri Kitao, and Giovanni L. Violante.** 2010. “Financing Medicare: A General Equilibrium Analysis.” *Demography and the Economy*, 333–366. University of Chicago Press.
- Autor, David H.** 2015. “Why Are There Still So Many Jobs? The History and Future of Workplace Automation.” *Journal of Economic Perspectives*, 29(3): 3–30.
- Bachmann, Ronald, Gökay Demir, Colin Green, and Arne Uhlendorff.** 2022. “The Role of Within-Occupation Task Changes in Wage Development.” Center for Research in Economics and Statistics Working Papers 2022-20.
- Baltagi, Badi H., Sergi Jiménez-Martín, José M. Labeaga, and Majid al Sadoon.** 2023. “Consistent estimation of panel data sample selection models.” *Econometrics and Statistics*.
- Banks, James, Richard Blundell, and Arthur Lewbel.** 1997. “Quadratic Engel Curves And Consumer Demand.” *The Review of Economics and Statistics*, 79(4): 527–539.
- Banks, James, Richard Blundell, Zoë Oldfield, and James P. Smith.** 2016. “House Price Volatility and the Housing Ladder.” In *Insights in the Economics of Aging. NBER Chapters*, 87–119. National Bureau of Economic Research, Inc.

- Barth, Daniel, Nicholas W. Papageorge, and Kevin Thom.** 2020. “Genetic Endowments and Wealth Inequality.” *Journal of Political Economy*, 128(4): 1474–1522.
- Baum-Snow, Nathaniel, and Lu Han.** 2024. “The Microgeography of Housing Supply.” *Journal of Political Economy*, 132(6): 1897–1946.
- Benzeval, Michaela, Apostolos Davillas, and Andrew M. Jones.** 2017. “The income-health gradient: evidence from self-reported health and biomarkers using longitudinal data on income.” Institute for Social and Economic Research ISER Working Paper Series 2017-03.
- Benzeval, Michaela, Apostolos Davillas, Meena Kumari, and Peter Lynn.** 2014. “Understanding Society: The UK Household Longitudinal Study.” Institute for Social and Economic Research Biomarker User Guide and Glossary.
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra.** 2017. “House Prices and Consumer Spending.” *The Review of Economic Studies*, 85(3): 1502–1542.
- Bijlsma, Michiel, and Remco Mocking.** 2017. “The Impact of House Price Shocks on the Savings of Dutch Homeowners and Renters.” CPB Netherlands Bureau for Economic Policy Analysis CPB Discussion Paper 346.
- Biro, Aniko, Márta Bisztray, João G. da Fonseca, and Tímea Laura Molnár.** 2023. “Accident-Induced Absence from Work and Wage Ladders.” Institute of Labor Economics (IZA) IZA Discussion Papers 16312.
- Black, Nicole, David W. Johnston, and Agne Suziedelyte.** 2017. “Justification bias in self-reported disability: New evidence from panel data.” *Journal of Health Economics*, 54: 124–134.
- Blau, David M., and Donna B. Gilleskie.** 2001. “Retiree Health Insurance and the Labor Force Behavior of Older Men in the 1990s.” *The Review of Economics and Statistics*, 83(1): 64–80.
- Blundell, Richard, and Stephen Bond.** 1998. “Initial conditions and moment restrictions in dynamic panel data models.” *Journal of Econometrics*, 87(1): 115–143.
- Blundell, Richard, and Steve Bond.** 2023. “Initial conditions and Blundell–Bond estimators.” *Journal of Econometrics*, 234: 101–110. Jubilee Issue - Celebrating our Fiftieth Anniversary: 1973–2023.
- Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French.** 2016. “The Dynamic Effects of Health on the Employment of Older Workers.” University of Michigan Retirement Research Center (MRRC) Working Paper 348.

- Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French.** 2021. “The impact of health on labour supply near retirement.” *The Journal of Human Resources*, 1217.
- Blundell, Richard, Margherita Borella, Jeanne Commault, and Mariacristina De Nardi.** 2020*a*. “Why Does Consumption Fluctuate in Old Age and How Should the Government Insure It?” National Bureau of Economic Research Working Paper 27348.
- Blundell, Richard, Monica Costa Dias, Robert Joyce, and Agnes Norris Keiller.** 2020*b*. “What has been happening to career progression?” Institute for Fiscal Studies Briefing Note BN301.
- Born, Benjamin, and Jörg Breitung.** 2016. “Testing for Serial Correlation in Fixed-Effects Panel Data Models.” *Econometric Reviews*, 35(7): 1290–1316.
- Bottazzi, Renata, Thomas Crossley, and Matthew Wakefield.** 2012. “Late starters or excluded generations? A cohort analysis of catch up in home ownership in England.” Institute for Fiscal Studies Working Papers W12/10.
- Bound, John.** 1991. “Self-Reported Versus Objective Measures of Health in Retirement Models.” *The Journal of Human Resources*, 26(1): 106–138.
- Bound, John, and Richard V. Burkhauser.** 1999. “Economic analysis of transfer programs targeted on people with disabilities.” In *Handbook of Labor Economics*. Vol. 3 of *Handbook of Labor Economics*, , ed. O. Ashenfelter and D. Card, Chapter 51, 3417–3528. Elsevier.
- Bound, John, Todd Stinebrickner, and Timothy Waidmann.** 2010. “Health, economic resources and the work decisions of older men.” *Journal of Econometrics*, 156(1): 106–129.
- Braun, R. Anton, Karen A Kopecky, and Tatyana Koreshkova.** 2017. “Old, Sick, Alone, and Poor: A Welfare Analysis of Old-Age Social Insurance Programmes.” *The Review of Economic Studies*, 84(2 (299)): 580–612.
- Broome, Molly, Adam Corlett, Sophie Hale, Charlie McCurdy, and Cara Pacitti.** 2023. “An intergenerational audit for the UK: 2023.” Resolution Foundation report.
- Brown, Alexandra Elizabeth.** 2023. “How should we model health as a dynamic process?” University of Oxford, Department of Economics Economics Series Working Papers 1023.
- Browning, Martin, Mette Gørtz, and Søren Leth-Petersen.** 2013. “Housing Wealth and Consumption: A Micro Panel Study.” *The Economic Journal*, 123(568): 401–428.

- Busch, Christopher.** 2020. “Occupational Switching, Tasks, and Wage Dynamics.” Unpublished working paper; available at [https://chrisbusch.eu/wp-content/uploads/2020/10/Busch\\_occ\\_2020.pdf](https://chrisbusch.eu/wp-content/uploads/2020/10/Busch_occ_2020.pdf).
- Campbell, John Y., and Joao F. Cocco.** 2007. “How do house prices affect consumption? Evidence from micro data.” *Journal of Monetary Economics*, 54(3): 591–621.
- Capatina, Elena.** 2015. “Life-cycle effects of health risk.” *Journal of Monetary Economics*, 74: 67–88.
- Caraballo, César, Rohan Khera, Philip G. Jones, Carole Decker, Wade Schulz, John A. Spertus, and Harlan M. Krumholz.** 2020. “Rates and Predictors of Patient Underreporting of Hospitalizations During Follow-Up After Acute Myocardial Infarction.” *Circulation: Cardiovascular Quality and Outcomes*, 13(7): e006231.
- Carrieri, Vincenzo, Andrew M. Jones, and Francesco Principe.** 2018. “Health shocks and labour market outcomes: evidence from professional football.” Health Econometrics and Data group, University of York Working Paper.
- Carrillo-Tudela, Carlos, and Ludo Visschers.** 2023. “Unemployment and Endogenous Reallocation Over the Business Cycle.” *Econometrica*, 91(3): 1119–1153.
- Carrillo-Tudela, Carlos, Bart Hobijn, Powen She, and Ludo Visschers.** 2016. “The extent and cyclical nature of career changes: Evidence for the U.K.” *European Economic Review*, 84: 18–41. European Labor Market Issues.
- Case, Anne, Angela Fertig, and Christina Paxson.** 2004. “The Lasting Impact of Childhood Health and Circumstance.” Princeton University, Woodrow Wilson School of Public and International Affairs, Center for Health and Wellbeing. Working Papers 246.
- Chandola, Tarani, and Nan Zhang.** 2017. “Re-employment, job quality, health and allostatic load biomarkers: prospective evidence from the UK Household Longitudinal Study.” *International Journal of Epidemiology*, 47(1): 47–57.
- Charles, Kerwin Kofi.** 2003. “The Longitudinal Structure of Earnings Losses among Work-Limited Disabled Workers.” *Journal of Human Resources*, 38(3): 618–646.
- Chetty, Raj, and Adam Szeidl.** 2016. “Consumption Commitments and Habit Formation.” *Econometrica*, 84: 855–890.
- Conti, Gabriella, and James J. Heckman.** 2010. “Understanding the Early Origins of the Education–Health Gradient: A Framework That Can Also Be Applied to Analyze Gene–Environment Interactions.” *Perspectives on Psychological Science*, 5(5): 585–605. PMID: 21738556.

- Cribb, Jonathan.** 2024. “IFS Election 2024: Homeownership for young adults has recovered to its 2010 level.” <https://ifs.org.uk/articles/homeownership-young-adults-has-recovered-its-2010-level>, Accessed: 5 June 2024.
- Cribb, Jonathan, and Polly Simpson.** 2018. “Barriers to homeownership for young adults.” Institute for Fiscal Studies book chapter.
- Cronqvist, Henrik, and Stephan Siegel.** 2015. “The Origins of Savings Behavior.” *Journal of Political Economy*, 123(1): 123–169.
- Crossley, Thomas F., Paul Fisher, and Hamish Low.** 2021. “The heterogeneous and regressive consequences of COVID-19: Evidence from high quality panel data.” *Journal of Public Economics*, 193: 104334.
- Dal Bianco, Chiara, and Andrea Moro.** 2022. “The Welfare Effects of Nonlinear Health Dynamics.” Center for Open Science SocArXiv dvhby.
- Dano, Anne Moller.** 2005. “Road injuries and long-run effects on income and employment.” *Health Economics*, 14(9): 955–970.
- Davillas, Apostolos, and Andrew M Jones.** 2020. “Ex ante inequality of opportunity in health, decomposition and distributional analysis of biomarkers.” *Journal of Health Economics*, 69: 102251.
- Davillas, Apostolos, Andrew M. Jones, and Michaela Benzeval.** 2019. “Chapter 22 - The Income-Health Gradient: Evidence From Self-Reported Health and Biomarkers in Understanding Society.” In *Panel Data Econometrics*, ed. Mike Tsionas, 709–741. Academic Press.
- Davillas, Apostolos, and Stephen Pudney.** 2020a. “Biomarkers as precursors of disability.” *Economics and Human Biology*, 36: 100814.
- Davillas, Apostolos, and Stephen Pudney.** 2020b. “Biomarkers, disability and health care demand.” *Economics and Human Biology*, 39: 100929.
- Davillas, Apostolos, and Stephen Pudney.** 2020c. “Using biomarkers to predict healthcare costs: Evidence from a UK household panel.” *Journal of Health Economics*, 73: 102356.
- Deaton, Angus, and John Muellbauer.** 1980a. “An Almost Ideal Demand System.” *The American Economic Review*, 70(3): 312–326.
- Deaton, Angus, and John Muellbauer.** 1980b. *Economics and Consumer Behavior*. Cambridge University Press.
- Deming, David J.** 2017. “The Growing Importance of Social Skills in the Labor Market.” *The Quarterly Journal of Economics*, 132(4): 1593–1640.

- De Nardi, Mariacristina.** 2024. “Health Inequality and Health Types.” *Sargan Lecture, presented at the Royal Economic Society Annual Conference, Belfast, United Kingdom.*
- De Nardi, Mariacristina, Eric French, and John B. Jones.** 2010. “Why Do the Elderly Save? The Role of Medical Expenses.” *Journal of Political Economy*, 118(1): 39–75.
- De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo.** 2019. “Non-linear Household Earnings Dynamics, Self-Insurance, and Welfare.” *Journal of the European Economic Association*, 18(2): 890–926.
- Department for Levelling Up, Housing & Communities.** 2023. “English Housing Survey.” Department for Levelling Up, Housing & Communities report.
- De Vera, Micole, Javier Garcia-Brazales, and Jiayi Lin.** 2023. “Early Labor Market Origins of Long-Term Mental Health and its Intergenerational Correlation.” Unpublished mimeo; [https://micoledevera.github.io/papers/DeVera\\_GBrazales\\_Lin-Manuscript-MentalHealth.pdf](https://micoledevera.github.io/papers/DeVera_GBrazales_Lin-Manuscript-MentalHealth.pdf).
- Diebolt, Jean, and Gilles Celeux.** 1993. “Asymptotic properties of a stochastic EM Algorithm for estimating mixing proportions.” *Communications in Statistics. Stochastic Models*, 9(4): 599–613.
- Disney, Richard, and John Gathergood.** 2018. “House Prices, Wealth Effects and Labour Supply.” *Economica*, 85(339): 449–478.
- Disney, Richard, John Gathergood, and Andrew Henley.** 2010. “House Price Shocks, Negative Equity, And Household Consumption In The United Kingdom.” *Journal of the European Economic Association*, 8(6): 1179–1207.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo.** 2018. “The Economic Consequences of Hospital Admissions.” *American Economic Review*, 108(2): 308–52.
- Emmerson, Carl, Paul Johnson, and Nick Ridpath.** 2024. “The Conservatives and the Economy, 2010–24.” Institute for Fiscal Studies Report 24/24.
- Engelhardt, Gary V.** 1994. “House Prices and the Decision to Save for Down Payments.” *Journal of Urban Economics*, 36(2): 209 – 237.
- Ermisch, J.F., J. Findlay, and K. Gibb.** 1996. “The Price Elasticity of Housing Demand in Britain: Issues of Sample Selection.” *Journal of Housing Economics*, 5(1): 64–86.
- Florisson, Rebecca, Paula Holland, Heather Taylor, and Melanie Wilkes.** 2022. “The changing workplace: Enabling disability-inclusive hybrid working.” The Work Foundation - Lancaster University Report.

- French, Eric.** 2005. “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour.” *The Review of Economic Studies*, 72(2): 395–427.
- French, Eric, and John Bailey Jones.** 2011. “The Effects of Health Insurance and Self-Insurance on Retirement Behavior.” *Econometrica*, 79(3): 693–732.
- French, Eric, and John Bailey Jones.** 2017. “Health, Health Insurance, and Retirement: A Survey.” *Annual Review of Economics*, 9(1): 383–409.
- Friedman, Jed, and James Levinsohn.** 2001. “The Distributional Impacts of Indonesia’s Financial Crisis on Household Welfare: A “Rapid Response” Methodology.” National Bureau of Economic Research Working Paper P8564.
- Gannon, Brenda, and Jennifer Roberts.** 2011. “Part-time work and health among older workers in Ireland and Britain.” *Applied Economics*, 43(30): 4749–4757.
- Garcia-Gomez, Pilar, Andrew M Jones, and Nigel Rice.** 2008. “Health effects on labour market exits and entries.” Univeristy of York Health Econometrics and Data Group Working Paper.
- García-Gómez, Pilar, Hans van Kippersluis, Owen O’Donnell, and Eddy van Doorslaer.** 2013. “Long-Term and Spillover Effects of Health Shocks on Employment and Income.” *Journal of Human Resources*, 48(4): 873–909.
- Groes, Fane, Philipp Kircher, and Iourii Manovskii.** 2014. “The U-Shapes of Occupational Mobility.” *The Review of Economic Studies*, 82(2): 659–692.
- Grossman, Michael.** 1972. “On the Concept of Health Capital and the Demand for Health.” *Journal of Political Economy*, 80(2): 223–255.
- Group, Biomarkers Definitions Working.** 2001. “Biomarkers and surrogate end-points: Preferred definitions and conceptual framework.” *Clinical Pharmacology & Therapeutics*, 69(3): 89–95.
- Guiso, Luigi, Monica Paiella, and Ignazio Visco.** 2005. “Do capital gains affect consumption? Estimates of wealth effects from Italian households’ behavior.” Bank of Italy, Economic Research and International Relations Area Temi di discussione (Economic working papers) 555.
- Guvenen, Fatih, Burhan Kuruscu, Satoshi Tanaka, and David Wiczer.** 2020. “Multidimensional Skill Mismatch.” *American Economic Journal: Macroeconomics*, 12(1): 210–44.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song.** 2021. “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica*, 89(5): 2303–2339.

- Halla, Martin, and Martina Zweimüller.** 2013. “The effect of health on earnings: Quasi-experimental evidence from commuting accidents.” *Labour Economics*, 24(C): 23–38.
- Halliday, Timothy.** 2008. “Heterogeneity, state dependence and health.” *Econometrics Journal*, 11: 499–516.
- Hansen, Lars Peter.** 1982. “Large Sample Properties of Generalized Method of Moments Estimators.” *Econometrica*, 50(4): 1029–1054.
- Harmon, Oskar R.** 1988. “The income elasticity of demand for single-family owner-occupied housing: An empirical reconciliation.” *Journal of Urban Economics*, 24(2): 173–185.
- Harris, Mark N., Xueyan Zhao, and Eugenio Zucchelli.** 2021. “Ageing Workforces, Ill-health and Multi-state Labour Market Transitions\*.” *Oxford Bulletin of Economics and Statistics*, 83(1): 199–227.
- Harris, Matthew C.** 2019. “The impact of body weight on occupational mobility and career development.” *International Economic Review*, 60(2): 631–660.
- Harris, Sarah E, Saskia P Hagenars, Gail Davies, W David Hill, David CM Liewald, Stuart J Ritchie, Riccardo E Marioni, International Consortium for Blood Pressure Genome-Wide Association Studies METASTROKE Consortium, International Consortium for Blood Pressure Genome-Wide Association Studies, CHARGE Consortium Aging, Longevity Group, CHARGE Consortium Cognitive Group, Cathie LM Sudlow, Joanna M Wardlaw, Andrew M McIntosh, Catharine R Gale, and Ian J Deary.** 2016. “Molecular genetic contributions to self-rated health.” *International Journal of Epidemiology*, 46(3): 994–1009.
- Hauck, Katharina, and Nigel Rice.** 2004. “A longitudinal analysis of mental health mobility in Britain.” *Health Economics*, 13(10): 981–1001.
- Heckman, James J.** 1979. “Sample Selection Bias as a Specification Error.” *Econometrica*, 47(1): 153–161.
- Heien, Dale, and Cathy Roheim Wessells.** 1990. “Demand Systems Estimation with Microdata: A Censored Regression Approach.” *Journal of Business & Economic Statistics*, 8(3): 365–371.
- Heiss, Florian, Steven F Venti, and David A Wise.** 2014. “The Persistence and Heterogeneity of Health among Older Americans.” National Bureau of Economic Research Working Paper 20306.
- HM Treasury.** 2023. “Autumn Statement.” available at <https://www.gov.uk/government/publications/autumn-statement-2023>.

- Hosseini, Roozbeh, Karen A. Kopecky, and Kai Zhao.** 2022. “The evolution of health over the life cycle.” *Review of Economic Dynamics*, 45: 237–263.
- Huckfeldt, Christopher.** 2022. “Understanding the Scarring Effect of Recessions.” *American Economic Review*, 112(4): 1273–1310.
- Hudomiet, Péter, Michael D Hurd, Susann Rohwedder, and Robert J Willis.** 2018. “The Effect of Physical and Cognitive Decline at Older Ages on Job Mismatch and Retirement.” National Bureau of Economic Research Working Paper 25229.
- Imrohorglu, Ayse, and Kai Zhao.** 2018. “The chinese saving rate: Long-term care risks, family insurance, and demographics.” *Journal of Monetary Economics*, 96: 33–52.
- Jolivet, Grégory, and Fabien Postel-Vinay.** 2020. “A Structural Analysis of Mental Health and Labor Market Trajectories.” IZA Institute of Labor Economics Discussion Paper Series DP13518.
- Jolly, Nicholas A.** 2013. “The impact of work-limiting disabilities on earnings and income mobility.” *Applied Economics*, 45(36): 5104–5118.
- Jones, Andrew M., Nigel Rice, and Francesca Zantomio.** 2020. “Acute health shocks and labour market outcomes: Evidence from the post crash era.” *Economics & Human Biology*, 36: 100811.
- Jones, Andrew M., Xander Koolman, and Nigel Rice.** 2006. “Health-Related Non-Response in the British Household Panel Survey and European Community Household Panel: Using Inverse-Probability-Weighted Estimators in Non-Linear Models.” *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 169(3): 543–569.
- Joyce, Robert, Matthew Mitchell, and Agnes Norris Keiller.** 2017. “The cost of housing for low-income renters.” Institute for Fiscal Studies Report .
- Jung, Juergen, and Chung Tran.** 2016. “Market inefficiency, insurance mandate and welfare: U.S. health care reform 2010.” *Review of Economic Dynamics*, 20: 132–159.
- Kapteyn, Arie, James P. Smith, and Arthur van Soest.** 2011. “Work Disability, Work, and Justification Bias in Europe and the United States.” *Explorations in the Economics of Aging*, 269–312. University of Chicago Press.
- Keane, Michael, Elena Capatina, and Shiko Maruyama.** 2020. “Health Shocks and the Evolution of Earnings over the Life-Cycle.” School of Economics, The University of New South Wales Discussion Papers 2018-14b.

- Khorunzhina, Natalia.** 2020. “Intratemporal nonseparability between housing and nondurable consumption: evidence from reinvestment in housing stock.” *Journal of Monetary Economics*.
- Kumari, Meena, Tarek Al Baghal, and Michaela Benzeval.** 2022. “A framework for the re-collection of biomarkers in Understanding Society at Wave 16.” Understanding Society, University of Essex 2022-10.
- Lee, Seungyeon.** 2023. “House prices, homeownership, and household consumption: Evidence from household panel data in Korea.” *Economic Modelling*, 126: 106355.
- Lewbel, Arthur, and Krishna Pendakur.** 2009. “Tricks with Hicks: The EASI Demand System.” *American Economic Review*, 99(3): 827–63.
- Lillard, Lee A., and Robert J. Willis.** 1978. “Dynamic Aspects of Earning Mobility.” *Econometrica*, 46(5): 985–1012.
- Lise, Jeremy, and Fabien Postel-Vinay.** 2020. “Multidimensional Skills, Sorting, and Human Capital Accumulation.” *American Economic Review*, 110(8): 2328–76.
- Li, Wenli, Haiyong Liu, Fang Yang, and Rui Yao.** 2016. “Housing over time and over the life cycle: a structural estimation.” *International Economic Review*, 57(4): 1237–1260.
- Low, Hamish, and Luigi Pistaferri.** 2015. “Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off.” *American Economic Review*, 105(10): 2986–3029.
- Lucas, Robert E, and Edward C Prescott.** 1974. “Equilibrium search and unemployment.” *Journal of Economic Theory*, 7(2): 188–209.
- MaCurdy, Thomas E.** 1982. “The use of time series processes to model the error structure of earnings in a longitudinal data analysis.” *Journal of Econometrics*, 18(1): 83–114.
- Maes, Sebastiaan, and Raghav Malhotra.** 2024. “Robust Hicksian Welfare Analysis under Individual Heterogeneity.” arXiv.org Papers 2303.01231.
- McCall, Brian P.** 1990. “Occupational Matching: A Test of Sorts.” *Journal of Political Economy*, 98(1): 45–69.
- Meghir, Costas, and Frank Windmeijer.** 1999. “Moment Conditions for Dynamic Panel Data Models with Multiplicative Individual Effects in the Conditional Variance.” *Annales d’Économie et de Statistique*, 55/56: 317–330.
- Meghir, Costas, and Luigi Pistaferri.** 2004. “Income Variance Dynamics and Heterogeneity.” *Econometrica*, 72(1): 1–32.
- Melcangi, Davide, and Silvia Sarpietro.** 2024. “Nonlinear Firm Dynamics.” Federal Reserve Bank of New York Staff Reports 1088.

- Mian, Atif, Kamalesh Rao, and Amir Sufi.** 2013. “Household Balance Sheets, Consumption, and the Economic Slump.” *The Quarterly Journal of Economics*, 128(4): 1687–1726.
- Michaud, Amanda, and David Wiczer.** 2018. “Occupational hazards and social disability insurance.” *Journal of Monetary Economics*, 96: 77–92.
- Miles, David, and Victoria Monro.** 2019. “UK house prices and three decades of decline in the risk-free real interest rate.” Bank of England Staff Working Paper 837.
- Mortensen, Dale T., and Christopher A. Pissarides.** 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *The Review of Economic Studies*, 61(3): 397–415.
- Muellbauer, John.** 2018. “Housing, Debt and the Economy: a Tale of Two Countries.” University of Oxford Department of Economics Discussion Paper Series 855.
- Mullahy, John.** 2024. “Analyzing health outcomes measured as bounded counts.” *Journal of Health Economics*, 95: 102875.
- Niedzwiedz, Claire L, Srinivasa Vittal Katikireddi, Aaron Reeves, Martin McKee, and David Stuckler.** 2017. “Economic insecurity during the Great Recession and metabolic, inflammatory and liver function biomarkers: analysis of the UK Household Longitudinal Study.” *Journal of Epidemiology & Community Health*.
- Noray, Kadeem.** 2020. “The Labor Market Value of Non-Cognitive Skills.” PhD diss. Harvard University.
- Nygaard, Vegard M.** 2021. “Causes and consequences of life expectancy inequality.” *Manuscript, University of Houston*.
- Oldfield, Z., J. Banks, P. Levell, and A. Leicester.** 2019. “Family Expenditure Survey and Living Costs and Food Survey Derived Variables, 1968-2017.” [data collection].UK Data Service. SN: 8583.
- Ozkan, Serdar.** 2017. “Preventive vs. curative medicine: A macroeconomic analysis of health care over the life cycle.” *Manuscript, University of Toronto*.
- O’Donnell, Owen, Eddy Van Doorslaer, and Tom Van Ourti.** 2015. “Chapter 17 - Health and Inequality.” In *Handbook of Income Distribution*. Vol. 2 of *Handbook of Income Distribution*, , ed. Anthony B. Atkinson and François Bourguignon, 1419–1533. Elsevier.
- Palumbo, Michael G.** 1999. “Uncertain medical expenses and precautionary saving near the end of the life cycle.” *The Review of Economic Studies*, 66(227): 395–421.

- Pashchenko, Svetlana, Ponpoje (Poe) Porapakkarm, and Mariacristina De Nardi.** 2017. “The Lifetime Costs of Bad Health.” National Bureau of Economic Research Working Paper 23963.
- Picchio, Matteo, and Michele Ubaldi.** 2022. “Unemployment and Health: A Meta-Analysis.” Institute of Labor Economics (IZA) IZA Discussion Papers 15433.
- Pitt, Mark, and Lung-Fei Lee.** 1986. “Microeconomic Demand Systems with Binding Nonnegativity Constraints: The Dual Approach.” *Econometrica*, 54(5): 1237–42.
- Poi, Brian P.** 2012. “Easy demand-system estimation with quads.” *Stata Journal*, 12(3): 433–446.
- Polinsky, A. Mitchell, and David T. Ellwood.** 1979. “An Empirical Reconciliation of Micro and Grouped Estimates of the Demand for Housing.” *The Review of Economics and Statistics*, 61(2): 199–205.
- Poterba, James M., Steven F. Venti, and David A. Wise.** 2017. “The asset cost of poor health.” *The Journal of the Economics of Ageing*, 9(C): 172–184.
- Prados, María José.** 2018. “Health And Earnings Inequality Over The Life Cycle: The Redistributive Potential Of Health Policies.” PhD diss. Columbia University, Department of Economics.
- Prinz, Daniel, Michael Chernew, David Cutler, and Austin Frakt.** 2018. “Health and Economic Activity Over the Lifecycle: Literature Review.” National Bureau of Economic Research Working Paper 24865.
- Pudney, Stephen, and Nicole Watson.** 2013. “If at first you don’t succeed? Fieldwork, panel attrition, and health-employment inferences in BHPS and HILDA.” Institute for Social and Economic Research (ISER), University of Essex 2013-27.
- Ray-Chaudhuri, Sam, and Tom Waters.** 2024. “Recent trends in and the outlook for health-related benefits.” Institute for Fiscal Studies Report 19 April.
- Ray, Ranjan.** 1983. “Measuring the costs of children: An alternative approach.” *Journal of Public Economics*, 22(1): 89–102.
- Rouwendal, Jan, and Rob Alessie.** 2002. “House Prices, Second Mortgages and Household Savings; An Empirical Investigation for the Netherlands, 1987-1994.” Tinbergen Institute Tinbergen Institute Discussion Papers 02-074/3.
- Roy, A. D.** 1951. “Some Thoughts on the Distribution of Earnings.” *Oxford Economic Papers*, 3(2): 135–146.
- Røed Larsen, Erling.** 2014. “The Engel curve of owner-occupied housing consumption.” *Journal of Applied Economics*, 17(2): 325–352.

- Saiz, Albert.** 2010. “The Geographic Determinants of Housing Supply.” *The Quarterly Journal of Economics*, 125(3): 1253–1296.
- Salis, Sergio, Stefan Speckesser, Lea Samek, and Paul Bivand.** 2021. “A study of work and health transitions: Analysis of Understanding Society.” Department for Work and Pensions DWP research report 994.
- Salvati, Francesca.** 2021. “Health Inequality, Labor Supply and Retirement Policies.” PhD diss. University College London, Department of Economics.
- Sanders, Carl.** 2012. “Skill Uncertainty, Skill Accumulation, and Occupational Choice.” Society for Economic Dynamics 2012 Meeting Papers 633.
- Sarafidis, Vasilis, and Neville Weber.** 2015. “A Partially Heterogeneous Framework for Analyzing Panel Data.” *Oxford Bulletin of Economics and Statistics*, 77(2): 274–296.
- Sheiner, Louise.** 1995. “Housing Prices and the Savings of Renters.” *Journal of Urban Economics*, 38(1): 94 – 125.
- Simonetti, Irene, Michele Belloni, Elena Farina, and Francesca Zantomio.** 2022. “Labour market institutions and long term adjustments to health shocks: Evidence from Italian administrative records.” *Labour Economics*, 79: 102277.
- Szerman, Christiane.** 2024. “The Labor Market Effects of Disability Hiring Quotas.” PhD diss. University College London, Department of Economics.
- Tanaka, Atsuko.** 2021. “The effects of sudden health reductions on labor market outcomes: Evidence from incidence of stroke.” *Health Economics*, 30(6): 1480–1497.
- van Horen, Neeltje, and Belinda Tracey.** 2022. “Help to Spend? The Housing Market and Consumption Response to Relaxing the Down Payment Constraint.” C.E.P.R. Discussion Papers CEPR Discussion Papers 16144.
- Wallenius, Panagiotis Margaris Johanna.** 2020. “Can Wealth Buy Health? A Model of Pecuniary and Non-Pecuniary Investments in Health.” Centre for Economic Policy Research Discussion Paper DP14796.
- Waters, Tom, and Thomas Wernham.** 2023. “Housing quality and affordability for lower-income households.” Institute for Fiscal Studies Report R260.
- Windmeijer, Frank.** 2005. “A finite sample correction for the variance of linear efficient two-step GMM estimators.” *Journal of Econometrics*, 126(1): 25–51.
- Yoshikawa, Hiroshi, and Fumio Ohtaka.** 1989. “An analysis of female labor supply, housing demand and the saving rate in Japan.” *European Economic Review*, 33(5): 997 – 1023.
- Zabel, Jeffrey E.** 2004. “The demand for housing services.” *Journal of Housing Economics*, 13(1): 16–35.

**Zhang, Lu.** 2019. “Do house prices matter for household consumption?” CPB Netherlands Bureau for Economic Policy Analysis CPB Discussion Paper 396.

**Zweifel, Peter, Friedrich Breyer, and Mathias Kifmann.** 2009. *Health economics [electronic resource]*. . 2nd ed. ed., Dordrecht ; New York:Springer.