THE UNIVERSITY OF CHICAGO

HETEROGENEOUS AND UNCERTAIN HEALTH DYNAMICS AND WORKING DECISIONS OF OLDER ADULTS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
ANGELA DENIS PAGLIERO

CHICAGO, ILLINOIS
JUNE 2021

To my dad.

TABLE OF CONTENTS

LIST OF FIGURES vi

LIST OF TABLES vii

ACKNOWLEDGEMENTS viii

ABSTRACT ix

1 HETEROGENEOUS AND UNCERTAIN HEALTH DYNAMICS AND WORKING DECISIONS OF OLDER ADULTS. 1

1.1 Introduction 1

1.2 Framework 9

1.3 Data and descriptive statistics 16

1.4 Health process with heterogeneous dynamics 22

1.5 Uncertain health dynamics and beliefs. 27

1.6 Working decisions as functions of beliefs about health 34

1.7 Reducing bias in initial beliefs 45

1.8 An information experiment: Blood-based biomarkers as signals of health . . 47

1.9 Conclusion 56

REFERENCES 59

APPENDICES 64

A Estimation of the summary measure of health h_{it} 64

B MLE results for health and survival 66

C Robustness: MLE results for self-assessed health 69

D Robustness: MLE with heteroskedastic errors ϵ_{it} 71

E Formula for $plive10_{it}$ 73

F Proof of identification of λ with ideal data 75

G Strategy for simulating survival expectations $plive10_{it}$ 77

H Probit likelihood and results 80
I Neural network details 84
J Additional biomarker results 88

LIST OF FIGURES

1.1 Summary health variable h_{it} by category of self-assessed health 19

1.2 Mean and variance of health by age 20

1.3 Health percentiles by age 20

1.4 Mean of health with age for individuals with different survival ages 20

1.5 Histogram of survival expectations to age 85 21

1.6 Mean and variance of health in models with different assumptions about slope heterogeneity and survival 26

1.7 Simulated moments of $plive10_{it}$ by uncertainty λ in data-generating process . 31

1.8 Average marginal effect of expected beliefs $\hat{\delta}_{it-1}$, and health h_{it-1} on the probability of working p_{it} for adults in their 50s 43

1.9 Average marginal effect of expected beliefs $\hat{\delta}_{it-1}$, and health h_{it-1} on the probability of working p_{it} for adults between 66 and 75 years old 44

1.10 Impulse-response function to a shift in prior beliefs eliminating overall bias b 46

1.11 Impulse-response function to reducing individuals' initial bias by half 47

1.12 Timing of the biomarker collection and information experiment 50

I1 Observed versus predicted probability of work 87

LIST OF TABLES

1.1 Summary statistics for health-related questions 18

1.2 Summary statistics for variables used in studying working decisions 23

1.3 MLE results on health and survival 25

1.4 SMM results on prior beliefs 34

1.5 Moments' fit 35

1.6 Probit results on probability of working 38

1.7 Average marginal effects on the probability of working and decomposition of the effects of a health shock 42

1.8 Information and interview-mode effects of biomarker experiment 52

1.9 Information and interview-mode effects by education level 53

1.10 Predicted survival expectations in a model with health and blood glucose as signals 56

A1 CFA results of health measurement 65

B1 MLE results on health and survival under different assumptions 68

C1 MLE results for SAH with and without a survival equation 70

D1 MLE results for health h_{it} allowing for heteroskedastic error ϵ_{it} 72

H1 Probit results on probability of working: main equation 82

H2 Probit results on probability of working: initial condition 83

I1 Hyperparameter space 86

I2 Loss and accuracy at 5^{th} iteration across 30 starting points 86

J1 Percentage of in-person interviews by wave and group 88

J2 Balance tests at wave 8 89

J3 Biomarkers experiment distinguishing bad vs good test results 91

J4 MLE results of health dynamics including biomarker information 92

ACKNOWLEDGEMENTS

I would like to thank my chair, Stéphane Bonhomme, and my advisors Alessandra Voena, Derek Neal, and David Meltzer. It has been an incredible privilege to learn from you, and to get your guidance and support. I will be forever grateful.

I am indebted to the Economics Department at the University of Chicago more generally, for receiving me as a student, supporting me financially, and teaching me so much over the years. It has been a dream.

I am also deeply grateful to my friends, Cristhian Seminario-Amez, Siying Cao, Kai-Wei Hsu, Kyeongbae Kim, Tesary Lin, and Wooyong Lee, who were such a big part of this experience. I am very lucky to have met you.

My dad, Miguel Denis, who unfortunately passed away, taught me to always try my best and to persevere, and he allowed me to dream big. I would definitely not be here without his unwavering support and unconditional love. I know you are proud, dad, wherever you are.

And last but definitely not least, I would like to thank my husband, Antonio Silva, who followed me in this crazy journey, cheered for me every step of the way, and believed in me even when I didn't. You are the biggest blessing of my life. I love you.

ABSTRACT

As the population ages, governments and international organizations are trying to lengthen the labor-force participation of older adults. For older adults, health is an important determinant of working decisions. In this paper, I introduce heterogeneity in health dynamics with age and argue uncertainty about health dynamics affects the working decisions of older adults. Using the Health and Retirement Study, I first show evidence of heterogeneity in health profiles with age. Second, I use subjective survival expectations to infer health beliefs in a Bayesian-learning framework. Third, I flexibly estimate how working decisions depend on those beliefs, using a neural-network approach that does not require additional structure. The results show beliefs have substantial negative bias. That is, on average, individuals incorrectly believe their health will deteriorate too fast. Furthermore, eliminating that bias would increase labor-force participation by up to 2 percentage points. In the last part of the paper, I look at a policy that could affect beliefs: the provision of information on blood-glucose and cholesterol levels. I take advantage of the randomization in the collection and provision of such information. The results show the information has only small effects on beliefs and working decisions, and consequently, policies with larger effects on beliefs are needed to delay retirement.

CHAPTER 1

HETEROGENEOUS AND UNCERTAIN HEALTH DYNAMICS AND WORKING DECISIONS OF OLDER ADULTS

1.1 Introduction

The population is aging rapidly. Worldwide, the median age was 40 years old in 2018 and is estimated to be 45 years old by 2050. And though the participation of older adults in the labor market has also been recently increasing, the number of older people out of the labor force who will need to be supported by each worker is projected to increase by around 40% between 2018 and 2050. This aging pattern puts considerable strain on public budgets, therefore, promoting employment at older ages has garnered large interest.[1] The success of policies promoting the employment of older adults depends on our correct understanding of the determinants of working decisions of this group, for whom health is an important factor. For older adults, health deteriorates naturally with aging, affecting retirement choices and expectations.[2] Yet, little is known about how heterogeneous health dynamics of older adults are and how this heterogeneity affects their working decisions.

This paper documents individual-level heterogeneity in health dynamics among older adults and studies how individuals' beliefs about their own health dynamics affect their working decisions. To do so, the paper proceeds in three parts. In the first part, I show evidence that health dynamics are indeed heterogeneous among older adults. That is, while some individuals see their health slowly deteriorating with age, other individuals see their health deteriorating much more rapidly. I argue this heterogeneity, which the literature has

[1] See statistics from OECD (2019). In 2015, the OECD adopted an agenda promoting employment at older ages, to protect living standards and public finances (OECD (2015)).

[2] For health effects on retirement choices, see, for example, Bound et al. (1999) and Maurer et al. (2011). For health effects on retirement expectations, see Dwyer and Mitchell (1999) and McGarry (2004).

mostly ignored, is an important factor in the working decisions of older adults. Furthermore, what matters for those decisions is how much individuals know about their own health profiles. Hence, in the second part of the paper, I study uncertainty in health dynamics by developing a Bayesian learning model in which individuals have beliefs about their own health profiles and update those beliefs as they see their health changing with age. I leverage data on survival expectations to infer these beliefs and to quantify how uncertain individuals are. Then, in the third part of the paper, I estimate the working decisions implied by an economic model that incorporates heterogeneous and uncertain health dynamics. I focus on the effects that health beliefs have on working decisions of older adults. Instead of following a structural estimation approach, I use machine-learning tools. A big limitation of this approach is that I cannot run counterfactual analyses. However, the tools do not require specifying the primitives of the model or adding almost any functional-form assumption. Thus, the results are robust to misspecification of those elements.

Using the Health and Retirement Study (HRS), the first part of the paper leverages the longitudinal nature of the data to estimate a dynamic model of health allowing for more general forms of heterogeneity.[3] In particular, I assume health is a persistent process with individual-level heterogeneity both in levels and in changes with age. The results show significant heterogeneity, in levels and in changes. Furthermore, the heterogeneity in changes helps explain the increasing variance of health with age, a pattern observed in the population but mostly ignored by traditional models of health.[4]

The panel estimates in the first part of the paper provide evidence of individual-level heterogeneity in health dynamics, but they do not address the question of how much individuals know about their own health profiles. In the second part of the paper, I study this question using a Bayesian learning model[5] with initial beliefs that allow for bias (through the mean)

[3]Most of the literature allows only for individual heterogeneity in health levels. See, for example, Contoyannis et al. (2004) and Heiss (2011).
[4]See, for example, Heiss (2011) and Heiss et al. (2014).
[5]The health process in this paper is similar to the income process studied by Guvenen (2007), who shows

and uncertainty (through the variance). Data on subjective survival expectations,[6] available in the HRS, allow me to identify the parameters governing these beliefs. Intuitively, future survival depends on future health; hence, expectations about future survival depend on beliefs about future health, and therefore on beliefs about health profiles. Thus, according to the model, survival expectations are a complex nonlinear function of health and health beliefs. Hence, I use simulated method of moments to estimate the parameters of those beliefs. Average survival expectations speak to bias in beliefs. Covariance between changes in health and changes in expectations speak also to uncertainty. To see this, note that, given a change in health, individuals update their survival expectations for two reasons: first, because the persistence of health implies future health is affected by a health change today, and second, because the uncertainty and the learning model imply beliefs are updated with a health change today. Moreover, the larger the persistence and the larger the uncertainty, the larger the change in survival expectations. Hence, moments of survival expectations are the key source for identification of beliefs. My results show individuals are uncertain, updating their beliefs over time, and they are negatively biased; that is, on average, they believe their health will deteriorate faster than the average rate in the population.[7]

The heterogeneity and uncertainty in health dynamics imply beliefs about health profiles enter the decisions of forward-looking individuals. In the third part of the paper, I study how these beliefs affect the working decisions of older adults. In particular, this step requires estimating the relationship between working decisions and all the information available to individuals at the moment they make those decisions, including their beliefs about their health profiles. Under the Bayesian assumptions of the learning model, including normality

that although learning of heterogeneous levels occurs fairly rapidly, learning of heterogeneous slopes with age is much slower.

[6]Survival expectations have been shown to have predictive power for individuals' survival and to be consistently updated with new health information. See, for example, Hurd et al. (2001), Hurd and McGarry (2002), and Smith et al. (2001).

[7]As discussed section 1.5, this result of negative bias in initial beliefs is consistent with the literature. See Elder (2013) and Ludwig and Zimper (2013).

of beliefs, those beliefs are summarized by their mean and variance. This section is similar in spirit to Arellano et al. (2017), who estimate a nonlinear policy rule for consumption nonparametrically, without specifying a full structural model. As in their case, I provide estimates of marginal quantities, here, marginal changes in the probability of working. One of the drawbacks of this approach is that by not fully specifying the structure of the model, both in terms of primitives and functional-form assumptions, we cannot perform policy counterfactuals. However, at the same time, the results on marginal effects are robust to misspecification of those elements. Besides robustness to misspecification, another attractive feature of this framework is that it can also be applied to study other outcomes that may depend on health beliefs, such as savings and health insurance of older adults. To the best of my knowledge, this paper is the first to study the effect of beliefs about heterogeneous health dynamics on the working decisions of older adults.

To flexibly estimate the policy rule for working decisions, I use neural networks. Neural networks are a tool within the machine-learning toolkit that, in the present context, generalize logit with a non-linear index (see Farrell et al. (2021), Hornik et al. (1989), Goodfellow et al. (2016)). To deal with the fact that some of the inputs are unobserved by the econometrician (mainly, the individual-level heterogeneity in initial beliefs), I use an iterative approach in the spirit of EM algorithms (Dempster et al. (1977)).

I discuss three results related to beliefs and working decisions of older adults. The first result shows beliefs matter in working decisions, and that expecting health to deteriorate more slowly is associated with larger probabilities of working. Furthermore, for individuals in their 50s who are not working, an interaction occurs between beliefs and health. The effects on working probabilities of both beliefs and health are larger for individuals who believe their health will deteriorate relatively slowly. These results suggest adjustment costs of finding a job are important in individuals' decisions about going back to work. This observation highlights an advantage of the current framework and the data-driven estimation method,

because this result is not a consequence of any assumption on the structure of the economic model. As mentioned by Arellano et al. (2017), economic structure could be added to this framework to conduct policy evaluation exercises. This result suggests that when adding such structure, adjustments costs of returning to the labor market should be included. Hence, the two approaches complement each other.

A second result is related to the resolution of uncertainty about health profiles and the precision of health as a signal. A health shock has two effects on working decisions: it affects working decisions by changing the stock of health through persistence, and it affects working decisions by changing beliefs about future health through changes in information about health profiles. I decompose the effect of a health shock into these two channels, namely, persistence and information, and find nearly all the effect goes through the persistence channel. Intuitively, this result comes from the signal-to-noise ratio of health being low, and it implies health by itself is not enough to resolve the uncertainty and correct the bias in beliefs.

In a third result, I simulate the impact of changing beliefs, by applying machine-learning tools to predict not only work but also assets and health insurance. I use those results to compare baseline working probabilities over time with probabilities after eliminating initial overall bias in beliefs. I find eliminating initial bias increases participation by 2 percentage points, an effect that lasts beyond traditional retirement ages.

Given that (i) individuals are uncertain about their health profiles, (ii) they have biased initial beliefs, (iii) health changes are not enough to resolve uncertainty, and (iv) beliefs matter for working decisions, a natural question that follows is: Can we provide additional information to individuals in order to correct their beliefs and affect their working decisions? In the last part of the paper, I look at this question in the context of an information experiment available in the HRS. Starting in 2006, the HRS collects and analyzes blood samples of their interviewees and informs them about their blood-glucose and cholesterol results. Although the implementation in the HRS was not designed as an information experiment, in

order to save costs, the blood sample is collected for a random half of the sample each wave, providing us with exogenous variation. A reduced-form analysis in the spirit of difference-in-differences[8] shows small and insignificant effects of this additional information on survival expectations and working decisions. I then analyze these data through the lens of the model. That is, I modify the learning model to include biomarker results as additional signals of health profiles. Consistently, this model-based analysis also shows small and insignificant results. The model, however, provides us with an interpretation for the results: the magnitude of this blood-based signal is too small.

Contribution to the literature. This paper is related to three strands of the literature. First, it is related to the literature studying health dynamics, a literature that consistently finds persistence and heterogeneity in health, both among the general population (Halliday (2008), Hernández-Quevedo et al. (2008), Contoyannis et al. (2004)) and among older adults (Heiss et al. (2009), Heiss (2011), Heiss et al. (2014), Lange and McKee (2011)). However, most of this literature allows only for limited heterogeneity. An exception is Halliday (2008), who allows for discrete types of multivariate heterogeneity, including heterogeneity in health changes with age. Contrary to my results, he finds only weak evidence of this heterogeneity. However, he focuses on a much younger population, whereas I focus on older individuals for whom health changes with age are prevalent. Thus, a first contribution of this paper is to highlight heterogeneity in health dynamics for older adults. An additional contribution to this literature is related to health measurement. Traditionally, health has been considered a latent variable measured with one binary variable (Halliday (2008), Hernández-Quevedo et al. (2008), Heiss et al. (2009), Heiss (2011)), though, more recently, several measures of health are being used (Heiss et al. (2014), Lange and McKee (2011), Blundell et al. (2017)). In this paper, I also use several measures of health to better capture the richness of health

[8]As discussed in section 1.8, the design needs to control also for changes in the interview mode.

and its dynamics, hence contributing in this direction.

Second, this paper is related to the literature on empirical learning. In a broad sense, the paper is related to the literature on the importance of beliefs for individuals' choices and economic outcomes.[9] More specifically, the paper is related to the literature studying individuals' learning of own unobserved heterogeneity, for example, regarding abilities (Stinebrickner and Stinebrickner (2014), Arcidiacono et al. (2016)), productivity (Arcidiacono et al. (2016)) and income profiles (Guvenen (2007), Guvenen and Smith (2014)). My paper is more closely related to Guvenen and Smith (2014), who study an income process with heterogeneous levels and heterogeneous growth rates. As in the case of health, the more flexible heterogeneity helps explain the income pattern of increasing variance over time. However, an important difference from that paper is the source of identification of profile uncertainty. Guvenen and Smith (2014) use consumption data to identify uncertainty in income profiles. Instead, I use data on expectations to identify uncertainty in health profiles. This difference is important because my goal is to study the effect of uncertainty regarding health dynamics on working decisions of older adults, and hence, using that outcome to identify beliefs would mean my results could suffer from misspecification issues. By using expectations data, my results are robust to such issues. I also allow for individuals to be biased overall in their initial beliefs, consistent with findings from the literature on survival expectations (see Elder (2013) and Ludwig and Zimper (2013)). Additionally, this paper contributes to a more recent literature on the provision of information and its effects on beliefs (see, e.g., Delavande and Kohler (2015), Wiswall and Zafar (2014), Bates (2020)).

Finally, the paper is related to the literature on health and other outcomes of older adults. Particularly, the paper is related to the literature studying the effects of health on work and

[9]Outcomes studied by this literature include occupational choices and college attrition (Breen and García-Peñalosa (2002), Arcidiacono et al. (2020), Arcidiacono et al. (2016)), labor supply of women and employment transitions (Gong et al. (2019), Conlon et al. (2018)), birth-control choice and risky sexual behaviors (Delavande (2008), Paula et al. (2014), Delavande and Kohler (2015)), and investment decisions (Delavande and Rohwedder (2011)).

retirement choices (Siddiqui (1997),McClellan (1998), Bound et al. (1999), French (2005), Disney et al. (2006), Zucchelli et al. (2010), Maurer et al. (2011)) and expectations (Dwyer and Mitchell (1999), McGarry (2004)). Although this literature considers future health as uncertain, it assumes a known stochastic process for health. On the contrary, this paper allows for a stochastic health process that is not fully known, introducing the role of health beliefs as an additional determinant of those decisions. More broadly, this paper is also related to a series of papers studying health-related outcomes for older individuals. These papers estimate structural models assuming discrete values for health with homogeneous transition probabilities. Examples include papers studying the effect of health insurance on retirement (French and Jones (2011), De Nardi et al. (2016a)), Social Security and labor supply (van der Klaauw and Wolpin (2008)), portfolio choice (Yogo (2016)), and long-term care (Ameriks et al. (2020), Lockwood (2018)). Though health is not the main explanatory variable of interest in these papers, the results here suggest beliefs about health may also play a role.

Organization. The paper proceeds as follows. Section 1.2 presents the framework, that is, an economic model of working decisions that incorporates heterogeneous and uncertain health dynamics. This framework underlies and gives context to the analysis in the rest of the paper. Section 1.3 describes the data. Section 1.4 provides evidence of heterogeneity in health dynamics, and section 1.5 provides evidence of uncertainty. Section 1.6 presents the main results for working decisions as a function of beliefs, and section 1.7 expands those results. Section 1.8 analyzes the information experiment available in the HRS. Section 1.9 concludes.

1.2 Framework

This paper introduces two elements into a standard model of labor-participation decisions in late life: individual-level heterogeneity in health dynamics and individuals' uncertainty regarding their own health profile. This section formalizes this idea and describes a framework in which older adults choose labor participation based on their health and on their beliefs about how their health will change with age. Let i denote an individual and let t denote his age. I focus on individuals 50 years and older and define t as 0 for age 50.

1.2.1 Health process with heterogeneous dynamics

Health is a dynamic process that, as people get older, naturally deteriorates in a heterogeneous way across individuals. In particular, I assume health is scalar and follows

$$h_{it} = \rho h_{it-1} + \alpha_i + \delta_i \cdot t + \epsilon_{it}. \tag{1.1}$$

The parameter $\rho \in (0, 1)$ captures persistence in health, α_i captures heterogeneous levels in health, δ_i captures heterogeneous changes in health with age, and ϵ_{it} represents health shocks. Both the persistence of health and its heterogeneity in levels are well-recognized elements of health in the literature, both among the general population (see, e.g., Hernández-Quevedo et al. (2008)) and among older individuals (see, e.g., Heiss et al. (2014)). The first novel element in this paper is to allow for heterogeneous slopes of health with age, δ_i. Larger values of h_{it} represent better health, and health decreases with age.

Throughout the paper, I assume health is exogenous. In a review of the literature on health, health insurance, and retirement, French and Jones (2017) mention much of the retirement literature assumes health is exogenous, and their model makes the same assumption. In a review of the literature on savings after retirement, De Nardi et al. (2016b)

conclude most of the studies on the effects of health care on health find small effects. A similar argument is made in French and Jones (2011). The exogeneity assumption implies we can estimate equation (1.1) without needing to model endogenous regressors.[10]

1.2.2 Uncertain health dynamics and beliefs

The second novel element is to allow for individuals to be uncertain about their own health dynamics. I assume individuals observe their health h_{it}, but they do not observe their health shocks ϵ_{it} nor their individual heterogeneity (α_i, δ_i). Given that health deteriorates in old age, I assume 50-year-old individuals do not know δ_i, which has not affected them before.[11] I assume they know their heterogeneous level α_i,[12] because they have observed their health for several decades.

Under uncertainty, rational individuals form beliefs about their health slopes δ_i (henceforth, slope beliefs) and update those beliefs as they see their health changing with age. In particular, I assume individuals are Bayesian learners, with initial beliefs (at age 50) about δ_i equal to $N(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$.[13] By further assuming health shocks ϵ_{it} are i.i.d. normally distributed, posterior beliefs in period t after observing health h_{it} are also normally distributed, $N(\hat{\delta}_{it}, \hat{\sigma}_t^2)$, with mean and variance defined recursively by

$$\frac{\hat{\delta}_{it}}{\hat{\sigma}_t^2} = \frac{\hat{\delta}_{it-1}}{\hat{\sigma}_{t-1}^2} + \frac{(h_{it} - \rho h_{it-1} - \alpha_i)t}{\sigma_\epsilon^2} \tag{1.2}$$

$$\frac{1}{\hat{\sigma}_t^2} = \frac{1}{\hat{\sigma}_{t-1}^2} + \frac{t^2}{\sigma_\epsilon^2}. \tag{1.3}$$

[10] The assumption is also relevant for the identification of beliefs, as discussed in section 1.5.

[11] This assumption is consistent with results from Halliday (2008), who studies health dynamics with discrete heterogeneity, using the Panel Study of Income Dynamics. He studies younger individuals, ages 22 to 60, and finds no heterogeneous slopes with age.

[12] This assumption can be generalized. In studying income profiles, Guvenen (2007) proposes a similar process with heterogeneous intercepts and slopes, both unknown. He finds the learning process for intercepts is much faster than the learning process for slopes.

[13] The assumption of common-prior variance across individuals is usual in the learning literature. See, for example, Guvenen (2007) and Arcidiacono et al. (2016). However, the assumption is important for the identification results provided later.

Equation (1.2) shows the posterior mean is a weighted average of the prior mean $\hat{\delta}_{it-1}$ and the signal derived from health h_{it}, with weights that depend on precision. The more certain an individual is to begin with (lower $\hat{\sigma}^2_{t-1}$), the more weight he gives to what he already knows, namely, the prior. The more precise health is as a signal (lower σ^2_ϵ), the more weight is given to its information. Equation (1.3) shows precision increases over time, and increases more when the signal is more precise, that is, when health is less noisy (lower σ^2_ϵ) and when individuals are older.

Conditional on health history, the key parameters determining beliefs are the parameters governing initial beliefs:

$$b = \mathbb{E}(\hat{\delta}_{i0} - \delta_i) \tag{1.4}$$

$$\lambda^2 = \frac{\hat{\sigma}^2_0}{Var(\delta_i)}. \tag{1.5}$$

The parameter b measures the bias in initial beliefs at the population level. If $b = 0$, individuals are overall unbiased, in the sense that $\mathbb{E}(\hat{\delta}_{i0}) = \mathbb{E}(\delta_i)$. If b is positive (negative), individuals are upward (downward) biased, and hence, they believe health deteriorates on average more slowly (faster) than the average rate. The parameter λ measures the degree of initial uncertainty individuals face regarding δ_i, which affects their amount of learning over time. If $\lambda = 0$, no uncertainty exists and therefore no learning. The larger the value of λ, the more uncertain individuals are and the more weight they give to new information. The Bayesian learning and normality assumptions allow me to reduce the dimensionality of the problem, giving structure to time-varying beliefs that are unobserved by the econometrician.

1.2.3 Embedding health uncertainty in a model of labor supply

In a life-cycle model, forward-looking individuals attempt to predict variables that will affect their future utility or their future set of options in order to choose their best current

action. The need for those predictions is given by the inherent uncertainty about many key variables. In this paper, I focus on working decisions of older adults and argue a key source of uncertainty for this group is related to their future health. In particular, I focus on uncertainty about health profiles with age, specifically δ_i, and study how beliefs about them, given by $N(\hat{\delta}_{it}, \hat{\sigma}_t^2)$, relate to their working decisions.

Consider a model where individual i must choose consumption c_{it} and labor participation p_{it} every period. I focus on the extensive margin of labor participation and assume p_{it} is a binary decision. The health of individual i is given by h_{it}, which follows equation (1.1). The main components of this life-cycle model are the following.

Preferences. Individual i's flow utility is given by a function U that depends on his participation and consumption decisions, p_{it} and c_{it}, as well as on his health h_{it}. Furthermore, preferences depend on past labor participation, for example, to reflect psychological costs of going back to work after retirement and adjusting to a new work environment. I summarize this dependence by allowing p_{it-1} to enter the utility function. Hence, flow utility is given by $U(p_{it}, c_{it}, h_{it}, p_{it-1})$. The individual discounts the future, and when he dies, his remaining assets a are left as a bequest.

Budget constraint. Let a_{it-1} denote individual i's assets at the end of period $t-1$. If the individual chooses to work, he receives labor income, which depends on his past labor income w_{it-1}, his health h_{it} due to the effects of health on productivity, and his past participation p_{it-1}, due to wage penalties of reentering the labor market after retirement. His assets at the end of the period depend also on his consumption choice, his other sources of income, including pension and social security, and other health-related costs.

Uncertainty. Individuals are uncertain about their future health, in part because of unpre-

dictable health shocks ϵ_{it}, and in part because they don't know their health slopes δ_i. They form beliefs about their slopes δ_i and update those beliefs as they see their health changing over time according to equations (1.2) and (1.3). Future wages are also uncertain, following a first-order Markov process.

Timing. At the beginning of period t, an individual must choose participation p_{it} and consumption c_{it} before health shocks are realized and health h_{it} is observed. Then, beliefs are updated. At the end period t, individual i may or may not die.

Information set. The information set of individual i at the beginning of period t is given by his history up to $t-1$ in terms of labor participation p_i^{t-1} (superscripts denote histories), consumption c_i^{t-1}, and health h_i^{t-1}, as well as labor income w_i^{t-1}. It also includes his known value α_i and his prior-beliefs parameters $\hat{\delta}_{i0}$ and $\hat{\sigma}_0^2$. The relevant information from this set can be summarized in his state variables, given by

$$\Omega_{it-1} \quad = \quad \left\{p_{it-1}, a_{it-1}, w_{it-1}, h_{it-1}, \hat{\delta}_{it-1}, \hat{\sigma}_{t-1}^2, \alpha_i\right\}.$$

Slope uncertainty implies δ_i does not belong to Ω_{it-1} but beliefs about δ_i do, with those beliefs summarized by $\hat{\delta}_{it-1}$ and $\hat{\sigma}_{t-1}^2$. Note I am assuming only heterogeneity in health; thus, no other individual-level heterogeneity is stated in Ω_{it-1}.

The solution to this problem is policy rules for labor participation p_{it} and consumption c_{it}, which are functions of the state variables and the parameters of the model θ (including the discount factor and parameters entering flow utility, health process, the budget constraint, and so on), which I omit for ease of notation. Focusing on p_{it}, which is the object of interest

in this paper,

$$\mathbb{P}(p_{it}=1|\Omega_{it-1}) \quad = \quad \mathbb{P}(p_{it}=1|p_{it-1},a_{it-1},w_{it-1},h_{it-1},\hat{\delta}_{it-1},\hat{\sigma}^2_{it-1},\alpha_i). \tag{1.6}$$

Similarly, policy rules for other decisions, including consumption and assets, can be written as functions of the state variables Ω_{it-1}. In the spirit of Arellano and Bonhomme (2016) and Arellano et al. (2017), the objective is to flexibly estimate this relation between participation decisions and its inputs, without adding the full structure required by structural models.

Equation (1.6) assumes these decisions are stochastic. Implicitly, I assume random taste shifters are affecting individuals' preferences. These taste shifters are part of the state variables, but they are unobserved by the econometrician. Hence, from the econometrician's point of view, the decision is stochastic, corresponding to a conditional choice probability problem.

With these elements, the model is a standard model of labor participation in late life and includes several channels through which health can play a role. First, health directly affects utility by changing the marginal utility of consumption and the disutility of work. Second, it enters the budget constraint via health-related costs and via effects on labor income due to changes in productivity. Third, health affects the probabilities of survival. The overall effect of health on individuals' participation decisions depends on all of these channels. The novel element in this paper is that beliefs about future health also play a role. They could have a positive or negative effect, depending on the relative importance of these channels in the individual's problem. For example, if an individual predicting better future health wants to work longer, the sign of beliefs would be positive. This case would arise if the dominant effect were the desire to save more given the longer life expectancy implied by better health. If an individual predicting worse future health wants to work longer, the sign would be negative.

This case would arise if the dominant effect were the desire to save more given the higher cost of future health care implied by worse health.

1.2.4 Objective of the paper

Under this framework, the objectives of the paper are the following:

(i) To document heterogeneity in health dynamics among older adults, particularly heterogeneity in δ_i.

(ii) To study older adults' beliefs about their health dynamics, in particular, to estimate their initial bias b and their initial uncertainty λ.

(iii) To examine whether these beliefs have an effect on working decisions of older adults, by studying the effect of marginal changes in beliefs on those decisions,[14]

$$\frac{\partial \mathbb{P}(p_{it} = 1|\Omega_{it-1})}{\partial \hat{\delta}_{it-1}}. \tag{1.7}$$

One goal of this paper is to estimate equation (1.7) flexibly, without imposing any additional structure on the model of labor supply (such as preferences, labor income process, and so on). A flexible estimation provides results that are robust to misspecification issues on that model. The paper uses a data-driven estimation method that allows me to achieve that flexibility. Furthermore, as discussed in section 1.6, this data-driven approach allows the data to suggest mechanisms that may be overlooked otherwise. Nevertheless, this framework could also be applied under a structural approach, by adding assumptions about the

[14] I focus on the marginal effect of the posterior mean $\hat{\delta}_{it-1}$ and not of the posterior variance $\hat{\sigma}^2_t$. The reason for this choice is that the posterior variance $\hat{\sigma}^2_{t-1}$ is common across individuals. Thus, I do not have variation in the data to separately identify its effects from the effects of age t, without relying on functional-form assumptions.

different elements in the model. A structural approach, on the other hand, has the advantage of allowing for interesting counterfactual analysis. Hence, the objective of the current approach is not to compete with structural models, but to complement them.

In this context of uncertain health dynamics, an additional interesting question is related to the dual role of health shocks ϵ_{it-1} in working decisions. On the one hand, a health shock ϵ_{it-1} affects h_{it-1}, which in turn affects h_{it} through persistence of the health process. This persistence effect disappears if $\rho = 0$. On the other hand, an uncertain individual cannot perfectly distinguish between ϵ_{it-1} and δ_i within h_{it-1}. Hence, the effect of a shock ϵ_{it-1} on h_{it-1} is partly interpreted as new information regarding δ_i, affecting beliefs $\hat{\delta}_{it-1}$. This information channel disappears if $\lambda = 0$. Using Bayes' rule, we can write,

$$\frac{d\mathbb{P}(p_{it}=1|\Omega_{it-1})}{d\epsilon_{it-1}} = \underbrace{\frac{\partial \mathbb{P}(p_{it}=1\Omega_{it-1})}{\partial h_{it-1}}}_{\text{persistence channel}} + \underbrace{\frac{\partial \mathbb{P}(p_{it}=1|\Omega_{it-1})}{\partial \hat{\delta}_{it-1}} \overbrace{\frac{(t-1)\hat{\sigma}^2_{t-1}}{\sigma^2_\epsilon}}^{\text{factor}}}_{\text{information channel}}, \tag{1.8}$$

where the factor term corresponds to the change in the posterior mean $\hat{\delta}_{it-1}$ given a marginal change in ϵ_{it-1}, and it is related to the signal-to-noise ratio of health as a signal. The term is larger when more uncertainty exists concerning the unknown δ_i and when the variance of the health shocks is smaller. How important these channels are in explaining the total effect of a health shock on working decisions of older adults is, then, an empirical question

1.3 Data and descriptive statistics

For this study, I use data from waves 4 to 12 of the Health and Retirement Study (2014) (HRS),[15] a longitudinal survey representative of the population 50 years and older in the US. This survey interviews individuals and their spouses every two years and includes several

[15]The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

measures of health, questions about expectations, information about labor participation and retirement, as well as income and wealth variables.[16] For most of the analysis, I use the RAND HRS Longitudinal File (2014).[17] In this section, I briefly describe the variables used in this study.

1.3.1 Data on health

The most common measure of health used in the literature is self-assessed health, an ordinal variable taking five values from very poor to excellent. It has been shown to correlate with several outcomes, including education, income, savings, retirement, and health insurance. Still, its limited range makes it not ideal in studying health dynamics with age. The HRS, however, provides a larger battery of health-related questions, which I exploit to construct a summary measure of health via factor analysis that I use throughout the paper. This approach of using several measures to construct a summary variable is not unique to this paper; see, for example, Heiss et al. (2014), Lange and McKee (2011), and Blundell et al. (2017). Table 1.1 presents summary statistics for these health-related questions and for the summary health measure. Note these measures reflect a health concept that is the relevant one for the working decisions of older adults, related to how individuals perceive their health in relation to their everyday activities. Appendix A provides details on the estimation of the summary measure h_{it} via factor analysis. The scale of h_{it} is set to be the inverse scale of the number of chronic conditions, which ranges from 0 to 7. That is, larger values of h_{it} represent better health, and an increase of one unit in h_{it} corresponds to one less chronic condition. Figure 1.1 shows a box plot for h_{it} per value of self-assessed health. Both measures are highly correlated, but h_{it} captures more variation than what we can capture with

[16] I exclude proxy interviews because these interviews do no ask questions about survival expectations.

[17] The RAND HRS Longitudinal File is an easy-to-use dataset based on the HRS core data. This file was developed at RAND with funding from the National Institute on Aging and the Social Security Administration.

Table 1.1: Summary statistics for health-related questions

Variable	Observations	Mean	SD	Min	Max
Number of chronic conditions	156,968	5.17	1.34	0	7
Self-assessed health	156,862	2.86	1.11	1	5
Body mass index (kg/m^2)	154,602	27.89	5.81	7	83
Eyesight in general	156,768	2.85	1.01	1	6
Eyesight at a distance	156,833	2.57	1.01	1	6
Eyesight up close	156,822	2.75	1.04	1	6
Hearing	156,869	2.63	1.09	1	5
Pain	156,550	0.63	0.97	0	3
Difficulties in ADLs regarding mobility	156,748	1.09	1.45	0	5
Difficulties in ADLs of large muscles	156,737	1.28	1.33	0	4
Difficulties in other ADLs	151,923	0.40	0.66	0	2
Summary health measure h_{it}	148,866	5.22	0.67	2.96	6.18

Note: Summary statistics for the health measures including the summary health measure. The sample comprises 30,657 individuals interviewed in person, in wave 4 or later, that are 50 years old or older. Chronic conditions include high blood pressure, heart attack, diabetes, stroke, lung disease, arthritis, and cancer. The categories for self-assessed health and hearing include 1. excellent, 2. very good, 3. good, 4. fair, 5. poor. These categories are also the same ones for eyesight variables, but those include alternative 6. legally blind. The categories for the level of pain are 0. no pain, 1. mail pain, 2. moderate, 3. severe. ADL stands for activities of daily living. ADLs regarding mobility include walk 1 block, several blocks, across room, climb one flight of stairs, several flight of stairs. ADLs involving large muscles include push or pull large object, sit for two hours, get up from chair, stoop kneel or crouch. Other ADLs include carry 10 lbs and reach arms.

a discrete measure. Furthermore, as I mention later, the heterogeneity in health dynamics is robust to the use of self-assessed health instead of the summary measure.

Figure 1.2 shows the mean and variance of health h_{it} by age.[18] Given the two years between waves, throughout this paper, I consider age as measured in two-year bins. These plots are the starting point for thinking about health for older adults: they show that with age, the average health in the population decreases while the variance of health in the population increases. This pattern of decreasing mean and increasing variance is robust to

[18]The standard errors in this figure, as well as the following results in this paper, need yet to be adjusted for the estimation of the summary health measure.

Figure 1.1: Summary health variable h_{it} by category of self-assessed health

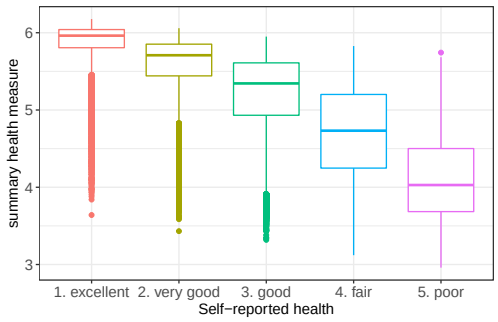

Note: Sample of 148,866 observations from Table 1.1.

sample composition and also holds for most of the individual measures. Similarly, Figure 1.3 shows percentiles of health by age, which also reflect an increasing variance over time. The pattern in these plots suggests a process with heterogeneous slopes with age, which I empirically investigate in section 1.4. Finally, Figure 1.4 shows the mean of health for groups of individuals surviving to different ages. The figure suggests survival bias, because cohorts of individuals surviving to older ages have better health than cohorts that may not survive that long. The relevance of addressing survival bias for older individuals is well recognized in the literature (see, e.g., Heiss et al. (2014)), and I address it also, as explained in section 1.4.

1.3.2 Data on subjective survival expectations

The HRS includes a battery of questions relative to subjective expectations, including subjective survival expectations, which I use in this paper. The question asks, *What is the percentage chance you will live to be (80, 85, 90, 95 or 100) or more?*, where the reference age is a function of the individual's age and the wave of the survey. This reference age is usually around 10 to 15 years into the future.[19] Survival expectations have been shown

[19]The HRS also includes a question on survival expectations to the age of 75. However, I do not use this variable for the main analysis, given that this question is only asked of individuals under 65 years old.

Figure 1.2: Mean and variance of health by age

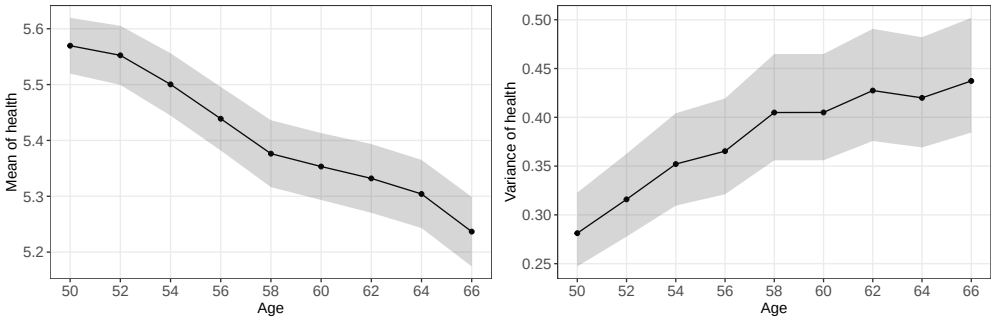

(a) Mean of h_{it} by age (b) Variance of h_{it} by age

Note: Results from a balanced sample of 433 individuals observed at 50 years with at least 9 consecutive waves. The bands represent the 95% confidence intervals.

Figure 1.3: Health percentiles by age

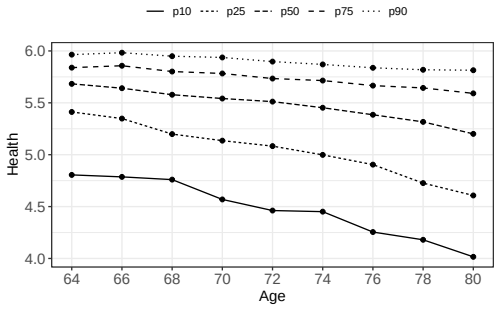

Note: Results from a balanced sample of 414 individuals observed at age 64 with at least 9 consecutive waves.

Figure 1.4: Mean of health with age for individuals with different survival ages

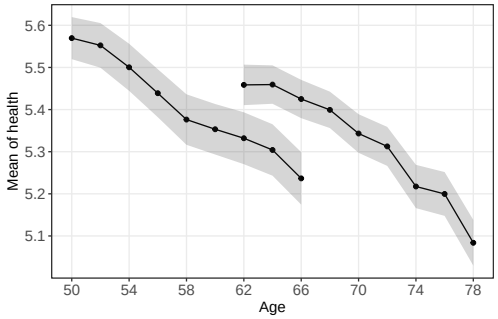

Note: Results from two balanced samples of individuals with at least 9 consecutive waves: 433 individuals observed from age 50, and 509 individuals observed from age 62. The bands represent 95% confidence intervals.

Figure 1.5: Histogram of survival expectations to age 85

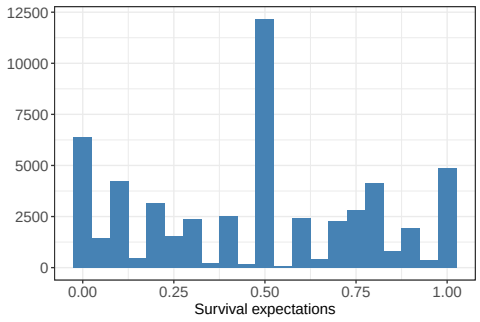

Note: Sample comprises 54,754 observations from individuals interviewed in person, in wave 4 or later, who are 50 years old or older, and who are asked for a reference age of 85 years old. The variable is rescaled to take values between 0 and 1 instead of 0 and 100.

to have predictive power for individuals' survival (Hurd et al. (2001), Hurd and McGarry (1995)) and to be consistently updated with new health information (Hurd and McGarry (2002), Smith et al. (2001)). Furthermore, survival expectations are correlated with several outcomes for older individuals. A histogram of the variable is shown in Figure 1.5. It is well known that this variable suffers from measurement error, including rounding and focal-point issues (Manski and Molinari (2010), Kleinjans and Van Soest (2014)). The model in this paper takes those issues into account.

1.3.3 Data on other variables related to working decisions

In this paper, the main objective is to study how beliefs regarding health profiles affect the working decisions of older adults. As described in section 1.2, doing so requires estimating the policy rule of participation p_{it} as a function of past participation p_{it-1}, health h_{it-1}, heterogeneity in health levels α_i, beliefs regarding health profiles $(\hat{\delta}_{it}, \hat{\sigma}_t^2)$, as well as other variables in the information set Ω_{it-1}, including assets a_{it-1} and labor income w_{it-1}. Table 1.2 presents summary statistics for these other variables in Ω_{it-1} that I use in section 1.6 in

Thus, using this variable would restrict my sample considerably.

predicting working outcomes of older adults.

1.4 Health process with heterogeneous dynamics

This section estimates a health process with heterogeneous intercepts and slopes. As Figure 1.4 suggests, for a population of older adults, we need to control for survival bias, which I address by jointly modeling the two processes, given the lack of a suitable instrument affecting survival chances but not health.

1.4.1 Empirical strategy

Let S_{it} be a binary variable for surviving up to the beginning of period t with $S_{i0} = 1$ and let the health and survival processes be given by

$$h_{it} = \rho h_{it-1} + \alpha_i + \delta_i \cdot t + \tau \cdot t^2 + \epsilon_{it}, \qquad \epsilon_{it} \text{ i.i.d. } N(0, \sigma_\epsilon^2) \tag{1.9}$$

$$S_{it} = \mathbb{1}\{\gamma h_{it-1} + \theta_0 + \theta_1 \cdot t + \theta_2' x_i + \eta_{it} \geq 0\} S_{it-1}, \qquad \eta_{it} \text{ i.i.d. } N(0,1) \tag{1.10}$$

with individual-level heterogeneity (α_i, δ_i),

$$\begin{pmatrix} \alpha_i \\ \delta_i \end{pmatrix} \Bigg| x_i, h_{i0} \sim N\left(\begin{pmatrix} \mu_\alpha + \nu_\alpha' x_i + \omega_\alpha h_{i0} \\ \mu_\delta + \nu_\delta' x_i + \omega_\delta h_{i0} \end{pmatrix}, \begin{bmatrix} \sigma_\alpha^2 & \phi \sigma_\alpha \sigma_\delta \\ \phi \sigma_\alpha \sigma_\delta & \sigma_\delta^2 \end{bmatrix} \right). \tag{1.11}$$

The health process is persistent, measured by the parameter ρ, and it has heterogeneous levels α_i and heterogeneous slopes with age δ_i. The survival process depends on age through the parameter θ_1, and it depends on health through the parameter γ. This dependence of survival on health allows us to take into account the survival bias observed in the data (see Figure 1.4). The health and survival shocks, ϵ_{it} and η_{it}, are assumed to be independent. Appendix B includes a specification allowing for survival to depend directly on individual-level

Table 1.2: Summary statistics for variables used in studying working decisions

Variable	Mean	SD	Min	Max
Panel (a)				
Age	66.26	7.49	52	80
Work	0.38	0.49	0	1
Female	0.52	0.5	0	1
Education: less than high school	0.20	0.40	0	1
Education: some college	0.55	0.50	0	1
White	0.84	0.37	0	1
Hispanic	0.06	0.24	0	1
Marital Status: married	0.70	0.46	0	1
Marital Status: separated or divorced	0.12	0.33	0	1
Marital Status: widow	0.14	0.35	0	1
Number of household members	2.15	1.03	1	12
Total number of years worked	39.79	9.17	20	68
Spouse works	0.28	0.45	0	1
Spouse has health insurance	0.17	0.38	0	1
Income from pension	6.08	50.49	0	10000
Income from Social Security	6.65	5.95	0	58.3
Wealth	366.51	730.98	-1585.01	10000
Health insurance: employer covering retirement	0.14	0.35	0	1
Health insurance: employer not covering retirement	0.07	0.25	0	1
Health insurance: employer (already 65)	0.17	0.37	0	1
Health insurance: government	0.47	0.5	0	1
Health insurance: other	0.11	0.31	0	1
Panel (b)				
Income from work	30.51	39.83	0	1190.68
Tenure	14.31	12.4	0	66.1
Self-employed	0.22	0.42	0	1
Occupation: managerial	0.16	0.36	0	1
Occupation: professional	0.21	0.4	0	1
Occupation: sales	0.12	0.32	0	1
Occupation: clerical	0.16	0.37	0	1
Occupation: services	0.14	0.35	0	1
Occupation: farming, mechanics, construction, operators	0.22	0.41	0	1
Occupation: FF.AA.	0.00	0.02	0	1
Job requires physical effort	0.17	0.38	0	1
Job requires lifting heavy loads	0.07	0.25	0	1
Job requires stooping or kneeling	0.13	0.34	0	1
Job requires good eyesight	0.68	0.47	0	1
Job involves lots of stress	0.16	0.37	0	1

Note: Summary statistics for the variables used in estimating working decisions in section 1.6. The sample consists of observations from 12,623 individuals who have participated in the labor market for at least 20 years, excluding missing values in any of these variables. Panel (a) comprises 48,607 observations, and panel (b) comprises 18,415 observations from working periods. Income and wealth variables are measured in thousands of 2002 dollars. Wealth variables are capped at $10 million.

heterogeneity, α_i and δ_i. However, those results indicate no such dependence. The variables in x_i are time-invariant binary variables for female, white, Hispanic, and an education level below high school graduation. These variables potentially affect health (through the individual-level heterogeneity) and survival. I also allow for the unobserved heterogeneity to depend on health h_{i0} (health at age 50) in order to address initial-conditions concerns.

Under these assumptions, the panel structure of the data identifies the distribution of α_i and δ_i. Let Θ be the set of parameters of this random-coefficients model.[20] I estimate these parameters by maximizing the likelihood:

$$\max_{\Theta} \sum_{i=1}^{N} log\Big(\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \prod_{t=1}^{T_i} \mathbb{P}\big(h_{it}, S_{it} | h_{it-1}, S_{it-1} = 1, x_i, \alpha, \delta\big) \cdot \phi(\alpha, \delta | x_i, h_{i0}) d\alpha d\delta \Big).$$

The full expression of this likelihood is included in appendix B.[21]

1.4.2 Results

I use a sample of 8,901 correlative observations from 1,671 individuals observed since they were 50 years old ($t = 0$). Over the span of the following eight waves, 112 of these individuals died. The main results are shown in Table 1.3 and full results are shown in appendix B. The table shows, first, heterogeneity in both the intercepts and the slopes of the health process, with positive and significant σ_α^2 and σ_δ^2. Second, these two sources of heterogeneity are uncorrelated, which implies knowing α_i does not provide additional information on δ_i. Health decreases with age, and the persistence of the health process is relatively low, with $\rho = 0.22$. The results in the appendix further show that health is worse for individuals with low levels of education, health decreases faster for white individuals, and probabilities of survival are higher on average for women and Hispanic individuals. Those results also show

[20] $\Theta = \{\rho, \tau, \sigma_\epsilon^2, \gamma, \theta_0, \theta_1, \theta_2, \mu_\alpha, \mu_\delta, \nu_\alpha, \nu_\delta, \omega_\alpha, \omega_\delta, \sigma_\alpha^2, \sigma_\delta^2, \phi\}$

[21] For estimation, I approximate the double integral by using $1,000$ draws from a bivariate normal distribution.

Table 1.3: MLE results on health and survival

	Symbol	Coefficient	Pvalue
Persistence	ρ	0.223	0.000
Mean* of α_i	μ_α	0.955	0.000
Mean* of δ_i	μ_δ	-0.057	0.018
SD of α_i	σ_α	0.235	0.000
SD of δ_i	σ_δ	0.043	0.000
$Corr(\alpha_i, \delta_i)$	ϕ	-0.033	0.714
SD of health shocks	σ_ϵ	0.266	0.000
Survival dependence on health	γ	0.583	0.001
Controls		Yes	
N alive observations		8,901	
N dead observations		112	
N individuals		1,671	
-Log likelihood		3,027.6	

Note: Main results of estimating equations (1.9), (1.10), and (1.11). Full set of results are shown in appendix B.

h_{i0} is correlated with α_i, but h_{i0} does not provide information on δ_i.

I want to emphasize two aspects of this model: the inclusion of heterogeneous slopes with age and the joint estimation with survival. To understand how these two aspects influence my results, I estimate two additional versions of the model: (i) one excluding the equation for survival but allowing for heterogeneous slopes with age, and (ii) another one assuming homogenous slopes with age but including an equation for survival. The results are in appendix B and show qualitatively similar results for the coefficients that are common across specifications. Their main difference is that ignoring slope heterogeneity increases the point estimate of the persistence parameter ρ by over 50% (from 0.22 to 0.37). However, a key takeaway is that these models achieve very different fits of health over time. This takeaway is more clearly seen in Figure 1.6, which repeats the exercise for a sample of individuals observed from 66 years old and plots the predicted mean and variance of health

Figure 1.6: Mean and variance of health in models with different assumptions about slope heterogeneity and survival

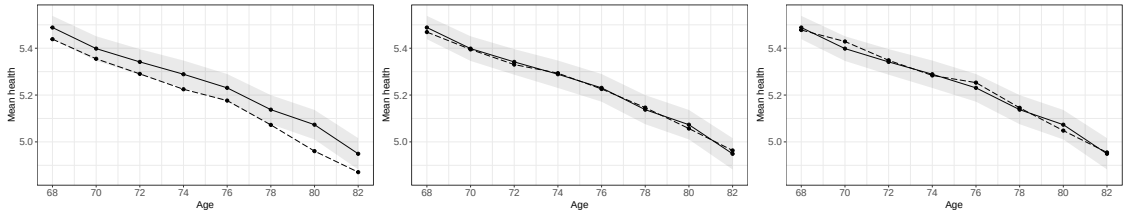

(a) Heterogeneous slopes without survival equation

(b) Heterogeneous slopes and survival equation

(c) Homogeneous slopes and survival equation

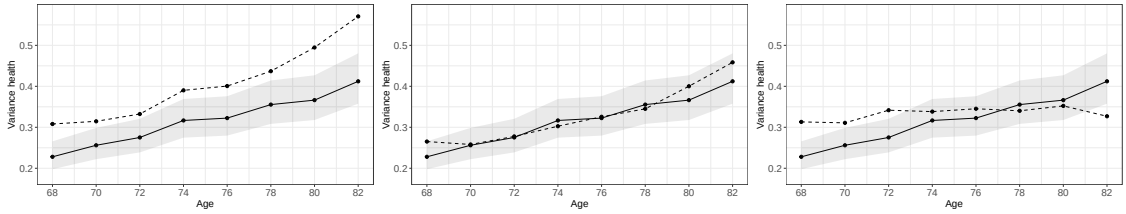

(d) Heterogeneous slopes without survival equation

(e) Heterogeneous slopes and survival equation

(f) Homogeneous slopes and survival equation

Note: The sample consists of 26,950 correlative observations from 7,301 individuals observed since they were 66 years old. Over the span of the following eight waves, 996 of them died. The figure plots data from 354 individuals with health in all 9 waves. The solid lines plot the health data and the dotted lines plot the predicted values of health in each model.

with age. The figure shows ignoring survival leads to a downward bias of average health and an upward bias of its variance, consistent with a model that includes the lower tail of the health distribution, which is dropped from the data as people die. The figure also shows that when ignoring slope heterogeneity, we predict a rather constant variance of health, contrary to what the data show. In that sense, these plots support a model with slope heterogeneity, though they don't discard alternative explanations for the increasing variance with age. As a robustness check, included in appendix D, I estimate a version with heteroskedastic error ϵ_{it}, allowing its variance to depend on age. The results show an increasing variance of health shocks does not explain away the heterogeneity in slopes δ_i.

Finally, I add two robustness checks included in the appendices. First, I estimate a similar model using self-assessed health instead of the constructed summary measure of health. The

results show the presence of heterogeneous slopes with age is robust to the use of this measure alone. Second, I estimate a version of the model adding the unobserved heterogeneity (α_i, δ_i) directly to the survival equation. The results show α_i and δ_i are not (jointly) significant; that is, I find no direct effect of heterogeneity in survival, once I condition on lagged health h_{it-1}. The lack of significance of δ_i has an additional advantage. It implies survival is not another signal for the unknown δ_i. If δ_i had a direct effect on survival, then, as in the case of health, survival would provide individuals with additional information. In that case, by being alive, individuals would learn something more about their heterogeneous slopes, and the Bayesian updating equations (1.2) and (1.3) would not be valid. The results in this exercise, with the lack of significance of δ_i on survival, say survival is not an additional signal for δ_i.

Overall, these results show novel evidence of heterogeneity in health profiles, in particular, in health slopes with age. To study the effects of this heterogeneity on individuals' working decisions, we need to know how much individuals know about their own slope δ_i, which I address next.

1.5 Uncertain health dynamics and beliefs

To study the effect of beliefs on labor-participation decisions of older adults, the main difficulty is that those beliefs are unobserved by the econometrician. The Bayesian learning model implies beliefs are updated over time using health, starting from initial beliefs, $N(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$. Hence, a key issue is the identification of those initial beliefs, in particular, the identification of the parameters b and λ. These parameters are defined by

$$\begin{aligned} b &= \mathbb{E}(\hat{\delta}_{i0} - \delta_i), \\ \lambda^2 &= \frac{\hat{\sigma}_0^2}{Var(\delta_i)}, \end{aligned}$$

and they measure how biased initial beliefs are and how much individuals at age 50 know about their slopes. Because the health process does not reveal slope beliefs, this section proposes the use of survival expectations, available in the HRS. Equation (1.10) implies survival is a health-related process. Therefore, expectations about future survival are related to expectations about future health; thus, they are related to slope beliefs.

1.5.1 Empirical strategy

The exact wording of the HRS question follows:

> [$plive10_{it}$] *What is the percentage chance you will live to be (80, 85, 90, 95 or 100) or more?*

where the reference age depends on the individual's age t at the time of the survey (and wave), and it is approximately 10 years in the future. Let s denote this reference age. Then, this question corresponds to

$$\begin{aligned} plive10_{it} &= \mathbb{P}(S_{is}=1|\Omega_{it}) = \prod_{l=t}^{s-1} \mathbb{P}(S_{il+1}=1|S_{il}=1,\Omega_{it}) \\ &= \prod_{l=t}^{s-1} \mathbb{P}(\gamma h_{il} + \theta_0 + \theta_1(l-1) + \theta_2' x_i + \eta_{il+1} \geq 0|\Omega_{it}), \end{aligned}$$

where we omitted the regressors in the survival equation (besides health) for ease of notation. Applying the equation for health (1.9) recursively, we can write

$$h_{il} = \underbrace{\rho^{l-t}h_{it} + \alpha_i \sum_{k=0}^{l-t-1} \rho^k}_{\text{known under } \Omega_{it}} + \underbrace{\delta_i \sum_{k=0}^{l-t-1} (l-k)\rho^k + \sum_{k=0}^{l-t-1} \rho^k \epsilon_{i(l-k)}}_{\text{unknown under } \Omega_{it}}.$$

From the view point of Ω_{it}, the second term is random, with a normal distribution that depends on $(\hat{\delta}_{it}, \hat{\sigma}_t^2)$ (and the parameters of the model). Because age-t beliefs depend on health history h_i^t and initial beliefs $N(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$, this second term is a function of λ and b.

Therefore, $plive10_{it}$ are complex non-linear functions of slope beliefs,

$$plive10_{it} \quad = \quad plive10_{it}(\alpha_i, h_{it}, \hat{\delta}_{it}, \hat{\sigma}_t^2, x_i) = plive10_{it}(\alpha_i, h_i^t, \hat{\delta}_{i0}, \hat{\sigma}_0^2, x_i). \tag{1.12}$$

The exact function is given in appendix E. Each period, individuals observe their health and update their beliefs regarding their unknown δ_i. This new information allows them to also update their expectations about their future health, and hence their expectations about future survival. Thus, slope beliefs, unobserved by the econometrician, are closely linked to survival beliefs, which are observed by the econometrician. Intuitively, the bias parameter b affects expected health and hence the average survival expectation. Thus, levels of survival expectation identify bias b. Next, I discuss identification of the uncertainty parameter λ.

In what follows, I assume $(\alpha_i, \delta_i, \hat{\delta}_{i0})$ are jointly normally distributed, with $Cov(\alpha_i, \hat{\delta}_{i0}) = Cov(\alpha_i, \delta_i)$ (which is zero according to the results in section 1.4). This assumption implies the information about δ_i contained in α_i is already incorporated in initial beliefs $\hat{\delta}_{i0}$.

Identification using subjective expectations about survival rates (ideal data)

The relation between survival expectations $plive10_{it}$ and the parameters governing beliefs, b and λ, is a complex one. To provide intuition, I start by discussing identification using ideal data, which I do not actually observe. This intuition carries out to the data available in the HRS, which I show numerically next. Let Ω_{it} be the information set of individual i after observing his health up to period t. Thus, $\alpha_i, \hat{\delta}_{it}, \hat{\sigma}_t^2 \in \Omega_{it}$.

Proposition 1.5.1 (Identification of λ) *Let the health and survival processes be given by equations (1.9) and (1.10), and assume individuals are Bayesian learners with prior beliefs*

about δ_i following $N(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$. Consider the subjective expectations about survival rates:

$$bsr_{itr} \equiv \mathbb{P}(S_{ir+1} = 1 | S_{ir} = 1, \Omega_{it}), \quad r \geq t$$

Then, conditional on bsr_{itt+1}, bsr_{itt+2}, and h_{it} (all in Ω_{it})

$$Cov(\Delta\Phi^{-1}(bsr_{it+1t+2}), \Delta h_{it+1}) = C_t Var(\Delta h_{it+1}),$$

where the time-varying constant C_t is increasing in λ.

The proof is in the appendix. The proposition says we can identify λ with enough longitudinal data on subjective expectations about these survival rates and health. The key equation behind this result,

$$\Delta_w \Phi^{-1}(bsr_{it+1r}) = \underbrace{\rho^{r-t-1}(h_{it+1} - \rho h_{it} - \alpha_i - \hat{\delta}_{it}(t+1))}_{\text{due to persistence}} + \underbrace{(\hat{\delta}_{it+1} - \hat{\delta}_{it}) \sum_{k=0}^{r-t-2} (r-k)\rho^k}_{\text{due to learning}},$$

shows individuals update their survival expectations for two reasons. The first reason is that health is a persistent process; thus, any change in health will have future repercussions on health and therefore on survival. Note that if $\rho = 0$, this channel disappears. The second reason is that learning implies a change in future predictions of health and therefore of survival. Note that if $\lambda = 0$, $\hat{\delta}_{it+1} = \hat{\delta}_{it}$, and this channel disappears.

Identification using subjective expectations about survival probabilities (HRS data)

We cannot use the previous result directly, because the HRS does not exactly measure subjective expectations about survival rates. However, Figure 1.7 shows the intuition of

proposition 1.5.1 extends to the available data. It shows the results of a simulation exercise. In the exercise, I first simulate individual-level heterogeneity (α_i, δ_i) and health h_{it}, according to equations (1.9), (1.10) and (1.11). Then, for different values of the uncertainty parameter λ, I simulate initial beliefs $(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$ assuming $b = 0$. I update those beliefs over time and construct $(\hat{\delta}_{it}, \hat{\sigma}_t^2)$ using the Bayesian updating equations and the simulated values of health. Finally, I use these simulated beliefs, to construct survival expectations $plive10_{it}$ according to equation (1.12). In Figure 1.7, the plots depict the uncertainty parameter λ in the x-axis, and a simulated moment in the y-axis. The six plots correspond to the six moments used later for estimation. The top row considers moments in levels, and the bottom row considers moments in differences. The figure clearly shows that, as before, the covariance between changes in health and changes in survival expectations depends on the underlying uncertainty.

Figure 1.7: Simulated moments of $plive10_{it}$ by uncertainty λ in data-generating process

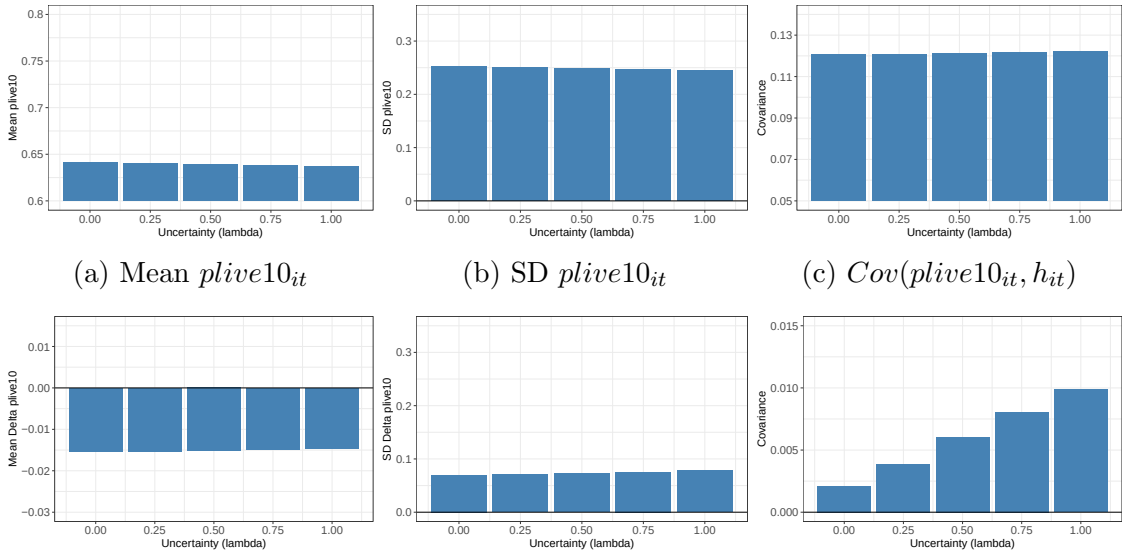

(a) Mean $plive10_{it}$ (b) SD $plive10_{it}$ (c) $Cov(plive10_{it}, h_{it})$

(d) Mean $\Delta plive10_{it}$ (e) SD $\Delta plive10_{it}$ (f) $Cov(\Delta plive10_{it}, \Delta h_{it})$

Note: Moments in simulated data following the structure of the available data in the HRS. The x-axis in each figure shows the value of the uncertainty parameter λ used in the data-generating process. In all cases, the bias parameter b is set to zero.

This model has two simplifying assumptions. First, the model assumes the health process is exogenous, with no choice variable that affects the evolution of health; that is, no investment is purposefully made in the form of health behaviors (e.g., exercising or smoking), and working decisions do not affect health. This assumption is not uncommon in the literature on labor market decisions among older individuals, and it emphasizes changes in health due to aging. By ruling out the possibility of individuals changing their behavior in order to affect their health, the strict exogeneity assumption implies the correlation between changes in health and changes in survival expectations are not confounded by changes in individuals' planned behaviors. Second, the model assumes health is the only or sufficient signal available to individuals. This assumption is partly addressed in the last section of the paper, where I look at another source of information that may shift beliefs.

Under these assumptions, $plive10_{it}$ is a function of initial beliefs $N(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$, heterogeneous intercept α_i, and health history up to t, $(h_{i0}, \dots h_{it})$. Hence, for any value of b and λ, I can use the estimated health process to simulate draws α_i and $\hat{\delta}_{i0}$, and then use those variables to simulate $plive10_{it}$.[22] I estimate the parameters governing initial beliefs, b and λ, by simulated method of moments (SMM). I use six moments, three in levels and three in differences, corresponding to the mean of $plive10_{it}$, its variance, and its covariance with h_{it}.[23] Details of the implementation are given in appendix G.

Subjective survival expectations are measured with substantive error, which is well established in the literature (see, e.g., Manski and Molinari (2010)). Similar to Kleinjans and Van Soest (2014), I allow for non-classical i.i.d. measurement error $\nu_{it} \sim N(\mu_{merror}, \sigma^2_{merror})$,

[22] The distribution of $\hat{\delta}_{i0}$ depends on b and λ. Hence, I first simulate α_i and δ_i conditional on health history $h_{i0}, \dots h_{iT_i}$, and then for a given value of b and λ, I draw $\hat{\delta}_{i0}$ conditional on α_i, δ_i, and h_{i0}.

[23] As described in appendix G, most individuals are first observed in sample at age t_0 older than 50, and I modify the simulation process for them accordingly. Overall, I target these six moments averaged across time for different subgroups of individuals, depending on the age t_0 I first observe them, for a total of 78 moments.

such that the observed survival expectations are given by

$$\widetilde{plive10}_{it} \quad = \quad \max\{\min\{plive10_{it} + \nu_{it}, 1\}, 0\}.$$

Note the measurement error shifts observed survival expectations by μ_{merror} on average. Similarly, the bias in initial beliefs b also shifts observed survival expectations. However, these two biases have different effects over time: the average shift due to measurement error is constant with age, given the i.i.d. assumption, whereas the average shift due to initial bias in beliefs is decreasing with age as individuals observe their health and update their beliefs. Thus, we can separately identify both effects.

1.5.2 Results

The estimation results presented in Table 1.4 show individuals face a sizable amount of uncertainty and a large amount of negative initial bias; that is, individuals believe their health will decay with age at a faster rate than what is actually true on average. In line with previous literature, subjective survival expectations are subject to large amounts of measurement error. Following Manski and Molinari (2010), I also estimate a version including rounding and find similar results. These results are consistent with previous evidence that finds that, on average, older adults up to 65 years old underestimate their chances of survival (Elder (2013), Ludwig and Zimper (2013)). Those papers also find adults 80 years and older overestimate their survival chances. My results similarly show overestimation at those ages, which is explained by measurement error. The fit of the results is shown in Table 1.5. Panel (a) shows the fit of the targeted moments using $plive10_{it}$, whereas panel (b) shows the fit of similar untargeted moments using survival expectations to age 75.[24]

[24]The HRS includes two questions on survival expectations every wave: $plive10_{it}$ asks for a reference age approximately 10 years ahead, and $plive75_{it}$ asks for a reference age equal to 75 years. However, this last question is only asked of individuals 65 or younger, limiting the sample; hence, I use it only here as a check.

Table 1.4: SMM results on prior beliefs

	Symbol	Coefficient	Lower bound	Upper bound
Uncertainty	λ	0.338	0.336	0.340
Bias	b	-0.061	-0.061	-0.060
Mean of measurement error	μ_{merror}	0.121	0.118	0.123
SD of measurement error	σ_{merror}	0.177	0.176	0.177

Note: Prior beliefs about slopes are unobserved $N(\delta_i + b, \lambda^2\sigma_\delta^2)$, depending on the bias b and uncertainty λ parameters, whereas subjective survival expectations $plive10_{it}$ are observed but measured with error. The estimation uses a subsample of 2,000 individuals with eight periods, chosen randomly for computational reasons. Moments are simulated using 20 draws of measurement error and 20 draws of unobserved heterogeneity. The bounds correspond to a 95% confidence interval, constructed using standard errors clustered at the individual level.

With these estimated parameters, I can simulate slope beliefs, which I use in the next section to study their effect on working decisions of older adults.

1.6 Working decisions as functions of beliefs about health

In the life-cycle model of labor participation p_{it} and consumption c_{it} outlined in section 1.2, an individual's dynamic problem is

$$
\begin{aligned}
V_t(\Omega_{it-1}) \quad = \quad & \max_{p_{it},c_{it}} \Bigg\{ \mathbb{E}\Bigg(U(p_{it}, c_{it}, h_{it}, p_{it-1}) \Bigg| \Omega_{it-1} \Bigg) + \\
& \qquad \beta \mathbb{E}\Bigg(S_{it+1} V_{it+1}(\Omega_{it}) + (1 - S_{it+1}) B(a_{it}) \Bigg| \Omega_{it-1}, p_{it}, h_{it} \Bigg) \Bigg\} \\
st. \qquad & \text{budget constraint,} \\
& \text{health (1.9) and survival (1.10) processes,} \\
& \text{and beliefs updating equations (1.2) and (1.3),}
\end{aligned}
$$

where $B(a_{it})$ is the utility perceived by leaving bequest a_{it}. In this problem, the policy rule for labor participation is a function of the state variables in the model. The novelty in this

Table 1.5: Moments' fit

(a) Targeted moments			
	Data moment	SE	Simulated moment
$\mathbb{E}(plive10)$	0.531	(0.00011)	0.538
$\mathbb{E}(plive10^2)$	0.371	(0.00012)	0.357
$\mathbb{E}(plive10 \cdot h)$	2.890	(0.00065)	2.957
$\mathbb{E}(\Delta plive10)$	-0.013	(0.00002)	-0.014
$\mathbb{E}((\Delta plive10)^2)$	0.070	(0.00003)	0.066
$\mathbb{E}(\Delta plive10 \Delta h)$	0.007	(0.00002)	0.007
(b) Other moments (not targeted)			
	Data moment	SE	Simulated moment
$\mathbb{E}(plive75)$	0.702	(0.00017)	0.806
$\mathbb{E}(plive75^2)$	0.556	(0.00021)	0.687
$\mathbb{E}(plive75 \cdot h)$	3.886	(0.00101)	4.469
$\mathbb{E}(\Delta plive75)$	-0.001	(0.00010)	0.018
$\mathbb{E}((\Delta plive75)^2)$	0.054	(0.00008)	0.042
$\mathbb{E}(\Delta plive75 \Delta h)$	0.006	(0.00005)	0.003

Note: Panel (a) uses the same sample used for estimation. Panel (b) uses a subsample of $1,247$ individuals up to 65 years old who are asked $plive75_{it}$ (*the percentage chance you will live to be 75*). Standard errors are clustered at the individual level.

paper is that those state variables include individuals' beliefs about their future health. These beliefs are the result of two key elements: heterogeneity in health dynamics and uncertainty about that heterogeneity. These elements imply beliefs about that heterogeneity -instead of just a common parameter- enter individuals' choices. In this section, I estimate the probability of working as a function of those state variables,

$$\mathbb{P}(p_{it}=1|\Omega_{it-1}) \quad = \quad \mathbb{P}(p_{it}=1|p_{it-1},a_{it-1},w_{it-1},h_{it-1},\hat{\delta}_{it-1},\hat{\sigma}^2_{t-1},\alpha_i). \tag{1.13}$$

By using the results from the previous section, we can simulate all of the state variables, and hence identify their effect on working decisions. Furthermore, by using survival expectations to identify and simulate beliefs, no additional assumption on the relation between beliefs and working decisions has been made. In particular, no restriction is imposed on the sign of the effect of $\hat{\delta}_{it-1}$ on working decisions. If individuals expecting better future health want to work longer, the sign would be positive. This case would arise if the dominant effect were the desire to save more, given the longer life expectancy implied by better health. If individuals expecting worse future health want to work longer, the sign would be negative. This case would arise if the dominant effect were the desire to save more given the higher cost of future health care implied by worse health. Note also that, conditional on states variables in Ω_{it-1}, survival expectations $plive10_{it-1}$ do not play an additional role in working decisions p_{it}.

1.6.1 Probit results on working decisions

To explore the relation between health beliefs and working decisions of older adults, I first estimate equation (1.13) using a probit approach, that is, assuming $\mathbb{P}(p_{it}=1|\Omega_{it-1}) = \Phi\big(\beta'\Omega_{it-1}\big)$. Note some of the input variables are unobserved by the econometrician, namely, heterogeneity in health level α_i and beliefs about slope heterogeneity, $\hat{\delta}_{it}$ and $\hat{\sigma}^2_t$. Conditional

on health history, these unobserved variables depend on individual-level heterogeneity, which is integrated out. See appendix H for details on the likelihood specification.

Table 1.6 presents the results of the probit estimation. As expected, the probability of working decreases with age and increases with better lagged health. Lagged work has a significant effect; the probability of working is larger for individuals who were working the previous period. This result confirms the dynamic aspect of the working decisions. Furthermore, the table shows that beliefs do matter for working decisions, with a positive and significant coefficient for $\hat{\delta}_{it-1}$. This positive sign implies expecting better health, that is, expecting health to deteriorate more slowly with age, is associated with larger probabilities of working.[25] The table also shows survival expectations $plive10_{it-1}$ are significant predictors of the probability of working, but that significance holds only while slope beliefs are not accounted for. This result is consistent with survival expectations reflecting individuals' beliefs about slope heterogeneity. Thus, once those beliefs are considered, survival expectations do not provide additional information.

Though interesting, these results assume a linear index for the probability of working, which is a strong assumption and is not justified by assumptions on the fundamentals of the model. Hence, the results may be inconsistent with the policy rule derived from the economic model. Thus, in what follows, I flexibly estimate the probability of working, without imposing this index linearity. I achieve that flexibility by using instead a neural-network approach.

1.6.2 Neural-network approach

Neural networks provide flexible tools for estimation (Goodfellow et al. (2016)). They are universal approximators, because they are capable of approximating any measurable function

[25] The assumptions of the learning model imply the posterior variance $\hat{\sigma}_t^2$ is constant across individuals of the same age t. Given that age is also a relevant determinant of working decisions, I don't have enough variation to disentangle these two effects separately; any results would be based on functional-form assumptions alone. Therefore, I focus instead on interpreting the effects of the posterior mean.

Table 1.6: Probit results on probability of working

		(1)		(2)		(3)	
		Coefficient	SE	Coefficient	SE	Coefficient	SE
age	$t-1$	-0.20***	(0.016)	-0.08***	(0.003)	-0.19***	(0.016)
lagged work	p_{it-1}	2.03***	(0.018)	2.03***	(0.019)	2.03***	(0.019)
lagged health	h_{it-1}	0.17***	(0.024)	0.26***	(0.033)	0.18***	(0.046)
heterogeneous intercept	α_i	0.24***	(0.036)	0.07	(0.046)	0.24***	(0.075)
beliefs mean	$\hat{\delta}_{it-1}$	1.93***	(0.249)			1.90***	(0.499)
beliefs var	$\hat{\sigma}^2_{t-1}/\sigma^2_\delta$	-13.85***	(2.048)			-13.33***	(2.102)
survival expectations	$plive10_{it-1}$			0.11***	(0.031)	0.01	(0.043)
Controls	other vars Ω_{it-1}	Yes		Yes		Yes	
N individuals		14,969		14,718		14,718	
N observations		58,040		55,592		55,592	

Note: Results of estimating equation (1.13) using a probit approach. Standard errors are clustered at the individual level.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

to any desired degree of accuracy (Hornik et al. (1989)). In the case of a binary outcome, and under some particular specifications, a neural network corresponds to a maximum likelihood estimation with logistic errors, where the probability of success is a complex non-linear index of the inputs. As mentioned by Farrell et al. (2021), we can think of neural networks as a type of non-parametric or sieve estimation whereby the basis functions are learned from the data, hence allowing for greater flexibility.

In this case, I also need to account for the fact that some of the input variables are unobserved by the econometrician. These unobserved variables are slope beliefs $(\hat{\delta}_{it-1}, \hat{\sigma}^2_t)$ and heterogeneous health levels α_i. Though they are time-varying variables, they can be written as functions of time-invariant unobserved variables $(\hat{\delta}_{i0}, \alpha_i)$ and the observed health path $(h_{i1}, \dots h_{iT_i})$ of each individual.[26] Thus, following a standard likelihood approach, I want to maximize the log of the likelihood integrating out this time-invariant unobserved heterogeneity. To do so, I follow the insight of EM-type algorithms (Dempster et al. (1977)).

[26]This relationship depends also on the parameters of the health process (ρ, σ^2_ϵ) and the parameters of beliefs (b and λ), but it does not depend on the parameters defining the relation between working decisions and state variables.

Let θ be the parameters governing an outcome variable, in this case, working decisions. When underlying heterogeneity exists, we estimate θ by maximizing a likelihood that integrates out that heterogeneity. In this context, EM-type algorithms provide us with two key insights. First, the parameter θ that maximizes the integrated log-likelihood also maximizes an alternative specification using the posterior distribution given the outcome variable. Formally, let η_i denote the vector of unobserved heterogeneity, $f(\eta_i)$ its prior distribution, and $f(\eta_i|p_i;\theta)$ its posterior distribution given the outcomes p_i. The first insight of EM-type algorithms is to note that

$$arg\max_{\theta} log \int \mathbb{P}(p_i|\eta_i;\theta)f(\eta_i)d\eta_i \quad = \quad arg\max_{\theta} \int log(\mathbb{P}(p_i|\eta_i;\theta))f(\eta_i|p_i;\theta)d\eta_i. \quad (1.14)$$

The expression on the right-hand side is simpler to use. However, because this posterior distribution depends on θ, it is unknown. Thus, the second insight of EM-type algorithms is to solve the problem for θ iteratively: in iteration k, the E step uses θ_{k-1} to update the posterior distribution of the heterogeneity, and the M step estimates θ_k by maximizing the right-hand side of equation (1.14), using that posterior.

I use this same iterative logic as a convenient implementation for maximizing the integrated likelihood under a neural-network approach. In this case, the E step is done by Markov chain Monte Carlo (MCMC) and provides draws from the posterior distribution of $(\alpha_i, \hat{\delta}_{i0})$ given working decisions p_i.[27] Those draws, along with individuals' health histories, are used to simulate the inputs $(\hat{\delta}_{it}, \hat{\sigma}^2_t, \alpha_i)$ and to expand the data. Then, the M step estimates θ by using a neural network on the expanded data.[28] I start this iterative process at an M step using an incomplete posterior: the distribution of $(\alpha_i, \hat{\delta}_{i0})$ conditional on the

[27] MCMC uses the likelihood of p_i given $(\alpha_i, \hat{\delta}_{i0})$ from the previous-iteration M step and the prior distribution of $(\alpha_i, \hat{\delta}_{i0})$.

[28] The standard EM algorithm is known to converge, as the likelihood increases in each step of the sequence. This convergence does not hold in this case, given the lack of uniqueness of the solution. Therefore, the approach is not aimed at getting at the unique solution, but as a convenient implementation.

health history $(h_{i1}, \ldots h_{iT_i})$ and the history of survival expectations $(plive10_{i1}, \ldots plive10_{iT_i})$. This distribution is incomplete because it does not condition on the working decisions, but it does already include the heterogeneity information contained in the health and expectations variables.

1.6.3 Neural-network results on working decisions

Following this strategy,[29] I estimate the probability of working conditional on the state variables Ω_{it-1}. This set includes past participation p_{it-1}, past health h_{it-1}, heterogeneous health levels α_i, and slope beliefs $\hat{\delta}_{it}, \hat{\sigma}_t^2$, which are the main interest in this paper. It also includes more traditional variables, listed in Table 1.2, including demographic variables, income, wealth, health insurance, and job characteristics. I restrict the analysis to a sample of individuals who are attached to the labor market, defined as individuals with at least 20 years of working experience. The loss and fit of the model is given in appendix I.

(1) Beliefs play a role in the participation decisions of older adults, with positive average marginal effects that are similar in orders of magnitude to the average marginal effects of health and assets.

Table 1.7 presents the effects of a marginal change in expected beliefs $\hat{\delta}_{it-1}$, health h_{it-1}, and assets a_{it-1}, respectively, on the probability of working, conditional on age and past participation p_{it-1}, averaged across individuals. The table shows that even though the effects are of different magnitudes and signs, they are similar in orders of magnitude. The same result holds in Figures 1.8 and 1.9, which show the marginal effects of beliefs $\hat{\delta}_{it-1}$ by deciles

[29] The results in this section come after running the iterative approach 5 times. These results are qualitatively similar to the results using the incomplete prior. This similarity is not unexpected given that the incomplete prior already incorporates the information on health and survival beliefs.

of health and beliefs for adults aged 52-59 and 66-75, respectively.

(2) For individuals in their 50s who are not working, an interaction exists between beliefs and health in their future participation decisions.

This result can be seen in Figure 1.8. The figure shows health has larger marginal effect on working probabilities for individuals with better beliefs, that is, for individuals who believe their health will deteriorate relatively slowly. A similar pattern is observed for the marginal effect of beliefs themselves. These results suggest adjustment costs of going back to work are important for the decisions of this group. These adjustment costs could be due to difficulties in finding jobs or in adapting to new work environments. The framework and data-driven approach used in this paper have the advantage of letting the data suggest mechanisms that may be otherwise overlooked. Overlooking important mechanisms is a source of misspecification in structural models. Hence, the approach in this paper complements structural models, by providing a flexible way to identify patterns in the data that suggest mechanisms to incorporate in such models.

(3) The total effect of a health shock ϵ_{it-1} on working decisions goes mostly through the persistence channel, with negligible effects through the information channel.

This result is shown in Table 1.7, which includes the decomposition of the effects of a health shock into a persistence channel and an information channel, according to equation (1.8). The persistence channel refers to the effect that a health shock ϵ_{it-1} has on h_{it} through h_{it-1} and the persistence of the health process. The information channel refers to the effect that a health shock ϵ_{it-1} has on beliefs $\hat{\delta}_{it-1}$, as individuals interpret h_{it-1} (and hence this health shock) as a health signal. According to equation (1.8), the total effect of a health

Table 1.7: Average marginal effects on the probability of working and decomposition of the effects of a health shock

Age	Average marginal effects				Decomposition of a health shock		
	Health h_{it-1}	Beliefs $\hat{\delta}_{it-1}$	Assets a_{1it-1}	Assets a_{2it-1}	Factor	Persistence channel	Information channel
$p_{it-1} = 0$							
52	0.056	0.028	-0.021	-0.043	0.003	1.00	0.00
54	0.049	0.024	-0.019	-0.039	0.006	1.00	0.00
56	0.043	0.021	-0.017	-0.034	0.009	1.00	0.00
58	0.038	0.018	-0.015	-0.030	0.011	0.99	0.01
60	0.033	0.016	-0.013	-0.027	0.013	0.99	0.01
62	0.028	0.013	-0.011	-0.023	0.014	0.99	0.01
64	0.022	0.010	-0.008	-0.018	0.015	0.99	0.01
66	0.019	0.009	-0.007	-0.015	0.015	0.99	0.01
68	0.015	0.007	-0.006	-0.013	0.014	0.99	0.01
70	0.013	0.006	-0.005	-0.010	0.014	0.99	0.01
72	0.010	0.004	-0.004	-0.008	0.013	0.99	0.01
74	0.008	0.003	-0.004	-0.007	0.012	1.00	0.00
$p_{it-1} = 1$							
52	0.010	0.011	0.000	-0.005	0.003	1.00	0.00
54	0.011	0.012	0.000	-0.006	0.006	0.99	0.01
56	0.013	0.013	0.000	-0.007	0.009	0.99	0.01
58	0.015	0.015	0.000	-0.008	0.011	0.99	0.01
60	0.017	0.017	0.000	-0.008	0.013	0.99	0.01
62	0.018	0.018	0.000	-0.009	0.014	0.99	0.01
64	0.020	0.020	0.001	-0.009	0.015	0.99	0.01
66	0.021	0.021	0.003	-0.009	0.015	0.99	0.01
68	0.021	0.021	0.004	-0.009	0.014	0.99	0.01
70	0.021	0.022	0.005	-0.009	0.014	0.99	0.01
72	0.022	0.022	0.005	-0.009	0.013	0.99	0.01
74	0.022	0.022	0.006	-0.009	0.012	0.99	0.01

Note: Assets a_1 are total assets excluding assets on retirement accounts, which are considered separately on variable a_2. The columns on persistence and information channels correspond to the terms in equation (1.8), expressed as a proportion of the total partial effect.

shock is a weighted sum of the effects through these two channels. Note the small values on the column Factor in Table 1.7, which imply a health shock has only a small effect on

Figure 1.8: Average marginal effect of expected beliefs $\hat{\delta}_{it-1}$, and health h_{it-1} on the probability of working p_{it} for adults in their 50s

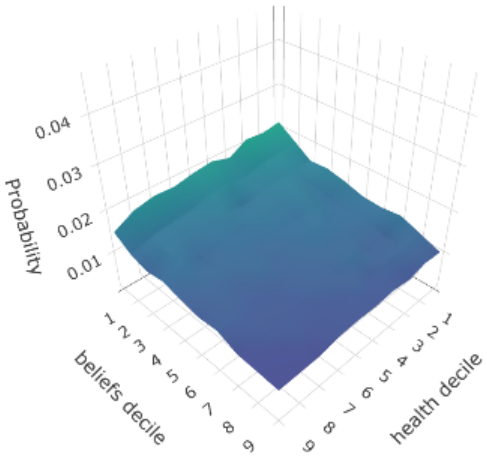

(a) Marginal change in $\hat{\delta}_{it-1}$ conditional on $p_{it-1} = 1$

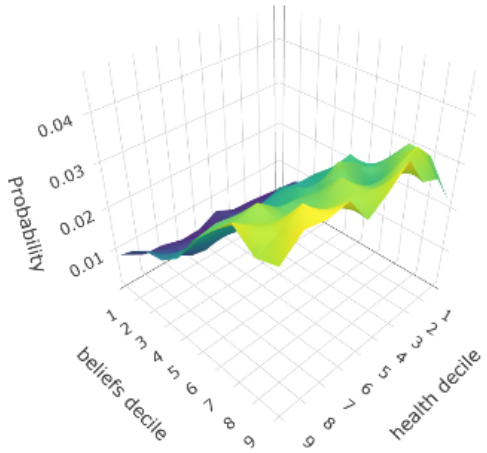

(b) Marginal change in $\hat{\delta}_{it-1}$ conditional on $p_{it-1} = 0$

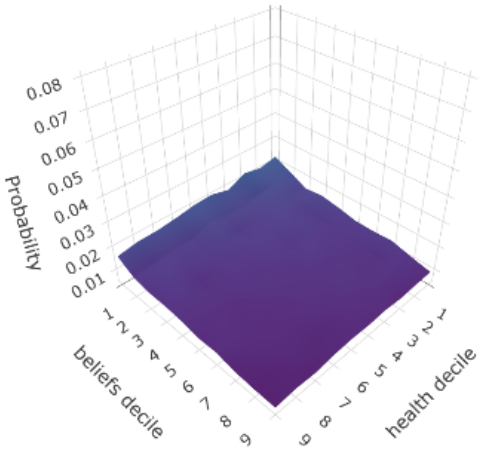

(c) Marginal change in h_{it-1} conditional on $p_{it-1} = 1$

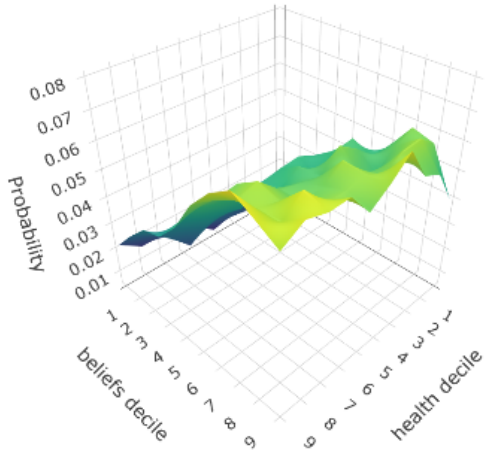

(d) Marginal change in h_{it-1} conditional on $p_{it-1} = 0$

Note: Each row corresponds to the average marginal effects with respect to $\hat{\delta}_{it-1}$ and h_{it-1}, respectively. The left column conditions on individuals who were working, $p_{it-1} = 1$, and the right column conditions on individuals who were not working, p_{it-1}=0, in the previous period. In each plot, the x- and y-axis correspond to deciles of health h_{it-1} and expected beliefs $\hat{\delta}_{it-1}$ for the corresponding subsample of the plot. The z-axis corresponds to the work response (probability). Note the range of the z-axis changes in each row.

beliefs $\hat{\delta}_{it-1}$ and therefore only a small effect through the information channel. This result highlights that even though individuals are uncertain and biased, to significantly affect their

Figure 1.9: Average marginal effect of expected beliefs $\hat{\delta}_{it-1}$, and health h_{it-1} on the probability of working p_{it} for adults between 66 and 75 years old

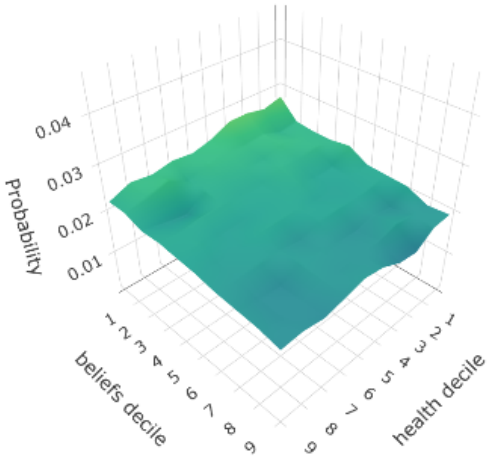

(a) Marginal change in $\hat{\delta}_{it-1}$ conditional on $p_{it-1} = 1$

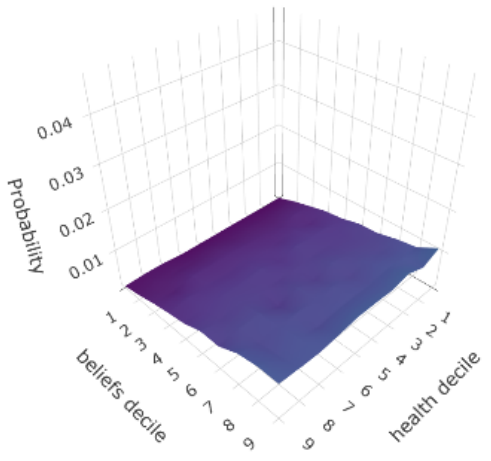

(b) Marginal change in $\hat{\delta}_{it-1}$ conditional on $p_{it-1} = 0$

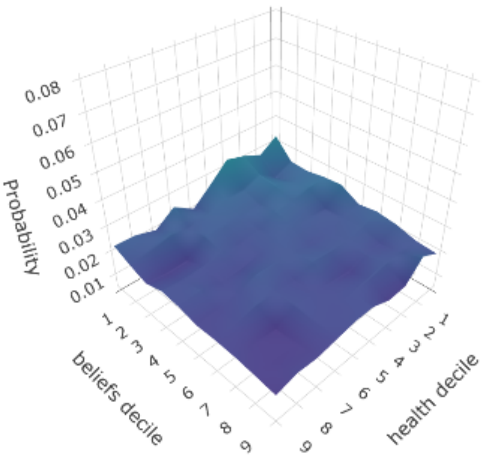

(c) Marginal change in h_{it-1} conditional on $p_{it-1} = 1$

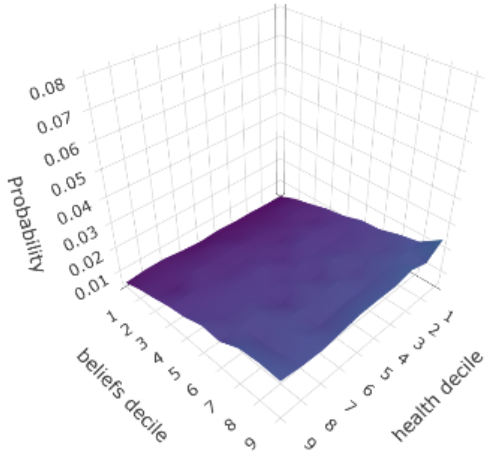

(d) Marginal change in h_{it-1} conditional on $p_{it-1} = 0$

Note: Each row corresponds to the average marginal effects with respect to $\hat{\delta}_{it-1}$ and h_{it-1}, respectively. The left column conditions on individuals who were working, $p_{it-1} = 1$, and the right column conditions on individuals who were not working, p_{it-1}=0, in the previous period. In each plot, the x- and y-axis correspond to deciles of health h_{it-1} and expected beliefs $\hat{\delta}_{it-1}$ for the corresponding subsample of the plot. The z-axis corresponds to the work response (probability). Note the range of the z-axis changes in each row.

decisions, we need large enough signals. Section 1.8 looks at one possible such policy: health information regarding blood glucose and cholesterol levels.

1.7 Reducing bias in initial beliefs

In this section, I study how labor participation would change if we could eliminate bias in initial beliefs. In particular, I look at two questions:

1. How much would labor participation change if initial beliefs were unbiased at the population level, that is, $\mathbb{E}(\hat{\delta}_{i0}) = \mathbb{E}(\delta_i)$?

2. How much would labor participation change if we could reduce each individual's bias in half, by closing the distance between $\hat{\delta}_{i0}$ and δ_i?

To look at these questions, I use an impulse-response-function approach. That is, I simulate working decisions under a sample's baseline scenario, and compare those predictions against the predictions simulated under each of these two potential changes in initial beliefs. The figures in this section present the response in terms of labor-participation decisions by age, given a change in initial beliefs. Over time, this change in initial beliefs translates into changes in posterior beliefs, labor-participation decisions, and decisions regarding assets and health insurance. The effects on these last two variables were also predicted using a neural-network approach. Note these exercises assume no other variable change in response to the change in initial beliefs or to the subsequent changes in participation, assets, or health insurance. Therefore, the exercises presented here are not exactly counterfactual analyses, but are interesting exercises as long as we are capturing the main choices.[30]

(1) Eliminating the bias in prior beliefs b would increase participation by more than 2 percentage points around the formal retirement age (66-67).

[30]The results presented in this section use the incomplete prior of the unobserved heterogeneity. As discussed in the previous section, this prior already accounts for the information in the health and survival expectations variables, and incorporating the additional information has only a minor effect.

Figure 1.10 shows the average change in the probability of working after eliminating the initial bias in prior beliefs, b. Note this effect has an inverted-U shape. In the early 50s, the effect is small given that individuals are still mostly working. But as people start to retire, the new beliefs imply larger probabilities of working that do not vanish completely over time and remain above 2 percentage points for individuals in their early seventies. Note that, in this sample, the average probability of working prior to the change in beliefs is 34% at age 66 and 17% at age 78; hence, the increment in the figure is not trivial. Furthermore, because this effect results from eliminating a misconception at the population level, it is an easier target policy that could be addressed by information campaigns, without the need to provide individual-specific information.

Figure 1.10: Impulse-response function to a shift in prior beliefs eliminating overall bias b

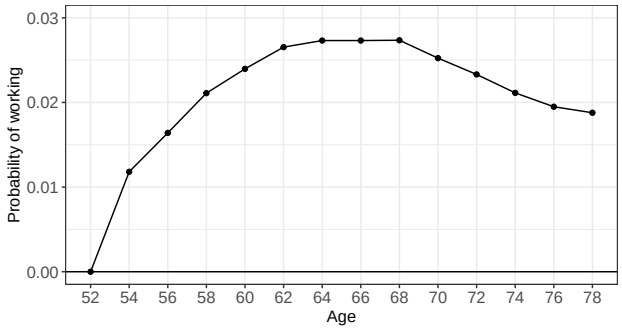

Note: Impulse-response function using the subsample of individuals used in estimation that are observed at 52 years old, corresponding to 1,184 individuals.

(2) Reducing the initial bias of each individual by half has a heterogeneous effect, with larger gains in the probability of working for individuals who are initially more biased.

Figure 1.11 shows this results, distinguishing by quartile of initial bias, $\hat{\delta}_{i0} - \delta_i$. Given the overall initial bias $b < 0$, most individuals are initially downward biased. Thus, reducing bias

in half per each individual means increasing initial beliefs for most of them, which translates into the effects being positive, as shown in the figure.

Figure 1.11: Impulse-response function to reducing individuals' initial bias by half

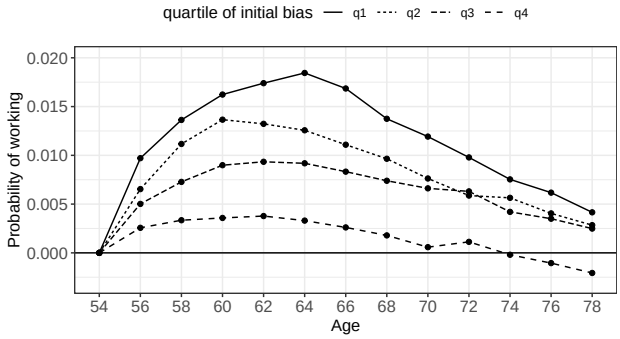

Note: Impulse-response function using the subsample of individuals used in estimation that are observed at 52 years old, corresponding to 1,184 individuals.

As a reference, using a structural model, French and Jones (2011) find raising the Medicare age from 65 to 67 leads individuals to work an additional 0.074 years over ages 60 to 69, whereas eliminating two years' worth of Social Security benefits increases time spent in the work force by 0.076 years.

1.8 An information experiment: Blood-based biomarkers as signals of health

1.8.1 Setup

The results on working decisions of older adults show beliefs matter, and expecting health to deteriorate more slowly is associated with larger probabilities of working. Furthermore, beliefs are initially biased, and eliminating that bias has non-trivial effects. Information campaigns providing better information can be a way to eliminate that bias. In this section, I

exploit a feature of the HRS and study the effects of an information shock on individuals' decisions: information on blood-based biomarkers. In 2006, the HRS introduced the collection of a blood sample for measuring biomarkers. With the blood sample, three biomarkers are measured and individuals are informed of their levels: HDL cholesterol (known as the good cholesterol), total cholesterol, and blood glucose hbA1c. The results are provided around a month after the survey has ended[31] (see Edwards (2018) for more details). These biomarkers are also included in other health surveys, including the REasons for Geographic and Racial Differences in Stroke study (REGARDS) and the National Health and Nutrition Examination Survey (NHANES), where the information is also provided to individuals. Studies using those biomarkers have found that new diagnoses through the surveys increase the number of doctor visits for Medicare beneficiaries (Myerson et al. (2018)), but they increase the fraction of patients with low uptake of ex-post medical treatment (Myerson et al. (2017)).

A key aspect in the introduction of these measures in the HRS is that, to control costs associated with their collection, the HRS randomly split the sample into two halves, and in each wave, the HRS collects these biomarkers in only one of those halves. Hence, this collection scheme provides us with an information experiment, that is, with exogenous variation in who receives this additional information. Note, however, that setting an experiment was not the intended goal of the HRS, and as such, this experiment is not ideal. An ideal experiment would include a control arm of individuals who get their blood taken but are not informed of their results. Still, the HRS collection scheme of biomarkers does provide us with exogenous variation that I use in this section. Another advantage of looking at this additional source of information is that it allows me to relax the assumption of health as the only (or sufficient) signal[32] and to use additional sources of variation when estimating the

[31] Two other biomarkers are measured: C-reactive protein (CRP), a general marker of systemic inflammation, and Cystatin C, an indicator of kidney functioning. However, individuals are not informed of their levels on these two biomarkers; hence, these results do not provide additional information to individuals.

[32] The signal analyzed here is provided exogenously to individuals. Hence, this paper does not address endogenous acquisition of information, which is left for future work.

effects of beliefs on working decisions of older adults.[33] However, the counterpart of using this experiment is that the information provided in the experiment is very small.

1.8.2 Reduced-form approach

I start by estimating the overall effect of receiving this information on individuals' survival expectations and working decisions. To that end, I use the experiment introduced with the biomarker collection in 2006 (wave 8), when the sample was randomly divided into two. To be able to generate this information, the experiment also introduces a difference in interview mode between the two groups, because collecting the blood sample requires an in-person interview.[34] The interview mode could have an effect on individuals' answers, in particular, on questions regarding opinions and expectations. Though potentially problematic, the timing of the information provision allows me to separately identify the interview-mode effect from the information effect of the biomarker results, because that information is only provided to individuals after the fieldwork has ended. Hence, individuals do not have the information in the wave when the blood is collected, but in the following wave.

Figure 1.12 presents the structure of the biomarker collection and the information experiment, and it helps us visualize the identification strategy. Note, first, that a difference-in-differences analysis using waves 7 and 8 returns the interview-mode effect. Second, a difference-in-differences analysis using waves 7 and 9 returns the interview-mode effect (with the opposite sign) plus the information effect of receiving the additional signal. Hence, we can identify the information effect by adding these two terms. Under the parallel-trends

[33] In inferring slope beliefs and using them to study their effect on labor-participation decisions in the previous sections, I only use cross-sectional variation given by differences in initial beliefs, conditional on health and survival-expectations histories.

[34] The HRS survey is usually conducted by phone, except for first interviews of new cohorts, people who request in-person interviews, and individuals residing in nursing homes. A shift to in-person interviews in 2004 also occurred in an attempt by the HRS to increase individuals' consent to link their survey responses with administrative data. These differences in interview mode are unimportant for the analysis as long as they are applied in the same way across the two groups.

assumption, the same idea holds if we construct these terms using wave 5 instead of wave 7.

Figure 1.12: Timing of the biomarker collection and information experiment

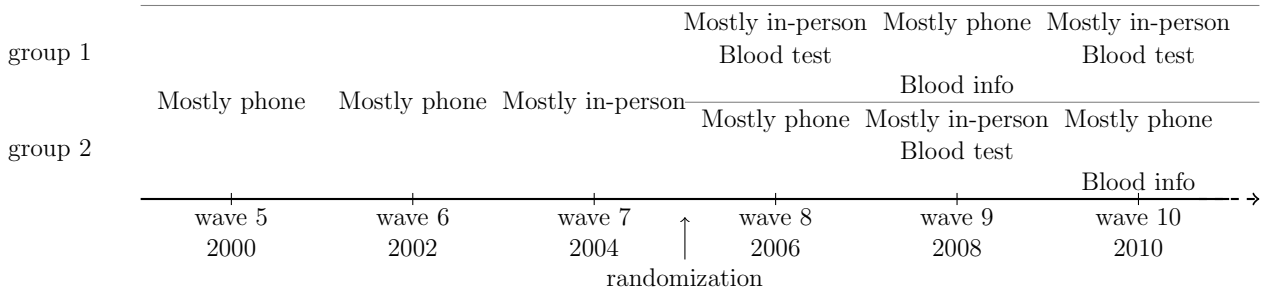

Therefore, I estimate the following equation:

$$y_{iw} = \beta_0 + \beta_1 d_{g_i} + \beta_{2w} d_w + \beta_{3w} d_{g_i} \cdot d_w + \epsilon_{iw}, \tag{1.15}$$

where i denotes an individual and w denotes a wave. I use w instead of t, because in this paper, t denotes age. I consider two dependent variables separately, survival expectations $plive10_{iw}$ and a binary of work p_{iw}. I estimate these equations using a balanced sample of individuals observed from waves 5 to 9.[35] d_{g_i} is a dummy for the group of individuals set for blood collection in wave 8 (group 1 in Figure 1.12, with group 2 as the reference category), and d_w are dummies for waves 6 to 9 (wave 5 is the reference category). Hence, the interview-mode effect is given by β_{3w_8}, and the information effect of the signal is given by $\beta_{3w_8} + \beta_{3w_9}$, where the interview-mode effects in each group cancel each other out. Parallel trends (before randomization) hold if $\beta_{3w_6} = \beta_{3w_7} = 0$, and randomization in the selection of the two groups implies $\beta_1 = 0$.

Table 1.8 presents the estimation results of equation (1.15) for both $plive10_{iw}$ and p_{iw}. When looking at the results for survival expectation, $plive10_{iw}$, the table shows the two

[35] I use only up to wave 9, because from wave 10 onward, the groups are no longer comparable, given that they have been provided information with different timing.

groups are similar and that pre-trends are parallel. The table also shows a positive and significant interview-mode effect of 1.77 percentage points and a similar but insignificant information effect of 1.36 percentage points. Though insignificant, this positive sign is aligned with what we already know about beliefs: on average, individuals' beliefs about health and survival are downward biased. Therefore, providing more information moves those expectations up. When looking at the results for working decisions, p_{iw}, the two groups are similar to begin with and have parallel pre-trends, but we find no significant effect of interview mode[36] or information. Overall, these results suggest the signal is not large enough to have a significant effect on expectations and decisions.

Table 1.9 presents the results separately by education level. It shows that for adults with a college degree, both the interview-mode and information effects are larger and significant when looking at survival expectations. For adults with less than a college degree, only the interview-mode effect is marginally significant (at 12%). When looking at working decisions, no significant effects —interview-mode or information effects— for either group are seen. These differences by education level suggest the ability to process the information matters, with more educated adults internalizing the provided information better. The effect on their working decisions is also larger though still not significant.[37]

Appendix J further decomposes group 1 into adults who receive a bad biomarker result versus those who do not. However, because we cannot make the same distinction in group 2,[38] we cannot identify information effects by the type of signal received (good or bad biomarker

[36]The lack of an interview-mode effect on working decisions is expected, given the more objective nature of working outcomes versus survival expectations.

[37]I run a similar regression with the number of doctor visits since the last interview as a dependent variable and find no effects (results not shown), neither interview-mode nor information effects, for either group. This result suggests the difference in survival expectations between these two groups is not explained by a different number of doctor visits. However, more educated individuals may still be better able to incorporate the new information with the help of their physicians, even if the number of doctor visits remains the same.

[38]One possibility would be to use the biomarker results in wave 9 to attempt the same distinction for group 2. However, an analysis using repeated biomarker results from future waves shows these results change over time, introducing noise when using results from wave 9 to assign wave 8 status for the second group.

Table 1.8: Information and interview-mode effects of biomarker experiment

		Survival expectation ($plive10_{iw}$)	Work decision (p_{iw})
Group 1	β_1	-0.47	0.00
Wave 6	β_{2w_6}	-1.42***	-0.07***
Wave 7	β_{2w_7}	-1.50***	-0.12***
Wave 8	β_{2w_8}	-6.41***	-0.16***
Wave 9	β_{2w_9}	-3.57***	-0.20***
Group 1, wave 6	β_{3w_6}	0.28	0.01
Group 1, wave 7	β_{3w_7}	-0.27	0.01
Group 1, wave 8	β_{3w_8} (a)	1.77**	0.01
Group 1, wave 9	β_{3w_9} (b)	-0.42	0.01
Constant	β_0	53.97***	0.49***
Observations		41,930	41,923
R-squared		0.004	0.021
Interview mode effect (a)		1.77**	0.01
Information effect (a)+(b)		1.36	0.02

Note: Results of estimating equation (1.15). The sample consists of $N = 8,386$ individuals with non-proxy interviews who are at least 50 years old in wave 8 and who give a valid answer to $plive10_{iw}$ every wave between waves 5 and 9. Seven of these observations do not have information on p_{iw}. Standard errors are clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.9: Information and interview-mode effects by education level

		Survival expectation ($plive10_{iw}$)		Work decision (p_{iw})	
		Less than college	College	Less than college	College
Group 1	β_1	-0.24	-1.38	0.01	-0.01
Wave 6	β_{2w_6}	-1.21**	-2.09**	-0.07***	-0.09***
Wave 7	β_{2w_7}	-1.44***	-1.72**	-0.12***	-0.12***
Wave 8	β_{2w_8}	-6.12***	-7.37***	-0.16***	-0.19***
Wave 9	β_{2w_9}	-3.22***	-4.70***	-0.20***	-0.22***
Group 1, wave 6	β_{3w_6}	-0.06	1.37	0.00	0.02
Group 1, wave 7	β_{3w_7}	-0.24	-0.33	0.01	0.01
Group 1, wave 8	β_{3w_8} (a)	1.29	3.31***	0.00	0.03
Group 1, wave 9	β_{3w_9} (b)	-1.12	1.82	0.01	0.00
Constant	β_0	52.42***	58.96***	0.45***	0.61***
Observations		31,815	10,115	31,810	10,113
R-squared		0.004	0.005	0.021	0.022
Interview mode effect (a)		1.29	3.31***	0.00	0.03
Information effect (a)+(b)		1.65	5.12**	0.01	0.04

Note: Estimation results are from equation (1.15). The sample consists of $N = 8,386$ individuals with non-proxy interviews who are at least 50 years old in wave 8 and who give a valid answer to $plive10_{iw}$ every wave between waves 5 and 9. Seven of these observations do not have information on p_{iw}. Standard errors are clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1

results). Still, this analysis is interesting because it shows older adults who receive bad results have lower survival expectations to begin with, suggesting they already knew at least some of this information. Consistently, by wave 7, people who later receive bad biomarker results also work less on average than those who receive good results.

1.8.3 Model-based approach

In this section, I use the learning model to re-assess the information experiment. I estimate survival expectations allowing for the biomarker information to be a second signal for health profiles. For these biomarkers to be a valid signal, being correlated with health is not enough; they must be correlated with δ_i. The appendix shows they are indeed. It presents the results of estimating an equation for health, similar to the equation of section 1.4, allowing for the distribution of the heterogeneity to depend on blood-glucose and cholesterol levels. The results show both the heterogeneous intercepts α_i and heterogeneous slopes δ_i are correlated with these particular biomarkers.

Hence, some individuals have two signals of δ_i: health h_{it} and biomarker results l_{it}. Let l_{it} be the blood-glucose level of individual i at age t,[39] and let the two signals be given by

$$\begin{aligned} h_{it} &= \rho h_{it-1} + \alpha_i + \delta_i \cdot t + \epsilon_{it} \\ l_{it} &= \tau_0 + \tau_1 h_{it-1} + \tau_2 \alpha_i + \tau_3 \delta_i \cdot t + \tau_4 \cdot t + \tau_5 \cdot x_i + \omega_{it}, \end{aligned}$$

where ω_{it} are i.i.d. and independent of health shocks ϵ_{it}. Bayes' rule implies the updating equations for the posterior mean and variance of beliefs are given by

$$\frac{\hat{\delta}_{it}}{\hat{\sigma}_t^2} = \frac{\hat{\delta}_{it-1}}{\hat{\sigma}_{t-1}^2} + \frac{(h_{it} - \rho h_{it-1} - \alpha_i)t}{\sigma_\epsilon^2} + \frac{(l_{it} - \tau_0 - \tau_1 h_{it-1} - \tau_2 \alpha_i - \tau_4 t - \tau_5 x_i)t \cdot \tau_3}{\sigma_\omega^2} \quad (1.16)$$

$$\frac{1}{\hat{\sigma}_t^2} = \frac{1}{\hat{\sigma}_{t-1}^2} + \frac{t^2}{\sigma_\epsilon^2} + \frac{\tau_3^2 t^2}{\sigma_\omega^2}. \quad (1.17)$$

[39] I focus on blood glucose because it is the biomarker more consistently related to slopes δ_i.

Equation (1.17) shows the posterior variance includes the information provided by health and by the biomarker results. As long as the biomarkers provide information about δ_i, that is, as long as $\tau_3 \neq 0$, having this additional signal increases the precision of posterior beliefs. Furthermore, the overall gain in precision depends on both τ_3 and σ_ω, because they determine the biomarkers' signal-to-noise ratio. Equation (1.16) shows the posterior mean of δ_i is a weighted average of the prior at that period, the signal provided by health, and the signal provided by the biomarker information. The weights depend on how uncertain individuals are to begin with, and on the precision of the information provided by each signal. Thus, to predict beliefs $\hat{\delta}_{it}$ and survival expectations $plive10_{it}$, a key issue is to measure the precision of the additional signal.

To measure that precision, I use the biomarker experiment in the spirit of Todd and Wolpin (2006). That is, to predict beliefs of group 1, I use parameters estimated using data from future waves of group 2. Specifically, I want to predict beliefs when biomarker information was first introduced between waves 8 and 9. By wave 8, only group 1 had their blood collected, and by wave 9, only group 1 had their biomarker information available as a second signal. I estimate the parameters governing the precision of that second signal, using individuals from group 2. They had their blood collected for the first time in wave 9, and they received their biomarker information before wave 10. Hence, I use their biomarker information and their survival expectations in waves 9 and 10 to estimate the parameters of the additional signal using simulated method of moments.[40] Using those parameters, I predict beliefs and survival expectations for group 1 in wave 9 (no second signal was available yet in wave 8). The randomness in the selection of the groups implies the parameters recovered by looking at group 2 must also represent the parameters governing the biomarker signal for group 1.

[40] In an alternative version, I use a maximum likelihood approach to jointly estimate health and biomarker results as a function of slope heterogeneity δ_i. I then use those parameters to predict slope beliefs and survival expectations. Under this alternative approach, I obtain qualitatively the same results as the ones from using SMM.

Table 1.10: Predicted survival expectations in a model with health and blood glucose as signals

	Number of observations	Predicted survival expectations		
		wave 8	wave 9	wave 9 - wave 8
Control (group 2)	4,852	45.8	45.4	-0.3
Treated (group 1)	5,357	44.8	44.9	0.1
Treated with bad blood glucose	552	39.1	38.5	-0.5
Treated with good blood glucose	3,649	46.0	46.3	0.3
Treated no blood glucose	1,156	43.8	43.7	-0.2

Note: The sample consists of $N = 10,209$ individuals with non-proxy interviews who are at least 50 years old in wave 8 and who provide a valid answer to $plive10_{iw}$ in waves 8 and 9. Survival expectations are predicted from a model with one signal for the control group (health) and two signals for the treated group (health and blood-glucose results). These two signals are assumed to be independent conditional on individual heterogeneity. The parameters determining the strength of blood glucose as a signal of δ_i come from an estimation using future values of the control group (waves 9 and 10)

The overall results on survival expectations by group are presented in Table 1.10. According to these results, the learning model suggests that by having the additional signal on health, group 1 increases their survival expectations between waves 8 and 9 by 0.4 percentages points more than the control group. This change in survival expectations is positive but negligible, consistent with the results in Table 1.9. Thus, though a valid signal for health profiles, biomarker results are only a small signal, and not enough to shift beliefs and significantly affect decisions.

1.9 Conclusion

This paper documents individual-level heterogeneity in health dynamics among older adults and studies how individuals' beliefs about their own health dynamics affect their working decisions. In the first part of the paper, I show evidence that health dynamics are

indeed heterogeneous. In particular, I show health is heterogeneous in the way it changes with age and that this heterogeneity helps explain the increasing variance of health with age, a pattern observed in the population but mostly ignored by traditional models of health. Motivated by this evidence, I turn to the question of how much individuals know about their own health profiles. I develop a Bayesian learning model in which individuals have beliefs about their own health profiles and update those beliefs as they see their health changing with age. Leveraging data on subjective survival expectations, I find individuals are uncertain and are negatively biased; that is, on average, they believe their health will deteriorate faster than the average rate in the population.

Using the results from the learning model, I infer individuals' beliefs about their own health profiles. In an economic model with heterogeneous and uncertain health dynamics, those beliefs are inputs in the policy rule for labor participation. I flexibly estimate this policy rule, using a neural-network approach. I find beliefs matter for working decisions, and that expecting health to deteriorate more slowly is associated with larger probabilities of working. Furthermore, for individuals in their 50s who are not working, an interaction exists between beliefs and health, suggesting adjustment costs of finding a job are important in individuals' decisions regarding going back to work. The framework and data-driven estimation approach imply this result is not a consequence of any additional structure imposed on the economic model. In an additional exercise, applying machine-learning tools to also predict assets and health insurance, I simulate the effects on participation over time of eliminating the initial bias in beliefs. I find labor-force participation would increase by up to 2 percentage points, an effect that lasts beyond traditional retirement ages. Taken together, these results suggest room exists for policies to affect labor-participation decisions by shifting individuals' beliefs about their future health.

Thus, in the last part of the paper, I look at one such policy: the provision of information on blood-glucose and cholesterol levels. I take advantage of the randomization of the

collection and provision of such information, and analyze the results using a reduced-form approach and a model-based approach. The results show the additional information has negligible effects on survival expectations and working decisions. This negligible result is due to a small effect of the information on beliefs, as shown by the model. Nevertheless, the fact that this particular information policy does not have an effect on beliefs and working decisions of older adults does not mean other policies could not have an effect. Such policies could include information policies aimed at correcting bias in beliefs about aggregate values in the population, or more individualized information.[41] In the case of the HRS, policies could include providing information about biomarkers on kidney function and systemic inflammation, as well as genetic information, all already collected in the survey but with results not shared with individuals. As a final note of caution, note this paper assumes no endogenous acquisition of health information (e.g., through preventive care), an interesting question that is left for future research.

[41] Information policies have been studied in other settings, for example, Delavande and Kohler (2015) and Bates (2020). Information policies have also been studied in the context of surveys, for example, Armona et al. (2018) and Wiswall and Zafar (2014).

REFERENCES

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6):2375–2451.

Arcidiacono, P., Aucejo, E., Maurel, A., and Ransom, T. (2016). College attrition and the dynamics of information revelation.

Arcidiacono, P., Hotz, V. J., Maurel, A., and Romano, T. (2020). Ex ante returns and occupational choice. *Journal of Political Economy*, 128(12):4475–4522.

Arellano, M., Blundell, R., and Bonhomme, S. (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica*, 85(3):693–734.

Arellano, M. and Bonhomme, S. (2016). Nonlinear panel data estimation via quantile regressions. *The Econometrics Journal*, 19(3):C61–C94.

Armona, L., Fuster, A., and Zafar, B. (2018). Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment. *The Review of Economic Studies*, 86(4):1371–1410.

Bates, M. (2020). Public and private employer learning: Evidence from the adoption of teacher value added. *Journal of Labor Economics*, 38(2):375–420.

Blundell, R., Britton, J., Dias, M. C., and French, E. (2017). The impact of health on labour supply near retirement. IFS Working Papers W17/18, Institute for Fiscal Studies.

Bound, J., Schoenbaum, M., Stinebrickner, T. R., and Waidmann, T. (1999). The dynamic effects of health on the labor force transitions of older workers. *Labour Economics*, 6(2):179 – 202.

Breen, R. and García-Peñalosa, C. (2002). Bayesian learning and gender segregation. *Journal of Labor Economics*, 20(4):899–922.

Conlon, J. J., Pilossoph, L., Wiswall, M., and Zafar, B. (2018). Labor market search with imperfect information and learning. Working Paper 24988, National Bureau of Economic Research.

Contoyannis, P., Jones, A. M., and Rice, N. (2004). The dynamics of health in the british household panel survey. *Journal of Applied Econometrics*, 19(4):473–503.

Cunha, F. and Heckman, J. J. (2008). Symposium on noncognitive skills and their development: Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of Human Resources*, 43(4):738 – 782.

De Nardi, M., French, E., and Jones, J. B. (2016a). Medicaid insurance in old age. *American Economic Review*, 106(11):3480–3520.

De Nardi, M., French, E., and Jones, J. B. (2016b). Savings after retirement: A survey. *Annual Review of Economics*, 8(1):177–204.

Delavande, A. (2008). Pill, patch, or shot? subjective expectations and birth control choice*. *International Economic Review*, 49(3):999–1042.

Delavande, A. and Kohler, H.-P. (2015). HIV/AIDS-related Expectations and Risky Sexual Behaviour in Malawi. *The Review of Economic Studies*, 83(1):118–164.

Delavande, A. and Rohwedder, S. (2011). Individuals' uncertainty about future social security benefits and portfolio choice. *Journal of Applied Econometrics*, 26(3):498–519.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38.

Disney, R., Emmerson, C., and Wakefield, M. (2006). Ill health and retirement in britain: A panel data-based analysis. *Journal of Health Economics*, 25(4):621 – 649.

Dwyer, D. S. and Mitchell, O. S. (1999). Health problems as determinants of retirement: Are self-rated measures endogenous? *Journal of Health Economics*, 18(2):173 – 193.

Edwards, R. (2018). If my blood pressure is high, do i take it to heart? behavioral effects of biomarker collection in the health and retirement study. *Demography*, 55:403–434.

Elder, T. E. (2013). The predictive validity of subjective mortality expectations: Evidence from the health and retirement study. *Demography*, 50(2):569–589.

Farrell, M. H., Liang, T., and Misra, S. (2021). Deep neural networks for estimation and inference. *Econometrica*, 89(1):181–213.

French, E. (2005). The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. *The Review of Economic Studies*, 72(2):395–427.

French, E. and Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3):693–732.

French, E. and Jones, J. B. (2017). Health, health insurance, and retirement: A survey. *Annual Review of Economics*, 9(1):383–409.

Gong, Y., Stinebrickner, R., and Stinebrickner, T. R. (2019). Marriage, children, and labor supply: Beliefs and outcomes. Working Paper 26334, National Bureau of Economic Research.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. http://www.deeplearningbook.org.

Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent? *The American Economic Review*, 97(3):687–712.

Guvenen, F. and Smith, A. A. (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica*, 82(6):2085–2129.

Halliday, T. J. (2008). Heterogeneity, state dependence and health. *The Econometrics Journal*, 11(3):499–516.

Health and Retirement Study (2014). (RAND HRS Longitudinal File, HRS Core and Biomarker Data) public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI.

Heiss, F. (2011). Dynamics of self-rated health and selective mortality. *Empirical economics*, 40(1):119–140.

Heiss, F., Börsch-Supan, A., Hurd, M., and Wise, D. A. (2009). *Pathways to Disability: Predicting Health Trajectories*, pages 105–150. University of Chicago Press.

Heiss, F., Venti, S. F., and Wise, D. A. (2014). The persistence and heterogeneity of health among older americans. *National Bureau of Economic Research*, (20306).

Hernández-Quevedo, C., Jones, A. M., and Rice, N. (2008). Persistence in health limitations: A european comparative analysis. *Journal of Health Economics*, 27(6):1472 – 1488.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359 – 366.

Hurd, M. D., McFadden, D., and Merrill, A. (2001). Predictors of mortality among the elderly. In David A., W., editor, *Themes in the Economics of Aging*, chapter 5, pages 171–198. University of Chicago Press.

Hurd, M. D. and McGarry, K. (1995). Evaluation of the subjective probabilities of survival in the health and retirement study. *The Journal of Human Resources*, 30:S268–S292.

Hurd, M. D. and McGarry, K. (2002). The predictive validity of subjective probabilities of survival. *The Economic Journal*, 112(482):966–985.

Kleinjans, K. and Van Soest, A. (2014). Rounding, focal point answers and nonresponse to subjective probability questions. *Journal of Applied Econometrics*, 29:567–585.

Kolenikov, S. (2009). Confirmatory factor analysis using confa. *Stata Journal*, 9(3):329.

Lange, F. and McKee, D. (2011). The evolution of latent health over the life course. Yale University.

Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–50.

Ludwig, A. and Zimper, A. (2013). A parsimonious model of subjective life expectancy. *Theory and Decision*, 75(4):519–541.

Manski, C. and Molinari, F. (2010). Rounding probabilistic expectations in surveys. *Journal of Business & Economic Statistics*, 28(2):219–231.

Maurer, J., Klein, R., and Vella, F. (2011). Subjective health assessments and active labor market participation of older men: Evidence from a semiparametric binary choice model with nonadditive correlated individual-specific effects. *The Review of Economics and Statistics*, 93(3):764–774.

McClellan, M. (1998). Health events, health insurance and labor supply: Evidence from the health and retirement survey. In David A., W., editor, *Frontiers in the Economics of Aging*, chapter 7, pages 301–349. University of Chicago Press.

McGarry, K. (2004). Health and retirement: Do changes in health affect retirement expectations? *The Journal of Human Resources*, 39(3):624–648.

Myerson, R., Colantonio, L., Safford, M., and Meltzer, D. (2017). Chasing the missing diagnoses: Exploring the consequences of removing barriers to information.

Myerson, R. M., Colantonio, L. D., Safford, M. M., and Huang, E. S. (2018). Does identification of previously undiagnosed conditions change care-seeking behavior? *Health Services Research*, 53(3):1517–1538.

OECD (2015). Recommendation of the council on ageing and employment policies, oecd/legal/0419. Technical Report 0419, OECD.

OECD (2019). *Working Better with Age*.

Paula, u. D., Shapira, G., and Todd, P. E. (2014). How beliefs about hiv status affect risky behaviors: Evidence from malawi. *Journal of Applied Econometrics*, 29(6):944–964.

RAND HRS Longitudinal File (2014). Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA.

Siddiqui, S. (1997). The impact of health on retirement behaviour: empirical evidence from west germany. *Health Economics*, 6(4):425–438.

Smith, V. K., Taylor, D. H., and Sloan, F. A. (2001). Longevity expectations and death: Can people predict their own demise? *American Economic Review*, 91(4):1126–1134.

Stinebrickner, R. and Stinebrickner, T. (2014). A major in science? initial beliefs and final outcomes for college major and dropout. *The Review of Economic Studies*, 81(1 (286)):426–472.

Todd, P. E. and Wolpin, K. I. (2006). Assessing the impact of a school subsidy program in mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility. *American Economic Review*, 96(5):1384–1417.

van der Klaauw, W. and Wolpin, K. I. (2008). Social security and the retirement and savings behavior of low-income households. *Journal of Econometrics*, 145(1):21 – 42.

Wiswall, M. and Zafar, B. (2014). Determinants of College Major Choice: Identification using an Information Experiment. *The Review of Economic Studies*, 82(2):791–824.

Yogo, M. (2016). Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. *Journal of Monetary Economics*, 80:17–34.

Zucchelli, E., Jones, A., Rice, N., and Harris, A. (2010). The effect of health shocks on labour market exits: Evidence from the hilda survey. *Australian Journal of Labour Economics*, 13(2):191–218.

APPENDICES

A Estimation of the summary measure of health h_{it}

Let M_{ijt} be the j-th observed measure of unobserved health h_{it}, $j = 1, \ldots 11$, described in Table 1.1. I assume a linear factor model structure, that is,

$$M_{ijt} = \mu_j + \lambda_j h_{it} + \epsilon^h_{ijt}, \tag{18}$$

where ϵ^h_{ijt} is a measurement error. The coefficients μ_j are called intercepts and the coefficients λ_j are called loadings. I assume these coefficients are invariant in age t. Given that h_{it} is not directly measured, its location and scale are not identified without further assumptions. Hence, I fix the intercept and loading of one of the measures, the number of chronic conditions[1], to 0 and 1 respectively.

I use confirmatory factor analysis (CFA) to estimate (18), assuming classical measurement errors that are normally distributed. Note, however, the model is identified under weaker assumptions (see Cunha and Heckman (2008)). Estimates of the latent health h_{it} are obtained by minimizing the generalized sum of squares deviations of the factor from their true values. The resulting formula can also be justified as an empirical Bayes estimator of the factor given a prior normal distribution (Kolenikov (2009)). Table A1 presents the results. The table shows all coefficients have the expected sign and are significant at 1%. The table also shows the percentage of the variance of each measure M_{ijt} that is explained by health h_{it}. Variables regarding difficulties in ADLs have the higher R-squared, consistent with its common use in the assignment of many health-related benefits, such as long-term care services provided by Medicaid. The values of h_{it} predicted are highly correlated with

[1]For the measurement system, I define the variable as 7 minus the number of chronic conditions, so larger values represent better health.

Table A1: CFA results of health measurement

Measure of health	Coefficients		R-squared
	Intercept	Loading	
Number of chronic conditions[(a)]	0	1	0.29
Self-assessed health	8.188	-1.027	0.44
Body mass index	37.278	-1.812	0.05
Eyesight in general	5.710	-0.549	0.15
Eyesight at a distance	5.177	-0.502	0.13
Eyesight up close	5.465	-0.523	0.13
Hearing	4.830	-0.424	0.08
Pain	4.792	-0.802	0.36
Difficulties in ADLs regarding mobility	9.398	-1.598	0.64
Difficulties in ADLs of large muscles	8.964	-1.475	0.63
Difficulties in other ADLs	3.812	-0.654	0.50

Note: (a) The first measure corresponds to 7 minus the number of chronic conditions, hence, larger values represent better health. For this variable, the intercept and loading are fixed to 0 and 1, respectively. All other coefficients are significant at 1%.

the values predicted by using principal component analysis instead.

B MLE results for health and survival

Consider the following health and survival processes,

$$\begin{aligned} h_{it} &= \rho h_{it-1} + \alpha_i + \delta_i \cdot t + \tau \cdot t^2 + \epsilon_{it} \\ S_{it} &= \mathbb{1}\{\gamma h_{it-1} + \iota_1 \alpha_i + \iota_2 \delta_i + \iota_3 \cdot t \cdot \alpha_i + \iota_4 \cdot t \cdot \delta_i + \theta_0 + \theta_1 \cdot t + \theta_2' x_i + \eta_{it}\} S_{it-1}, \end{aligned}$$

where ϵ_{it} is i.i.d $N(0, \sigma_\epsilon^2)$ and η_{it} is i.i.d. $N(0,1)$, independent of ϵ_{it}. Furthermore,

$$\begin{pmatrix} \alpha_i \\ \delta_i \end{pmatrix} \Big| x_i, h_{i0} \sim N\left(\begin{pmatrix} \mu_\alpha + \nu_\alpha' x_i + \omega_\alpha h_{i0} \\ \mu_\delta + \nu_\delta' x_i + \omega_\delta h_{i0} \end{pmatrix}, \begin{bmatrix} \sigma_\alpha^2 & \phi \sigma_\alpha \sigma_\delta \\ \phi \sigma_\alpha \sigma_\delta & \sigma_\delta^2 \end{bmatrix} \right).$$

Let Θ be the set of parameters of this random-coefficients model. The likelihood corresponds to

$$\max_{\Theta} \sum_{i=1}^{N} log\Big(\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \prod_{t=1}^{T_i} \mathbb{P}\big(h_{it}, S_{it} | h_{it-1}, S_{it-1} = 1, x_i, \alpha, \delta\big) \cdot \phi(\alpha, \delta | x_i, h_{i0}) d\alpha d\delta \Big),$$

where,

$$\begin{aligned} \mathbb{P}\big(h_{it}, S_{it} = 1 | h_{it-1}, S_{it-1} = 1, x_i, \alpha, \delta\big) &= \phi\big(h_{it} - \rho h_{it-1} - \alpha - \delta \cdot t - \tau \cdot t^2\big) \\ &\cdot \Phi\big(\gamma h_{it-1} + \iota_1 \alpha_i + \iota_2 \delta_i + \iota_3 \cdot t \cdot \alpha_i + \iota_4 \cdot t \cdot \delta_i + \theta_0 + \theta_1 \cdot t + \theta_2' x_i\big) \\ \mathbb{P}\big(S_{it} = 0 | h_{it-1}, S_{it-1} = 1, x_i, \alpha, \delta\big) &= \\ & 1 - \Phi\big(\gamma h_{it-1} + \iota_1 \alpha_i + \iota_2 \delta_i + \iota_3 \cdot t \cdot \alpha_i + \iota_4 \cdot t \cdot \delta_i + \theta_0 + \theta_1 \cdot t + \theta_2' x_i\big). \end{aligned}$$

Table B1 presents the MLE results of estimating these equations under 3 different set of assumptions. Columns (3) and (4) present the main results. Columns (1) and (2) present the results of an MLE estimation of health only, ignoring survival. Columns (5) and (6) present

the results of an MLE estimation of health and survival, but assuming no heterogeneity in health slopes, that is, $\delta_i = \delta$[1]. Thus, for these columns, $\sigma_\delta = \phi = \omega_\delta = \iota_2 = \iota_4 = 0$. Note also the equation for survival includes direct effects of the individual heterogeneity (α_i, δ_i), hence it allows us to test for these direct effects. The results imply these direct effects are not (jointly) significant; hence, survival does not provide additional information on δ_i.

[1]The model allows for differences in slopes by observed heterogeneity, but it does not allow for differences in slopes by unobserved heterogeneity.

Table B1: MLE results on health and survival under different assumptions

	Heterogeneous slopes without survival eq		Heterogeneous slopes with survival eq		Homogeneous slopes with survival eq	
	Coefficient (1)	Pvalue (2)	Coefficient (3)	Pvalue (4)	Coefficient (5)	Pvalue (6)
ρ	0.225	0.000	0.223	0.000	0.366	0.000
τ	0.001	0.087	0.001	0.119	0.001	0.108
μ_α	0.968	0.000	0.955	0.000	0.781	0.000
$\nu_{\alpha female}$	-0.029	0.132	-0.029	0.131	-0.024	0.163
$\nu_{\alpha white}$	0.026	0.338	0.027	0.335	0.018	0.458
$\nu_{\alpha hispanic}$	0.004	0.909	0.005	0.889	-0.001	0.973
$\nu_{\alpha less\_HS}$	-0.134	0.000	-0.134	0.000	-0.120	0.000
ω_α	0.599	0.000	0.603	0.000	0.492	0.000
μ_δ	-0.060	0.012	-0.057	0.018	-0.051	0.000
$\nu_{\delta female}$	0.006	0.146	0.006	0.136	0.005	0.198
$\nu_{\delta white}$	0.015	0.007	0.015	0.008	0.013	0.011
$\nu_{\delta hispanic}$	0.010	0.196	0.010	0.199	0.006	0.390
$\nu_{\delta less\_HS}$	-0.003	0.677	-0.003	0.624	0.001	0.896
ω_δ	0.000	0.956	0.000	0.962		
σ_α	0.235	0.000	0.235	0.000	0.212	0.000
σ_δ	0.042	0.000	0.043	0.000		
ϕ	-0.030	0.741	-0.033	0.714		
σ_ϵ	0.266	0.000	0.266	0.000	0.285	0.000
γ			0.583	0.001	0.640	0.000
ι_1			-0.277	0.334	-0.422	0.125
ι_2			0.044	0.986		
ι_3			0.029	0.306	0.036	0.287
ι_4			0.241	0.601		
θ_0			0.529	0.326	0.514	0.336
θ_1			-0.178	0.136	-0.193	0.092
$\theta_{2female}$			0.259	0.002	0.255	0.002
θ_{2white}			0.019	0.847	0.029	0.758
$\theta_{2hispanic}$			0.317	0.079	0.311	0.078
θ_{2less_HS}			-0.106	0.305	-0.114	0.267
N alive observations	8,901		8,901		8,901	
N dead observations	0		112		112	
N individuals	1,671		1,671		1,671	
-LL	2,498.6		3,027.6		3,067.6	

Note: Standard errors are clustered at the individual level.

C Robustness: MLE results for self-assessed health

Let h_{it}^{SAH} denote the 1 to 5 self-assessed health (SAH) measure, rescaled so that larger values represent better health. In this section, I estimate a model similar to the main model in the text, but using h_{it}^{SAH} instead of the summary measure h_{it}. Let $\tilde{h}_{it}^{SAH}$ be the latent health variable for h_{it}^{SAH}. Consider the following equations,

$$
\begin{aligned}
\tilde{h}_{it}^{SAH} &= \rho h_{it-1}^{SAH} + \alpha_i + \delta_i \cdot t + \tau \cdot t^2 + \epsilon_{it}, \quad \epsilon_{it} \sim N(0,1), \quad t \geq 1 \\
S_{it} &= \mathbb{1}\{\gamma h_{it-1} + \theta_0 + \theta_1 \cdot t + \theta_2' x_i + \eta_{it}\} S_{it-1}, \qquad \eta_{it} \text{ i.i.d. } N(0,1) \\
h_{it}^{SAH} &= \begin{cases} 1 & \text{if} \quad \tilde{h}_{it}^{SAH} \leq 0 \\ 2 & \text{if} \quad 0 < \tilde{h}_{it}^{SAH} \leq O_2 \\ 3 & \text{if} \quad O_2 < \tilde{h}_{it}^{SAH} \leq O_3 \\ 4 & \text{if} \quad O_3 < \tilde{h}_{it}^{SAH} \leq O_4 \\ 5 & \text{if} \quad O_4 < \tilde{h}_{it}^{SAH} \end{cases}
\end{aligned}
$$

with individual-level heterogeneity (α_i, δ_i),

$$
\begin{pmatrix} \alpha_i \\ \delta_i \end{pmatrix} \Big| x_i, h_{i0} \sim N\left(\begin{pmatrix} \mu_\alpha + \nu_\alpha' x_i + \omega_\alpha h_{i0} \\ \mu_\delta + \nu_\delta' x_i + \omega_\delta h_{i0} \end{pmatrix}, \begin{bmatrix} \sigma_\alpha^2 & \phi \sigma_\alpha \sigma_\delta \\ \phi \sigma_\alpha \sigma_\delta & \sigma_\delta^2 \end{bmatrix} \right).
$$

This system of equations is similar to the system of equations in the main text, but replace h_{it} for discrete h_{it}^{SAH}.

Table C1 presents the MLE results of estimating these equations. The table shows that in this case too there is evidence of slope heterogeneity, that is, $\sigma_\delta > 0$. Thus, heterogeneity in health dynamics is robust to using SAH instead of the summary measure of health used in the main analysis.

Table C1: MLE results for SAH with and without a survival equation

	Without survival eq		With survival eq	
	Coefficient	Pvalue	Coefficient	Pvalue
ρ	0.230	0.000	0.230	0.000
τ	0.012	0.000	0.012	0.000
μ_α	-1.168	0.000	-1.185	0.000
$\nu_{\alpha female}$	-0.006	0.939	-0.005	0.951
$\nu_{\alpha white}$	0.236	0.010	0.242	0.009
$\nu_{\alpha hispanic}$	-0.265	0.048	-0.266	0.047
$\nu_{\alpha less\_HS}$	-0.612	0.000	-0.603	0.000
ω_α	1.148	0.000	1.151	0.000
μ_δ	-0.057	0.158	-0.054	0.182
$\nu_{\delta female}$	0.030	0.085	0.029	0.089
$\nu_{\delta white}$	-0.008	0.696	-0.009	0.647
$\nu_{\delta hispanic}$	0.060	0.040	0.060	0.040
$\nu_{\delta less\_HS}$	0.020	0.378	0.019	0.406
ω_δ	-0.043	0.000	-0.043	0.000
σ_α	0.970	0.000	0.970	0.000
σ_δ	0.137	0.000	0.137	0.000
ϕ	-0.258	0.004	-0.257	0.004
γ			0.402	0.000
θ_0			1.371	0.000
θ_1			-0.101	0.000
$\theta_{2female}$			0.164	0.043
θ_{2white}			0.034	0.711
$\theta_{2hispanic}$			0.404	0.018
θ_{2less_HS}			-0.076	0.457
O_2	1.713	0.000	1.712	0.000
$O_3 - O_2$	1.711	0.000	1.711	0.000
$O_4 - O_3$	2.062	0.000	2.063	0.000
N alive observations	8,901		8,901	
N dead observations	0		112	
N individuals	1,671		1,671	
-LL	8,985.2		9,502.0	

Note: Standard errors are clustered at the individual level.

D Robustness: MLE with heteroskedastic errors ϵ_{it}

In this appendix, I estimate the health and survival processes defined in equations (1.9), (1.10) and (1.11), except that I allow for heterokedastic errors in the health equation, such that, $Var(\epsilon_{it}) = \sigma_\epsilon^2 + t \cdot \sigma_{t\epsilon}^2$. Table D1 presents the results of estimating these equations by MLE. The table shows allowing for increasing variance of health shocks does not explain away heterogeneity in health slopes δ_i.

Table D1: MLE results for health h_{it} allowing for heteroskedastic error ϵ_{it}

	Coefficient	Pvalue
ρ	0.225	0.000
τ	0.001	0.088
μ_α	0.961	0.000
$\nu_{\alpha female}$	-0.03	0.122
$\nu_{\alpha white}$	0.027	0.330
$\nu_{\alpha hispanic}$	0.003	0.928
$\nu_{\alpha less\_HS}$	-0.134	0.000
ω_α	0.601	0.000
μ_δ	-0.059	0.015
$\nu_{\delta female}$	0.006	0.139
$\nu_{\delta white}$	0.015	0.008
$\nu_{\delta hispanic}$	0.010	0.193
$\nu_{\delta less\_HS}$	-0.003	0.661
ω_δ	0.000	0.986
σ_α	0.234	0.000
σ_δ	0.042	0.000
ϕ	-0.025	0.776
σ_ϵ	0.266	0.000
$\sigma_{\epsilon t}$	0.000	1.000
γ	0.494	0.000
θ_0	-0.103	0.707
θ_1	-0.083	0.000
$\theta_{2female}$	0.244	0.005
θ_{2white}	0.025	0.793
$\theta_{2hispanic}$	0.248	0.263
θ_{2less_HS}	-0.096	0.345

Note: Standard errors are clustered at the individual level.

E Formula for $plive10_{it}$

Let s denote the reference age asked in $plive10_{it}$. By definition,

$$\begin{aligned} plive10_{it} &= \mathbb{P}(S_{is}=1|\Omega_{it}) = \prod_{l=t}^{s-1}\mathbb{P}(S_{il+1}=1|S_{il}=1,\Omega_{it}) \\ &= \prod_{l=t}^{s-1}\mathbb{P}(\gamma h_{il}+\theta_0+\theta_1(l-1)+\theta_2'x_i+\eta_{il+1}\geq 0|\Omega_{it}), \end{aligned}$$

where

$$h_{il} = \underbrace{\rho^{l-t}h_{it}+\alpha_i\sum_{k=0}^{l-t-1}\rho^k}_{\text{known under }\Omega_{it}} + \underbrace{\delta_i\sum_{k=0}^{l-t-1}(l-k)\rho^k+\sum_{k=0}^{l-t-1}\rho^k\epsilon_{i(l-k)}}_{\text{unknown under }\Omega_{it}}.$$

Then,

$$\mathbb{P}(S_{il+1}=1|S_{il}=1,\Omega_{it}) = \mathbb{P}(\gamma h_{il}+\eta_{il+1}\geq 0|\Omega_{it}) = \Phi\left(\frac{M_{itl}}{W_{tl}^{1/2}}\right)$$

where

$$\begin{aligned} M_{itl} &= \gamma\left(\rho^{l-t}h_{it}+\alpha_i\sum_{k=0}^{l-t-1}\rho^k+\hat{\delta}_{it}\sum_{k=0}^{l-t-1}(l-k)\rho^k\right)+\theta_0+\theta_1(l-1)+\theta_2'x_i \\ W_{tl} &= \gamma^2\hat{\sigma}_t^2\left(\sum_{k=0}^{l-t-1}(l-k)\rho^k\right)^2+\gamma^2\sigma_\epsilon^2\sum_{k=0}^{l-t-1}\rho^{2k}+1 \end{aligned}$$

Note that M_{itl} and W_{tl} are functions of h_{it}, α_i, $\hat{\delta}_{it}$, $\hat{\sigma}_t^2$, x_i and parameters of the model. Hence,

$$plive10_{it} = \prod_{l=t}^{s-1}\Phi\left(\frac{M_{itl}}{W_{tl}^{1/2}}\right) = plive10_{it}(\alpha_i, h_{it}, \hat{\delta}_{it}, \hat{\sigma}_t^2, x_i).$$

Furthermore, beliefs at age t are a function of prior beliefs at age 50 ($t = 0$), the heterogeneity in levels α_i, and the health history up to that point h_i^t (and parameters of the model). The exact formulas come from applying the Bayesian updating equations recursively. First, for the posterior variance,

$$\frac{1}{\hat{\sigma}_t^2} = \frac{1}{\hat{\sigma}_{t-1}^2} + \frac{t^2}{\sigma_\epsilon^2} \Rightarrow \frac{1}{\hat{\sigma}_t^2} = \frac{1}{\hat{\sigma}_0^2} + \frac{1}{\sigma_\epsilon^2} \sum_{l=1}^{t} l^2$$

We can also rewrite the Bayesian updating equation for the posterior mean as

$$\hat{\delta}_{it} = (1 - tK_t)\hat{\delta}_{it-1} + K_t(h_{it} - \rho h_{it-1} - \alpha_i - \tau t^2) \tag{19}$$

where $K_t = \frac{t\hat{\sigma}_t^2}{\sigma_\epsilon^2}$. Moreover, K_t satisfies that $(1 - tK_t)K_{t-1} = \frac{t-1}{t}K_t$. Using this property and equation (19) recursively, we can write

$$\hat{\delta}_{it} = \hat{\delta}_{i0} \prod_{l=1}^{t} (1 - lK_l) + \sum_{l=1}^{t} \frac{l}{t} K_t(h_{il} - \rho h_{il-1} - \alpha_i - \tau l^2)$$

Noting that K_t is a function of $\hat{\sigma}_0^2$, σ_ϵ^2 and t, we conclude

$$plive10_{it} = plive10_{it}(\alpha_i, h_i^t, \hat{\delta}_{i0}, \hat{\sigma}_0^2, x_i).$$

F Proof of proposition 1.5.1

Identification of λ with ideal data

We could identify λ with longitudinal information on beliefs about survival rates[1],

$$bsr_{itr} = \mathbb{P}(S_{ir+1} = 1 | S_{ir} = 1, \Omega_{it}) = \mathbb{P}(\gamma h_{ir} + \eta_{ir+1} \geq 0 | \Omega_{it}).$$

From the equation for health (1.9),

$$h_{ir} = \rho^{r-t} h_{it} + \alpha_i \sum_{k=0}^{r-t-1} \rho^k + \delta_i \sum_{k=0}^{r-t-1} (r-k)\rho^k + \sum_{k=0}^{r-t-1} \rho^k \epsilon_{ir-k}$$

Hence,

$$h_{ir}|\Omega_{it} \sim N\left(\rho^{r-t} h_{it} + \alpha_i \sum_{k=0}^{r-t-1} \rho^k + \hat{\delta}_{it} \sum_{k=0}^{r-t-1} (r-k)\rho^k, \hat{\sigma}_t^2 \left(\sum_{k=0}^{r-t-1} (r-k)\rho^k \right)^2 + \sigma_\epsilon^2 \sum_{k=0}^{r-t-1} \rho^{2k} \right)$$

Defining

$$w_{tr} = \frac{1}{\gamma} \sqrt{\gamma^2 \hat{\sigma}_t^2 \left(\sum_{k=0}^{r-t-1} (r-k)\rho^k \right)^2 + \gamma^2 \sigma_\epsilon^2 \sum_{k=0}^{r-t-1} \rho^{2k} + 1}$$

we can write

$$\begin{aligned} \Delta_w \Phi^{-1}(bsr_{it+1r}) &\equiv w_{t+1r} \Phi^{-1}(bsr_{it+1r}) - w_{tr} \Phi^{-1}(bsr_{itr}) \\ &= \rho^{r-t-1} \big(h_{it+1} - \alpha_i - \hat{\delta}_{it}(t+1) \big) + (\hat{\delta}_{it+1} - \hat{\delta}_{it}) \sum_{k=0}^{r-t-2} (r-k)\rho^k \quad (20) \end{aligned}$$

[1]For ease of notation, in this section I ignore the quadratic term for age in the health equation and the demographic terms in the survival equation.

We denote individual i's perceived innovation in health at period t as

$$\hat{\zeta}_{it} = h_{it} - \mathbb{E}(h_{it}|\Omega_{it-1}) = h_{it} - \rho h_{it-1} - \alpha_i - \hat{\delta}_{it-1} \cdot t$$

and note that the Bayesian updating formulas can be rewritten as

$$\begin{aligned} \hat{\delta}_{it} &= \hat{\delta}_{it-1} + K_t \cdot \hat{\zeta}_{it} \\ \hat{\sigma}_t^2 &= (1 - K_t \cdot t)\hat{\sigma}_{t-1}^2 \end{aligned}$$

where $K_t = \frac{\hat{\sigma}_{t-1}^2 \cdot t}{\hat{\sigma}_{t-1}^2 \cdot t^2 + \sigma_\epsilon^2} = \frac{t\hat{\sigma}_t^2}{\sigma_\epsilon^2}$, $K_t \leq 1$. Then, we can write equation (20) as

$$\begin{aligned} \Delta_w \Phi^{-1}(bsr_{it+1r}) &= \left(\rho^{r-t-1} + K_{t+1} \sum_{k=0}^{r-t-2} (r-k)\rho^k\right)\hat{\zeta}_{it+1} \\ &= \left(\rho^{r-t-1} + K_{t+1} \sum_{k=0}^{r-t-2} (r-k)\rho^k\right)(h_{it+1} - w_{tt+1}\Phi^{-1}(bsr_{itt+1})) \end{aligned}$$

Then, for $r = t+2$, conditional on h_{it}, bsr_{itt+1} and bsr_{itt+2} (belonging to Ω_{it}),

$$Cov(\Delta\Phi^{-1} bsr_{it+1t+2}, \Delta h_{it+1}) = \underbrace{\frac{(\rho + (t+2)K_{t+1})}{w_{t+1t+2}}}_{C_t} \cdot Var(\Delta h_{it+1})$$

Finally,

$$\begin{aligned} \frac{\partial C_t}{\partial \lambda} &= \frac{1}{w_{t+1t+2}^2}\left[\frac{(t+2)(t+1)}{\sigma_\epsilon^2}\frac{\partial \hat{\sigma}_{t+1}^2}{\partial \lambda} w_{t+1t+2} - \frac{1}{2w_{t+1t+2}}(t+2)^2 \frac{\partial \hat{\sigma}_{t+1}^2}{\partial \lambda}\left(\rho + (t+2)(t+1)\frac{\hat{\sigma}_{t+1}^2}{\sigma_\epsilon^2}\right)\right] \\ &= \frac{t+2}{w_{t+1t+2}^3}\frac{\partial \hat{\sigma}_{t+1}^2}{\partial \lambda}\left[\frac{(t+1)}{\sigma_\epsilon^2}(\hat{\sigma}_{t+1}^2(t+2)^2 + \sigma_\epsilon^2 + 1/\gamma^2) - \frac{1}{2}(t+2)\left(\rho + (t+2)(t+1)\frac{\hat{\sigma}_{t+1}^2}{\sigma_\epsilon^2}\right)\right] \\ &= \frac{t+2}{w_{t+1t+2}^3}\frac{\partial \hat{\sigma}_{t+1}^2}{\partial \lambda}\left[\frac{(t+1)}{2\sigma_\epsilon^2}\hat{\sigma}_{t+1}^2(t+2)^2 + \frac{t+1}{\gamma^2\sigma_\epsilon^2} + \underbrace{(t+1) - \frac{1}{2}(t+2)\rho}_{>0}\right] \geq 0 \end{aligned}$$

■

G Strategy for simulating survival expectations $plive10_{it}$

To estimate the bias b and uncertainty λ parameters, I use simulated method of moments comparing empirical moments of observed $plive10_{it}$ with simulated moments of $plive10_{it}$. The simulated moments come from $plive10_{it}$ being a function of individual-level heterogeneity α_i, health history h_i^t, initial beliefs $(\hat{\delta}_{i0}, \hat{\sigma}_0^2)$ and demographic characteristics x_i. The exact expression of this function is derived in appendix E. In this expression, α_i and $\hat{\delta}_{i0}$ are random variables unobserved by the econometrician, but with a know distribution, given b and λ.

Let t_0 denote the age an individual is first observed in the data, and let T denote the age an individual is last observed in the data. These values are individual-specific, but I omit the index i for ease of notation. The simulation strategy depends on the age an individual is first observe, t_0.

Case $t_0 = 0$, individuals first observed in data at age 50

In this case, the health history relevant for beliefs, that is, the health history starting at 50 years old is fully observed. Then, the simulation strategy is straightforward:

1. Draw (α_i, δ_i) conditional on $h_{i0}, \dots h_{iT}$ (which follows a known normal distribution).
2. For a given b and λ,
   (a) Set $\hat{\sigma}_0^2 = \lambda^2 \sigma_\delta^2$.
   (b) Draw $\hat{\delta}_{i0}$ conditional in $\alpha_i, \delta_i, h_{i0}$ (which follows a known normal distribution given b and λ).
   (c) Use α_i, h_i^T, x_i, $\hat{\delta}_{i0}$, and $\hat{\sigma}_0^2$ to set $plive10_{it}$ (according to the formula in section E).

Case $t_0 > 0$, individuals first observed in data at age older than 50

In this case, we only observe $h_{it_0}, \dots h_{iT}$. Moreover, the prior mean $\hat{\delta}_{it_0}$ at that point is not random conditional on b and λ, because survival up to the point depends on past health, and therefore on health profiles. Instead, it satisfies,

$$\hat{\delta}_{it_0} = K_{t_0}(\lambda)\Bigg[-\rho^{t_0} h_{i0} - \alpha_i \sum_{k=0}^{t_0-1} \rho^k + \delta_i \Bigg(\frac{1}{t_0} \sum_{l=1}^{t_0-1} l^2 - \sum_{k=1}^{t_0-1} (t_0 - k)\rho^k \Bigg)$$
$$-\rho T_{i1} + T_{i2} \frac{1}{t_0} + \Bigg(h_{it_0} - \gamma \sum_{k=0}^{t_0-1} (t_0 - k)^2 \rho^k \Bigg) \Bigg] + \hat{\delta}_{i0} \frac{\sigma_\epsilon^2}{\lambda^2 \sigma_\delta^2} \frac{K_{t_0}(\lambda)}{t_0} \tag{21}$$

where

$$T_{i1} = \sum_{l=1}^{t_0-1} \rho^{t_0-1-l} \epsilon_{il}, \qquad T_{i2} = \sum_{l=1}^{t_0-1} l \epsilon_{il}$$

are random variables, and $K_{t_0}(\lambda)$ is constant across individuals depending on both λ and t_0. According to this expression,

$$\hat{\delta}_{it_0} = \hat{\delta}_{it_0}(\underbrace{h_{i0}, \alpha_i, \delta_i, T_{i1}, T_{i2}, \hat{\delta}_{i0}}_{\substack{\text{unobserved by} \\ \text{the econometrician}}}, h_{it_0}; \lambda)$$

Hence, we can simulate $\hat{\delta}_{it_0}$ by simulating $(h_{i0}, \alpha_i, \delta_i, T_{i1}, T_{i2}, \hat{\delta}_{i0})$ and using (21) to define $\hat{\delta}_{it_0}$. However, being alive at t_0 further restricts the distribution of $(h_{i0}, \alpha_i, \delta_i, T_{i1}, T_{i2})$. The distribution of this vector conditional on observed health history and conditional on surviving up to t_0 has no closed-form solution. Hence, I use Markov chain Monte Carlo (MCMC) to get these conditional draws. In this case, the simulation strategy is the following:

1. Draw $(h_{i0}, \alpha_i, \delta_i, T_{i1}, T_{i2})$ conditional on $h_{it_0} \dots h_{iT}, S_{it_0} = 1$ by MCMC.

2. For a given b and λ,

(a) Set $\hat{\sigma}^2_{t_0} = \hat{\sigma}^2(\lambda^2, \sigma^2_\delta, t_0)$ (defined by the Bayesian updating equation for the posterior variance).

(b) Draw $\hat{\hat{\delta}}_{i0}$ conditional on $\alpha_i, \delta_i, h_{i0}$ (which follows a known normal distribution given b and λ).

(c) Use $\hat{\hat{\delta}}_{i0}$ and $(h_{i0}, \alpha_i, \delta_i, T_{i1}, T_{i2})$ to construct $\hat{\hat{\delta}}_{it_0}$ according to (21).

(d) Use α_i, $h_{it_0}, \ldots h_{iT}$, x_i, $\hat{\hat{\delta}}_{it_0}$, and $\hat{\sigma}^2_{t_0}$ to set $plive10_{it}$ (according to a modification of the formula in section E, starting at t_0).

Overall, I target moments of averages across time for sub samples of individuals with different values of t_0.

H Probit likelihood and results

The working decision of individual i at age t depends on the information he has at that moment, including his age t, past participation p_{it-1}, past health h_{it-1}, beliefs about health slopes $(\hat{\delta}_{it-1}, \hat{\sigma}^2_{t-1})$ and individual-level heterogeneity α_i. His decision also depends on his assets a_{it-1}, past labor income w_{it-1}, and demographic variables, all of which I denote together as x_{it-1}.

$$\begin{aligned}\mathbb{P}(p_{it}=1|\Omega_{it-1}) &= \Phi\Bigg(\beta_0+\beta_{0t}t+\beta_1 h_{it-1}+\underbrace{\beta_2\hat{\delta}_{it-1}+\beta_3\hat{\sigma}^2_{t-1}+\beta_4\alpha_i}_{\substack{\text{unobserved to the}\\ \text{econometrician}}}+\beta_5 p_{it-1}+\beta_6 x_{it-1}\Bigg)\\ &= \Phi\Big(\beta'\Omega_{it-1}\Big)\end{aligned}$$

Conditional on Ω_{it-1}, the likelihood of p_{it} is

$$L^c_{it} = \Phi\Big(\beta'\Omega_{it-1}\Big)^{p_{it}}\cdot\Big(1-\Phi\Big(\beta'\Omega_{it-1}\Big)\Big)^{1-p_{it}}$$

Let t_0 be the age at which individual i is first observed in the data[1]. Then, according to the economic framework discussed in section 1.2, the likelihood of the vector $(p_{it_0+1},\dots p_{iT})$ conditional on p_{it_0} and $(t_0, T, h_{it_0},\dots h_{iT},\hat{\delta}_{it_0},\dots\hat{\delta}_{iT},\hat{\sigma}^2_{t_0},\dots\hat{\sigma}^2_T,\alpha_i,x_{it_0},\dots x_{iT})$ is

$$L^c_i = \prod_{t=t_0+1}^{T} L^c_{it}$$

I address the initial condition problem by modeling the initial condition p_{it_0} as a function of $(t_0, h_{it_0}, \hat{\delta}_{i0}, \hat{\sigma}_0, \alpha_i, x_{i0})$

$$\mathbb{P}(p_{it_0}=1) = \Phi\Big(\gamma_0+\gamma_{0t}t_0+\gamma_1 h_{it_0}+\gamma_2\hat{\delta}_{it_0}+\gamma_3\hat{\sigma}^2_{t_0}+\gamma_4\alpha_i+\gamma_6 x_{it_0}\Big)=\tilde{L}^c_{it_0}$$

[1]Note t_0 and T are individual specific, though I omit that index for ease of notation.

This last equation is not derived from the economic model, but it is assumed as a way of approximating the conditional density of p_{it_0}. Then, the likelihood of observing $(p_{it_0}, \dots p_{iT})$, conditional on $(t_0, T, h_{it_0}, \dots h_{iT}, \hat{\delta}_{it_0} \dots \hat{\delta}_{iT}, \hat{\sigma}_{t_0} \dots \hat{\sigma}^2_T, \alpha_i, x_{it_0}, \dots x_{iT})$ is given by

$$L^c_i = \tilde{L}^c_{it_0} \prod_{t=t_0+1}^{T} L^c_{it}$$

However, this likelihood is conditional on variables that are unobserved by the econometrician, namely, slopes beliefs $\{\hat{\delta}_{it}, \hat{\sigma}^2_t\}^T_{t=t_0}$ and heterogeneity in health levels α_i. These time-varying unobserved variables can be written as a function of time-varying observed health $(h_{it_0} \dots h_{iT})$ and time-invariant unobserved variables, namely, beliefs at t_0 $(\hat{\delta}_{it_0}, \hat{\sigma}^2_{t_0})$ and α_i. Hence, I write instead the likelihood of $(p_{it_0}, \dots p_{iT})$, conditional on $(t_0, T, h_{it_0} \dots h_{iT}, plive10_{it_0}, \dots plive10_{iT}, x_{it_0} \dots x_{iT})$ (observed by the econometrician), integrating out this time-invariant unobserved heterogeneity,

$$L_i = \int L^c_i \cdot f(\alpha_i, \hat{\delta}_{it_0} | t_0, T, h_{it_0}, \dots h_{iT}, plive10_{it_0} \dots plive10_{iT}, x_{it_0}, \dots x_{iT})$$

where I used that $\hat{\sigma}^2_t$ is constant for individuals of the same age. Note I added in the conditional set $plive10_{it0}, \dots plive10_{iT}$. These variables do not enter the economic model, and hence the probability of working, but they provide information on individuals slopes beliefs. This formulation assumes no other unobserved heterogeneity at the i-level.

The distribution within the integral has no closed form solution, given that surviving up to t_0 adds additional restrictions on the distribution of the underlying individual heterogeneity. Hence, in practice, I approximate this integral using draws from this distribution gotten by Markov chain Monte Carlo (MCMC). Tables H1 and H2 present the full set of results of this integrated probit. They also include a specification using survival expectations $plive10_{it-1}$ instead of slope beliefs $(\hat{\delta}_{it-1}, \hat{\sigma}^2_{t-1})$, and a specification using both, survival expectations and

slope beliefs.

Table H1: Probit results on probability of working: main equation

	(1)		(2)		(3)	
	Coefficient	SE	Coefficient	SE	Coefficient	SE
Main equation						
intercept	-0.564	(0.294)	-2.445	(0.098)	-0.693	(0.297)
$t-1$	-0.196	(0.016)	-0.082	(0.003)	-0.192	(0.016)
work	2.032	(0.018)	2.031	(0.019)	2.034	(0.019)
health	0.169	(0.024)	0.261	(0.033)	0.175	(0.046)
educ LHS	-0.032	(0.020)	-0.034	(0.021)	-0.032	(0.022)
MS married	-0.030	(0.040)	-0.014	(0.041)	-0.012	(0.041)
MS divorce	0.053	(0.043)	0.064	(0.044)	0.069	(0.045)
MS widow	0.012	(0.045)	0.029	(0.046)	0.028	(0.046)
Q1 income	-0.283	(0.026)	-0.294	(0.027)	-0.290	(0.027)
Q2 income	-0.165	(0.022)	-0.168	(0.023)	-0.165	(0.023)
Q3 income	-0.105	(0.020)	-0.112	(0.020)	-0.108	(0.020)
Q1 wealth	0.176	(0.024)	0.181	(0.025)	0.187	(0.025)
Q2 wealth	0.112	(0.022)	0.112	(0.022)	0.117	(0.022)
Q3 wealth	0.027	(0.020)	0.025	(0.021)	0.027	(0.021)
female	-0.037	(0.015)	-0.048	(0.016)	-0.036	(0.016)
α_i	0.244	(0.036)	0.074	(0.046)	0.243	(0.075)
$\hat{\delta}_{it-1}$	1.933	(0.249)			1.903	(0.499)
$\hat{\sigma}^2_{t-1}/\sigma^2_\delta$	-13.854	(2.048)			-13.335	(2.102)
$plive10_{it-1}$			0.114	(0.031)	0.007	(0.043)

Note: Standard errors are clustered at the individual level.

Table H2: Probit results on probability of working: initial condition

	(1)		(2)		(3)	
	Coefficient	SE	Coefficient	SE	Coefficient	SE
Initial condition						
intercept	-2.840	(0.417)	-1.583	(0.138)	-2.779	(0.419)
t_0	-0.107	(0.022)	-0.163	(0.004)	-0.106	(0.022)
health	0.481	(0.040)	0.549	(0.058)	0.448	(0.083)
educ LHS	-0.059	(0.032)	-0.040	(0.033)	-0.038	(0.033)
MS married	-0.276	(0.063)	-0.297	(0.063)	-0.288	(0.063)
MS divorce	0.055	(0.068)	0.045	(0.068)	0.051	(0.069)
MS widow	0.023	(0.072)	0.008	(0.072)	0.012	(0.073)
Q1 income	-1.201	(0.045)	-1.227	(0.045)	-1.218	(0.046)
Q2 income	-0.677	(0.039)	-0.708	(0.039)	-0.703	(0.039)
Q3 income	-0.413	(0.035)	-0.426	(0.035)	-0.421	(0.035)
Q1 wealth	0.709	(0.043)	0.695	(0.043)	0.703	(0.044)
Q2 wealth	0.512	(0.039)	0.507	(0.039)	0.513	(0.039)
Q3 wealth	0.249	(0.037)	0.253	(0.037)	0.255	(0.037)
female	-0.09	(0.025)	-0.097	(0.026)	-0.079	(0.026)
α_i	0.200	(0.057)	0.057	(0.076)	0.249	(0.126)
$\hat{\delta}_{it_0}$	1.473	(0.383)			2.238	(0.788)
$\hat{\sigma}^2_{t_0}/\sigma^2_{\delta}$	8.775	(2.992)			9.279	(3.081)
$plive10_{it_0}$			-0.016	(0.047)	-0.135	(0.065)

Note: Standard errors are clustered at the individual level.

I Neural network details

A neural network is comprised of units arranged into layers: an input layer, hidden layers, and an output layer. The units of the first layers are the inputs or observed variables of the problem. In each subsequent layer, units are a transformation of a linear combination of the units in the previous layer. The weights in the linear combinations are chosen to minimize a loss function.

For a binary outcome p, the output layer has 2 units. Let V_0 and V_1 denote last layer's units pre-transformation (as non-linear functions of the inputs). The transformation at the last layer corresponds to $s_j = \frac{e^{V_j}}{e^{V_0}+e^{V_1}}$, $j = 0, 1$ (softmax activation function for 2 categories) and the loss function (cross-entropy) corresponds to

$$-\sum_{\text{obs}} \big\{ \mathbb{1}(p=0) log(s_0) + \mathbb{1}(p=1) log(s_1) \big\}.$$

Hence, a neural network for a binary outcome is a generalization of a logit with a flexible non-linear index.

This flexibility, however, implies that the optimization problem is non-convex and may have multiple local minima, so a few techniques are usually applied: weight regularization, ensemble of results from multiple starting values, and search of hyperparameters. The algorithm uses gradient descent and back propagation to find the weights in an efficient and fast way.

In this paper, I apply neural networks to panel data

$$\max \sum_{i,t} log\big(\mathbb{P}(p_{it}|x_{it})\big)$$

The inputs x_{it} are state variables in the Bellman equation, Ω_{it-1}. However, some of the

inputs are unobserved latent variables: slope beliefs $(\hat{\delta}_{it-1}, \hat{\sigma}^2_{t-1})$ and heterogeneity in health levels α_i. Conditional on health history h_i^t, these latent variables can be subsumed in time-invariant unobserved $(\alpha_i, \hat{\delta}_{i0}) \equiv \eta_i$, with η_i included in x_{it}.

The objective is to maximize the log likelihood after integrating out this unobserved and time-invariant heterogeneity

$$\max \sum_i log \int \prod_t \mathbb{P}(p_{it}|x_{it}) f(\eta_i) d\eta_i$$

which is a difficult object to work with. Hence, I use instead a key insight from the EM-algorithm, that is,

$$arg\max_\theta \sum_i log \int \mathbb{P}(p_i^T|x_i^T;\theta) f(\eta_i) d\eta_i = arg\max_\theta \sum_i \int log\mathbb{P}(p_i^T|x_i^T;\theta) \underbrace{f(\eta_i|p_i^T;\theta)}_{\substack{\text{unknown} \\ \text{posterior}}} d\eta_i \quad (22)$$

and solve this problem iteratively: given θ_{k-1}

1. Get draws of η_i from the posterior distribution $f(\eta_i|p_i^T;\theta_{k-1})$ by MCMC

2. Estimate θ_k by using a neural network approach in the augmented data

$$\max \sum_i \sum_{\text{draws}} \sum_t log\mathbb{P}(p_{it}|x_{it})$$

corresponding to an approximation of the right-hand side of equation (22)

As mentioned before, this iterative approach is used as a convenient implementation, but given the lack of unique solution to the problem there is no convergence result. As a starting point, I use the posterior distribution $f(\eta_i|h_i^T, plive10_i^T)$, which should already incorporate a substantial amount of information about η_i. The results presented here confirm that intuition: after 5 iterations the results are not qualitatively different.

For estimation, I perform the following steps:

Table I1: Hyperparameter space

Hyperparameter	Space
Depth	$\{3, 5, 8\}$
Width	$\{3, 5, 8\}$
Regularization	$L_1 : \{1e^{-15}, 1e^{-10}, 1e^{-5}\}$
	$L_2 : \{1e^{-15}, 1e^{-10}, 1e^{-5}\}$
Epochs	up to 200

Table I2: Loss and accuracy at 5^{th} iteration across 30 starting points

	Mean	Median	SD
Loss	0.313	0.312	0.003
Accuracy	0.883	0.883	0.0005

1. I split the sample in an estimation and validation sample (80% and 20% of the individuals respectively). Using the estimation sample and one draw from the incomplete prior, I estimate a neural network for several combinations of hyperparameters. Table I1 show the hyperparameter space considered.

2. I choose the hyperparameters of the neural network as those that minimize the predicted loss in the validation sample. The values chosen are depth 3, width 3, L_1 regularization $1e^{-15}$, and 33 epochs.

3. Using that structure, I apply the iterative approach describe earlier 5 times. In each iteration, I average the results across 30 starting points. The loss and accuracy of the last iteration is presented in table I2.

Figure I1: Observed [dashed] versus predicted [solid] probability of work at 5^{th} iteration

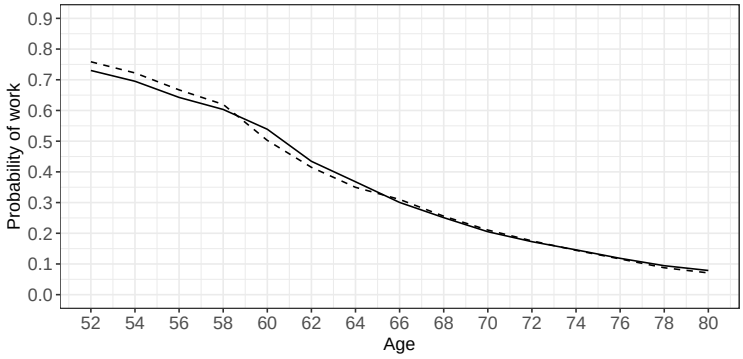

J Additional biomarker results

Besides the collection of a blood sample for measuring biomarkers, the change in collection mode for the selected group also introduced more detailed measures of health, including physical measures, and a saliva sample for DNA analysis. The physical measures include blood pressure and pulse, lung function, hand grip strength, balance test, timed walk test, height, weight, and waist circumference. These variables are valuable measures of health, but I do not include them in this paper given that they are measured only every two waves. Furthermore, their value as signals of health is limited given that, on one hand, they reflect aspects of health already experienced by individuals in their everyday life, and on the other, the results of the measures are immediately communicated to individuals before asking them about their survival expectations.[1]

Distribution of in-person interviews

Table J1: Percentage of in-person interviews by wave and group

wave	group 2	group 1
wave 5	8.7	8.3
wave 6	10.5	10.8
wave 7	75.7	74.2
wave 8	17.8	94.5
wave 9	95.3	21.0

Note: The sample consists of $N = 8,386$ individuals with non-proxy interviews who are at least 50 years old in wave 8, and who give a valid answer to $plive10$ every wave between waves 5 and 9.

[1] In that sense, part of what am denoting interview-mode effect could reflect differences in information given by these results. I expect those effects to be low given that individuals experience most of them in their everyday lives.

Balance test

Table J2: Balance tests at wave 8

		Mean per group		Difference	
	N obs.	$g_i = 2$	$g_i = 1$	coeff.	p-value
female	8,386	.611	.608	.002	.819
age	8,386	68.8	68.6	.177	.302
race: white	8,385	.875	.879	-.004	.543
race: black	8,385	.099	.096	.003	.661
race: other	8,385	.026	.025	.002	.658
hispanic	8,386	.045	.046	-.002	.742
education: less than highschool	8,386	.199	.183	.016	.058
education: highschool	8,386	.336	.345	-.008	.425
education: some college	8,386	.227	.228	-.001	.933
education: college	8,386	.237	.245	-.007	.44
$plive10$	8,386	47.6	48.9	-1.3	.057
number doctor visits	8,145	9.851	10.06	-.208	.571
diagnosis of HBP	8,382	.556	.561	-.005	.668
diagnosis of heart condition	8,381	.241	.234	.007	.437
diagnosis of stroke	8,382	.075	.064	.011	.05
medication for HBP	8,283	.547	.547	0	.967
medication for diabetes (oral meds)	8,335	.141	.143	-.002	.833
medication for diabetes (insulin)	8,335	.038	.039	-.001	.731
medication for cholesterol	8,374	.439	.435	.004	.696
work	8,384	.323	.336	-.013	.197

Note: The sample consists of $N = 8,386$ individuals with non-proxy interviews who are at least 50 years old in wave 8, and who give a valid answer to $plive10$ every wave between waves 5 and 9.

The 1.3 percentage-points difference in survival expectations between the two groups is also captured in table 1.8, and as I mentioned before, I interpret it as caused by differences in interview mode between those two groups, given that no significant differences are found when the interview mode is also similar.

Results distinguishing bad biomarkers results

I use the biomarker results of wave 8 to further distinguish group 1 in two subgroups: individuals whose biomarker results are within normal levels (good results) and those whose results are outside normal levels (bad results). Hence, I estimate the following equation

$$y_{iw} = \beta_0 + \beta_1 d_{g_i} + \beta_2 d_{b_i} + \beta_{3w} d_w + \beta_{4w} d_{g_i} \cdot d_w + \beta_5 d_{g_i} \cdot d_w \cdot d_{b_i} + \epsilon_{iw} \quad (23)$$

where as before, d_{g_i} is a dummy for group 1 (those who get their blood collected in wave 8), and d_w are dummies for waves. The new variable d_{b_i} is a dummy for the subgroup of individuals in group 1 that get bad results in any of the 3 tests. Receiving a bad results corresponds to having a total cholesterol equal or above 240 mg/dL, HDL cholesterol below 40 mg/dL, or blood glucose hbA1c equal or above 6.4%. Note that in this equation, the interpretation of the coefficients is not the same as in equation (1.15). For example, β_1 is now comparing the individuals in group 1 who get good results versus all individuals in group 2, whether or not their (unobserved) test results are good or bad. Thus, β_1 is not a fair comparison. Consequently, the interest in this equation lies not on the comparison between groups 1 and 2, but on comparing the differences between group 1 individuals that receive good versus bad results.

Table J3 presents the results of estimating this equation. The results suggest the information contained on bad test results is at least partially known by individuals themselves, as they have lower survival expectations even before receiving this information, and their labor participation is also decreasing ahead of time.

Table J3: Biomarkers experiment distinguishing bad vs good test results

		Survival expectations $plive10_{iw}$	Working decisions p_{iw}
group 1	d_{g_i}	-0.39	-0.01
group 1, bad results	d_{b_i}	-0.37	0.04**
wave 6	d_{w6}	-1.42***	-0.07***
wave 7	d_{w7}	-1.50***	-0.12***
wave 8	d_{w8}	-6.41***	-0.16***
wave 9	d_{w9}	-3.57***	-0.20***
group 1, wave 6	$d_{g_i} \cdot d_{w6}$	0.58	0.01
group 1, wave 7	$d_{g_i} \cdot d_{w7}$	0.15	0.02*
group 1, wave 8	$d_{g_i} \cdot d_{w8}$	2.23***	0.02*
group 1, wave 9	$d_{g_i} \cdot d_{w9}$	-0.05	0.02
group 1, bad results, wave 6	$d_{g_i} \cdot d_{b_i} \cdot d_{w6}$	-1.25	-0.01
group 1, bad results, wave 7	$d_{g_i} \cdot d_{b_i} \cdot d_{w7}$	-1.75*	-0.04**
group 1, bad results, wave 8	$d_{g_i} \cdot d_{b_i} \cdot d_{w8}$	-1.94*	-0.05***
group 1, bad results, wave 9	$d_{g_i} \cdot d_{b_i} \cdot d_{w9}$	-1.56	-0.03
Constant		53.97***	0.49***
Observations		41,930	41,923
R-squared		0.005	0.021
% of group 1 individuals with bad results		12.29	12.30

Note: Estimation results from equation (23). The sample consists of $N = 8,386$ individuals with non-proxy interviews who are at least 50 years old in wave 8, and who give a valid answer to $plive10_{iw}$ every wave between waves 5 and 9. Of these, 7 observations do not have information on p_{iw}. Standard errors are clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1

Correlation between biomarker results and unobserved slopes δ_i

In order for the biomarker results to be valid signals, they must be correlated with δ_i. I checked this by restimating the equation of health dynamics allowing for the mean heterogeneity to depend on biomarker results. That is, I include binary variables indicating values out of range, averaged across waves. Table J4 show that there is indeed this correlation.

Table J4: MLE results of health dynamics including biomarker information

	Coefficient	Pvalue
ρ	0.189	0.000
γ	0.002	0.057
σ_ϵ	0.264	0.000
μ_α	4.450	0.000
$\nu_{\alpha female}$	-0.128	0.000
$\nu_{\alpha white}$	0.108	0.002
$\nu_{\alpha hispanic}$	0.022	0.640
$\nu_{\alpha less\_HS}$	-0.354	0.000
$\nu_{\alpha cohorte1}$	-0.059	0.040
$\tau_{\alpha Total\_chol}$	0.170	0.000
$\tau_{\alpha HDL}$	-0.030	0.467
$\tau_{\alpha HBP}$	-0.161	0.001
μ_δ	-0.053	0.000
$\nu_{\delta female}$	0.005	0.271
$\nu_{\delta white}$	0.010	0.073
$\nu_{\delta hispanic}$	0.008	0.326
$\nu_{\delta less\_HS}$	0.003	0.579
$\nu_{\delta cohorte1}$	0.005	0.247
$\tau_{\delta Total\_chol}$	-0.005	0.396
$\tau_{\delta HDL}$	-0.010	0.076
$\tau_{\delta HBP}$	-0.028	0.000
σ_α	0.442	0.000
σ_δ	0.040	0.000
ϕ	-0.057	0.336
N observations	7,768	
N individuals	1,344	
-LL	4,223.2	