Mortality Beliefs and Saving Decisions: The Role of Personal Experiences *

Frederik Horn‡

First version: April 3, 2022
This version: June 30, 2023

Abstract

This paper is the first to non-experimentally establish a causal relationship between households' mortality beliefs and subsequent saving and consumption decisions. Motivated by prior literature on the effect of personal experiences on individuals' expectation formation, I exploit the death of a close friend as an exogenous shock to the salience of mortality of a household. Using data from a large household panel, I find that the death of a close friend induces a significant reduction in saving rate of 2.2 percentage points which persist over the following 5 years. I augment the life-cycle model of consumption by the experienced-based learning model and quantify the impact of this personal experience on mortality beliefs. Even though the shock has no material impact on a household's situation, I find a quantitatively large initial reduction in expected survival probability of 7.1 percent.

Keywords: Household finance, Mortality beliefs, Belief formation, Personal experiences, Household saving, Life-cycle model

JEL Codes: D14, D15, G41, G51

*I thank Dániel Kárpáti (Discussant), Christine Laudenbach, Laurent Calvet, and the participants of the 6th SAFE Household Finance Workshop, the participants of the Research in Behavioral Finance Conference (RBFC), Petra Vokata, and the participants of the Annual Meeting of the Swiss Society for Financial Market Research (SGF), Piera Bello (Discussant), Tabea Bucher-Koenen, and the participants of the ZEW Conference on Ageing and Sustainable Finance, Marie-Hélène Broihanne (Discussant) and the participants of the 38th International conference of the French Finance Association, Rawley Heimer, Oliver Spalt, Konrad Stahl, and the University of Mannheim for valuable comments.

‡University of Mannheim, Mannheim Business School, Finance Area, Chair of Financial Institutions, 68161 Mannheim, L9 1-2, E-mail: *frederik.horn@uni-mannheim.de*

1 Introduction

Beliefs are a crucial part of every economic model. In particular, mortality beliefs affect a wide range of economic decisions like insurance choices, healthcare planning, and most notably saving and consumption decisions. Even though the theoretical relationship between longevity expectations and the saving rate is well established, there is little empirical research showing that individuals in fact consider mortality beliefs in their financial decision making. It is difficult to demonstrate a causal link between mortality beliefs and saving decisions both due to endogeneity concerns and a lack of data. Mortality beliefs are typically correlated with the socioeconomic status of an individual, which itself is highly correlated with financial decision making. Similarly, health shocks tend to both entail a lowered life expectancy as well as out-of-pocket expenses. In this paper, I exploit the death of a close friend as a shock to an individual's mortality beliefs. This plausible exogenous shock allows me to causally demonstrate that more pessimistic mortality beliefs translate into lower saving rates.

Recent evidence suggests that personal experiences are an important component of the belief formation process (Malmendier & Nagel, 2011, 2016; D'Acunto et al., 2021). However, it is challenging to quantify the impact of personal experiences on the belief formation process as beliefs are inherently difficult to observe and personal experiences often affect both beliefs as well as the economic situation of a household. I augment the classic life-cycle model of consumption and saving by the experienced-based learning model of Malmendier (2021) to quantify the impact of this non-material personal experience on the belief formation process. The staggered but rare nature of my shock allows me to isolate the impact of one personal experience on a household's economic outcomes from which I can deduce the impact on the belief formation process.

Hence, in my paper I address two questions. Do individuals consider mortality beliefs in their saving decisions? How large is the impact of a personal experience on the belief formation process? To answer these questions, I use a long-running representative panel covering the Australian population to exploit the death of a close friend as an exogenous shock to the mortality beliefs of an individual. The survey covers around 17,000 Australians each year since 2001. The data set is unique in that it collects detailed information on a household's saving and consumption behavior, a plethora of information on the socio-economic status and attitudes of a household as well as whether a close friend died in the previous year. First, I establish a causal relationship between mortality beliefs and saving decisions. Utilizing the death of a close friend as an exogenous shock to mortality beliefs, I find that the shock reduces the saving rate by a 2.2 percentage points. Considering the non-material nature of the shock, the effect size is considerable. Furthermore, this reduction in saving rate persists for the 5 following years. This suggests that it is not driven by a short-term emotional reaction but rather induced by a more long-term change in mortality beliefs. I utilize two self-reported proxies for a household's saving behavior to establish the robustness of the main findings. I find that survey participants report less regular saving habits and a significantly shorter saving horizon following the shock.

On top of that, the data allows me to explicitly link the death of a close friend to a subsequent significant decrease in subjective longevity expectations reported by the households. Furthermore, I strengthen this link by establishing that the effect on the saving rate is driven by households with a weak bequest motive. These analyses demonstrates that the exogenous shock works through the intended channel of more pessimistic mortality beliefs. The data allows me to break the effect on the saving rate down into consumption subcategories. This analysis reveals that the reduction in saving rate is not caused by increased concerns about one's own health as health expenditure is barely affected. On the contrary, consumption of leisure related items like alcohol or meals eaten out increases the most. Moreover, the results are not driven by bequests of the deceased friend, drastic life changes, or reductions in income.

Second, I use the life-cycle model of consumption and saving augmented by the experienced-based learning model of Malmendier (2021) to derive two unique predictions which I test empirically. On the one hand, the agent's age crucially determines how strongly she should be affected by the shock. Each new experience makes up a larger proportion of the set of relevant experiences for younger agents and thereby they are more strongly affected by them. Indeed, I find that younger individuals reduce their saving rate three times more than older individuals (3.5 versus 1.2 percentage points). On the other hand, the canonical life-cycle model predicts that the agent's reaction to the shock crucially depends on her risk-aversion. Intuitively, more risk-averse agents should react less to an increase in longevity risk. I find that more risk-loving households reduce their saving rate by 3.2 percentage points whereas more risk-averse households only lower their saving rate by 1.2 percentage points. These results suggest that the experience-based learning model in the context of the life-cycle model helps to understand how personal experiences are incorporated into the belief formation process.

Third, I quantify the impact of the shock on mortality beliefs in the context of the canonical life-cycle model of consumption. For that purpose, I use the augmented life-cycle model to structurally estimate both the impact of the personal experience on mortality beliefs as well as the parameters that govern how fast the shock fades out of the set of relevant experiences. I find that depending on an agent's risk aversion the death of a close friend leads to a reduction in expected probability of surviving to the next period of 1.1 percent to 13.8 percent. This reduction in expected survival probability slowly attenuates to zero over the following 6 years. The magnitude of the effect is quantitatively large considering the non-material nature of the death of a close friend. On top of that, I estimate that the parameter λ that governs how fast the experience fades out of memory ranges from 1.3 to 1.7. This is in line with estimates of Malmendier and Nagel (2011) who find estimates ranging from 1.4 to 1.9 in the vastly different domain of stock returns.

Overall, these results establish a causal link between mortality beliefs and households' saving decisions. An exogenous shock to mortality beliefs induces a significant reduction in saving behavior. I provide evidence that experience-based learning has a quantitatively large impact on the belief formation process. Moreover, my results suggest that the shape of the weighting function governing how fast experiences fade out of memory is similar across domains.

My paper adds to the academic literature exploring the effect of mortality beliefs on saving and investment decisions. This literature goes back to Hamermesh (1985) who elicits subjective survival probabilities and discusses the implications for household saving models. Since then, several papers attempt to link mortality beliefs to saving decisions (Hurd et al., 2004; Puri & Robinson, 2007; De Nardi et al., 2010; Post & Hanewald, 2013; Jarnebrant & Myrseth, 2013; Spaenjers & Spira, 2015). In particular, Spaenjers and Spira (2015) try to rule out endogeneity concerns by instrumenting an individual's subjective survival probabilities with the death of their parents. My paper goes a step further by removing associations of hereditary illnesses and bequest issues from the equation. The death of a close friend should not be correlated with ones own genetic history as well as should not result in significant windfall gains due to bequests. Furthermore, most of the aforementioned papers utilize health and retirement studies and therefore focus on older individuals. Conversely, my paper covers a representative sample of the Australian population, which includes households at all stages of life. This facilitates the generalizability of my results and provides additional insights into younger households for whose lifetime utility these financial decisions matter the most.[1]

More broadly, I contribute to the literature investigating the role of personal experiences in financial decision making and expectation formation. In general, these studies find that individuals overweight personal experiences in the expectation formation process. This has been shown in a variety of contexts like IPOs (Kaustia & Knüpfer, 2008), investments in 401(k)s (Choi et al., 2009), financial risk taking (Malmendier & Nagel, 2011; Knüpfer et al., 2017; Bernile et al., 2017), inflation expectations (Malmendier & Nagel, 2016), household leverage (Kalda, 2020), house price expectations (Kuchler & Zafar, 2019; Bailey et al., 2018), and unemployment rate expectations (Kuchler & Zafar, 2019). My paper adds to this literature by demonstrating that personal experiences also play an important role for the belief formation process in the domain of mortality. Furthermore, I am able to quantify the impact of one personal experience on beliefs. Thus, I gauge the importance of personal experiences for financial outcomes beyond purely demonstrating a link.

Finally, this paper is closely related to the seminal work by Heimer, Myrseth, and Schoenle (2019). They argue that young individuals underestimate survival whereas older individuals overestimate survival. The authors hypothesize that younger individuals overweight rare events due to them being salient. Hence, the salience of death distorts mortality beliefs and subsequently crucially affects optimal household decision-making. My paper contributes direct evidence that

[1]There is also recent concurrent work by Kárpáti (2022) who exploits genetic testing to establish a causal link between objective mortality beliefs and a wide range of financial outcomes in a representative Dutch dataset.

mortality salience affects mortality beliefs and thereby financial decision-making. Furthermore, my findings might provide a possible link between personal experiences and the overweighting of rare events for the young. Younger individuals are more likely in relative terms to die due to such rare events. Hence, their friends learn about these events and subsequently overweight the likelihood of such an event happening to themselves.

This paper is structured as follows. Section 2 outlines the canonical life-cycle model and derives the importance of survival probabilities in that context. Furthermore, I adapt the experience-based learning model and demonstrate how the personal experience affects mortality beliefs over time. Section 3 describes the data and introduces the identification strategy. Section 4 presents and discusses the main empirical results of my paper. In section 5, I structurally estimate the impact of the shock on mortality beliefs. Finally, section 6 shows robustness checks and section 7 concludes.

2 Theoretical Framework

2.1 Life-cycle Consumption Model

I set up a classic life-cycle model to demonstrate the importance of mortality expectations for the consumption and saving decision (e.g. Deaton, 1991; Hubbard et al., 1995). For details regarding the model setup refer to appendix B1. In the model, a representative household maximizes its expected lifetime utility. The household receives stochastic labor income each period and decides how much to allocate to consumption and the remainder is allocated to saving. I assume that there is only one asset with a risk-free rate of R. Furthermore, each household lives a maximum of T periods and is assumed to exhibit a power utility function. This gives rise to the following maximization problem:

$$\max \ \mathbb{E}[\sum_{t=1}^{T} \beta^{t-1} (\prod_{j=0}^{t-2} \mathbb{E}(s_j)) u(c_t)] \tag{1}$$

where c_t is a household's consumption, β a time discount factor, and $\mathbb{E}(s_j)$ the expected probability of survival to period j. One can rewrite this problem in recursive form as a Bellman equation:

$$\nu_t(m_t) = \max_{c_t} \ u(c_t) + \beta \mathbb{E}(s_{t+1}) \mathbb{E}[(p_{t+1}/p_t)^{1-\rho} \nu_{t+1}(m_{t+1})] \tag{2}$$

with:

$$m_{t+1} = m_t - c_t + Y_{t+1} \tag{3}$$

where m_{t+1} is the available resources next period that could be potentially used for consumption and Y_{t+1} is next period's labor income. Furthermore, p_t is the permanent labor income in period t, and ρ is the coefficient of relative risk aversion of a power utility function. Taking the derivative gives rise to the following first order condition:

$$0 = u^{'}(c_t) - \beta \mathbb{E}(s_{t+1}) \mathbb{E}[R(p_{t+1}/p_t)^{-\rho} \nu_{t+1}(m_{t+1})] \tag{4}$$

Solving for c_t yields the following optimal consumption in t:

$$c_t^* = (\beta \mathbb{E}(s_{t+1}))^{-1/\rho} (\mathbb{E}[\cdot])^{-1/\rho} \tag{5}$$

Even though there is no analytical solution to this problem, it is straightforward to see from the optimal consumption equation that a decrease in survival probability leads to an increase in consumption and thereby to a reduction in the savings rate. In this paper, I argue that the death of a close friend increases the salience of death for an individual. Subsequently, she becomes more pessimistic about her mortality beliefs, resulting in a lower survival rate s_{t+1}. Thus, c_t^* increases and mechanically the saving rate decreases. Intuitively, the agent does not defer her consumption as much if there is a certain probability that she will not survive to the next period.

Largely following Cocco, Gomes, and Maenhout (2005), I calibrate this model to the Australian panel. For illustrative purposes, I solve it numerically for (1) survival rates taken from the Australian Bureau of Statistics and (2) agents that hold 5 percent more pessimistic *expected* survival probabilities relative to the objective survival probabilities.

[Insert Figure 1 about here.]

Figure 1 shows from upper left to lower right the survival probabilities, average consumption, average saving rate, and wealth accumulation of the simulated households over the life-cycle. The black line displays the results for the simulation with the objective survival probabilities, and the red line displays the agents with pessimistic expectations about their survival probabilities. The upper right panel demonstrates that pessimistic mortality beliefs result in overconsumption in younger years. However, at around age 50 the agents with distorted beliefs are starting to underconsume as their previous saving rate does not lead to a sufficient capital stock to comfortably smooth consumption in later years. The lower right panel clearly shows that the pessimistic agents accumulate a lot less wealth over their lifespan which results in a reduced consumption in retirement.

In conclusion, mortality beliefs clearly have important implications for an agent's saving behavior in the context of a life-cycle model. An agent who is more pessimistic about her survival has an unambiguously lower saving rate, all else equal. However, there is little empirical evidence that causally links mortality beliefs to saving decisions. This paper addresses the gap. In the next part, I propose how a shock to mortality beliefs induced by the death of a close friend translates into a change in survival rates in the context of an experienced-based learning model.

2.2 Mortality Belief Formation

I adapt the experience-based learning model of Malmendier, Pouzo, and Vanasco (2020) to put a more rigorous structure on how the death of a close friend affects an agent's mortality beliefs. The agent experiences the death of a close friend which translates into a negative shock to her mortality beliefs. In the context of the life-cycle model, this means a reduction in the expected survival rate in that period. In each period, the agent weighs these past periods depending on how long ago they have occured and forms the expectation about her survival rate for the current period. I argue that expectations about the probability of surviving to the next period are given by the following equation:

$$\mathbb{E}_t(s_{t+1}) = \Gamma_t(X, a) + \sum_{k=0}^{t} w(\lambda, k, t) M_{t-k} + \epsilon_t \tag{6}$$

where Γ_t is the baseline probability of surviving to the next period for an individual at age a with a vector of personal characteristics X. These personal characteristics could include whether she is a smoker, has a chronic health condition, or is working in an unsafe occupation. $w(\lambda, k, t)$ is the weight the agent assigns to the personal experience M that occurred k years before year t and λ governs the shape of that weighting function. ϵ_t is the idiosyncratic error of an individual when forming expectations which is normally distributed with mean zero. I use the weighting function proposed by Malmendier et al. (2020):

$$w(k, \lambda, t) = \frac{(t+1-k)^{\lambda}}{\sum_{k\prime=0}^{t}(t+1-k')^{\lambda}} \tag{7}$$

where w is the weight an agent at t assigns to the personal experience experienced k periods ago. The parameter λ determines the weight of more recent compared to less recent experiences. As agents rarely experience the death of a close friend, mortality beliefs will become gradually more optimistic after the initial negative shock as long as $\lambda > 0$. Hence, one should observe an initially strong drop in the saving rate which is attenuated in the following periods.

3 Data and Methodology

3.1 Data

I employ data from the Household, Income and Labour Dynamics in Australia (HILDA) survey. HILDA is a household panel study surveying around 17,000 Australians each year beginning in 2001. Table 1 shows summary statistics for a variety of variables of interest. As the HILDA panel aims to survey a representative sample of the Australian population, it is not surprising that the sample consists of 50 percent women, the average age lies around 37, and the average income equals

75,426 Australian dollar with the median only being roughly 60,000 Australian dollar.

[Insert Table 1 about here.]

My main dependant variable is an individual's saving behavior. I use three measures to elicit an individual's savings decision. First, I directly calculate the savings rate from the consumption stated by households in the survey. Beginning with the fifth wave of the panel, individuals are asked about their annual expenditure covering a wide range of items[2]. These items include for example groceries, alcohol, clothing, pharmaceuticals, and many others. For a comprehensive list refer to Table 6 in the Appendix. Following Dynan, Skinner, and Zeldes (2004), I calculate the saving rate as one minus the sum all of these expenditures divided by the total after-tax income reported by the household. Furthermore, I exclude any household-year observation where the household received any windfall payments to ensure that the results are not driven by received inheritances. Finally, I winsorize at the 3 percent level to ensure that outliers are not driving the results. Yet, the results do not depend on the winsorized percentage.

The average saving rate is 54 percent, which is significantly higher than official statistics by the Australian Bureau of Statistics. This is due to consumption elicited by the panel only covers non-durable consumption and even there might not comprehensively cover all subareas. However, there is little reason to believe that my calculated savings rate is systematically biased across individuals. Figure 2 shows the average saving rate by age. The graph displays the typical hump-shaped age profile (e.g. Guvenen, 2007; Aguiar & Hurst, 2013) which suggests that the aggregated consumption represents a sensible proxy for a household's saving rate.

Second, participants are asked *"Which of the following statements comes closest to describing your (and your family's) savings habits?"*. The predefined answer range from *"don't save: usually spend more than income"* to *"save regularly by putting money aside each month"*. Third, participants are asked about their saving horizon with possible answers ranging from *"the next week"* to *"more than 10 years ahead"*.

[Insert Figure 2 about here.]

My main independent variable of interest is the death of a close friend dummy. It equals 1 if the participant states that a close friend died within the last 12 months before the survey. Unconditionally, 11 percent of respondents experienced such an event in the previous year. This seemingly large percentage is in line with the percentage elicited by the Australian Bureau of Statistics for the General Social Survey (Liu et al., 2019). The perceived life expectancy is measured by the question *"How likely do you think it is that you will live to be 75 or more?"* where people aged older than 65 are asked how likely it is that they live for 10 more years. The answers range

[2] If several members of the household provided answers, the responses were averaged by HILDA.

from *"Very likely"* to *"Very unlikely"* on a four point ordinal scale. On average, individuals are optimistic about their life expectancy with around 45 percent of respondents stating that it is very likely that they will live to 75. Only around 12 percent of individuals respond that it is unlikely or very unlikely that they are going to live to 75. Furthermore, I elicit an individual's risk aversion using the question *"Are you generally a person who is willing to take risks or are you unwilling to take risks?"*. The answers range from 0 to 10 where I rescale the answers such that a higher value indicates a higher level of risk aversion. On average, the distribution is centered around the value of 5 with a standard deviation of around 2.5.

For all regressions on household level, I exclude households where it is likely that financial decision making is done independently by household members, but the consumption behavior is still aggregated on household level. These include for example siblings living together or shared flats. If there is a couple living in the household, I require both partners to report the death of a close friend as the financial decision-making is not easily attributable to one of the two. Next, I describe the identification strategy I employ in this paper.

3.2 Identification

My identification strategy is based on the idea that the death of a close friend represents an exogenous negative shock to an individual's mortality beliefs. This is rooted in the literature on how personal experiences affect an individual's beliefs in a wide range of economic contexts (e.g. Malmendier & Nagel, 2011; Kuchler & Zafar, 2019). At the same time, using the death of a close friend as a shock to the mortality beliefs of an individual has two advantages over previous attempts that utilize the death of a family member (e.g. Spaenjers & Spira, 2015). First, the death of a non-relative should not affect the financial situation of an individual. It is rare that a deceased individual bequests a meaningful amount of wealth to a friend rather than her family members. Second, the death of parents or siblings often contains information about an individual's own hereditary health risks. Hence, the effect should not be driven by a response to a signal about one's own health. It could be argued that the death of a close friend represents a signal about the health consequences about an individual's own lifestyle. However, I will show in later parts that the effect is most pronounced for demographics where this is highly unlikely.

Furthermore, using panel data allows me to abstract from personal characteristics that have been shown to affect the financial decision making of an individual like income (Imbens et al., 2001; Dynan et al., 2004) or financial literacy (Calvet et al., 2007; Van Rooij et al., 2011). Thus, I estimate the staggered differences-in-differences models both for the average effect and for event studies. For the average effect I use the following regression model:

$$S_{it} = \beta FD_{i,t} + \gamma_t + \delta_i + \epsilon_{it} \tag{8}$$

where S_{it} represents the saving rate of either an individual or a household depending on the respective unit of observation in year t. FD is an indicator variable equal to one for each year after the death of a close friend was reported. For couples, this indicator variable turns one as soon as both partners reported the death of a close friend. Finally, γ_t are age fixed effects and δ_i either person or household fixed effects. Hence, the β captures the average change in saving rate of treated households after the shock compared to untreated households. Furthermore, I also explore the dynamics around the shock to test for pretrends and to better understand the reaction over the following years. Hence, I estimate the following regression model:

$$S_{it} = \sum_{k=-5}^{k=5} \beta_k FD_{i,k} + \gamma_t + \delta_i + \epsilon_{it} \tag{9}$$

where $FD_{i,k}$ are time dummies relative to the death of a close friend ranging from 5 years before to 5 years after. Hence, β_k captures the change in saving rate of treated households in the years around the event compared to untreated households.

4 Empirical Results

4.1 Impact of the Shock on Saving Behavior

First, I establish that the exogenous shock to mortality beliefs indeed has an impact on the saving behavior of a household. Column 1 of table 2 reports the results of regressing the household's saving rate on a indicator variable equal to one in all periods following the death of a close friend. All regressions include both household as well as age fixed effects. Furthermore, I cluster standard errors on household level to account for auto-correlation (Bertrand et al., 2004). I find that the death of a close friend reduces the saving rate on average by 2.2 percentage points. This effect is highly significant at the 1 percent level. This result suggests that the death of a close friend induces more pessimistic mortality beliefs which translate into a lower saving rate.

[Insert Table 2 around here.]

Furthermore, I explore the saving rate dynamics around the shock. Columns 2 exhibits the results of regressing the saving rate on 5 pre-treatment dummies and 5 post-treatment dummies. Figure 3 visualizes the regression results. Prior to the shock, there is no significant pretrend observable. However, as soon as the death of a close friend occurs households instantly reduce their saving rate by around 2 percentage points. Over the following 5 years, this effect attenuates to 1 percentage point. One potential concern could be that the death of a close friend induces a strong emotional reaction which results in an immediate increase in expenditure to distract oneself from the event. This could lead to a mechanical short-term increase in expenditure which is not caused by more pessimistic mortality beliefs. However, this concern becomes highly unlikely given that there is a persistent long-term reaction to the shock observable over the following 5 years.

To address potential concerns associated with staggered differences-in-differences estimators as raised by Baker, Larcker, and Wang (2022), I implement the estimator proposed by Sun and Abraham (2021) and the stacked regression estimator as in Cengiz, Dube, Lindner, and Zipperer (2019). These estimators only include never-treated or last-treated households in the control group and thereby create a "clean" control group. Columns 3 and 4 demonstrate that the results of the alternative estimators barely deviate from the OLS estimates. Again, the initial reduction in saving rate is around 1.9 percentage points which is highly significant at the 1 percent level.

[Insert Figure 3 around here.]

Furthermore, I exploit two additional proxies for a household's saving behavior to establish that the shock induces a reduced inclination to save. I regress the *Saving Habit* and *Saving Horizon* variables on an indicator variable equal to one if the death of a close friend was reported in that period. I conduct the analyses on the level of an individual as the survey elicits these variables at this aggregation level. Crucial for these regressions is the timing of the death of a friend dummy. When I regress saving habit on the death of a friend dummy, I lag the variable as saving habit represents a backward looking persistent behavior. Thus, I avoid that the event, namely the death of a friend, and the self-reported saving behavior overlap. Conversely, the saving horizon is a forward looking variable describing future behavior. Hence, there is no need to lag the death of a friend dummy as the shock to the salience of death and the described behavior are sufficiently separated.

[Insert Table 3 around here.]

Columns 1 and 3 of table 3 show that the shock both reduces the self-reported saving habit as well as the individual's saving horizon. Yet, the impact on the latter is not statistically significant at the 10 percent level. This is not surprising as older individuals are not likely to adjust their saving horizon as they approach death. Hence, in columns 2 and 4 I repeat the analysis for working age individuals. Indeed, the shock induces a statistically significant reduction in the reported saving horizon of the younger subsample. Overall, these additional results strengthen the argument that the death of a close friend represents an exogenous negative shock to an individuals mortality beliefs which results in a lower saving rate. Especially, the finding that individuals significantly reduce their saving horizon suggests that they hold more pessimistic mortality beliefs.

In conclusion, these findings suggests that the death of a close friend represents a negative exogenous shock to mortality beliefs and that a shift in mortality beliefs has an impact on saving behavior. Yet, at this point it is not possible to definitely conclude that the shock works through the intended channel of mortality beliefs. Hence, in the next sections I exclude possible alternative channels and directly link the shock to a reduction in mortality beliefs.

4.2 Expenditure Subcategories

One possible explanation for the reduction in saving rate could be that the shock prompts individuals to be concerned about their own health which would result in increased health care spending. However, my data allows me to test for this concern explicitly. Thus, I explore which components of consumption increase most following the shock. I cluster the various consumption subcategories elicited by the HILDA survey into three groups: leisure related expenditure, expenditure on necessities, and health and insurance related expenditure. For details refer to table 6.

[Insert Table 4 around here.]

Table 4 reports the results of regressing the saving rate as well as the expenditure categories on the friend of a death indicator variable. Columns 2 indicates that following the shock the expenditure on leisure related items increases by 0.6 percentage points which is highly significant at the 1 percent level. Similarly, columns 3 and 4 show that the shock increases expenditure on necessities and health related items by 1.2 and 0.2 percentage points, respectively. Both coefficients are highly significant at the 1 percent level. Moreover, the table reports the percentage of each of these expenditure categories of overall expenditure. Relating the regression coefficients to the unconditional expenditure percentage reveals that expenditure on leisure is affected the most by the shock as it increases by 2.8 percent relative to the baseline. Conversely, the expenditure on healthcare related items is affected the least as it increases only by 1.6 percent relative to the baseline.

Overall, these findings indicate that the reaction to the shock is not driven by households massively increasing their expenditure on health related items. Treated households rather increase their consumption of leisure related items like cigarettes, alcohol, and meals eaten out. These are consumption items that tend to be detrimental to one's health. Hence, it is unlikely that concerns about one's health induced by the shock are responsible for the large reduction in saving rate.

4.3 Mortality Beliefs

The necessary condition for the death of a close friend being a plausible shock is that it in fact has a negative impact on mortality beliefs. The HILDA panel allows me to explicitly test for this link. I utilize the question *"How likely is it that you are going to live to 75?"*. The question is asked only three times with each being 4 years apart. Yet, it is possible to conduct some basic analyses to demonstrate that the death of a close friend actually affects an individual's life expectancy. Furthermore, I can replicate the finding of previous papers that mortality beliefs have a strong impact on saving decisions (e.g. Heimer et al., 2019). Figure 4 plots the distribution of answers to the life expectancy question by age bins. Overall, individuals are optimistic about their survival probability until the age of 75. This is justified as 75 is significantly lower than the current life expectancy in Australia. Comparing the distribution of answers for the 20 to 35 year old with the answers of the 45 to 60 year old might provide some evidence for a similar pattern as reported by Heimer et al. (2019). Younger individuals also appear to be slightly pessimistic about their survival rates compared to their older counterparts. Conversely, the above 75 year old individuals might be slightly optimistic about their survival as a significant portion is reporting that it is "Very Likely" or "Likely" to live to 75. Yet, the exact interpretation of the findings depends on the perception of the question by participants.

[Insert Figure 4 about here.]

Columns 1 and 2 of table 7 display the results of regressing the likelihood to live to 75 on the death of a close friend either in the same period or in the previous period. I run OLS regressions with individual and age fixed effects. Standard errors are clustered at the individual level. Thus, I elicit the within person change in stated survival probability due to the exogenous shock. Column 1 shows that the death of a close friend has a significant impact on an individual's mortality beliefs. On average, the shock reduces the stated likelihood to live to 75 category by 0.027. This coefficient is statistically significant at the 5 percent level. In addition, column 2 indicates that there is still a negative impact on next period's stated life expectancy. However, the effect size is halved and the statistical significance is low. Yet, considering the limited power of these tests due to the small sample size and the inclusion of individual fixed effects the reaction is considerable. Overall, this analysis demonstrates that such a shock to the salience of death has a significant negative effect on life expectancy. These findings provide further evidence that the previous results that a friend's death translates into less saving and more consumption is driven by changes in mortality beliefs.

[Insert Table 7 about here.]

Next, I establish that mortality beliefs have a significant impact on saving behavior. Previous literature suggests that mortality beliefs are correlated with the saving rate (e.g. Post & Hanewald, 2013). The challenge with these results is that both mortality beliefs and saving rate are strongly correlated with observable and unobservable factors like income, health, and financial literacy. I go one step further by including person and age fixed effects when regressing the saving rate on life expectancy. Thus, I explore the within person change in saving behavior following a change in mortality beliefs. Columns 3 and 4 of table 7 exhibit the results of regressing the saving rate on the likelihood to live to 75 variable. On average, going from one category to a higher category increases the saving rate by 0.5 percentage points. This is statistically significant at the 5 percent level. Similarly, a positive change in the previous period increases next period's saving rate by 0.5 percentage points as well. This coefficient is still statistically significant at the 10 percent level. Yet, this is not conclusive evidence that mortality beliefs causally affect saving behavior. It would be for example possible that an individual falls ill which both affects mortality beliefs negatively and might induce increased spending on health care related expenditure. This is the reason I exploit in the previous section the exogenous shock to mortality beliefs induced by the death of a close friend.

An agent's bequest motive should play a significant role in her saving decision if indeed the death of a close friend represents a negative shock to mortality beliefs. If an agent considers bequests to be a part of her utility function, the reduction in saving rate in response to the shock should be less pronounced. Thus, I proxy for the bequest motive with the parenthood status of households.

[Insert Table 5 about here.]

Table 5 shows the results of regressing the saving rate on the death of a close friend indicator variable depending on whether households have children. Columns 1 and 2 demonstrate that childless households reduce, on average, their saving rate by 4.7 percentage points which is highly statistically significant at the 1 percent level. Conversely, parents reduce their saving rate, on average, by only 1.5 percentage point which is less than half of the effect size of childless individuals. This disparity indicates that households consider bequest motives in their response to a close friend dying which suggests that mortality beliefs are negatively affected by the shock. Yet, the reduced effect size might be caused by parents having less leeway in financial matters as they have to provide for their children. Hence, columns 3 and 4 present the findings for the sample of parents depending on whether their child is still part of the household or not. Indeed, parents having their child living with them do not react to the shock. Households that do not having a child living with them reduce the saving rate by 1.5 percentage points. This effect is statistically significant at the 1 percent level. However, the coefficient is half the coefficient of the childless households whereas childless households only have a 10 percent higher saving rate. Hence, households seem to consider bequests when confronted with the death of a close friend even though the effect on the saving rate is not fully mitigated by having a child to bequeath to.

In conclusion, the findings demonstrate that the shock works through the intended channel. Consistent with the literature on the effect of personal experiences on expectation formation (e.g. Malmendier & Nagel, 2016; Kuchler & Zafar, 2019), the agent overweights the likelihood of the rare event happening due to its salience. Thus, she irrationally forms too pessimistic mortality expectations which in turn translate into a lower saving rate. In the next section, I test further predictions that arise from the experience-based learning model.

4.4 Additional Predictions of the Model

After establishing a significant link between mortality beliefs and saving decisions, I turn to the question in which way the salience of death affects mortality beliefs and subsequently saving decisions. The model introduced in section 2 allows me to test two predictions how the shock to mortality beliefs should affect the saving rate. First, younger individuals should be more strongly affected by the shock than older individuals. Second, the life cycle consumption model predicts a stronger impact of mortality beliefs for less risk-averse individuals.

4.4.1 The Role of Age

Following the argument of Malmendier (2021), the experience of the death of a close friend should have a more pronounced effect on the beliefs of younger individuals. Intuitively, younger individuals have experienced less relevant events such that a new event constitutes a larger weight in their set of experiences and thereby in their expectation formation process. Subsequently, the change in saving behavior should be more drastic for younger individuals. Furthermore, younger individuals on average have younger friends. The causes of death of younger individuals tend to be suicides, crimes, and accidents (c.f. Online Appendix) which cannot be anticipated. This should result in a more sharp updating of beliefs.

Hence, I split the sample along the median adult age of 50 and regress the saving rate on the death of a close friend indicator variable for each of the subsamples separately. Columns 1 of table 8 exhibits the results for the younger households whereas columns 2 display the results for the older households. Columns 1 and 2 reveal that the shock reduces the saving rate of older households by only 1.2 percentage points whereas the impact on younger households is three times as large at 3.5 percentage points. The coefficients are statistically significant at the 5 percent and 1 percent level, respectively.

[Insert Table 8 about here.]

In conclusion, these findings are consistent with two not necessarily mutually exclusive explanations. On the one hand, the shock represents a larger part of younger individuals' set of experience. On the other hand, the shock is more surprising for younger individuals as their friends tend to be younger and experience non-natural causes of death. Hence, the shock would induce a stronger emotional reaction. Yet, both explanations would be consistent with the predictions of the experience-based learning model of Malmendier (2021).

4.4.2 Risk Aversion

As described earlier, the optimal consumption in period t is given by:

$$c_t^* = (\beta s_{t+1})^{-1/\rho}(\mathbb{E}[\cdot])^{-1/\rho} \tag{4}$$

One parameter that crucially determines the size of the effect of a shock to mortality beliefs on consumption is the risk aversion ρ. Everything else equal, households with lower risk aversion should increase their consumption more. Intuitively, high risk aversion households react less to the increased uncertainty surrounding their own survival. I use the question *"On a scale from 0 to 10, are you generally a person who is willing to take risks or are you unwilling to take risks?"* to elicit an individual's risk aversion. Next, I rescale the variable such that a high value indicates a high level of risk aversion. Finally, I split the sample into a high and a low risk aversion group. For each of these groups I separately run fixed effects regressions eliciting the long-term impact of a friend's death on a household's saving decisions.

Column 3 of table 8 shows that the high risk aversion households reduce their saving rate in response to the shock by 1.2 percentage points which is only statistically significant at the 10 percent level. Conversely, column 4 reveals that the low risk aversion households reduce their saving rate about three times as much by 3.2 percentage points which is highly significant at the 1 percent level. Overall, these findings are consistent with the predictions of the life-cycle model. High risk aversion households react less strongly to the increase in survival risk compared to low risk-aversion households.

5 Structural Estimation

In the final part of this paper, I structurally estimate the reduction in expected survival rate implied by the saving rate response and the parameter λ that governs how fast the personal experience fades out of the set of experiences relevant for the belief formation process.

5.1 Empirical Saving Rate Response

In a first step, I revisit the dynamics of the reduction in saving rate around the death of a close friend. One issue with the previous estimation of the dynamics around the shock might be that the post event period is contaminated by further shocks like another death of a close friend, or entering or exiting the work force. Moreover, I require the reduction in saving rate following the shock compared to the average previous saving rate of a household rather than compared to untreated households. Hence, I create a sample of treated households that are between 25 and 65 years old. In case that a household experiences several shocks in close temporal proximity, I reset, in the spirit of the EBL model, the event time to zero. The new shock makes the issue salient again. On top of that, I require that at least the first 5 years after the shock are not missing. Then, I estimate the following regression model for this sample:

$$S_{it} = \sum_{k=-6}^{k=7} \beta_k FD_{i,k} + \gamma_t \times \tau_t + \epsilon_{it} \tag{10}$$

where $FD_{i,k}$ is an indicator variable equal to one in period k relative to the death of a close friend, γ_t are age fixed effects, and τ_t are year fixed effects. I include age times year fixed effects to average out age and cohort effects. Importantly, this estimation differs from table 2 as it does not compare the reduction in saving rate of the treated households to the untreated households. In this regression, I compare the reduction in saving rate around the shock to the saving rate outside of the event window.

[Insert Figure 5 about here.]

Figure 5 displays the effect decay after the shock. The households strongly reduce their saving rate following the shock. This initial reaction attenuates back to zero over the following 6 years. This result is in line with the experience-based learning model which predicts that the personal experience gets less weight in the belief formation process as it moves farther into the past. Intuitively, the experience fades out of memory. In the next section, I use these reductions in saving rate to back out the model implied associated reduction in expected survival probability. Based on these changes in expected survival probabilities over the event window, I estimate the decay parameter λ which governs the shape of the weighting function in the belief formation process.

5.2 Estimation Problem

There are two parameters of interest I cannot observe in the data: the actual reduction in expected survival rate induced by the shock and the decay parameter λ. In a first step, I estimate the implied reduction in survival rate associated with the estimated coefficients in figure 5. I can back out the implied drop in expected survival rate consistent with the observed impact on the saving rate using the model set up in part 2. Hence, I minimize the absolute difference between the relative reduction in saving rate estimated in that figure and the relative reduction in saving rate given a reduction in survival rate in the life-cycle model simulations.

$$\min_{\Delta s_{e+1}} \mid \Delta S_e(\Delta s_{e+1}) - \Delta \hat{S}_e \mid \tag{11}$$

where $\Delta \hat{S}_e$ is the relative reduction in saving rate estimated in table **??** for event time e and $\Delta S_e(\Delta s_{e+1})$ is the relative reduction in saving rate given the reduction in expected survival rate Δs_{e+1} implied by model simulations, where s_{e+1} is the subjective probability of surviving to period $e + 1$.

The coefficients of figure 5 represent the average reduction in saving rate following the death of a close friend across the sample. Moreover, these coefficients are net of age and cohort effects as the regression model includes age times year fixed effects. Hence, when simulating the shock to expected survival probabilities in the life-cycle model I assign it to the age of 49 which is roughly the average age at which the death of a close friend occurs in my sample. One assumption I have to make for this analysis concerns the agents' expectations about the survival probability before the shock. I assume that previous to the shock all agents hold objective mortality beliefs. That means they act according to the survival rates taken from the Australian Bureau of Statistics. This is reasonable as previous research has shown that, on average, individual's longevity expectations are in line with actual survival patterns (Smith et al., 2001).

Second, I estimate the decay parameter given the implied reduction in survival rates. This is possible by recognizing that the change in expected survival rate is given by:

$$\Delta E[s_t] = \left[\Gamma_t(a, X) + \sum_{k=0}^{t} w(\lambda, k, t) M_{t-k}\right] - \left[\Gamma_{t-1}(a, X) + \sum_{k=0}^{t-1} w(\lambda, k, t-1) M_{t-k}\right] \tag{12}$$

It is crucial to recognize that in the first period following the shock the new experience receives a weight of 1 in the set of experiences as it is the only relevant experience in this domain. Moreover, by the construction of my sample, the agents do not experience further shocks in all following periods. Hence, all following M are equal to zero:

$$\Delta E[s_t] = \Gamma_t(a, X) - \Gamma_{t-1}(a, X) + (w(\lambda, k, t) - w(\lambda, k, t-1)) M_{t=k} \tag{13}$$

The change in baseline survival probability $\Gamma_t(a, X) - \Gamma_{t-1}(a, X)$ is close to zero from one period to the next. Hence, I am left with:

$$\Delta E[s_t] = (w(\lambda, k, t) - w(\lambda, k, t-1)) M_{t=k} \tag{14}$$

where $M_{t=k}$ is the initial reduction in expected survival rate following the shock. Thus, the change in weights is just equal to the change in survival rate divided by the initial reduction in expected survival rate. Given that I estimate the implied reduction in expected survival rate in the first step of the estimation procedure and the initial weight of the experience is equal to 1, it is straightforward to calculate the weights implied by the empirical reduction in expected survival probability. Finally, this allows me to estimate the decay parameter λ that minimizes the squared difference between the implied weights by the empirical results and the theoretical weights:

$$\min_{\lambda} \; (\mathbf{w}(t, \lambda, e) - \hat{\mathbf{w}}(t, e))'(\mathbf{w}(t, \lambda, e) - \hat{\mathbf{w}}(t, e)) \quad \forall\, t = e \in [0, 7] \tag{15}$$

where $\hat{\mathbf{w}}$ is the vector of weights of the t-periods ago event from the relative reduction in Δs_{e+1} estimated from formula (11) and $\mathbf{w}$ is the vector of weights implied by the above formula for a given λ. For details regarding the exact estimation process, please refer to appendix B3.

5.3 Results

Table 9 shows the reduction in expected survival rate implied by the empirically observed reduction in saving rate in the 6 years following the shock. As mentioned in section 4.3.1 the agent's reaction to the shock strongly depends on her risk aversion ρ. Hence, I estimate the reduction in expected survival probability for a range of reasonable risk aversion specifications.

[Insert Table 9 about here.]

Depending on the level of risk aversion, the initial reduction in expected survival probability implied by the observed reduction in saving rate ranges from 1.1 percent to 13.8 percent. Even at a reasonable level of risk aversion of 3 (Chetty, 2006), the observed reduction in saving rate implies a reduction of survival probability of 7.1 percent. In the next year, the relative reduction in expected survival probability is still at 4.3 percent. Over the following five years, this initial reduction in survival probability attenuates to zero. These effects are considerable given that the objective survival probability at age 49 is 99.79 percentage points. Hence, a reduction of 7.1 percent suggests that the expected survival rate drops to 92.7 percentage points directly following the shock.

Next, the last row in table 9 displays the decay parameter λ associated with the attenuating reaction to the shock. The findings show that the estimated λ does not strongly depend on the agent's risk aversion. This is not surprising as it estimated from the changes in expected survival probability from one period to the next rather than from levels. The coefficient estimates range from 1.3 to 1.7. This λ estimate is in the range of the estimates of Malmendier and Nagel (2011) which lie between 1.3 and 1.9. In conclusion, my estimations reveal that the personal experience of the death of a close friend has a quantitatively large impact on a household's mortality beliefs. This is surprising given the non-material nature of the shock. On top of that, the weighting function that governs how this personal experience is incorporated into the belief formation process over time exhibits a similar shape as Malmendier and Nagel (2011). This finding is interesting as my paper explores the completely different domain of mortality beliefs.

6 Robustness

In this section, I address two potential concerns that could explain the observed reduction in saving rate following the death of a close friend. These alternative mechanisms are related to the shock but do not work through the channel of mortality beliefs becoming more pessimistic. Households could take some drastic life choices that affect the composition or work situation of their household. Building on that, there might be unobserved events induced by the shock that lead to a drastic reduction in income which then mechanically reduces the saving rate as consumption might be sticky.

6.1 Other Events

First, the psychology literature asserts that mortality salience changes the timing of conceiving a child. Specifically, individuals that face a mortality salience shock perceive the ideal point of time to bear a child to be earlier (Wisman & Goldenberg, 2005; Fritsche et al., 2007). If individuals in my sample had an increased probability of getting a child following the mortality salience shock, it might mechanically increase consumption and thereby reduce the saving rate. To test for this channel, I simply regress a dummy variable that indicates a child birth in the previous year on the death of a close friend dummy lagged by 1 and 2 periods to account for the 9 months a pregnancy takes. Column 1 and 2 in table 10 demonstrate that the death of a close friend does not increase the likelihood to conceive a child. If anything, it reduces the probability of such an event, even though the economic significance of the coefficient is negligible.

[Insert Table 10 about here.]

Second, the death of a close friend could lead to a drastic change in priorities in ones life. One could imagine that somebody quits her well-paying job to pursue a more fulfilling career. To address this issue, in columns 3 to 4 in table 10 I regress a dummy indicating a change in occupation on the death of a friend dummy. In columns 3, I regress on the same period change whereas in columns 4 the death of friend dummy is lagged. The results show that there does not seem to be neither an immediate nor a delayed reaction concerning an individual's job situation. Last, an individual might feel inclined to reduce her working hours in response to the death of a close friend to enjoy more free time. Thus in columns 5 and 6, I regress an individual's weekly hours worked on the death of a close friend dummy. However, the hours worked only increase on average by 0.06 following this shock which is both economically as well as statistically negligible.

In conclusion, there is no evidence for an indirect channel through which the death of a close friend induces a reduction in the saving rate. The shock to the salience of death neither leads to an increase in childbearing nor to significant changes to one's professional life. This analysis strengthens the idea that the shock to mortality beliefs has a direct effect on the consumption and saving decisions of a household.

6.2 Changes in Income

In this section, I even go a step further and demonstrate that the reduction in survival rate does not depend on a reduction in incoming following the shock. For that purpose, I repeat the analyses of table 2 for a subset of households who experience a non-negative change in income in the next one, two, three, or four years following the death of a close friend.

[Insert Table 11 about here.]

Table 11 displays the results of this analysis. Column 1 demonstrates that focusing on a subset of households that experience non-negative income shocks following the death of a close friend barely changes the estimates. Again, the saving rate is reduced by 2.2 percentage points (compared to the 2.2 percentage points of the full sample) which is highly significant at the 1 percent level. Similarly, columns 2 to 4 reveal that the effect size barely depends on the window in which I require a positive change in a household's income. Even for the subsample of households that experience a non-negative change in income over the next 4 years, the shock induces, on average, a reduction in saving rate of 1.8 percentage points which is highly significant at the 1 percent level. Overall, this robustness test shows that the lower saving rate following the shock is not primarily driven by a reduction in saving rate and sticky consumption patterns.

7 Conclusion

My paper exploits an exogenous shock to the salience of death to causally link mortality beliefs to a household's saving decisions. I show that the death of a close friend has a significant negative impact on both life expectancy as well as a household's saving rate. The impact persists over several years and cannot be explained by adverse health outcomes, bequests, or drastic lifestyle changes. Furthermore, I augment the canonical life-cycle model of consumption by the experience-based learning model of Malmendier et al. (2020). Based on this theoretical framework, I quantify the impact of the shock on beliefs as well es structurally estimate the associated parameter λ that governs how fast the experience fades out of memory. I find that even though the shock has no impact on the household's material situation, it massively affects a household's mortality beliefs. Moreover, the decay parameter λ is in line with previous estimates.

It is crucial to understand whether and how subjective mortality beliefs affect the financial planning of households as miscalibrations can lead to large lifetime utility losses due to undersaving for retirement. My results suggest that individuals do in fact consider mortality beliefs in their consumption-saving decisions apart from possible covariates like health, financial literacy, or wealth. Moreover, my paper demonstrates the importance of personal experiences in forming beliefs as even a non-material shock like the death of a close friend has a substantial impact on beliefs.

My results have important implications for both household finance as well as more generally for how economic expectations are formed. From a household finance point of view, my findings indicate that subjective mortality beliefs are an important component when evaluating the empirical fit of life-cycle models. Taking survival rates as purely exogenous parameters might severely distort model outcomes. Moreover, my results quantify the importance of personal experiences in the expectation formation process. My findings are in accordance with the neuroscientific foundations for experience-based learning proposed by Malmendier (2021). Individuals overweight recent shocks to longevity expectations in their financial decision making and subsequently overadjust their saving rate. This suggests that life-time experiences can distort the financial decision-making of large parts of the population. The importance of personal experiences in forming beliefs might even exacerbate inequalities. One could imagine a situation where individuals of lower socioeconomic status are more often affected by negative experiences like becoming unemployed which translates into more pessimistic beliefs and even less optimal financial decision making.

References

Aguiar, M., & Hurst, E. (2013). Deconstructing life cycle expenditure. *Journal of Political Economy*, *121*(3), 437–492.

Bailey, M., Cao, R., Kuchler, T., & Stroebel, J. (2018). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, *126*(6), 2224–2276.

Baker, A. C., Larcker, D. F., & Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, *144*(2), 370–395.

Bernile, G., Bhagwat, V., & Rau, P. R. (2017). What doesn't kill you will only make you more risk-loving: Early-life disasters and ceo behavior. *The Journal of Finance*, *72*(1), 167–206.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, *119*(1), 249–275.

Calvet, L. E., Campbell, J. Y., & Sodini, P. (2007). Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, *115*(5), 707–747.

Carroll, C. D., Kaufman, A. M., Kazil, J. L., Palmer, N. M., & White, M. N. (2018). The Econ-ARK and HARK: Open Source Tools for Computational Economics. In Fatih Akici, David Lippa, Dillon Niederhut, & M. Pacer (Eds.), *Proceedings of the 17th Python in Science Conference* (p. 25 - 30). doi: 10.25080/Majora-4af1f417-004

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, *134*(3), 1405–1454.

Chetty, R. (2006). A new method of estimating risk aversion. *American Economic Review*, *96*(5), 1821–1834.

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2009). Reinforcement learning and savings behavior. *The Journal of Finance*, *64*(6), 2515–2534.

Cocco, J. F., Gomes, F. J., & Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, *18*(2), 491–533.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica*, *59*(5), 1221–1248.

De Nardi, M., French, E., & Jones, J. B. (2010). Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, *118*(1), 39–75.

Dynan, K. E., Skinner, J., & Zeldes, S. P. (2004). Do the rich save more? *Journal of Political Economy*, *112*(2), 397–444.

D'Acunto, F., Malmendier, U., Ospina, J., & Weber, M. (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy*, *129*(5), 1615–1639.

Fritsche, I., Jonas, E., Fischer, P., Koranyi, N., Berger, N., & Fleischmann, B. (2007). Mortality salience and the desire for offspring. *Journal of Experimental Social Psychology*, *43*(5), 753–762.

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

Hamermesh, D. S. (1985). Expectations, life expectancy, and economic behavior. *The Quarterly Journal of Economics*, *100*(2), 389–408.

Heimer, R. Z., Myrseth, K. O. R., & Schoenle, R. S. (2019). Yolo: Mortality beliefs and household finance puzzles. *The Journal of Finance*, *74*(6), 2957–2996.

Hubbard, R. G., Skinner, J., & Zeldes, S. P. (1995). Precautionary saving and social insurance. *Journal of Political Economy*, *103*(2), 360–399.

Hurd, M. D., Smith, J. P., & Zissimopoulos, J. M. (2004). The effects of subjective survival on retirement and social security claiming. *Journal of Applied Econometrics*, *19*(6), 761–775.

Imbens, G. W., Rubin, D. B., & Sacerdote, B. I. (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American Economic Review*, *91*(4), 778–794.

Jarnebrant, P., & Myrseth, K. O. R. (2013). Mortality beliefs distorted: Magnifying the risk of dying young.

Kalda, A. (2020). Peer financial distress and individual leverage. *The Review of Financial Studies*, *33*(7), 3348–3390.

Kaustia, M., & Knüpfer, S. (2008). Do investors overweight personal experience? Evidence from IPO subscriptions. *The Journal of Finance*, *63*(6), 2679–2702.

Knüpfer, S., Rantapuska, E., & Sarvimäki, M. (2017). Formative experiences and portfolio choice: Evidence from the finnish great depression. *The Journal of Finance*, *72*(1), 133–166.

Kuchler, T., & Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance*, *74*(5), 2491–2542.

Kárpáti, D. (2022). Household finance and life-cycle economic decisions under the shadow of cancer. *Working paper*.

Liu, W.-M., Forbat, L., & Anderson, K. (2019). Death of a close friend: Short and long-term impacts on physical, psychological and social well-being. *PloS one*, *14*(4), e0214838.

Malmendier, U. (2021). Experience effects in finance: Foundations, applications, and future directions. *Review of Finance*, *25*(5), 1339–1363.

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

Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, *131*(1), 53–87.

Malmendier, U., Pouzo, D., & Vanasco, V. (2020). Investor experiences and financial market dynamics. *Journal of Financial Economics*, *136*(3), 597–622.

Post, T., & Hanewald, K. (2013). Longevity risk, subjective survival expectations, and individual saving behavior. *Journal of Economic Behavior & Organization*, *86*, 200–220.

Puri, M., & Robinson, D. T. (2007). Optimism and economic choice. *Journal of Financial Economics*, *86*(1), 71–99.

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

Spaenjers, C., & Spira, S. M. (2015). Subjective life horizon and portfolio choice. *Journal of Economic Behavior & Organization*, *116*, 94–106.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199.

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

Wisman, A., & Goldenberg, J. L. (2005). From the grave to the cradle: Evidence that mortality salience engenders a desire for offspring. *Journal of Personality and Social Psychology*, *89*(1), 46.

Figures

Figure 1: This figure shows the average wealth, consumption, saving rate, and perceived survival probabilities of the simulated life-cycle model. Each panel plots the solution for a household with objective survival probabilities (black) and a household with more pessimistic survival probabilities (red).

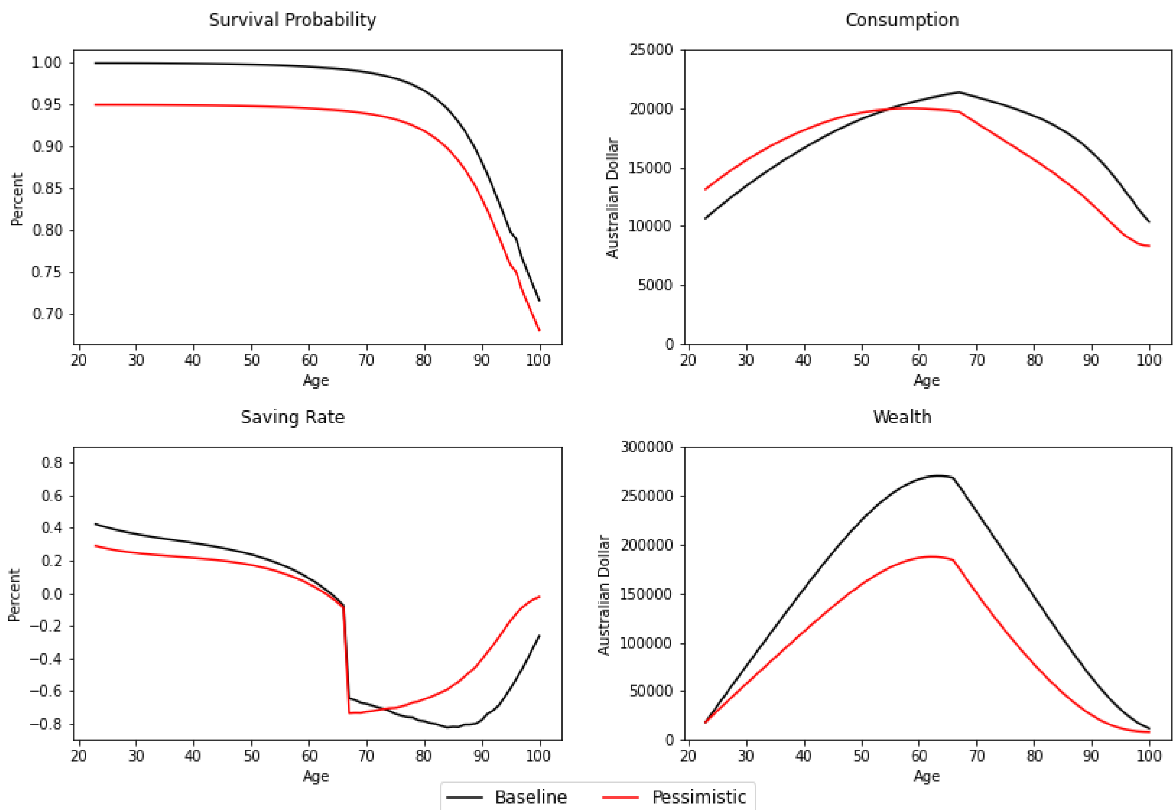

Figure 2: This figure shows the average household saving rate by age. For the left figure, the age of the first member of the household in the sample is chosen. For the right figure, the age of the most senior member of the household is chosen.

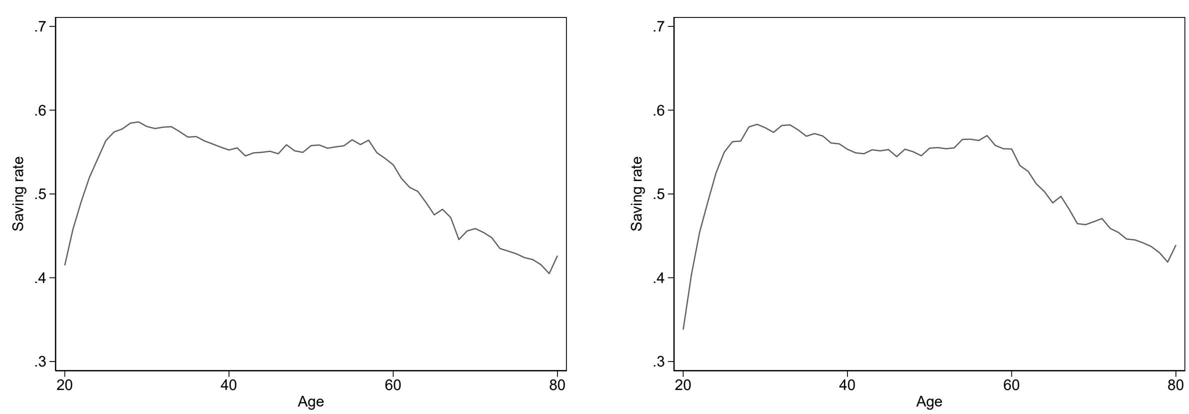

Figure 3: This figure plots the point estimates of column 2 of table 2. The bars around the point estimate indicate the 95 percent confidence intervals.

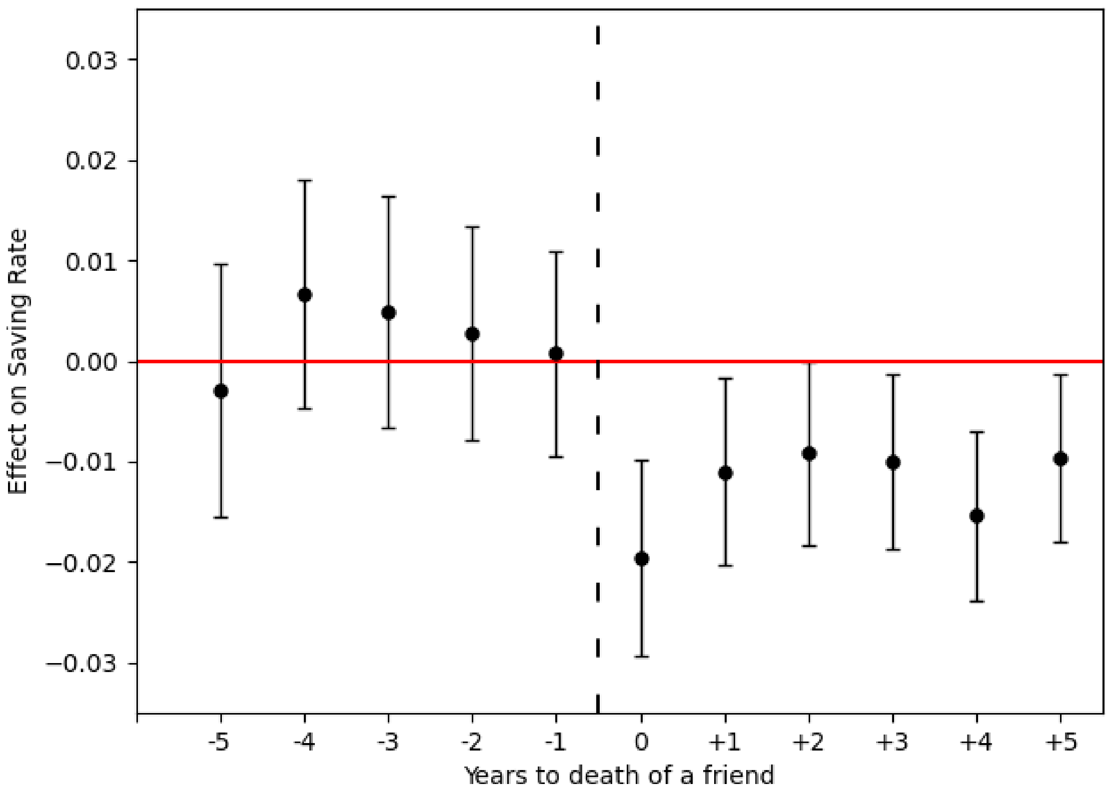

Figure 4: This figure shows the distribution of answers to the question "How likely that you will live to 75 or at least 10 more years?" for age bins of 5 years. Individuals older than 65 are asked instead "How likely that you will live ten more years?".

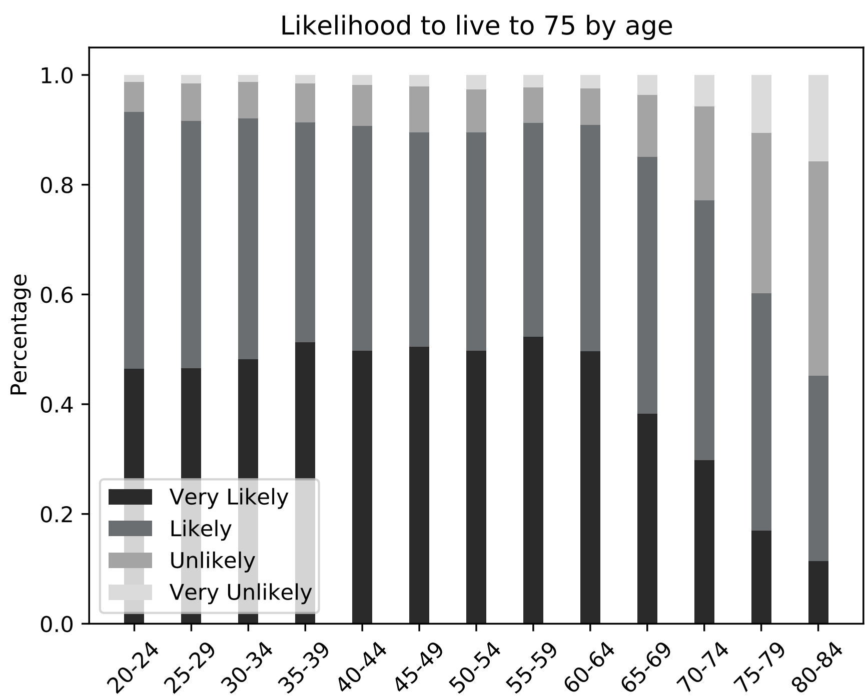

Figure 5: This figure shows the reduction in saving rate around the death of a close friend. The reference group is the saving rate outside of the event window. The bars indicate 95% confidence intervals adjusted for standard error clustering on household level.

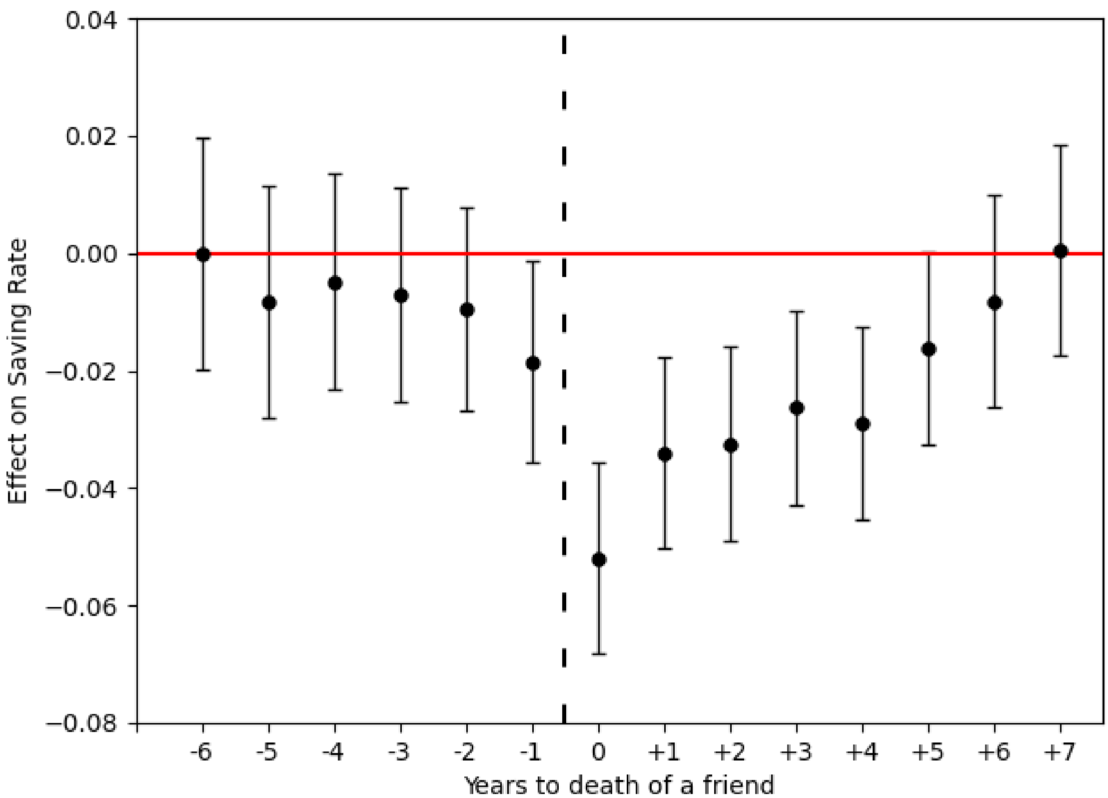

Tables

Table 1: This table presents the summary statistics for the HILDA panel for the years 2001 to 2019. The upper panel shows the variables on individual level whereas the lower panel shows the variables on a household level. Columns 1 to 4 display the mean, median, standard deviation and number of observations for the whole sample.

	Mean	Median	SD	Observations
Individual level				
Female	0.51	1	0.50	387,010
Age	36.99	36	22.39	380,262
Death friend	0.11	0	0.31	242,743
Live to 75?	3.30	3	0.75	46,549
Saving habit	3.33	3	1.21	143,393
Saving horizon	2.87	3	1.53	143,000
Risk aversion	5.36	5	2.47	253,549
Coldness	2.18	2	1.33	19,8235
Household level				
Income (in AUD)	75,426.30	59,535	71,560.32	158,276
Saving rate	0.54	0.62	0.26	114,439
Fun expenditure	0.09	0.07	0.07	120,708
Necessities expenditure	0.25	0.21	0.14	121,259
Health expenditure	0.05	0.04	0.04	117,766

Table 2: This table shows the results from regressing the saving rate on the death of a close friend indicator variable. In column 1, I regress the saving rate on an indicator variable equal to one if the shock occurred in any previous period. In Columns 2 to 4, I regress the saving rate on indicator variables equal to one in the 10 years surrounding the shock. In columns 1 and 2, I run OLS regressions. In column 3 and 4, I use the Sun & Abraham (2021) estimator and the Cengiz et al. (2019) estimator, respectively. All regressions includes household and age fixed effects. Standard errors are clustered by household level, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Saving Rate			
	OLS	OLS	Sun & Abraham (2021)	Cengiz et al. (2019)
Friend Death	-0.022***			
	(-5.42)			
Friend Death (-5)		-0.003	-0.001	-0.003
		(-0.45)	(-0.19)	(-0.40)
Friend Death (-4)		0.007	0.008	0.007
		(1.16)	(1.24)	(1.21)
Friend Death (-3)		0.005	0.006	0.005
		(0.82)	(0.89)	(0.91)
Friend Death (-2)		0.003	0.005	0.003
		(0.51)	(0.84)	(0.61)
Friend Death (-1)		0.001	0.003	0.001
		(0.14)	(0.60)	(0.18)
Friend Death (t=0)		-0.020***	-0.017***	-0.019***
		(-3.96)	(-3.36)	(-3.88)
Friend Death (+1)		-0.011**	-0.010**	-0.011**
		(-2.30)	(-2.01)	(-2.21)
Friend Death (+2)		-0.009**	-0.007	-0.009**
		(-1.98)	(-1.37)	(-2.00)
Friend Death (+3)		-0.010**	-0.008*	-0.010**
		(-2.25)	(-1.80)	(-2.33)
Friend Death (+4)		-0.015***	-0.014***	-0.016***
		(-3.56)	(-3.07)	(-3.64)
Friend Death (+5)		-0.010**	-0.008*	-0.010**
		(-2.29)	(-1.79)	(-2.38)
Household FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Observations	98,946	100,218	100,218	966,539
Adjusted R^2	0.462	0.462	0.463	0.465

Table 3: This table shows the results from regressing the *Saving Habit* or *Saving Horizon* variable on the death of a close friend indicator variable. In columns 1 and 2, I regress the *Saving Habit* on an indicator variable equal to one if the shock occurred in the previous year. In Columns 3 and 4, I regress the *Saving Horizon* variable on an indicator variables equal to one in the year of the shock. Columns 2 and 4 display the results for the subsample of individuals that are 65 years or younger. All regressions includes person and age fixed effects. Standard errors are clustered by person level, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Saving Habits		Saving Horizon	
	Full Sample	Younger than 65	Full Sample	Younger than 65
Friend Death(t-1)	-0.023**	-0.030**		
	(-2.16)	(-2.25)		
Friend Death(t)			-0.019	-0.031**
			(-1.59)	(-2.12)
Person FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Observations	123,540	102,506	99,823	80,906
Adjusted R^2	0.454	0.455	0.458	0.456

t statistics in parentheses

Table 4: This table shows the results of regressing saving rate and consumption components on a dummy variable that is equal to one in each period following the death of a close friend. Column 1 shows the effect on the overall saving rate. Columns 2 to 4 group the consumption components into the categories leisure, necessities, and health and insurance. I estimate OLS regressions with household and age fixed effects. Standard errors are clustered by household, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Saving Rate	Expenditure on Leisure	Expenditure on Necessities	Expenditure on Health
Friend Death	-0.022***	0.006***	0.012***	0.002***
	(-5.42)	(5.77)	(5.01)	(2.78)
Household FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Percentage of overall expenditure		21%	67%	12%
Observations	98,946	104,384	104,858	101,911
Adjusted R^2	0.462	0.494	0.468	0.545

t statistics in parentheses

Table 5: This table shows the results of regressing the saving rate on the death of a close friend indicator variable splitting the households along their parenthood status. Columns 1 and 2 display the results for parents and childless individuals, respectively. Columns 3 and 4 present the results for parents where the child does not live in the household and parents living with a child. I estimate OLS regressions with household and age fixed effects. Standard errors are clustered on household level, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Saving Rate		Saving Rate	
	Parent	Childless	Child not in HH	Child in HH
Friend Death	-0.015***	-0.047***	-0.015**	0.002
	(-3.16)	(-4.74)	(-2.25)	(0.36)
Household FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Observations	73,012	23,241	35,132	37,261
Adjusted R^2	0.454	0.507	0.458	0.432

t statistics in parentheses

Table 6: This table shows the elicited consumption categories that I aggregate to calculate a household's total consumption. I cluster the categories into leisure related expenditure, expenditure on necessities, and health and insurance related expenditure.

Category	**Expenditure on ...**
Leisure	Alcohol, Cigarettes, Meals eaten out, Men's clothing, Women's clothing
Necessities	Groceries, Public transport and taxis, Children's clothing, Telephone rent and calls, Internet charges, Utilities, Car repairs and maintenance, Education fees, Motor vehicle fuel
Health and Insurance	Private health insurance, Other insurances, Medicines, prescriptions and pharmaceuticals, Health practitioners

Table 7: This table shows the results of regressing (1) the likelihood to live to 75 on the death of a close friend dummy and (2) the saving rate on the likelihood to live to 75. In columns 2 and 4, the independent variable is lagged by one year. I estimate OLS regressions with individual and age fixed effects. Standard errors are clustered by individual level, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Likelihood live to 75	Likelihood live to 75	Saving Rate	Saving Rate
Friend Death(t=0)	-0.027**			
	(-1.99)			
Friend Death(t=-1)		-0.011		
		(-0.82)		
Likelihood live to 75(t=0)			0.005**	
			(2.00)	
Likelihood live to 75(t=-1)				0.005*
				(1.83)
Person FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Observations	34,554	32,608	36,246	34,117
Adjusted R^2	0.513	0.519	0.367	0.372

t statistics in parentheses

Table 8: This table shows the results of regressing the saving rate on an indicator variable equal to one for each period following the death of a close friend splitting households along age and risk aversion. Columns 1 and 2 display the results for households younger and older than 50, respectively. Columns 3 and 4 present the findings for high and low risk aversion households, respectively. I estimate OLS regressions with household and age fixed effects. Standard errors are clustered on the household level, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Saving Rate			
	Age < 50	Age > 50	High ρ	Low ρ
Friend Death	-0.035***	-0.012**	-0.012*	-0.032***
	(-5.47)	(-2.14)	(-1.89)	(-3.33)
Household FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Observations	49,617	48,870	31,875	18,660
Adjusted R^2	0.458	0.469	0.459	0.456

t statistics in parentheses

Table 9: This table shows the relative reduction in survival rate implied by the estimated reduction in saving rate. The rows represent the time periods relative to the death of a close friend. Each column displays the results for a different coefficient of risk aversion ρ. The final row shows the fitted decay parameter λ.

	$\rho = 1$	$\rho = 2$	$\rho = 3$	$\rho = 4$	$\rho = 5$
Period 0	0.011	0.039	0.071	0.107	0.138
Period 1	0.007	0.023	0.043	0.070	0.093
Period 2	0.007	0.022	0.040	0.064	0.085
Period 3	0.006	0.020	0.034	0.050	0.065
Period 4	0.006	0.018	0.030	0.050	0.072
Period 5	0.003	0.008	0.012	0.022	0.031
Period 6	0.002	0.003	0.009	0.008	0.015
λ	**1.302**	**1.622**	**1.698**	**1.602**	**1.477**

Table 10: This table shows the results of regressing various life choices on the death of a close friend dummy. Column 1 and 2 display the findings for the birth of a child dummy, columns 3 and 4 for the change in occupation dummy, and columns 5 and 6 for the reported average hours worked. I estimate OLS regressions with person and age fixed effects. Standard errors are clustered on the individual level, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Birth of child	Birth of child	Change in occupation	Change in occupation	Hours worked	Hours worked
Friend Death(t)			-0.003		0.059	
			(-0.73)		(0.58)	
Friend Death(t-1)	-0.002**			-0.001		0.059
	(-2.14)			(-0.36)		(0.55)
Friend Death(t-2)		0.001				
		(0.71)				
Person FE	YES	YES	YES	YES	YES	YES
Age FE	YES	YES	YES	YES	YES	YES
Observations	196,760	175,118	139,533	130,499	150,163	133,061
Adjusted R^2	0.110	0.114	0.146	0.148	0.630	0.632

t statistics in parentheses

Table 11: This table shows the results of regressing the saving rate on an indicator variable equal to one in each period following the death of a close friend for a subsample of households that experience a psoitive change in income in the next 1, 2, 3, or 4 years following the shock. I estimate OLS regressions with household and age fixed effects. Standard errors are clustered by household, and *, **, and *** denote statistical significance at the p < 10%, p < 5%, and p < 1% levels, respectively.

	Saving Rate			
	Next year pos. inc. change	Next 2 years pos. inc. change	Next 3 years pos. inc. change	Next 4 years pos. inc. change
Friend Death	-0.022***	-0.020***	-0.020***	-0.018***
	(-4.83)	(-4.49)	(-4.61)	(-4.18)
Household FE	YES	YES	YES	YES
Age FE	YES	YES	YES	YES
Observations	90,401	91,425	92,263	92,375
Adjusted R^2	0.461	0.466	0.464	0.463

t statistics in parentheses

Appendix A - Variable Descriptions

Variable	Description
Female	Indicator variable equal to 1 if participant is female, 0 otherwise.
Age	Age of participant.
Income	Yearly disposable income from all sources. Households with windfall income are excluded.
Saving rate	One minus the sum of self-reported non-durable consumption divided by yearly disposable income from all sources.
Saving habit	Which of the following statements comes closest to describing your (and your family's) saving habits? *1 Don't save: usually spend more than income* *2 Don't save: usually spend about as much as income* *3 Save whatever is left over - no regular plan* *4 Spend regular income, save other income* *5 Save regularly by putting money aside each month*
Saving horizon	In planning your saving and spending, which of the following time periods is most important to you ? *1 The next week* *2 The next few months* *3 The next year* *4 The next 2 to 4 years* *5 The next 5 to 10 years* *6 More than 10 years ahead*
Fun expenditure	Sum of non-durable expenditure on leisure related categories (c.f. table 6) divided by income.
Necessities expenditure	Sum of non-durable expenditure on necessity related categories (c.f. table 6) divided by income.
Health expenditure	Sum of non-durable expenditure on health and insurance related categories (c.f. table 6) divided by income.

Continued on next page

Friend Death(t)	Indicator variable equal to one if the individual reports the death of a close friend in period t, and zero otherwise.
Friend Death	Indicator variable equal to one for each period following the death of a close friend, and zero otherwise.
Likelihood to live to 75	How likely that you will live to 75 or at least 10 more years? *1 Very likely* *2 Likely* *3 Unlikely* *4 Very unlikely*
Risk aversion	Are you generally a person who is willing to take risks or are you unwilling to take risks? *0 Very willing to take risks* ... *10 Unwilling to take risk*

Appendix B - Model and Estimation Details

B1 - Canonical Life-cycle Model Setup

An agent maximizes her lifetime utility. Let t be the agent's adult age and T the maximum number of periods the agent lives. Then the agent faces the following maximization problem:

$$\max\ \mathbb{E}[\sum_{t=1}^{T} \beta^{t-1} (\prod_{j=0}^{t-2} s_j) u(c_t)]$$

where c_{it} is the consumption of agent i at age t, β is the discount factor, and most importantly s_j is the agent's probability to survive from period $j-1$ to j. I do not consider bequest motives and assume u to represent a power utility function. Each period the agent decides how much of his income to consume and the remainder is saved at a fixed rate of R.

Labor Income Process. During an agent's working age, she receives an exogenously given stochastic labor income Y:

$$log(Y_{it}) = f_t + \zeta_{it} + \epsilon_{it}$$

where f_t is a function representing the deterministic component of labor income at age t and ϵ_{it} is an idiosyncratic shock to labor income which is distributed $N(0, \sigma_\epsilon^2)$. ζ_{it} constitutes a persistent shock to labor income:

$$\zeta_{it} = \zeta_{i,t-1} + u_{it}$$

where u_{it} is $N(0, \sigma_u)$ distributed and uncorrelated with ϵ_{it} and all shocks are uncorrelated across households. After the agent reaches the age of 65, she enters retirement and her labor income becomes deterministic. It is given by the last working period's permanent income multiplied by a replacement factor.

Optimization Problem. All real variables are normalized by the permanent labor income P_t to reduce the dimensionality of the state space to 1. I denote all normalized variables by lower case letter. Each period, the agent has a certain amount of cash-on-hand which is the sum of her savings and savings returns and her labor income:

$$m_{it} = y_{it} + w_{it}$$

where w_{it} is given by:

$$w_{it} = R(w_{i,t-1} + y_{i,t-1} - c_{i,t-1})$$

The agent maximizes (B1) under all of these conditions. The Bellman equation is given by:

$$\nu_{it}(m_{it}) = \max_{c_{it}} \ u(c_{it}) + \beta s_{i,t+1}\mathbb{E}[(p_{i,t+1}/p_{it})^{1-\rho}\nu_{i,t+1}(m_{i,t+1})]$$

There is no analytical solution to this problem. Hence, the policy functions are solved numerically.

B2 - Solving the Model

The model is solved by backward induction. The solution for the last period is trivial as the agent consumes all of her remaining wealth. Hence, in the second to last period one can plug in the indirect utility function for next period's value function. Based on this, it is possible to derive a consumption function that gives the optimal level of consumption given a certain level of wealth (cash-on-hand). Furthermore, one can derive the value function for the second to last period. To obtain the solution for all periods, one iterates backwards from the last to the first period.

Unfortunately, there is no analytical solution to the maximization problem. In practice, to reduce computational load I construct a discrete grid of possible cash-on-hand levels and find the optimal level of consumption for each of these grid points. Finally, the grid points are interpolated to construct the consumption function. For the graphs, I simulate the outcomes for 5000 agents and average over outcomes[3].

B3 - Structural Estimation

I estimate the implied reduction in survival rate and the associated decay in effect based on the reduction in saving rate observed in the data following the death of a close friend. I do not directly observe the impact of the shock on the survival rate. However, the rareness of the event of a close friend dying greatly reduce the complexity of the problem: (1) The initial shock represents 100% of the set of experiences. Hence, I can normalize all further effects by the initial shock. (2) The initial shock remains the only component of the set of relevant experiences as the agent is not exposed to any new experiences. Thus, I can directly compare the subsequent changes in survival rate to the initial reduction in survival rate to elicit the weight of the first experience in these later periods.

I take this intuition to the empirical results. In a first step, I estimate the corresponding drop in perceived survival rate associated with the reduction in saving rate estimated from the data. For that purpose, I fit the survival rate separately for each period after the shock. I simulate the saving rate for a list of relative reductions in survival rate from 0.3 to 0 in steps of 0.001. Then, I select the relative reduction in survival rate that corresponds to the survival rate estimated in that period in Table 5. This gives rise to a list of relative reductions in survival rate for each of the

[3]For setting up and solving the model, I utilize the *Heterogeneous Agents Resources and toolKit (HARK)* by Carroll et al. (2018)

seven periods following the mortality beliefs shock. I repeat this procedure for a list of coefficients of relative risk aversion ranging from 1 to 5. In a second step, I estimate the λ that fits the implied reductions in survival rate best. First, I calculate the weights of the period 0 experience for all 6 periods following the initial shock for a grid of λ ranging from 0 to 5 in steps of 0.01. Then, I find the squared distance between the in the previous step calculated weights and the implied reductions in survival rate which gives me the best fitting λ. Finally, I make sure this represents a global minimum.