CANCER AND PORTFOLIO CHOICE: EVIDENCE FROM NORWEGIAN REGISTER DATA

Technical report

Trond Doeskeland
Jens Kvaerner

THINK FORWARD INITIATIVE

CANCER AND PORTFOLIO CHOICE:
EVIDENCE FROM NORWEGIAN REGISTER DATA [*]

TECHNICAL REPORT

Trond Doeskeland and Jens Kvaerner [†]

October, 2020

Abstract

We examine the personal investment decisions of 70,000 Norwegian couples following a cancer diagnosis, including choices made after the loss of a partner. To create a control group that minimizes the likelihood that the results are confounded by lifestyle factors, we rely on couples who are diagnosed with the same disease, but a few years later. Our difference-in-difference estimates show that a cancer diagnosis reduces households' willingness to take risk with their financial wealth. Widowhood decreases stock market participation by a factor of ten relative to non-fatal cancers. The effect depends on the relative importance of the deceased individual's human capital. There is no link between prior experience with cancer and current personal investment strategies.

Keywords: Health, Cancer, Financial Decision Making

1. Introduction

Cancer is a health crisis that affects people of all genders, both young and old, rich and poor. More than half of all those aged under 65 will eventually experience cancer (Sasieni, Ahmad, & Ormiston-Smith, 2015). Cancer can have a major impact on earnings, mortality, family dynamics, roles, priorities and mental health. Given its prevalence and its impact on variables typically tied to personal investment decisions, such as whether or not to participate in the stock market, and the risk profile of uncertain financial investments, it is important to understand how and when cancer influences investment decisions. In this study, we estimate the financial portfolio responses of 70,000 Norwegian households over a period in which one member of the household is diagnosed with cancer.

In order to maximize the welfare benefits of financial markets, it is critical to ensure that households invest effectively (Campbell, 2006). In this regard, our research potentially has several practical implications for financial advisors and households that will experience adverse health outcomes in the future. For example, regarding financial advisors, it is important to understand what triggers changes in personal investment decisions during such difficult times. Regarding households, more information about what other people do in similar vulnerable situations may provide guidance on how to proceed.

When a family member is diagnosed with cancer, many things happen at the same time, and all members of the family are affected in some way. Portfolio theory asserts that differences in preferences or circumstances cause heterogeneity in risk-taking (Curcuru, Lucas, Heaton, & Moore, 2009). Heterogeneity in circumstances encompasses a wide range of potential explanatory factors including households' future earning profiles and family

composition (Bodie, Merton, & Samuelson, 1992; Viceira, 2001; Cocco, Gomes, & Maenhout, 2005). Earlier literature shows that leukemia, lymphomas, lung, brain, bone, colorectal, and head-and-neck cancers in particular result in major reductions in employment and earnings (Syse, Tretli, & Kravdal, 2008).

Normative portfolio theory advocates households to take all assets, including pension and labor earnings, into account when choosing the risk profile of their financial wealth (Bodie, Merton, & Samuelson, 1992). Households with safe and high labor earnings are recommended to take more risks than households with riskier and lower earnings. The normative portfolio theory is in line with the folk wisdom often advocated by financial advisors that the fraction of financial wealth in safe assets should equal the investor's age. Following this train of thought, we expect that, on average, households will take less risks after a cancer diagnosis. In addition, we expect households to reduce risk more dramatically in the case of cancers that have a major impact on future family earnings.

We have three main findings. First, we investigate the consequences of a change in health status, with a corresponding impact on survival probability. We show that a cancer diagnosis affects risk-taking both in terms of the likelihood of participating in the stock market and the composition of financial wealth among the participants. Second, by investigating the consequences of mortality, we find that fatal cancers increase the probability that the surviving spouse will exit the stock market by a factor of ten relative to non-lethal cancer. Lethal cancers decrease human family-wealth and increase the debt-to-income ratio of the household. As mentioned above, standard theory and conventional folk wisdom both suggest that these factors make the household more

cautious. In other words, there is no need for explanations such as shift in risk aversion or other behavioral biases to explain our results. Finally, we use the historical records of cancer diagnoses in our data to investigate whether personal experience with cancer many years ago is associated with the current portfolio choice. We find that the negative short-to-medium term effects fade out over time.

Taken together, our findings suggest that cancer leads to short-term changes in circumstances that influence portfolio choice. In the long-term, these imbalances wash out, and diagnosed members of households therefore choose similar portfolios to non-diagnosed peers. In the same way that patients that recover from cancer will be more like cancer-free households on many socio-demographic variables, the portfolio choice of households that have experience with cancer will be similar to non-diagnosed households.

The empirical literature on the role of health and family composition on portfolio choice has struggled with two main challenges. The first is the unavailability of large-scale household-level data on health, family links, and financial outcomes. The second obstacle is the difficulty of isolating the causal effects of health crises in the presence of complex dynamics and systematic measurement error. In this context, systematic measurement error may reflect a justification bias, which induces respondents to report a worse subjective level of health to justify their current economic status (McGarry, 2004). As a result, solely relying on subjective measures of health status and self-reported measures of wealth might render the coefficient estimates difficult to interpret. In this research we tackle the above-mentioned challenges. First, we construct a large-scale dataset by merging individual medical records with administrative data on financial and demographics for the entire Norwegian population from 2005 to 2013. The availability of our large-scale dataset has the potential to uncover new facts about health and financial decision-making. For example, we can test

whether the importance of the diagnosed individual's labor income to family consumption prior to the health crisis affects ex-post financial outcomes. Second, we develop an identification strategy that controls non-parametrically for unobserved confounding variables. For example, smoking can affect both cancer and financial choices. By identifying the effect of treatment from the difference in outcomes between households affected today and households affected a few years later, we difference out lifestyle factors that may differ between the treatment and control group. Following both treatment and controls for many years before and after the event year, we illustrate parallel trends in outcomes before both cancer diagnosis and mortality, which implies that the identifying assumption holds. In the following, we provide details about our three key findings outlined above, and place them in context.

A cancer diagnosis increases the probability of exiting the stock market by about 1 percent. Diagnosed individuals that stay in the stock market hold about the same portfolio as non-diagnosed peers. More severe diagnoses measured by the probability of survival tend to increase these precautionary responses. In general the relatively limited effects of cancer on portfolio choice hold, regardless of the relative importance of the diagnosed individual's income, and across the income-to-financial wealth distribution.

These findings relate to the empirical literature on health and financial decision-making. Using data from the Health and Retirement Study, Rosen & Wu (2004) are first to note that health is a significant cross-sectional predictor of risky asset holdings; households in poor health hold a smaller share of their financial portfolios in risky assets. Follow-up studies, generally find that the effect of moving into the lowest self-reported health category on portfolio choice is negligible and statistically weak after controlling for individual fixed-effects and changes in financial wealth (Love & Smith, 2010; Fan & Zhao, 2009; Berkowitza & Qiu, 2006). With regard to these

papers, our causal estimates are in accordance with Rosen & Wu (2004).

Our next set of results examines the effect of mortality from cancer on portfolio choice. The estimated magnitudes are now economically consequential and highly statistically significant. First, almost every sixth household that suffers a loss from cancer will also exit the stock market. Second, the surviving individuals reduce their allocation to risky assets by 7 percent relative to the counterfactual household. In line with normative portfolio theory, the estimated effects depend on the diagnosed individual's share of the total household income: The larger the share, the more important the diagnosed individual's income is for total household consumption, and the higher the likelihood is that the surviving spouse exits the stock market.

Apart from being of practical importance to households and financial advisors, our new insights into the causal relationships and mediating pathways between cancer and financial portfolio decisions can provide guidance to economic models of intra-household financial decision-making (Hong & Rios-Rull, 2006; Love & Smith, 2010; Addoumy, Kungz, & Morales, 2016). In this regard, we provide several causal moments these models can target. For example, our results indicate that the sensitivity of couples' risk bearing capacity with respect to changes in the family composition depends crucially on the income shares of the individuals in the household. This suggests that the analyst should model income shares carefully. Simplifying assumptions such as equal income shares may point to misleading model predictions.

Medical research shows adverse effects related to mental and physical health after cancer treatment (Miller & Massie, 2006). A growing trend in finance literature shows that heterogeneity in risk-taking can be attributed to personal experience with fluctuations in macroeconomic conditions (Andersen, Hanspal, & Nielsen, 2019; Malmendier & Nagel, 2011; Rantapuska, Knupfer, & Sarvimaki, 2017). We contribute to this literature by examining the relationship between personal experience with cancer and personal investment styles many years after the diagnosis. Our results show that couples who experienced cancer more than five years ago are indistinguishable from couples that will experience cancer in the next year. In that sense, experience with cancer appears different from a bad experience with macroeconomic conditions.

Our results fit into the growing literature that occupies the intersection between health and finance. Recent contributions include, Gupta, Morrison, Fedorenko, & Ramsey (2018) who find that cancer patients that are homeowners use house equity to smooth consumption. Heimer, Myrseth, & Schoenle (2019) show that biased mortality expectations can explain both the low saving rates at the beginning of the life-cycle and high saving rates among retirees. Koijen & Nieuwerburgh (2019) propose that life insurance companies have incentives to pay for life-extending medical treatments.

The outline of the rest of the paper is as follows. In Section 2 we elaborate on the data, the institutional setting, and the research design. We present the results in Sections 3. Finally, we provide a summary of our conclusions from the study.

2. Data sources and the institutional setting

In this section we describe the dataset and the institutional setting of Norway. We describe the saving norms, the social security system, the cancer treatment and the insurance market. We include a particular focus on the extent to which the results generalize to other OECD countries.

2.1 Data sources

The original data come from the Cancer Registry of Norway (CRN), the Norwegian Tax Registry (NTR), and Statistics Norway. Each of these registers uses data provided by several minor registers. The CRN has registered all cancer cases in Norway since 1953. Mandatory reporting from clinicians, pathologists, and death certificates ensures high-quality data on the date of diagnosis, the patient's age at diagnosis, gender, tumor location (International Classification of Diseases, 10th revision (ICD-10)), and stage at diagnosis (local, regional, distant, or unknown). Records are linked to the Cause of Death Registry maintained by Statistics Norway.

We complement the cancer data with financial and demographic data from the Norwegian Tax Registry and Statistics Norway. The Norwegian Tax Administration is responsible for collecting income and wealth taxes in Norway. By law, employers, banks, and public agencies are obliged to disclose personal information on income and wealth to the Tax Administration. The tax return includes all sources of income, as well as detailed information on wealth and debt. Individuals are accountable for the information provided in their tax return, and the submission of inaccurate information is punishable by Norwegian law. Another characteristic of our dataset is that we have detailed information about the family structure. The tax authority collects information on the complete wealth holdings for all households at the end of every year. For wealth tax purposes, the household can allocate wealth in a way that gives the lowest wealth tax. Thus, there are no incentives for tax-motivated asset allocation within the household. In the event of a death, most surviving spouses retain a life interest in the undivided estate, which is typically not distributed to heirs before the surviving spouse dies.

Our dataset is similar in content to that used by Kvaerner (2017). The main differences lie in the unit of analysis and the dependent variables. Kvaerner (2017) examines how a cancer diagnosis affects saving choices of two-generational households, in which the oldest generation is restricted to single individuals aged above 59. In contrast, we study whether a cancer diagnosis and mortality influence portfolio choice over the full life-cycle.

2.2 Institutional setting

The effect of a cancer diagnosis on stock market participation and asset allocation may depend on the external environment. Therefore, to assess the external validity of our study, we compare Norway to other OECD countries in terms of dimensions that can affect the relationship between cancer and portfolio choice.

Our first point is related to portfolio choice over the life-cycle and the "saving environment" in general. Overall, in Norway, households' portfolio choices are similar to most other countries while the aggregate net saving rate is higher. Much of the saving concerns long-term pension saving. The Norwegian pension system rests on three pillars: public pensions, occupational pensions, and private savings. The public pension system is a pay-as-you-go income-based pension system. The contribution to the income-based pension is 18.1% of an

individual's income, though the income is capped (in 2013, the cap is NOK 605 thousand or $105 thousand). Pension benefits are indexed to the average wage growth of the economy. In addition to public pensions, the Norwegian workforce is entitled to occupational pensions. Agreements between labor unions and employer organizations are broad and inclusive. Additionally, in Norway, there is a shift from defined benefit pension to defined contribution for the private sector. However, for the period investigated, most households have a defined benefit occupational pension. Effectively, the return on the public and occupational pension is similar to that of a real bond. In practice, however, households can only influence private savings, which is the focus of our analysis. For private savings, the asset composition is similar to other OECD countries. In particular, the stock market participation rate is at the same level as comparable countries (Ke, 2018). Among participants, the average share invested in stocks is 32%; similar numbers prevail in other European countries. In summary, the "saving environment" in Norway is similar to most OECD countries.

Our second point is related to the cost and availability of cancer treatment. Cancer can lead to higher medical expenditure, which absorbs financial wealth. In this view, cancer expenditure is a type of undiversifiable background risk that induces precautionary saving and prompts the choice of safer portfolios. In Norway, standardized cancer treatment is promptly available and practically free of charge (Cherny, Sullivan, Torode, Saar, & Eniu, 2016). Private treatment was almost non-existent in the period (Fiva, Hoegelandb, Roenningb, & Syse, 2014). The medical expense risk related to treatment in Norway is thus similar to most other European countries but different from, for example, the U.S. However, in the U.S., many households have employer-provided health insurance while working, and Medicare in retirement, which covers basic health expenses. A related point is out-of-pocket (OOP) medical expenditure unrelated to treatment (e.g., follow-ups). These costs vary significantly between nations. For example, countries with higher public funding of healthcare report lower OOP medical expenditure and frequent barriers to service access. On the other hand, in countries where households face higher OOP health expenses unrelated to treatment, we observe more developed markets for supplemental health insurance, which washes out some of the potential country-specific precautionary saving. In summary, from the viewpoint of financial decision-making, dealing with cancer in Norway is similar to having cancer in most OECD countries (e.g., Germany) and for American households with good private insurance coverage.

The final point relates to payouts from insurances. These payouts come at the same time as changes in fundamental variables that themselves may influence portfolio choice, including future earnings. Unfortunately, in the data, we cannot observe whether individuals are policyholders, and therefore cannot say anything about whether policyholders make different financial decisions than non-policyholders. In the case of non-lethal cancers, in our sample period, relatively few individuals hold some kind of private insurance that pay out an amount in the case of cancer. Among policyholders, most hold the default option, which provides a one-time payment of about NOK 250 thousand (about $43 thousand) if they are diagnosed with cancer. A related insurance is insurance against so-called "critical illnesses", which covers a set of predefined diseases including cancer. At the beginning of our sample in 2007, approximations based on official statistics suggest that less than 0.02 percent of diagnosed individuals have some kind of insurance against cancer (i.e., cancer prevalence times fraction of population that are policyholders). In the bottom panel on the right-hand side in Figure 1, we see that the ratio between income and financial wealth is stable before and after the cancer. This indicates that the payout from private insurance is negligible.

In the case of lethal cancers, the prevalence of life insurance is at approximately the same level as for

other OECD countries. The payout from life insurance is likely to drive down the risky share because it is paid in cash. However, there is no reason why payouts would cause stock market exit, which is what we find in the data. In the bottom panels in Figure 1, we see that the ratio between income and financial wealth after cancer is developing differently than risky share and the participation rate. This indicates that there is no mechanical relation between life insurance and portfolio choice. In conclusion, it is unlikely that payouts from insurances heavily influence the average effect of cancer on portfolio choice.

3. Methodology

In estimating the effect of adverse life events such as cancer on portfolio choice, we set up an identification strategy that controls non-parametrically for unobserved confounding variables following (Fadlon & Nielsen, 2019). For example, smoking may affect both cancer and financial choices. Another example is risk aversion, which may influence both the portfolio choice and, for example, how long one lives with cancer before being diagnosed. The basic idea is to recover the effect of cancer from the difference in outcomes between households affected by the same shock at different points in time, which in turn would "difference out" such confounding..

We define the treatment group from households affected by a cancer diagnosis in 2007, 2008, and 2009. Similarly, we define the control group from households affected by a cancer diagnosis in 2010, 2011, and 2012. We follow the treatment group from three years before the diagnosis until two years after the diagnosis. For the control group, we follow them from six years to one year before the diagnosis. The basic idea is to then use the households of individuals diagnosed in 2010 as counterfactual for those diagnosed in 2007, and so on. We recover the treatment effect from the difference in outcomes between the two groups over the period 2007, 2008, and 2009.

The treatment and the control group are either based on all diagnoses, conditional on expected survival, or ex-post mortality status. Regarding the effect of mortality on portfolio choices, we compare the portfolio behavior of the surviving partner with that of couples where one individual in the household will die from cancer three years later. The relative year only depends on the year of death and not the diagnosis year. Since all diagnosed individuals in households

in the sample now live with cancer, the difference in outcomes isolates the effect of mortality from the effect of cancer. Ensuring homogeneity in exposure within the sample is important to maximize the likelihood that we "difference out" potential confounding.

The research design removes changes in the asset holdings that come from changes in asset prices, and therefore solve potential problems with portfolio inertia. To see this, it is useful to decompose the observed portfolios for the households in the treatment and control group, after those in the treatment group have been diagnosed with cancer, into two components: A passive component, which reflects the portfolio the household would have in the absence of any active change, and an active component, which reflects the active change in the portfolio (for example the sale of stock A). Under the identifying assumption that the exact timing of a cancer diagnosis is random within three years, the initial portfolios of households in the two groups are identical. Consequently, the passive component of their portfolio will be the same after individuals in the households in the treatment group have been diagnosed with cancer. As a result, any observed differences in the portfolios between the two groups reflect active portfolio decisions.

To quantify the effects, as advocated by Angrist & Pischke (2008), we rely on a linear probability model, which is much easier to interpret than the non-linear specification. A related benefit associated with the linear specification is that it provides an easy way of testing for (and interpreting of) moderating relationships. By moderating relationships, we refer to variables that influence the relationship between portfolio choice and the shock resulting from either a cancer diagnosis or mortality. Specifically, we

estimate the following two-period difference-in-difference regression:

$$y_{i,t} = a_i + bPost_{i,t} + \gamma_i Post_{i,t} \times Treat_i + \lambda X_{i,t} + e_{i,t},$$

where $\gamma_i = \gamma_0 + \gamma_1 m_i$. Here m_i is the moderating variable that we postulate affects the strength of the relationship between a dependent and an independent variable. The two parameters of interest are the average effect of the treatment, γ_0, and the incremental effect of the moderating variable, γ_1, for a given level of treatment. The other variables are standard: $Treat_i$ is a dummy variable that takes the value of one if the household member is in the treatment group and zero otherwise, $Post_{i,t}$ takes the value of one if relative time is greater or equal to zero but less than three, $X_{i,t}$ is a vector of controls, which always includes $m_i \times Post_{i,t}$ and $m_i \times Treat_i$. All regressions include household fixed effects and standard errors are clustered at the household level.

3.1 Variable construction

We now define the variables we use and justify the restrictions we impose on the sample. Apart from the cancer data, the variables we use follow mainstream research in household finance closely (Calvet, Campbell, & Sodini, 2007; Rantapuska, Knupfer, & Sarvimaki, 2017).

Portfolio choice is reflected both in the decision to participate in the stock market and the relative share of financial assets invested in risky assets. We refer to households holding a positive amount of risky assets as stock market participants. Risky financial wealth is defined as the sum of directly owned stocks and equity mutual funds. Safe financial wealth refers to bank account balances and bond mutual funds. Financial wealth is the sum of the two. Our second dependent variable is the risky share, defined as the ratio of risky financial wealth to financial wealth.

Our chosen unit of analysis concerns couples for two reasons. First, marital status and children almost certainly affect household resources, preferences,

and background risk, and these elements lie at the heart of financial decision making (Hong & Rios-Rull, 2006). Second, cancer affects not just the person that becomes ill but also the family of the diagnosed individual. We define couples as individuals registered as married, or as cohabitant. When we consider couples, we include same-sex couples as well as opposite-sex couples. For couples, we aggregate individual assets to the household level. Household demographics are based on the individual that is not diagnosed with cancer. If the spouse dies, we follow the surviving spouse in the subsequent years. In some specifications, we exclude widows or widowers starting in the year after the loss of the spouse. The purpose of excluding these households is to isolate the effect of losing a spouse from the effect of cancer on portfolio choice.

For cancer diagnosis, we distinguish between two primary outcomes. The first is to be diagnosed with non-lethal cancer. For non-lethal cancers, we differentiate between high and low expected survival rates based on so-called relative survival rates. Relative survival is a measure of the excess mortality and it is defined as the observed survival in a patient group, divided by the survival of a comparable group in the general population, with respect to key factors affecting survival, including age, gender and calendar year of investigation. Importantly, relative survival rates depend on the cancer type, gender, and the cancer stage, and are easily accessible from various public cancer information sources. For example, in Norway, the five-year relative survival for lung cancer discovered in Stage 1 is about 41% for men and 51% for women, whereas it is about 16% for men and 21% for women in Stage 2, and only 1% for men and 2% for women if discovered in stage 3. The second outcome is lethal cancers. In the case of lethal cancers, we study the portfolio choice of the surviving partner. Separate analyses for cancers that are non-lethal and lethal allow us to distinguish between the effect of being diagnosed with cancer and the mortality effect.

Guided by portfolio choice theory in incomplete markets, we expect that a household's lifetime earnings will affect the relationship between treatment and the outcome variables. We define income as the sum of gross salary income and pensions. We define relative income in the household as the income of the diagnosed spouse divided by total income. The larger the relative income, the higher the relative loss in human capital resulting from the cancer diagnosis. We complement relative income with the income-to-financial wealth ratio. To ascertain the possibility of surviving men making different financial choices than surviving women, we add gender as an interaction variable. We also include children as an interaction variable, which captures the effect of bequest motives on the relationship between cancer and portfolio choice. Finally, we perform separate regressions using proxies for being financially constrained as outcome variables.

To arrive at a final sample, we restrict the sample to households in which one individual was diagnosed with cancer between 2007 and 2012. Table (1) presents summary statistics for the control and treatment group.

Insert Table 1 here

The average householder in our sample is about 60 years old, with 2.4 children. The average household income is NOK 770 thousand ($100 thousand). About 50 percent of the sample participates in the stock market. This number is higher than for all households since the cancer sample is older than the entire Norwegian population. The average risky share is 15 percent. The last column reports the standardized mean difference (SMD), which is the most commonly used statistic to examine the balance of covariate distribution between treatment groups. It is calculated as the difference in means between the treatment and the control group scaled by the square root of the sum of the group variance. Because SMD is independent of the unit of measurement, it allows comparison between variables with different units of measurement. A rule-of-thumb is that standardized covariate mean differences of a matched dataset should be below 0.1 (Branson, 2019). We see that apart from the relative income variables, all our variables satisfy the rule of thumb criteria.

4. Results

We present our results in three main segments. First, we analyze the effect of cancer on portfolio choice. Second, we estimate the effect of mortality on portfolio choice. The main set of results includes average treatment effects and tests for moderating relationships. Third, we investigate whether personal experience with cancer many years ago is associated with the current portfolio choice.

4.1 Cancer effects

We start by investigating the average effect of cancer on portfolio choice. First we estimate the total effect of cancer before we try to isolate the effect of the cancer diagnosis from the effect of death. Table 2 presents the results from estimating Equation (1) without any interaction term. The mean effect of cancer on portfolio choice is given by the coefficient on the variable $Post_{i,t} \times Treat_i$.

Insert Table 2 here

Columns 1 and 2 report the results using all diagnoses without any restrictions on survival. We find that a cancer diagnosis reduces the average stock market participation rate by 0.8 percentage points, or 1.6 percent relative to the average participation rate before the diagnosis. Average risky shares decreases with 0.3 percentage points, or 2.0 percent relative to the average risky share before the event. To shed light on the underlying mechanisms behind the effect, we identify two sources from which the effect can originate; the cancer diagnosis in itself, and/or as a result of the death of one individual in the family. As we will see, there is a major difference between these two events.

4.2 The effect of the cancer diagnosis

In this section we investigate the effect of receiving the diagnosis on portfolio outcomes. Column 3-8 in

Table 2 presents the estimates restricted to the sample in which both partners in the household are alive at the beginning of the year. As expected, the results are quite similar to the average effect reported in columns 1 and 2 reported above as most individuals survive. Moreover, interestingly, columns 5 to 8 indicate that severe cancer diagnoses (i.e., low probability of survival) are associated with a higher likelihood of stock market exit and a lower risky share than less severe cancer diagnoses.

Figure 1 displays the results graphically. The broken line with triangles shows the behavior of the treatment group while the line with circles displays the counterfactual outcome. The top plots show the average participation rate and risky share over the event-window. Our falsification tests are based on inspection of parallel trends prior to the diagnosis year for the treatment group. The parallel trends in outcomes displayed in Figure 1 provide evidence that our main results are not driven by general trends, or realizations of other confounders, over the event-window by means of which we follow the households. Similar choices prior to treatment gives support to the identifying assumption of the random nature of cancer diagnosis within a few years. In conclusion, our findings are robust to falsification tests.

Insert Figure 1 here

To conclude the first set of the results, all our empirical evidence points to the fact that a cancer diagnosis in itself has a relatively minor effect on portfolio choice. In relation to previous findings, the empirical results are similar to papers based on U.S data. For example, Rosen & Wu (2004) find that having bad health reduces the probability of owning risky assets by 1.7 percentage points. The similarity between our causal effect estimates and previous

findings based primarily on survey data from different countries provides external validity.

4.3 The effect of mortality caused by cancer

The next set of results focuses on the effect of mortality shock on portfolio choice resulting from losing a partner to cancer. In this analysis, all households are diagnosed with cancer, so that the difference in outcomes isolates the effect of mortality from the effect of cancer. Table 3 presents the results.

The effect of mortality on portfolio choice is both economically and statistically significant. As shown in column 1, the mean effect of mortality on stock market participation is 7.9 percentage points. Relative to the average participation rate before the diagnosis, this is a reduction of about 15 percent. In other words, almost every sixth household that loses one individual to cancer exits the stock market. Column 2 shows that allocation to risky assets is 1.1 percent points lower for households that have lost their spouse. In light of the baseline risky share of 15 percent, the coefficient estimate implies a 7 percent decrease in the risky share.

Insert Table 3 here

Figure 2 displays these results graphically. The top plots show the average participation rate and risky share around the loss of the spouse. Again, the parallel trends in outcomes prior to the diagnosis year suggest that our findings are robust to placebo and falsification tests. The main takeaway is that losing a partner to cancer has a significant impact on all outcomes.

Insert Figure 2 here

To investigate potential sources of heterogeneity in the response to losing the spouse to cancer, we develop three measures of human wealth and formally test the moderating factors in Columns 3 to

8 in Table 3. The motivation for focusing on human wealth comes from portfolio theory, which advocates that households should take all assets, including pension and labor earnings, into account when choosing the risk profile of their financial wealth (Bodie, Merton, & Samuelson, 1992). Households with safe and high labor earnings are recommended to take more risks than households with riskier and lower earnings and vice versa.

The first interaction variable is the relative importance of the deceased individuals' income. The higher the relative income, the more important the individual's income is for the household. To facilitate the interpretation, we standardize the relative income to have a mean zero and a standard deviation of one. In line with basic theory, the likelihood that the surviving individual exits the stock market is significantly higher if the deceased individual is the primary-earner. The coefficient on the interaction term with Relative income reveals that one standard deviation increase in the relative income of the deceased individual leads to a further 2.3 percentage points lower stock market participation. Along the intensive margin, relative income also decreases the risky share.

The second interaction variable is the households' income-to-financial wealth ratio. Columns 5 and 6 show the income-to-financial wealth ratio before the diagnosis, which is also normalized to have zero mean and a standard deviation of one, which does not alter the relationship between mortality and financial risk-taking.

The third interaction variable is gender. The gender dummy variable takes the value of one if the non-diagnosed individual is male. The result reported in column 7 shows that surviving males are less likely to exit the stock market than surviving females. In our sample, the sum of gross salary and pension income for women is about 40 percent lower than for men. The results for gender are thus consistent with that of relative income. An alternative explanation (not

mutually exclusive) for why the effect of a death appears to depend on the gender of the surviving spouse is that women are more risk-averse than men (Barber & Odean, 2001).

Columns 9 and 10 investigate whether children affect the relationship between widowhood and portfolio choice. Overall, the results reveal that having children tends to increase the likelihood of a stock market exit.

The results presented above provide novel empirical evidence on financial decision making in the transition within family structures. Love & Smith (2010) use a life-cycle model to investigate the impact of changes in family composition on optimal asset allocation. They show that at all ages, widowhood leads to a decline in the optimal portfolio share, with the most significant movements for working-age women. Widowhood causes a sudden drop in the income-to-wealth ratio which, in turn, makes future consumption streams more sensitive to fluctuations in the rate of return on risky assets. Addoumy, Kungz, & Morales (2016) study portfolio choice in a life-cycle model similar to Love & Smith (2010), but endogenize marital status. In their model, after divorce, risky asset holdings in the female's (male's) single portfolio decrease (increase) relative to the married individual's portfolio. Part of the gender difference in the portfolio adjustment following the transition from being married to being single reflects differences in ex-ante risk aversion between women and men. Our findings are mostly in line with Love & Smith (2010).[2]

4.4 Long-term cancer effect
The next set of results focuses on the effect of a mortality shock on portfolio choice resulting from losing a partner to cancer. Bearing in mind that all households are diagnosed with cancer, the

difference in outcomes thus isolates the effect of mortality from the effect of cancer.

Several papers show that heterogeneity in revealed risk-taking between individuals can be attributed to experiences. Malmendier & Nagel (2011) and Rantapuska, Knupfer, & Sarvimaki (2017) document that personal experiences with negative macroeconomic conditions tend to reduce risk-taking many years after the experience. Kaustia & Knupfer (2012) and Hoffmann & Post (2017) relate personal histories in the stock market to cross-sectional variation in risk-taking. Andersen, Hanspal, & Nielsen (2019) show that even a relation to others with negative experiences in the stock market can depress risk-taking. However, little is known about the underlying mechanisms behind the relation between experiences and risk-taking. We complement the above-cited studies by estimating the association between having personal experience with cancer and subsequent portfolio choice using the following cross-sectional regression:

$$y_i = a + \sum_{\tau=-10}^{2} \gamma(\tau)1\{\text{Diagnosisyear}_i = \tau\} + \lambda X_i + e_i.$$

Here, τ denotes the number of years that have passed since the cancer diagnosis. For example, $\tau = -10$ means that the outcome variable is measured ten years after the diagnosis year. Similarly, $\tau=2$ means that the outcome variable is measured two years prior to the diagnosis year. One year before the diagnosis, $\tau=1$ is the reference year. The set of coefficients $\gamma(\tau)$ provides an estimate of the relationship between having dealt with cancer and portfolio choice. For reasons of simplicity, we estimate the regression for the calendar year 2010 only. Again X_i is a vector of controls, which includes a third-order polynomial in age, and the cross product of gender, education, as well as the number of children living in the household. Figure 3 presents the

subsequent portfolio decisions. Overall, we found little evidence for such moderating relationships.

---

[2] We also estimated whether the same variables influenced the relationship between cancer and

results from the estimation. The plot on the left-hand side shows the results for stock market participation. The plot on the right-hand side shows the plot for the risky share.

Insert Figure 3 here

Personal experience with cancer is negatively related to stock market participation and risk-taking in the short- to medium-term, but the association becomes weaker over time. In particular, the financial risk-taking among couples with direct experience with cancer in the last few years is lower than the participation rate among couples that will experience cancer in the next year. In contrast, the participation rate among couples that experienced cancer 10 years ago is indistinguishable from couples that will experience cancer in the next year. The last coefficient in the participation plot represents the falsification test. It shows that the average participation rate among those that will be diagnosed with cancer in two years is identical to those who will be diagnosed with cancer in one year.

In summary, our results indicate that personal experience with cancer is not associated with risk-taking in the long-run. According to American Society of Clinical Oncology (ASCO), the medical community considers many cancers "cured" when doctors cannot detect cancer 5 years after diagnosis. Our finding, that individuals who have been cancer-free for between five and ten years have the same investment styles as individuals with similar demographics but who have not had cancer, are in line with the forward looking assumption that underlies most of economic modeling.

5. Conclusion and implications

Having direct experience with a health risk at some time in life is becoming more common. Two important reasons are longer life expectancy and globalization. Longer life expectancy due to advancement in biotechnical research and treatment options means that more people will fully recover from previously fatal health crises. Globalization, traveling, and international trade will continue to increase the pace at which new infectious diseases spread (for example SARS and COVID-19). Given the prevalence of health crises and their influence on variables typically tied to individuals' and households' financial decisions, it is important to understand how, when, and for whom such negative life events affect financial choices, which is the focus of our research.

The empirical literature on the role of health and family composition on portfolio choice has struggled with two main challenges. The first is the unavailability of large-scale household-level data on health, family links, and financial outcomes. The second obstacle is the difficulty of isolating the causal effects of health crises in the presence of complex dynamics and systematic measurement error. In this context, systematic measurement error may reflect a justification bias, which induces respondents to report a more severe subjective level of health to justify their current economic status. As a result, solely relying on subjective measures of health status and self-reported measures of wealth is unsatisfactory. We tackle the above-mentioned challenges in our research.

The negative life events in our study are cancer diagnosis and widowhood. Both events can influence labor market outcomes and individuals' well-being with potential spillover effects on financial decision-making. Both events are also common and affect everyone, men and women, young and old, rich and poor. Due to our individual-level data, we are able to investigate how financial decision-making after cancer and widowhood depends on variables such as family income, and financial commitments such as debt, prior to the diagnosis.

We document three empirical facts related to the effect of cancer on financial decision making. First, we show that a cancer diagnosis affects portfolio choice. Second, we show that lethal cancers increase the probability of a stock market exit for the surviving spouse by a factor of ten relative to non-lethal cancers. We establish these results by exploiting the fact that the timing of a cancer diagnosis is as good as random over three years. As a result, we can identify the effect of treatment on the outcome variable from the difference in behavior between someone affected today and someone affected by the same crisis a few years later. Third, we show that personal experience with cancer more than five years ago is uncorrelated with current portfolio choice. Overall, our findings suggest that cancer leads to short-term changes in circumstances such as human capital and liquidity that influences participation in the stock market. In the long-term, these changes vanish and the portfolio choice is similar to cancer-free households.

In order to maximize the welfare benefits of financial markets, it is critical to ensure that households invest effectively (Campbell, 2006). In this regard, our results have practical implications for financial advisors and households that will experience adverse health outcomes in the future. For example, for financial advisors, it is important to understand what triggers the surviving spouse to exit the stock market. For households, more information about

what other people do in similar vulnerable situations may provide guidance on how to proceed.

Our research can be expanded in several academic and non-academic directions. The first step is to evaluate the external validity of this study, that is, to investigate whether our results apply outside Norway and across various diseases. Second, as already indicated, it would be interesting to explore not only whether negative health shocks affect financial choices, but also why they do so. We hypothesize four key channels: Mortality, wages, medical expenditure, and the ability to work. Third, our ongoing research shows that negative health-related lifetime events can lead to a major reduction in family income and increase the likelihood of liquidity constraints. Building on these insights, it would be interesting to investigate the effect of health crises on personal bankruptcy. Here one could, for example, identify the key variables that explain variation in realized bankruptcy for a given health crisis (e.g., cancer). Finally, it would be interesting to build an economic model that helps to understand the quantitative importance of our findings and whether current insurance products are sufficient to fully insure against the financial consequences of these health crises. An important goal for future research in this regard is to help policymakers and insurance companies to come up with standardized health insurance products which are automatically optimized throughout life, as retirement accounts often are.

6. References

Addoumy, J. M., Kungz, H., & Morales, G. (2016). Limited Marital Commitment and Household Portfolios. Working Paper.

Andersen, S., Hanspal, T., & Nielsen, K. M. (2019). Once bitten, twice shy: The power of personal experiences in risk taking. Journal of Financial Economics, 132, 97-117.

Angrist, J. D., & Pischke, J.-S. (2008, 12). Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press.

Barber, B. M., & Odean, T. (2001). Boys will be Boys: Gender, Overconfidence, and Common Stock Investment. The Quarterly Journal of Economics, 116, 261-292.

Berkowitza, M. K., & Qiu, J. (2006). A further look at household portfolio choice and health status. Journal of Banking & Finance, 30, 1201_1217.

Bodie, Z., Merton, R. C., & Samuelson, W. F. (1992). Labor Supply Flexibility and Portfolio Choice in a Life-Cycle Model. Journal of Economic Dynamics and Control, 16, 427-449.

Branson, Z. (2019). Is My Matched Dataset As-If Randomized, More, or Less? Unifying the Design and Analysis of Observational Studies. Working Paper.

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, 707-747.

Campbell, J. Y. (2006). Household Finance. Journal of Finance, 61, 1553-1604.

Cherny, N., Sullivan, R., Torode, J., Saar, M., & Eniu, A. (2016). ESMO European Consortium Study on the availability, out-of-pocket costs and accessibility of antineoplastic medicines in Europe. Annals of Oncology, 27, 1423-1443.

Cocco, J. F., Gomes, F. J., & Maenhout, P. J. (2005). Consumption and Portfolio Choice over the Life Cycle. Review of Financial Studies, 18, 491-533.

Coile, C., & Milligan, K. (2009). How household portfolios evolve after retirement: The effect of aging and health shocks. Review of Income and Wealth, 55, 226-248.

Curcuru, S., Lucas, D., Heaton, J., & Moore, D. (2009). Heterogeneity and portfolio choice: Theory and evidence in Ait-Sahalia Yacine, and Lars Hansen, ed.: Handbook of Financial Econometrics: Tools and Techniques. North-Holland.

Fadlon, I., & Nielsen, T. H. (2019). Family Health Behaviors. American Economic Review, 109, 3162-91.

Fan, E., & Zhao, R. (2009). Health status and portfolio choice: Causality or heterogeneity? Journal of Banking & Finance, 33, 1079-1088.

Fiva, J. H., Hoegelandb, T., Roenningb, M., & Syse, A. (2014). Access to treatment and educational inequalities in cancer survival. Journal of Health Economics, 36, 98-111.

Gupta, A., Morrison, E., Fedorenko, C. R., & Ramsey, S. (2018). Home Equity Mitigates

the Financial and Mortality Consequences of Health Shocks: Evidence from Cancer Diagnoses. working paper.

Heimer, R. Z., Myrseth, K. O., & Schoenle, R. S. (2019). YOLO: Mortality Beliefs and Household Finance Puzzles. Journal of Finance, 74, 2957-2996.

Hoffmann, A., & Post, T. (2017). How return and risk experiences shape investor beliefs and preferences. Accounting & Finance, 57, 759-788.

Hong, J. H., & Rios-Rull, J.-V. (2006). Social Security, Life Insurance and Annuities for Families. Working Paper.

Kaustia, M., & Knupfer, S. (2012). Peer performance and stock market entry. Journal of Financial Economics, 104, 321-338.

Ke, D. (2018). Cross-Country Differences in Household Stock Market Participation: The Role of Gender Norms. AEA Papers and Proceedings, 108, 159-162.

Koijen, R. S., & Nieuwerburgh, S. V. (2019). Combining Life and Health Insurance. Quarterly Journal of Economics.

Kramer, L., & Weber, J. M. (2012, 3). This is Your Portfolio on Winter. Soc Psychol Pers Sci, 3. doi:10.1177/1948550611415694

Kvaerner, J. S. (2017, 6). Intergenerational Altruism: Estimates Based on News About Expected Mortality. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2985465

Love, D. A., & Smith, P. A. (2010). Does health affect portfolio choice? Health Economics, 19, 1441-1460.

Malmendier, U., & Nagel, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? The Quarterly Journal of Economic, 126, 376-416.

McGarry, K. (2004). Health and Retirement: Do Changes in Health Affect Retirement Expectations? The Journal of Human Resources, 39, 628-648.

Miller, K., & Massie, M. J. (2006). Depression and anxiety. Cancer Journal, 12, 388-97.

Montazerhodjat, V., Weinstock, D. M., & Lo, A. W. (2016). Buying Cures Versus Renting Health: Financing Health Care with Consumer Loans. Science Translational Medicine, 327, 1-7.

Rantapuska, E., Knupfer, S., & Sarvimaki, M. (2017). Formative experiences and portfolio choice: Evidence from the Finnish Great Depression. Journal of Finance, 72, 133-166.

Rosen, H., & Wu, S. (2004). Portfolio choice and health status. Journal of Financial Economics, 72, 457-484.

Sasieni, P., Ahmad, A. S., & Ormiston-Smith, N. (2015). Trends in the lifetime risk of developing cancer in Great Britain: comparison of risk for those born from 1930 to 1960. British Journal of Cancer, 112, 943-947.

Syse, A., Tretli, S., & Kravdal, O. (2008, 9). Cancer's impact on employment and earnings -a population-based study from Norway. Journal of Cancer Survivorship, 2, 149-158.

Viceira, L. M. (2001). Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income. Journal of Finance, 433-470.

7. Tables and Figures

Table 1: Descriptive statistics

	Control Group				Treatment Group				SMD
	N	Median	Mean	SD	N	Median	Mean	SD	
Family income	38,8	509,1	564,7	359,6	27,6	516,9	569,6	313,3	0,010
Relative income	38,6	0,526	0,515	0,222	27,4	0,526	0,520	0,216	0,016
Real Estate	38,8	270,4	286,0	254,3	27,6	274,9	287,6	243,8	0,004
Safe Financial Wealth	38,8	185,8	431,8	1137,6	27,6	192,6	441,8	1464,9	0,005
Debt	38,8	337,3	734,4	1894,1	27,6	354,9	753,5	3724,9	0,005
Risky Financial wealth	38,8	1,1	138,4	730,8	27,6	1,6	142,2	1732,9	0,002
Age	38,8	60,0	60,1	11,1	27,6	60,0	60,2	10,9	0,004
Number of Children	38,8	2,000	2,371	1,179	27,6	2,000	2,372	1,195	0,001
High School	38,6	1,000	0,747	0,435	27,4	1,000	0,743	0,437	-0,006
Any college	38,6	0,000	0,253	0,435	27,4	0,000	0,257	0,437	0,006
Male	38,8	0,000	0,397	0,489	27,6	0,000	0,403	0,490	0,009
Risky Share	38,8	0,007	0,152	0,246	27,6	0,008	0,153	0,246	0,003
Stock Market Particpation	38,8	1,000	0,526	0,499	27,6	1,000	0,531	0,499	0,007
HtM dummy	38,8	0,000	0,187	0,390	27,6	0,000	0,182	0,386	-0,009
DtY dummy	38,8	0,000	0,080	0,271	27,6	0,000	0,080	0,271	0,000

Note: The table reports summary statistics of demographic and financial variables for the treatment and the control group two years before the treatment group will be diagnosed with cancer (i.e., for the relative year -2). The unit of analysis is the household, defined as married and with a valid ID for their spouse, or as individuals registered with a valid cohabitant ID. Individual assets are aggregated to the household level, and household demographics (age, education, sex, etc.) are based on the individual that is not diagnosed with cancer. N refers to the number of observations and reported in thousands. Any college is an indicator variable that takes the value of one if the individual has a college education or higher. Family income refers to the sum of gross salary income and pension plus net capital income and total government transfers. Relative income is defined as the income of the diagnosed spouse divided by total income. Risky financial wealth is defined as the sum of stocks and mutual funds, and Safe financial wealth refers to bank account balances and money market funds. Real estate is the book value of real estate assets, and total liabilities are the sum of mortgages and other outstanding credit. Stock Market Particpation is a binary variable that takes the value of one if the risky share is non-zero, and the risky share is the relative share of financial assets invested in risky financial wealth. ´´Hand-to-Month" (HtM) refers to households with cash holdings two years prior to the diagnosis less than 1/15th of gross income. The classification is binary. In the same vein, the debt overhang dummy variable (DtY) takes the value of one if as gross debt - financial wealth, exceeds three times gross income. Financial variables are measured in NOK thousands. The exchange rate for US dollars to Norwegian kroner as of 31 Dec 2009 was $ 0.1731. Standardized mean difference (SMD) is defined as the mean in the treatment group minus the mean in the control group divided by the square root of the sum of the variance of the two groups.

Table 2: The effect of cancer on portfolio choice

Outcome	I	α	I	α	I	α	I	α
TreatxPost	-0.0085***	-0.0032**	-0.0063**	-0.0030**	-0.0080**	-0.0016	-0.0155*	-0.0097**
	(0.0026)	(0.0014)	(0.0026)	(0.0014)	(0.0039)	(0.0020)	(0.0093)	(0.0047)
Num. obs.	376694	376694	371057	371057	154060	154060	43348	43348
Adj. R2	0.8208	0.6852	0.8214	0.6855	0.8260	0.6913	0.8265	0.6792
Excl. Singles	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Prognosis Severity	All	All	All	All	RS>0.75	RS>0.75	RS<0.25	RS<0.25

Note: The coefficient of interest, TreatxPost shows the effect of a cancer diagnosis on the outcome variable, measured over relative years: t=0,1,2. I is a dummy variable that takes the value of one if household i is a stock market participant at time t and is otherwise zero. α is the risky share. RS stands for relative survival rate. Standard errors are clustered at the individual level. The notation ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$ mean significant at the 1, 5, 10 percent level, respectively. All regressions include age controls and individual fixed effects.

Table 3: The effect of mortality on portfolio choice

| Outcome	I	α	I	α	I	α	I	α
TreatxPost	-0.079***	-0.0107**	-0.097***	-0.006	-0.09***	-0.015***	-0.06**	0.015
	(0.0088)	(0.0043)	(0.011)	(0.0055)	(0.0111)	(0.005)	(0.0279)	(0.015)
TreatxPostxm	-0.023***	-0.0137***	0.0119	-0.0040	0.0381**	0.015	-0.020	-0.028*
	(0.0086)	(0.0052)	(0.010)	(0.0052)	(0.0190)	(0.01)	(0.0294)	(0.0160)
Num. obs.	43681	43681	32685	32685	70572	70572	70572	70572
Adj. R2	0.8266	0.6714	0.8361	0.7025	0.8385	0.6891	0.8384	0.6892
Interaction	Relinc	Relinc	Income / FW	Income / FW	Gender	Gender	Children	Children

Note: The first coefficient of interest, TreatxPost shows the mean effect of losing a spouse on the outcome variable, measured over relative years: t=0,1,2. The second coefficient of interest, TreatxPostxm, shows the effect interaction variable has on the relationship between the outcome variable and widowhood. All interaction variables are measured two years prior to the diagnosis (i.e., in relative year t=-2). Relinc is short for relative income, which is defined as the income of the diagnosed spouse divided by total income and age is the age of the surviving spouse. Relinc is normalized to have mean zero and a standard deviation of one. Income / FW measures total household income normalized by total financial wealth. Gender is a dummy variable that takes the value of one if the surviving spouse is male and children is a dummy variable that takes the value of one if the household has children. I is a dummy variable that takes the value of one if household i is a stock market participant at time t and is otherwise zero. α is the risky share. Standard errors are clustered at the individual level. The notation ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$ mean significant at the 1, 5, 10 percent level, respectively. All regressions include age controls, individual fixed effects and all cancer types.

**Figure 1: Financial outcomes around cancer diagnosis**
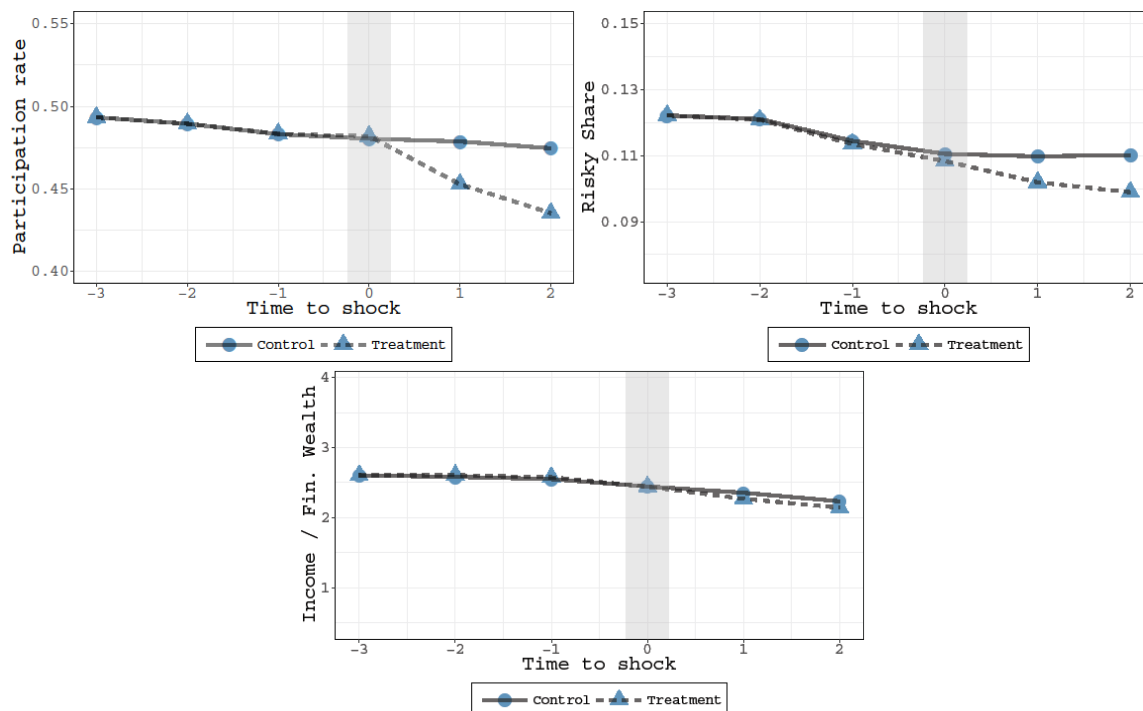
Note: The y-axis shows the dependent variable of the household, and the x-axis shows the years around the diagnosis year. Survival indicates that we look at the sample of surviving cancer households. For the treatment group, period zero is the year of the diagnosis while for the control group period zero is three years before the diagnosis year. The exchange rate for US dollars to Norwegian kroner as of 31 December 2009 was 0.1731$.

Figure 2: Financial outcomes around mortality
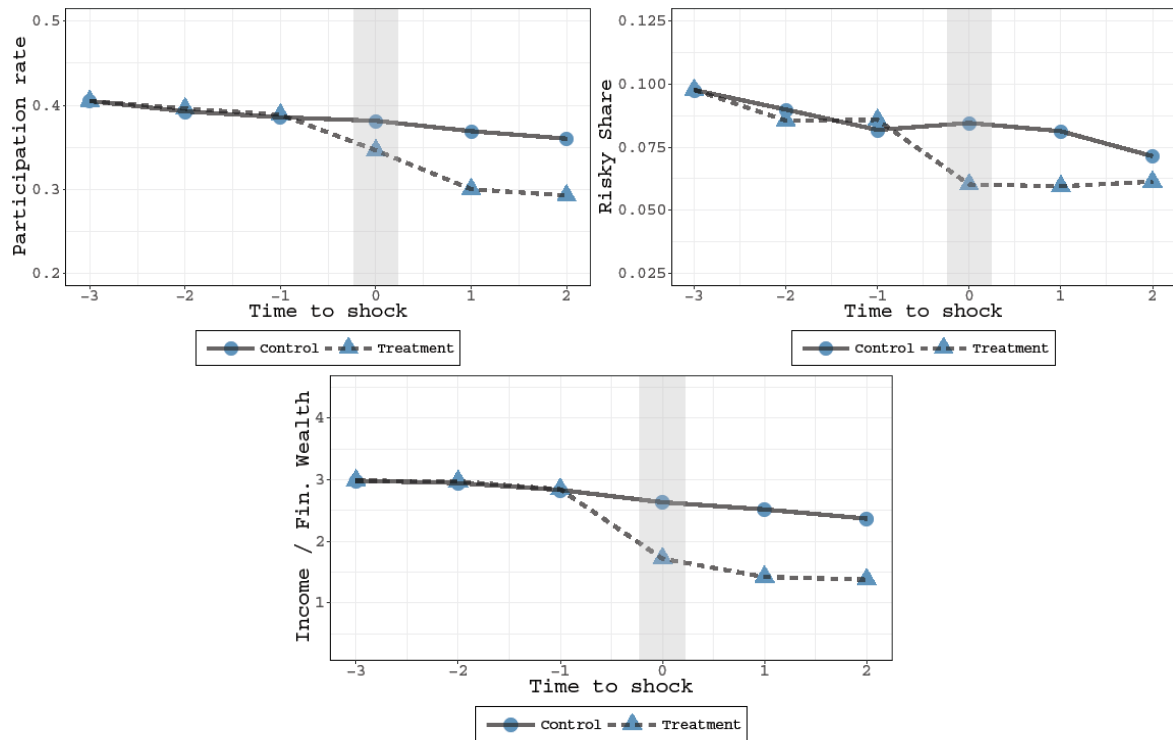
Note: The y-axis shows the dependent variable of the household, and the x-axis shows the years around the year of the death of the spouse. Death indicates that we look at the sample where one in the household died. For the treatment group, period zero is the year the spouse dies while for the control group period zero is three years before the spouse dies. The exchange rate for US dollars to Norwegian kroner as of 31 December 2009 was 0.1731$.

Figure 3: Long run financial outcomes after cancer
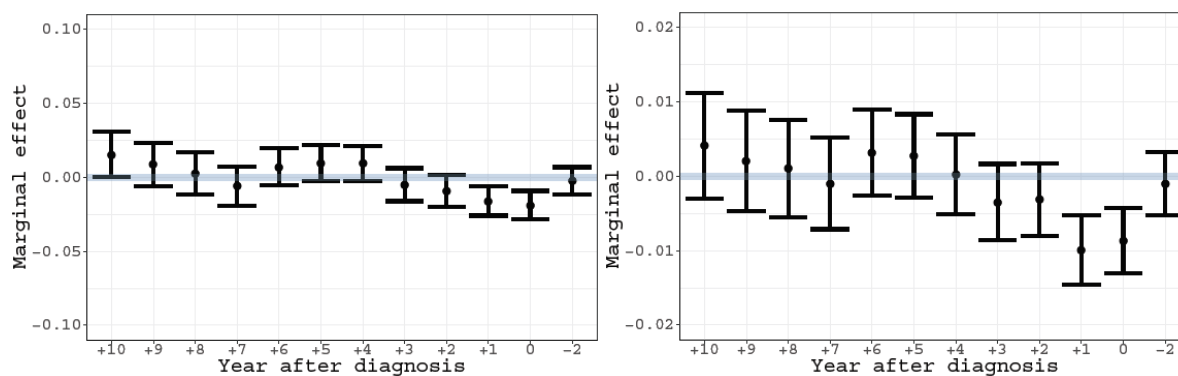
Note: The plots display the regression coefficient in a cross-sectional regression in 2010 on a set of household characteristics and year-dummies since the cancer diagnosis. Each dot represents the coefficient estimate for a particular relative year dummy variable. For example, the dot on +10 is the coefficient on the dummy variable that takes the value of one if it is 10 years since the individual was diagnosed with cancer in 2010. The omitted year diagnosis year is 2011 and the dot on -2 is the coefficient on the dummy variable that takes the value of one if the individual is diagnosed with cancer in two years.

The authors

Jens Kvaerner
Assistant Professor
Tilburg University
E-mail: J.S.Kvaerner@tilburguniversity.edu

Trond Doeskeland
Professor
Norwegian School of Economics
E-mail: Trond.Doskeland@nhh.no