

David Sturrock
Stefan Groot
Jan Möhlmann

22/29

Working paper

Wealth, gifts and estate planning at the end of life

Wealth, gifts and estate planning at the end of life*

David Sturrock[†] Stefan Groot[‡] Jan Möhlmann[§]

July 28, 2022

Abstract

We quantify the size and tax-responsiveness of financial transfers made to heirs before death. The wealth of singles (including widows) with children declines substantially in anticipation of death. 92% of this decline is explained by transfers to children, while long-term care copayments can explain the remainder. Tax-reducing transfers are made across the wealth distribution and are very responsive to the tax rate. Using bunching and difference-in-bunching estimation and exploiting a reform to gift taxation, we estimate Frisch elasticities of gifts to the net-of-tax rate between 9, for those giving around €27,000, and 1, for those giving around €125,000.

Keywords: Deathbed giving, inter-vivos transfers, bequests

JEL Codes: D12, D14, D64, H26

*Funding from the Economic and Social Research Council (through Research Grant number ES/V001248/1 and through the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP) (ES/M010147/1)) is gratefully acknowledged. We thank participants at seminars at Netherlands Bureau for Economic Policy Analysis, University of Cambridge, UCL, Princeton University, and Institute for Fiscal Studies and at the Netspar International Pension Workshop (2022), Royal Economic Society Conference (2022) and the Annual Meeting of the Society of Economics of the Household (2022). We are also grateful to Richard Blundell, Rowena Crawford, Carl Emmerson, Eric French, Wojciech Kopczuk, Rory McGee, Esteban Garcia-Miralles, Cormac O’Dea, Eduard Suari-Andreu and Owen Zidar for helpful comments. Any errors are our own. Correspondence to david.sturrock.17@ucl.ac.uk.

[†]University College London and Institute for Fiscal Studies

[‡]Rabobank

[§]Netherlands Bureau for Economic Policy Analysis

1 Introduction

Understanding the behavioural mechanisms driving intergenerational transfers is crucial for transfer taxation policy (Piketty and Saez, 2013) and the design of insurance at older ages (De Nardi et al. (2010); De Nardi and Yang (2016); Lockwood (2018)). In most advanced economies, the systems of estate, inheritance and lifetime gift taxation offer opportunities for households to reduce their tax burden by passing on wealth before the point of death (Poterba (2001); McGarry (2001)). To the extent that this incentive is exploited, it may change the optimal level or progressivity of the system of estate or inheritance taxation and increase the efficiency gains from moving to a system that taxes lifetime receipt of gifts and inheritances in a unified system (Mirrlees et al., 2011a).

A key question is therefore the extent to which elderly households engage in ‘estate planning’ behaviour whereby they begin to pass on wealth in the months and years before their death in order to reduce their tax burden. More generally, ‘estate planning’ behaviour of this sort would constitute evidence of an active motivation to pass on wealth, as opposed to bequests being a by-product of precautionary saving resulting from medical expenditure or longevity risk (so-called ‘accidental bequests’).

In this paper, we use intergenerationally-linked administrative panel data on wealth and taxable gifts from the Netherlands to estimate the decline in wealth, size of financial gifts made, and elasticity of gifts to taxation, over the four years before death. In the Netherlands, inheritances and the annual flow of gifts are taxed separately according to a progressive schedule with exempt amounts. Someone anticipating leaving a taxable inheritance (for giving to children, that is an inheritance over around €20,000) can increase the after-tax transfers received by their heirs by passing on some wealth before death. We find that single individuals with children make substantial gifts to their heirs in anticipation of death and that these transfers are highly responsive to incentives to transfer wealth as a gift while alive rather than as an inheritance at death.

In the first part of the paper, we quantify the decline in wealth and size of transfers made,

in anticipation of death. To identify changes in wealth and gifts attributable to anticipatory behaviour, we exploit variation in the timing of death conditional on initial characteristics, using a matched ‘event study’ or ‘dynamic differences-in-differences’ design. Our treatment group consists of all individuals resident in the Netherlands who died between 2013 and 2015. We match each individual in the treatment group to someone of a similar age, wealth level, income level, health status, and household composition four years before the death of the treated individual, but who died four years later than the treated individual. We then compare the evolution of the wealth, gifts, and wealth of heirs of the individuals in the treatment group to that of the individuals in the matched control group, over the four years leading up to the death of the treated individuals. The idea is that, when matched, the treatment and control groups have the same survival prospects, but the treatment group contains those individuals who experience a health shock leading to their death in the subsequent four years. We interpret our event study estimates as the effect of additional anticipatory behaviour that occurs three, two, one or zero years before the year of death, compared to that present four or more years before death.

We find that the wealth of single individuals (including those formerly in couples) is 6.9% lower by the start of year of their death, as a result of proximity to death. By contrast, the wealth of couples is 1.5% higher by the year of the first death of a couple member. These changes in wealth are larger, and begin further in advance of death, for those who experience a decline in health further in advance of their death. This suggests that news about health and therefore mortality precipitates changes in behaviour or health-related costs.

We demonstrate that the vast majority of the substantial decline in wealth of singles is explained by transfers to children, while long-term care copayments play a minority role. First, we show that the wealth decline of singles is primarily driven by those with children. Second, we find that at the beginning of the year of their death, 92% of the wealth decline of singles with children can be explained by a contemporaneous rise in their children’s wealth. Third, we find that dynasty wealth (the sum of parents’ and their children’s wealth) does

not change significantly as death approaches. Even for those in the top quartile of initial wealth, 66% of the decline in parents' wealth is explained by a rise in their children's wealth, with long-term care co-payments able to explain the residual decline in wealth.

In the second part of the paper, we use data on taxable gifts to examine the motivations for giving. Taxable gifts increase as death approaches for singles in the top half of the wealth distribution who have children, and are very responsive to taxation. Quantitatively, taxable gifts are less than a tenth of the value of the increase in children's wealth, indicating that the vast majority of giving is below the tax-exempt threshold and so tax-reducing for those who will leave substantive bequests. We use bunching estimation and, exploiting a reform to the gift taxation system, differences-in-bunching estimation, to estimate Frisch elasticities of gifts to the net-of-tax rate. We estimate elasticities of between 9, for those giving gifts of around €27,000, and 1, for those giving gifts of around €125,000. These elasticities represent a high degree of willingness to respond to temporary tax incentives to make gifts when death is near. Looking within families, we find that less-wealthy children receive only slightly larger gifts than their wealthier siblings, consistent with a 'warm glow' from giving and a modest role for altruism.

Our interpretation of our results is that there is a significant transfer of wealth from singles to their children in the final years before death, that this is the main driver of the decline in wealth for this group as death approaches, and that these transfers are highly responsive to tax incentives. For those in couples there is no increase in transfers made in advance of death, and we infer that the spending of couples reduces in the final years before one member of the couple dies. These findings are consistent with a model of households who wish to pass on wealth to their children and to minimise their tax burden when doing so, but who, in the face of uncertainty over future longevity or spending needs, or a desire to keep control of wealth, hold onto their wealth until they believe their death to be imminent. Couples in this framework do not pass on their wealth to their children as the longevity of the surviving member is still uncertain. How do we explain the different consumption responses

amongst singles and couples? If health shocks lead to a reduction in longevity expectations while at the same time reducing the marginal utility of consumption, this can lead couples to save for the future, when the marginal utility of the surviving spouse is higher. But for singles, there is no such reason to save for the future.

This paper makes two major contributions to the literatures on estate planning and taxation of wealth transfers. First, we provide direct evidence of transfers to heirs in anticipation of death by examining taxable gifts, and explicitly quantify how much of the decline in wealth before death is explained by transfers, as opposed to other potential drivers such as increased medical or long-term care spending, declines in income, or increases in consumption. This is important because long-term care co-payments and reductions in income following health shocks can be substantial, even in the Netherlands (García-Gómez et al., 2013). Therefore, on the face of it, these might explain the decline in wealth that we see. Previous studies of the dynamics of wealth near death have not quantified the role of transfers to heirs as we do. While Jones et al. (2020) draw on questions about medical expenses and bequests to non-spousal heirs, they do not use data that enable them to quantify inter-vivos transfers. Kvaerner (2017) examines the dynamics of the wealth of children around the time of death of their parents but is focused on parents diagnosed with cancer, who may face quite different information and changes in circumstances than those near death as a whole. We are able to demonstrate significant financial transfers in advance of death when looking at the whole population of those dying in a three-year period. Given the large sample size this yields, we are able to show that estate planning behaviour is not confined to the very rich but is the key driver of wealth declines of singles across the top half of the wealth distribution.

Our second major contribution is to show that the timing of transfers is highly responsive to tax incentives. This is possible because our dataset, uniquely in the literature, contains data on taxable gifts (i.e. annual gifts received from parents if they are larger than around €5000). We demonstrate that most estate planning behaviour involves tax-free gifts and, given the low threshold for inheritance tax, is likely tax-reducing and consistent with tax

avoidance as a motivation for giving. For the first time, we quantify the responsiveness of gifts to tax incentives to give in advance of death, estimating large Frisch elasticities of between 1 and 9. These findings demonstrate that tax avoidance through estate planning is substantial and responsive to policy. This behaviour is therefore important to consider in the design of systems of transfer taxation.

We make a further methodological contribution by showing the importance of using individual-level measures of length of illness, rather than cause of death alone, when categorising deaths as sudden (and unanticipated) or not-sudden (and potentially anticipated) in the context of estate planning. Our dataset contains insured health spending on each individual as well as cause of death. We show that those who had elevated health spending further in advance of their death began to change their behaviour sooner and showed larger changes in wealth by the year of their death. This is consistent with the findings and interpretation of Kopczuk (2007), whereby health shocks lead to revisions in longevity expectations and changes in behaviour and those who are ill for longer have more time to act and hence have a greater decline in their wealth by death. We show that using standard classifications of cause of death codes into ‘sudden’ and ‘non-sudden’ understates the degree of estate planning behaviour because some of those who die of causes classified as ‘sudden’ are already in deteriorating health well before their death.

The paper proceeds as follows. Section 2 gives a short summary of the relation of our paper to the existing literature. In section 3 we set out the institutional context and in section 4 we describe the dataset we construct. Section 5 sets out our empirical strategy for identifying the dynamics of wealth and transfers in anticipation of death and section 6 sets out our results of this analysis. Section 7 sets out our analysis of the drivers of transfers to heirs, including the estimation of the elasticity of gifts to taxation. In section 8 we conclude.

2 Related literature

A number of papers have examined changes in wealth around the time of death and the role of estate planning specifically. Kopczuk (2007) compared the size of the estates of those who died after a lengthy illness to those who died more suddenly. Finding that the latter left estates that were 10%-18% larger in value, and arguing that income and medical expense shocks could not explain the magnitude of difference, he concluded that wealthy individuals engage in ‘estate planning’ before death. Suari-Andreu et al. (2019) and Erixson and Escobar (2020) combine administrative data on cause of death with administrative data on wealth and estates in the Netherlands and Sweden, respectively, to implement a similar research design. Comparing those who died of ‘sudden’ causes of death to those who died of ‘non-sudden’ causes of death, Suari-Andreu et al. (2019) find that wealth at the start of the year of death is around 5% lower for both couples and singles in the latter group. However, in their sample of married individuals, Erixson and Escobar (2020) find no difference in estate size between those who died suddenly and those who died after a period of illness.

Within the literature looking at the dynamics of wealth around the time of death, our paper is methodologically related to two studies also using an event-study designs. Kvaerner (2017) examines the change in wealth that happens around the time of a diagnosis of cancer, finding that singles diagnosed with cancer respond by reducing their wealth by 30% to 50% but the wealth of couples does not respond. The decline in singles’ wealth is mirrored by an increase in the wealth of their children. Jones et al. (2020) examine the evolution of wealth over the final six years preceding death, comparing those who died to a control group observed over the same years but who died later. Jones et al. (2020) find that wealth declines before death for both singles and couples and that this fall can be explained by increased medical expenses and (in the case of couples) bequests to non-spousal heirs.

Our findings also contribute to the literature on estate, inheritance and gift taxation. Piketty and Saez (2013) shows that the steady-state elasticity of bequests to taxation is one crucial determinant of optimal inheritance taxation. Kopczuk and Slemrod (2001) and

Goupille-Lebret and Infante (2018) find modest responsiveness of accumulated wealth and estates to taxation. Glogowsky (2020) uses a bunching estimator to find a short-run elasticity of gifts to transfer taxation of below 0.1. However, this paper is in the context of a system of joint taxation of gifts and inheritances without any tax incentive to make transfers as gifts rather than bequests. McGarry (2001), Bernheim et al. (2004) and Joulfaian (2005) show that inter-vivos gifts appear responsive to the relative taxation of gifts and bequests, exploiting variation in tax rates across states and over time. Our evidence from bunching at gift tax kink points provides direct evidence of responsiveness to tax considerations and allows us to estimate an elasticity of gifts to taxation. Our large elasticity is likely driven by shifting of transfers between tax years and between gifts and bequests, meaning these are important margins to consider in tax design.

This paper also relates to a wider literature examining the drivers of wealth dynamics at older ages. Estimating structural lifecycle models of the retirement period of life, De Nardi et al. (2010), Ameriks et al. (2011) and Lockwood (2018) find that bequest motives are important for fitting observed wealth holdings of singles. De Nardi et al. (2021) find that bequest motives are important for explaining the behaviour of couples and high lifetime income singles. Banks et al. (2019) consider the role of medical expenses in driving retirement wealth trajectories in the UK and USA. Our paper provides more direct, quasi-experimental evidence that transfers to heirs are valued by much of the elderly population, including those with modest wealth, and that they are quantitatively important to asset dynamics near death. Our finding that richer parents transfer a larger share of their wealth to their children is consistent with the literature's finding that bequests are a luxury good (De Nardi (2004); Lockwood (2018)). We are able to look more closely into the nature of this bequest motive. Those with children engage most strongly in bequest-planning behaviour and make only slightly larger taxable gifts to their poorer children. This is consistent with role for altruism in driving the transfers of some households but suggests that most households can be modelled as receiving a 'warm glow' from the net-of-tax transfers they make to their heirs.

The fact that couples do not engage in such giving suggests a strong role for precautionary motives.

3 Institutional context

In this section we set out some of the details of policies in the Netherlands that are relevant to wealth and transfers near the end of life.

3.1 Inheritance and gift taxation in the Netherlands

The Netherlands has separate, but related, taxes for inheritances and *inter vivos* gifts. There are tax-free exemptions which differ across gifts and inheritances, but both types of transfers use the same progressive tax rates. Inheritance tax is levied on the net value of that part of an estate that is received by an individual.¹ Gift tax is levied on the annual flow of gifts received from an individual. Above the tax-free exemption level the ‘low’ tax rate is applicable to taxable transfers received up to a threshold (€121,296 in 2015), with the ‘high’ rate payable on further sums. The exemptions and tax rates depend on the relationship between giver and receiver and are shown for 2015 in Table 1. This system with two brackets has been in place since 2010, with the same rates and similar thresholds applicable in each year since then. Before 2010, the system had the same form but with a greater number of tax brackets, and different rates applicable, ranging from 5% to 27%.² This change in the structure of the

¹Taxable estates consist of all worldwide assets of those who were Dutch residents before passing away. Estates of Dutch nationals remain taxable until 10-years after emigration. Inheritances that are received from foreign residents are untaxed. Any gifts made in the final 180 days before death are also counted as part of the deceased’s estate. Before 2018, community of assets applied automatically upon marriage and a half of jointly-held assets are treating as the estate of a decedent with surviving spouse. Intestacy rules state that if the deceased leaves no will estates are divided equally between each child and a surviving spouse, if any. Up to 50% of these shares can be enforced by any disinherited child. Actual transfer of assets to children can be deferred until after the death of a surviving spouse.

²For full details of the inheritance and gift tax systems see Groot et al. (2019). Various other exemptions may be applicable. For example, for children between 18 and 40 there is a one-time exemption available of 25,322 euros, which can be increased to 52,752 if the gift is used for a study or for owner-occupied housing. Bequests that consist of a firm or substantial stake in a firm ($\geq 5\%$ of total shares) can be eligible for an additional exemption of €1,055,022 and 83% above. Bequests that are made to recognized charities are fully exempt from taxation.

transfers taxation system is used in the difference-in-bunching estimation set out in Section 7 and the rates and thresholds for the taxation of gifts from parents to children under both tax ‘regimes’ are set out in Appendix A.

Table 1: Exemptions and tax rates for inheritance and gift taxation in 2015

Relationship of receiver to giver	Exemption threshold for inheritance tax (Euros)	Exemption threshold for gift tax (Euros)	Low tax rate	High tax rate
Partner	633,014	2,111	10%	20%
Child	20,047	5,277	10%	20%
Grandchild	20,047	2,111	18%	36%
Parent	47,477	2,111	30%	40%
Other	2,111	2,111	30%	40%

Note: The exemption thresholds give the amount of inheritance and gifts that can be passed on from receiver to giver without any inheritance or gift tax being payable. For all relationships between giver and receiver, tax is payable on that part of inheritances and gifts that are above the tax-free threshold. In 2015, the first €121,296 above the tax-free threshold was taxable at the ‘low’ rate. The ‘high’ tax rate was payable on that part of inheritances or gifts that exceed the tax-free threshold plus €121,296.

Importantly for this paper, the fact that a progressive system is applied separately to gifts received each year and to inheritances means that spreading a given level of transfers over a number of years can reduce the total of tax due. By spreading transfers over time, the yearly exemption for gifts can be used in each year and the amount of transfers falling into a higher tax bracket can be reduced. If there was a tax levied instead on lifetime transfers received, as proposed by Mirrlees et al. (2011b), for example, such tax avoidance would not be possible.

A second important feature is that while inheritance tax is levied on the part of the estate left to a surviving partner, married couples are assumed by the law to hold all assets jointly. Couples therefore cannot in general reduce inheritance tax liability for the surviving partner

by transferring assets between themselves.

3.2 Health and long-term care insurance

A uniform insurance for regular medical costs is mandatory for Dutch households. For most medical procedures, no co-payments are required. Annually, the first 375 euros (in 2015) of medical expenses are paid out of pocket. The mandatory insurance has an annual premium (depending on the insurance provider), but is mostly financed by income-dependent contributions collected by the tax authority. At the point of use, payments for medical care are therefore minimal.

Long-term care, such as that in a residential care home or nursing home setting, is publicly funded but can require co-payments that depend on the income and the net wealth of the person receiving care. For the purpose of these co-payments, net wealth is based only on financial wealth, real estate not used as the primary home, other assets and debt other than the mortgage on the primary home. A ‘low’ co-payment applies when the care is received at home, for short-term stays in residential care (up to four months) and for long-term stays in residential care while a partner is still living at home. In other cases the ‘high’ co-payment regime applies. In 2020, the ‘low’ co-payment ranged between 168 euros and 882 euros per month and the ‘high’ co-payment ranged between 0 and 2,419 euros per month. When the maximum co-payment is not yet reached, the annual marginal contribution of net wealth is 0.4% in the low regime and 4% in the high regime. Co-payments therefore have the potential to lead to substantial reductions in wealth.

4 Data and sample

This paper uses anonymized administrative micro data which are made available by Statistics Netherlands (CBS). We start by selecting all individuals who lived in the Netherlands and died between 2007 and 2019. For these individuals we use personal characteristics including

age, gender and household type and merge this with fiscal data at the household level on income, wealth and gifts as well as data on health-spending, long-term care co-payments and cause of death.

Wealth is registered at the household level, as partners declare their combined wealth. The wealth measure that we use is total net wealth excluding pension wealth (i.e. the sum of all property wealth and debts, all financial assets and debts and all business assets and debts). Wealth values are generally either directly reported by financial institutions or are independent valuations. Wealth is measured on the first of January of each year. For further details of the wealth measure see Appendix B. For each individual, we use data on family relationships to construct a measure of their children’s household wealth by combining the wealth of the households of any children that they have and who live in the Netherlands. We construct a measure of their grandchildren’s household wealth analogously.

Our income measure is pre-tax income (including social transfers and pension income) aggregated to the household level. Gifts are aggregated to the level of the giver’s household. We can only observe gifts for which a tax statement was filed. Filing a tax statement is only required for gifts above the yearly exemptions or for gifts that use a one-time exemption (see footnote 2 for an overview of the various exemptions).

Finally, we merge in variables for the individual’s annual insured health spending, annual long-term care co-payments, and International Classification of Diseases code for cause of death (ICD-10).

For the event-study analysis, the treatment group is drawn from individuals for whom all data is available for the seven years up to and including the year of death and who could be matched to someone who died 4 years later than them. This implies selecting individuals who died between 2013 and 2015. The control group for the event study is drawn from the individuals who died between 2017 and 2019 and for whom all data is available for the seven years up to and including the fourth year prior to their death.³ For the bunching analysis,

³There are exceptions in the case of insured health spending data, which is available only from 2009, and long-term care co-payments, data which is available only from 2015. As will become clear when we introduce

we draw on all observations of individuals who died between 2007 and 2019 and who made at least one gift to a child over the period 2007 to 2016 (the period for which the gift data is available).⁴

Table 2 shows the descriptive statistics for the treatment group used in the event-study analysis. We note that the wealth distribution is very skewed and that the median single has very modest net wealth, of less than €24,000.

Table 2: Summary statistics for key variables in the treatment group sample

	Mean	25th percentile	Median	75th percentile
<i>Singles (N=106,566)</i>				
Age	78.0	73	80	86
Annual taxable gifts (,000)	1.3	0	0	0
Insured health spending (,000)	3.5	.1	.6	2.9
Annual income (,000)	22.5	14.3	17.5	25.2
Net wealth (,000)	137.7	6	24.1	167.9
<i>Couples (N = 155,587)</i>				
Age	70.8	64	72	79
Annual taxable gifts (,000)	1.9	0	0	0
Insured health spending (,000)	5.2	.1	.9	4.4
Annual income (,000)	47.7	25.0	35.2	57.6
Net wealth (,000)	274.4	23.5	144	324.3

Note: Monetary amounts are in Euros. Sample is those individuals who died in 2013, 2014 and 2015 and were included in the matched sample, as described in section 5.

our methodology, we do not require the health spending data for 2007 and 2008 and our method accounts for the period for which co-payment data is available.

⁴We could have restricted the sample for the bunching analysis to include only those included in the event-study analysis but this would have reduced the precision of our estimates without obvious benefits.

5 Empirical strategy for identifying the dynamics of wealth and transfers near death

Identifying estate planning - that is to say wealth transfers to heirs that occur near death in the anticipation of death - presents a number of empirical challenges. The first is to quantify the change in wealth that is associated with proximity to death, disentangling this from age and time effects, and other factors that impact wealth and will be changing as death approaches. It is known that those who are wealthier throughout their lives tend to live longer (Attanasio and Emmerson, 2003). Therefore, the probability of death at a given age and time is likely endogenous to an individual's wealth level, and a simple comparison of the wealth levels of those who do, and who do not, die at a given age, or in a given time period, would likely capture factors other than the fact that death is approaching.

One approach has been to exploit variation in the time available to individuals to conduct estate planning and to compare the estates of those who died after a period of prolonged illness to the estates of those who died suddenly and to attribute the difference to estate planning (Kopczuk (2007); Erixson and Escobar (2020)). This type of research design relies upon having data that can be used to identify whether (and potentially for how long) individuals had a health condition that made their death more likely. For reasons we set out in the next section, standard cause of death classifications alone do not appear to cleanly identify those for whom death arrived 'suddenly' and so we do not believe this design is appropriate given our data.

Our approach is to exploit variation in the timing of death, conditional on characteristics more than four years before the time of death. We employ a matched 'event study' or 'dynamic differences-in-differences' design, where we compare the change in the wealth of a group of individuals over the final four years before their death to the change in wealth of a second group of individuals who died four years later (see Fadlon and Nielsen (2019) for an example of this event-study with matched control group design and Jones et al. (2020) for an

application to the dynamics of wealth in anticipation of death). Our difference-in-differences approach accounts for any differences in wealth levels between the two groups in the baseline year (the period four years before the death of the individual who died). It remains possible that, even in the counterfactual scenario where they were not close to death, the wealth of those who are in fact close to death would change at a different rate than those who die at a later point. For example, it could be that those with lower wealth or lower income tend to spend down their wealth more rapidly at older ages than those with higher wealth or income, even in the absence of proximity to death, due to greater discounting of the future. An association between wealth profiles and mortality of this kind would bias the results of an estimation that compared the change in wealth of all those who died in a given period to the change in wealth of all those who died four years later. To alleviate this concern, we conduct an initial matching step, where we attempt to match all individuals who died in a given time period to an individual with similar characteristics in the baseline year but who died four years later. We then conduct the difference-in-differences estimation using this matched sample.

A necessary condition for identification is that had our treatment group died four years later than they actually died, their outcomes would have evolved in the same way as for those in the control group. Given our matching approach, this is implied by an assumption of parallel trends of treatment and control groups in the counterfactual where both were four or more years away from death, conditional on the variables that we match on. This assumption is plausible to the extent that we accept that, conditional on covariates used for matching four years before the death of the treated observations, whether an individual dies in four years or eight years time is as-if randomly assigned. This would be the case if, for example, four years before the death of the treated individuals, the treatment and control groups faced the same probabilities of receiving a health shock that would lead to their death within four years, and the treatment group contains those individuals for whom such a shock was then realised.

A second key assumption of our approach is that the effects of anticipation of death begin at most four years before the beginning of the year of death. In reality there will be some individuals who consider themselves to be close to death more than four years before they die. This could happen in the case of a prolonged terminal illness, for example. If this anticipation affects behaviour, it would contaminate our control group and be differenced out from our estimates, most likely biasing them towards zero. There is no theoretical restriction to how far in advance of death anticipation effects could occur and so any restriction here will be to some extent arbitrary. However, we show in Section 6.1 that treatment and control groups follow parallel trends in the ‘pre-periods’ five and six years before death. We experimented with allowing effects of anticipation of death up to five years in advance of death and this only marginally increased the size of our estimated effects. We are therefore confident both that any attenuation due to contamination of the control group will be very small and that we capture all, or almost all, of the effect of proximity to death in our estimates.⁵

We emphasise two related points. First, our research design does not require that individuals predict exactly when they will die or that all individuals have increased expectations of their own death as their own death approaches. There will be some individuals in our treatment group for whom their death comes unexpectedly, whereas for others it is preceded by a deterioration in health that indicates reduced longevity. The difference in the evolution of wealth between our treatment and control groups will be driven by any greater prevalence in health shocks and associated changes in wealth amongst the former compared to the latter. Second, this means that our estimand is an average of the change in wealth attributable to proximity to death over the population of those who die, including some for whom there is no anticipatory behaviour. This makes our estimand different to others in the literature that either use sudden deaths as a control group or focus in on those who die after a lengthy

⁵We note that we could have chosen a control group that died more than four years later than the treatment group to reduce the likelihood of any such contamination of the control group and expand the time horizon of our analysis. However, we judge that this may have increased the possibility that these groups were then more different in other unobserved ways and would also have reduced our sample size. There also appears to be little gain from doing so, either in reducing contamination or in capturing effects at longer time horizons.

illness. We view it as a virtue that our estimates pertain to the average estate planning behaviour across the whole population as this informs us whether this planning is sizeable in the aggregate and therefore of import for policy.

The second main empirical challenge is to separate the change in wealth in advance of death into parts attributable to financial transfers to children and parts attributable to other effects such as changes in income, medical or long-term care expenses or consumption. We address this in the following way. First, in order to identify the proportion of the change in parents' wealth that is due to transfers to their children, we analyse the effect of proximity to death on the wealth of children of our sample members, on 'dynasty' wealth (the sum of an individual's wealth and the household wealth of all of their children), and on taxable gifts. These approaches are complimentary: the gifts data is a direct measure of these transfers, but transfers below the exemption threshold are not covered in our dataset. Analysis of children's wealth and dynasty wealth provides a comprehensive measure of transfers made, under the additional assumption that children do not experience a change in their income, spending decisions, or returns to wealth as a result of their parent being close to death. Second, we are able to directly test for and rule out any substantive change in income as death approaches. Third, as set out in section 3, there is very little private medical expenditure in the Netherlands, so we can safely assume that the effect of proximity to death on medical spending and therefore wealth is essentially zero. Fourth, we are able to use the long-term care co-payment data to construct an upper bound for the contribution of care co-payments to the decline in wealth. As we have care co-payment data from 2015 onwards only, we are unable to implement our matched event-study design for this outcome. Instead, we obtain an upper bound for the effect of co-payments on wealth based on the cumulative amount of co-payments made over the four years up to the start of 2019, for those who died in that year. Starting at a low level four years before death, care co-payments increase rapidly as death approaches and so we believe we are likely to obtain a reasonably tight upper bound. We treat any residual change in wealth not explained by these prior factors as being driven

by a change in consumption.

5.1 Matching

We now give the details of our matching process. We begin with all individuals who died in 2013, 2014 or 2015 and who had positive wealth at the beginning of the year four years before the year in which they died. We refer to these individuals as ‘treated’ individuals. For each treated individual, we identify the group of individuals who died four years later and who had the same values of the following characteristics four years prior to the year in which the treated individual died: age (in 2-year categories), household type (single or couple), wealth quantile (2 percentage point quantiles), income decile, and insured health spending quintile.⁶ Where this process finds a unique individual, they are assigned to the control group. Where there are multiple such individuals, one is picked at random. If no individual is found with the same characteristics, we relax the match on age to 5-year groups and the match on wealth to 5 percentage point quantiles. If there is still no match for a treated individual, they are dropped from the treatment group. Matching is without replacement, so that there is the same number of unique individuals in the matched treatment and control groups. This process matches 81% of treated individuals to a control individual.

5.2 Specification

Using the matched sample, we conduct our event study estimation as follows. We re-define the year of death of an individual in the control group to be equal to their actual year of death minus four years. We then pool all observations of our treatment and control groups over the final 6 years up to and including the year of their death and estimate specifications of the following form with OLS:

⁶Quantiles are taken within each year.

$$y_{i,t} = \alpha_i + \sum_{\substack{\tau=-6 \\ \tau \neq -4}}^0 \beta_{\tau}^C \cdot \mathbb{1}\{t = t_{Di} + \tau\} + \sum_{\substack{\tau=-6 \\ \tau \neq -4}}^0 \beta_{\tau}^{ES} \cdot \mathbb{1}\{t = t_{Di} + \tau\} \cdot T_i + \epsilon_{i,t} \quad (1)$$

where $y_{i,t}$ is the outcome e.g. log household net wealth, t_{Di} is i 's time of death and T_i is an indicator for being in treatment group. We exclude the indicator variable for four years before death, $\tau = -4$, which is therefore the reference period.

The coefficients $\{\beta_{\tau}^C\}_{\tau=-6}^{\tau=0}$ recover the average change in the outcome in the control group, relative to the reference period, by number of years until death. Our coefficients of interest are the event-study coefficients $\{\beta_{\tau}^{ES}\}_{\tau=-6}^{\tau=0}$, which recover the *difference* between the change in the treatment and change in the control group, for each year until death. These coefficients are interpreted as the effect of proximity to death on the outcome. Following Abadie and Spiess (2021), we cluster standard errors at the level of the matched pair to account for the matching step when making inference.⁷

As is set out in the following section, we examine heterogeneity in effects by performing estimation on subgroups defined by characteristics that we match on. For example, all of our estimation is carried out separately for singles and couples and we also show results split by wealth quartiles four years before death. As we will go on to focus on individuals with children, we perform a second match where we include being a parent as a covariate upon which we match.⁸

Finally, to examine taxable gifts as an outcome on a comparable basis to log wealth, we define the ‘gift share’ as cumulative taxable gifts made over the period since the start of the

⁷Clustering at the level of the matched pair gives very similar standard errors to clustering at the individual level.

⁸We match on being a parent only when we turn to examine child wealth in order to increase precision and so that log child wealth is defined for all those in our control group. We do not match on whether or not the individual is a parent when first analysing the wealth of all individuals because the curse of dimensionality leads to a sharp reduction in the match rates for those without children. The lack of any pre-trends in the event-study estimates for the wealth of non-parents is reassuring that any matching of parents to non-parents is unlikely to bias our results.

year four years before the year of death, as a percentage of wealth plus the cumulative gifts:

$$GiftShare_{i,t} = \frac{\sum_{\tau=t_{D_i}-4}^{t-1} TaxableGifts_{i,t}}{W_{i,t} + \sum_{\tau=t_{D_i}-4}^{t-1} TaxableGifts_{i,t}} \quad (2)$$

where $TaxableGifts_{i,t}$ are the taxable gifts made by individual i in year t and $W_{i,t}$ is individual i 's wealth at the start of year t . This variable defines the percentage reduction in wealth that can be attributed to taxable gifts given.


6 Results: the dynamics of wealth and transfers

We now set out our main results. We first analyse the change in wealth that happens as death approaches before turning to examine the change in wealth of the children of those who die, the change in the dynasty's (parents' plus children's) wealth, and the change in taxable gifts. In the main text we focus on results for the effects on log wealth. The wealth distribution is highly skewed and so the analysis of wealth in levels has the potential to be more affected by a small number of observations. Results using levels of wealth are reported in the appendix.

6.1 Effect of proximity to death on wealth

Figure 1 shows the mean levels of log wealth in our treatment and control groups and the 'event study' or 'dynamic difference-in-differences' estimates for single individuals. Log wealth is declining for both groups over the period (covering 2009 to 2015, during which house prices declined in the Netherlands). From 3 years before death, a difference opens up, with the treatment group holding 0.7% less wealth (we use the approximation that log point deviations can be interpreted as percentage differences). This difference expands to 3.1% by the beginning of the year before the year of death and then reaches 6.9% by the beginning of the year of death.

Figure 1: Levels of log wealth and event-study estimates for single individuals

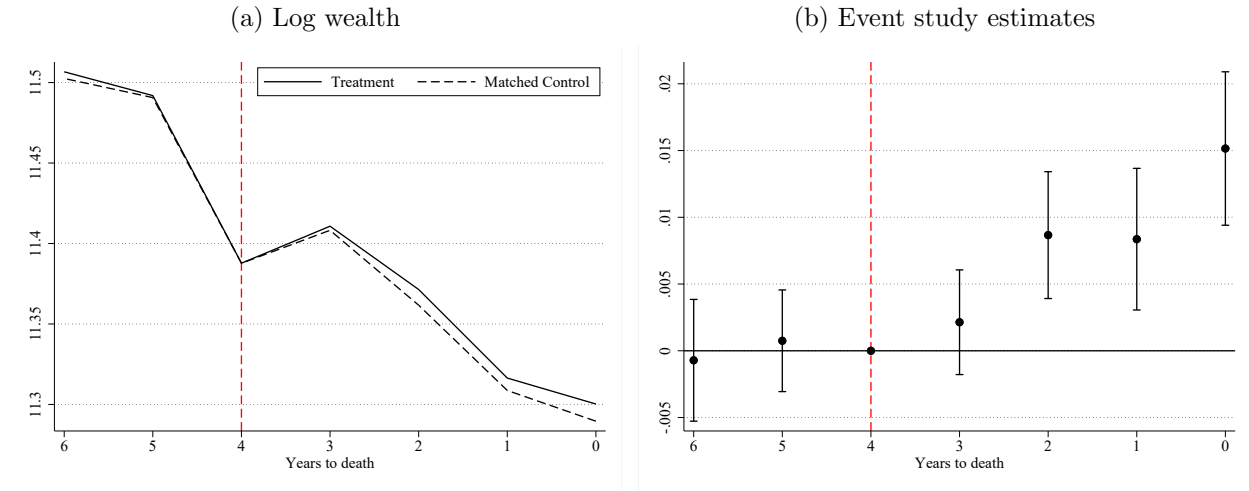


Note: Panel (a) shows the mean values of log wealth of singles, by years until death, for both the treatment and matched control groups. The treatment group consists of those who died between 2013 and 2015 and had positive wealth 4 years before death. The control group is matched based on characteristics 4 years before death. Panel (b) shows the event study estimate for the difference between the treatment and control group in the change in their log wealth relative to the base period 4 years before death. We show 95% confidence intervals, with standard errors clustered at the match level. Sample size is 213,132 individuals.

We show the mean levels of log wealth and the event study estimates for the two periods prior to the base period, five and six years before death. This shows a very similar path of wealth among the treatment and control groups. This is consistent with negligible differences in anticipation effects more than 4 years before death and suggests that our estimates capture all quantitatively important effects of proximity to death. This similarity in pre-trends is also encouraging evidence that our matching method creates two groups who, in the absence of proximity to death, would see their wealth change in the same way.

Figure 2 shows the main results for couples. Again, wealth is declining for both groups over time, but among couples wealth declines more slowly for those approaching death. By the year of death, wealth is 1.5% higher as a result of being closer to death. Again, we find no significant differences in pre-trends.

Figure 2: Levels of log wealth and event-study estimates for couples




Note: Panel (a) shows the mean values of log wealth of those in couples, by years until death, for both the treatment and matched control groups. The treatment group consists of those who died between 2013 and 2015 and had positive wealth 4 years before death. The control group is matched based on characteristics 4 years before death. Panel (b) shows the event study estimates for the difference between the treatment and control group in the change in their log wealth relative to the base period 4 years before death. We show 95% confidence intervals, with standard errors clustered at the match level. Sample size is 311,174 individuals.

6.1.1 Wealth changes as a response to health shocks

We next present evidence that supports our interpretation of the difference-in-difference estimates as capturing the effects of the arrival of health shocks and deterioration in mortality prospects. Our ultimate interpretation will be that these health shocks bring news about mortality and lead to a change in behaviour. Figure 3 shows our results splitting our sample by the year in which they first had elevated insured health spending. We define elevated health spending as having annual spending above the 75th percentile for individuals who died in that year. This implies spending of around €10,000 or more. We can see this as a proxy for when an individual has received news about a deterioration in their health. Even though this is an imperfect proxy, the pattern in the results is quite clear: singles who have higher health spending from an earlier point in time see their wealth decline earlier, and by a larger amount by the year of their death. While singles who had elevated health spending

Figure 3: Event-study estimates for the effect on log wealth, split by how many years in advance of death the individual first had high healthcare costs



Note: Panel (a) shows, for single individuals, the event study estimates for log wealth, splitting the sample by the year that the individual first had elevated healthcare costs. Those who always had elevated healthcare costs four years before death or who never had elevated healthcare costs are not shown. We show 95% confidence intervals, with standard errors clustered at the match level. 88,652 individuals. Panel (b) shows the equivalent figure for couples. 162,056 individuals.

beginning three years before their death had reduced their wealth by 11% by the year of their death, those who did not have elevated spending until the year they died had no significant change in their wealth by that point. In the case of couples, the increase in wealth is greater for those who first had elevated health spending further in advance of death, although the differences between groups are not statistically significant.

We emphasise that out-of-pocket health costs are very minimal in the Netherlands and so hospital spending will not have a significant *direct* impact on wealth. However, health shocks could be correlated with a range of changes that could reduce wealth, such as entry into long-term care and increases in associated co-payments, a decline in income, or an increase in consumption. It is therefore important that we provide direct evidence on the role of transfers to heirs in explaining the wealth decline and disentangle this from any other potential drivers of the decline in wealth.

6.1.2 The importance of individual-level measures of suddenness of death

Before turning to the results for the wealth of children, we set out our methodological contribution to understanding how best to identify the effects of anticipatory behaviour in advance of death. Our dataset contains the ICD-10 classification of cause of death. This classification has been used by Andersen and Nielsen (2011) and Andersen and Nielsen (2017) to identify a subset of deaths as due to sudden - and hence unexpected - causes of death. Using this data, we classify each individual in our dataset as dying from a ‘sudden’ or ‘non-sudden’ cause, as defined in Andersen and Nielsen (2011). We then examine the evolution of health spending and wealth in advance of death for those who died of a ‘sudden cause’, comparing them both to those who died of ‘non-sudden’ causes of death and those who died 4 years later than them. The purpose of this is to assess whether those individuals whose deaths were classified as ‘sudden’, based on their ICD-10 cause of death, experience any change in health or wealth that may be indicative of worsening health and anticipatory behaviour in the years before their death.

We find that average health spending rises before the year of death for those who die of ‘sudden’ causes of death. Though much smaller than the rise in spending for those who die of causes classified as ‘non-sudden’, this suggests that some of those who die of causes defined as sudden already had declining health before the year of their death (see Appendix C.2 for details). This is plausible when we consider that some of the causes classified as ‘sudden’ include surgical complications, falls, and heart attacks. It seems possible, or even likely, that some individuals experiencing these events would already be in poor health and have faced a relatively high chance of death even in the absence of the events that were the proximate cause of their death.

The fact that those who die of ‘sudden’ causes are in declining health before their death appears to matter for the evolution of their wealth. Using our matched event-study design, we find that those who die of a ‘sudden’ cause of death see a decline in their wealth as death approaches, compared to a group who died four years later. Further, we show that

estimating the size of estate planning behaviour as the difference between the final wealth of those who die of causes of death classified as ‘sudden’ and ‘non-sudden’ understates the degree of estate planning of singles in our setting (see Appendix C.2 for details).


These findings suggest that it is important to use individual-level information that can more closely capture how long an individual is in ill-health before their death, if exploiting the ‘sudden deaths’ research design. For example, direct measurement of time spent in ill-health is used in Kopczuk (2007) and Erixson and Escobar (2020), while Suari-Andreu et al. (2019) remove some individuals from the ‘sudden death’ group if they were previously hospitalised with the condition that led to their death. It is not clear that exploiting cause of death alone is well-suited for classifying deaths as anticipated or unanticipated.

We attempted to find a group of individuals whose cause of death was very likely unanticipated and unrelated to their health by examining those who died in transport accidents (ICD-10 codes V01-V99). Appendix Figure 12 shows that the group who died in transport accidents did not see any marked increase in their health spending until the year of their death. However, the sample size of transport accident deaths is too small to yield precise estimates of the difference in the evolution in wealth between this and other groups. Therefore, this is not a fruitful avenue for analysis. For these reasons, we believe that the method we have adopted, which exploits differences in the timing of death, may be useful for future research in this area.

6.2 Quantifying the role of transfers to heirs in advance of death

As a step towards assessing whether the declines in wealth in advance of death may be as a result of transfers to children, we show the estimated effects of proximity to death on wealth, splitting our sample into those who have children and those who do not. Figure 4 shows that the effects found amongst singles are much larger, and start further in advance of death, for those who have children. For singles with children, wealth begins to decline three years before death and is 8.7% lower by the beginning of the year of death. For those without

Figure 4: Event-study estimates for the effect on log wealth, for parents, and non-parents



Note: Panel (a) shows, for single individuals, the event study estimates for log wealth for the whole sample, and for those with and without children. We show 95% confidence intervals, with standard errors clustered at the match level. 158,227 individuals with children; 54,905 individuals without children. Panel (b) shows equivalent results for couples. 277,456 couples with children; 33,718 couples without children.

children, wealth first shows a small (not statistically significant) rise of 1.2% 2 years before death, before declining to be 3.1% lower by the start of the year of death. For couples, we do not observe statistically significant differences between those with and without children, but our estimates for those without children are imprecise because only a small proportion of couples do not have any children.

These results are consistent with a larger behavioural change for singles with children than those without children. This findings is in contrast to Hurd (1987), who found that those with and without children did not display differential motives to leave bequests. We note that the modest decline in wealth for those without children could be consistent with some motive to pass on wealth to others, such as siblings or other relatives.⁹ However we do not explore this further.

⁹This would be consistent with Kopczuk and Lupton (2007), who find evidence of a bequest motive for those with and without children.

6.2.1 Children’s wealth, dynasty wealth, and taxable gifts

We now turn to quantify how much of the decline in the wealth of singles with children over their final years is explained by transfers to their children. We do this by using our empirical framework to examine the effect of proximity to death on three outcomes: the wealth of children, dynasty wealth (sum of the parent’s and children’s wealth), and taxable gifts.

Under the assumption that children’s wealth is affected by their parents’ proximity to death only through an effect on transfers from parents to children, estimating the effect of parents’ proximity to death on children’s wealth recovers the size of transfers made. We can then estimate the size of the parents’ wealth decline explained by this effect on transfers by comparing the effect on children’s wealth to the effect on parents’ wealth or by examining the effect on dynasty wealth. If transfers explain all of the wealth decline, then children’s wealth will rise by a corresponding amount and dynasty wealth will be unchanged.


The taxable gifts data allows us to directly examine transfers made before death. Although taxable transfers do not include those gifts made below the tax-free threshold, examining this outcome allows us to check whether our results based on children’s wealth are corroborated qualitatively and, by comparing with the results for children’s wealth, to estimate how much of the increase in transfers is due to an increase in taxable gifts and how much is due to giving below the tax-free threshold.

Any decline in parents’ wealth that is not explained by transfers to children must be attributable to some other change in income or other expenditure. We show in Appendix C.4 that we can rule out any changes in parents’ annual income due to proximity to death of greater than 0.3% for singles and greater than 0.7% for couples. Given that private expenditure on health care is very small, we also rule out any substantive change in this component of spending. We are able to place an upper bound on the share of the wealth decline that can be explained by increased long-term care co-payments.

Figure 5 sets out our results. In order to make the percentage changes in parents’, children’s and dynasty wealth and share of wealth given as gifts comparable, we multiply

both the estimated effects on parents' wealth and the share of their wealth given as gifts by the mean level of dynasty wealth divided by the mean level of parents' wealth. We rescale the estimates for children's wealth analogously. We see that mirroring the reduction in singles' wealth of 8.7% by the year of death (which is 1.4% when expressed as a share of dynasty wealth), there is an increase in the wealth of their children of the same magnitude. There is no statistically significant change in dynasty wealth at any time horizon. Our point estimate is that the effect on children's wealth explains 92% of the parental wealth decline at the start of the year of death. Taxable gifts are estimated to have reduced singles' wealth by 0.7%, or 8% of the magnitude of increase in children's wealth. We therefore infer that over 90% of the value of gifts made is below the tax-free threshold.

Figure 5: Event-study estimates of the effect on singles with children's log wealth, children's log wealth, dynasty log wealth and taxable gifts, as a share of dynasty wealth



Note: Figure shows for singles with children, the event study estimates for the effect on parents' log wealth, children's log wealth, dynasty log wealth and taxable gifts, all expressed as a percentage of dynasty wealth. We show 95% confidence intervals, with standard errors clustered at the match level. 152,756 individuals.

The extent of the decline in parents' wealth and rise in children's wealth differs signif-

icantly by the initial wealth level of the parent. We do not find an effect of proximity to death on the wealth of those parents who were in the bottom wealth quartile four years before their death. Given that the average initial wealth of this group was around €3000, this is unsurprising. However, by the year of their death, the decline in wealth for those in the second, third and top quartiles is 9.2%, 13.1% and 16.7%, respectively (Appendix Figure 12 shows the full set of event study estimates split by initial wealth quartile).

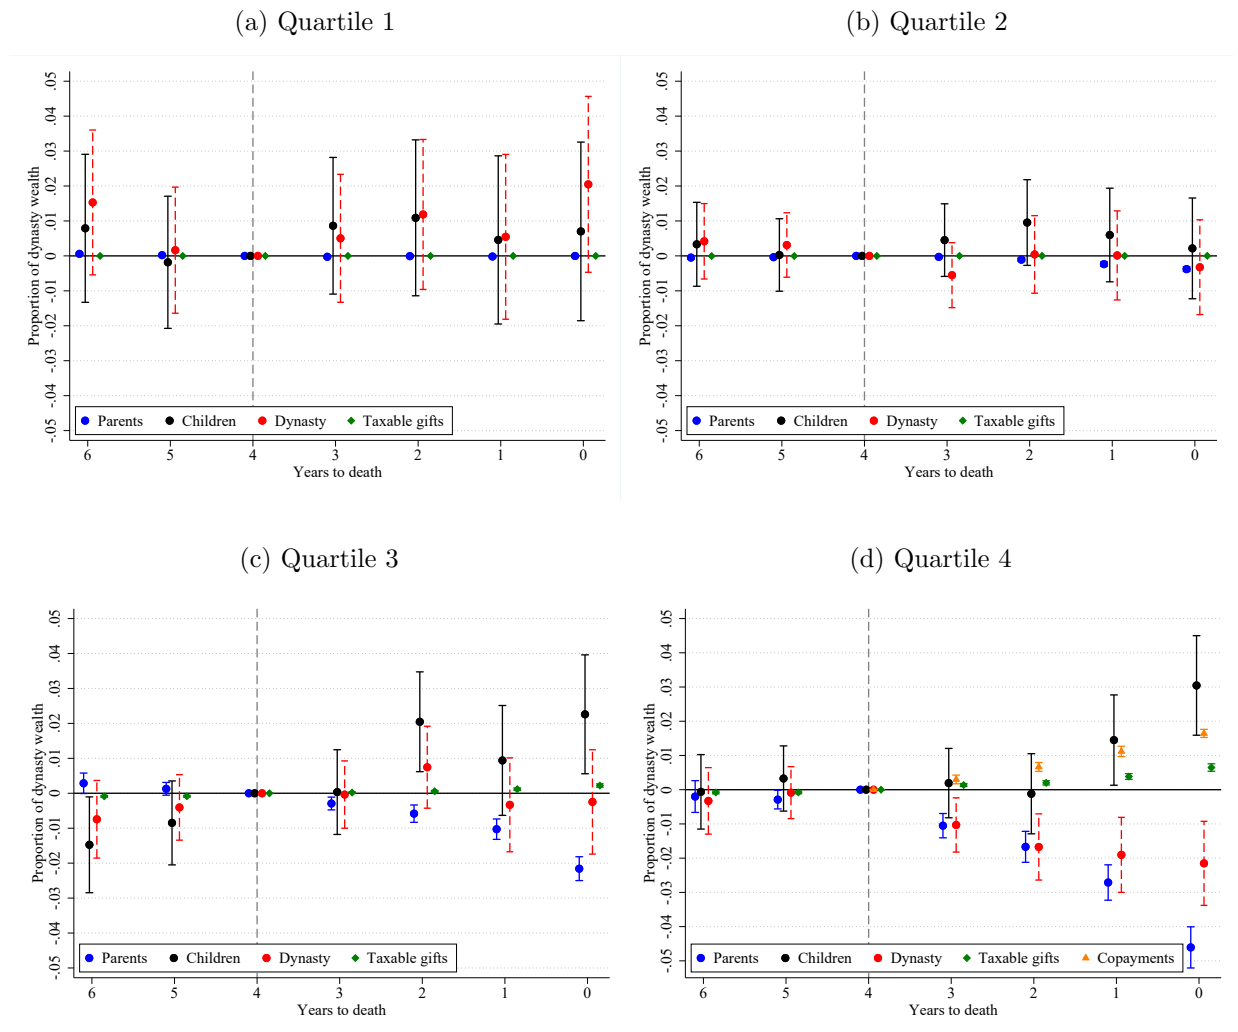
Figure 6 shows the estimated effects of proximity to death on wealth of parents, their children, their dynasty, and on parents' taxable gifts, split by the parents' initial wealth quartile and expressed as a percentage of dynasty wealth. In the bottom quartile, there is no significant change in any outcome as death approaches. In the second quartile, there is a decline in the wealth of parents but we do not detect any effect on the wealth of their children or their dynasty or see a rise in taxable gifts. For this group, parents' wealth is too small (averaging €22,000 four years before death) relative to changes in children's wealth for us to assess whether the decline in wealth reflects transfers to children, but any transfers cannot be very large. In the third quartile, there is a significant increase in children's wealth, which can fully explain the decline in parents' wealth, and a significant rise in taxable gifts. Our point estimate is that dynasty wealth does not change as death approaches, and our findings are consistent with transfers to children explaining the entirety of the decline in parents' wealth. This is strong evidence of estate planning behaviour for those in the third quartile, whose wealth is less than €168,000. The estimated effect on taxable gifts for this group is 9.8% of the size of the increase in children's wealth, indicating that most giving is below the tax-free threshold.

For parents in the top wealth quartile, there is a substantial increase in their children's wealth. This explains 66% of the decline in parents' wealth and dynasty wealth declines by 2%. In the wealthiest quartile, transfers to children therefore appear to be the primary but not only explanation for parents' wealth declining in advance of death. Taxable transfers explain 21.2% of the rise in children's wealth, implying significant transfers below the tax-

free amount, consistent with tax-minimising behaviour.¹⁰ Appendix C.6 shows that wealth declines by a similarly large proportion over the final years for those in the top 10% and top 1% and that transfers to children also explain the majority, but not all, of this wealth decline for these very wealthiest.

¹⁰Quantitatively, the rise in children's wealth that is not explained by taxable gifts could all be explained by gifts below the tax-free amount. In all quartiles, if we assume that, in each of the final four years before their year of death, parents who don't report a taxable gift make a gift equal to the tax-free threshold to each of their children, this would more than fully account for the rise in children's wealth.

Figure 6: Event-study estimates of the effect on singles with children’s log wealth, log children’s wealth and log dynasty wealth, by quartile of wealth



Note: Each panel shows, for single individuals with children, the event study estimates for the effects on log wealth, log wealth of children, log wealth of the dynasty, and cumulative taxable gifts as a share of wealth, split by the parents’ wealth quartile 4 years before death. Panel (d) also shows the upper bound on the effect on cumulative care copayments as a share of wealth. Effects for each outcome are expressed as a share of dynasty wealth. We show 95% confidence intervals, with standard errors clustered at the match level. 152,756 individuals.

6.2.2 Explaining the residual decline in wealth

What other expenditures could explain the residual decline in the wealth of parents in the top quartile? Giving to grandchildren is one possibility. However, in Figure 13 in the appendix

we show that the decline in dynasty wealth for the top quartile group is unchanged with the addition of the wealth of grandchildren to our dynasty wealth measure.

Copayments for long-term care can be substantial (up to 4% of net wealth per year) and are more likely to occur in advance of death. They are therefore a potential explanation for the remaining decline in wealth. In panel (d) of Figure 6, we show an estimate of an upper bound on the effect of care copayments on wealth (the estimation of this upper bound is set out in Appendix D). For those singles in the top wealth quartile, we estimate that care copayments over the final four years before the year of death can explain at most a 6% reduction in the wealth of single parents in the top quartile, equivalent to just under 2% of dynasty wealth. This upper bound is equal to the decline in dynasty wealth and so we cannot rule out that care copayments explain all of the wealth decline that is not attributable to transfers to children.

7 The drivers of transfers made before death

Given that transfers to heirs explain the vast majority of the decline in wealth of singles in advance of death, we examine what drives these transfers.

While the evidence shown thus far is consistent with a tax-avoidance motivation for giving, there are other reasons why transfers might increase when mortality looks to become imminent. For example, individuals might get more enjoyment from giving wealth when they are alive, and can see the impact it has on the receiver or exert more control over how the gift is used, compared to giving as a bequest (Kopczuk, 2007). We therefore explicitly estimate the responsiveness of giving to taxation using a bunching estimation approach.

Whether or not transfers respond to tax incentives, there are multiple possible reasons why people may value making transfers to heirs in the first place. For example, the literature has suggested that transfers may be made as part of within-family insurance (Attanasio et al., 2018), in exchange for care services (Bernheim et al., 1985), because of altruism towards

heirs (Altonji et al. (1992); Altonji et al. (1997)), or, simply, because of a direct “warm glow” enjoyment of transfers given. Disentangling the role of each of these motivations is beyond the scope of this paper. However, we assess whether there is evidence that transfers are altruistically motivated by testing whether parents make larger transfers to children with lower wealth.

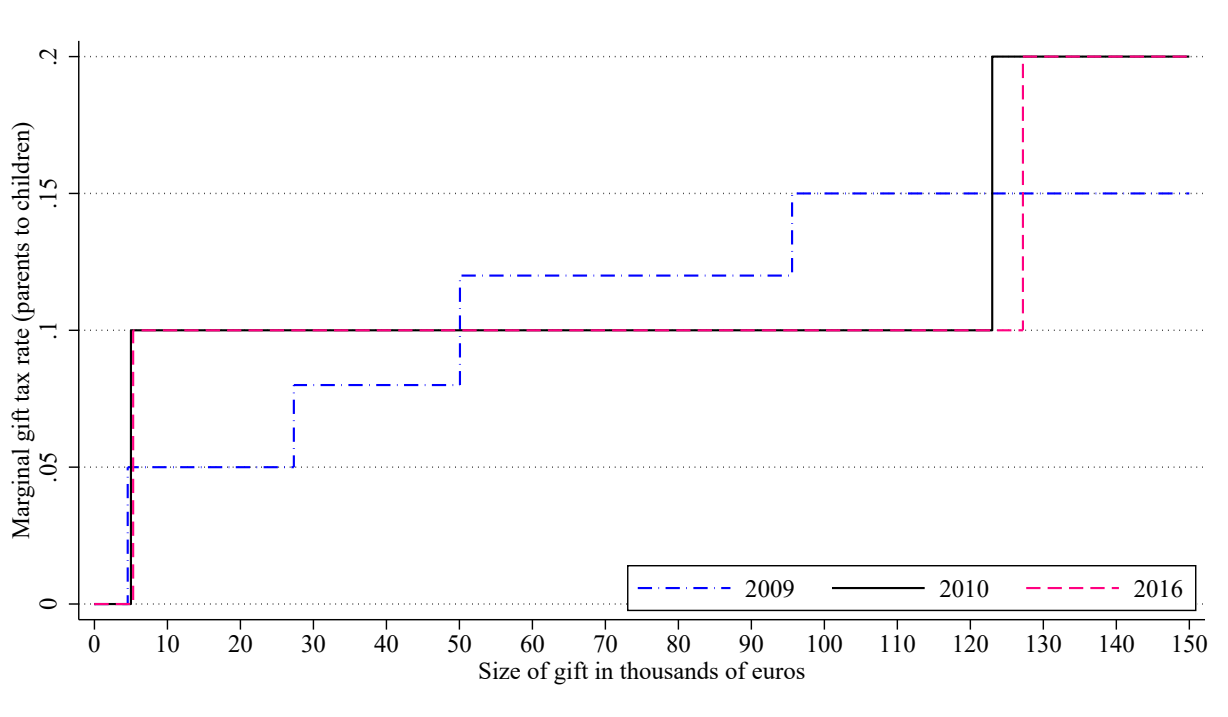
7.1 Estimating the tax elasticity of gifts

We estimate elasticities of gifts to the net-of-tax rate at four different ‘kink’ points in the gift tax schedule: the first three kink points in the pre-2010 system and the single kink point in the post-2010 system.¹¹ As set out in Section 3, before 2010, there were 7 different gift tax rates. At the first 3 kink points, the marginal tax rate changed from 5% to 8%, from 8% to 12% and from 12% to 15%, respectively. There are too few gifts given around the higher kink points to make meaningful estimates using these. Since 2010, there has been just one kink point in the gift tax system, where the rate changes from 10% to 20%, and we make estimates using this kink.

Figure 7 illustrates the tax schedule for gifts of up to €150,000 from parents to children, in 2009, 2010 and 2016. Outside of the ‘regime’ change from 2009 to 2010, nominal thresholds were updated slightly from year-to-year such and their value did not tend to vary significantly in real terms. The full set of thresholds for each year is set out in Appendix A,

¹¹There could be considered to be kinks at the tax-free threshold but we do not observe gifts below this threshold so we cannot use it to estimate an elasticity and we use ‘first kink’ to refer to the first kink involving two strictly positive tax rates.

Figure 7: Gift tax rate schedule for gifts from parents to children of up to €150,000, in 2009, 2010 and 2016



Note: Figure shows the main rate of gift tax applicable, excluding any special exemptions, in 2009, 2010 and 2016.

7.1.1 Theoretical framework

We employ the bunching estimation approach developed by Saez (2010) and Chetty et al. (2011). The fundamental idea behind this approach is that the more responsive are agents' transfers to the gift tax rate, the more that we should expect to see them choosing to give gifts at “kink points” in the tax schedule i.e. immediately below the level at which the marginal tax rate increases.

Saez (2010) showed that for a population of agents with smoothly-distributed convex preferences over two goods and facing a linear budget set with a concave “kink”, an elasticity with respect to the net-of-tax rate can be inferred from the excess mass of agents whose choices are located at the kink. The “excess mass” is defined as the additional mass of agents at the kink, compared to the counterfactual where there was no change in the marginal tax

rate at that point. If the change in tax rate at the kink is treated as small or if preferences are iso-elastic, we can infer a compensated elasticity. If the change in tax rate is not small or we are unwilling to assume iso-elastic preferences, then the elasticity recovered is a combination of the compensated and uncompensated elasticities. If preferences are heterogeneous in the population then the elasticity estimated is a local average elasticity for those who choose a level of gifts at the kink point. In a dynamic context, the elasticity estimated is a Frisch elasticity i.e. the response to a tax change when holding constant the marginal utility of wealth. For a review of bunching estimation, see Kleven (2016).

Appendix E formally sets out how the Saez (2010) framework can be applied in the case of gifts.¹² The implication of the bunching theory is that if we can recover the distribution of gifts in the no-kink counterfactual then we can infer the excess mass at the kink point, which in turn allows us to infer the compensated elasticity of gifts to the net-of-tax rate for individuals giving gifts at the kink. The elasticity that we estimate will represent the elasticity of gifts to the net-of-tax rate when holding the marginal utility of wealth constant, and so can be interpreted primarily as capturing willingness to shift gifts over time, or to shift between gifts and bequests, in response to tax incentives.

7.1.2 Sample

We perform bunching estimation using gifts given from parents to children over the final years before death. Specifically, we take those who die between 2007 and 2019 and use the observations from 2007 to 2016, the period for which the gifts data is available. The sample of gifts that we use for our bunching estimation is all gifts from these individuals to their children, made up to 6 years before the start of the year of death of the individual. The unit of observation is a parent-child-year combination, as tax is levied based on the annual flow received by one individual from another. We drop gifts that claim any special exemptions such as the once-in-life tax free gift or the temporary special exemptions for housing or

¹²This setup mirrors that in Glogowsky (2020).

education. As these gifts pay a zero rate of gift tax, their inclusion would confound estimation of the counterfactual distribution of gifts. We drop any gifts of over €200,000.

7.1.3 Estimating the counterfactual distribution of gifts

The main empirical challenge for our bunching estimation is to estimate the counterfactual distribution of gifts near the kink points, in the case where there was no kink and the lower tax rate applied throughout the region. We take two complementary approaches. The first is to estimate the counterfactual distribution by fitting a flexible polynomial to the observed density, excluding a window around the kink point, using the method of Chetty et al. (2011). The counterfactual is the fitted values from this estimation, excluding the contribution of the dummies for the excluded region around the kink, which therefore capture the bunching mass. The second approach is to use as a counterfactual the density from years in which there was not a kink at the point in the distribution being examined. We call this a “difference-in-bunching” approach.

Fitted polynomial approach: We follow the method of Chetty et al. (2011). We begin by putting the data into €1,000 ‘bins’, indexed j . For each year and bin we calculate the count of gifts. We denote the count for bin j in year t as $c_{j,t}$, the lowest value of gift in bin j as z_j , and the lowest value of gift in the same bin as the kink point as z^* . For the kink in the 2010-2016 regime, we estimate the following specification with OLS

$$c_{j,t} = \alpha_t + \sum_{p=1}^P \beta_e \cdot (z_j - z^*)^e + \sum_{\tau=2010}^{\tau=2016} \sum_{l=g_-}^{g_+} \gamma_{l,t} \cdot \mathbb{1}\{z_j - z^* = l\} \cdot \mathbb{1}\{t = \tau\} + \phi \cdot \mathbb{1}\{z_j \in \mathbb{X}\} \quad (3)$$

where \mathbb{X} is a set of round numbers.¹³ The first term and first summation term in this specification is the polynomial in distance from the kink point, with a year-specific constant,

¹³These round numbers are (in thousands of euros): 2, 4, each multiple of 5 between 5 and 45 and each multiple of 10 between 50 and 150. We also include in \mathbb{X} the bin within which the limit for the one-time-in-life exemption sits in case this is a ‘focal’ point for giving. This is between €22,048 (in 2007) and €25,526 (in 2016).

α_t . The second term is the series of dummies for the ‘excluded region’ around the kink point, which pick up the excess gifts near the kink. We control for a set of round numbers where we see ‘heaping’ of responses. The reason for conducting estimation at the level of the year-and-bin, rather just the bin is that the round numbers fall within different bins in different years due to the uprating of the kink point. As suggested by Kleven (2016), these ‘round number’ effects can therefore be separately identified from excess bunching at the kink point and also do not confound estimation of the polynomial terms.

We define the excess number gifts near the kink as $\hat{B} \equiv \sum_{\tau=2010}^{\tau=2016} \sum_{l=g_-}^{g_+} \hat{\gamma}_{l,t}$. We can then obtain an initial estimate of the excess mass at the kink point as

$$\hat{b} = \frac{\hat{B}L}{\sum_{\tau=2010}^{\tau=2016} \sum_{l=g_-}^{g_+} \hat{c}_{l,t}}, \quad (4)$$

where L is the number of bins in the excluded region.

The above estimation does not account for the fact that the actual distribution has a lower mass to the right of the kink point than does the counterfactual distribution, due to the fact that the excess mass at the kink is drawn from there. In order to account for this, we perform the adjustment as in Chetty et al. (2011). We estimate the counterfactual distribution, $\{\hat{c}_{j,t}\}$, as the fitted values, excluding the contribution of the dummies in the excluded region, from the regression

$$c_{j,t} \left(1 + \mathbb{1}\{z_{j,t} > g_+\} \frac{\hat{B}}{\sum_{\tau=2010}^{\tau=2016} \sum_{l>g_+} c_{l,t}} \right) = \alpha_t + \sum_{p=1}^P \beta_e \cdot (z_j - z^*)^e + \sum_{\tau=2007}^{\tau=2009} \sum_{l=g_-}^{g_+} \gamma_{l,t} \cdot \mathbb{1}\{z_j - z^* = l\} \cdot \mathbb{1}\{t = \tau\} + \phi \cdot \mathbb{1}\{z_j \in \mathbb{X}\}. \quad (5)$$

On the left-hand-side, a term now multiplies the counts in bins above the excluded region such that the total count of gifts in the estimated counterfactual is equal to the observed count. This is equivalent to assuming that the observed distribution above the kink point is equal to the counterfactual scaled down by some ratio. The solution to this equation is

found using an iterative procedure, starting from the estimates without the adjustment. In practice, this adjustment makes little difference to our estimates. Bootstrapped standard errors are obtained by means of random resampling of regression residuals.

There are a number of decisions to make about the precise implementation of the above estimation. The first is the interval of bins to be used in the estimation. We include in the estimation sample all bins at or above 30 bins below the kink point.¹⁴ In the case of the 2007-2009 regime, this means the region above the kink also includes other kinks and so we also include as controls the series of dummies for the excluded regions around these kinks. Second is the choice of size of excluded region around the kink (i.e. choices of g_- and g_+). Based on visual inspection, this is set to the region from 5 bins below the bin containing the kink point to 1 bin above the bin containing the kink point for each kink in the 2007-2009 regime. For the 2010-2016 regime this is set to the region from 6 bins below the bin containing the kink point to 1 bin above. Finally, following the bunching literature, we use a polynomial of degree seven, though our choice here does not appear to affect our results.

Difference-in-bunching estimation: The difference-in-bunching estimation is much simpler than the polynomial estimation, in theory and in practice. We assume that the counterfactual distribution is given by the distribution in a year in which the kink is not present. To estimate bunching at the kinks in the 2007-2009 regime, we use the observed distribution from 2009 to obtain the actual distribution and the observed distribution from 2010 to obtain the counterfactual distribution. Frequencies of gifts at any particular bin may differ across years because the tax rates applicable are different or because of other changes in economic conditions. We assume that the *relative* density across the region around the kink is the same across years. Specifically, define the counterfactual count in bin j as

$$\tilde{c}_j = c_{j,2010} \frac{\sum_{l < g_-} c_{l,2009}}{\sum_{l < g_-} c_{l,2010}} \quad (6)$$

¹⁴For the second kink, we use only the 20 bins below so as not to include the region of the first kink.

where g_- is again the lowest bin that is part of what is considered to be the bunching region. The counterfactual distribution is the distribution from 2010, scaled proportionally so that it has the same mass in the bins below the bunching region as does the 2009 distribution.¹⁵ The excess mass at the kink is then defined as

$$\hat{b}_{dib} = \frac{L \sum_{l=g_-}^{g_+} (c_{j,2009} - \tilde{c}_j)}{\sum_{l=g_-}^{g_+} \tilde{c}_j} \quad (7)$$

where L is again the number of bins in the bunching region.


7.1.4 Results of bunching estimation

Figure 8 graphically shows the result of our estimation for the kink in the 2010-2016 tax system. We see a large excess mass in the bin where the kink point falls.¹⁶ The estimated counterfactual is not smooth because it picks up the non-smoothness in the actual distribution that is accounted for by the heaping of gifts at round numbers.

¹⁵As in the polynomial estimation, we limit the region used in estimation to 30 bins below the kink point (or 20 in the case of the second kink in the 2007-09 regime).

¹⁶There is also an excess mass in the bin 5 bins below the kink point. These individuals are located close to the point equal to the kink point minus the tax-free allowance. The tax system is generally communicated by listing the amounts above the tax-free allowance that can be given before the higher tax rate is paid. It is possible that people misunderstand where the kink is located or believe that they should report only the taxable component of their gift. Either way, it seems that these individuals are attempting to locate at the kink point and should be treated as bunchers.

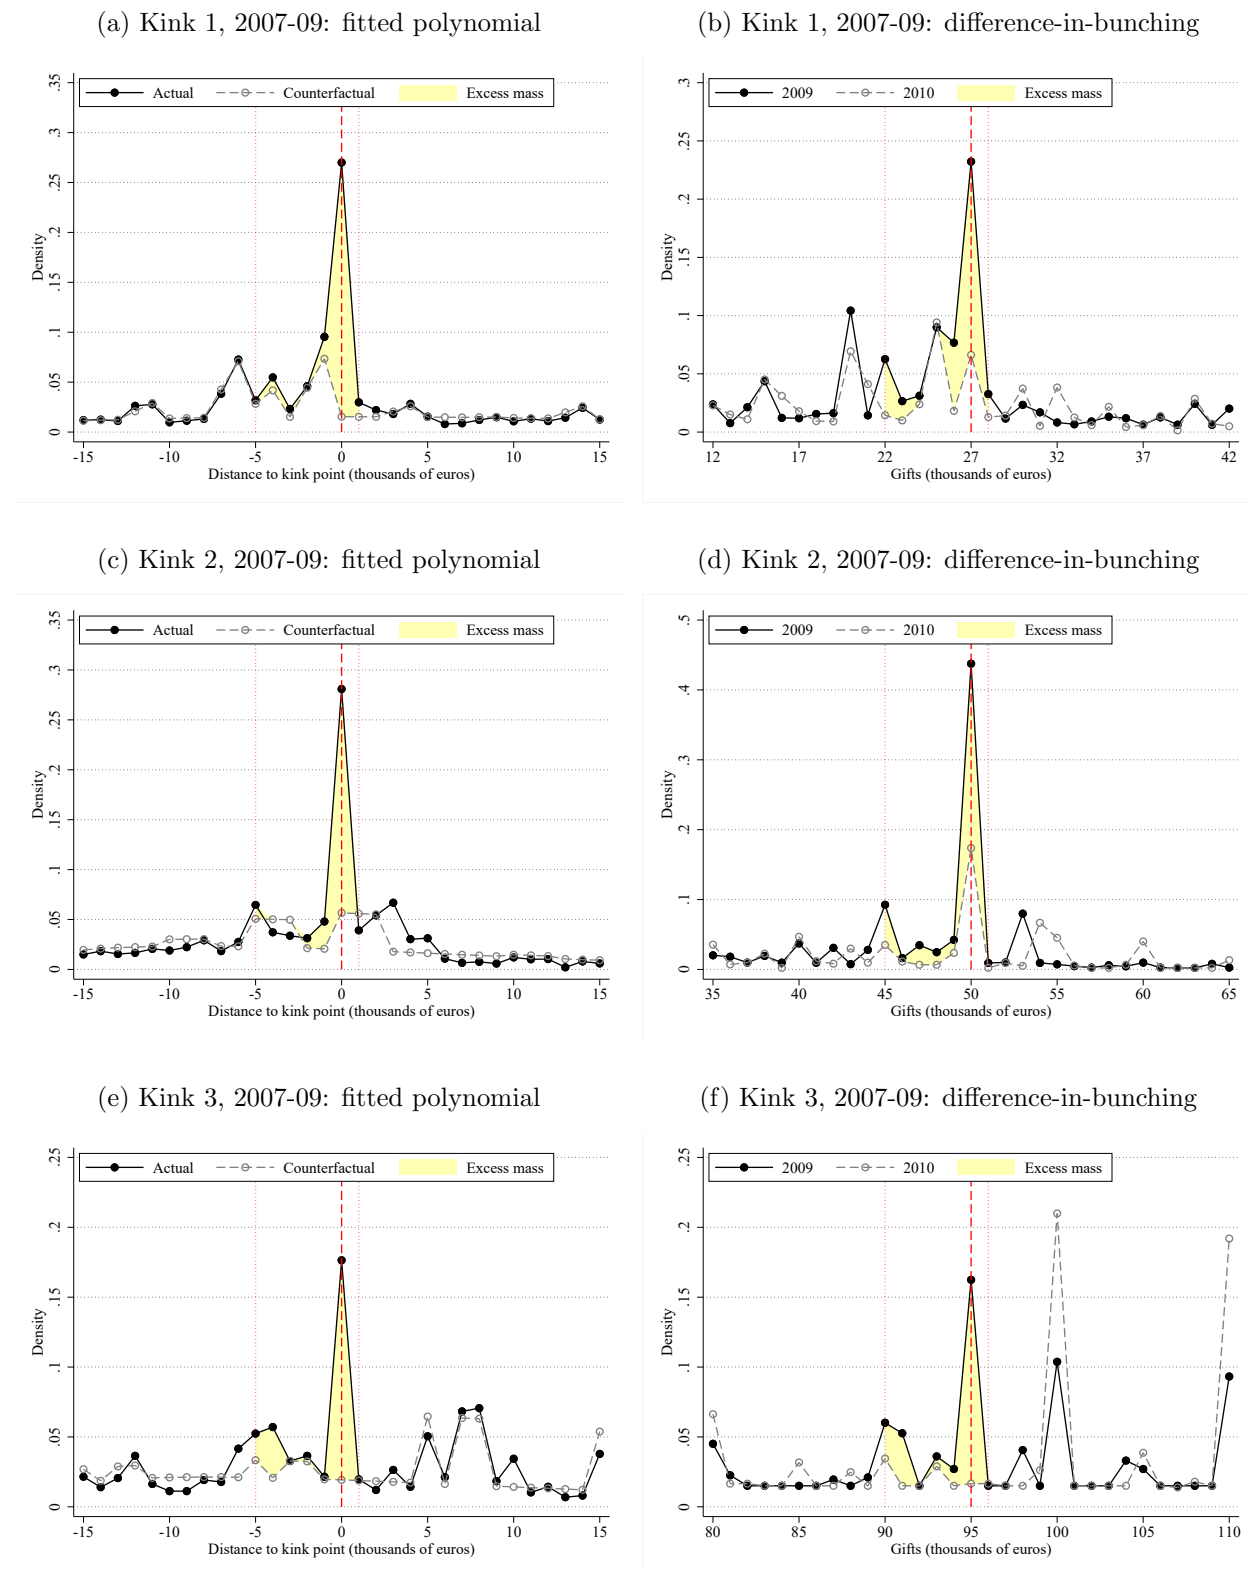
Figure 8: Actual and counterfactual densities for the kink in the 2010-16 regime



Note: The red dashed line indicates the bin containing the kink point. The dotted red lines indicate the bounds of the excluded region around the kink. The shaded region shows the excess mass at the kink.

Figure 9 shows graphically the results for the first three kinks in the 2007-09 tax system. Panels (a), (c) and (e) show the estimates using the polynomial estimation method. Panels (b), (d) and (f) show the estimates using the ‘difference-in-bunching’ method. From the second set of panels, we can directly see heaping of gifts at round numbers, even when they are not a kink point, such as €20,000 and €25,000, which are not kink points in 2009 or 2010, and at €50,000, which is not a kink point in 2010. Each of the ‘difference-in-bunching’ figures also shows how the density of gifts immediately reduces upon the elimination of the three kink points, between 2009 and 2010. Reassuringly, when we compare the estimates made by the two methods at each kink point, we see that the shape of the densities and extent of bunching is similar.

Figure 9: Actual and counterfactual densities for the first three kinks of the 2007-09 regime



Note: The red dashed line indicates the bin containing the kink point. The dotted red lines indicate the bounds of the excluded region around the kink. For disclosure reasons, data points are lower-windrosised if based on fewer than 10 gifts. However, the exact density is used to calculate the elasticities in all cases.

With our actual and counterfactual distributions and bunching density estimates, we calculate the elasticity using method m at kink point k as follows:

$$e_{k,m} = \frac{\log(1 + 1000 \times \hat{b}_{k,m}/G_k)}{1 - (\tau_2 - \tau_1)/(1 - \tau_2)} \quad (8)$$

where G_k is the mean level of the kink point and τ_1 and τ_2 are the marginal rates below and above the kink point, respectively.

The elasticity estimates are shown in Table 3. In general, we obtain large elasticities indicating a high degree of responsiveness to incentives to shift giving over time. The elasticity estimates are similar across the two methods and both show the highest estimates - 9.32 for the polynomial method and 6.90 for the difference-in-bunching method - at the first kink in the 2007-09 tax regime (located at around €27,000). The estimated elasticity declines to around 3-4 for the kink points located at around €50,000 and €95,000 and is lowest, at 0.85, for the kink in the 2010-2016 system, located at around €125,000. This implies that the highest responsiveness to gift taxation, at least in terms of the Frisch elasticities being estimated here, is found for those making gifts at lower levels. Of course, considered relative to the average gifts made across the wider sample, gifts at the lowest kink point are still very substantial sums.

Table 3: Elasticity estimates from bunching estimation

Method	2007-09: kink 1	2007-09: kink 2	2007-09: kink 3	2010-16: sole kink	Observations
Polynomial	9.32 (0.15)	2.51 (0.12)	2.52 (0.15)	0.85 (0.26)	37,161
Diff-in-Bunch	6.90	3.56	4.49		18,111

Note: The first row contains the elasticity estimates based on the counterfactual estimated using the fitted polynomial approach. Standard errors for these estimates, obtained using a bootstrap resampling procedure are reported in parentheses. The second row contains the elasticity estimates based on the difference-in-bunching method.

We note that bunching estimation has been criticised on the basis that frictions in relation

to the choice being made mean that the degree of bunching does not straightforwardly reveal an elasticity to taxation. For example, many people cannot precisely adjust their earnings in order to sit at tax kink points. This criticism does not seem very pressing in the case of gifts. When divisible financial assets are being gifted, there are very minimal frictions to adjusting the size of gifts. A second challenge to bunching estimation is made by Blomquist et al. (2021), who show that bunching estimation recovers elasticities only under (potentially implicit) assumptions about the way that preferences are distributed for agents making choices near the kink. The issue is that the distribution of choices can be made consistent with various combinations of elasticities and distributions of preference heterogeneity. Using a polynomial specification implicitly assumes that preference heterogeneity is not generating the bunching of gifts at the kink. The fact that we can use difference-in-bunching estimation, where we observe the full distribution of gifts around the kink point in a year where there is no change in tax rate at that level, is reassuring that the bunching pattern is unlikely to be driven by a form of preference heterogeneity that drives a spike in that region. For this to be the case, it would also have to be that preferences change at the same time as the tax reform in such a way as to eliminate the bunching.

7.2 Transfers as partly altruistic

A simple model of altruism whereby the utility of each of a parent's children enters into the parent's utility function (potentially with less weight than the parents' own utility) would predict that parents will make larger transfers to children with higher marginal utility of wealth.

Proxying relative marginal utility of wealth with relative wealth, we examine the relationship between measures of a child's wealth position relative to their siblings and both (a) the change in a child's wealth and (b) the size of taxable gifts given from a parent to their child. We measure children's relative wealth position by their wealth rank among their siblings and using indicator variables for whether the child is the wealthiest among their

siblings and for whether they are the least wealthy among their siblings.

We pool all observations of children of singles in our treatment and control groups. First, we estimate a triple difference version of our main event-study. The unit of observation is the child household. The explanatory variables are as per Eq. (1), with the addition of an interaction of the indicator variables for the number years until death with the measure(s) of relative child wealth. Second, we take all observations of children of singles in our treatment group and estimate an OLS regression where our outcome of interest is the amount of taxable gifts received by the child from their parent and the explanatory variables are the measure(s) of relative child wealth and parent fixed effects.

For both the triple difference estimation and the analysis of gifts, we estimate three specifications with different measures of child wealth position. Our first measure is the child's wealth rank among their siblings. In the second specification, we include an indicator variable for being the most wealthy child of the parent. This variable captures the size of the gifts or change in wealth relative to children that are not the wealthiest among their siblings and is therefore identified from variation among families with at least two children. In the third specification, we include indicator variables for being the least wealthy child of the parent and for being the most wealthy child of the parent. These variables capture the size of gifts or change in wealth relative to 'middle' children and are therefore identified from variation among families with at least three children.

Table 4 shows the results. For the triple difference specification, we report the coefficient on the interaction between the treatment group indicator, being at the start of the year of the death of the parent, and the measure of child wealth position. The first specification shows that while being one rank higher in terms of wealth in a family is associated with a 0.5 percentage point smaller increase in wealth over the final four years before a parent's death, this is not significantly different from zero. The second and third specifications show that the least and most wealthy children tend to see larger increases and smaller increases, respectively, in wealth as their parents' death approaches but again the differences are not

statistically significant. Turning to the regressions with taxable gifts as an outcome, we find that being one rank higher in terms of wealth is associated with receiving €15.30 less in gifts per year from the parent, over the years before their death. Being the wealthiest child of the parent is associated with receiving around €50 per year less in gifts. These are small differences compared to the mean taxable gift received of around €500 and it is perhaps not surprising that effects of this size could not be detected in the triple-difference specification.¹⁷

These results provide some evidence that, at least for some parents, transfers may be driven by altruism. However, transfers do not vary substantially between more or less wealthy children. This is consistent with some notions of equity between siblings or parents who value the net-of-tax transfers made (i.e. a joy-of-giving motive), irrespective of their impact on their children’s utility. This evidence is in line with other findings of equalisation of transfers across children (McGarry (2016); Wilhelm (1996)) and with the structural lifecycle literature which models a joy-of-giving bequest motive (Kopczuk and Lupton (2007); Lockwood (2018)).

8 Conclusion

In this paper, we have examined the dynamics of wealth and financial transfers to heirs in advance of death. Single individuals (including widows) who died saw an additional 6.9% decline in their wealth over the final four years of their life compared to the control group. This decline was larger for those singles who had children. By contrast, couples in which one member died saw a relative increase in their wealth of 1.5%. The decline in wealth of singles with children is predominantly explained by a rise in the wealth of their children. The residual wealth can be explained by increased copayments for long-term care.

Transfers to children are strongly motivated by tax avoidance considerations. We showed

¹⁷The two sets of regressions also test slightly different hypotheses. The triple difference estimates whether the giving near death *relative to a control group* is larger or smaller for less or more wealthy children. The analysis of gifts asks whether gifts are larger or smaller in absolute terms for more or less wealthy children. These could differ if giving in anticipation of death was distributed differently across more or less wealthy children than was giving in general.

Table 4: Relationship between change in child wealth or parents' annual taxable gifts and child's wealth position

	Child log wealth triple-diff			Taxable gifts (Euros)		
	(1)	(2)	(3)	(4)	(5)	(6)
Child wealth rank	-0.005 (0.005)			-15.30*** (4.59)		
Most wealthy child		-0.008 (0.008)	-0.003 (0.009)		-49.36*** (14.08)	-58.14*** (11.10)
Least wealthy child			0.009 (0.010)			-14.20 (12.06)
Constant	10.38*** (0.007)	11.19*** (0.003)	11.46*** (0.004)	520.86*** (9.31)	509.34*** (5.56)	518.42*** (6.09)
Parent fixed effects	No	No	No	Yes	Yes	Yes
Child fixed effects	Yes	Yes	Yes	No	No	No
Observations	325,392	325,392	325,392	103,174	103,174	103,174

Note: Statistical significance at the 1% level is denoted by ***. Standard errors are clustered at the parent level. Number of observations is the number of children in specifications (1), (2) and (3) and the number of parents in specifications (4), (5) and (6). The triple difference results report the coefficients for the interaction between the treatment group indicator, the indicator for being at the start of the year of the parent's death and, respectively, child wealth rank and indicator variables for most wealthy child and least wealthy child.

that most of the rise in children's wealth is attributable to the making of gifts below the tax-free allowance (which are very likely to reduce the tax payable compared to giving as a bequest). We estimated large Frisch elasticities of gifts to the next-of-tax rate between 9, for those giving around €27,000, and 1, for those giving around €125,000. These represent a high degree of willingness to shift giving over time in response to tax incentives.

These results are consistent with the behaviour of households who value transfers received by their children, have precautionary motives from longevity risk, and health-dependent

utility from consumption. In the face of a tax incentive to spread gifts over multiple years, a shock to longevity can lead single households to increase giving. However, when one member of a couple experiences a health shock, the precautionary motives for holding onto wealth may not be much reduced, as the partner's longevity remains uncertain, and so gifts do not increase. If the health shock decreases the marginal utility of current consumption of the household, then this could lead couples to defer consumption to after the first member of the couple has died, and so increase their wealth in advance of the first death. A simple theoretical model making these predictions is set out in Appendix F.

We have therefore seen that there is substantial estate planning behaviour, even among individuals of moderate wealth and that those close to death are very responsive to tax incentives to shift giving forward in time. Tax design should take into account the potential for substantive and widespread use of any tax avoidance opportunities through giving before death.

References

- Abadie, A. and J. Spiess (2021). Robust Post-Matching Inference. Journal of the American Statistical Association 0(0), 1–37.
- Altonji, J., F. Hayashi, and L. Kotlikoff (1992). Is the Extended Family Altruistically Linked? Direct Tests Using Micro Data. American Economic Review 82(5), 1177–1198.
- Altonji, J. G., L. J. Kotlikoff, and F. Hayashi (1997). Parental Altruism and Inter Vivos Transfers: Theory and Evidence. Journal of Political Economy 105(6), 1121–1166.
- Ameriks, J., A. Caplin, S. Laufer, and S. V. Nieuwerburgh (2011). The Joy of Giving or Assisted Living? Using Strategic Surveys To Separate Bequest and Precautionary Motives. The Journal of Finance LXVI(13105).
- Andersen, S. and K. M. Nielsen (2011). Participation constraints in the stock market: Evidence from unexpected inheritance due to sudden death. Review of Financial Studies 24(5), 1667–1697.
- Andersen, S. and K. M. Nielsen (2017). Fire Sales and House prices: Evidence from Estate Sales Due to Sudden Death. Management Science 63(1), 201–212.
- Attanasio, O., C. Meghir, and C. Mommaerts (2018). Insurance in extended family networks.
- Attanasio, O. P. and C. Emmerson (2003). Mortality, health status, and wealth. European Economic Review 1(4), 821–850.
- Banks, J., R. Blundell, P. Levell, and J. P. Smith (2019). Life-cycle consumption patterns at older ages in the United States and the United Kingdom: Can medical expenditures explain the difference? American Economic Journal: Economic Policy 11(3), 27–54.
- Bernheim, B. D., R. J. Lemke, and J. K. Scholz (2004). Do estate and gift taxes affect the timing of private transfers? Journal of Public Economics 88(12), 2617–2634.

- Bernheim, B. D., A. Shleifer, and L. H. Summers (1985). The Strategic Bequest Motive. Journal of Political Economy 93(6), 1045–1076.
- Blomquist, S., W. K. Newey, A. Kumar, and C. Y. Liang (2021). On bunching and identification of the taxable income elasticity. Journal of Political Economy 129(8), 2320–2343.
- Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records. Quarterly Journal of Economics 126(2), 749–804.
- De Nardi, M. (2004). Wealth inequality and intergenerational links. Review of Economic Studies 71(3), 743–768.
- De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. Journal of Political Economy 118(1), 39–75.
- De Nardi, M., E. French, J. B. Jones, and R. McGee (2021). Why Do Couples and Singles Save During Retirement? Federal Reserve Bank of Richmond Working Papers 21(09), 1–65.
- De Nardi, M. and F. Yang (2016). Wealth inequality, family background, and estate taxation. Journal of Monetary Economics 77, 130–145.
- Erixson, O. and S. Escobar (2020). Deathbed tax planning. Journal of Public Economics 185, 104170.
- Fadlon, I. and T. H. Nielsen (2019). Family health behaviors. American Economic Review 109(9), 3162–3191.
- García-Gómez, P., H. van Kippersluis, O. O’Donnell, and E. van Doorslaer (2013). Long-term and spillover effects of health shocks on employment and income. Journal of Human Resources 48(4), 873–909.

- Glogowsky, U. (2020). Behavioral responses to inheritance and gift taxation: Evidence from Germany. Journal of Public Economics 193(October), 104309.
- Goupille-Lebret, J. and J. Infante (2018). Behavioral Responses to Inheritance Tax: Evidence from Notches in France. Journal of Public Economics 168, 21–34.
- Groot, S., M. Lever, and J. Möhlmann (2019). Effect van erfenissen en schenkingen op vermogensongelijkheid en de rol van belastingen. Technical report, Centraal Planbureau.
- Hurd, M. D. (1987). Savings of the Elderly and Desired Bequests. American Economic Review 77(3), 298–312.
- Jones, J., M. De Nardi, E. French, R. McGee, and R. Rodgers (2020). Medical Spending, Bequests, and Asset Dynamics around the Time of Death. Economic Quarterly 106(04), 135–157.
- Joulfaian, D. (2005). Choosing between gifts and bequests: How taxes affect the timing of wealth transfers. Journal of Public Economics 89(11-12), 2069–2091.
- Kleven, H. J. (2016). Bunching. Annual Review of Economics 8, 435–464.
- Kopczuk, W. (2007). Bequest and tax planning: Evidence from estate tax returns. Quarterly Journal of Economics 122(4), 1801–1854.
- Kopczuk, W. and J. P. Lupton (2007). To leave or not to leave: The distribution of bequest motives. Review of Economic Studies 74(1), 207–235.
- Kopczuk, W. and J. Slemrod (2001). The Impact of the Estate Tax on Wealth Accumulation and Avoidance Behavior Chapter. In W. G. Gale, J. R. Hines, and J. Slemrod (Eds.), Rethinking Estate and Gift Taxation, pp. 299–349. Brookings Institution Press.
- Kvaerner, J. (2017). How Strong Are Bequest Motives? Ssrn (November).

- Lockwood, L. M. (2018). Incidental Bequests : Bequest Motives and the Choice to Self-Insure Late-Life Risks. American Economic Review 108(9), 2513–2550.
- McGarry, K. (2001). Inter Vivos Transfers or Bequests? Estate Taxes and the Timing of Parental Giving, Volume 14. MIT Press.
- McGarry, K. (2016). Dynamic aspects of family transfers. Journal of Public Economics 137, 1–13.
- Mirrlees, J., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba (2011a). Tax by Design: The Mirrlees Review. Oxford University Press.
- Mirrlees, J., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba (2011b). The Mirrlees Review: Conclusions and Recommendations for Reform. Fiscal Studies 32(3), 331–359.
- Piketty, T. and E. Saez (2013). A Theory of Optimal Inheritance Taxation. Econometrica 81(5), 1851–1886.
- Poterba, J. (2001). Estate and gift taxes and incentives for inter vivos giving in the US. Journal of Public Economics 79(1), 237–264.
- Saez, E. (2010). Do Taxpayers Bunch at Kink Points? American Economic Journal: Economic Policy 2(August), 180–212.
- Suari-Andreu, E., R. Van Ooijen, R. Alessie, and V. Angelini (2019). Giving with a warm hand: Evidence on estate planning and bequests.
- Wilhelm, M. (1996). Bequest Behavior and the Effect of Heirs ' Earnings : Testing the Altruistic Model of Bequests. American Economic Review 86(4), 874–892.

A Online Appendix: Tax rates and threshold for gifts from parents to children

Table 5: Exemptions and tax rates for gifts from parents to children (2007 to 2016)

<i>Regime 1</i>								
Tax bracket	Rate	Threshold from which applicable (Euros, by year)						
		2007	2008	2009				
1	5%	4,412	4,479	5,556				
2	8%	26,463	26,861	27,319				
3	12%	45,508	49,237	50,075				
4	15%	92,593	93,983	95,582				
5	19%	180,765	183,478	186,598				
6	23%	357,108	362,466	368,629				
7	27%	886,134	889,427	914,719				

<i>Regime 2</i>								
Tax bracket	Rate	Threshold from which applicable (Euros, by year)						
		2010	2011	2012	2013	2014	2015	2016
Lower	10%	5,000	5,030	5,030	5,141	5,229	5,277	5,304
Upper	20%	123,000	123,738	120,738	123,395	122,443	126,573	127,207

Note: Threshold amounts are in Euros.

B Online Appendix: The measurement of net wealth

Net wealth is measured at the level of the household on the first of January of each year. Our measure of net wealth includes the following components: Financial assets (savings, bonds, portfolio investments); Real estate used as the primary home; Real estate not used as the primary home; Other assets; Debt associated with the primary home; Debt other than the

mortgage on the primary home; Owned firms without legal entity; Substantial ownership of legal entity listed below. The last category includes firms without a legal entity or firms with legal entity that are substantially (at least 5%) owned by the household. Second pillar pension wealth is not included in net wealth but this is not directly accessible by households.


The value of financial assets and debt are directly reported by Dutch financial institutions to the tax authority. When these assets or debts are not held at Dutch financial institutions, households have to self-report them. For real estate, the value is determined annually by municipalities and reported to the tax authority. The valuation is based on recent sales date of comparable real estate. Other assets are always self-reported and include for example loans to other households. Usable objects like cars or boats and art that is not used for investment purposes, are exempt from tax and not included in the definition of net wealth. A personal income tax is levied on an imputed return on net wealth, with the imputed return varying by asset type (i.e. there is an effective wealth tax with a level varying by wealth type). The actual return is not fiscally relevant. This means that the tax authority has an interest in making sure the quality of the wealth data is high. For ownership of firms without legal entity or where the household owns at least 5% of the firm, the information on the wealth is based on the equity as reported in the balance sheet of the underlying firm. The firm equity is distributed to households based on their share of ownership in the firm. Note that the valuations of assets in the balance sheet are subject to accounting standards and may not always reflect the actual economic value of the assets of the firm.

C Online Appendix: Supplementary results

C.1 Results for the level of wealth

Figure 10 shows the main results for the level of wealth.

Figure 10: Levels of wealth and event-study estimates for singles and couples



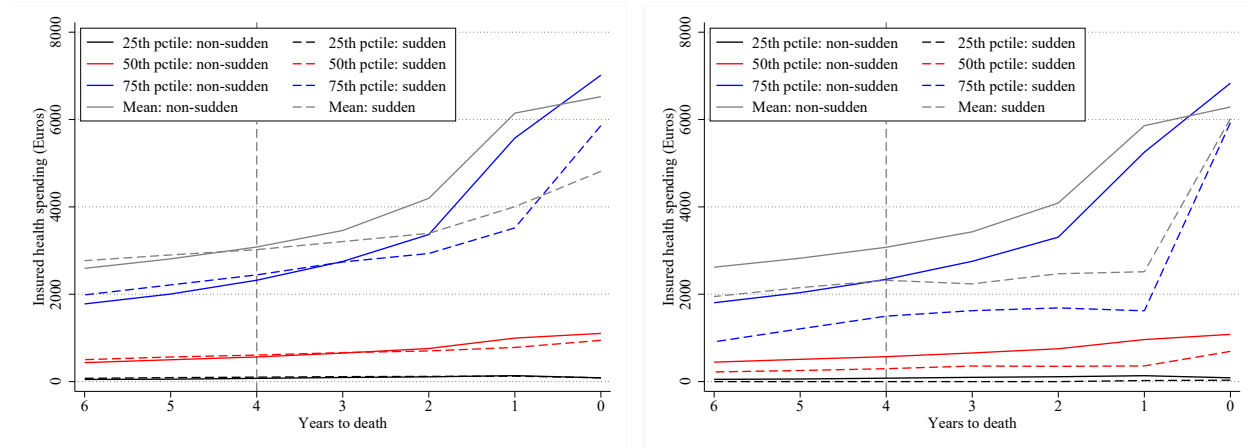
Note: Panel (a) shows the mean values of wealth of singles, by years until death, for both the treatment and matched control groups. Panel (b) shows the event study estimate for the difference between the treatment and control group in the change in their wealth relative to the base period 4 years before death. Panels (c) and (d) show the equivalent figures for couples. All measurements are in Euros. We show 95% confidence intervals, with standard errors clustered at the match level. Sample size is 213,132 singles and 311,174 couples.

C.2 The evolution of health and wealth in advance of deaths defined as ‘sudden’ and ‘non-sudden’, based on ICD-10 cause-of-death classification

Figure 11 shows selected percentiles of health spending by years until death, for ‘sudden’ and ‘non-sudden’ deaths, based on two cause-of-death definitions of ‘sudden death’ that use the ICD-10 classification of causes of death. One definition follows Andersen and Nielsen (2011). The second definition categorises transport accident deaths (ICD-10 codes V01-V99) as sudden and all others as non-sudden. Health spending increases over the years before death at the mean, median and 75th percentile for those who die of non-sudden causes, using either definition. Using the Andersen and Nielsen (2011) definition, there is a rise in health spending also in the sudden death group. For those who die of transport accidents, we see no rise in health spending until the year of their death (of course significant health spending may be incurred at the point of a sudden death).

Figure 11: Selected percentiles of health spending by years until death, for ‘sudden’ and ‘non-sudden’ deaths, based on two cause-of-death definitions of ‘sudden death’

(a) Andersen and Nielsen (2011) definition of ‘sudden’ (b) Transport accidents definition of ‘sudden’




Note: Both panels show selected percentiles of the distribution of insured health spending by years until deaths, splitting the sample into those who died of ‘sudden’ and ‘non-sudden’ causes of death. Panel (a) uses the definition of ‘sudden’ from Andersen and Nielsen (2011). Panel (b) defines a ‘sudden’ death as one caused by a transport accident (ICD codes V01-V99). All measurements are in Euros. 325,449 individuals, including 46,720 sudden deaths in panel (a) and 1,036 sudden deaths in panel (b)

Figure 12 shows the difference in the evolution of wealth between those who die of ‘sudden’ causes and those who died later, or who died of a ‘non-sudden’ cause. In each case, we take the group who died of a sudden cause and match them to control observation using the method described in the main text. The first comparison (black markers) is between those who died of sudden causes as defined by Andersen and Nielsen (2011) and those who died four years later. Wealth declines for those who died of sudden causes, compared to those who died later indicating some anticipatory behaviour for in the ‘sudden’ group. The second comparison (blue markers) is between those who died of sudden and non-sudden causes (as defined by Andersen and Nielsen (2011)). The sudden group has higher wealth than the ‘non-sudden’ group, 2 and 1 years before death. However, the estimated difference is at most 2% of wealth and is insignificant by the year of death, suggesting that the anticipatory behaviour captured in (a) leads this method to significantly understate the effects of estate

planning. The third comparison (red markers) is between those who died of sudden and non-sudden causes (defining sudden death as caused by transport accidents). In this case small sample of those dying in transport accidents leads to very imprecise estimates and the difference in wealth between the two groups is not statistically significant from zero at the 5% level.

Figure 12: Event-study estimates for changes in singles' log wealth, comparing 'sudden' deaths to a control group composed of those who died 4-years later or who died of 'non-sudden' causes




Note: Figure shows, for single individuals, 3 sets of event study estimates for log wealth. We compare (a) those who died of 'sudden' causes (Andersen and Nielsen, 2011 definition) to a control group who died four years later (48,698 individuals); (b) Those who died of 'sudden' causes to a group who died of 'non-sudden' causes (Andersen and Nielsen, 2011 definition, 50,128 individuals). (c) Those who died of 'sudden' causes to a group who died of 'non-sudden' causes (Transport accidents definition, 580 individuals). We show 95% confidence intervals, with standard errors clustered at the match level.

C.3 Results for single individuals with children’s log wealth, split by quartile of wealth four years before death

Figure 13 shows the results for log wealth for singles with children, splitting by quartile of wealth four years before death. The decline in wealth is larger in percentage terms for those with higher levels of initial wealth.

Figure 13: Event-study estimates of the effect on singles with children’s log wealth by initial quartile of wealth




Note: Figure shows, for single individuals with children, the event study estimates for log wealth, split by wealth quartile 4 years before death. We show 95% confidence intervals, with standard errors clustered at the match level. 152,756 individuals.

C.4 Effect of proximity to death on income

Figure 14 shows the estimated effect of proximity to death on income, split by singles and couples, using the matched event-study method of the main text. We do not show results for the year of death because those who die do not experience a full year and so mechanically see a large decline in their income. We can rule out any substantive changes in income for

both singles and couples as death approaches (although the change for couples is statistically different from zero, it is very small in magnitude).

Figure 14: Event-study estimates for the effect of proximity to death on log income for singles and couples

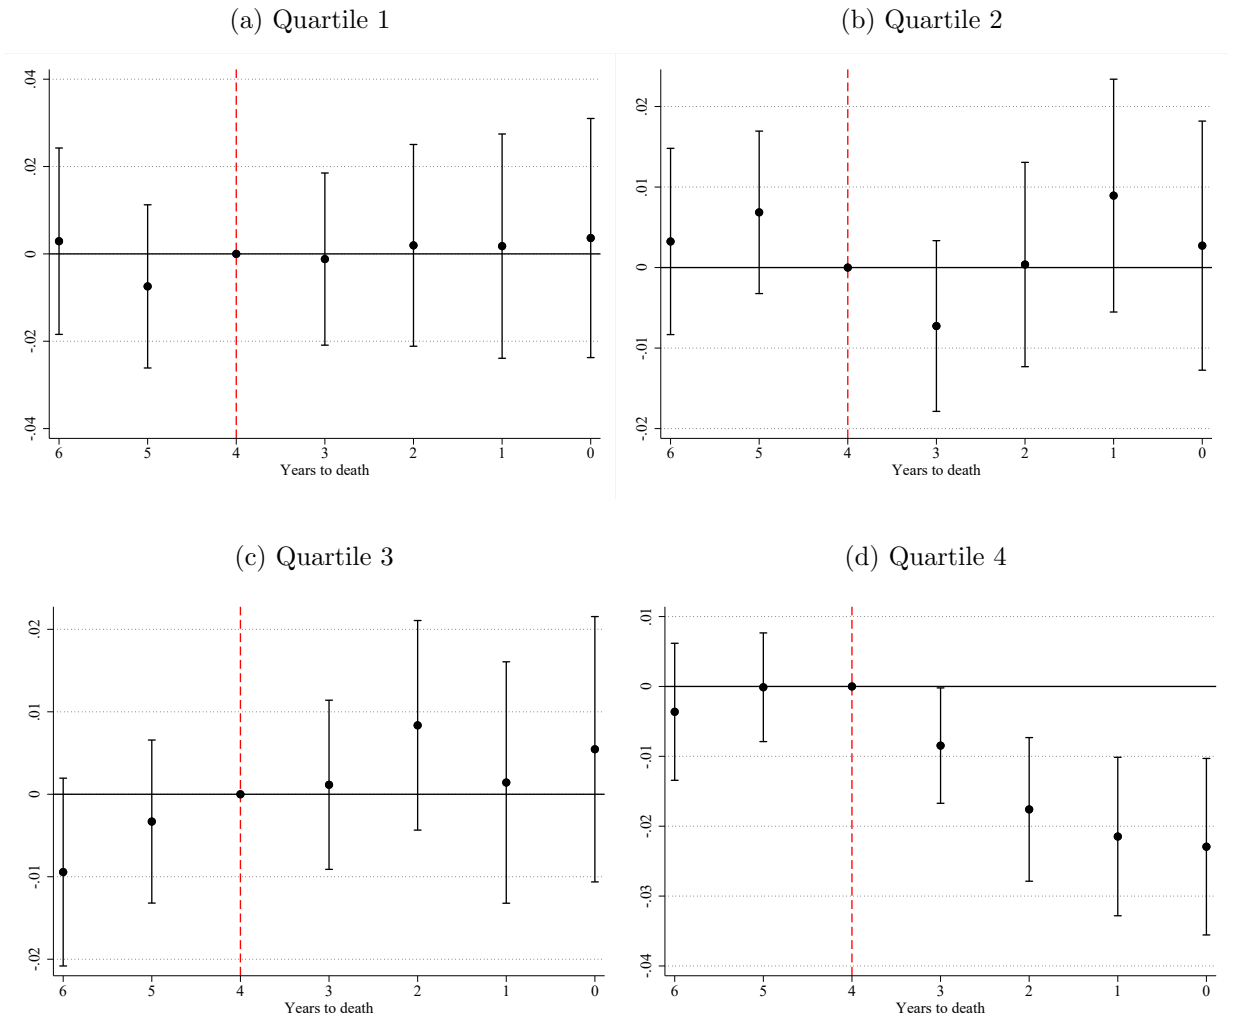


Note: Figure shows the event study estimates for the effect on log income for singles and couples. We show 95% confidence intervals, with standard errors clustered at the match level. 213,132 singles and 311,174 couples. Year of death is omitted as income in that year is affected by death but does not affect wealth measured at the start of that year.

C.5 Results including grandchildren’s wealth within the definition of dynasty wealth

Figure 15 shows the estimated effects of proximity to death on dynasty wealth, split by initial quartile of parents’ wealth, when grandchildren’s wealth is included in the dynasty wealth measure. We find results consistent with those in section 6: there is no significant change in dynasty wealth across the lower three quartiles of the wealth distribution but dynasty wealth for the riches quartile of singles with children declines by 2% by the year of death.

Figure 15: Event-study estimates of the effect on singles with children’s log dynasty (including grandchildren) wealth, by quartile of wealth




Note: The four panels show estimates split by the quartile of wealth of the parents, four years before death. Each panel shows, for single individuals with children, the event study estimates for the effects on log dynasty wealth where dynasty wealth is the sum of parents’ wealth, and the wealth of the households of all of their children and grandchildren. We show 95% confidence intervals, with standard errors clustered at the match level. 152,756 individuals.

C.6 Results for single individuals with children’s log wealth, log children’s wealth and log dynasty wealth, for selected quantiles, at the start of the year of death

Figure 16 shows the event-study estimates for log wealth, log children’s wealth and log dynasty wealth for singles with children, for selected quantiles of wealth. All effects are expressed as a proportion of dynasty wealth.

Figure 16: Event-study estimates of the effect on singles with children’s log wealth, log children’s wealth and log dynasty wealth, by initial quantile of wealth



Note: Figure shows, for single individuals with children, the event study estimates for log wealth, log children’s wealth and log dynasty wealth. We show estimates for the whole sample, and those who were in the wealthiest 25%, 10% and 1%, four years before death. We show 95% confidence intervals, with standard errors clustered at the match level. 152,756 individuals.

D Online Appendix: Estimation of an upper bound on the effect of long-term care copayments on wealth

We calculate the upper bound on the effect of copayments the decline in wealth before death as follows. As we only have care copayment data from 2015 onwards, we cannot pursue our difference-in-difference strategy to estimate the increase in copayments due to proximity to death. Instead, we construct an upper bound using cumulative copayment spending over the final four years before the year of death of the individuals in the matched control group who die in 2019 (i.e. the individuals who are matched to those dying in 2015). For these individuals, we observe copayment spending from four years before the start of the year of their death onwards. For each individual and each year from 2016 to 2019, we create an outcome equal to the sum of copayments from the beginning of 2015 to the beginning of that year, divided by the sum of the cumulative copayments and the wealth at the start of that year. That is, we define the ‘copayment share’ as:

$$CopaymentShare_{i,t} = \frac{\sum_{\tau=2015}^{\tau=t-1} Copayment_{i,t}}{W_{i,t} + \sum_{\tau=2015}^{\tau=t-1} Copayment_{i,t}} \quad (9)$$

where $Copayment_{i,t}$ are the copayments made by individual i in year t and $W_{i,t}$ is individual i 's wealth at the start of year t . We estimate the mean level of this outcome for 3, 2, 1 and 0 years before the year of death.

In essence, this estimates the simple difference in cumulative copayments over the four years before death, expressed as a percentage of wealth plus cumulative copayments. This estimate is an upper bound on the percentage reduction in wealth due to copayment spending in advance of death for two reasons. First, not all of the use of care is attributable to proximity to death. There are some individuals who would have care needs even in the situation where they were not approaching death. Second, some of the copayment spending is a replacement for other spending rather than an addition to it. For example, some of

those going in to residential care will reduce their spending on rent, bills and food.

E Online Appendix: Bunching theory

We can formally apply the Saez (2010) framework as follows.¹⁸ Suppose that there is a continuum of agents, each endowed with wealth y and a child to whom they can give gifts. Denote the pre-tax level of gifts given as g and the after-tax gifts received by their child as c . They face a piecewise linear tax schedule $T(g)$ such that $c = g - T(g)$. Suppose that each agent has a utility function $u(c, g)$ that is concave and increasing in the after-tax gifts their child receives (because of altruism, for example) and concave and decreasing in the level of pre-tax gifts made (because of a convex opportunity cost of giving gifts, as these reduce future consumption or bequests, for example).

For ease of exposition, consider the case with iso-elastic preferences. Suppose that utility is of the form

$$u_i = c - \frac{\phi_i}{1 + \frac{1}{e}} \left(\frac{g}{\phi_i} \right)^{1 + \frac{1}{e}} \quad (10)$$

where the first term is the benefit from post-tax gifts and the second term is a convex cost of increasing pre-tax gifts. With a constant marginal tax rate (i.e. $T(g) = \tau \cdot g$), if the optimal choice of gifts is interior then it is $g^* = \phi(1 - \tau)^e$. The parameter e is therefore equal to the net-of-tax elasticity of gifts.¹⁹ Assume that the parameter e is homogeneous in the population. The parameter ϕ can capture a taste for gifts. If we assume that there is a smooth distribution of ϕ_i in the population and that y is large enough that choices are interior, then there will be a smooth distribution of choices of gift, which we denote $h_0(g)$. Now consider a two-bracket linear gift tax of the form

$$T(g) = \tau g + \Delta\tau(g - k)\mathbb{1}(g > k). \quad (11)$$

¹⁸This setup mirrors that in Glogowsky (2020).

¹⁹To see this, consider a small change in the linear tax rate from τ to $\tau + \Delta\tau$ that leads to a change of choice of gifts from g^* to $g^* + \Delta g$. From the optimal gift choice, we have that $\frac{g^* + \Delta g}{g^*} = \left(\frac{1 - \tau}{1 - \tau - \Delta\tau} \right)^e$ and therefore $e = -\frac{\ln(1 + \Delta g/g^*)}{\ln(1 - \Delta\tau/(1 - \tau))}$. When $\Delta\tau$ is small and so Δg is also small, we have $e \approx \frac{\Delta g/g^*}{\Delta\tau/(1 - \tau)}$.

This is a schedule with a marginal rate of τ on gifts up to the level k and a marginal rate of $\tau + \Delta\tau$ on gifts above k . The point k is a kink point in the tax schedule. For individuals with taste for gifts such that $\frac{k}{(1-\tau)^\epsilon} \leq \phi_i \leq \frac{k}{(1-\tau+\Delta\tau)^\epsilon}$, their optimal choice in response to this schedule will be to locate at the kink. Within this group, those individuals for whom it is the case that $\phi_i = \frac{k}{(1-\tau+\Delta\tau)^\epsilon}$ make the same choice of gift under the two-bracket system as they would if faced with the linear tax schedule with constant marginal rate $\tau + \Delta\tau$ and wealth equal to $\tilde{y} = y + \Delta\tau \cdot k$. The difference in the choice of gift of this ‘marginal buncher’ under the two-bracket schedule compared to the schedule with single marginal rate τ , which we will denote Δg , is therefore equal to their compensated response to a tax rate change of $\Delta\tau$.

The key insight of the bunching approach is that the observed excess mass of individuals located at the kink point can be used to recover the compensated response of the marginal buncher, Δg , and therefore infer the compensated elasticity. To see this, note first that those individuals for whom $\phi_i \leq \frac{k}{(1-\tau)^\epsilon}$ would make the same choice, locating at or below the kink, under the single and two-bracket systems. Second, note that the marginal buncher is the individual for whom, of those located at the kink point under the two-bracket system, the difference between the optimal choices under the single and two-bracket systems is greatest. Consequently, the additional mass of gifts that are made at the kink point under the two-bracket system compared to the single-bracket system is equal to the mass of gifts made in the interval $(k, k + \Delta z)$ under the single bracket system. This yields the relationship

$$B \equiv h_1(k) - h_0(k) = \int_{g=k}^{k+\Delta g} h_0(g) dg \approx h_0(k) \Delta g \quad (12)$$

where $h_1(g)$ is the density function for gifts under the two-bracket system, and B is defined as the ‘excess mass’ of gifts at the kink point, relative to the single-bracket counterfactual. The final approximation is sometimes employed on the basis of assuming that the counterfactual density is constant over the interval $(g, g + \Delta g)$. If we can estimate B and $h_0(g)$ we can

therefore infer Δg and the compensated elasticity. When elasticities are heterogeneous across individuals, we obtain the average response of gifts to the difference in tax rates, integrating over the distribution of elasticities for the marginal bunchers. Provided the change in the tax rate is small, we then infer the local average elasticity for the marginal bunchers (Kleven, 2016).²⁰ When preferences are not assumed to be iso-elastic, income effects can be assumed to be small, and the elasticity to approximate the compensated elasticity, if the difference between the lower and higher tax rate, and therefore difference in gifts for the marginal buncher, is small.

F Online Appendix: Simple model of wealth and gifts

In this section we set out a simple theoretical model to ground our interpretation of our empirical analysis. The idea is to set out some predictions about the effect of shocks to survival expectations and the marginal utility of consumption on consumption, gifts and wealth of singles and couples. We present a simple two-period model in which the agent is either a single or a couple household.

F.1 Single households

Single households are endowed with a level of wealth, W_0 . They face a probability of survival to the second period, denoted s . In the first period, singles can use their wealth to fund consumption, denoted c_1 , gifts, denoted g_1 , or future wealth, denoted W_1 . Wealth saved is left as a bequest, denoted b_2 , in the case where they die before period two. In the case where they survive to period two, wealth can be used for consumption, denoted c_2 , or gifts, denoted g_2 . Any wealth not used for consumption or gifts in period two is left as a bequest, b_2 , at the end of period two. Singles receive utility from consumption and gifts in the first period, and from bequests if they die before the second period, and from consumption, gifts

²⁰As discussed by Kleven (2016), if the tax rate change is large, then aggregation bias is possible as the elasticity at the marginal response may diverge from the average elasticity for the marginal bunchers.

and bequests in the second period if they survive to that period. For simplicity, we assume no discounting and no return to wealth. Utility is assumed to be linear and additive in the log of each period's consumption, gifts and bequests. Expected utility for a single household is given by,

$$U^s(c_1, c_2, g_1, g_2, b_1, b_2; \alpha, \phi, s) = \alpha \cdot \log(c_1) + \phi \cdot \log(g_1) + (1 - s) \cdot \phi \cdot \log(b_1) \\ + s \cdot \alpha \cdot \log(c_2) + s \cdot \phi \cdot \log(g_2) + s \cdot \phi \cdot \log(b_2) \quad (13)$$

where α is a marginal utility of consumption shifter and ϕ denotes the taste for gifts and bequests. We assume both of these parameters are strictly positive. The budget constraints for a single household are:

$$c_1 + g_1 + W_1 \leq W_0 \quad (14)$$

$$c_2 + g_2 + b_2 \leq W_1 \quad (15)$$

$$b_1 = W_1 \quad (16)$$

A key feature of this very simple model is that there are diminishing returns to making intergenerational transfers at any one point in time. This is modelled through the use of the log function, which should be seen as representing the progressive taxation of transfers. This gives an incentive to spread giving across gifts and bequests in multiple periods. The second key feature is that there is uncertainty over longevity. If the agent knew that they would die with certainty after the first period, they would divide their desired level of intergenerational transfers equally between gifts and bequests. The possibility that they could live on, and hence have further consumption needs, introduces a trade-off between precautionary savings in the face on longevity and a desire to make transfers in an efficient way.

The solution for the period one variables is given by

$$c_1 = \frac{\alpha \cdot W_0}{\alpha + 2\phi + s \cdot \alpha + s \cdot \phi} \quad (17)$$

$$g_1 = \frac{\phi \cdot W_0}{\alpha + 2\phi + s \cdot \alpha + s \cdot \phi} \quad (18)$$

$$W_1 = \frac{[s + s \cdot \alpha + s \cdot \phi]W_0}{\alpha + 2\phi + s \cdot \alpha + s \cdot \phi}. \quad (19)$$

We take the first derivative with respect to s to examine the impact of a change in survival expectations:

$$\frac{dc_1}{ds} = -\frac{\alpha(\alpha + \phi)}{[\alpha + 2\phi + s \cdot \alpha + s \cdot \phi]^2}W_0 < 0 \quad (20)$$

$$\frac{dg_1}{ds} = -\frac{-\phi(\alpha + \phi)}{[\alpha + 2\phi + s \cdot \alpha + s \cdot \phi]^2}W_0 < 0 \quad (21)$$

$$\frac{dW_1}{ds} = \frac{(\alpha + \phi)^2}{[\alpha + 2\phi + s \cdot \alpha + s \cdot \phi]^2}W_0 > 0 \quad (22)$$

Eq. (20) tells us that period one consumption is decreasing in survival probability, i.e. a negative shock to survival chances increases present consumption. Eq. (21) tells us that gifts are decreasing in survival probability, i.e. a negative shock to survival chances increases gift giving when alive. Eq. (22) tells us that wealth held into the second period is increasing in survival probability, i.e. a negative shock to survival decreases wealth held into the future.

We next allow the marginal utility of consumption to be a function of the survival probability in order to consider a health shock that leads to a change in marginal utility of consumption and survival chances. To model this, we let the marginal utility of consumption shifter be a function of survival probability, denoted $\alpha(s)$. Taking the derivative of the

functions for period one choice variables with respect to s , we then obtain:

$$\frac{dc_1}{ds} = \frac{\phi \cdot \alpha'(s)(2 + s) - \alpha(s)^2 - \alpha(s)\phi}{[\alpha + 2\phi + s \cdot \alpha + s \cdot \phi]^2} W_0 \quad (23)$$

$$\frac{dg_1}{ds} = -\frac{\phi[\alpha'(s)(1 + s + \phi) + \alpha(s)]}{[\alpha + 2\phi + s \cdot \alpha + s \cdot \phi]^2} W_0 \quad (24)$$

$$\frac{dW_1}{ds} = \frac{(\alpha + \phi)^2 - \alpha'(s)\phi}{[\alpha + 2\phi + s \cdot \alpha + s \cdot \phi]^2} W_0 \quad (25)$$

If we assume that the marginal utility of consumption is increasing in the agent's survival probability, i.e. $\alpha'(s) > 0$, then Eq. (24) tells us that gifts are decreasing in the survival probability. Eq. (23) and Eq. (25) tell us that provided $\alpha'(s)$ is sufficiently small, first-period consumption is decreasing in the survival probability and wealth is increasing in the survival probability. If we alternatively assume that $\alpha'(s) < 0$, then consumption and wealth are decreasing and increasing, respectively, in the survival probability. In this case, gifts will be decreasing in the survival probability so long as $\alpha'(s)$ is sufficiently small.

Table 6 summarises the predictions of our model for the response of current consumption, current gifts and next-period wealth to a negative shock to the probability of survival. Pertinently, given our empirical results, our model yields the prediction that if the marginal utility of consumption and survival chances are associated with each other, but the rate of change of marginal utility with survival chances is sufficiently small, a shock that leads to a reduction in survival probability leads the agent to increase their current consumption and current gifts and decrease the level of wealth that they hold into the future.

F.2 Couple households

For couples, we consider the implications of uncertainty over the survival of the first member of the couple. We assume that the couple is endowed with a level of wealth, W_0 . They face a probability that both members of the couple survive to the second period, denoted s . In the first period, couples can use their wealth to fund consumption, denoted c_1 , gifts, denoted g_1 ,

Table 6: Summary of model predictions of response of first-period variables to negative shock to survival expectations

$\alpha'(s)$	Singles			Couples		
	c_1	g_1	W_1	c_1	g_1	W_1
+, small	↑	↑	↓	↑	↑	↓
+, large	↓	↑	↑	↓	↑	↑
-, small	↑	↑	↓	↑	↑	↓
-, large	↑	↓	↓	↑	↓	↓

or future wealth, denoted W_1 . In period two, wealth can be used for consumption or gifts. We assume that the marginal utility of consumption depends on the number of members of the couple alive. Second-period choices may therefore depend on how many members of the couple household are alive. We denote consumption and gifts as c_2^A and g_2^A , respectively, in the case where both survive, and as c_2^B and g_2^B if a member of the couple dies. Utility is received from consumption and gifts and of a form analogous to that for singles. Expected utility for a couples household is given by,

$$\begin{aligned}
 U^s(c_1, c_2, g_1, g_2; \alpha, \phi, \beta) = & \alpha \cdot \log(c_1) + \phi \cdot \log(g_1) + s \cdot \alpha \cdot \log(c_2^A) \\
 & + (1 - s) \cdot \beta \cdot \log(c_2^B) + s \cdot \phi \cdot \log(g_2^A) + (1 - s) \cdot \phi \cdot \log(g_2^B) \quad (26)
 \end{aligned}$$

where α is a marginal utility of consumption when two members of the couple are alive and β is the marginal utility of consumption when one member of the couple is alive and ϕ again denotes the taste for gifts. We assume that these parameters are strictly positive.

The budget constraints for a couple household are:

$$c_1 + g_1 + W_1 \leq W_0 \quad (27)$$

$$c_2^A + g_2^A \leq W_1 \quad (28)$$

$$c_2^B + g_2^B \leq W_1 \quad (29)$$

The solution for the period one variables is given by

$$c_1 = \frac{\alpha \cdot W_0}{\alpha + \phi + \beta + s(\alpha - \beta)} \quad (30)$$

$$g_1 = \frac{\phi \cdot W_0}{\alpha + \phi + \beta + s(\alpha - \beta)} \quad (31)$$

$$W_1 = \frac{[\beta(1 - s) + s \cdot \alpha]W_0}{\alpha + \phi + \beta + s(\alpha - \beta)}. \quad (32)$$

Allowing the first-period marginal utility to depend on the survival probability, we take the first derivative with respect to s :

$$\frac{dc_1}{ds} = \frac{\alpha'(s)(\phi + (1 - s)\beta) - \alpha(s)^2}{[\alpha + \phi + \beta + s(\alpha - \beta)]^2} W_0 \quad (33)$$

$$\frac{dg_1}{ds} = -\frac{\phi[\alpha'(s)(1 + s) + \alpha(s)]}{[\alpha + \phi + \beta + s(\alpha - \beta)]^2} W_0 \quad (34)$$

$$\frac{dW_1}{ds} = \frac{\alpha(s)^2 + \phi \cdot \alpha(s) + s \cdot \alpha'(s) \cdot \phi - \alpha'(s)\beta(1 + s)}{[\alpha + \phi + \beta + s(\alpha - \beta)]^2} W_0 \quad (35)$$

From these equations, we note the following implications. If $\alpha'(s)$ is positive and large, then a decrease in survival probability can lead to a decline in consumption. Here, the household defers consumption given the possibility not only of leaving gifts instead of consuming but also because the future marginal utility of consumption may rise when the first member of the couple dies and the marginal utility of consumption rises. We can see that if $\alpha'(s) > 0$, a sufficiently large β will guarantee that consumption reduces in response to a negative shock to the survival probability. Gifts will increase in response to a negative shock to the

survival probability (unless $\alpha'(s)$ is large and negative). Finally, wealth will increase if $\alpha'(s)$ is positive and large or if β is sufficiently large compared to α . This means that our model can rationalise a rise in couples' wealth in response to a negative shock to the survival of the first member of the couple. We can see this as representing a case where the couple's marginal utility of consumption falls substantially when the first member of the couple experiences a health shock and their survival probability decreases. In response, the couple holds back on consumption in order to spend more after the first member of the couple has died.