

IFN Working Paper No. 1570, 2026

Does Family Caregiving Complement Formal Care? Daughters and Parental Survival After a Severe Health Shock

Elin Colmsjö, Francesco Ruggieri och Matteo Saccarola

Does Family Caregiving Complement Formal Care?

Daughters and Parental Survival After a Severe Health Shock

Elin Colmsjö*

Francesco Ruggieri[†]

Matteo Saccarola[‡]

September 21, 2026

Abstract

We study how informal family caregiving and formal care interact after severe parental health shocks, and whether they operate as substitutes or complements. Using Danish administrative data, we compare parents of only children who experience a stroke at the same age and differ only in the gender of their child. Before the stroke, parents of daughters and parents of sons follow similar health trajectories, with comparable levels of pre-existing conditions and healthcare utilization. After the stroke, a survival gap opens and grows, and parents of daughters are 1.5 percentage points more likely to be alive by year five. Over the same period, daughters adjust more than sons along several margins: they become more likely to live near their parents, less likely to be married, and have fewer children, consistent with greater involvement in informal caregiving. Crucially, parents of daughters use *more* physical rehabilitation, an elective type of formal care that requires sustained engagement. Rather than crowding out formal provision, daughters appear more likely to help parents navigate and use professional services. Informal and formal care can therefore operate, at least in part, as complements in the response to severe health shocks.

*Research Institute of Industrial Economics (IFN), Stockholm, Sweden. E-mail: elin.colmsjo@ifn.se.

[†]University of Virginia. E-mail: ruggieri@virginia.edu.

[‡]University of Chicago. E-mail: msaccarola@berkeley.edu.

We are grateful to Abigail Adams, Luigi Butera, Jeanne Commault, Manasi Deshpande, Stefano DellaVigna, Michael Greenstone, Juanna Joensen, Amir Kermani, Magne Mogstad, Enrico Moretti, and Dmitry Taubinsky for helpful comments. We also thank audiences at the 2024 Public-Labor Lunch in the Department of Economics at the University of Chicago, the 2024 Haas Seminar in Real Estate, the 2025 UC Berkeley Psychology and Economics Seminar, the 2025 UC Berkeley Finance Student Workshop, the 2026 Northwestern University Micro Gender Student Seminar, and the 2026 Bergen Macro Workshop for their feedback. Saccarola acknowledges financial support from the NBER Fellowship in Behavioral Macroeconomics and the Fisher Center for Real Estate and Urban Economics. Colmsjö acknowledges financial support from Handelsbankens forskningsstiftelser (P23-0186).

1 Introduction

Life expectancy has risen steadily across OECD countries, increasing the demand for long-term care for older adults (Colombo et al., 2011). Family members remain central to the provision and coordination of care. Across OECD countries, 60 percent of older adults requiring care report receiving only informal care, most often provided by close relatives (Rocard and Llana-Nozal, 2022). The demands on families are particularly acute when elderly parents experience nonfatal health shocks that generate persistent needs for assistance over extended periods (Ciccarelli and Van Soest, 2018).¹ Governments have sought to reduce this burden by expanding public long-term care systems and subsidizing private long-term care insurance, often with the explicit objective of relieving the caregiving responsibilities of working-age relatives. These policies generally presume that formal and informal care are substitutes.²

Yet, from a theoretical perspective, formal and informal long-term care need not be substitutes. First, preferences over who provides care may limit substitutability. If care recipients or their relatives place a sufficiently high value on receiving non-specialized care from relatives, even fully subsidized access to formal services may not induce families to replace informal with formal care. Second, formal and informal care may exhibit a degree of complementarity. Access to high-skill formal care often requires informal inputs, such as identifying providers, arranging appointments, providing transportation, coordinating care, and encouraging care recipients to seek treatment. Family members are particularly well positioned to provide these inputs, so greater family involvement can increase rather than reduce the use of formal care.

In this paper, we examine whether and to what extent formal long-term care and informal

¹Online Appendix Figure A.1 shows that informal care take-up in Europe and Denmark is 22 and 33 percent at baseline, rising to 43 and 46 percent among groups exposed to severe health events, and that a large share of this support is provided by adult children.

²This substitution motive is reflected in the design of several major reforms. Japan’s 2000 long-term care insurance program was explicitly framed as the “socialization of care,” made eligibility independent of the availability of family support, and provided services rather than cash benefits to family caregivers (Campbell and Ikegami, 2000; Hanaoka and Norton, 2008). Spain’s 2007 Dependency System created a universal entitlement to home-based and residential services and was explicitly designed to reduce the burden on family caregivers and facilitate women’s continued employment (Spijker and Zuera, 2020). South Korea’s 2008 long-term care insurance reform similarly established a need-based entitlement to home and institutional care, with relieving the caregiving burden on families among its statutory objectives (Kwon, 2009).

family caregiving are substitutes or complements. We focus on how the availability of informal support within the family shapes responses to severe nonfatal health shocks experienced by older parents. Our setting is Denmark, a country in which healthcare coverage is universal and tax-financed and long-term care is primarily allocated on the basis of need and provided free of charge at the point of access. We exploit well-documented heterogeneity in the availability of informal care by the gender composition of adult children. A large literature shows that daughters provide substantially more informal care to aging parents than sons (Grigoryeva, 2017; Bolin et al., 2007), and we show that this source of heterogeneity is unrelated to various measures of parental health prior to the shock.

To identify how family responses to parental health shocks vary with the gender of children, we focus on strokes, severe health events whose timing is difficult to predict (Fadlon and Nielsen, 2019). We compare the evolution of outcomes between parents of only children who experience a stroke at the same age and differ only by whether they have a daughter or a son. Relative to well-established event-study designs in this literature, which exploit variation in the timing of health shocks across families (e.g., Fadlon and Nielsen, 2019, 2021; Frimmel et al., 2025), our design has two key advantages.

First, we compare parents who experience a stroke at the same age, rather than using parents whose stroke occurs later as controls for those whose stroke occurs earlier. This limits concerns that differences in outcomes reflect selection associated with the timing of the shock, particularly because the risk and severity of stroke rise sharply with age (Skajaa et al., 2021). It also aligns families at similar stages of the children’s life cycle. Children’s circumstances, including where they live, whether they are married, and how many children they have, evolve systematically with age and may independently affect parental outcomes.

Second, because our design does not rely on variation in the timing of the stroke, it does not require families to be unresponsive in anticipation of the event. We require only that any anticipatory responses be the same, in expectation, in families with a daughter and families with a son. Prior to the stroke, the two groups exhibit similar levels and trends across a broad range

of health measures, including preexisting conditions (cardiovascular disease, cancer, dementia, diabetes, and mental health conditions), health behaviors (smoking, alcohol use, and obesity), and prior healthcare utilization. These patterns support interpreting subsequent differences in parental survival and care use as responses attributable to heterogeneity in the availability of informal care.

We begin by documenting a striking result: among parents who experience a stroke at the same age, the mortality consequences of the event differ by the gender of their child. Five years after the stroke, parents of daughters are 1.5 percentage points more likely to be alive than parents of sons (p-value <0.01). This difference is sizable relative to baseline stroke mortality: approximately 37 percent of first-time stroke patients in Denmark die within five years (Kristensen et al., 2025), implying that the daughter gradient corresponds to 4.1 percent of the five-year mortality rate among stroke patients. The gradient is small in the year of the stroke and widens steadily over the subsequent five years. Importantly, we find no immediate differences in inpatient care or emergency room visits following the stroke, margins over which patients have little discretion. Together, these patterns make it unlikely that the survival gap is driven by residual differences in underlying health or the initial severity of the stroke.

We next examine informal caregiving as a potential channel underlying the survival gradient. We begin with the differential response of daughters' and sons' residential location to a parental stroke. Specifically, we measure geographic proximity using an indicator for whether the child lives within 10 kilometers of the parent. We find that the effect of the stroke on geographic proximity is 1.6 percentage points larger for daughters than for sons (3.2 percent of the pre-stroke mean). In other words, following the stroke, daughters are more likely than sons to live near their sick parent. Because geographic proximity facilitates frequent, hands-on assistance and responses to day-to-day needs, we interpret it as a proxy for the provision of informal care.

We then examine differential effects on children's own family formation and fertility. A parental stroke has a more negative effect on daughters' probability of being married: the effect is 2.0 percentage points lower for daughters than for sons (3.5 percent of the pre-stroke

mean). Similarly, the effect on fertility is more negative for daughters, who have, on average, 0.053 fewer children relative to sons following the shock (3.6 percent of the pre-stroke mean). Taken together, these results show that, following a parental stroke, daughters adjust more than sons along several important life cycle margins. These responses are consistent with greater involvement in day-to-day caregiving and provide a potential mechanism through which daughters may contribute to improved parental survival.

Finally, we examine whether greater availability of informal care crowds out formal care or instead facilitates its use. We focus on forms of treatment that require sustained engagement after a stroke. Four years after the stroke, the effect on publicly reimbursed spending on physical rehabilitation is \$44 larger per year for parents with daughters than for parents with sons. This difference corresponds to 21 percent of the raw increase in rehabilitation spending among parents with sons. Rehabilitation is a central component of recovery from a severe and potentially disabling event such as a stroke, but accessing and continuing treatment may require information, transportation, scheduling, and encouragement (Winstein et al., 2016). Taken together, these results show that greater informal caregiving does not simply substitute for formal provision. Instead, daughters appear to facilitate parents' use of rehabilitation, pointing to an important complementarity between informal and formal inputs in the response to severe health shocks.

Our results contribute to three strands of literature. First, we build on work on the substitutability of formal and informal long-term care. Recent research shows that informal care crowds out demand for its formal counterpart (Mommaerts, 2018; Barczyk and Kredler, 2019; Coe et al., 2023; Mommaerts, 2025), and additional studies confirm that families provide most care even when formal care is heavily subsidized (Gørtz et al., 2023; Gruber and McGarry, 2023).³ Bolin et al. (2007) use cross-country European data to show that informal and formal home care act as substitutes, while informal care appears to complement doctor and hospital visits. We add to this literature by exploiting variation in informal care availability through child gender rather than in the price of formal care; by extending the outcome from utilization to parental mortality;

³Earlier reduced-form evidence shows that family caregiving lowers nursing home entry and formal home care utilization in the US (Charles and Sevak, 2005; Van Houtven and Norton, 2004, 2008) and in Europe (Bonsang, 2009).

and by studying Denmark, a universal care setting where formal care is free at the point of use, so any residual role for informal care is unlikely to reflect financial substitution.

Second, we contribute to a growing literature on the spillover effects of health shocks within families, particularly adult children. In Denmark, Fadlon and Nielsen (2019) show that severe nonfatal health events lead spouses and adult children to adjust their own health investments. In a follow-up article, Fadlon and Nielsen (2021) examine spousal labor supply responses to similar shocks; Jensen and Zhang (2026) document persistent labor market and mental health consequences of parental death, and Altındağ et al. (2026) find that parental dementia has little effect on adult children’s labor market or physical health outcomes but substantial mental health effects concentrated among daughters. In Taiwan, Chen et al. (2025) document changes in marriage and fertility following old-age health shocks. We complement this line of work by broadening the set of measured responses to include non-monetary margins, including residential mobility, and analyzing the heterogeneity of these responses based on the adult child’s gender.

Finally, our results contribute to a broader literature on gender roles and the intra-household provision of unpaid labor. A large body of work has documented a persistent “child penalty” in women’s earnings and employment following childbirth (Olivetti and Petrongolo, 2017; Kleven et al., 2019, 2025). A second line of work establishes that such penalties can extend beyond parenthood, showing that daughters bear a disproportionate share of eldercare responsibilities, with consequences for their employment and earnings. In Austria, Frimmel et al. (2025) find that adult daughters bear the largest labor supply cost of parental strokes, an effect attenuated by formal care market expansion. Similar gendered caregiving costs appear across US and European settings (Van Houtven et al., 2013; Skira, 2015; Løken et al., 2017; Fahle and McGarry, 2022; Bergeot et al., 2024; Ramirez Lizardi et al., 2025). Gender norms also shape household location decisions: Jayachandran et al. (2024) show that when couples in Germany and Sweden move across commuting zones, men’s earnings rise more than women’s, consistent with households placing more weight on the man’s income. We build on the existing evidence by documenting that this “old-parent penalty” extends beyond labor market outcomes: daughters facing a parental health shock

reduce or delay their own fertility and relocate closer to the affected parent.

2 Institutional Background and Data

This section provides institutional details on long-term elderly care in Denmark, describes our data sources and sample construction, and explains how we identify parental health shocks in the administrative records. Appendix Section A provides additional details on data sources, variable construction, and data cleaning.

2.1 Institutional Setting

Denmark provides a universal, tax-financed system of health and long-term elderly care, with responsibility for older adults divided between two layers of government. Acute and specialist medical care is governed by the Health Act (*Sundhedsloven*) and provided by regional authorities through public hospitals. Long-term care and everyday assistance are regulated by the Social Service Act (*Service-loven*) and delivered by municipalities, which are legally obliged to assess needs and provide services accordingly.

Acute care and rehabilitation Acute treatment for stroke is free of charge, with no co-payments for inpatient care or medically necessary rehabilitation. Regional providers deliver early rehabilitation during and immediately after the hospital stay, while municipalities provide subsequent community-based rehabilitation. Patient out-of-pocket costs are limited to prescription co-payments and non-essential services.

Long-term care Once the acute phase ends, responsibility shifts almost entirely to municipalities. Municipalities must provide personal care and practical assistance according to assessed need, as well as training to maintain functional capacity. For individuals whose needs cannot be met at home, municipalities must offer places in nursing homes or assisted-living facilities. All

services are financed through local taxation and free at the point of use; nursing home residents pay only rent and basic living expenses and eligibility depends solely on functional need.

Family involvement and informal care Danish law imposes no formal caregiving obligation on adult children. In practice, however, municipal in-home care is delivered in short, scheduled visits targeted at specific tasks, leaving activities such as supervision, transportation, medication management, and provider coordination outside the formal system. Spouses and adult children often supply this complementary informal care, particularly for individuals with moderate or fluctuating impairments who may not meet thresholds for extensive municipal support.

The scope for informal care is further amplified by substantial cross-municipal variation in publicly provided care (Jensen and Lolle, 2013), with informal care estimated to account for one-fifth to one-third of total long-term care expenditure in Denmark (Gørtz et al., 2023). Two formal mechanisms allow relatives to take on caregiving roles: (1) full-time municipal employment when the alternative would be institutionalization; (2) a care allowance for relatives of the terminally ill. However, both cover only a small share of care episodes.

2.2 Registry Data and Key Definitions

Our analysis draws on comprehensive Danish administrative registers covering individual financial, demographic, and medical records. All records are linked at the individual level using unique, anonymized personal identifiers from the Civil Registration System (CPR).⁴ We describe each data source and key variable definitions in turn.

Definition of Stroke We identify health events using the *Landspatientregisteret* (National Patient Register, LPR), which records every inpatient contact with the Danish hospital system between 1977 and 2017.⁵ Following Skajaa et al. (2021), we define a stroke as any hospitalization assigned an ICD-10 code in the range I60–I64. To avoid confounding from recurrent strokes, we

⁴CPR identifiers are the Danish equivalent of U.S. Social Security numbers.

⁵Coverage of outpatient care for somatic (i.e., non-psychiatric) conditions is also near universal starting in 1994 (Schmidt et al., 2015).

focus on *first-ever* strokes. At the family level, we define the shock year as the calendar year in which either parent experiences their first stroke, whichever occurs first.⁶

To analyze the effect on mortality, we identify each parent's date of death using the Danish Cause-of-Death Register, maintained by the National Board of Health (*Sundhedsstyrelsen*). This register is constructed from official death certificates issued by attending physicians and records the date and cause of death for every deceased Danish resident.

Healthcare Utilization Beyond identifying stroke events, we use the LPR to construct the measures of hospital utilization and pre-existing health conditions described below. To capture healthcare use outside the hospital system, we complement the LPR with the Health Insurance Register (*Sygeskriftingsregistret*, SSSY), which records services provided under the Danish public health insurance scheme by provider specialty, together with the gross fee paid by the public health insurance for each service. The fee excludes any patient co-payment, so our expenditure measures capture publicly reimbursed spending rather than total treatment costs or patients' out-of-pocket payments. For services that are only partly subsidized, such as ordinary physiotherapy, the recorded fee therefore understates the total cost of treatment.⁷

We measure healthcare utilization using both registers. From the LPR, we classify each hospital contact into three groups: (i) *inpatient* admissions, corresponding to contacts with at least an overnight admission; (ii) *emergency-department (ER)* contacts; and (iii) *outpatient (non-ER)* contacts, comprising all remaining non-emergency ambulatory consultations. We aggregate contact counts to the person-year level. From the SSSY, we construct three further person-year outcomes: total expenditure across all provider specialties, expenditure on general practitioner (GP) contacts, and expenditure on physiotherapy and related long-term-care rehabilitative services. We winsorize each expenditure variable annually at the 0th and 99th percentiles, deflate to 2015

⁶The Danish National Patient Register used ICD-8 coding from 1977 to 1993 and transitioned to ICD-10 in 1994. For stroke events recorded prior to 1994, we identify strokes using ICD-8 codes 430–434 and 436, which correspond to the ICD-10 range used in the main sample period. Because our analysis sample is restricted to first strokes occurring in 1997 or later, the ICD-8 records are used only to screen out parents with a prior stroke and thus to enforce the first-stroke restriction.

⁷SSSY data are available for 2005–2020; for all analyses involving SSSY-based outcomes, we restrict the sample to observations within this window.

Danish kroner using the consumer price index data provided by *Statistics Denmark*, and convert to US dollars at the average exchange rate prevailing in 2015.⁸

Pre-Existing Health Conditions To characterize each parent’s health prior to the shock, we construct indicators for a set of pre-existing conditions from LPR diagnosis codes recorded before the calendar year of the parent’s first stroke. We flag cardiovascular disease, cancer, diabetes, obesity, alcohol-related diagnoses, dementia, and any mental health diagnosis. Because smoking status is not directly recorded in the register, we additionally flag chronic respiratory disease (COPD) as a proxy for smoking, and we construct a composite indicator equal to one if the parent has any of chronic respiratory disease, diabetes, obesity, or an alcohol-related diagnosis. Appendix Table A.2 reports the complete list of ICD-10 codes underlying each health indicator.

Household Financial and Demographic Information Annual income and balance sheet data for all Danish residents come from administrative tax records maintained by SKAT. These records are highly reliable: income and financial accounts are subject to third-party reporting, and tax evasion on such items is minimal (Kleven et al., 2011). All income variables are measured at the level of the individual taxpayer. When referring to total income we measure income pre-tax before labor market contributions, comprising labor income, public transfers, property income, and other taxable income attributable to the individual. Because Denmark provides generous public transfers for caregiving and extended leaves from work, total income may understate the labor market consequences of a shock, as affected individuals can substitute toward transfer income while reducing market work. To isolate effects on children’s labor market prospects, we therefore complement our main analysis with specifications that restrict attention to labor earnings.⁹

Demographic information is obtained from the Civil Registration System (CPR), which records

⁸We use an exchange rate of 6.72 DKK/USD, corresponding to the average Danmarks Nationalbank reference rate over 2015.

⁹Income variables are winsorized, deflated, and converted to US dollars following the same procedure described for healthcare expenditures in Section 2.

household composition and demographic characteristics on a quarterly basis. We link parents to their children through the CPR and classify a parent as having one child if exactly one child is registered to them around the time of the stroke. Educational attainment comes from the Ministry of Education (*Undervisningsministeriet*), which records each individual’s highest completed level of education and professional qualifications.

Household Locations, Geographic Proximity, and Nursing Home Residence We construct annual residential histories for all parents and children in our sample using CPR address records. Geographic distance between parent and child is measured as the distance between the centroids of their respective postal codes of residence. Within this same residential history construction, we further identify parents residing in a nursing home using an address-based imputation method developed by Statistics Denmark (DST). Specifically, all persons registered at an address flagged as a nursing home are treated as nursing home residents in that year.¹⁰

2.3 Data Cleaning and Final Sample

Our sample consists of Danish parents of exactly one child whose first stroke occurred between 1997 and 2017.¹¹ We link both the stroked parent and their child to administrative income and demographic records from Statistics Denmark. The initial dataset contains 957,409 yearly observations covering 59,415 unique families, each consisting of a stroked parent and their only child. We then apply the following restrictions sequentially, in the order reported in Appendix Table A.1.

¹⁰The algorithm first extracts all citizens aged 80+ from the Population Register and identifies each remaining person’s registered address; addresses housing more than six persons aged 80+, and where the elderly-to-young age composition exceeds 75 percent, are flagged as nursing homes. The procedure has been cross-validated against reported municipal data by DST, the Danish Health Data Authority address list, and prior years’ verified data. Source: <https://www.dst.dk/da/TilSalg/data-til-forskning/generelt-om-data/dokumentation-af-data/hoejkvalitetsvariable/aeldreserviceindikator---imputerede-plejehjemsbeboere--afsluttet->

¹¹We observe all strokes recorded in the registry from 1977 onward and use this full history to identify each parent’s first stroke. Individuals whose first stroke occurred before 1997 are excluded from the sample, since our main income and demographic outcomes are only available at high quality from 1996 onward and outcomes must be observed in the year prior to the stroke for our empirical strategy.

First, we drop the few families in which the parent’s recorded date of death precedes the date of the stroke. Second, since we are interested in the effect of a family’s first unanticipated health shock, we keep at most one stroked parent per child, retaining the parent whose first stroke occurred earliest when both parents were stroked (this restriction removes 3.4 percent of unique families). We then apply three restrictions so that each remaining observation can be effectively used in our empirical strategy: the stroked parent must appear in the administrative registers both in the year prior to the stroke ($r = -1$) and in the year of the stroke ($r = 0$); we retain only person-years in which both the parent and the child appear in the registers; and, when a family’s remaining observations contain gaps, we keep only the contiguous spell that contains the stroke year.

Finally, we impose two age restrictions. First, we require the child to be at least 23 years old in the year of the stroke (Frimmel et al., 2025), ensuring that every child is already an adult (aged 18 or older) five years before the stroke, the earliest relative period in our main event window. This restriction excludes 11.8 percent of the remaining families. Second, we restrict the sample to parents who are between ages 60 and 90 at the time of the stroke, excluding a further 15.1 percent of families. This restriction is motivated by statistical power: because our empirical strategy estimates effects nonparametrically by age at stroke, observations in the sparsely populated tails of the age distribution would otherwise generate estimates that are too imprecise to be informative.

Our main analysis sample is an unbalanced panel comprising 619,195 observations from 40,728 unique families. Appendix Table A.1 reports the number of observations and unique families remaining after each sample restriction.

3 Identification and Estimation of Target Parameters

In this section, we provide details on our empirical strategy, with particular attention to the assumptions required to identify our target parameter. Unlike standard difference-in-differences designs in this literature, which estimate the average effect of a health shock by exploiting vari-

ation in its timing across families (e.g., Fadlon and Nielsen, 2019, 2021; Frimmel et al., 2025), our target parameter is the difference in this effect between parents with only a daughter and parents with only a son. To identify it, we do not rely on variation in the timing of health shocks across families. Instead, we exploit variation in the gender composition of children among parents who experience a stroke at the same age and have exactly one adult child. We compare the evolution of outcomes between parents whose only child is a daughter and parents whose only child is a son, holding age at stroke fixed while allowing the event to occur in different calendar years. Relative to designs based on the timing of health shocks, this approach has two key advantages.

First, all comparisons are made within age at stroke. As in Massenkoff and Rose (2024), we organize the comparisons around age rather than calendar time because age is a particularly important determinant of both stroke incidence and severity. The probability of stroke increases monotonically with age (Skajaa et al., 2021), and parents who experience strokes early may be negatively selected on unobserved characteristics, such as health habits, preventive behavior, or risk preferences, that also affect subsequent outcomes. Moreover, several outcomes of adult children, including residential proximity to their parents, marital status, and fertility, vary systematically over the life cycle. Comparing parents whose strokes occur at the same age therefore limits confounding from age-related variation in both parental health and children’s outcomes.

Second, because we do not use parents who experience a stroke later as controls for parents who experience it earlier, we do not need to impose the standard no anticipation condition (Abbring and Van Den Berg, 2003). Families may respond to an impending stroke before it occurs. We require only that any such anticipatory response be the same, in expectation, in families with only a daughter and in families with only a son (Assumption 1 below).

Potential Outcomes Model and Identifying Assumptions. We observe an unbalanced panel of parents indexed by i over years of age indexed by a . For each parent and age, we define a binary treatment indicator $D_{i,a} \in \{0, 1\}$ that equals one if the parent has experienced a stroke by that age. In the rare cases in which a parent experiences multiple strokes, we consider only the first

event. We therefore model treatment as an absorbing state, such that $D_{i,s} = 1 \implies D_{i,w} = 1$ for all $w > s$. Under this structure, an individual's treatment path can be summarized by the treatment cohort G_i , defined as the age at which the parent first experiences a stroke. We denote the set of treatment cohorts by $\mathcal{G}$.

We restrict the analysis sample to parents who eventually experience a stroke. Although we index potential outcomes by treatment cohort, we do not use variation in stroke timing across cohorts to identify the target parameter. Instead, we compare parents with different child gender compositions within the same age-at-stroke cohort.

We examine heterogeneity by the gender composition of the parent's children. We define a family composition indicator $F_i \in \{0, 1\}$ that equals one if the parent has only a daughter and zero if the parent has only a son. Family composition is measured before the stroke and treated as time-invariant. Because all parents in our sample experience a stroke at age 60 or older, their children's gender composition is predetermined well before the health shock.

We consider a potential outcomes model in which outcomes are indexed by the timing of the parental stroke, while family composition is treated as a covariate. Specifically, for each $g \in \mathcal{G}$, $Y_{i,a}(g)$ denotes parent i 's potential outcome at age a under the treatment path in which the stroke occurs at age g .

We first require any anticipatory responses to a future stroke to be the same, in expectation, across family composition types.

Assumption 1 (No Differential Anticipation). *For any age a and pair of treatment cohorts (g, h) such that $a < g < h$,*

$$\mathbb{E}[Y_{i,a}(g) - Y_{i,a}(h) | G_i = g, F_i = 0] = \mathbb{E}[Y_{i,a}(g) - Y_{i,a}(h) | G_i = g, F_i = 1]$$

Assumption 1 adapts the standard no anticipation condition (Abbring and Van Den Berg, 2003) to our setting. It does not require the absence of all anticipatory responses to an impending stroke. Instead, it requires any such response to be the same, in expectation, in families with only

daughters and families with only sons.

The limited predictability of the exact timing of a stroke further supports this assumption. Among parents of similar age and with comparable long-run health characteristics, neither parents nor their adult children are likely to know precisely when the event will occur. This limits the scope for anticipatory responses and, in particular, for such responses to vary systematically with the gender composition of the children (Newman-Toker et al., 2014; Arch et al., 2016; Fadlon and Nielsen, 2019, 2021; Frimmel et al., 2025).

We also impose a common trends assumption adapted to the within-cohort comparisons underlying our design. Among parents who experience a stroke at the same age, we require counterfactual outcome changes under a delayed stroke to be mean-independent of family composition.

Assumption 2 (Common Counterfactual Trends). *For any age a , positive age difference k , and any pair of cohorts (g, h) such that $a - k < g \leq a < h$,*

$$\mathbb{E}[Y_{i,a}(h) - Y_{i,a-k}(h) \mid G_i = g, F_i = 0] = \mathbb{E}[Y_{i,a}(h) - Y_{i,a-k}(h) \mid G_i = g, F_i = 1]$$

For compactness, define $\Delta_{a-k,a}Y_i(h) \equiv Y_{i,a}(h) - Y_{i,a-k}(h)$ as the change in potential outcomes between ages $a - k$ and a under the counterfactual treatment path in which the parent experiences a stroke at age h . Assumption 2 requires this counterfactual change to be the same, in expectation, for families with only daughters and families with only sons whose parent actually experiences a stroke at age g . In other words, absent the stroke at age g , outcomes in the two family types would have evolved similarly.

Target Parameters and Estimands For any age-at-stroke cohort $g \in \mathcal{G}$, relative year $r \geq 0$, and counterfactual stroke age $h > g + r$, define the family-specific average treatment effects on the treated (ATTs) as

$$\text{ATT}_{g+r}(g, f; h) \equiv \mathbb{E}[Y_{i,g+r}(g) - Y_{i,g+r}(h) \mid G_i = g, F_i = f] \quad (1)$$

where $f \in \{0, 1\}$. Because $h > g+r$, the parent has not yet experienced a stroke at age $g+r$ under the counterfactual treatment path indexed by h . The index h therefore denotes a counterfactual delayed-stroke path for parents whose observed treatment cohort is $G_i = g$. It does not index an observed comparison group and, therefore, parents with $G_i = h$ are not used to identify the cohort- g effect.

Our target parameter is the difference between the family-specific treatment effects:

$$\theta_{g+r}(g; h) \equiv \text{ATT}_{g+r}(g, 1; h) - \text{ATT}_{g+r}(g, 0; h) \quad (2)$$

Thus, $\theta_{g+r}(g; h)$ measures how the average effect of a parental stroke differs between families with only daughters and families with only sons, i.e., it captures heterogeneity in the average effect of the stroke by family composition.

Under Assumptions 1 and 2, we identify this parameter as follows. For each relative period $r \geq 0$ and age-at-stroke cohort $g \in \mathcal{G}$:

1. Restrict the sample to units belonging to treatment cohort $G_i = g$.
2. Further restrict the sample to units observed at both ages $a = g - 1$ and $a = g + r$.
3. In this subsample, estimate the saturated regression

$$Y_{i,a} = \alpha_{g,r} + \beta_{g,r}F_i + \gamma_{g,r}\mathbb{I}[a = g + r] + \delta_{g,r}(F_i \times \mathbb{I}[a = g + r]) + U_{i,a} \quad (3)$$

The interaction coefficient $\delta_{g,r}$ equals the difference-in-differences estimand

$$\mathbb{E}[Y_{i,g+r} - Y_{i,g-1} \mid G_i = g, F_i = 1] - \mathbb{E}[Y_{i,g+r} - Y_{i,g-1} \mid G_i = g, F_i = 0]$$

Under Assumptions 1 and 2, $\delta_{g,r}$ identifies $\theta_{g+r}(g; h)$ for any counterfactual stroke age $h > g + r$.

In practice, we account for secular changes in stroke incidence and healthcare technology

by including calendar-year-specific intercepts.¹² When examining outcomes measured for the stroked parent’s children, we additionally include intercepts for the adult child’s age to account for life-cycle patterns in earnings, location choice, family formation, and fertility (Moffitt, 1984; Keane and Wolpin, 2010). We therefore modify the regression in step 3 as follows:

$$Y_{i,a} = \beta_{g,r} F_i + \gamma_{g,r} \mathbb{I}[a = g + r] + \delta_{g,r} (F_i \times \mathbb{I}[a = g + r]) + \lambda_{g,r,t} + \eta_{g,r,z} + V_{i,a} \quad (4)$$

where $\lambda_{g,r,t}$ and $\eta_{g,r,z}$ denote calendar-year-specific and child-age-specific intercepts, respectively.

For each relative period $r \geq 0$, we aggregate the cohort-specific estimates using convex weights, as in Callaway and Sant’Anna (2021). Letting $\bar{a}$ denote the oldest age observed in the sample, we define

$$\delta_r \equiv \sum_{g \in \mathcal{G}} \delta_{g,r} \mathbb{P}(G_i = g \mid G_i + r \leq \bar{a}). \quad (5)$$

The weight assigned to cohort g is its relative share among the cohorts for which horizon r is observed within the sample window. Larger cohorts therefore receive proportionally greater weight, while cohorts treated too late to be observed at horizon r receive zero weight.

To assess the plausibility of the identifying assumptions, we adapt the pre-trend checks proposed by Callaway and Sant’Anna (2021) by using a universal base period. Specifically, for each treatment cohort $g \in \mathcal{G}$ and pre-treatment relative period $r < -1$, we repeat the preceding steps and estimate the corresponding coefficient $\delta_{g,r}$. Relative period $r = -1$ is omitted and serves as the universal base period. Under Assumptions 1 and 2, the population coefficients satisfy $\delta_{g,r} = 0$ for every $r < -1$. The pre-treatment coefficients therefore assess a joint implication of the two identifying assumptions. We aggregate them using the same convex weighting procedure applied to the estimates for nonnegative relative periods.

¹²Over the past three decades, the age-standardized incidence of first-time ischemic stroke in Denmark rose through the late 1990s, peaked around 2002, and subsequently fell below its initial level, with the more recent decline concentrated primarily among people aged 70 or older (Yafasova et al., 2020; Skajaa et al., 2021).

Statistical Inference We cluster standard errors at the parent level to allow for arbitrary correlation in the regression residuals within parents over time.

Validation of the Identifying Assumptions Assumptions 1 and 2 are not directly testable. However, some of their implications can be examined. As discussed in the preceding paragraph, we assess their plausibility in our setting by testing the statistical significance of the pre-stroke coefficients. This approach is not informative for an outcome such as death because, for a parent to experience a stroke, they must be alive. Therefore, $Y_{i,a} = 0 \mid G_i = g$ with probability one at any age $a < g$. As a result, for the death outcome, the common trends assumption simplifies to

$$\mathbb{E}[Y_{i,a}(h) \mid G_i = g, F_i = 0] = \mathbb{E}[Y_{i,a}(h) \mid G_i = g, F_i = 1]$$

for any age a and any pair of cohorts (g, h) such that $g \leq a < h$. Although this assumption is also untestable, we assess its plausibility by comparing the means of several health-related outcomes in the year preceding the stroke between parents with only one daughter and parents with only one son.

Table 1 compares the raw means of parental demographic and health characteristics across the two family types. Parents of a son are, on average, 0.4 years older than parents of a daughter (75.7 versus 75.3) and marginally less likely to have a partner (51.1 versus 51.9 p.p.), while the share of mothers is similar across the two groups (52.9 versus 52.1 p.p.). Stroke type is statistically indistinguishable across family types: ischemic, hemorrhagic, and unspecified strokes occur at nearly the same rates among parents of sons and parents of daughters. Pre-stroke health characteristics are similarly balanced. The prevalence of diabetes and indicators of smoking, alcohol use, and obesity is nearly identical across the two groups, as is the prevalence of dementia (1.8 versus 1.8 p.p.) and of any prior mental health diagnosis (4.6 versus 4.6 p.p.). The two groups also have very similar healthcare utilization in the pre-period across all three measures: inpatient admissions, ER visits, and outpatient visits. Cancer prevalence is slightly higher among parents of daughters (11.3 versus 10.5 p.p.), while nursing home residence is slightly lower (2.5 versus 2.8 p.p.).

These differences are small in magnitude and point in opposite directions in terms of predicted mortality risk. Taken together, the results indicate that the two groups are highly comparable along the observable dimensions most directly related to post-stroke mortality and healthcare use, supporting the plausibility of the assumption that pre-stroke health levels—not merely their trajectories—are similar across the two groups.

Our design compares parents who experience a stroke at the same age. We therefore complement the raw comparisons in Table 1 with a regression-based balance analysis. Continuing to restrict the sample to the year preceding the stroke, we standardize each health-related outcome by its standard deviation and regress it on the family composition indicator F_i . We first report unadjusted estimates and then include age-at-stroke and stroke-year intercepts. Figure 1 presents the resulting coefficients. After accounting for the age and calendar year of the stroke, parents in the two family composition groups remain virtually indistinguishable in their observed pre-stroke health and healthcare use. Across the 19 outcomes, the coefficients from both specifications are small, ranging from -0.016 to 0.028 standard deviations. In the adjusted specification, only 1 of the 19 coefficients is statistically significant at the five percent level. Nor do the point estimates exhibit a systematic directional pattern: 9 are positive and 10 are negative.

4 Results

We report the main results below, following the identification strategy described in Section 3. Online Appendix B.2 complements this analysis by reporting the corresponding conditional mean estimates separately for parents of daughters and parents of sons.

4.1 The Daughter Gradient in Parental Survival


Figure 2 plots the difference in cumulative post-stroke mortality between daughter-only and son-only families, the target parameter defined in equation (2). Five years after the stroke, parents in daughter-only families are 1.5 percentage points more likely to be alive than parents in son-

Table 1: Pre-Stroke Characteristics: Son vs. Daughter (one-child families)

	Full Sample	Only Son	Only Daughter
<i>Parent Demographics and Financial Characteristics</i>			
Age at Stroke, Years	75.5 (7.87)	75.7 (7.81)	75.3 (7.93)
1(Female)	0.525 (0.499)	0.529 (0.499)	0.521 (0.500)
1(Has a Partner)	0.515 (0.500)	0.511 (0.500)	0.519 (0.500)
Parent's Total Household Income, USD 000s	31.6 (37.6)	31.3 (40.8)	32.0 (33.4)
<i>Stroke Type at Onset</i>			
1(Ischemic)	0.420 (0.494)	0.417 (0.493)	0.424 (0.494)
1(Hemorrhagic)	0.131 (0.338)	0.132 (0.338)	0.131 (0.337)
1(Unspecified)	0.448 (0.497)	0.451 (0.498)	0.445 (0.497)
<i>Parent Health Diagnoses at Stroke</i>			
1(Cardiovascular)	0.349 (0.477)	0.346 (0.476)	0.351 (0.477)
1(Cancer)	0.109 (0.311)	0.105 (0.307)	0.113 (0.316)
1(COPD)	0.048 (0.215)	0.050 (0.217)	0.047 (0.212)
1(Diabetes)	0.051 (0.221)	0.051 (0.220)	0.052 (0.223)
1(Obesity)	0.003 (0.058)	0.003 (0.055)	0.004 (0.061)
1(Alcohol-Related)	0.022 (0.148)	0.022 (0.146)	0.023 (0.149)
1(COPD, Diabetes, Obesity, or Alcohol)	0.117 (0.321)	0.117 (0.321)	0.117 (0.322)
1(Dementia)	0.018 (0.133)	0.018 (0.133)	0.018 (0.132)
1(Any Mental Health)	0.046 (0.209)	0.046 (0.209)	0.046 (0.210)
<i>Parent Healthcare Utilization</i>			
Nr. Inpatient Hospital Admissions	0.532 (1.17)	0.540 (1.18)	0.523 (1.16)
Nr. ER Visits	0.230 (0.576)	0.232 (0.578)	0.228 (0.573)
Nr. Outpatient Hospital Visits	1.00 (1.59)	0.994 (1.58)	1.02 (1.60)
Total Practice Sector Fees, USD	572 (586)	566 (579)	578 (594)
General Practitioner Fees, USD	274 (254)	276 (255)	273 (253)
Physiotherapy Fees, USD	53 (316)	51 (309)	54 (323)
1(Lives in Nursing Home)	0.027 (0.161)	0.028 (0.164)	0.025 (0.157)
<i>Child Demographics and Financial Characteristics</i>			
Age at Parent's Stroke, Years	43.6 (7.97)	43.9 (7.88)	43.3 (8.06)
1(Lives within 10km of Parent)	0.492 (0.500)	0.507 (0.500)	0.476 (0.499)
1(Child is Employed)	0.839 (0.368)	0.836 (0.370)	0.842 (0.365)
Child's Total Household Income, USD 000s	56.1 (29.2)	55.8 (31.5)	56.4 (26.2)
1(Child is Married)	0.564 (0.496)	0.549 (0.498)	0.582 (0.493)
Nr. Children	1.46 (1.12)	1.41 (1.15)	1.52 (1.08)
Observations	40,728	21,950	18,778

Notes: This table reports pre-stroke means of parental demographic and health characteristics for the analysis sample of one-child families and for the two subgroups defined by the sex of the child. Column (1) reports the mean for the full sample of one-child families. Column (2) reports the mean for parents whose only child is a son (S). Column (3) reports the mean for parents whose only child is a daughter (D). Means are computed in the calendar year preceding the parent's first stroke. Standard deviations are reported in parentheses. Monetary values are deflated to 2015 Danish kroner and converted to USD at the 2015 average exchange rate. Sample restrictions are described in Section 2.


Figure 1: Balance of Pre-Stroke Health and Healthcare Utilization, by Family Composition



Notes: The figure reports coefficients from separate linear regressions of each pre-stroke health diagnosis and healthcare utilization measure listed in Table 1 on the family composition indicator F_i , which equals one for one-daughter families. All regressions are estimated using outcomes measured in the year preceding the stroke. Each outcome is standardized by its standard deviation, so the coefficients are expressed in standard deviation units and are comparable across measures with different scales. Gray circles report unadjusted estimates, while red diamonds report estimates from specifications that include age-at-stroke and stroke-year intercepts. Horizontal bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The vertical line at zero represents the null hypothesis of no difference between one-daughter and one-son families.

only families. To put this magnitude in perspective, approximately 37 percent of first-time stroke patients in Denmark die within five years (Kristensen et al., 2025). Thus, the observed daughter gradient corresponds to 4.1 percent of the five-year mortality rate among stroke patients. The gap is small in the year of the stroke and widens steadily thereafter.

Figure 2: The Daughter-Son Gap in Parental Mortality Following a Parental Stroke




Notes: The figure plots estimates of the target parameter defined in equation (2), which measures the difference in the average effect of a stroke between one-daughter and one-son families. The outcome is an indicator for whether the parent has died by the observation year. The estimator, described in Section 3, compares the evolution of cumulative mortality between one-daughter and one-son families in which the parent experiences a stroke at the same age, while accounting for calendar time. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The reported pre-mean is the average value of the outcome one year before the stroke.

A concern is that the daughter mortality gradient may reflect differences in stroke severity, which could generate differential use of acute healthcare services after the event. To assess this possibility, we compare inpatient hospital admissions and emergency room visits among parents of daughters and parents of sons following the stroke. Panel (a) of Figure 3 shows that the daughter gradient in inpatient service utilization is not statistically different from zero at conventional significance levels. Panel (b) displays a similar pattern for the number of emergency room visits. The absence of a difference in this relatively nondiscretionary margin of care provides further ev-

idence that the post-stroke mortality gap is not driven by differences in the severity of the initial health shock.

Figure 3: The Daughter-Son Gap in Inpatient Hospital Admissions and ER Visits



Notes: The figure plots estimates of the target parameter defined in equation (2), which measures the difference in the average effect of a stroke between one-daughter and one-son families. In panel (a), the outcome is the annual number of inpatient hospital admissions. In panel (b), the outcome is the annual number of emergency room visits. The estimator, described in Section 3, compares the evolution of these outcomes between one-daughter and one-son families in which the parent experiences a stroke at the same age, while accounting for calendar time. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The reported pre-means are the average values of the respective outcomes one year before the stroke.

4.2 The Daughter Gradient in the Supply of Informal Care

We next turn to the provision of informal care, which, by definition, is not directly observed in our data. We therefore examine several outcomes that plausibly proxy for the provision of care and assistance to a parent following a nonfatal health shock. Across these outcomes, we find consistent evidence of a sizable daughter gradient in the supply of informal care.

4.2.1 Residential Proximity

We first consider residential proximity as a margin of adjustment to parental health shocks. Living closer to a parent makes it easier to provide informal care, coordinate formal care, or respond to

emergencies.

Figure 4 plots the daughter-versus-son gap in the effect of a parental stroke on the probability that the adult child lives within 10 kilometers of the affected parent, measured using the distance between postal code area centroids (Section 2). Daughters and sons show no differential trends in proximity before the stroke. After the stroke, daughters increase their proximity by approximately 1.6 percentage points more than sons, a differential of about 3.2 percent relative to the pre-stroke baseline proximity rate of 49.2 percent (reported in Table 1). The gap emerges gradually over the five years following the stroke, consistent with residential adjustments materializing through relocation decisions that unfold over time. Overall, this pattern indicates that proximity acts as an enabling margin for the informal caregiving provided by daughters.


4.2.2 Labor Supply

Panel (a) of Figure 5 shows that sons in one-son families reduce their earnings by more than daughters in one-daughter families following a parental stroke. Five years after the event, the differential is approximately \$836 per year—roughly 1.6 percent of the pre-stroke earnings baseline in one-son families (\$50,722)—and statistically significant at conventional levels ($p\text{-value} < 0.05$). Panel (b) shows no differential effect on the probability of employment, indicating that the response operates on the intensive margin (hours or wages within employment) rather than the extensive margin of labor market participation.

4.2.3 Family Formation and Fertility

A final margin of adjustment is the adult child's own family formation. Figure 6 shows no differential trends in marriage or fertility between daughters and sons before the parental stroke. Both outcomes diverge measurably thereafter. Five years after the event, daughters are approximately 2.0 percentage points less likely to be married than sons, and have accumulated approximately 0.053 fewer children on average. Both results indicate that daughters are more likely to trade off their own family formation against the caregiving and residential commitments documented

Figure 4: The Daughter-Son Gap in Geographic Proximity to the Stroked Parent




Notes: The figure plots estimates of the target parameter defined in equation (2), which measures the difference in the average effect of a stroke between one-daughter and one-son families. The outcome is an indicator for whether the child resides within 10 kilometers of the parent. Distance is measured between the centroids of the child's and parent's postal codes of residence. The estimator, described in Section 3, compares the evolution of residential proximity between one-daughter and one-son families in which the parent experiences a stroke at the same age, while accounting for calendar time. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The reported pre-mean is the average value of the outcome one year before the stroke.

above.

Taken together, these results indicate that sons and daughters adjust to parental health shocks along different margins. Their labor supply and earnings responses are similar, but daughters exhibit larger residential adjustments and greater declines in fertility. Overall, daughters appear more likely to respond through nonpecuniary forms of support, particularly location choices that facilitate the provision of informal care.

Figure 5: The Daughter-Son Gap in Labor Market Outcomes




Notes: The figure plots estimates of the target parameter defined in equation (2), which measures the difference in the average effect of a stroke between one-daughter and one-son families. In panel (a), the outcome is the adult child's annual wage income, measured in thousands of U.S. dollars. In panel (b), it is an indicator for whether the adult child is employed. The estimator, described in Section 3, compares the evolution of these outcomes between one-daughter and one-son families in which the parent experiences a stroke at the same age, while accounting for calendar time and the adult child's age. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The reported pre-means are the average values of the respective outcomes one year before the stroke.

4.3 Does Family Caregiving Complement Formal Care?

The evidence in Section 4.2 points to greater caregiving by daughters as a plausible contributor to the daughter mortality gradient. We next examine how family involvement relates to parents' use of formal care. Family support may replace some professional services while facilitating access to others, particularly treatments requiring transportation, coordination, and sustained engagement. Substitution and complementarity may therefore coexist across different types of care. To explore these possibilities, we compare post-stroke use of several formal care services between parents of daughters and parents of sons.

Panels (a) and (b) of Figure 7 show limited evidence of differential responses in outpatient hospital use and nursing home residence by child gender. In Panel (a), the estimated daughter-son differences in non-emergency outpatient hospital visits are generally positive at longer horizons,

Figure 6: The Daughter-Son Gap in Marriage and Number of Children




Notes: The figure plots estimates of the target parameter defined in equation (2), which measures the difference in the average effect of a stroke between one-daughter and one-son families. In panel (a), the outcome is an indicator for whether the adult child is married or in a registered partnership. In panel (b), it is the adult child's number of children, equivalently the parent's number of grandchildren. The estimator, described in Section 3, compares the evolution of these outcomes between one-daughter and one-son families in which the parent experiences a stroke at the same age, while accounting for calendar time and the adult child's age. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The reported pre-means are the average values of the respective outcomes one year before the stroke.

but the confidence intervals include zero. Panel (b) likewise shows no statistically significant differential change in nursing home residence, although residence rates rise substantially in both groups following the stroke (Online Appendix Figure A.6, Panel (b)). These estimates provide no clear evidence of substitution along these margins, but their precision is insufficient to rule out economically meaningful reductions in formal care use.

The key evidence comes from physiotherapy and related rehabilitation services, reported in Panel (d). Parents whose only child is a daughter experience larger post-stroke increases in the use of these services than parents whose only child is a son, with statistically significant differences at multiple horizons before year five. These earlier estimates indicate a sustained difference in rehabilitation use during the years following the stroke. Rehabilitation typically requires repeated appointments and continued engagement, creating a role for family members in arranging transportation, coordinating treatment, and encouraging attendance. The findings are therefore

consistent with daughters facilitating access to and continued use of formal care. Panel (c) also shows positive post-stroke differences in total practice sector fees, with the rehabilitation results providing the clearest evidence consistent with complementarity.

Figure 7: The Daughter-Son Gap in Parental Healthcare Utilization



Notes: The figure plots estimates of the target parameter defined in equation (2), which measures the difference in the average effect of a stroke between one-daughter and one-son families. In panel (a), the outcome is the annual number of non-emergency outpatient hospital visits. In panel (b), it is an indicator for whether the affected parent resides in a nursing home. Panel (c) uses the parent's total annual physician fees, measured in U.S. dollars, while panel (d) uses the parent's total annual fees for physiotherapy and long-term care services, also measured in U.S. dollars. The estimator, described in Section 3, compares the evolution of these outcomes between one-daughter and one-son families in which the parent experiences a stroke at the same age, while accounting for calendar time. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level. The reported pre-means are the average values of the respective outcomes one year before the stroke.

Taken together, these patterns suggest that daughters help parents access and sustain formal care, potentially contributing to the observed survival advantage.

5 Conclusion

Long-term care for older adults is an increasingly important challenge in developed countries. Governments have responded by expanding formal care systems, often with the explicit objective of reducing families' caregiving responsibilities and under the presumption that formal and informal care are substitutes. This paper provides evidence that, following a severe parental health shock, the relationship between the two may instead involve an important degree of complementarity.

Using Danish administrative data, we compare parents of only children who experience a stroke at the same age but differ in the gender of their child. We document a substantial survival gradient: five years after the stroke, parents of daughters are 1.5 percentage points more likely to be alive than parents of sons. This difference is not accompanied by differential use of inpatient care or emergency room services, two relatively nondiscretionary margins of healthcare use, making differences in stroke severity an unlikely explanation. Instead, daughters become more likely to live near their parents and experience larger changes in their own family trajectories: they are less likely to marry and have fewer children after the shock. These findings indicate that the provision of informal care imposes meaningful costs on caregivers' lives.

Crucially, we show that parents of daughters receive more publicly reimbursed physical rehabilitation, a type of formal care that requires sustained engagement after the acute phase of a stroke. Four years after the stroke, the effect on publicly reimbursed spending on physical rehabilitation is \$44 larger per year for parents with daughters than for parents with sons. This difference corresponds to 21 percent of the raw increase in rehabilitation spending among parents with sons. This result suggests that daughters help parents access and continue formal treatment rather than simply replacing professional services with care provided at home. Taken together,

the findings indicate that family caregiving and formal care operate, at least in part, as complements in shaping recovery and survival after a severe health shock. More broadly, policies that expand formal coverage may be less effective for older adults with limited family support unless they also address the informational, logistical, and coordination barriers that determine whether formal services are actually used.

References

- Abbring, Jaap H. and Gerard J. Van Den Berg**, “The Nonparametric Identification of Treatment Effects in Duration Models,” *Econometrica*, September 2003, 71 (5), 1491–1517.
- Altındağ, Onur, Jane Greve, and Yulya Truskinovsky**, “Family Spillovers of Dementia,” Technical Report, National Bureau of Economic Research 2026.
- Arch, Allison E., David C. Weisman, Steven Coca, Karin V. Nystrom, Charles R. Wira, and Joseph L. Schindler**, “Missed Ischemic Stroke Diagnosis in the Emergency Department by Emergency Medicine and Neurology Services,” *Stroke*, March 2016, 47 (3), 668–673.
- Barczyk, Daniel and Matthias Kredler**, “Long-Term Care Across Europe and the United States: The Role of Informal and Formal Care,” *Fiscal Studies*, 2019, 40 (3), 329–373.
- Bergeot, Julien, Irene Ferrari, and Ya Gao**, “The Effect of Parental Health Shocks on Living Arrangements and Employment,” *Health Economics*, 2024, 33 (12), 2798–2837.
- Bolin, Kristian, Bjorn Lindgren, and Petter Lundborg**, “Informal and Formal Care Among Single-Living Elderly in Europe,” Technical Report, Tinbergen Institute Discussion Paper 2007.
- Bonsang, Eric**, “Does Informal Care from Children to Their Elderly Parents Substitute for Formal Care in Europe?,” *Journal of Health Economics*, 2009, 28 (1), 143–154.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, December 2021, 225 (2), 200–230.
- Campbell, John Creighton and Naoki Ikegami**, “Long-Term Care Insurance Comes to Japan,” *Health Affairs*, 2000, 19 (3), 26–39.
- Charles, Kerwin Kofi and Purvi Sevak**, “Can Family Caregiving Substitute for Nursing Home Care?,” *Journal of Health Economics*, 2005, 24 (6), 1174–1190.

- Chen, Kuan-Ming, Maxwell Kellogg, and Kuan-Ju Tseng**, “Family Trajectories and the Burden of Care in the Aftermath of Old-Age Health Shocks,” *SSRN Electronic Journal*, 2025.
- Ciccarelli, Nicola and Arthur Van Soest**, “Informal Caregiving, Employment Status and Work Hours of the 50+ Population in Europe,” *De Economist*, September 2018, 166 (3), 363–396.
- Coe, Norma B., Gopi Shah Goda, and Courtney Harold Van Houtven**, “Family Spillovers and Long-Term Care Insurance,” *Journal of Health Economics*, 2023, 90, 102781.
- Colombo, Francesca, Ana Llana-Nozal, Jérôme Mercier, and Frits Tjadens**, *Help Wanted? Providing and Paying for Long-Term Care* OECD Health Policy Studies, Paris: OECD Publishing, 2011.
- Fadlon, Itzik and Torben Heien Nielsen**, “Family Health Behaviors,” *American Economic Review*, September 2019, 109 (9), 3162–3191.
- and —, “Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records,” *American Economic Journal: Applied Economics*, July 2021, 13 (3), 1–30.
- Fahle, Sean and Kathleen McGarry**, “How Caregiving for Parents Reduces Women’s Employment: Patterns Across Sociodemographic Groups,” in Lisa F. Berkman and Beth C. Truesdale, eds., *Overtime: America’s Aging Workforce and the Future of Working Longer*, Oxford University Press, 2022.
- Frimmel, Wolfgang, Martin Halla, Jörg Paetzold, and Julia Schmieder**, “Health of Parents, Their Children’s Labor Supply, and the Role of Migrant Care Workers,” *Journal of Labor Economics*, 2025, 43 (3), 803–841.
- Gørtz, Mette, Bent Jesper Christensen, and Nabanita Datta Gupta**, “Long-Term Care in Denmark,” NBER Working Paper 31889, National Bureau of Economic Research 2023.

- Grigoryeva, Angelina**, “Own Gender, Sibling’s Gender, Parent’s Gender: The Division of Elderly Parent Care Among Adult Children,” *American Sociological Review*, 2017, 82 (1), 116–146.
- Gruber, Jonathan and Kathleen McGarry**, “Long-Term Care in the United States,” NBER Working Paper 31881, National Bureau of Economic Research 2023.
- Hanaoka, Chie and Edward C. Norton**, “Informal and Formal Care for Elderly Persons: How Adult Children’s Characteristics Affect the Use of Formal Care in Japan,” *Social Science & Medicine*, 2008, 67 (6), 1002–1008.
- Jayachandran, Seema, Lea Nassal, Matthew J Notowidigdo, Marie Paul, Heather Sarsons, and Elin Sundberg**, “Moving to Opportunity, Together,” *National Bureau of Economic Research Working Paper Series*, 2024, No. 32970.
- Jensen, Mathias Fjællegaard and Ning Zhang**, “Effects of Parental Death on Labor Market Outcomes,” *American Economic Review*, 2026, 116 (5), 1811–1836.
- Jensen, Per H. and Henrik Lolle**, “The Fragmented Welfare State: Explaining Local Variations in Services for Older People,” *Journal of Social Policy*, April 2013, 42 (2), 349–370.
- Keane, Michael P. and Kenneth I. Wolpin**, “The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic, and White Women,” *International Economic Review*, 2010, 51 (3), 851–892.
- Kleven, Henrik, Camille Landais, and Gabriel Leite-Mariante**, “The Child Penalty Atlas,” *The Review of Economic Studies*, October 2025, 92 (5), 3174–3207.
- , —, and **Jakob Egholt Sogaard**, “Children and Gender Inequality: Evidence from Denmark,” *American Economic Journal: Applied Economics*, October 2019, 11 (4), 181–209.
- Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez**, “Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark,” *Econometrica*, 2011, 79 (3), 651–692.

- Kristensen, Frederik Pagh Bredahl, Helene Matilde Lundsgaard Svane, Kristina Lauge-sen, Sofie Kejlberg Al-Mashhadi, Diana Hedevang Christensen, Henrik Toft Sørensen, and Nils Skajaa**, “Risk of Mortality and Recurrence After First-Time Stroke Among Patients with Type 2 Diabetes: A Danish Nationwide Cohort Study,” *European Stroke Journal*, 2025, 10 (1), 190–197.
- Kwon, Soonman**, “The Introduction of Long-Term Care Insurance in South Korea,” *Eurohealth*, 2009, 15 (1), 28–29.
- Løken, Katrine V., Shelly Lundberg, and Julie Riise**, “Lifting the Burden: Formal Care of the Elderly and Labor Supply of Adult Children,” *Journal of Human Resources*, 2017, 52 (1), 247–271.
- Massenkoff, Maxim and Evan K. Rose**, “Family Formation and Crime,” *American Economic Journal: Applied Economics*, October 2024, 16 (4), 444–483.
- Moffitt, Robert**, “Profiles of Fertility, Labour Supply and Wages of Married Women: A Complete Life-Cycle Model,” *The Review of Economic Studies*, 1984, 51 (2), 263–278.
- Mommaerts, Corina**, “Are Coresidence and Nursing Homes Substitutes? Evidence from Medicaid Spend-Down Provisions,” *Journal of Health Economics*, 2018, 59, 125–138.
- , “Long-Term Care Insurance and the Family,” *Journal of Political Economy*, 2025, 133 (1).
- Newman-Toker, David E., Ernest Moy, Ernest Valente, Rosanna Coffey, and Anika L. Hines**, “Missed Diagnosis of Stroke in the Emergency Department: A Cross-Sectional Analysis of a Large Population-Based Sample,” *Diagnosis*, June 2014, 1 (2), 155–166.
- Olivetti, Claudia and Barbara Petrongolo**, “The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries,” *Journal of Economic Perspectives*, 2017, 31 (1), 205–230.

Ramirez Lizardi, Eduardo, Elisabeth Fevang, Knut Røed, and Henning Øien, “Effects of Health Shocks on Adult Children’s Labor Market Outcomes and Well-Being,” *Health Economics*, 2025, 34 (10), 1804–1820.

Rocard, Eileen and Ana Llana-Nozal, “Supporting Informal Carers of Older People: Policies to Leave No Carer Behind,” Technical Report 140, OECD Publishing, Paris 2022.

Schmidt, Morten, Sigrun Alba Johannesdottir Schmidt, Jakob Lynge Sandegaard, Vera Ehrenstein, Lars Pedersen, and Henrik Toft Sørensen, “The Danish National Patient Registry: A Review of Content, Data Quality, and Research Potential,” *Clinical Epidemiology*, November 2015, p. 449.

Skajaa, Nils, Kasper Adelborg, Erzsébet Horváth-Puhó, Kenneth J. Rothman, Victor W. Henderson, Lau Casper Thygesen, and Henrik Toft Sørensen, “Nationwide Trends in Incidence and Mortality of Stroke Among Younger and Older Adults in Denmark,” *Neurology*, March 2021, 96 (13).

Skira, Meghan M., “Dynamic Wage and Employment Effects of Elder Parent Care,” *International Economic Review*, 2015, 56 (1), 63–93.

Spijker, Jeroen and Pilar Zueras, “Old-Age Care Provision in Spain in the Context of a New System of Long-Term Care and a Lingered Economic Crisis,” *Journal of Population Ageing*, 2020, 13 (1), 41–62.

Van Houtven, Courtney Harold and Edward C. Norton, “Informal Care and Health Care Use of Older Adults,” *Journal of Health Economics*, 2004, 23 (6), 1159–1180.

— **and** — , “Informal Care and Medicare Expenditures: Testing for Heterogeneous Treatment Effects,” *Journal of Health Economics*, 2008, 27 (1), 134–156.

— , **Norma B. Coe, and Meghan M. Skira**, “The Effect of Informal Care on Work and Wages,” *Journal of Health Economics*, 2013, 32 (1), 240–252.

Winstein, Carolee J., Joel Stein, Ross Arena, Barbara Bates, Leora R. Cherney, Steven C. Cramer, Frank Deruyter, Janice J. Eng, Beth Fisher, Richard L. Harvey, Catherine E. Lang, Marilyn MacKay-Lyons, Kenneth J. Ottenbacher, Sue Pugh, Mathew J. Reeves, Lorie G. Richards, William Stiers, and Richard D. Zorowitz, “Guidelines for Adult Stroke Rehabilitation and Recovery: A Guideline for Healthcare Professionals from the American Heart Association/American Stroke Association,” *Stroke*, 2016, 47 (6), e98–e169.

Yafasova, Adelina, Emil Loldrup Fosbøl, Mia Nielsen Christiansen, Naja Emborg Vinding, Charlotte Andersson, Christina Kruuse, Søren Paaske Johnsen, Gunnar Hilmar Gislason, Christian Torp-Pedersen, Lars Køber, and Jawad Haider Butt, “Time Trends in Incidence, Comorbidity, and Mortality of Ischemic Stroke in Denmark (1996–2016),” *Neurology*, 2020, 95 (17), e2343–e2353.

Online Appendix

Table of Contents

A	Data	37
A.1	Sample Restrictions	37
A.2	ICD-10 Code Definitions	38
A.3	Household Data Details	39
A.4	Data Citations	40
B	Additional Empirical Results	41
B.1	Descriptive Patterns in Informal Care Use in Europe	41
B.2	Estimates of Raw Conditional Outcome Means Around the Event	42

A Data

A.1 Sample Restrictions

Table A.1: Sample Restrictions

Restriction	Observations	Unique Families
Initial sample	957,409	59,415
Inconsistent death date	957,383	59,413
First stroked parent per child	922,512	57,415
Parent observed at stroke	900,178	54,692
Parent and child observed in person-year	870,353	54,692
Contiguous spell containing event	855,557	54,378
Child's age at stroke: 23 or older	747,544	47,958
Parent's age at stroke: 90 or younger	722,887	46,400
Parent's age at stroke: 60 or older	619,195	40,728
Final analysis sample	619,195	40,728

Notes: The table reports the numbers of observations and families remaining after each successive sample restriction. Restrictions are applied sequentially at the family level: if any child of a shocked parent fails a restriction, the entire family is excluded for the given calendar year. Sample restrictions are described in Section 2.

A.2 ICD-10 Code Definitions

Table A.2 reports the ICD-10 codes used to define the stroke event itself and each pre-existing health condition. All diagnoses are drawn from the National Patient Register (LPR; Section 2).

Pre-existing condition indicators equal one if the parent has any hospital contact from 1996 onward whose primary diagnosis carries the listed code, recorded before the calendar year of the first stroke.

Table A.2: ICD-10 Code Definitions

Condition	ICD-10 Code(s)
<i>Stroke event classification</i>	
Stroke (any subtype)	I60–I64
Hemorrhagic Stroke	I60–I62
Ischemic Stroke	I63
Unspecified Stroke	I64
<i>Parent pre-existing conditions</i>	
Any cardiovascular diagnosis	I (all codes)
Cancer (incl. in-situ neoplasms)	C00–C97, D00–D09
COPD / chronic respiratory disease	J40–J44
Diabetes	E10–E14
Obesity	E66
Alcohol-related diagnosis or liver fibrosis/cirrhosis	F10; K70, K74; K86.0
COPD, diabetes, obesity, or alcohol-related diagnosis	(union of above four)
Dementia and other degenerative neurological disease	F00–F03; G30, G31
Any mental health diagnosis	F (all codes)

Notes: Stroke events are identified from the primary diagnosis of hospital contacts. For contacts recorded before 1994 under ICD-8, codes 430–432 are classified as hemorrhagic, 433–434 as ischemic, and 436 as unspecified stroke. ICD-10 chapter I (circulatory system) contains the stroke codes themselves; because the cardiovascular-disease indicator is measured strictly before the calendar year of the stroke, the stroke event does not mechanically satisfy this condition.

A.3 Household Data Details

Definition of Household We follow Statistics Denmark and define a household as either a single person or a couple (married, in a registered partnership, or cohabiting under CPR-based criteria) together with their cohabiting children. Children are included in the parental family if they live at the same address, are under the age of 25, unmarried, without children, and not part of a cohabiting couple. Cohabiting couples are identified either (i) if they have a child together registered in the CPR, or (ii) if they do not have children together, they are of opposite sex, not closely related, have an age difference of less than 15 years, and constitute the only two adults in the household.

Aggregation of Individual Administrative Records at the Household Level Financial and income data are reported at the individual level in personal tax records and must therefore be aggregated to the household level for computation of household-level measures of income in Table 1.

Practice Sector Service Classification Our measure of fees for physiotherapy and related long-term-care rehabilitative services sums the fees recorded in the Health Insurance Register (SSSY) under the following provider specialty codes: 51 and 61 (physiotherapy), 58 (rehabilitation, *genoptræning*), 62 (physiotherapy free of charge, *vederlagsfri fysioterapi*), 64 (chiropractic care for the chronically ill), and 65 (riding physiotherapy free of charge, *vederlagsfri ridefysioterapi*).


A.4 Data Citations

- Statistics Denmark, Befolkningen (BEF; Population Register)
- Statistics Denmark, Indkomst (IND; Income from tax returns)
- Statistics Denmark, Uddannelser (UDDA; Education)
- Statistics Denmark, Boligtællingen (BOL; Housing census)
- Statistics Denmark, Ejendomme salgsoplysninger (EJSA; Property sales information)
- Statistics Denmark, Ejendomsskatter (EJSK; Property taxes)
- Statistics Denmark, Ejere af ejendomme (EJER; Property owners)
- Statistics Denmark, Detaljeret lønmodtagerdata fra e-Indkomst (BFL; Detailed employee data from e-Income)
- Statistics Denmark, Landspatientregistret (LPR; National Patient Register)
- Statistics Denmark, Døde i Danmark (DOD; Deaths in Denmark)
- Statistics Denmark, Boligfinansiering (REAL; Housing finance)

B Additional Empirical Results

B.1 Descriptive Patterns in Informal Care Use in Europe

Figure A.1: Informal Care in Europe and Denmark



Notes: The figure compares informal care receipt by reported stroke history in Europe (Panel (a)) and Denmark (Panel (b)). Within each stroke-history group, separate bars show the shares receiving *any* informal care, care from any adult child, and care from any daughter. These categories overlap and are not additive. Data are from the Survey of Health, Ageing and Retirement in Europe (SHARE), waves 1-9 (2004-2022). Panel (a) pools SHARE countries that are EU members. The any-care and any-child-care rates are constructed from the harmonized main file across waves 1-9, excluding wave 3 (SHARELIFE). The daughter-care rates use the wave-8 child roster and therefore refer to a different survey sample. Reported sample sizes are: Europe, $N_{\text{no stroke}} = 136,885$ and $N_{\text{ever stroke}} = 12,380$; Denmark, $N_{\text{no stroke}} = 5,630$ and $N_{\text{ever stroke}} = 555$.

B.2 Estimates of Raw Conditional Outcome Means Around the Event

This section reports estimates of the conditional mean of outcomes around the parental stroke. For each realization of the family composition indicator $f \in \{0, 1\}$, age-at-stroke cohort $g \in \mathcal{G}$, and relative year $r = a - g$, we proceed as follows:

1. Restrict the sample to units belonging to age-at-stroke cohort $G_i = g$.
2. Further restrict the sample to units for whom the family composition indicator is $F_i = f$.
3. Within this subsample, estimate the conditional mean

$$\mu_{g+r}(g, f) \equiv \mathbb{E}[Y_{i,g+r} | G_i = g, F_i = f] \quad (\text{A.1})$$


To obtain a single conditional mean for each relative period r , we aggregate the cohort-specific estimates across age-at-stroke cohorts using convex weights, as in Callaway and Sant’Anna (2021). Letting $\bar{a}$ denote the oldest age observed in the sample, we define

$$\mu_r(f) = \sum_{g \in \mathcal{G}} \mu_{g+r}(g, f) \times \mathbb{P}(G_i = g | G_i + r \leq \bar{a}) \quad (\text{A.2})$$

As in the difference-in-differences specification in the main text, the weight assigned to cohort g is its relative share among the cohorts for which horizon r is observed within the sample window.


Importantly, we use the same cohort weights for $f = 0$ and $f = 1$. That is, for each age-at-stroke cohort and relative year, the weights do not depend on whether parents with daughters are more or less prevalent than parents with sons. We therefore interpret $\mathbb{P}(G_i = g | G_i + r \leq \bar{a})$ as composition-invariant convex weights used to aggregate the conditional means for both family composition groups. This ensures that differences in the aggregated conditional means reflect differences in outcomes within age-at-stroke cohorts rather than differences in the cohort composition of the two groups.

Figure A.2: The Daughter-Son Gap in Inpatient Hospital Admissions and ER Visits




Notes: The figure reports estimates of the target parameter defined in equation (A.1), the conditional mean of each outcome by year relative to the parental stroke, separately for parents with a daughter and parents with a son. As described in Appendix Section B.2, conditional means are first estimated separately for each age-at-stroke cohort and relative year and then aggregated across cohorts within each relative year using convex weights. In panel (a), the outcome is the annual number of inpatient hospital admissions; in panel (b), it is the annual number of emergency room visits. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the parent level. Pre-means correspond to the average value of each outcome one year before the stroke.

Figure A.3: The Daughter-Son Gap in Geographic Proximity to the Stroked Parent




Notes: The figure reports estimates of the target parameter defined in equation (A.1), the conditional mean of the outcome by year relative to the parental stroke, separately for parents with a daughter and parents with a son. As described in Appendix Section B.2, conditional means are first estimated separately for each age-at-stroke cohort and relative year and then aggregated across cohorts within each relative year using convex weights. The outcome is an indicator for whether the child resides within 10 kilometers of the parent. Distance is measured between the centroids of the child's and parent's residential postal codes. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the parent level. The pre-mean corresponds to the average value of the outcome one year before the stroke.

Figure A.4: The Daughter-Son Gap in Labor Market Outcomes




Notes: The figure reports estimates of the target parameter defined in equation (A.1), the conditional mean of each outcome by year relative to the parental stroke, separately for parents with a daughter and parents with a son. As described in Appendix Section B.2, conditional means are first estimated separately for each age-at-stroke cohort and relative year and then aggregated across cohorts within each relative year using convex weights. In panel (a), the outcome is the adult child's annual wage income, measured in thousands of U.S. dollars; in panel (b), it is an indicator for whether the adult child is employed. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the parent level. Pre-means correspond to the average value of each outcome one year before the stroke.

Figure A.5: The Daughter-Son Gap in Marriage and Number of Children



Notes: The figure reports estimates of the target parameter defined in equation (A.1), the conditional mean of each outcome by year relative to the parental stroke, separately for parents with a daughter and parents with a son. As described in Appendix Section B.2, conditional means are first estimated separately for each age-at-stroke cohort and relative year and then aggregated across cohorts within each relative year using convex weights. In panel (a), the outcome is an indicator for whether the adult child is married or in a registered partnership. In panel (b), it is the adult child's number of children, equivalently the parent's number of grandchildren. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the parent level. Pre-means correspond to the average value of each outcome one year before the stroke.

Figure A.6: The Daughter-Son Gap in Parental Healthcare Utilization



Notes: The figure reports estimates of the target parameter defined in equation (A.1), the conditional mean of each outcome by year relative to the parental stroke, separately for parents with a daughter and parents with a son. As described in Appendix Section B.2, conditional means are first estimated separately for each age-at-stroke cohort and relative year and then aggregated across cohorts within each relative year using convex weights. In panel (a), the outcome is the annual number of non-emergency outpatient hospital visits. In panel (b), it is an indicator for whether the affected parent resides in a nursing home. Panel (c) uses the parent's total annual physician fees, measured in U.S. dollars, while panel (d) uses the parent's total annual fees for physiotherapy and long-term care services, also measured in U.S. dollars. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the parent level. Pre-means correspond to the average value of each outcome one year before the stroke.