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

Elin Colmsjö
University of Copenhagen & IFN

Francesco Ruggieri
University of Virginia

Matteo Saccarola
UC Berkeley

July 31, 2026

Abstract

We study how informal family care shapes survival and formal healthcare use after severe health shocks. 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 are on 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 1.6 percentage point more likely to be alive by year five. This is not accompanied by differences in inpatient care or emergency room visits, margins over which patients have little discretion, making differential stroke severity an unlikely explanation. Instead, daughters are more likely to live near their parents, are less likely to marry, and have fewer children after the shock, indicating that the burden of informal care affects caregivers' own family trajectories. Crucially, parents of daughters use more physical rehabilitation, an elective type of formal care that requires sustained engagement. This suggests that daughters do not merely substitute for formal care but help parents access and continue it: informal and formal care operate, at least in part, as complements in shaping post-stroke survival.

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 Haas Seminar in Real Estate, UC Berkeley Psychology and Economics Seminar, UC Berkeley Finance Student Workshop, Northwestern University Micro Gender Student Seminar, and 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.

1 Introduction

Life expectancy has risen steadily across OECD countries, increasing the demand for long-term care for older adults (Colombo et al., 2011). Although formal care institutions play an important role, family members often 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 Llena-Nozal, 2022). These pressures are especially pronounced when elderly parents experience nonfatal health shocks that generate persistent needs for assistance over extended periods (Ciccarelli and Van Soest, 2018). 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.[1]

Yet the take-up of formal care remains limited even when services are heavily subsidized.[2] This pattern may partly reflect a strong preference for care provided by relatives.[3] It may also arise if formal coverage does not translate into effective access. Obtaining formal services may require information, applications, transportation, scheduling, and coordination—inputs that family members are particularly well positioned to provide. Informal support may therefore complement the use of formal care, even when the two substitute for one another in the direct provision of assistance. As a result, older adults with the weakest family support, and thus potentially the

[1]This 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 Zueras, 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).

[2]In Denmark, where access to care is universal, three-quarters of elderly care recipients rely exclusively on informal help, and informal care accounts for between one-fifth and one-third of total long-term care costs (Gørtz et al., 2023). In the United States, roughly 70 percent of elderly care recipients receive only informal home care, and informal care accounts for over one-third of total long-term care costs (Gruber and McGarry, 2023).

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

greatest need for formal insurance, may also be the least able to benefit from it.

In this paper, we study how the availability of informal care within the family shapes responses to severe nonfatal health shocks experienced by older parents. We examine three sets of outcomes: the provision of informal care, the use of formal long-term care, and parental survival after the shock. Because informal care is typically unobserved, we use the gender composition of adult children as a proxy for its availability. A large literature documents 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 plausibly unrelated to parental health prior to the shock.

To identify how family responses to parental health shocks vary with the gender composition of children, we focus on strokes, severe health events whose timing is difficult to predict. This makes the occurrence of strokes an ideal setting to study family care provision following changes in health status since they are unlikely to be anticipated (Fadlon and Nielsen, 2019). We compare the evolution of several outcomes among parents who experience a stroke at the same age and have the same number of children but differ in the gender composition of those children. Before the stroke, parents of daughters and parents of sons exhibit similar levels and trends across a broad range of health measures. These include preexisting medical conditions, such as cancer, dementia, diabetes, and depression; health behavior indicators, including smoking, alcohol use, and obesity; and prior healthcare utilization. This similarity supports a causal interpretation of subsequent differences in parental survival and care use.

We begin by documenting a striking result: among parents who experience a stroke at the same age, the mortality consequences of the event vary with the gender of their children. Five years after the stroke, parents of daughters are 1.6 percentage points (p-value$<$0.01) more likely to be alive than parents of sons. To put this magnitude in perspective, first-ever stroke in Denmark carries a five-year cumulative mortality of approximately 37 percentage points (Kristensen et al., 2025). The observed daughter gradient thus corresponds to 4.3 percent of the excess mortality attributable to the stroke itself. The daughter gradient is statistically significant but small in the

year of the stroke and widens steadily over the subsequent five years. This dynamic pattern suggests that residual unobserved heterogeneity in stroke incidence or initial severity is unlikely to account for the result.

To further assess the role of stroke severity, we examine two types of healthcare use over which patients have little discretion: inpatient days and emergency room visits. Both respond primarily to emergencies and severe medical complications. If parents of daughters experienced systematically less severe strokes than parents of sons, we would expect them to use fewer of these services after the event. Reassuringly, we find no systematic difference in either margin, before or after the stroke.

We next examine informal care as a plausible channel through which the survival gradient emerges. First, following a parental stroke, daughters are 1.8 percentage points more likely than sons to live near their parent (3.6 percent of the pre-stroke mean). Geographic proximity is an important prerequisite for providing frequent, hands-on assistance and responding to day-to-day needs. Second, severe parental health shocks alter daughters' marriage trajectories more than those of sons: daughters are 2.0 percentage points less likely to be married after the stroke (3.7 percent of the pre-stroke mean). This result suggests that caregiving responsibilities may delay or disrupt major life-course transitions. Third, daughters have, on average, 0.047 fewer children following the shock (3.4 percent of the pre-stroke mean). The fertility response indicates that the consequences of parental illness extend beyond immediate caregiving decisions and may reshape daughters' own family formation. Taken together, these nonmonetary adjustments suggest that daughters place other aspects of their lives on hold to provide more of the day-to-day care that may contribute to improved parental survival.

Finally, we examine whether the availability of informal care crowds out formal care or instead facilitates its use. To study this margin of adjustment, we consider longer-term forms of physical care that often require sustained engagement after a stroke. We find that parents with daughters spend, on average, $20.64 more per year on physical rehabilitation than parents with sons who experience a stroke at the same age. 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, the results suggest that daughters do not simply replace formal care. Instead, they appear to help parents access and engage with it, implying that informal and formal care operate as complements in shaping post-stroke survival.

Our results contribute to three strands of literature. First, we build on work on the substitutability of formal and informal long-term care. Recent US quasi-experimental work exploits Medicaid rules and long-term care insurance to show that informal care crowds out demand for the formal product (Mommaerts, 2018, 2025; Coe et al., 2023; Barczyk and Kredler, 2019), and additional studies confirm that families provide most care even when formal care is heavily subsidized (Gruber and McGarry, 2023; Gørtz et al., 2023).[4] (Bolin et al., 2007) use cross-country European data to show that informal and formal home care act as substitutes, while informal care appear 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 formal care price; by extending the outcome from utilization to parental mortality; 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 cannot reflect financial substitution.

Second, we contribute to a literature on adult children's responses to parental health shocks. Fadlon and Nielsen (2019, 2021) show that Danish spouses and adult children adjust health investments and labor supply after severe parental events, and Jensen and Zhang (2026) document lasting labor market and mental health consequences of parental death in Denmark. Altındağ et al. (2026) likewise use Danish administrative data to show that parental dementia produces no measurable labor market or physical health responses in adult children but large mental health spillovers concentrated among daughters. Finally, Chen et al. (2024) document marriage and fertility responses to old-age health shocks in Taiwan. We complement this line of work by broadening the set of measured responses to include non-monetary margins, including residen-

[4]Earlier 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).

tial mobility.

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 (Kleven et al., 2018; Olivetti and Petrongolo, 2017). 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. Frimmel et al. (2025) finds that in Austria 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 (Løken et al., 2017; Van Houtven et al., 2013; Skira, 2015; Fahle and McGarry, 2022; Bergeot et al., 2024; Ramirez Lizardi et al., 2025). We build on the existing evidence by documenting an "old-parent penalty": 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 instititutional 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 B provides additional details on data sources, variable construction, and 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 (*Serviceloven*) 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 the municipality. 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—full-time municipal employment when the alternative would be institutionalization, and a care allowance for relatives of the terminally ill—though both cover only a small share of care episodes.

2.2 Registry Data and Key Definitions

Our analysis draws on comprehensive Danish administrative registers linking individual financial, demographic, and medical records. We describe each data source and key variable definitions in turn.

Medical Data and 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. Because our identification requires that shocked families do not anticipate the stroke, our health shock of interest is a parent's *first* stroke: we define the shock year as the calendar year in which the first parent in a family experiences their first stroke.[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 the stroke event itself, 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 as well, we complement the LPR with the Health Insurance Register (*Sygesikringsregistret*, SSSY), which records every service provided under the Danish public health insurance scheme, classified by provider specialty, together with the provider's gross fee for that service –the sum of the public subsidy and any patient co-payment,

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

[6] The 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 430–438, which correspond to the ICD-10 range used in the main sample period. Because our analysis sample is restricted to first strokes occurring in 1996 or later, the ICD-8 records are used only to determine whether a parent had a prior stroke and thus to enforce the first-stroke restriction; they do not directly affect the analysis.

and thus neither the patient's out-of-pocket cost nor the public sector's expenditure alone.[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.

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.[8] We flag cardiovascular disease, cancer, diabetes, obesity, and alcohol-related diagnoses, along with a set of mental health conditions: dementia, depression, bipolar disorder, anxiety, psychotic disorders, sleep disorders, and any mental health diagnosis. Because smoking status is not directly recorded in the register, we additionally flag chronic respiratory disease as a proxy for smoking. 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

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

[8] Restricting to diagnoses recorded before the calendar year of the stroke avoids mechanically classifying the stroke admission's own diagnosis code as a pre-existing condition: a stroke is itself coded under ICD-10 chapter I (circulatory system), the same chapter used to define our cardiovascular-disease indicator.

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. We winsorize each nominal variable annually at the 1st and 99th percentiles, deflate to 2015 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.[9]

Demographic information is obtained from the Civil Registration System (CPR), which records household composition and demographic characteristics on a quarterly basis. Educational attainment comes from the Ministry of Education (*Undervisningsministeriet*), which records each individual's highest completed level of education and professional qualifications.

All records are linked at the individual level using unique, anonymized CPR identifiers.[10]

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): all persons registered at an address flagged as a nursing home are treated as nursing-home residents in that year.[11]

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

[10] CPR identifiers are the Danish equivalent of U.S. Social Security numbers.

[11] The 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.

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.[12] We link both the stroked parent and their child to administrative income and demographic records from Statistics Denmark. The initial dataset contains 859,792 yearly observations covering 52,614 unique stroked parents. We apply the following restrictions at the parent (family) level.

First, since we are interested in the effect of a parent'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. We then apply three further restrictions so that each remaining observation is usable for estimation: 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 restrict to parents whose age at the stroke lies between 60 and 90, which drops 24.7percent of the remaining parents. We impose this last restriction for statistical power: our identification strategy non-parametrically estimates effects separately by age at stroke, so the sparsely populated cells at the tails of the age distribution would otherwise be estimated too imprecisely to be informative.

Our main analysis sample is an unbalanced panel of 575,716 observations covering 39,254 unique stroked parents (and equivalently 39,254 unique children). Appendix Table A.1 reports the number of observations dropped by each restriction.

[12] We 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 dropped from the sample, since our main income and demographic outcomes are only available at high quality from 1996 onward. The 1996 stroke cohort is excluded because outcomes must be observed in the year prior to the stroke for our difference in differences approach, which falls outside the available data window.

3 Identification and Estimation of Target Parameters

To identify heterogeneous responses to parental health shocks by family composition, we exploit variation in the gender mix of children among parents who experience a stroke at the same age. Similarly to Massenkoff and Rose (2024), we compare equally aged individuals, irrespective of calendar time, rather than individuals who experience a stroke in the same calendar year, irrespective of their age. We do so because age is a salient predictor of stroke incidence and severity, and parents who experience strokes later in life may differ systematically from parents who experience strokes earlier. The probability of stroke increases monotonically with age (Skajaa et al., 2021), and parents who experience strokes at unusually early ages may be negatively selected on unobservable characteristics that also influence stroke timing (for example, health habits, prevention behavior, or risk preferences).

Reassuringly, as widely documented in the medical literature (Newman-Toker et al., 2014; Arch et al., 2016) and in prior empirical work in this area (Fadlon and Nielsen, 2019, 2021; Frimmel et al., 2025), the exact timing of a stroke is largely unpredictable, at least among individuals of similar age and with comparable long-term health characteristics, such as obesity status or smoking behavior. This feature alleviates concerns about dynamic selection into treatment that could otherwise invalidate the standard assumptions required for difference-in-differences designs to identify interpretable causal parameters (Marx et al., 2024).

Potential Outcomes Model and Identifying Assumptions Formally, we observe an unbalanced panel of parents indexed by i over several years of age indexed by a. For each parent and age, we define a binary treatment indicator $D_{i,a} \in \{0, 1\}$ that records whether the parent has experienced a stroke by that age. In the rare cases in which a parent experiences more than one stroke, we focus on the first event. Treatment is therefore modeled as an absorbing state, so that $D_{i,s} = 1 \implies D_{i,w} = 1$ for all $w > s$. This structure allows us to summarize each individual's treatment path by the treatment cohort G_i, defined as the age at which the parent first experiences a stroke. We denote the set of all treatment cohorts by $\mathcal{G}$.

Our sample includes only parents who experience a stroke at some point in their lives. Thus, our empirical strategy does not rely on a never-treated group. This restriction reflects the concern that never-treated parents may differ from eventually treated parents in underlying health and other unobserved characteristics, making them a less suitable comparison group. Although we use standard difference-in-differences terminology, our design does not use variation in stroke timing across cohorts for identification. Instead, all comparisons are between parents who experience a stroke at the same age.

We examine one main source of heterogeneity: 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 female children and zero if the parent has only male children. Family composition is measured before the stroke and treated as time-invariant, as all parents in our sample experience a stroke after age 60.

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

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

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

Assumption 1 adapts the standard no anticipation condition (Abbring and Van Den Berg, 2003) to a setting without a never-treated group. Importantly, it does not rule out anticipatory responses to an impending stroke. Instead, it requires any such anticipatory response to be the same, in expectation, for families with only daughters and families with only sons. In our context, this assumption is plausible if the timing of severe parental health events is largely unpredictable and families do not differentially anticipate a stroke depending on the gender composition of

their children.

We also impose a common trends assumption adapted to our setting. Specifically, among parents who experience a stroke at a given age, we require counterfactual outcome trends under a delayed stroke to be mean-independent of family composition.

Assumption 2 (Common 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 scenario in which the parent experiences a stroke at age h. In economic terms, Assumption 2 requires $\Delta_{a-k,a}Y_i(h)$ to be the same, in expectation, for all-daughter and all-son families whose parent actually experiences a stroke at age g.

Target Parameters and Estimands For any $r \geq 0$ and any counterfactual stroke age $h > g + r$, define the family-specific average treatment effect on the treated as

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

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. We emphasize that the index h denotes a counterfactual late-stroke path for parents whose observed treatment cohort is $G_i = g$. It does not denote an observed comparison group, i.e., parents with $G_i = h$ are not used to identify the cohort-g effect. Our target parameter is the difference between the family-specific effects:

$$\theta_{g,r}(h) \equiv \mathrm{ATT}_{g,r}(1; h) - \mathrm{ATT}_{g,r}(0; h) \tag{2}$$

Thus, $\theta_{g,r}(h)$ measures the differential response to a parental stroke between families with only daughters and families with only sons. It measures heterogeneity in the average effects of strokes by family composition.

Under Assumptions 1 and 2, this parameter is identified as follows. For any relative period $r \geq 0$ and treatment cohort $g \in \mathcal{G}$:

1. Restrict the sample to units belonging to treatment cohort $G_i = g$.

2. Restrict the sample to units observed at both ages $a = g - 1$ and $a = g + r$.

3. In this subsample, estimate the saturated regression

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

The coefficient $\delta_{g,r}$ is equal to 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 $\theta_{g,r}(h)$ for any cohort $h > g + r$. In practice, we account for secular trends in stroke incidence and healthcare technology by including calendar-year-specific intercepts.[13] When examining outcomes related to the stroked individual's children, we additionally include flexible intercepts for the child's calendar age to account for life cycle trends in earnings, location choice, family formation, and fertility (Moffitt, 1984; Keane and Wolpin, 2010). Formally, we modify the regression in step 3 as follows:

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

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

[13] Over 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).

tively.

Fixing a relative period $r \geq 0$, and letting $\overline{a}$ be the last available age in the sample, we aggregate cohort-specific estimates using convex weights, as in Callaway and Sant'Anna (2021):

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

The weight $\mathbb{P}(G_i = g \mid G_i + r \leq \overline{a})$ is the relative share of cohort g among all cohorts for which relative period r is observed within the sample window, i.e., those for which $g + r \leq \overline{a}$. Larger cohorts therefore receive proportionally more weight, and 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) using a universal base period. Specifically, for each pre-treatment relative period $r < -1$ and treatment cohort $g \in \mathcal{G}$, we follow the same steps as for nonnegative relative periods and estimate the corresponding coefficient $\delta_{g,r}$. The period $r = -1$ is omitted because it serves as the reference period, i.e., the "universal base". Under Assumptions 1 and 2, the population coefficients satisfy $\delta_{g,r} = 0$ for all $r < -1$. These pre-trend coefficients therefore assess an implication of the two identifying assumptions jointly. The aggregation step is analogous to that used for nonnegative relative periods.

Statistical Inference We cluster standard errors at the parent level to account for within-family correlation of residuals 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, 0.5 years older than parents of a daughter (75.8 versus 75.3) and marginally less likely to have a partner (50.9 versus 51.8 p.p.), while the share of mothers is similar across the two groups (52.3 versus 51.5 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 (2.6 versus 2.4 p.p.) and of any prior mental health diagnosis (5.8 versus 5.7 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 (13.1 versus 12.3 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 individuals who experience a stroke at the same age. We therefore refine the raw comparisons reported in Table 1 using linear regression. Specifically, continuing

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

	Full Sample	Only Son	Only Daughter
Parent Demographics and Financial Characteristics			
Parent Age at Stroke, Years	75.5 (7.96)	75.8 (7.90)	75.3 (8.02)
1(Parent is Female)	0.519 (0.500)	0.523 (0.499)	0.515 (0.500)
1(Parent Pre-Stroke Partner)	0.513 (0.500)	0.509 (0.500)	0.518 (0.500)
Parent Total Income, USD	33,302 (41,195)	32,743 (26,120)	33,955 (53,664)
Stroke Type at Onset			
1(Ischemic Stroke)	0.442 (0.497)	0.440 (0.496)	0.445 (0.497)
1(Hemorrhagic Stroke)	0.141 (0.348)	0.142 (0.349)	0.140 (0.347)
1(Unspecified Stroke)	0.417 (0.493)	0.418 (0.493)	0.415 (0.493)
Parent Health Diagnoses at Stroke			
1(Parent Cardiovascular Diagnosis)	0.105 (0.307)	0.108 (0.311)	0.102 (0.302)
1(Parent Cancer Diagnosis)	0.127 (0.333)	0.123 (0.328)	0.131 (0.338)
1(Parent COPD Diagnosis)	0.0554 (0.229)	0.0576 (0.233)	0.0528 (0.224)
1(Parent Diabetes Diagnosis)	0.0580 (0.234)	0.0579 (0.234)	0.0580 (0.234)
1(Parent Obesity Diagnosis)	0.00344 (0.0585)	0.00322 (0.0566)	0.00370 (0.0607)
1(Parent Alcohol-Related Diagnosis)	0.0251 (0.157)	0.0245 (0.154)	0.0259 (0.159)
1(Parent Smoking, Alcohol, Diabetes, or Obesity)	0.132 (0.339)	0.133 (0.340)	0.131 (0.337)
1(Parent Dementia Diagnosis)	0.0250 (0.156)	0.0256 (0.158)	0.0243 (0.154)
1(Parent Any Mental Health Diagnosis)	0.0576 (0.233)	0.0578 (0.233)	0.0573 (0.232)
Parent Health-Care Utilization			
Nr. Parent Inpatient Hospital Admissions	0.519 (1.09)	0.526 (1.10)	0.510 (1.08)
Nr. Parent ER Visits	0.241 (0.629)	0.242 (0.632)	0.238 (0.625)
Nr. Parent Outpatient Hospital Visits	0.973 (1.40)	0.966 (1.40)	0.983 (1.41)
Parent Total Practice-Sector Fees, USD	598 (620)	591 (612)	605 (628)
Parent GP Fees, USD	287 (269)	288 (270)	286 (268)
Parent Physiotherapy Fees, USD	55.8 (338)	54.2 (330)	57.5 (347)
1(Parent Lives in Nursing Home)	0.0271 (0.162)	0.0284 (0.166)	0.0255 (0.158)
Child Demographics and Financial Characteristics			
Child Age at Parent's Stroke, Years	43.6 (8.32)	43.9 (8.21)	43.2 (8.45)
Child Wage Income, USD	45,021 (32,030)	49,803 (34,998)	39,452 (27,141)
Nr. Grandchildren	1.45 (1.12)	1.40 (1.15)	1.51 (1.08)
1(Child Married)	0.559 (0.497)	0.544 (0.498)	0.577 (0.494)
1(Child Lives within 10km of Parent)	0.492 (0.500)	0.505 (0.500)	0.477 (0.499)
1(Child Employed)	0.837 (0.369)	0.834 (0.372)	0.841 (0.366)
Observations	39,254	21,134	18,120

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.

to restrict the sample to the year preceding the stroke, we standardize each health-related outcome by its own standard deviation and regress it on the family composition indicator F_i, both without controls and after adding age-at-stroke and stroke-year intercepts. Figure 1 presents the resulting coefficients: after accounting for the age and calendar year at which the stroke occurs, parents in the two family composition groups remain virtually indistinguishable in terms of their observed pre-stroke health and healthcare use. Across all 16 outcomes and both specifications, the standardized coefficients remain very small in magnitude (between -0.024 and 0.026 standard deviations), are largely not statistically significant (only 1 of 16 is significant at the 5 percent level, with age-at-stroke and stroke-year intercepts), and point estimates exhibit no clear pattern favoring either family type (7 of 16 coefficients are positive and 9 are negative).

4 Results

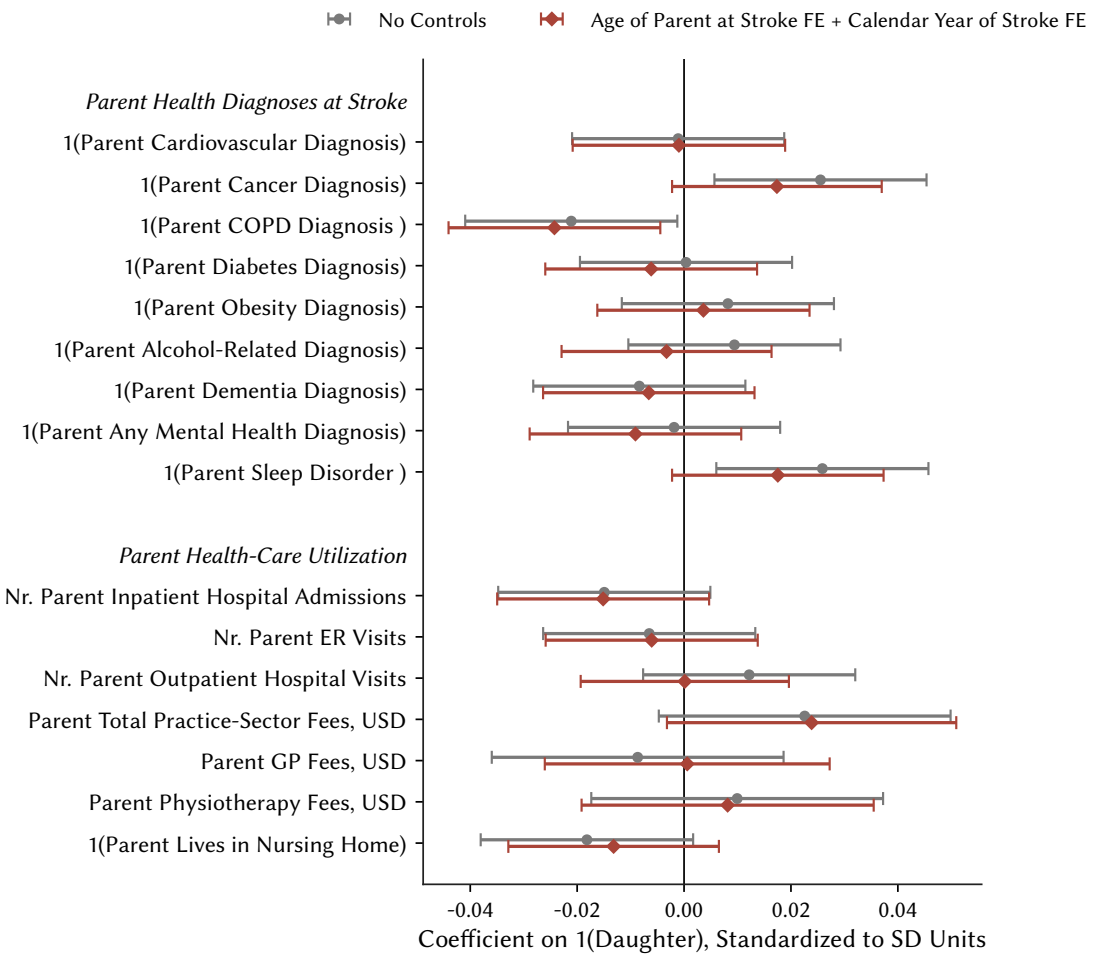
We report the main results below, following the identification strategy described in Section 3. Online Appendix A.2 complements this analysis by reporting difference-in-differences estimates of the average effect of a parental stroke on the level of each outcome.

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 1.6 percentage points more likely to be alive than parents in son-only families. To put this magnitude in perspective, a first-ever stroke in Denmark carries a five-year cumulative mortality of approximately 37 percent (Kristensen et al., 2025). The observed daughter gradient thus corresponds to 4.3 percent of the excess mortality attributable to the stroke itself. The gap is small in the year of the stroke and widens steadily thereafter.

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

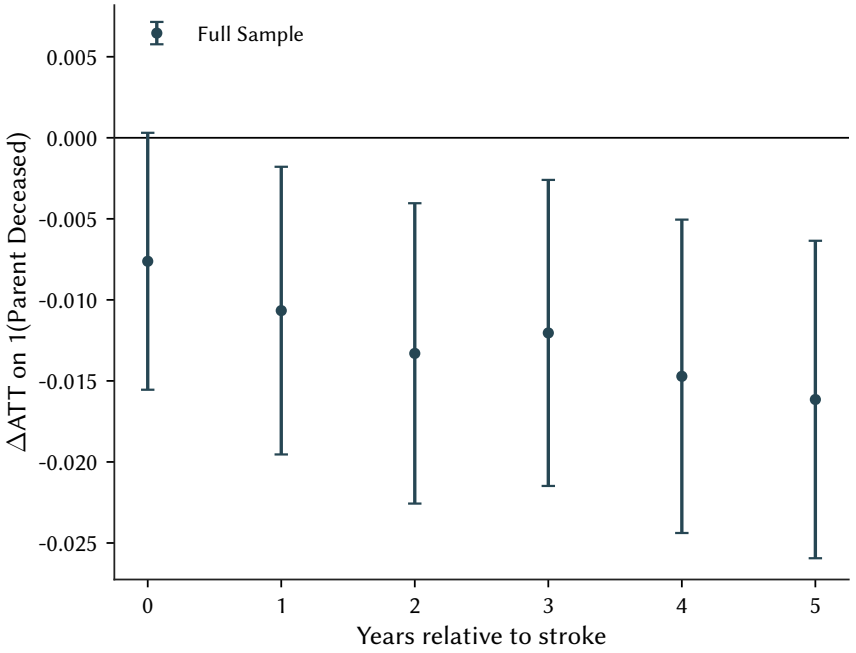
Figure 1: Balance of Pre-Stroke Health and Healthcare Utilization, by Family Composition



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.

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 whether the parent visits an emergency room. The absence of a difference in this relatively nondiscretionary margin of care provides fur-

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



Notes: The figure plots estimates of the target parameter in equation (2), using as the outcome 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.

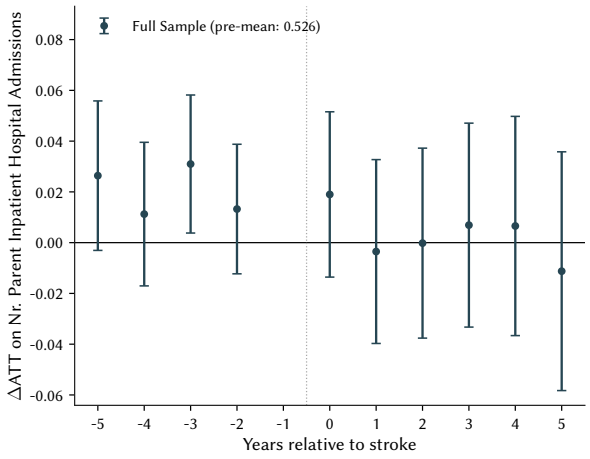
ther evidence that the post-stroke mortality gap is not driven by differences in the severity of the initial health shock.

4.2 The Daughter Gradient in the Supply of Informal Care

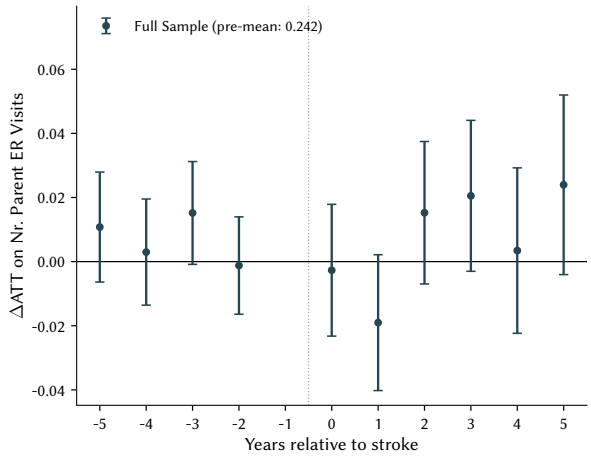
We next turn to the provision of informal care, which, by definition, is not directly observed in our data. We therefore examine several outcomes that plausibly proxy for the provision of care and assistance to a parent following a nonfatal health shock. Across these outcomes, we find consistent evidence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

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

(a) Hospital Admissions



(b) ER Visits



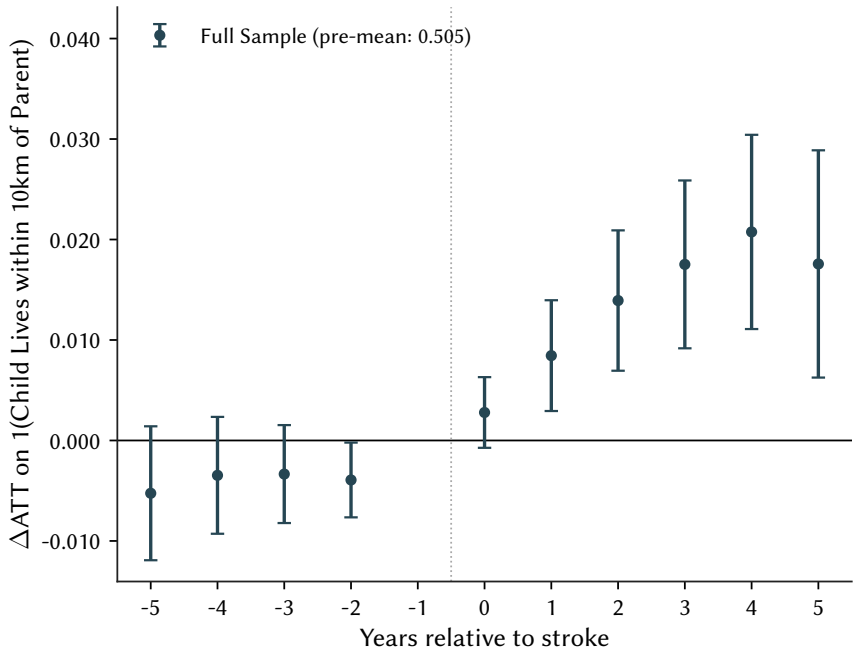
Notes: The figure plots estimates of the target parameter in equation (2). In panel (a), the outcome is the annual number of inpatient hospital admissions. In panel (b), the outcome is an indicator for whether the parent visits an emergency room at least once during the observation year. The estimator, 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.

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 postal-code centroid distance (Section 2). Daughters and sons show no differential trends in proximity before the stroke. After the stroke, daughters increase their proximity by approximately 1.8 percentage points more than sons, a differential of about 3.6percent relative to the pre-stroke baseline proximity rate of 49.2percent (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.

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

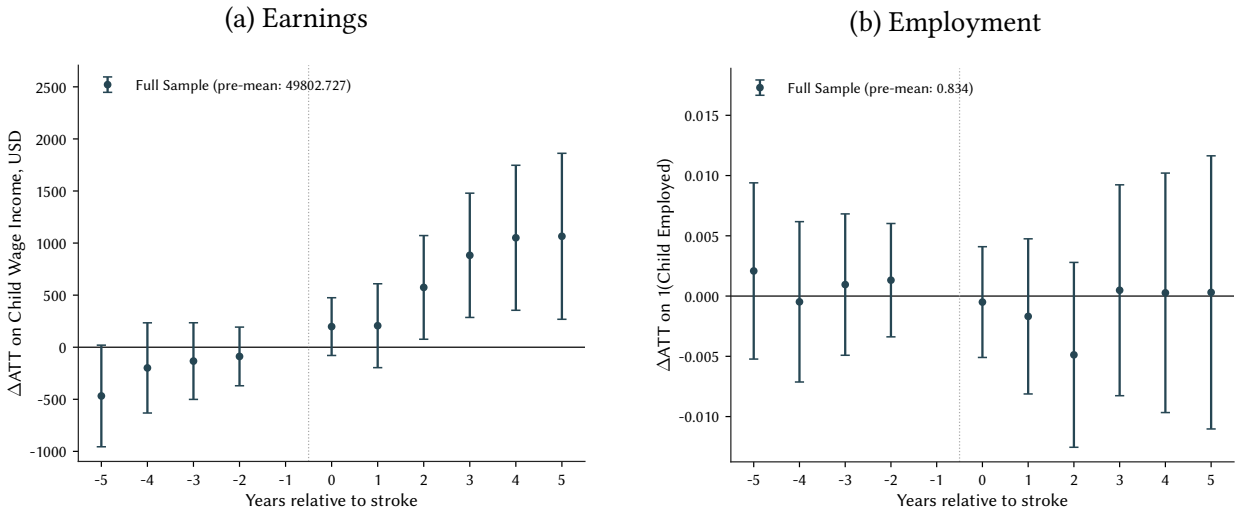


Notes: The figure plots estimates of the target parameter in equation (2), using as the outcome an indicator for whether the child resides within 10 kilometers of the parent. Distance is measured between the centroids of the child's and parent's postal codes of residence. The estimator, described in 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.

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 $1,065 per year—roughly 2.1percent of the pre-stroke earnings baseline in one-son families ($49,803)—and statistically significant at conventional levels (p-value$<$0.01). 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.

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



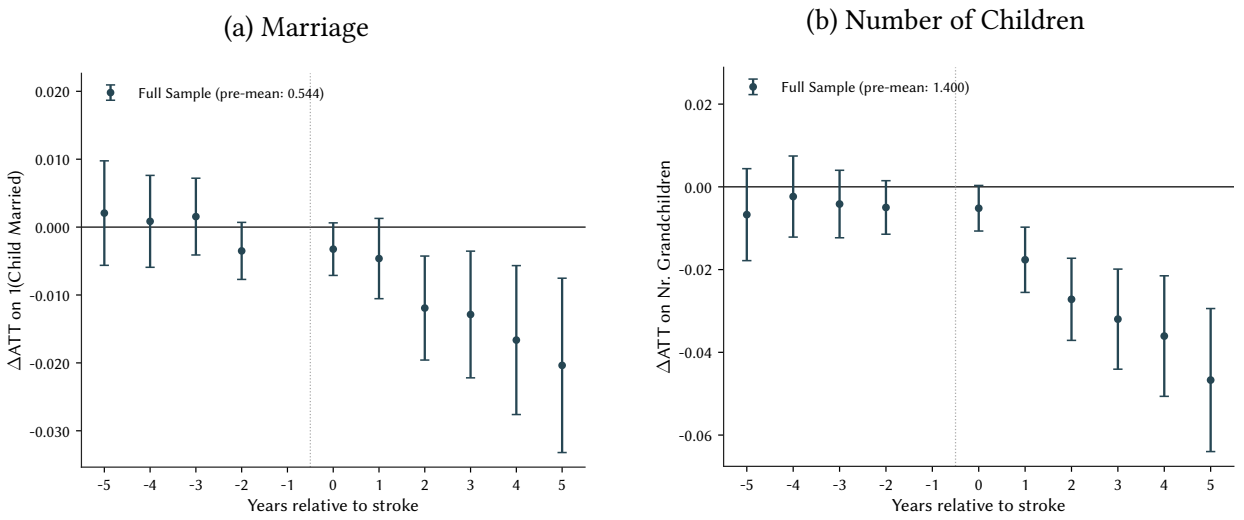
Notes: The figure plots estimates of the target parameter in equation (2). In panel (a), the outcome is the adult child's annual wage income, measured in U.S. dollars. In panel (b), it is an indicator for whether the adult child is employed. The estimator, described in Section 3, compares the evolution 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.

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 0.047 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 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 6: The Daughter-Son Gap in Marriage and Number of Children



Notes: The figure plots estimates of the target parameter in equation (2). 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.

4.3 Does Family Caregiving Complement Formal Care?

The evidence in Section 4.2 points to greater caregiving by daughters as a plausible source of the daughter mortality gradient. This raises a central question: does family caregiving replace formal care, or does it help parents obtain it? The two mechanisms yield opposite predictions. Substitution implies greater formal care use among parents of sons, whereas complementarity implies greater use among parents of daughters.

Our findings reported in Figure 7 are consistent with complementarity. Panel (a) shows that parents in one-daughter families accumulate somewhat more non-overnight hospital contacts following the stroke than parents in one-son families, indicating that daughters channel their parents into the medical system rather than shielding them from it. Panel (b) reports no differential response in nursing-home residence, despite an average post-stroke increase of roughly 10.1 percentage points (p-value$<$0.001) in this outcome (Online Appendix Figure A.5, Panel (b)). This rules out the substitution of informal caregiving for institutional care.

Panels (c) and (d) report higher total doctor fees, driven entirely by increased spending on physiotherapy and long-term care — the components of formal care most directly tied to post-stroke rehabilitation. Within five years of the stroke, only-daughter parents accumulate approximately $24.33 more in total doctor fees per year than only-son parents. Online Appendix Figure A.5 shows that the average 5-year level effect of a parental stroke on these fees is approximately $212.56, so daughter families utilize roughly 11.4percent more medical services than son families in the years following the shock.

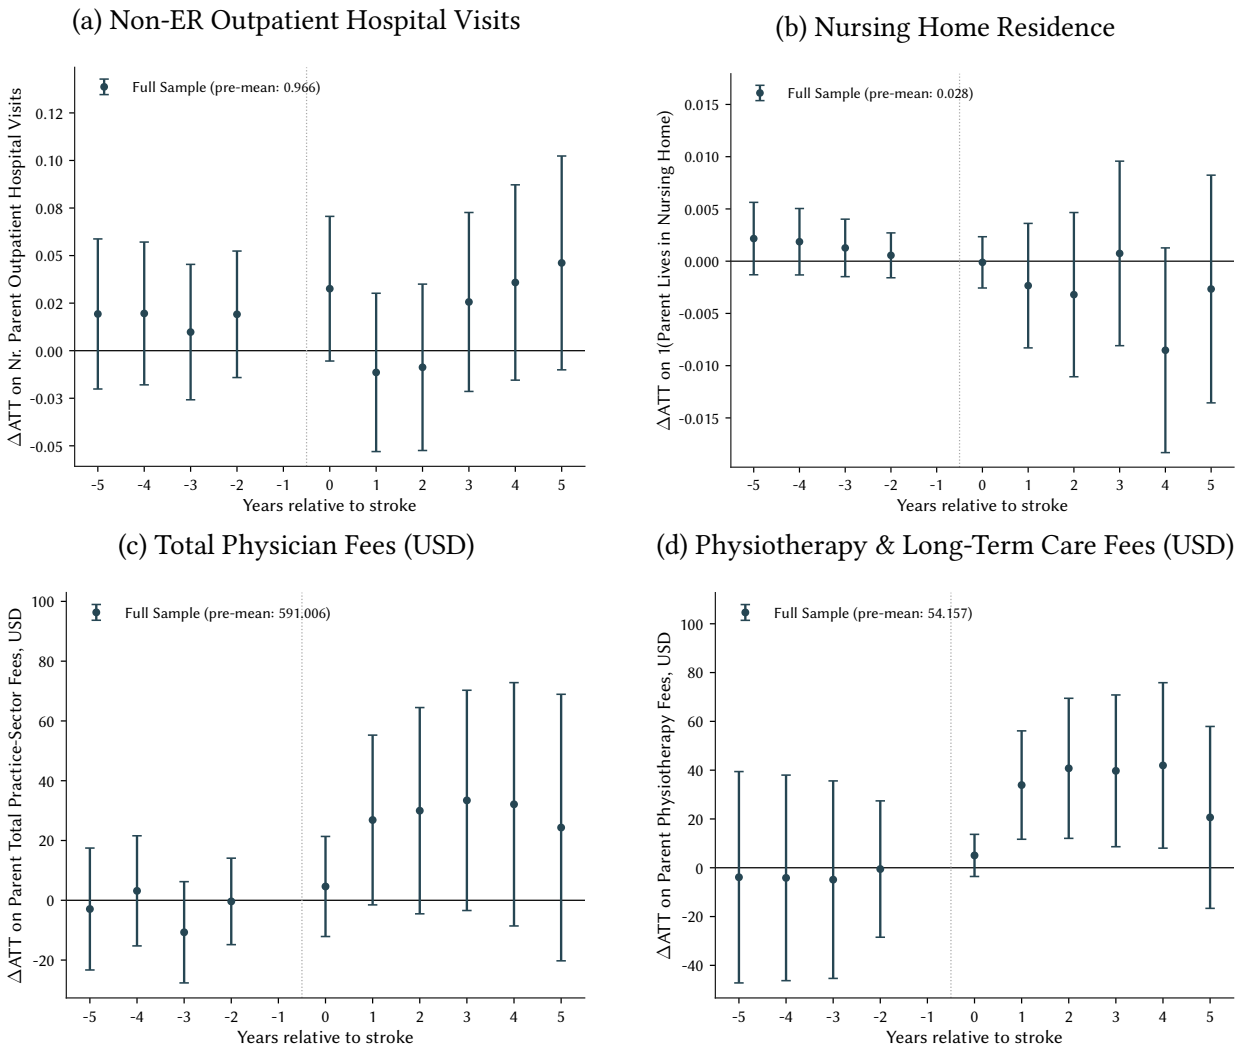
Taken together, these patterns indicate that daughters improve parental survival not by replacing the healthcare system but by mobilizing it.

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 1.6 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.

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



Notes: The figure plots estimates of the target parameter in equation (2). In panel (a), the outcome is an indicator for whether the adult child has any non-emergency outpatient contact during the observation year. 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.

Crucially, we show that parents of daughters spend more on physical rehabilitation, a type of formal care that requires sustained engagement after the acute phase of a stroke. By year five, their annual rehabilitation spending is $20.64 higher than that of otherwise comparable

parents of 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*, 2024.

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 Llena-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.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky, "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income," *The Quarterly Journal of Economics*, March 2024, *139* (2), 1321–1395.

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.

Houtven, Courtney Harold Van 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.

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 Jakob Egholt Søgaard, "Children and gender inequality: Evidence from Denmark," *NBER working paper*, 2018, (w24219).

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 Laugesen, 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.

Lizardi, Eduardo Ramirez, 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.

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.

Marx, Philip, Elie Tamer, and Xun Tang, "Parallel Trends and Dynamic Choices," *Journal of Political Economy Microeconomics*, February 2024, *2* (1), 129–171.

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.

Rocard, Eileen and Ana Llena-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 Lingering Economic Crisis," *Journal of Population Ageing*, 2020, *13* (1), 41–62.

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 Additional Empirical Results **35**

A.1 Descriptive Mortality Figures 35

A.2 Level Effects of the Stroke: Difference-in-Differences Estimates 36

B Data Appendix **43**

B.1 Sample Restrictions 43

B.2 ICD-10 Code Definitions 44

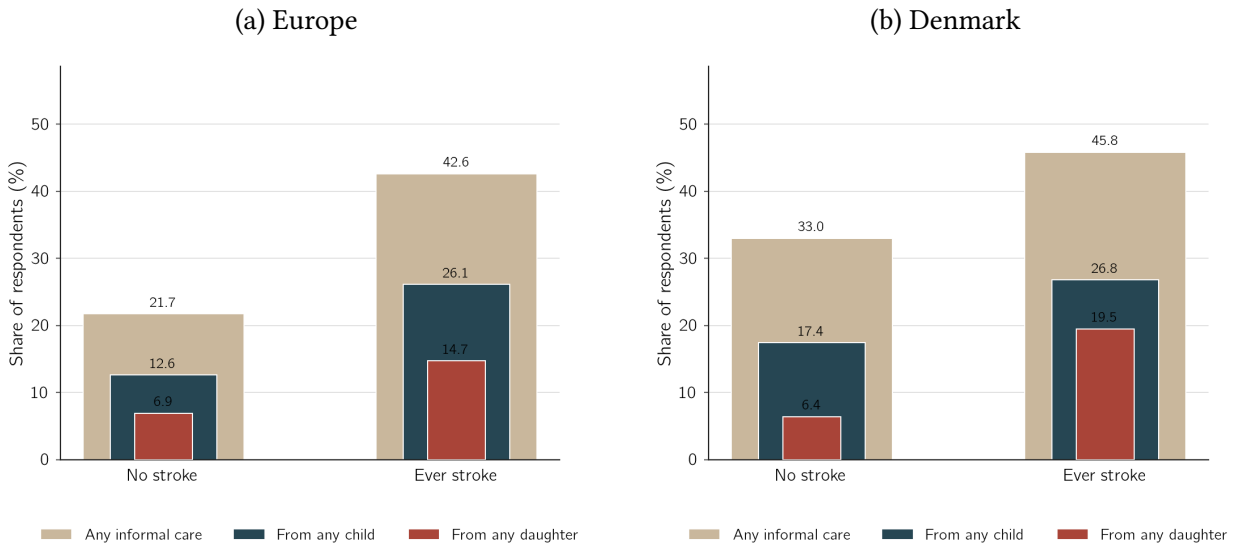
B.3 Household Data Details 45

B.4 Data Citations 45

A Additional Empirical Results

A.1 Descriptive Mortality Figures

Figure A.1: Informal Care in the Europe and Denmark



Notes: The figure decomposes informal-care take-up by whether the respondent ever reported a stroke, for Europe (Panel a) and Denmark (Panel b). Within each bar, the outermost segment is the share receiving *any* informal care, the middle segment the share receiving care from any adult child, and the innermost segment the share receiving care from any daughter. Data are from the Survey of Health, Ageing and Retirement in Europe (SHARE), waves 1–9 (2004–2022). Panel (a) pools all SHARE EU-member countries. The any-care and from-any-child rates are constructed from the harmonised main file across waves 1–9 (excluding wave 3, SHARELIFE); the from-any-daughter rate is taken from the wave-8 child roster. Sample sizes: Europe $N_{\text{no stroke}} = 136{,}885$, $N_{\text{ever stroke}} = 12{,}380$; Denmark $N_{\text{no stroke}} = 5{,}630$, $N_{\text{ever stroke}} = 555$.

A.2 Level Effects of the Stroke: Difference-in-Differences Estimates

The target parameter in our main analysis is the difference between average effects of a stroke experienced by parents of daughters versus parents of sons. To identify the corresponding intercepts, i.e., the average *level* effects of parental strokes, we consider a difference-in-differences design that compares parents experiencing strokes at different ages. Our design therefore relies on the comparison between currently treated and not-yet-treated units, similarly to, among others, Fadlon and Nielsen (2019), Golosov et al. (2024), and Massenkoff and Rose (2024).

This design relies on the following no anticipation assumption.

Assumption 3 (No Anticipation)**.** *For any age a and pair of treatment cohorts (g, h) such that $a < g < h$, $Y_{i,a}(g) = Y_{i,a}(h)$ with probability one.*

This assumption corresponds to the standard no anticipation condition (Abbring and Van Den Berg, 2003) adapted to a setting without a never-treated group. In our context, Assumption 3 rules out forward-looking behavior by parents or adult children who might otherwise change their behavior in anticipation of a parental health shock. In practice, this assumption is plausible if the timing of severe parental health events is largely unpredictable.

We also impose a common trends assumption with one important modification. Specifically, we restrict the comparison window between early- and late-treated cohorts by only using the first available κ not-yet-treated comparison cohorts, in a similar fashion to Fadlon and Nielsen (2019) and Massenkoff and Rose (2024). As we discuss in more detail below, these modifications ensure that the not-yet-treated group is comparable to the currently treated group in terms of life-cycle stage and any unobservable determinants of stroke timing that may correlate with health behavior and socioeconomic status.

Assumption 4 (Common Trends)**.** *For any age a, positive age difference k, and any pair of cohorts (g, h) such that $a - k < g \leq a < h$ and $h \leq t + \kappa$,*

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

Here, $G_i = g$ denotes the early-treated cohort and $G_i = h$ denotes the late-treated cohort. 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 along the late treatment path. In economic terms, Assumption 4 requires $\Delta_{a-k,a}Y_i(h)$ to be the same, in expectation, for the early- and late-treated cohorts. Age a is after treatment for the early cohort but before treatment for the late cohort, so the assumption requires the counterfactual evolution of outcomes for the early cohort to be captured by the observed pre-treatment evolution of the late cohort. This restriction is imposed only over horizons for which the late cohort's treatment is not too distant from age a ($h \leq a + \kappa$).

The common trends assumption in designs that rely on not-yet-treated units becomes more plausible when the time span between $G_i = g$ and $G_i = h$ is small. As discussed in the main text, parents who experience strokes later in life may differ systematically from parents who experience strokes earlier. We address this issue by restricting the time span between comparison cohorts to a moving window that uses the first available κ not-yet-treated cohorts as the control group.

Target Parameters and Estimands. Given the identifying assumptions, our target parameters are cohort-specific Average Treatment Effects on the Treated (ATTs), defined as

$$\mathrm{ATT}_{g,r} \equiv \mathbb{E}[Y_{i,g+r}(g) - Y_{i,g+r}(h) \mid G_i = g], \tag{A.1}$$

where $r \geq 0$ denotes relative time and $G_i = h$ denotes any not-yet-treated cohort within the comparison window, i.e., $g + r + 1 \leq t < h \leq g + r + \kappa$.

Under Assumptions 3 and 4, this parameter can be identified using the following steps. For any relative period $r \geq 0$ and treatment cohort $g \in \mathcal{G}$:

1. Restrict the sample to units belonging either to the treated cohort $G_i = g$ or to not-yet-treated cohorts within the comparison window $G_i \in \{g + r + 1, \ldots, g + r + \kappa\}$.

2. Restrict the sample to units observed in both ages $a = g - 1$ and $a = g + r$.

3. In this subsample, estimate the saturated regression

$$Y_{i,a} = \alpha_{g,r} + \beta_{g,r}\mathbb{I}[G_i = g] + \gamma_{g,r}\mathbb{I}\left[a = g + r\right] + \delta_{g,r}\mathbb{I}\left[G_i = g, a = g + r\right] + U_{i,a}. \quad \text{(A.2)}$$

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

$$\mathbb{E}[Y_{i,g+r} - Y_{i,g-1} \mid G_i = g] - \mathbb{E}[Y_{i,g+r} - Y_{i,g-1} \mid g + r + 1 \leq G_i \leq g + r + \kappa].$$

4. Under Assumptions 3 and 4, this difference-in-differences estimand identifies $\mathrm{ATT}_{g,r}$.

As in the main text, we modify step 3 by including calendar-year-specific intercepts as well as indicators for adult children's age:

$$Y_{i,a} = \alpha_{g,r} + \beta_{g,r}\mathbb{I}[G_i = g] + \gamma_{g,r}\mathbb{I}\left[a = g + r\right] + \delta_{g,r}\mathbb{I}\left[G_i = g, a = g + r\right] + \gamma_{g,r,t} + \eta_{g,r,z} + V_{i,a}. \quad \text{(A.3)}$$

Finally, we aggregate cohort-specific estimates using convex weights, as in the main analysis.

To assess the plausibility of the common trends assumption, we adapt the pre-trend checks proposed by Callaway and Sant'Anna (2021) in two ways. First, we use a universal base period. Second, we hold fixed the composition of the not-yet-treated comparison group across all relative time periods. Specifically, for each pre-treatment relative period $r < 0$ and treatment cohort $g \in \mathcal{G}$, we proceed as follows:

1. Restrict the sample to units belonging to cohorts $G_i = g$ or $G_i \in \{g + 1, \ldots, g + \kappa\}$.

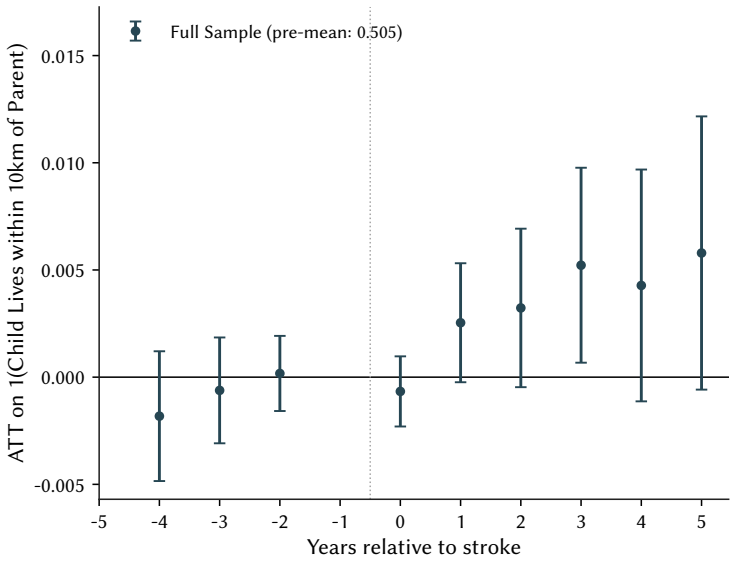
2. Restrict the sample to units observed in both periods $a = g - 1$ and $a = g + r$.

3. Estimate $\delta_{g,r}$ using specification (A.3).

The aggregation step is analogous to the case of nonnegative relative periods.

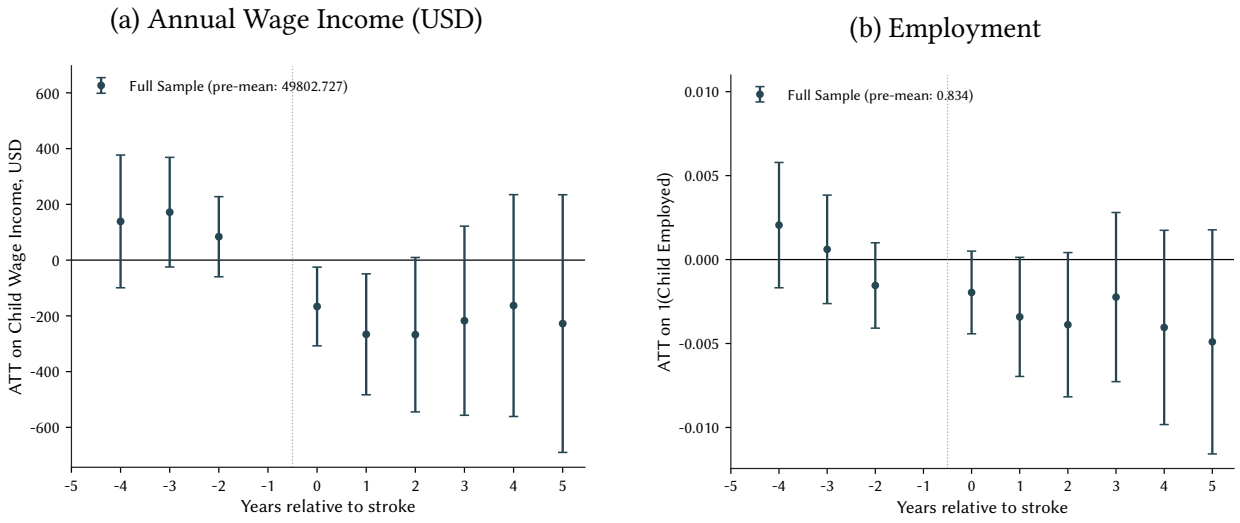
Figures A.2–A.4 report the estimated target parameters for the outcomes considered in the main analysis, grouped by residential proximity, parental healthcare utilization, adult child labor supply, and adult child family formation.

Figure A.2: Difference-in-Differences Estimates: Residential Proximity to Parent



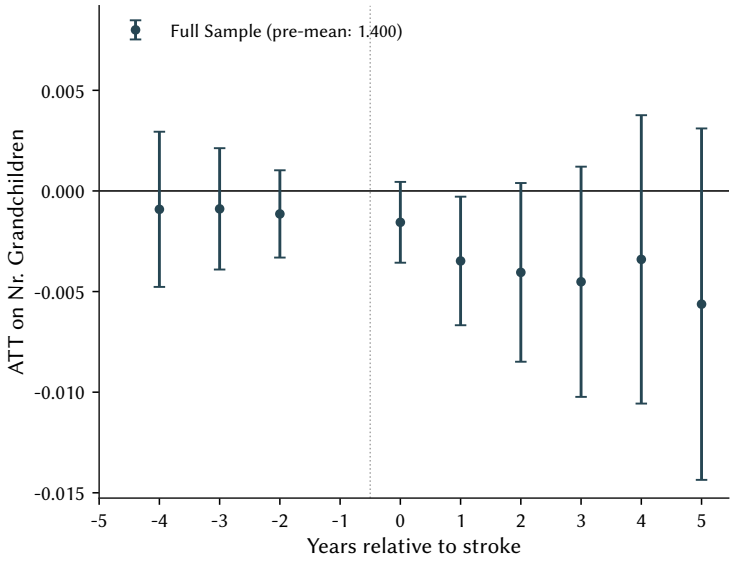
Notes: The figure plots estimates of the target parameter in equation (A.1), using as the outcome an indicator for whether the child resides within 10 kilometers of the parent. Distance is measured between the centroids of the child's and parent's postal codes of residence. The estimator, described in Section 3, compares families in which the parent experiences a stroke at a given age with families in which the parent experiences a stroke at a later age. The comparison group is restricted to the five nearest not-yet-treated cohorts ($\kappa = 5$), 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.

Figure A.3: Difference-in-Differences Estimates: Adult Child Labor Supply



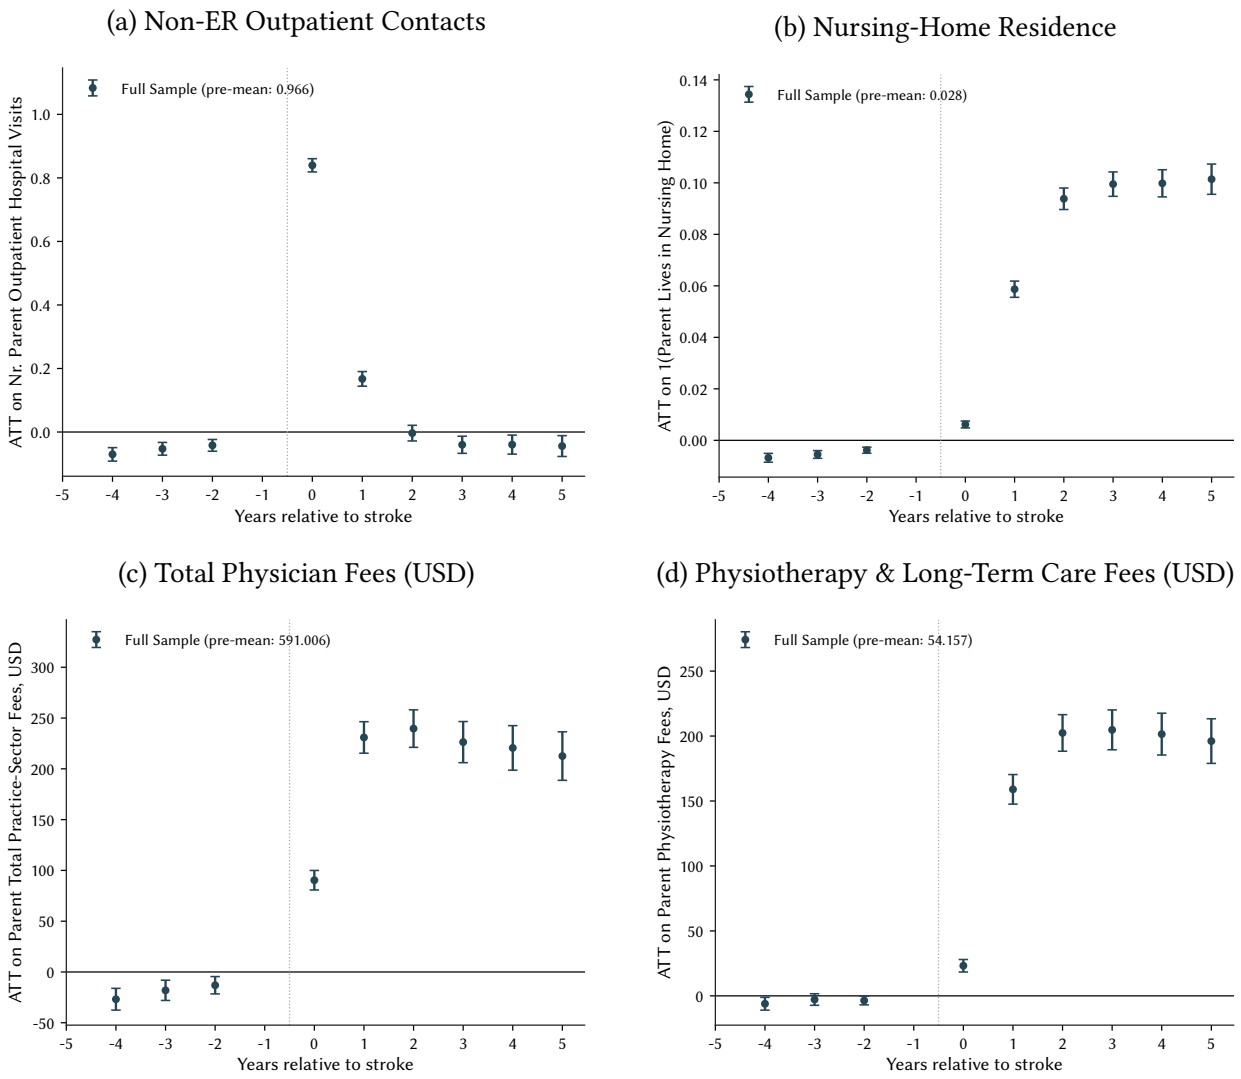
Notes: The figure plots estimates of the target parameter in equation (A.1). In panel (a), the outcome is the adult child's annual wage income, measured in U.S. dollars. In panel (b), it is an indicator for whether the adult child is employed. The estimator, described in Section 3, compares families in which the parent experiences a stroke at a given age with families in which the parent experiences a stroke at a later age. The comparison group is restricted to the five nearest not-yet-treated cohorts ($\kappa = 5$), 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.

Figure A.4: Difference-in-Differences Estimates: Adult Child Fertility



Notes: The figure plots estimates of the target parameter in equation (A.1), using as the outcome the adult child's number of children, equivalently the parent's number of grandchildren. The estimator, described in Section 3, compares families in which the parent experiences a stroke at a given age with families in which the parent experiences a stroke at a later age. The comparison group is restricted to the five nearest not-yet-treated cohorts ($\kappa = 5$), 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.

Figure A.5: Difference-in-Differences Estimates: Parental Healthcare Utilization



Notes: The figure plots estimates of the target parameter in equation (A.1). In panel (a), the outcome is an indicator for whether the adult child has any non-emergency outpatient contact during the observation year. 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 families in which the parent experiences a stroke at a given age with families in which the parent experiences a stroke at a later age. The comparison group is restricted to the five nearest not-yet-treated cohorts ($\kappa = 5$), while accounting for calendar time. Vertical bars denote 95 percent confidence intervals, and standard errors are clustered at the parent level.

B Data Appendix

B.1 Sample Restrictions

Table A.1: Sample Restrictions

Restriction	Observations	Unique Children	Unique Parents
Initial sample	859,792	52,614	54,285
First stroked parent per child	831,487	52,614	52,614
Parent observed at stroke	829,456	52,412	52,412
Parent and child observed in person-year	802,954	52,412	52,412
Contiguous spell containing event	790,182	52,113	52,113
Parent age at stroke 60–90	575,716	39,254	39,254
Final analysis sample	575,716	39,254	39,254

Notes: Table reports the number of families (shocked parents) dropped at 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 dropped for the given calendar year. Sample restrictions are described in Section 2.

B.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 associated with the listed code at any point before the calendar year of the first stroke.

Table A.2: ICD-10 Code Definitions

Condition	ICD-10 Code(s)
Stroke event classification	
Stroke (any subtype)	I60–I64
Hemorrhagic Stroke	I60–I62
Ischemic Stroke	I63
Unspecified Stroke	I64
Parent pre-existing conditions	
Any cardiovascular diagnosis	I (all codes)
Cancer	C00–C97, D00–D09
COPD / chronic respiratory disease	J40–J44
Diabetes	E10–E14
Obesity	E66
Alcohol-related diagnosis	F10; K70, K74; K86.0
Smoking, alcohol, diabetes, or obesity	(union of above)
Dementia	F00–F03; G30, G31
Depression	F32–F34
Bipolar disorder	F30, F31
Anxiety disorder	F40–F48
Psychotic disorder	F20–F29
Sleep disorder	F51; G47
Any mental health diagnosis	F (all codes)

Notes: 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.

B.3 Household Data Details

Household Definition 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 income terciles in the heterogeneity analysis. All other variables are measured at the individual level.

B.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)