cemfi

working paper
2612

The Burden of Informal Care: Does Looking After Aging Parents Weaken Social Ties?

Zhenghao Chen
Yarine Fawaz
Pedro Mira

August 2026

CENTRO DE ESTUDIOS
MONETARIOS Y FINANCIEROS
Casado del Alisal 5, 28014 Madrid, Spain
www.cemfi.es

CEMFI Working Paper No. 2612
August 2026

The Burden of Informal Care: Does Looking After Aging Parents Weaken Social Ties?

Abstract

Informal care provides the bulk of long-term care in many European countries, yet its social costs for caregivers remain poorly understood. Using panel data from the Survey of Health, Ageing, and Retirement in Europe (SHARE) and an instrumental-variable design based on parental poor health, we estimate the causal effect of intensive caregiving on the social connectedness of European women aged 50–70. We find that caregiving erodes the quality of social ties without changing their number: network satisfaction and a sizeadjusted social connectedness quality index both fall significantly. Effects are concentrated in countries with limited formal long-term care provision, where the first stage is strong. Complementary evidence shows that associational activities predict long-run survival, connecting the social erosion we identify to mortality risks documented in the epidemiological literature. Policies addressing caregiver burden should account for social isolation as a distinct cost.

JEL Codes: I12, I14, J14.

Keywords: Informal care, social connectedness, mortality, aging, long-term care.

Zhenghao Chen
UCL
zhenghao.chen.24@ucl.ac.uk

Yarine Fawaz
CEMFI and UAM
fawaz@cemfi.es

Pedro Mira
CEMFI
mira@cemfi.es

Acknowledgement

We thank Erik Ortiz Covarrubias for his assistance with the exploration of the literature. Jérôme Schoenmaeckers, Giorgia Menta, María Noel Pi Alperin and other participants at the LISER-SHARE Luxembourg SHARE Workshop. We also thank Jérôme Adda, Elena Bassoli, and other participants at the RF Berlin Workshop on Economics of Aging. We thank Luisa Licker, Peter Eibich, and other participants at the CINCH Workshop on Mental Health for helpful comments and suggestions. We thank Juan Vargas, Shanker Satyanath, and other participants at the 2026 Petralia Workshop. The authors declare that they have no financial conflicts of interest or other competing interests related to this research.

1 Introduction

Population ageing is reshaping care systems across Europe. As life expectancy rises and fertility declines, mounting pressure falls on long-term care (LTC) systems whose coverage remains limited and deeply unequal, particularly in Southern and Eastern Europe, where formal care provision is scarce and families shoulder most of the burden (Colombo et al., 2016; European Commission, 2022). In this context, families remain the backbone of elderly care provision. Informal caregiving – unpaid support provided by relatives, most often daughters – provides the primary source of long-term care in settings where formal services are limited. Middle-aged women aged 50–65 are the most exposed group, frequently providing intensive daily care while juggling employment and family responsibilities.

A substantial body of research documents the consequences of this burden for caregivers' health and labour market outcomes. Causal studies confirm that intensive caregiving worsens mental health and reduces labour force participation, working hours, and wages, with costs borne most acutely by women in low-formal-care environments (Crespo and Mira, 2014; Eibich, 2023; Schmitz and Westphal, 2015,1). Section 2 reviews this literature in detail.

One important dimension of caregiver well-being has received comparatively little attention: social connectedness. Social ties are foundational to health. They provide emotional support, buffer stress, and supply practical assistance; richer social networks predict lower depression, better coping with adversity, and reduced mortality risk (Fawaz and Mira, 2023; Holt-Lunstad and Smith, 2015; Thoits, 2011). Caregiving, however, is time-intensive and emotionally demanding. It crowds out social interaction, limits community involvement, and can foster isolation (Le and Ibuka, 2023; Zwar et al., 2020). Yet the available evidence is largely correlational. The causal impact of intensive caregiving on the structure and quality of social networks remains an open question.

This paper provides the first causal evidence on what we call the social erosion mechanism – the impact of intensive informal caregiving on social network quality, network satisfaction, and participation in associational activities. We use panel data from SHARE on women aged 50–70 with at least one surviving parent, and an instrumental variable strategy that uses an indicator for whether the respondent's parent is currently in poor health as a shock to caregiving demand (Crespo and Mira, 2014; Eibich, 2023). The sample is restricted to respondents whose parents are in non-poor health at their first observation in the panel, so the instrument captures the arrival of poor parental health after that baseline. All regressions control for the baseline outcome (measured at the daughter's first wave in the panel, before any subsequent parental health event can have occurred), so identification comes from post-baseline variation rather than levels: we therefore identify changes in social network quality rather than its level. In particular, the exclusion restriction only needs to hold conditional on the daughter's baseline social connectedness – the parental health state must affect the *change* in her social life relative to that baseline only through caregiving. We do not need the stronger assumption that parental health and social connectedness are unrelated in the cross-section: daughters of sicker parents may plausibly differ in their baseline social ties, but any such differences are absorbed by the baseline outcome control.

Credible identification requires a strong first stage, and that condition is met only in countries where formal long-term care is scarce and poor parental health actually triggers intensive caregiving. We follow the OECD (Colombo et al., 2016) and define this group as countries where above 22.5% of women aged 45–64 provide informal care for 20 or more hours per week: a threshold that corresponds to a natural gap in the cross-country distribution and selects Spain, Italy, Greece, Portugal, Poland, Croatia, Bulgaria, Malta, and Slovakia. We refer to this group throughout as *High-IC* (high informal care) countries. In this group, the Kleibergen-Paap first-stage F-statistic is approximately 36. In countries with more generous formal care provision, where the shock is absorbed by the LTC system, the first stage is near zero; we do not report IV estimates for those groups and treat pooled estimates as a reference. The estimated effect is a local average treatment effect (LATE) for women whose caregiving is triggered by poor parental health in low-formal-care environments, the group most exposed to the burden of care.

A distinctive feature of our empirical analysis is the use of SHARE's Social Network module (waves 4, 6, 8, and 9), one of the richest network data collections available in a large-scale household panel. Unlike surveys that proxy connectedness through role relationships such as marital status or co-residence, SHARE uses a name generator: respondents identify up to six confidants with whom they discussed important matters in the past year and report, for each, their geographic proximity, contact frequency, and emotional closeness. This design allows us to separate two distinct dimensions of social connectedness that are typically bundled together: network *quality* (how close, frequent, and intimate existing ties are) and network *quantity* (how many ties one has). Following Litwin and Stoeckel (2015), we build a quality index that expresses proximity, contact frequency, and emotional closeness as shares of social network links, so that quality is measured separately from size.

We find that intensive caregiving erodes the quality dimensions of social life while leaving the sheer number of ties broadly unchanged: network satisfaction and the social connectedness quality index both fall significantly ($p < 0.05$), while the extensive margin of associational participation has the expected sign but is not significant on average. Our results point to a thinning of existing ties rather than a contraction of the network. Heterogeneity within High-IC suggests that different subgroups might absorb the burden on different margins: still-working daughters seem to lose the quality of their close ties, while retired daughters might instead lose participation in associational activities. A direct test of the "family effect" concern – that the parental poor-health indicator could affect the daughter's social life through emotional distress rather than through caregiving – exploits the sharp cross-country variation in our first stage: reduced-form effects on social network outcomes are near zero in countries where the first stage collapses (Northern and Continental Europe), and negative and significant only in High-IC countries, where the first stage is strong. This is the pattern expected if caregiving, not grief, drives the social erosion. We return to magnitudes, this test, and the heterogeneity evidence in Section 5.

The main results survive a wide battery of sensitivity checks (Section 6). Weak-instrument-robust Anderson-Rubin inference rejects the null at 5% for both SN satisfaction and SC (quality), confirming the Wald conclusions. Restricting the sample to ages 50–64 (the OECD reference range) leaves the two quality outcomes larger in absolute value and significant, while the activity coefficient shrinks further towards zero. Dropping the co-residence control (potentially endogenous to caregiving) leaves all three point estimates essentially unchanged. Including parents in the network index – a mechanically inflated construction – yields a larger negative effect on SC (quality), confirming that our parent-exclusion rule produces conservative estimates. Decomposing the instrument shows it is effectively identified off maternal poor health: the maternal-only instrument gives $F \approx 37$ and significant estimates on all three outcomes, while the paternal-only version has $F \approx 4.5$ and cannot be used for inference, consistent with gender norms that assign filial care to daughters and with fathers typically having a surviving spouse who absorbs the care load. Adopting a supply-side country criterion (public LTC expenditure $< 1.5\%$ of GDP) instead of the OECD demand-side measure keeps activity participation significant but SN satisfaction and SC (quality) lose significance in that noisier grouping; a stricter $< 1\%$ cutoff recovers SC (quality). Varying the informal-care prevalence threshold between 15% and 30% leaves SN satisfaction significant across the plausible range and shows activity participation is sensitive to which mid-tier countries (Romania, Sweden, Greece, Poland, Bulgaria) are included.

As complementary evidence on the health relevance of the mechanism, we show in a dedicated section that participation in associational activities is associated with lower long-run mortality. Following the approach used in Fawaz and Mira (2023), both the social predictor and the full battery of health and socioeconomic controls are measured and controlled for at baseline, with mortality followed prospectively over the subsequent waves – by construction at least two years away from baseline. This temporal separation between predictors and outcome, together with the detailed baseline control set (frailty, chronic conditions, functional limitations, cognition, and depressive symptoms), moves the estimates closer to a causal reading than a purely contemporaneous association, though we do not claim strict causality. The exercise connects the social erosion we identify to the mortality risks documented in the social epidemiology literature.

The remainder of the paper is structured as follows. Section 2 reviews the related literature and describes the institutional context. Section 3 presents the data and measurement framework. Section 4 lays out the identification strategy. Section 5 reports the main causal estimates and the family-effect test. Section 6 presents a systematic sensitivity analysis. Section 7 presents the time-reallocation mechanism. Section 8 provides complementary survival evidence. Section 9 concludes.

2 Background and Related Literature

2.1 Institutional context of long-term care in Europe

Population ageing and rising dependency ratios have placed long-term care (LTC) at the center of social policy debates. While formal LTC systems exist across Europe, their coverage and generosity differ markedly. A pronounced North–South (and East) gradient persists. Northern countries (Denmark, Sweden, Finland) feature relatively generous, universal LTC systems – public spending often exceeds 2.5–3% of GDP, accompanied by extensive home care and institutional services (Colombo et al., 2016). In contrast, Southern countries (Italy, Spain, Greece) typically rely on minimal, means-tested public aid, with public expenditure around 0.6% of

GDP. These countries, alongside many Eastern European states (which share similarly limited formal LTC infrastructure and high reliance on informal caregiving), constitute one end of the spectrum. Continental states (Germany, France, Austria, the Netherlands, Belgium, Switzerland) fall in between, with mixed financing systems and intermediate levels of coverage (Crespo and Mira, 2014).

This institutional variation has direct implications for family caregiving: where formal LTC is scarce, families (especially women) become the primary caregivers. As Crespo and Mira (2014) emphasize, this interacts with gendered patterns of labour market attachment (stronger in Northern than in High-IC countries) to produce varying opportunity costs of caregiving. These cross-regional differences provide a natural backdrop for studying both the prevalence of informal care and its consequences for caregivers.

To illustrate this gradient in practice, Figure 1 reports the share of daughters aged 50–70 who provide any care to a parent, and the subset who provide daily (intense) care, using SHARE waves 4 to 9. It is apparent that cross-country variation is large and that cross-country patterns of provision of any care and intense care are different from one another. Focusing on the latter, the prevalence of intense (daily) care aligns with the institutional variation and OECD classification discussed above. Where formal LTC is scarce (eg Southern Europeand and some Eastern European countries), families fill the gap, often through daughters, whereas in countries with generous formal provision (eg Scandinavian countries) the reliance on daily informal care is rare. The pattern of prevalence of daily care also broadly aligns with the High-IC classification based on OECD reports.

2.2 Health effects of informal caregiving

A large body of evidence links caregiving to adverse health outcomes. Meta-analyses confirm associations with higher stress, depression, and poorer physical health (Bauer and Sousa-Poza, 2015; Pinquart and Sörensen, 2003; Roth et al., 2015; Vitaliano et al., 2003). More recent studies seek to establish causality. Matching approaches show increased depression and lower life satisfaction among caregivers compared to matched controls (Bom et al., 2019; Bom and Stöckel, 2021; Schmitz and Westphal, 2015,1). Instrumental variables (IV) strategies, exploiting parental health deterioration, parental death, or sibling structure, identify negative effects on mental health and subjective well-being (Coe and Van Houtven, 2009).

Results are heterogeneous. Care burdens weigh more heavily on women, who are both more likely to provide care and more affected by it (Abrahamsen and Grøtting, 2023; Bom et al., 2019; Colombo et al., 2016; de Zwart et al., 2017; Simard-Duplain, 2022; Wang et al., 2024). Intensity matters as well: high-frequency or long-duration care amplifies the adverse impacts on mental and physical health (Bom and Stöckel, 2021; Schmitz and Westphal, 2017; Stöckel and Bom, 2022; Van Houtven et al., 2013). Recent reviews note that while mental health effects are robustly negative, evidence on physical health is more mixed; some studies even suggest short-run improvements in self-assessed health due to comparison effects when spending time with sicker relatives (Bom et al., 2019; Di Novi et al., 2015; Eibich, 2023). Long-term studies, however, document persistent adverse consequences for psychological well-being (Schmitz and Westphal, 2015; Stöckel and Bom, 2022), and mechanisms such as social isolation and reduced

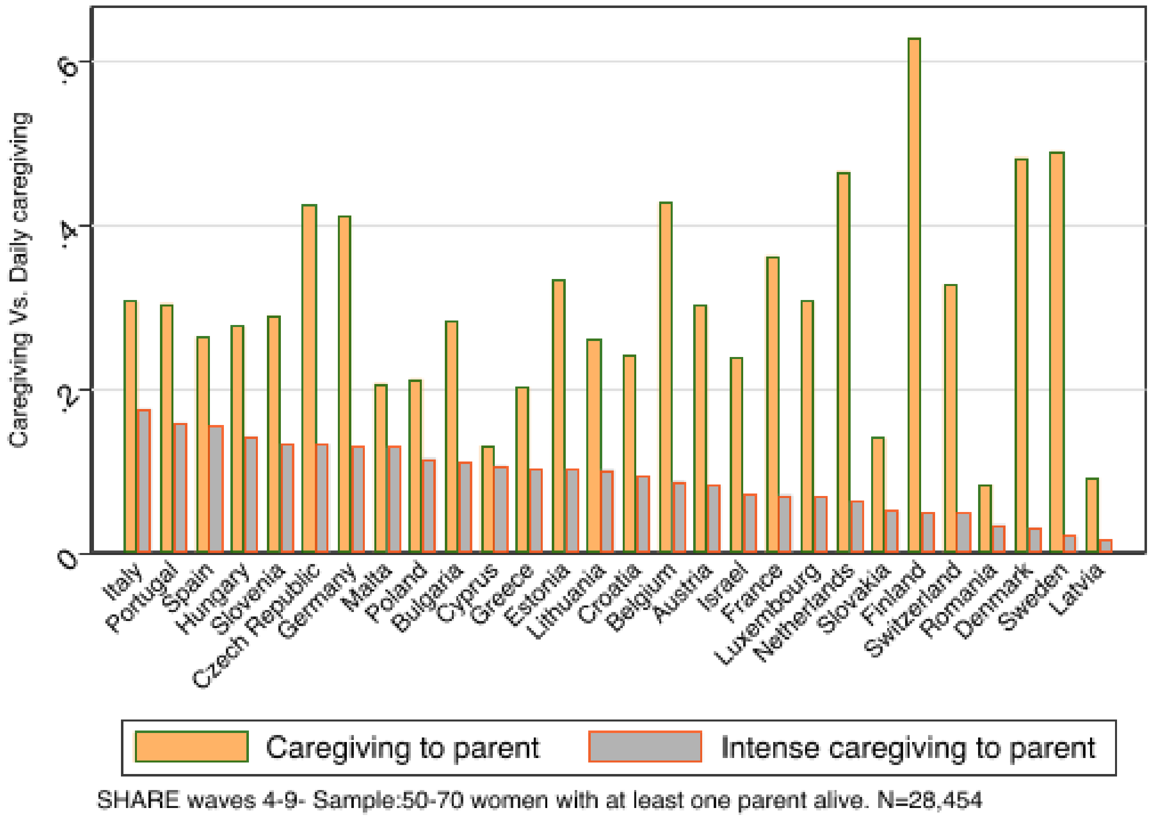

Fig. 1. Informal Caregiving across Countries

Notes: Share of daughters aged 50–70 providing any care (dark bars) and daily (intense) care (light bars) to a parent, by country. Source: SHARE waves 4–9.

leisure are increasingly recognized as potential channels (Le and Ibuka, 2023; Zwar et al., 2020).

2.3 Labor market effects

A second large literature examines how caregiving affects labor supply. Studies consistently show that caregiving can reduce participation, working hours, and wages, though magnitudes vary across contexts (Bolin et al., 2008; Casado-Marín et al., 2011; Crespo and Mira, 2014; Heitmueller, 2007; Kolodziej et al., 2018; Simard-Duplain, 2022). Instrumental variables analyses find that parental health shocks reduce daughters' employment probabilities, with effects concentrated in Southern Europe where public LTC is scarce (Crespo and Mira, 2014). Some evidence also points to increased early retirement or transitions out of full-time work (Eibich, 2023; Van Houtven et al., 2013). Long-term penalties include lower earnings trajectories and wage scarring (Schmitz and Westphal, 2017).

Heterogeneity again plays a central role. The costs are largest for women, especially those at the margin of labor force participation or in low-wage jobs (Crespo and Mira, 2014). Institutional differences also matter: the negative employment effects of caregiving are negligible in generous Nordic systems but substantial in Mediterranean settings (Crespo and Mira, 2014).

2.4 Social connectedness as a missing dimension

Despite extensive work on health and labor, one crucial margin of caregiver well-being, *social connectedness*, remains understudied. Networks of close ties are important buffers against stress and powerful predictors of morbidity and mortality (Fawaz and Mira, 2023; Holt-Lunstad and Smith, 2015; Thoits, 2011). Yet caregiving may constrain social interaction and foster isolation. Two recent studies point in this direction: Zwar et al. (2020) find increases in loneliness after transitioning into caregiving in German panel data (fixed-effects, no instrument), and Le and Ibuka (2023) use an FE-IV design on Japanese in-law caregivers – instrumenting caregiving with the in-laws' care needs – and detect reductions in perceived social support and time in community meetings. In both studies, social connectedness is captured by coarse indicators that essentially proxy loneliness – binary questions about whether the respondent has "someone to rely on" or reports feeling lonely – rather than by direct measures of network composition. Our paper contributes on two fronts. First, the SHARE Social Network module – a name generator listing up to six confidants along with their proximity, contact frequency, and emotional closeness, complemented by a satisfaction score and by measures of associational participation – yields substantially richer network outcomes and lets us separate network *quality* from network *quantity*, a distinction the binary indicators used in prior work cannot deliver. Second, our cross-country design provides a direct test of the exclusion restriction: because the parental health shock triggers intensive caregiving only in High-IC countries – and virtually not at all in Northern and Continental Europe – reduced-form effects on social outcomes in the low-first-stage groups isolate any residual *direct* effect of parental illness (grief, distress) on the daughter's social life. This family-effect falsification is unavailable in single-country designs.

Moving from correlation to causation faces a specific selection challenge: which daughters step in as intensive carers when a parent's health deteriorates is far from random. Geographic proximity to the parent, sibling availability, employment status, and family norms all shape who takes up care – and these same factors may independently relate to changes in social connectedness through channels that are difficult to sign a priori.

Figure 2 plots average SC quality against the share of daughters providing intense care, one observation per country. The relationship is positive: countries with higher caregiving prevalence tend to have higher average network quality. This reflects the fact that the same cultural factors driving high caregiving intensity – stronger family norms, geographic proximity between generations, closer intergenerational ties – also produce emotionally closer and more proximate personal networks. As a result, simple cross-country comparisons may point in the opposite direction from the causal effect: they would suggest caregiving is associated with better network quality, whereas the causal evidence (Section 5) will show it erodes it.

This cross-country confounding also illustrates why a within-person identification strategy is necessary. Conditioning on past social connectedness, as we do throughout, removes much of the pre-existing heterogeneity, but residual selection driven by time-varying determinants of caregiving uptake remains. A credible causal estimate requires a source of variation in caregiving that is orthogonal to these selection factors. We exploit parental health shocks for this purpose, and use SHARE's Social Network module (one of the richest egocentric network data collections available in a large-scale panel) to measure social connectedness along dimensions of both quality and quantity, a distinction our results show to be consequential: it is network quality, not the number of ties, that intensive caregiving erodes.

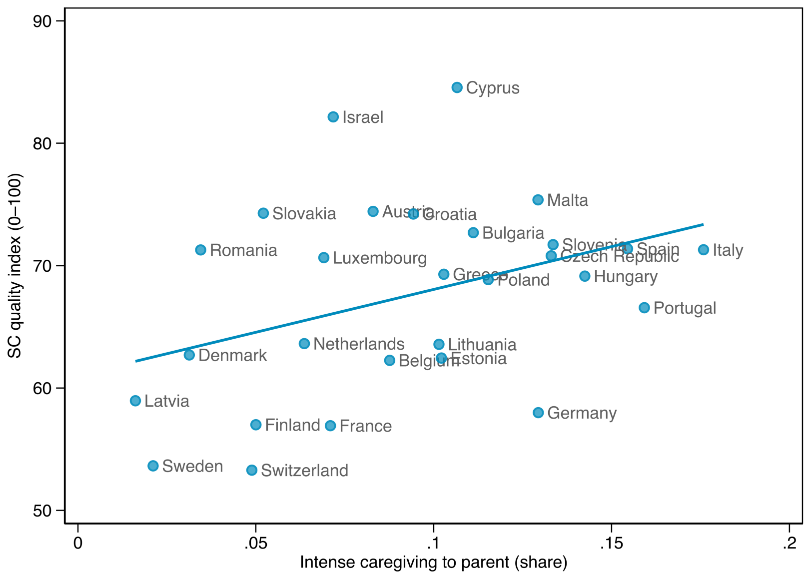

Fig. 2. Descriptive relationship between caregiving intensity and network quality

Notes: Each dot represents a country. The x-axis shows the share of daughters aged 50–70 providing intense care to a parent; the y-axis shows the average SC quality index – the average share of non-parental confidants who are geographically close, in frequent contact, and emotionally close (0–100; defined formally in Section 3) – among daughters in the same country and age group. The line is a linear fit. Source: SHARE waves 4, 6, 8, and 9.

3 Data and Measurement

3.1 Data Source: SHARE

We use the Survey of Health, Ageing and Retirement in Europe (SHARE), a multidisciplinary longitudinal survey of individuals aged 50 and over. SHARE covers health, socioeconomic status, employment, social support, and social networks, and is representative of the community-dwelling older population in 28 European countries and Israel. Data have been collected biennially since 2004. Our analysis uses waves 1 to 9 (2004–2023), excluding the retrospective life-history wave (wave 3), to construct the shock, and waves 4, 6, 8 and 9 for the analysis.

3.2 Sample selection

The analytical sample consists of daughters aged 50 to 70 with at least one living parent. We begin at age 50 because women in this age group are most likely to face parental care demands, while extending the sample up to age 70 increases statistical power and is appropriate given our focus on social connectedness rather than labor supply. Women older than 70 rarely have a surviving parent, while women younger than 50 are less likely to be observed in caregiving roles.

Including women aged 65 to 70 raises two distinct issues, which we treat separately. The first is that these women are largely retired and may differ in how their networks respond to a caregiving shock, given that peripheral ties tend to contract at the retirement transition (Kauppi et al., 2021; Settels and Böckerman, 2025). To address this heterogeneity concern, Sensitivity A verifies that our main results are robust to restricting the estimation sample to ages 50–64, the range used in the OECD intensive care prevalence measure. The second is that retirement itself, along with labor-force exit and co-residence with the parent, may be a direct response to the same parental health shock that triggers caregiving. To the extent that this is the case, effects on the daughter's network that operate through these adjustments are part of the causal effect of caregiving that our IV design identifies, not a source of confounding. We return to this point when interpreting the local average treatment effect in Section 5.2.

Starting from the full SHARE dataset, we apply standard cleaning procedures, restrict to women aged 50–70 with at least one surviving parent whose health was not already poor at baseline, and limit the analysis to the four waves in which the Social Network module was fielded (waves 4, 6, 8, and 9). Because our specification includes a baseline outcome control, we require each respondent to appear in at least two Social Network waves so that we can observe a pre-shock baseline value of the outcome and a subsequent post-baseline value. This requirement, together with missing values in some lagged covariates, yields a final estimation sample of 5,352 person-wave observations for the participation-in-any-activity regression (5,358 for network satisfaction, 5,366 for the social connectedness quality index), pooled across waves and corresponding to 4,493 unique women (1,315 person-waves and 1,118 women in the High-IC subsample). Attrition at each step is documented in Table A.1.

Two possible perspectives are available in SHARE: one is to use the daughters' sample, which records caregiving activities as reported by daughters themselves, together with basic information on parental health; the other is to use the parents' sample, which provides richer information on the health of the parents, including recent illnesses and hospitalizations. The latter would be valuable for instrument construction but does not contain systematic information on daughters' social networks, so given our focus we rely on the daughters' sample.

Table A.2 reports descriptive statistics for the final estimation sample of 5,352 person-wave observations, separately for non-caregivers (non-CG, N = 4,777) and daily caregivers (CG, N = 575). On average, daughters in our sample have about 12.3 years of education, three quarters are married, and they have roughly two children and live in households of 2.4 persons. About 1.6 grandchildren and one female sibling are reported (around 91 per cent have at least one sibling). Co-residence with a parent is relatively rare in the all countries, at about 9 percent on average, but this masks a sharp difference across groups: only 6 percent of non-caregivers live with a parent, compared to 37 percent of daily caregivers. The column labelled p reports the p-value from a two-sample t-test of equality of means; statistically significant differences between the two groups motivate the use of an instrumental variable strategy.

3.3 Informal caregiving and parental poor health

Informal caregiving in SHARE is recorded through questions on whether respondents provided help to their parents in the past twelve months. Help may include personal care such as dressing or bathing, practical household support such as shopping or repairs, or assistance with paperwork and administrative tasks. For parents living in the same household, caregiving is automatically coded as daily, since such situations typically involve sustained daily assistance. For parents living outside the household, respondents are asked about the frequency of help (daily, weekly, monthly, less often) and about the number of hours devoted.

Following established practice (Crespo and Mira, 2014; Eibich, 2023), we define intense caregiving as providing daily help to one's parent. This threshold captures the situations in which caregiving represents a substantial and regular time burden and is most likely to interfere with other aspects of life. The panel design of SHARE, combined with our baseline restriction on non-poor parental health, lets us use post-baseline variation in the parental poor-health indicator to identify take-up of intense caregiving.

Our instrument is constructed from the parent's overall health as reported by the respondent (the daughter) on a five-point Likert scale ranging from "excellent" to "poor." This is a child-reported measure of parental health, not the parent's own self-assessment: the parent is not a SHARE respondent, and the daughter answers this question in a dedicated module on her parents. We restrict the sample to daughters whose parents were not already in poor health at their first Social Network wave observation, and define the instrument as an indicator for whether a parent is in poor health at wave t (see Section 4 for the formal definition and further discussion). Given the baseline restriction, the indicator captures longitudinal variation in the parent's health status observed across the daughter's subsequent panel waves.

In our estimation sample, about 10 percent of daughters provide daily care to their parent (see Table A.2). The parental poor-health indicator equals 1 in roughly 27 percent of person-wave observations. These figures highlight that while intense caregiving is a relatively infrequent but important event, spells of parental poor health are much more common and potentially provide substantial exogenous variation in caregiving demand.

3.4 Social network outcomes

A distinctive feature of SHARE is the inclusion of a Social Network (SN) module, first introduced in wave 4 and repeated in waves 6, 8, and 9. This was the first large-scale cross-national survey of ageing to adopt a name-generator approach for eliciting personal networks. Rather than inferring ties from role relationships such as spouse, children, or co-resident relatives, SHARE asks respondents to name the individuals with whom they most often discussed important matters in the past twelve months. Up to six names can be provided. For each named confidant, detailed information is collected on relationship type, geographic proximity, contact frequency, and emotional closeness.

This design allows us to separate two dimensions of social connectedness that surveys relying on role relationships cannot distinguish: network *quality* (the relational density of existing ties) and network *quantity* (the number of ties). Our results show that intensive caregiving erodes quality but not quantity – pointing to a thinning of existing ties rather than a contraction of the network.

SC (size-adjusted) quality. Our primary structural outcome captures relational quality independently of network size. For each of three dimensions (geographic proximity, contact frequency, and emotional closeness), we compute the share of non-parental confidants with that property and average across the three shares.[1] Building on Litwin and Stoeckel (2015), this construction separates relational quality from network size: a respondent with two very close, nearby, frequently-seen friends scores identically to one with four equally close and proximate friends, whereas a raw count would rank them differently. The index ranges from 0 to 100.

SN size. We measure network quantity as the raw count of non-parental confidants (0–6). This is the dimension most commonly reported in the social network literature and serves as a natural benchmark: a caregiver who loses time for socialising might be expected to maintain fewer ties. As reported in Section 5, caregiving does not significantly reduce network size in our sample, which makes the quality result more striking – it is the relational quality of existing ties, not their number, that is eroded.

Figure 3 plots the distributions of both measures at Wave 4. SC quality is spread across the full range with a concentration at the upper end, reflecting that most respondents report at least some close, nearby, frequently-contacted ties. SN size is concentrated at 1–3 non-parental confidants, with few respondents naming more than four.

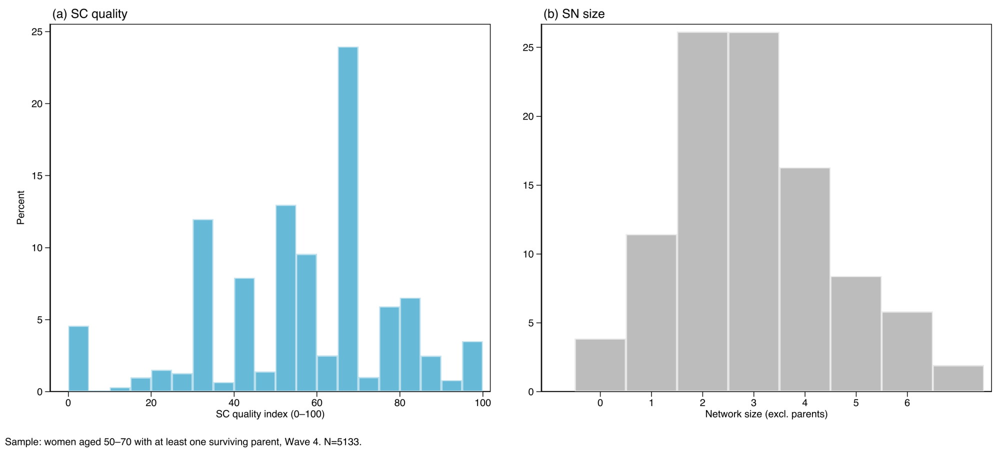

Fig. 3. Distributions of SC quality and SN size at baseline

Notes: Wave 4 distributions in the estimation sample (women aged 50–70 with at least one surviving parent). Left: SC quality index (0–100), computed as the average share of non-parental confidants who are geographically close, in frequent contact, and emotionally close, following Litwin and Stoeckel (2015). Right: SN size, the raw count of non-parental confidants (0–6). $N = 5{,}133$. Source: SHARE wave 4.

Network satisfaction and associational activities. We use two further outcomes that capture the evaluative and behavioural dimensions of social life. Network satisfaction is recorded on a 0–10 scale and asks respondents directly how satisfied they are with their social network. Participation in associational activities (sports clubs, religious groups, political organisations, charitable work, and similar) is a binary indicator of whether the respondent took part in any such activity in the past month. These two outcomes, together with SC (quality), are the three primary outcomes of the paper.

Parent exclusion. Throughout, we exclude parents from the network roster when construct-

[1] Formally, SC quality$_i = \frac{1}{3}\left(\frac{\text{prox}_i}{n_i} + \frac{\text{freq}_i}{n_i} + \frac{\text{close}_i}{n_i}\right) \times 100$, where n_i is the number of non-parental confidants and each numerator counts those with the given property. When $n_i = 0$ the index is set to zero.

ing quality and size measures. When a daughter takes up intense caregiving, she necessarily spends more time with her parent, which could mechanically increase the likelihood of naming them as a close confidant – and also increase measured proximity, frequency, and closeness. Including parents would therefore risk conflating the treatment with the outcome. By restricting the roster to non-parental ties, our estimates capture the effect of caregiving on the daughter's relationships with friends, other family members, and acquaintances. As a sensitivity check (Sensitivity H), we verify that results hold when parents are retained in the index at both t and $t-1$; the effect on SC (quality) is somewhat larger in absolute value. A plausible mechanism is that a severely ill parent may drop out of the confidant roster once their health deteriorates, or receive a lower rating on emotional closeness (the dimension we find most affected by caregiving, see Figure 5). Excluding parents from the baseline index avoids conflating this direct effect of the health shock on the network with the effect of caregiving itself.

Satisfaction with one's social network is a global question that cannot be re-computed excluding specific individuals. We use it as recorded. That it shows the same directional pattern as the parent-exclusive structural measures provides additional reassurance that results are not driven by parent inclusion or exclusion.

We complement these structural measures with participation in associational activities and time use. Associational activities include sports clubs, religious groups, political organisations, and charitable work; we record both whether the respondent participated in any activity and the number of distinct activities. Time allocation is drawn from the Time Expenditure module, introduced in wave 8, and captures minutes devoted to paid work, sleep, leisure, and caregiving.

4 Empirical Strategy

We estimate the causal effect of intense informal caregiving on social connectedness using a two-stage least squares (2SLS) approach, following Crespo and Mira (2014). The central challenge is that caregiving is endogenous: daughters who provide daily care may differ systematically from those who do not, in both observable and unobservable ways. Table A.2 illustrates these differences on observables. Daily caregivers are significantly more likely to co-reside with their parent (37 vs. 6 per cent, $p < 0.001$), less likely to be married (73 vs. 76 per cent, $p < 0.05$), and have fewer children (1.82 vs. 2.00, $p < 0.001$) and fewer siblings (81 vs. 92 per cent have at least one, $p < 0.001$). Education and household size do not differ significantly across groups. These systematic differences underscore the difficulty of identifying causal effects using simple comparisons, and motivate the use of an instrumental variable approach exploiting plausibly exogenous variation in caregiving induced by the onset of poor parental health.

A further feature of our specification is that we always include the outcome measured at the respondent's baseline observation, $Y_{i,0}$, as a control. Individuals differ widely in their levels of connectedness due to many factors that are difficult to observe or measure, such as personality traits, sociability, or long-standing family arrangements. While some of these differences can be partially captured by observables, the list of available covariates is inevitably limited. Controlling for the baseline outcome provides a way to account for this heterogeneity, so that identification comes from within-individual changes rather than from cross-sectional variation.

We use the baseline outcome $Y_{i,0}$, defined as the value of Y observed at the respondent's first wave in the panel (when her parent is by construction not in poor health), rather than the contemporaneous lag $Y_{i,t-1}$. The reason is that by the time we observe a respondent in a later wave t, the shock may already have occurred, and $Y_{i,t-1}$ may already reflect the effect of caregiving; conditioning on the lag would then absorb part of the treatment effect. The baseline outcome, in contrast, is predetermined with respect to any subsequent shock and does not have this bad-control problem.

The specification is nonetheless demanding: because we condition on the baseline outcome, only within-individual changes after baseline are available to be explained by caregiving. Most respondents in our sample contribute either their baseline observation alone or their baseline plus one to three later Social Network waves; we pool these observations without modelling the wave-gap explicitly. This choice improves power at the cost of not separating short-run from longer-run responses. Section A.8 in the Appendix shows that the main results are essentially unchanged under alternative specifications for the outcome control ($Y_{i,t-1}$ instead of $Y_{i,0}$, no SES lags, no lagged outcome at all, no co-residence control).

4.1 Instrumental variable

Our instrument is based on the parent's health status, assessed and reported by the respondent on a five-point scale from "excellent" to "poor." We restrict the sample to daughters whose parents (mother and father, when both are alive) are not in poor health at the daughter's first wave in the panel; those whose parent already reports poor health at baseline are excluded, since for them we cannot separate the effect of a subsequent poor-health episode from pre-existing poor health.

Given the baseline restriction, we define an indicator P_{it} for whether either parent is in poor health at wave t:

$$P_{it} = \begin{cases} 1 & \text{if Parental Health}_{i,t} = \text{Poor}, \\ 0 & \text{otherwise.} \end{cases} \qquad [1]$$

Because the sample has been selected on good parental health at baseline, this indicator captures longitudinal variation in the parent's health, observed across the daughter's subsequent panel waves. If a parent's reported health improves after an initial poor-health episode, we retain $P_{it} = 1$ in later waves. Previous poor-health episodes are informative about the parent's ongoing health, and the daughter's time allocation, household arrangements, and expectations regarding future care demand do not typically reset after a single wave of apparent recovery. In addition, a subsequent improvement in the parent's reported health may itself partly reflect the caregiving already being provided, so resetting the instrument on such improvements would risk contaminating it with the outcome.

4.2 Econometric specification

The first-stage equation is:

$$\text{Daily Care}_{it} = \pi \cdot P_{it} + \phi \cdot Y_{i,0} + X'_{i,t-1}\gamma + \alpha_t + u_{it}, \qquad [2]$$

where Daily Care_{it} is an indicator for whether daughter i provides daily care to any of her parents at time t, P_{it} is the parental poor-health indicator defined in Equation [1], $Y_{i,0}$ is the baseline outcome (value at the respondent's first observation in the panel), $X_{i,t-1}$ is a vector of predetermined covariates measured at $t-1$, and α_t are dummies for year of interview, cohort, and country-specific cohort trends.

The second-stage equation is:

$$Y_{it} = \beta \cdot \widehat{\text{Daily Care}}_{it} + \lambda \cdot Y_{i,0} + X'_{i,t-1}\delta + \alpha_t + \varepsilon_{it}, \qquad [3]$$

where Y_{it} denotes social connectedness and other social network outcomes. Controlling for $Y_{i,0}$ means that β identifies the effect of caregiving on the deviation of Y_{it} from its baseline level, and the exclusion restriction need hold only conditional on that baseline value. Since $Y_{i,0}$ is predetermined with respect to any subsequent shock, this control removes pre-existing individual heterogeneity without absorbing part of the treatment effect.

The vector of covariates $X_{i,t-1}$ includes age, education, marital status, number of children and grandchildren, household size, siblings and female siblings, an indicator for co-residence with a parent, quartiles of household income and net wealth, and an indicator for living in a rural area. All controls are measured in the prior wave to keep them predetermined with respect to the current caregiving decision. We deliberately restrict the set of controls to variables that are plausibly exogenous to caregiving given a lag: adding income, wealth, and rural residence at $t-1$ addresses potential SES confounding without introducing new endogeneity, and their inclusion leaves the first stage and the main second-stage estimates essentially unchanged (see Section 5). We deliberately *exclude* lagged employment from the baseline specification, even though it is available in SHARE. Employment is one of the main channels through which caregiving operates (Crespo and Mira, 2014; Van Houtven et al., 2013), and daughters anticipating parental care demands may reduce or exit work *before* the health shock is recorded; conditioning on lagged employment would therefore risk absorbing part of the treatment effect (a "bad control" in the sense of Angrist and Pischke, 2009). A robustness check that adds lagged employment as a control – currently reported for the time-use outcomes in Table A.7 (Appendix) – leaves those estimates essentially unchanged.

Our IV strategy relies on three standard assumptions for causal interpretation:

1. *Relevance.* A parent being in poor health must significantly increase the likelihood that daughters provide daily informal care. We show below that the first-stage coefficient on P_{it} is large and precisely estimated.

2. *Exclusion restriction.* A critical assumption of our IV strategy is that the parental poor-health indicator affects the daughter's social network outcomes at time t exclusively through its impact on her caregiving provision, conditional on the daughter's baseline

social connectedness. The key threat is the direct psychological impact of parental illness (sometimes called the "family effect" (Crespo and Mira, 2014)), whereby emotional distress, anxiety, or anticipatory grief could directly cause withdrawal from social life, independent of the time burden of caregiving.

Several features of our design mitigate this concern. First, because we condition on the baseline outcome $Y_{i,0}$, the exclusion restriction requires only that the health shock is independent of Y_{it} conditional on $Y_{i,0}$. The instrument therefore identifies the effect of caregiving on the deviation of social connectedness from its baseline level; any persistent individual-level component of Y that was already present at baseline is absorbed by the control. Second, our outcomes include behavioral measures such as participation in associational activities, which are more plausibly constrained by time than by emotional state. Third, we exclude parents from the network roster, so mechanical substitution cannot drive the results.

We provide a sharper empirical test by exploiting cross-country heterogeneity in Section 5.3. In countries where the parental health shock generates almost no caregiving response (Northern and Continental Europe, first-stage F below one), any reduced-form effect on social network outcomes must reflect the family effect rather than caregiving. By contrast, in High-IC countries where the first stage is strong ($F \approx 36$), reduced-form effects capture a mixture of both channels. We show that reduced-form effects on social outcomes are near zero in the North and Continental, and negative and significant only in High-IC countries: exactly the pattern expected if caregiving, not the family effect, is the mechanism.

3. *Monotonicity.* A parent being in poor health should weakly increase, but never decrease, the likelihood that daughters provide daily care. That is, there are no "defiers" who would stop providing care because their parent's health worsened.

A further potential concern is differential attrition: if daughters whose parent enters poor health are more likely to drop out of the survey before the next Social Network wave (for instance, because caregiving disrupts their engagement with the survey), our estimates could be biased by selective sample composition. We test this directly by regressing an indicator for dropping out of the estimation sample at the next SN wave on P_{it}, conditional on being in the sample at wave t. The coefficient is small (-0.031, SE $= 0.027$ in the pooled sample; -0.033, SE $= 0.051$ in High-IC countries) and statistically indistinguishable from zero, with a sign that goes against differential attrition. We conclude that selective attrition is not a concern for our identification strategy.[2]

Under these conditions, the coefficient β in equation [3] can be interpreted as the local average treatment effect (LATE) of providing daily care on social connectedness, for the subpopulation of daughters whose caregiving responds to poor parental health (Imbens and Angrist, 1994). We return to the implications of this LATE interpretation (including who the compliers are and why the IV estimate may differ from OLS) after presenting the main results.

[2] Attrition is measured among estimation-sample observations at waves 4, 6, and 8 (wave 9 is excluded as the last SN wave in our data, so no follow-up is possible). Controls and fixed effects match the main specification.

4.3 First Stage

Table 1 reports the first-stage estimates from equation [2], where the dependent variable is an indicator for providing daily care to one's parent. The key regressor is the parental health instrument.

Across columns (1)–(4) we report estimates from Equation [2] with progressively richer sets of controls. The sample is held fixed across all specifications to facilitate comparability. As described in Section 3, this consists of daughters aged 50–70 with a surviving parent not already in poor health at baseline, observed in at least two consecutive waves of the Social Network module.

In all specifications, the parental poor-health indicator P_{it} is a strong predictor of caregiving. In the pooled sample, the coefficient on P_{it} is around 0.06, with a tightly estimated standard error of 0.010–0.011, and the first-stage F-statistic ranges from 25.7 to 28.7 across control specifications. In the High-IC subsample (Table A.6), the coefficient is roughly 0.15 and the F-statistic ranges from 28.5 in the specification without controls to 35.8 in the main specification (column 4), a substantially stronger first stage that reflects the greater responsiveness of caregiving to parental health where formal long-term care is scarce. These results confirm the relevance of our instrument and the strength of identification.

Figure 4 illustrates the first stage graphically. The probability of providing daily care is markedly higher in waves when the parent is in poor health (as reported by the daughter) than in waves when the parent is not: the difference is roughly 6 percentage points in the pooled sample and 15 percentage points in High-IC countries. This shift is visible across the age range of 50 to 70 and confirms that parental poor health is a strong predictor of caregiving behaviour. Beyond age 65, the two lines converge and the estimates become noisier, reflecting the declining share of daughters with at least one surviving parent at those ages rather than a weakening of the instrument; Sensitivity A, which restricts the sample to ages 50–64, delivers quantitatively similar results. We focus on intense (daily) caregiving, since, consistent with Crespo and Mira (2014)'s time-allocation model, it is this form of care that represents a meaningful burden competing with other uses of time and that is most relevant for our hypotheses. For completeness, we show that a first stage also exists for non-intense caregiving, but the relationship is notably weaker.

As Crespo and Mira (2014) emphasize, instruments based on parental health are attractive precisely because they capture exogenous increases in the demand for care. The intuition is straightforward: when parents are in good health, few daughters provide daily help, but when health deteriorates, a sizeable group of compliers is induced into caregiving. The jump in caregiving probability at the onset of poor health in our figure corresponds to the mass of compliers. Eibich (2023) shows that many alternative instruments (siblings, parental survival) face validity or weak instrument concerns, particularly for men, and that parental health-based shocks are the most strongly predictive and theoretically coherent. In line with this reasoning, our estimated first-stage F-statistics are in the range of 30–36, alleviating concerns about weak instruments and supporting the use of parental health shocks to identify the causal impact of caregiving on social connectedness.

A notable feature of our first stage is its sharp heterogeneity across country groups. Ta-

Table 1: First Stage estimates- Dep. variable: Intense caregiving

	(1)	(2)	(3)	(4)
Parental health shock	0.058***	0.056***	0.056***	0.055***
	(0.011)	(0.011)	(0.011)	(0.010)
Basic controls	No	Yes	Yes	Yes
Siblings	No	No	Yes	Yes
Coresidence	No	No	No	Yes
F-Stat	26.9	25.9	25.7	28.7
R2	0.07	0.07	0.08	0.18
Observations	5366	5366	5366	5366

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort trend X country FE. Dependent variable is a dummy for daily caregiving to father or mother. Basic controls: age, married, education, number of children and grandchildren, size of household. All covariates are lagged. The sample is held fixed across all columns. F-statistics vary across specifications because the Kleibergen-Paap F-stat measures instrument strength conditional on the included controls; adding covariates changes the residual variation in the instrument after partialling out, which can raise or lower the F-stat even with an identical sample.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ble 5 (Panel A) reports first-stage estimates separately for Northern, Continental, Eastern, Southern, and High-IC countries. The Kleibergen-Paap F-statistic is close to 36 in the High-IC group (the set of countries defined by the OECD high-intensity criterion) while it is essentially zero in Northern Europe ($F \approx 0.0$) and in Continental Europe ($F \approx 0.2$). These numbers fall well below any reasonable threshold for credible 2SLS inference, so we do not report IV estimates for the Northern and Continental subsamples; doing so would produce unreliable, potentially heavily biased estimates.[3] This pattern is by design: parental health shocks generate strong behavioral responses only in settings where formal long-term care is limited and families shoulder the burden directly. The concentration of identifying variation in High-IC countries

[3] Even the conventional Staiger-Stock (1997) F > 10 rule of thumb has been shown to be insufficient in clustered settings (Lee et al., 2022; Montiel Olea and Pflueger, 2013a). Montiel Olea and Pflueger (2013a) propose an effective F-statistic (which adjusts for clustering and heteroskedasticity) with a threshold of 23.1 for 10% maximum size distortion at the 5% level. In a just-identified IV model (one instrument, one endogenous regressor), the Montiel Olea and Pflueger (2013a) effective F-statistic equals the Kleibergen-Paap rk Wald F, which already accounts for clustering. Our effective F-statistic in the High-IC main specification is approximately 36, comfortably above the 23.1 threshold, and Wald 95% confidence intervals for our two significant primaries (SC quality and SN satisfaction) exclude zero. Anderson-Rubin weak-instrument-robust tests confirm the Wald conclusions: for SC (quality), AR $F(1, 1121) = 5.57$ ($p = 0.019$); for SN satisfaction, AR $F(1, 1119) = 5.13$ ($p = 0.024$); for any activity, AR $F(1, 1120) = 2.10$ ($p = 0.147$).

has two important implications. First, it reinforces the relevance condition: our instrument is strongest precisely where it is most theoretically meaningful. Second, it provides the basis for the exclusion-restriction test described above: the near-zero first stage in the North and Continental implies that any reduced-form effects in those countries cannot be attributed to caregiving, and thus provides a falsification benchmark.

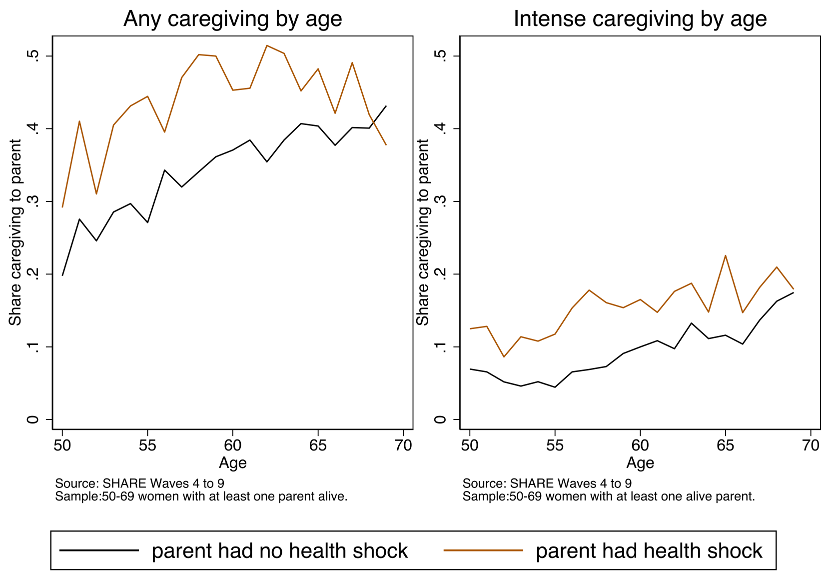

Fig. 4. Probability of Caregiving by age

Notes: Share of daughters providing any care (left panel) and daily care (right panel) to a parent, by age, separately for daughters whose parent experienced a health shock (orange line) and those whose did not (blue line). Source: SHARE waves 4–9.

5 Empirical Results

We present results in five steps. We first document large cross-country heterogeneity in the first stage, which motivates focusing on High-IC countries where informal care intensity is highest and the instrument is most powerful. We then present our main causal estimates for High-IC countries, followed by full-sample results for context. We then turn to a robustness exercise addressing the family effect channel. We close by examining which specific network dimensions are most sensitive to caregiving.

5.1 Cross-Country Heterogeneity in the First Stage

To explore cross-country heterogeneity, we restrict attention to those contexts where informal caregiving is most intensive. The choice of countries follows a transparent selection rule based on OECD data on informal care intensity (see Colombo et al. (2016) for the full report and data). Specifically, we use the OECD indicator "Share of informal carers aged 45–64 providing high-intensity informal care, defined as 20 hours or more per week, women". The European average of this measure is 17.9 percent, with most Northern and Continental countries lying well below this value (e.g. Austria 15.1, Germany 14.1, Netherlands 8.7, France 8.0, Belgium 9.5). We set the threshold strictly above 22.5 percent, the value recorded for Sweden – the highest-intensity country outside the South and East – thereby isolating the group of countries where high-intensity caregiving is systematically more prevalent. This rule yields Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia. These countries represent the institutional contexts in which informal caregiving remains particularly demanding, and where the social costs of caregiving are most likely to materialize.

Table 2: First Stage estimates- By clusters of countries- Dep. variable: Intense caregiving

	(1) All	(2) Southern	(3) Continental	(4) Northern	(5) Eastern	(6) High-IC
Parental health shock	0.055***	0.143***	0.007	0.001	0.098***	0.150***
	(0.010)	(0.030)	(0.015)	(0.014)	(0.021)	(0.025)
F-Stat	28.7	22.3	0.2	0.0	20.9	35.8
R2	0.18	0.24	0.18	0.12	0.18	0.21
Observations	5366	849	2285	588	1580	1317

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort X country FE. Dependent variable is a dummy for daily caregiving to father or mother. Basic controls: age, married, education, number of children and grandchildren, size of household. All covariates are lagged.
Southern countries include: Spain, Italy, Greece, Portugal, Cyprus. Continental countries include: Austria, Germany, the Netherlands, France, Switzerland, Belgium, Luxemburg. Northern countries include: Sweden, Denmark, Finland. Eastern countries include: Czech Republic, Poland, Hungary, Slovenia, Estonia, Croatia, Lithuania, Bulgaria, Latvia, Romania, and Slovakia. Countries in last column include: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia.
The 'All' column includes all countries in the main analysis sample, including Israel, but Israel is not assigned to any regional group and is therefore excluded from columns (2)–(5).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 applies our preferred first stage specification across clusters of countries. In the all-countries pooled sample (column 1), parental health shocks increase the probability of daily caregiving by about 6 percentage points, consistent with the results shown in Table 1. When broken down by country groups, however, large differences emerge. In Southern countries (column 2), the effect more than doubles to about 14 percentage points, while in Eastern countries (column 5) the effect is about 10 percentage points. By contrast, the estimates for Continental Europe (column 3) are close to zero (about 1 pp, not statistically significant), and those for Northern Europe (column 4) are essentially zero, reflecting the stronger reliance on formal care services in those regions.

The final column combines High-IC countries (those exceeding the OECD threshold for high-intensity informal care). In this group, a parental health shock raises the probability of daily caregiving by about 15 percentage points. This is the strongest first stage we obtain for a cluster of countries, with an F-statistic close to 36, and it lines up with the ex-ante criterion described earlier. These results show that the identifying variation in our IV strategy comes from countries where caregiving responsibilities are heavy, consistent with institutional differences in long-term care provision.

Table A.6 in the Appendix reports the first stage estimates restricted to the set of countries identified by our OECD-based selection rule, using the same progressive-controls layout as Table 1. Relative to the all-countries sample, the coefficient on the parental health shock is much larger – around 0.15 across specifications, roughly three times the effect in the all countries (about 0.06). The associated F-statistics (around 30–36) remain well above conventional thresholds, and the R^2 values are also higher than in the pooled sample. This pattern aligns with the institutional context: in settings where formal long-term care provision is more limited, parental health shocks translate more directly into intense caregiving by daughters.

5.2 Main Results: High-IC Countries

Table 3: Reduced form estimates, High-IC countries – SC (quality), SN Satisfaction, Activities, and SN Size

	(1) SC (quality)	(2) SN Satisfaction	(3) Any activity	(4) SN size
Parental health shock	-3.075**	-0.168**	-0.036	-0.064
	(1.303)	(0.074)	(0.025)	(0.085)
Mean	71.79	9.13	0.21	2.54
R2	0.177	0.107	0.171	0.280
N	1,317	1,315	1,315	1,317

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort X country FE. Controls: age, married, education, number of children and grandchildren, size of household, siblings, female siblings, coresidence, income quartile, wealth quartile, rural residence, and the baseline outcome $Y_{i,0}$. All other covariates are measured at $t-1$.
Countries selected include: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 reports reduced form estimates for High-IC countries across four outcomes. The two quality dimensions of social connectedness are both negative and statistically significant: a parental poor-health indicator is associated with a drop of 3.08 points on SC (quality) (column 1, $p < 0.05$) and a decline of 0.17 points on network satisfaction (column 2, $p < 0.05$). The reduced-form effect on the probability of participating in any associational activity (column 3) is -0.036 and does not reach conventional significance. Given the mean satisfaction of 9.1, the SC quality and satisfaction effects correspond to economically meaningful magnitudes. Column 4 shows the reduced-form effect on SN size: the coefficient is indistinguishable from zero. The pattern confirms that the identifying variation is concentrated in High-IC countries and that the shock operates on quality rather than on quantity or on the extensive margin of activity participation.

Table 4 contrasts OLS and 2SLS estimates for the three primary outcomes in High-IC countries, together with SN size as a quantity benchmark (column 4). Panel A (OLS) shows that conditional on past outcomes, the associations are close to zero for SC (quality) and SN

Table 4: OLS vs. 2SLS estimates, High-IC countries – SC (quality), SN Satisfaction, Activities, and SN Size

	Panel A: OLS			
	(1) SC (quality)	(2) SN Satisfaction	(3) Any activity	(4) SN size
Intense caregiving to parent	-0.391	0.0744	0.0448	0.114
	(1.826)	(0.0906)	(0.0325)	(0.108)
	Panel B: 2SLS			
Intense caregiving to parent	-20.49**	-1.112**	-0.244	-0.428
	(9.362)	(0.526)	(0.175)	(0.569)
F-test (IV)	35.54	36.08	34.46	35.60
Hausman p-value	0.02	0.01	0.08	0.32
Observations	1317	1315	1315	1317

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort X country FE. Controls: age, married, education, number of children and grandchildren, size of household, siblings, female siblings, coresidence, income quartile, wealth quartile, rural residence, and the baseline outcome $Y_{i,0}$. All other covariates are measured at $t-1$.
High-IC countries: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia.
The F-test is the Kleibergen-Paap rk F-statistic. The Hausman p-value tests endogeneity of intensive caregiving.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Satisfaction, and slightly positive for activities – consistent with positive selection into caregiving that masks the underlying causal effect.

Panel B presents the 2SLS results. Intensive caregiving reduces SC (quality) by 20.5 percentage points ($p < 0.05$) and network satisfaction by 1.11 points on a 0–10 scale ($p < 0.05$). The point estimate for participation in any associational activity is -0.24, of the expected sign but not statistically significant at conventional levels ($p = 0.16$); we return to this outcome below and in the heterogeneity analysis, where it becomes clearer that the activities response is driven by specific subgroups. The first stage is strong across all columns (F-statistics between 34 and 36) and Hausman tests confirm endogeneity for SC (quality) and SN satisfaction ($p < 0.05$). Column 4 shows the benchmark quantity result: the 2SLS coefficient for SN size is -0.43 ($p = 0.45$), far from significance, with a Hausman p-value of 0.32 indicating no detectable causal effect. Intensive caregiving erodes the quality of social ties but not their number – pointing to a thinning of existing relationships rather than a contraction of the network.

The sign reversal from close-to-zero OLS to sharply negative 2SLS for SC (quality) reflects positive selection: women who become intense caregivers in the absence of a parental health shock tend to have better-quality networks, generating upward bias in OLS. The IV corrects for this: among compliers induced into caregiving by the shock, network quality falls markedly.[4]

[4] Appendix Tables A.4 and A.5 report reduced-form and 2SLS estimates for the full pooled sample. Network satisfaction reaches significance in the pooled 2SLS (-1.68, $p < 0.05$), while SC (quality) and any activity have the expected sign but larger standard errors than in the High-IC-only sample: pooling over contexts where the

Because caregiving may itself trigger downstream adjustments in labor-force participation, retirement decisions, and co-residence with the parent, the LATE we identify is the total causal effect of intensive caregiving on the network, inclusive of any pathways operating through these adjustments; it is not intended to isolate a pure time-crowding effect conditional on labor-market and living-arrangement status.

5.3 Testing the Family Effect Channel

A key concern for our IV strategy is the so-called "family effect": the parental health shock may affect daughters' social connectedness directly through emotional distress (grief, anxiety, or anticipatory loss) rather than solely through the time and resource demands of caregiving. If so, the exclusion restriction would be violated and our 2SLS estimates would conflate two distinct channels.

We test this directly by exploiting the large cross-country variation in our first stage. In countries where the parental health shock generates almost no caregiving response (primarily Northern and Continental Europe), any reduced-form effect of the shock on daughters' outcomes must reflect a direct family effect rather than caregiving. By contrast, in High-IC countries where the first stage is strong, reduced-form effects capture a mixture of caregiving and any direct channel.

Table 5 reports reduced-form estimates of the parental health shock on five outcomes (depressive symptoms, SN satisfaction, SC (quality), participation in any associational activity, and satisfaction with activities) separately by country group, alongside the corresponding first-stage coefficients for reference.

The results show a clear dissociation. For depressive symptoms (EURO-D), the parental health shock has a significant positive effect across most country groups, *including Continental Europe* (0.405, $p < 0.01$), where the first stage is near zero ($\hat{\pi} = 0.007$, F$= 0.2$). This confirms that a genuine family effect exists for mental health: witnessing a parent's health deteriorate raises daughters' depression independently of whether they actually take up caregiving. Notably, this also implies that EURO-D is not a clean instrument-validation outcome – its reduced-form response combines a direct emotional channel with any caregiving-mediated effect.

Crucially, however, the picture is entirely different for our social network outcomes. In Continental Europe, the reduced-form effect on SN satisfaction is essentially zero (-0.037) and statistically indistinguishable from zero, and the same holds for Northern Europe, where the first stage is virtually null. By contrast, the reduced-form effect on SN satisfaction is negative and statistically significant in Southern countries (-0.199^{**}) and in the High-IC group (-0.168^{**}) – precisely the contexts where caregiving take-up is highest.

The pattern is even sharper for SC (quality), our primary structural measure of the network. In Continental Europe the reduced-form point estimate is actually positive (+0.697) and non-significant, and it is also non-significant in North (-0.892) and (marginally) in East (-2.12^{*}). In High-IC countries, by contrast, it is clearly negative and significant (-3.08^{**}). The Southern point estimate is negative but not significant (-1.46), reflecting the smaller sample and the

first stage is near zero dilutes estimated effects without adding identification. In all specifications, the SN size column is insignificant, consistent with the quality-not-quantity finding in Table 4.

noisier South-only first stage. Taking Panels C and D together, both social-network outcomes respond only where caregiving actually takes place: they are indistinguishable from zero in country groups with a null first stage, and negative and significant in the High-IC group. This is exactly the pattern that a caregiving-mediated effect predicts and that a pure family-effect story rules out. Effects on our social-network outcomes are thus concentrated in the country group defined by intensive-care prevalence rather than by geography.

Table 5: Reduced-form effects of parental health shock by country group

	(1) All	(2) South	(3) Continental	(4) North	(5) East	(6) High-IC
Panel A: First stage (dep. var.: Daily caregiving)						
	(1) All	(2) South	(3) Continental	(4) North	(5) East	(6) High-IC
Parental health shock	0.055***	0.143***	0.007	0.001	0.098***	0.150***
	(0.010)	(0.030)	(0.015)	(0.014)	(0.021)	(0.025)
F-Stat	28.7	22.3	0.2	0.0	20.9	35.8
N	5,366	849	2,285	588	1,580	1,317
Panel B: Mental health (dep. var.: EURO-D depressive symptoms)						
Parental health shock	0.384***	0.528***	0.405***	0.122	0.398***	0.477***
	(0.064)	(0.181)	(0.093)	(0.171)	(0.120)	(0.144)
Mean	2.40	2.27	2.57	2.13	2.28	2.27
N	5,309	843	2,263	583	1,560	1,304
Panel C: SN satisfaction (dep. var.: SN satisfaction, 0–10)						
Parental health shock	-0.092**	-0.199**	-0.037	-0.141	-0.093	-0.168**
	(0.036)	(0.092)	(0.054)	(0.093)	(0.072)	(0.074)
Mean	9.05	9.10	8.87	9.32	9.19	9.13
N	5,358	848	2,280	588	1,578	1,315
Panel D: SC (quality) (dep. var.: SC (quality))						
Parental health shock	-0.841	-1.460	0.697	-0.892	-2.123*	-3.075**
	(0.638)	(1.601)	(0.978)	(1.801)	(1.250)	(1.303)
Mean	66.75	73.39	62.91	59.87	70.74	71.79
N	5,366	849	2,285	588	1,580	1,317
Panel E: Any associational activity (dep. var.: participation in any activity, 0–1)						
Parental health shock	-0.019	-0.038	-0.008	-0.043	-0.018	-0.036
	(0.014)	(0.032)	(0.022)	(0.041)	(0.025)	(0.025)
Mean	0.481	0.224	0.606	0.754	0.342	0.213
N	5,352	848	2,277	585	1,578	1,315
Panel F: Satisfaction with activities (dep. var.: satisfaction, 0–1 scale)						
Parental health shock	-0.020***	-0.044**	-0.008	-0.003	-0.019	-0.043***
	(0.006)	(0.019)	(0.007)	(0.012)	(0.012)	(0.015)
Mean	0.826	0.762	0.840	0.888	0.814	0.776
N	5,321	840	2,263	585	1,572	1,305

Notes: Standard errors in parentheses, clustered at individual level. All regressions include country, year, cohort FE, and country × cohort FE. Controls: age, married, education, number of children and grandchildren, household size, siblings, female siblings, coresidence, income quartile, wealth quartile, rural residence, and the baseline outcome $Y_{i,0}$. All other covariates are measured at $t-1$. Country groups: South = Spain, Italy, Greece, Portugal, Cyprus; Continental = Austria, Germany, Netherlands, France, Switzerland, Belgium, Luxembourg; North = Sweden, Denmark, Finland; East = Czech Republic, Poland, Hungary, Slovenia, Estonia, Croatia, Lithuania, Bulgaria, Latvia, Romania, Slovakia. High-IC = Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, Slovakia (OECD high-intensity threshold). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The direct emotional response to parental illness does affect mental health across all contexts, but it does *not* erode social networks in countries where daughters rarely provide care. Social network erosion is concentrated exactly where caregiving is most prevalent, consistent with time pressure and behavioral constraints – rather than grief – as the operative channel.

Panels E and F extend the same test to associational activities. In Continental Europe, the reduced-form effects on participation in any activity and on satisfaction with activities are small and statistically indistinguishable from zero (-0.008 and -0.008, respectively). In Northern Europe the coefficients are also near zero. In the High-IC group, the reduced-form effect on any-activity participation is -0.036 with the expected sign but not significant at conventional levels; the reduced-form effect on satisfaction with activities is -0.043 ($p < 0.01$), consistent with a genuine caregiving-driven decline in the quality of engagement even when participation itself is not sharply crowded out. The pattern is consistent with time demands of caregiving affecting subjective engagement with activities more sharply than the extensive margin of participation.

Our IV is therefore valid for all social and behavioral outcomes that are the focus of this paper. For mental health outcomes such as EURO-D, by contrast, the exclusion restriction is likely imperfect: the large reduced-form effects on depression in Continental Europe (where the first stage is near zero) reveal that parental illness affects mental health through channels other than caregiving. Our SN and activities results are free from this concern.

5.4 Zooming In: Which Network Dimensions Are Affected?

Figure 5 decomposes SC (quality) into its three share-based components for High-IC countries: geographic proximity, contact frequency, and emotional closeness. All three point estimates are negative, but the effect is concentrated in emotional closeness: caregivers rate a significantly smaller share of their confidants as emotionally close, while the proximity and frequency components are negative but imprecisely estimated. This pattern suggests that intensive caregiving primarily erodes the perceived quality of existing ties – felt intimacy – rather than reducing their geographic reach or the regularity of contact.

Appendix Figure A.2 probes whether the drop in the emotional-closeness share reflects a change in *network composition* or a change in *felt closeness within existing ties.* We regroup confidants into four mutually exclusive relationship types – immediate kin (spouse and children), extended kin (siblings, in-laws, and other relatives, excluding parents), friends, and other non-kin (colleagues, neighbours, and other contacts) – and estimate 2SLS effects on the *share* of each type in the network of non-parental confidants (looking at shares rather than counts neutralises any composition change driven mechanically by the total network size). Two directional patterns emerge: the share of friends falls by roughly 8 percentage points (from a baseline mean of about 12%), while the share of immediate kin rises by roughly 10 percentage points (from a baseline of about 73%); extended kin and other non-kin shares move only marginally. Individually none of these coefficients reaches conventional significance, but the joint shift is coherent: caregivers' networks become more concentrated on immediate family and less on friends. Two non-exclusive interpretations are consistent with this pattern: caregiving may reduce the maintenance of elective friendships, or the parental health shock may draw the daughter closer to

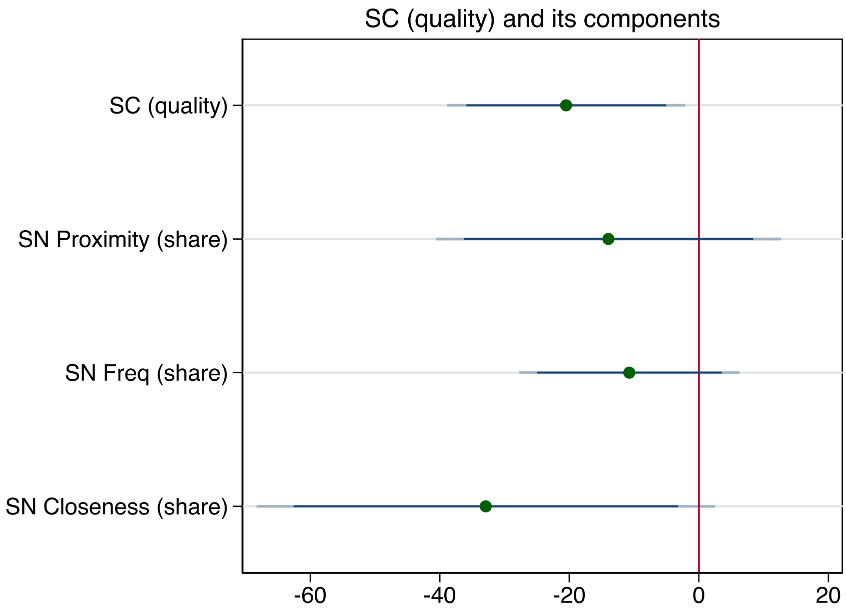

Fig. 5. 2SLS estimates: SC (quality) and components, High-IC countries

Notes: 2SLS estimates for High-IC countries. Inner bars show 90% CI; outer lines show 95% CI. SC (quality) is the share-based quality index; components are the share of confidants who are geographically close, in frequent contact, and emotionally close. High-IC countries: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia.

her immediate family. We cannot separate these mechanisms with the data at hand, and treat the composition evidence as corroborating rather than conclusive. Appendix Figure A.1 replicates the SC (quality) decomposition for the all countries; the closeness pattern persists, though estimates are less precise outside High-IC countries.

5.5 Who bears the social cost?

The baseline results speak to the average High-IC caregiver. Three dimensions of heterogeneity within this group are theoretically motivated. The first is whether the caregiver has siblings with whom she can share the care burden. A natural first cut is by *any* siblings, but the no-sibling subsample is too small to identify a LATE: only about 100 daughters in High-IC report no siblings at all, and the first stage collapses. We therefore consider a sharper and more theoretically apt cut: whether the caregiver has a *female* sibling. In High-IC countries, gender norms assign filial care primarily to daughters (see Sensitivity D, where identification runs off maternal poor health), so a sister is the relevant substitute; a brother, less so. The no-female-sibling subgroup is substantially larger ($N = 486$) and delivers a strong first stage ($F = 20.2$). The second dimension is whether the caregiver is employed: employed daughters have less slack in their time endowment, so the same caregiving load competes with a tighter set of alternative uses of time. The third is age: the sample spans 50–70, a range that includes both women still in the labour force and women who have exited it. Because the composition of a caregiver's discretionary time changes at retirement, the *which* margin of social life gets crowded out may differ. The sibling and employment splitters are measured at the respondent's baseline observation, so they are fully predetermined with respect to any subsequent parental poor-health event; splitting on the lagged value $t-1$ would risk conditioning on a variable that has itself already responded to the shock (a daughter anticipating care may have exited work at $t-1$). The age split uses age at t, which is arithmetically predetermined. Table 6 reports 2SLS estimates for the three primary outcomes separately by these subgroups.

Table 6: Heterogeneity within High-IC: 2SLS estimates by subgroup

	SC (quality)	SN Satisfaction	Any activity	N	K-P F
Baseline (all High-IC)	-20.49**	-1.11**	-0.24	1,317	35.5
	(9.36)	(0.53)	(0.17)		
No female siblings	-23.96*	-1.18*	-0.23	486	20.2
	(12.74)	(0.70)	(0.21)		
Has female siblings	-15.40	-0.63	-0.21	829	21.5
	(12.15)	(0.65)	(0.24)		
Employed at baseline	-27.13***	-1.28**	-0.15	633	26.1
	(10.41)	(0.62)	(0.21)		
Not employed at baseline	-1.33	-0.96	-0.40	683	9.6
	(17.12)	(1.05)	(0.33)		
Age 50–64 at t	-23.68**	-1.41**	-0.08	1,007	24.7
	(10.75)	(0.65)	(0.20)		
Age 65–70 at t	-2.15	0.54	-0.75*	243	8.6
	(21.15)	(1.19)	(0.45)		

Notes: Each row is a separate 2SLS estimation on the indicated subgroup of the High-IC sample. Sibling and employment splitters are measured at the respondent's baseline observation, i.e. her first wave in the panel, so that the split is predetermined w.r.t. any subsequent parental poor-health event. The age split uses age at t, which is arithmetically predetermined. Same specification as the baseline (Table 4): baseline outcome $Y_{i,0}$ and the full covariate vector included. Standard errors clustered at the individual level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. K-P F is the Kleibergen–Paap first-stage F-statistic for the SC (quality) regression (very similar across outcomes within a subgroup).

Three patterns emerge, and together they point to what might be a two-channel erosion story that depends on the caregiver's stage of life.

Employment. Employed daughters appear to show substantially larger effects on both subjective and structural quality margins: they lose 1.28 points of network satisfaction (vs. 1.11 in the baseline, $p < 0.05$) and roughly 27 points of SC (quality) on a 100-point scale (vs. 20.5 in the baseline, $p < 0.01$). Non-employed daughters show essentially no effect on SC (quality) (-1.33, NS) and a smaller negative effect on SN satisfaction (-0.96, NS). The first stage in the non-employed cell is weak ($F = 9.6$), so the point estimates there are best read as descriptive; the contrast that matters is with the employed subgroup, where identification is strong ($F = 26$). Substantively, women who are already non-employed at baseline may be less likely to see intense caregiving crowd out labour supply – there is little labour supply left to crowd out – which would be consistent with a time-pressure reading of the effect.

Age. The age split seems to reveal a mirror image. Among daughters aged 50–64 at wave

t, quality outcomes move sharply – SC (quality) -23.68^{**} and SN satisfaction -1.41^{**} with $F = 24.7$ – while the activity margin barely moves (-0.08, NS). Among daughters aged 65–70 at t, the pattern appears to flip: SC (quality) and SN satisfaction shrink to near zero (and slightly the "wrong" sign for satisfaction), while any-activity participation drops by -0.75 ($p = 0.10$ in Wald), which would represent a very large effect on a base rate of about 21%. The 65–70 subgroup has a weak first stage ($F = 8.6$, $N = 243$), so Wald inference is unreliable. We therefore complement it with the Anderson-Rubin test, which is valid regardless of first-stage strength; because our design is just-identified, AR reduces to the reduced-form F-test of the instrument. The reduced form on any-activity in the 65–70 subgroup is -0.121 ($SE = 0.057$, $p = 0.034$), so the null of no caregiving effect on activities is rejected at 5% under weak-instrument-robust inference. Given the small subsample, this reading is suggestive rather than conclusive.

Sibling. The sibling split shows the same directional pattern as the baseline – point estimates are larger in absolute value among daughters without a female sibling than among those with one – and reaches marginal significance for the no-female-sibling group on both SC (quality) (-24.0^{*}) and SN satisfaction (-1.18^{*}). Both sibling subsamples have first-stage F-statistics slightly below the Montiel-Pflueger 23.1 threshold applied to the main specification ($F = 20.2$ for no female siblings, $F = 21.5$ for has female siblings), so the point estimates in this row should be read as suggestive rather than sharply identified.

The age and employment splits, read jointly, might describe two erosion channels that depend on where the caregiver sits in the life cycle. Younger, still-working caregivers may absorb the burden through the quality of their existing ties: closeness, satisfaction, and the perceived depth of the network appear to fall, while the extensive margin of associational participation moves little because it is already low. Older, mostly-retired caregivers seem to absorb the burden through the extensive margin: their social lives could run more through associational engagement – clubs, religious groups, volunteering – which caregiving might crowd out, while the quality of their remaining close ties appears unaffected. The sibling result would complement both channels: identification runs through the female line, and gender norms may bind hardest where there is no sister with whom to share the load, consistent with the maternal-vs-paternal decomposition in Sensitivity D.

A within-High-IC split on rural residence also collapses the first stage ($F < 1$ in the rural subgroup, $N = 391$), so we do not report it.

6 Sensitivity Analysis

This section evaluates the robustness of our main causal estimates. We organize the checks around distinct threats to internal validity and design choices. Some (the AR inference test, the age restriction, and threshold variation) are conventional robustness checks where a materially different result would be cause for concern. Others (notably the maternal versus paternal instrument decomposition and the EURO-D exercise) are not designed to replicate the baseline; they illuminate which sub-populations drive identification and which channels carry the effect. Results throughout refer to the High-IC sample. Full coefficient tables appear in the

Appendix. Figures 6, 7, and 8 provide a visual summary for our three significantly-affected primary outcomes – SC (quality), SN satisfaction, and satisfaction with activities – plotting the 2SLS coefficient (with 90% and 95% confidence intervals) for each check.

6.1 Inference: Weak-Instrument Robustness

We first assess whether conventional Wald inference is reliable given our first stage. The Kleibergen-Paap F-statistic is approximately 36 in High-IC countries, well above commonly cited thresholds. To confirm this, we compute Anderson-Rubin (AR) confidence intervals, which are valid regardless of first-stage strength, alongside the Montiel-Pflueger effective F (Montiel Olea and Pflueger, 2013b). The AR test inverts the reduced-form regression rather than relying on the first-stage estimate, making it robust to arbitrary weakness of the instrument.

For SN satisfaction, the 2SLS estimate is -1.11 with Wald 95% CI $[-2.12, -0.11]$ and Anderson-Rubin (AR) test $F(1, 1119) = 5.13$ ($p = 0.024$). For SC (quality), the estimate is -20.5 with Wald 95% CI $[-38.4, -2.6]$ and AR test $F(1, 1121) = 5.57$ ($p = 0.019$). Both intervals exclude zero and both AR tests reject the null at 5%. For any associational activity, the point estimate is -0.24 with Wald 95% CI $[-0.58, 0.09]$ (includes zero) and AR test $F(1, 1120) = 2.10$ ($p = 0.147$), also failing to reject at conventional levels. The Montiel-Pflueger effective F is approximately 36 in the High-IC sample, well above the critical value of 23.1 for 10% worst-case size distortion (Montiel Olea and Pflueger, 2013b). Weak-instrument-robust and Wald conclusions therefore coincide.

6.2 Sample Restrictions

Age cap at 64. The OECD informal care prevalence measure is defined for women aged 45–64. Including women aged 65–70 in our baseline extends the age range and raises statistical power, but risks confounding the caregiving effect with retirement-induced network changes (reduced peripheral ties, restructured leisure). Sensitivity A restricts the estimation sample to daughters aged 50–64. In High-IC countries, the Kleibergen-Paap F-statistic falls from approximately 36 to 25 (reflecting the smaller sample, $N = 1{,}005$ vs. 1,317) but remains robust. The 2SLS estimate for SN satisfaction is -1.41^{**} and for SC (quality) -23.7^{**} – both slightly larger in absolute value than the baseline and clearly significant. The activity estimate falls to -0.08 (NS), essentially zero; the baseline point estimate of -0.24 was already not significant, and the further shrinkage suggests the small negative signal on activities was partly concentrated among the 65–70 age band. The network quality and satisfaction findings are therefore robust to restricting to the OECD reference age range; the activity result was already noisy in the baseline and cannot be pinned down here.

Threshold stability. The 22.5% threshold corresponds to a natural gap in the cross-country distribution of informal care intensity: Romania (22.2%) and Sweden (22.5%) end the lower cluster; the next country, Malta, sits at 24.4%. Sensitivity G estimates the model for eight thresholds ranging from $> 15\%$ to $> 30\%$, which progressively add or remove countries in a predictable way.[5]

[5] Moving the threshold *down*: >22% adds Romania (22.2%) and Sweden (22.5%); >20% further adds Lithua-

SN satisfaction is negative and statistically significant at the 10% level for all thresholds from >18% upward; significance strengthens to 5% at >22% and >22.5% – the thresholds that exclude Romania and Sweden but retain the core High-IC group – before returning to 10% for stricter cutoffs. Activity participation is significant at >15% and >18%, drops below conventional thresholds when Estonia, Slovenia, and Lithuania are added (>20%, >22%), re-emerges at the baseline and >25%, and again loses significance when Greece, Poland, and Bulgaria are dropped (>27%, >30%). The point estimate for activities remains stable around −0.3 throughout. The 22.5% threshold is therefore both empirically motivated by the distributional gap and consistent with a wide range of neighboring choices, though the activity result is sensitive to which mid-tier countries (Romania, Sweden) or smaller High-IC countries (Greece, Poland, Bulgaria) are included.

6.3 Specification Checks

Co-residence control. A lagged indicator of whether the respondent lives with a parent is included as a control in the baseline to absorb the correlation between residential proximity and caregiving. However, it may be a "bad control": daughters who take up intensive caregiving often move in with the parent (or vice versa), so controlling for co-residence risks absorbing part of the treatment effect. Sensitivity E drops this control from the specification. In High-IC countries, the 2SLS estimate for SN satisfaction is -1.10^{**} (vs. -1.11^{**} baseline), for SC (quality) -20.4^{**} (vs. -20.5^{**} baseline), and for any activity −0.24 (vs. −0.24 baseline). Estimates are virtually unchanged; there is no evidence that the baseline results were inflated by co-residence bad-control bias.

Parent exclusion from the network index. In the baseline, we exclude parents from the confidant roster when constructing SC (quality) and network size, to avoid the mechanical channel whereby a daughter who begins caregiving is more likely to name her parent as a close confidant. Sensitivity H constructs SC (quality) including parents in both the current and lagged observation. If the effect were mechanical, including parents should amplify the measured decline (more network entries mechanically shift when the parent's status changes). Indeed, the 2SLS estimate for SC (quality) is -22.1^{**} (vs. -20.5^{**} baseline) – larger in absolute value, as expected. SN satisfaction (-1.11^{**}) and activity participation (−0.24) are essentially unchanged. This confirms that our parent-exclusion rule produces conservative estimates: the true effect on the full confidant network is at least as large.

Controlling for mental health (EURO-D): a partial-effect exercise. Section 5.3 established that the parental health shock raises daughters' depressive symptoms significantly – including in Continental Europe, where caregiving barely responds. Because EURO-D at time t is itself affected by the instrument, adding it as a second-stage control is a potential bad-control: it may absorb part of the caregiving effect that runs through depression, and could even introduce collider bias if depression is a common consequence of both the health shock and unobserved determinants of social networks. For this reason, we do not treat this exercise as a conventional

nia (21.7%), Slovenia (20.4%), and Estonia (20.3%); >18% adds Latvia (18.1%); >15% adds Austria (15.1%). Moving the threshold *up*: >25% drops Malta (24.4%), leaving eight countries; >27% additionally drops Greece (25.8%), Poland (26.0%), and Bulgaria (26.7%), leaving only Italy, Portugal, Slovakia, Croatia, and Spain. No country lies between 27% and 30%, so >30% produces the same five-country sample as >27%.

robustness check. We include it nonetheless, labeled explicitly as a partial-effect estimate, because the result is informative: in High-IC countries, the 2SLS coefficients with EURO-D controlled are -23.5^{**} for SC (quality), -1.00^{*} for SN satisfaction, and -0.21 for any activity (F-stat ≈ 32). Both structural outcomes remain sizeable and significant; the effect on SC (quality) is even somewhat larger once depressive symptoms are held fixed. The social erosion we document is not operating primarily through the depression channel; it reflects a direct behavioral constraint on time and energy for social engagement.

6.4 Alternative Sample of Countries

The baseline High-IC classification uses a demand-side measure: the OECD share of women aged 45–64 providing intensive informal care (20+ hours per week). Sensitivity F adopts a supply-side alternative: countries where public long-term care expenditure is below 1.5% of GDP. This selects a different and larger group ($N = 1{,}954$, $F \approx 36$) that partially overlaps with High-IC but includes some Continental European countries not in the baseline group.

The results are mixed. Activity participation remains significant (-0.39^{**}). However, SN satisfaction (-0.46) and SC (quality) (-12.3) lose statistical significance, though the point estimates remain negative. Using a stricter threshold of 1.0% of GDP (which selects a smaller, more closely overlapping set with $F = 26.9$), SC (quality) recovers significance (-19.3^{*}). We interpret this pattern honestly: the demand-side measure (OECD intensive care prevalence) appears to more precisely capture the cultural and institutional context in which daughters substitute for absent formal care, generating a cleaner first stage and stronger instrument relevance. The supply-side measure is noisier – low LTC spending need not imply high informal care intensity if there are other substitutes or if need is lower. The robustness to the strict supply-side threshold provides partial reassurance.

6.5 Instrument Source: Maternal vs. Paternal Poor Health

Sensitivity D decomposes the parental poor-health indicator into maternal and paternal components. In High-IC countries, the maternal instrument (indicator for the mother being in poor health, as reported by the daughter, given the baseline restriction on maternal health) generates Kleibergen-Paap F-statistics between 37 and 38 across outcomes, indistinguishable from the baseline. The 2SLS estimates based on the maternal instrument alone are -1.23^{**} for SN satisfaction, -14.1^{*} for SC (quality), and -0.27^{*} for activity participation. The paternal instrument, by contrast, generates an F-statistic of only about 4.5 in High-IC countries – well below the level at which 2SLS inference is reliable – and the corresponding estimates are highly imprecise. This finding is not a weakness of the identification: it reflects a substantive pattern consistent with two complementary mechanisms. First, gender norms in Southern and Eastern European societies assign primary filial caregiving to daughters, particularly toward mothers (Crespo and Mira, 2014). Second, when a father is in poor health, his wife is typically still alive and well enough to absorb the bulk of the care load; the daughter is therefore less likely to become the primary caregiver in response. Maternal poor health, by contrast, leaves no equivalent substitute within the household, placing the care burden directly on daughters.

Our instrument is in practice identified off maternal poor health, and this is the channel most relevant to the social-norm argument that motivates our High-IC classification.

What does not work: the strong-shock instrument. Sensitivity B restricts the instrument to parents entering poor health from good health or better (child-reported parental health ≤ 3), excluding fair-to-poor transitions. This raises the severity threshold but drastically reduces variation: the Kleibergen-Paap F-statistic falls to 4–5 in High-IC countries (and to 2.4 pooled). 2SLS estimates with such a weak first stage are unreliable and are not interpreted causally. The difficulty is partly mechanical – high-IC countries already tend to have parents in worse baseline health, leaving fewer observations who start from "good or better" health – and partly reflects that smaller, incremental episodes (fair $\rightarrow$ poor) are meaningful triggers of caregiving in our sample. We do not treat this as a failed robustness check but as confirmation that our baseline instrument captures the relevant margin.

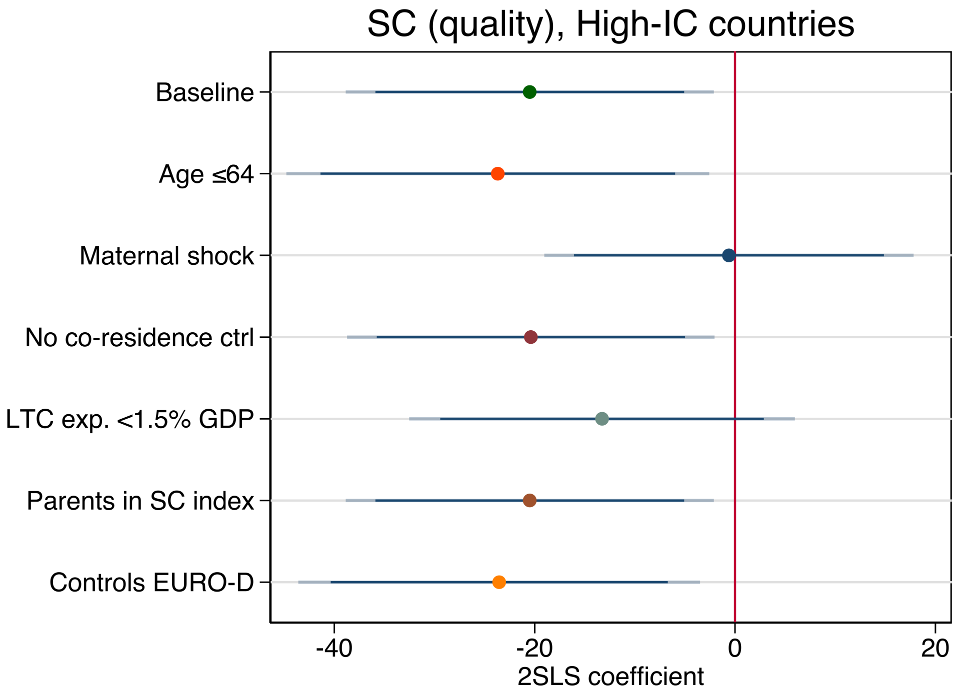

Fig. 6. Sensitivity of 2SLS estimates – SC (quality), High-IC countries

Notes: Each row shows the 2SLS coefficient and 90%/95% CI for one specification change, High-IC sample. Rows: *Baseline* – parental health shock instrument, co-residence control included, ages 50–70, OECD demand-side country definition; *Age ≤64* – sample restricted to daughters 50–64; *Maternal shock* – instrument is maternal health shock alone; *No co-residence ctrl* – lagged co-residence with parent control excluded; *LTC exp. <1.5% GDP* – supply-side country definition; *Parents in SC index* – SC (quality) index includes parents in the confidant roster; *Controls EURO-D* – contemporaneous depressive symptoms added as control (partial-effect estimate; EURO-D is potentially endogenous to caregiving – see text). Checks B (strong shock, F≈5) and D paternal (F≈3) excluded.

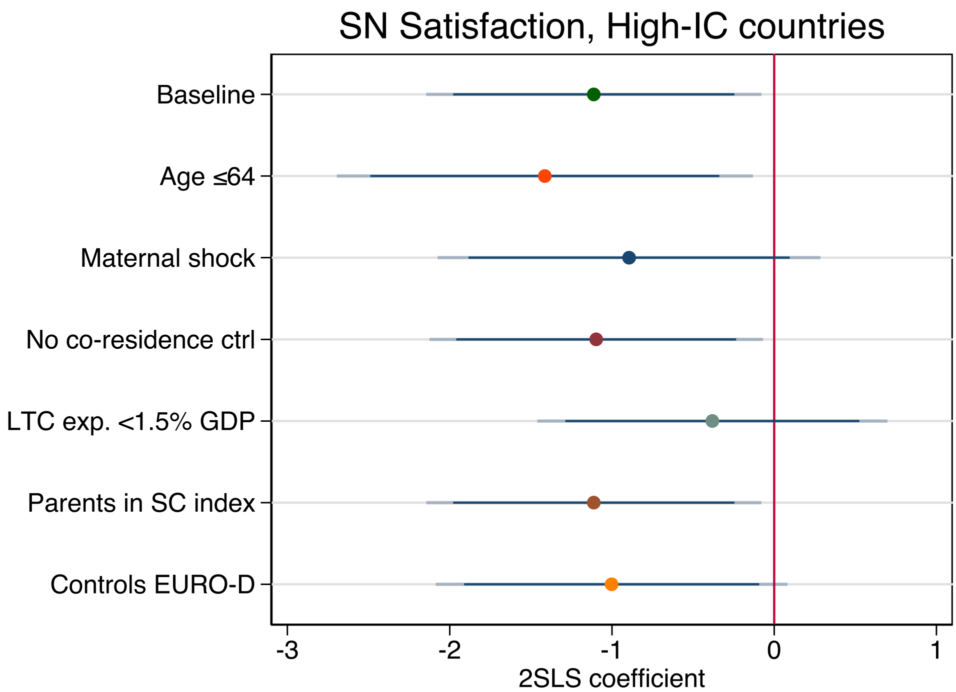

Fig. 7. Sensitivity of 2SLS estimates – SN Satisfaction, High-IC countries

Notes: Same specification as Figure 6 for the outcome "SN satisfaction".

7 Complementary Evidence on Social Erosion

The preceding analysis established that intense caregiving erodes SC (quality) and network satisfaction, with emotional closeness declining most sharply in High-IC countries, and reduces participation in associational activities. We now extend these findings in two directions. First, we decompose the binary activity result into activity types, counts, and satisfaction, gauging the breadth and depth of behavioral social erosion. Second, we trace the time-reallocation mechanism directly using wave-8 time-use data, documenting how caregiving reshapes daughters' daily schedules. These two analyses play distinct roles: the activity decomposition extends and contextualises the primary outcome, while the time-use results speak to the mechanism through which caregiving crowds out social life.

7.1 Caregiving and Associational Activities

Section 5 reported a binary outcome – whether the respondent participates in any associational activity – as one of the three primary outcomes. Here we decompose that result by activity type and examine count and satisfaction measures alongside the participation indicator. All outcomes have been standardized to facilitate comparison across different scales.

We examine participation in associational activities, which represents an important dimension of social networks. Since wave 4, SHARE respondents have been asked about their participation in various social activities over the past twelve months. These activities include voluntary or charity work, attending educational or training courses, going to sports, social, or

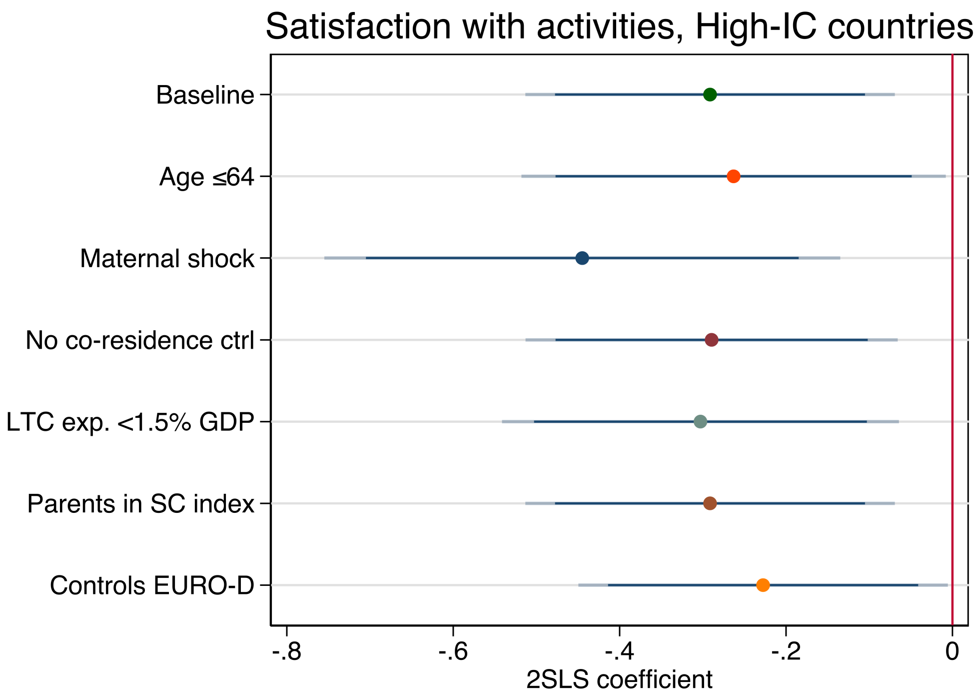

Fig. 8. Sensitivity of 2SLS estimates – Satisfaction with activities, High-IC countries

Notes: Same specification as Figure 6 for the outcome "satisfaction with associational activities", a subjective measure of engagement quality that parallels SN satisfaction.

other clubs, and participating in political or community-related organizations.

As emphasized in Fawaz and Mira (2023), following Steptoe et al. (2013), participation in social activities is a central element of social isolation indices and one of the strongest predictors of health outcomes in later life.[6] Beyond its role as a marker of isolation, associational participation is also widely used in the social capital literature as a behavioral proxy for otherwise hard-to-observe norms of trust and reciprocity. Following Alesina and La Ferrara (2000), we rely on measures of whether respondents engage in voluntary or charity work, attend a sport, social, or other type of club, participate in political or community organizations. [7] As Putnam (1995) argues, face-to-face civic engagement of this type generates the externalities associated with social capital. In this sense, our analysis of associational activities captures a dimension of social connectedness that is simultaneously central to healthy ageing and to the production of social capital.

Figure 9 presents 2SLS estimates by activity type for High-IC countries. Coefficients are in standard deviation units, facilitating comparison across the heterogeneous original scales (0–10 for satisfaction, 0–4 for counts, and 0–1 for binary indicators).

[6] Fawaz and Mira (2023) show that exposure to social isolation at baseline is associated with a 20–30% increase in the mortality hazard. Among the three items used in their isolation index (living alone, no participation in associational activities, and infrequent contact with children), the absence of social activities was found to drive almost the entire association with mortality.

[7] Participation in religious activities is not considered, since this information was only collected in the early SHARE waves. Likewise, solitary pastimes such as reading or doing puzzles are excluded, as they do not reflect the social dimension of networks.

The results point to consistent negative effects. Caregiving reduces the likelihood of participation in social or civic activities, with the largest declines in volunteering and charity work, as well as in overall activity counts. Satisfaction with activities also falls substantially, by about one standard deviation, suggesting that the subjective dimension of engagement is sensitive to caregiving pressures. Caregiving erodes social connectedness in part by crowding out associational participation – the activities that sustain broader social ties and civic engagement.

We stress that this standardization was applied only to associational activities. For social network outcomes we retain the original SHARE scales, and for time allocation we report effects in minutes, so that magnitudes remain directly interpretable.

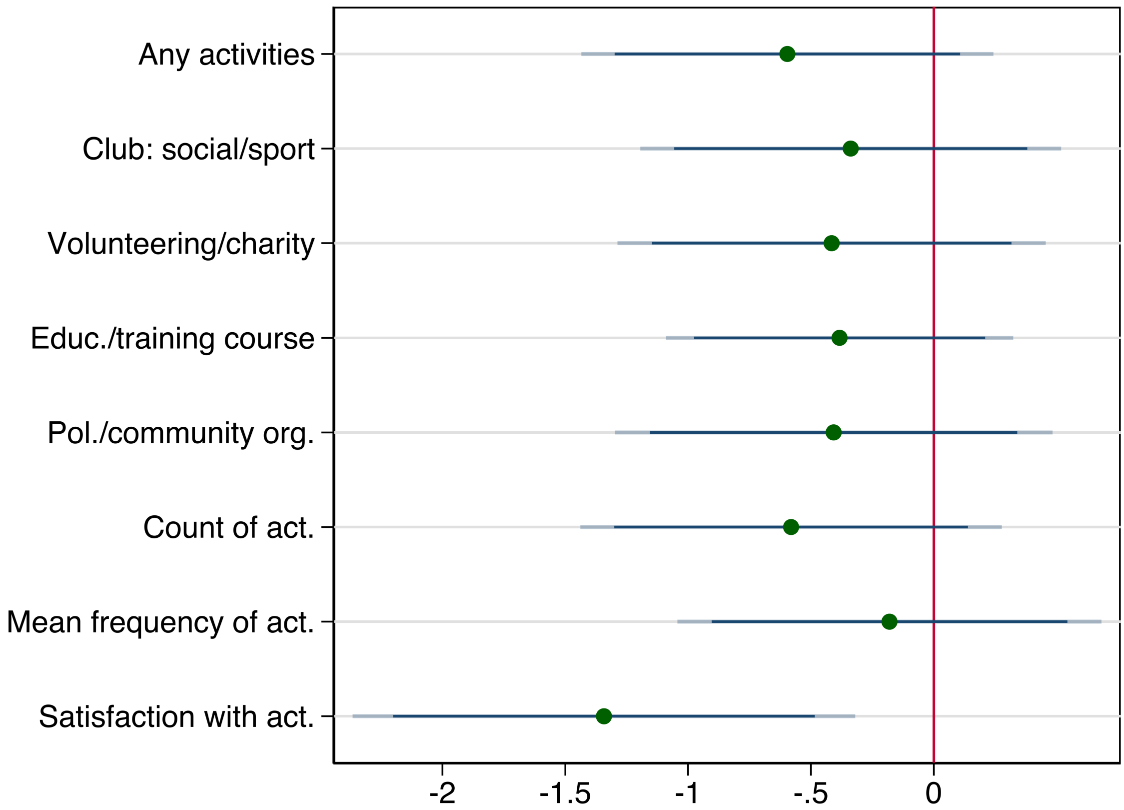

Fig. 9. 2SLS estimates of impact of caregiving on activities, High-IC countries

Notes: 90% and 95% CI are plotted along the corresponding coefficient. High-IC countries: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia. Outcomes were standardized to mean 0, SD 1.

To gauge magnitudes, in High-IC countries the point estimate for the probability of participating in any associational activity implies a reduction of about 24 percentage points, and the count of activities falls by approximately 0.35 activities (relative to a mean of 0.28 in the estimation sample). Both extensive-margin measures are directionally consistent but not statistically distinguishable from zero in the main specification. Satisfaction with activities, by contrast, declines by about 0.30 points on a 0–1 scale (equivalent to 3 points on the original 0–10 scale), and this intensive-margin effect is precisely estimated ($p < 0.01$). The pattern suggests that intensive caregiving affects the perceived quality of civic engagement more sharply than the objective count or any-participation margin: among daughters who remain active, engagement thins in depth and quality even when the raw participation rate does not detectably change.

7.2 The Time-Reallocation Mechanism

The time use module in wave 8 of SHARE records respondents' allocation of time on the day prior to the interview, asking about hours and minutes spent in specific activities such as paid work, caregiving, leisure, and sleep. Although this information was not yet available when Crespo and Mira (2014) developed their model, it speaks directly to the mechanism they emphasized: caregiving competes with other uses of time, and daughters reallocate hours between market work, care provision, and leisure in response to parental health shocks. Their theoretical framework formalized the daughter's time endowment as being split between employment, caregiving, and leisure, and predicted that daily caregiving represents a significant burden in this allocation. The SHARE time use data now allow us to observe these trade-offs directly. By comparing the time spent on alternative activities between caregivers and non-caregivers, we can provide empirical evidence on how caregiving displaces other uses of time, particularly social and leisure activities that sustain connectedness.

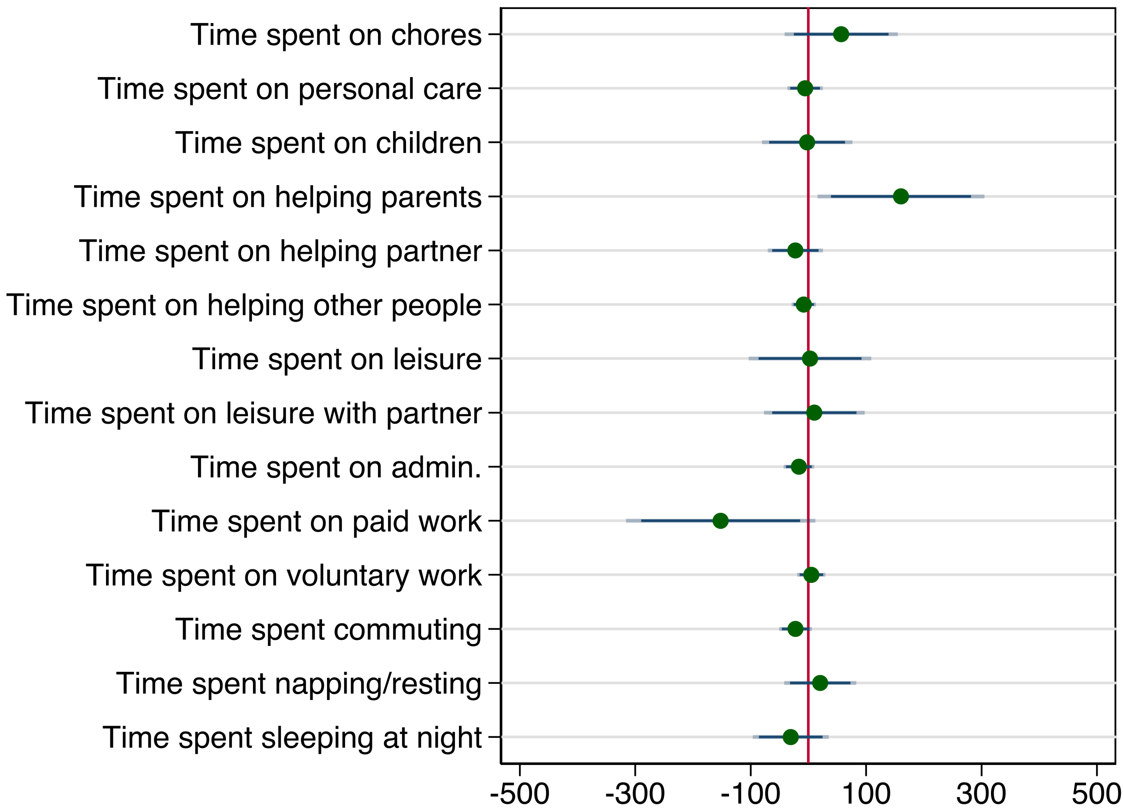

Fig. 10. 2SLS estimates of impact of caregiving on time allocation, High-IC countries

Notes: 90% and 95% CI are plotted along the corresponding coefficient. High-IC countries: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia. Time is expressed in minutes.

Figure 10 plots the 2SLS estimates for High-IC countries. The time spent helping parents rises by 160 minutes per day (about 2.7 hours, $p < 0.05$), confirming that the instrument successfully shifts caregiving intensity. The picture for other activities is consistent with the Crespo–Mira time-pressure mechanism: paid work falls by 152 minutes ($p < 0.10$) and sleep by about 30 minutes (not significant), while leisure changes little. Table A.7 in the Appendix reports the full coefficient table with exact standard errors, and checks robustness to controlling for lagged employment status. Roughly speaking, an increase in daily caregiving of about 2.7 hours is offset by a comparable reduction in paid work, with smaller adjustments elsewhere – tightening the time constraint in ways that limit opportunities for social and civic engagement.

7.3 Summary

Our results provide a consistent picture of the social costs of caregiving. First, the first-stage evidence shows that parental health shocks are strong predictors of daughters' entry into daily caregiving, particularly in High-IC countries where formal long-term care provision is scarce. Second, the reduced form and 2SLS estimates indicate that caregiving affects the quality of social networks more than their size. While the number of confidants remains broadly stable (and in some specifications even slightly larger), emotional closeness to confidants declines and satisfaction with one's network falls. Appendix estimates also suggest some shift in network composition toward more peripheral ties, though these composition effects do not reach conventional significance. Focusing only on size metrics would therefore miss the more meaningful ways in which caregiving reshapes social connectedness.

Third, the complementary evidence from associational activities and time use corroborates and extends these findings. Caregiving crowds out participation in civic and social activities and reduces the satisfaction derived from them. Time-use data are consistent with the time-pressure mechanism: daughters reallocate time away from paid work into parental care, with large but imprecise reductions in sleep; leisure is stable. These estimates are suggestive rather than conclusive given the small time-use sample, but their direction and pattern are consistent with time pressure as the operative channel of social erosion.

Caregiving induced by parental health shocks leaves women with networks that may not look smaller in size but are weaker in terms of closeness, composition, and engagement. These findings highlight why assessments of connectedness must go beyond headcounts of ties to capture their quality, composition, and the activities that sustain them.

8 Corroborating Evidence: Social Engagement and Survival

As a final piece of corroborating evidence, Appendix B reports Kaplan–Meier survival curves and Cox proportional hazard models linking baseline social engagement to ten-year survival among women aged 50–70 in SHARE. To highlight the mortality-relevance of the dimensions we study, Table 7 juxtaposes, for each of eight social network dimensions, the 2SLS effect of intensive caregiving on that dimension (Panel A, High-IC countries) and the corresponding Cox hazard ratio for baseline exposure (Panel B, all SHARE women 50–70 with >24-month follow-up, conditioning on a detailed battery of health controls including frailty, chronic conditions, functional limitations, cognition, and depressive symptoms).

Table 7: The mortality link: 2SLS effect and Cox hazard ratio side by side

	SN Sat.	SN Quality	Closeness	Any Activity	Size	Freq.	Proximity	SC Index
Panel A: 2SLS effect of intense caregiving, High-IC countries								
2SLS coef	-1.112**	-20.494**	-32.902*	-0.244	-0.428	-10.741	-13.960	-0.240
	(0.526)	(9.362)	(18.029)	(0.175)	(0.569)	(8.671)	(13.566)	(0.315)
Panel B: Cox hazard ratios, baseline at Wave 4, follow-up to Wave 9								
Cox HR	1.040	0.998	0.999	0.782***	0.955*	0.999	1.000	0.934
	(0.028)	(0.002)	(0.001)	(0.073)	(0.025)	(0.001)	(0.001)	(0.054)

Notes: Panel A reports the 2SLS coefficient of intensive caregiving on each social-network dimension (High-IC countries, main specification of Table 4). Panel B reports the exponentiated coefficient (hazard ratio) from a Cox proportional hazards model on time-to-death for SHARE women aged 50–70, using the baseline (Wave 4) value of the dimension as the sole social predictor, with the same rich health/socio-economic controls as in Appendix Table A.10. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Reading the two panels together shows which dimensions are both eroded by caregiving and relevant for long-run survival. Two dimensions have a significantly negative 2SLS coefficient at the 5% level (SN satisfaction and SC quality); their Cox hazard ratios are essentially one and not statistically distinguishable from unity, so any pathway from caregiving to mortality through them is muted. The Any-Activity indicator behaves oppositely: its 2SLS coefficient is negative but not significant at conventional levels (the direction is consistent with a caregiving-induced decline in participation), while its Cox hazard ratio is by far the largest in the table (HR $= 0.782$, $p < 0.01$), meaning that at baseline having any associational activity is strongly protective against ten-year mortality. So the behavioural margin is the one that links caregiving to mortality: the 2SLS point estimate implies a 24 percentage-point reduction in participation, which, taken at face value against the Cox HR, corresponds to a roughly 6% increase in the mortality hazard.[8]

9 Conclusion

This paper contributes to the literature on informal caregiving by documenting its impact on social connectedness and, ultimately, on caregivers' health. Using rich panel data from SHARE and an IV design based on parental health shocks, we provide new causal evidence on how caregiving shapes the social lives of European women aged 50–70.

Our results show that, in the pooled sample, effects on the standard social network indexes are modest. Yet once we exploit cross-country heterogeneity, a clearer picture emerges: in High-IC countries, where formal long-term care provision is limited and the burden of family care is greatest, intense caregiving significantly erodes the qualitative dimensions of networks. Satisfaction with one's network and the perceived closeness of ties decline substantially, even if the number of confidants remains unchanged. This demonstrates that looking only at size-based measures risks missing the more meaningful ways in which caregiving reshapes connectedness.

[8] The 2SLS point estimate on "Any Activity" is not statistically distinguishable from zero in the High-IC sample, so this back-of-the-envelope calculation should be read as indicating the magnitude of the mediation channel rather than as a formal test. The remaining dimensions in Panel B (network size, frequency, proximity, SC index) have hazard ratios close to one, indicating that they do not carry a mortality-relevant signal at baseline.

Beyond the network module, we find consistent effects on other dimensions of social life. Caregiving lowers satisfaction with associational activities: precisely the forms of engagement that sustain civic ties and social capital. The time-use data corroborate this mechanism: daughters reallocate large amounts of time from paid work into caregiving, with smaller reductions in sleep and little change in leisure, tightening the time constraint on behaviorally structured social engagement. These findings are consistent with a social-erosion mechanism through which caregiving affects well-being, operating primarily through time constraints on behavioral engagement. While our analysis focuses on the acute effects of transitioning into care, an important avenue for future research is to examine whether these social costs accumulate or intensify with the duration of caregiving spells, a question that will require even larger panel datasets to credibly answer.

Corroborating evidence links one dimension of the mechanism to long-run health: baseline participation in any associational activity predicts ten-year survival even after conditioning on detailed health controls (HR = 0.782, $p < 0.01$), while subjective network quality (HR = 0.998) and network satisfaction (HR = 1.040) do not. Among the dimensions we study, the behavioural margin (any-activity participation) is the one that carries the mortality gradient: the 2SLS coefficient is negative in the expected direction and, taken at face value against the Cox hazard ratio, is consistent with about a 6% increase in ten-year mortality risk.

Our findings provide evidence that informal caregiving carries social costs that go beyond the well-documented labor market and mental health consequences. These costs are especially acute in contexts where family caregiving substitutes for formal long-term care, and they operate through mechanisms (such as participation in social activities) that have direct implications for mortality. Recognizing and addressing these social costs should be central to policies aimed at supporting caregivers in ageing societies.

References

Abrahamsen, S. A. and Grøtting, M. W. (2023). Formal care of the elderly and health outcomes among adult daughters. *Health Economics*, 32(2):436–461.

Alesina, A. and La Ferrara, E. (2000). Participation in heterogeneous communities. *The quarterly journal of economics*, 115(3):847–904.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Bauer, J. M. and Sousa-Poza, A. (2015). Impacts of informal caregiving on caregiver employment, health, and family. *Journal of population Ageing*, 8:113–145.

Bolin, K., Lindgren, B., and Lundborg, P. (2008). Your next of kin or your own career?: Caring and working among the 50+ of europe. *Journal of health economics*, 27(3):718–738.

Bom, J., Bakx, P., Schut, F., and Van Doorslaer, E. (2019). Health effects of caring for and about parents and spouses. *The Journal of the Economics of Ageing*, 14:100196.

Bom, J. and Stöckel, J. (2021). Is the grass greener on the other side? the health impact of providing informal care in the uk and the netherlands. *Social Science & Medicine*, 269:113562.

Casado-Marín, D., García-Gómez, P., and López-Nicolás, Á. (2011). Informal care and labour force participation among middle-aged women in spain. *SERIEs*, 2:1–29.

Coe, N. B. and Van Houtven, C. H. (2009). Caring for mom and neglecting yourself? the health effects of caring for an elderly parent. *Health economics*, 18(9):991–1010.

Colombo, F., Llena-Nozal, A., Mercier, J., and Tjadens, F. (2016). Help wanted. *Providing and paying for long-term care*, 1.

Crespo, L. and Mira, P. (2014). Caregiving to elderly parents and employment status of european mature women. *Review of Economics and Statistics*, 96(4):693–709.

de Zwart, P., Bakx, P., and van Doorslaer, E. (2017). Will you still need me, will you still feed me when i'm 64? the health impact of caregiving to one's spouse. *Health Economics*, 26(S2):127–138.

Di Novi, C., Jacobs, R., and Migheli, M. (2015). The quality of life of female informal caregivers: from scandinavia to the mediterranean sea. *European journal of population*, 31(3):309–333.

Eibich, P. (2023). Instrumental variable estimates of the burden of parental caregiving. *The Journal of the Economics of Ageing*, 26:100467.

European Commission (2022). The 2021 long-term care report. https://ec.europa.eu/social/main.jsp?langId=en&catId=89&furtherNews=yes&newsId=10382. Accessed: 2024-06-18.

Fawaz, Y. and Mira, P. (2023). Social isolation, health dynamics, and mortality: evidence across 21 european countries. *Journal of Population Economics*, 36(4):2483–2518.

Fried, L. P., Tangen, C. M., Walston, J., Newman, A. B., Hirsch, C., Gottdiener, J., Seeman, T., Tracy, R., Kop, W. J., Burke, G., et al. (2001). Frailty in older adults: evidence for a phenotype. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, 56(3):M146–M157.

Heitmueller, A. (2007). The chicken or the egg?: Endogeneity in labour market participation of informal carers in england. *Journal of health economics*, 26(3):536–559.

Holt-Lunstad, J. and Smith, T. B. e. a. (2015). Loneliness and social isolation as risk factors for mortality: A meta-analytic review. *Perspectives on Psychological Science*, 10(2):227–237.

Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.

Kauppi, M., Virtanen, M., Pentti, J., Aalto, V., Kivimäki, M., Vahtera, J., and Stenholm, S. (2021). Social network ties before and after retirement: a cohort study. *European Journal of Ageing*, 18(4):503–512.

Kolodziej, I. W., Reichert, A. R., and Schmitz, H. (2018). New evidence on employment effects of informal care provision in europe. *Health Services Research*, 53(4):2027–2046.

Le, D. D. and Ibuka, Y. (2023). Understanding the effects of informal caregiving on health and well-being: heterogeneity and mechanisms. *Social Science & Medicine*, 317:115630.

Lee, D. S., McCrary, J., Moreira, M. J., and Porter, J. (2022). Valid t-ratio inference for IV. *American Economic Review*, 112(10):3260–3290.

Litwin, H. and Stoeckel, K. J. (2015). Social network, activity participation, and cognition: A complex relationship. *Research on Aging*, 38:76–97.

Montiel Olea, J. L. and Pflueger, C. (2013a). A robust test for weak instruments. *Journal of*

Business & Economic Statistics, 31(3):358–369.

Montiel Olea, J. L. and Pflueger, C. (2013b). A robust test for weak instruments. *Journal of Business & Economic Statistics*, 31(3):358–369.

Pinquart, M. and Sörensen, S. (2003). Differences between caregivers and noncaregivers in psychological health and physical health: a meta-analysis. *Psychology and aging*, 18(2):250.

Putnam, R. D. (1995). Bowling alone: America's declining social capital. *Journal of Democracy*, 6:65–78.

Roth, D. L., Fredman, L., and Haley, W. E. (2015). Informal caregiving and its impact on health: A reappraisal from population-based studies. *The Gerontologist*, 55(2):309–319.

Schmitz, H. and Westphal, M. (2015). Short-and medium-term effects of informal care provision on female caregivers' health. *Journal of health economics*, 42:174–185.

Schmitz, H. and Westphal, M. (2017). Informal care and long-term labor market outcomes. *Journal of health economics*, 56:1–18.

Settels, J. and Böckerman, P. (2025). Impact of retirement on quality of life: Role of changes in social network involvement. *Research on Aging*, 47(7-8).

Simard-Duplain, G. (2022). Heterogeneity in informal care intensity and its impact on employment. *Journal of Health Economics*, 86:102647.

Steptoe, A., Shankar, A., Demakakos, P., and Wardle, J. (2013). Social isolation, loneliness, and all-cause mortality in older men and women. *Proceedings of the National Academy of Sciences*, 110(15):5797–5801.

Stöckel, J. and Bom, J. (2022). Revisiting longer-term health effects of informal caregiving: Evidence from the uk. *The Journal of the Economics of Ageing*, 21:100343.

Thoits, P. A. (2011). Mechanisms linking social ties and support to physical and mental health. *Journal of Health and Social Behavior*, 52(2):145–161.

Van Houtven, C. H., Coe, N. B., and Skira, M. M. (2013). The effect of informal care on work and wages. *Journal of health economics*, 32(1):240–252.

Vitaliano, P. P., Zhang, J., and Scanlan, J. M. (2003). Is caregiving hazardous to one's physical health? a meta-analysis. *Psychological bulletin*, 129(6):946.

Wang, J., Zhang, M., and Zhang, Y. (2024). The effect of informal caregiving on caregivers' healthcare utilization in china. *Applied Economics*, 56(37):4405–4419.

Zwar, L., König, H.-H., and Hajek, A. (2020). Psychosocial consequences of transitioning into informal caregiving in male and female caregivers: findings from a population-based panel study. *Social Science & Medicine*, 264:113281.

A Appendix Tables

Table A.1: Sample construction

Restriction	Observations remaining	Deleted
Full SHARE sample (raw)	616,415	–
Drop wave 3 and COVID waves 82/91	477,612	138,803
Minor SN cleaning (parent-exclusion consistency)	477,532	80
Keep waves ≥ 4	409,258	68,274
Keep women	230,906	178,352
Keep age 50–70	128,618	102,288
Mother or father alive	44,334	84,284
Exclude mothers already in poor health at baseline	39,685	4,649
Exclude fathers already in poor health at baseline	38,839	846
Social Network module available (waves 4/6/8/9)	22,093	16,746
Require two SN waves (baseline outcome)	8,016	14,077
Exclude missing lagged covariates	5,352	2,664
Final estimation sample	5,352	–

Table A.2: Descriptive statistics: Main outcomes and Covariates

	All countries			High-IC countries		
	Non-CG	CG	p	Non-CG	CG	p
Primary outcomes						
SC (quality)	66.57	68.44	0.048	72.03	70.63	0.385
SN Satisfaction	9.06	9.02	0.471	9.12	9.20	0.366
Any activities	0.49	0.42	0.002	0.21	0.24	0.301
Instrument						
Parental health shock	0.25	0.39	0.000	0.24	0.49	0.000
Covariates						
Age	60.15	61.71	0.000	59.92	61.33	0.000
Years of Education	12.35	12.09	0.141	11.39	11.32	0.812
Married	0.76	0.73	0.095	0.83	0.78	0.110
Number of Children	2.00	1.83	0.001	1.90	1.67	0.005
Household Size	2.32	2.41	0.041	2.69	2.75	0.511
Number of grandchildren	1.68	1.78	0.315	1.37	1.28	0.593
Whether has Siblings	0.92	0.81	0.000	0.93	0.85	0.000
Number of female siblings	1.09	0.83	0.000	1.09	0.83	0.003
Coresidence	0.06	0.38	0.000	0.11	0.47	0.000
Income quartile	2.75	2.61	0.005	2.63	2.57	0.503
Wealth quartile	2.75	2.78	0.519	2.78	2.62	0.073
Rural residence	0.36	0.35	0.842	0.30	0.30	0.984
N (non-CG / CG)	4777 / 575			1111 / 204		

Note: Means by caregiving status (Non-CG: non-intense caregiver; CG: daily caregiver to a parent) in the full estimation sample (N = 5,352 person-wave observations, of which 575 are daily caregivers) and in High-IC countries (N = 1,315, of which 204 are daily caregivers). Sample: participation-in-any-activity regression sample (tightest of the three main outcomes). Column p reports the p-value from a two-sample t-test of equality of means. All covariates are measured in the prior wave.

Table A.3: Descriptive statistics: Secondary outcomes

	All countries			High-IC countries		
	Non-CG	CG	p	Non-CG	CG	p
Network structure (secondary)						
SC (quantity)	2.11	2.02	0.012	1.92	1.89	0.622
SN size (excl. parents)	3.02	2.79	0.001	2.55	2.54	0.916
Associational activities						
Any activities	0.49	0.42	0.003	0.21	0.25	0.206
Club: social/sport	0.30	0.24	0.001	0.09	0.13	0.049
Volunteering/charity	0.20	0.22	0.303	0.08	0.13	0.044
Educ./training course	0.21	0.15	0.000	0.06	0.04	0.369
Pol./community org.	0.06	0.06	0.536	0.04	0.05	0.499
Count of activities	0.77	0.66	0.004	0.26	0.34	0.072
Freq.: almost every day	0.10	0.14	0.054	0.09	0.13	0.283
Satisfaction with activities (0–1)	0.83	0.81	0.040	0.78	0.78	0.770
N (non-CG / CG)	4790 / 576			1112 / 205		

Note: Means by caregiving status in the same samples as Table A.2. SC quality is the average share of non-parental confidants who are geographically close, in frequent contact, and emotionally close (0–100 scale). Satisfaction with activities is rescaled to the unit interval. Column p reports the p-value from a two-sample t-test of equality of means.

Table A.4: Reduced form estimates, all countries – SC (quality), SN Satisfaction, Activities, and SN Size

	(1) SC (quality)	(2) SN Satisfaction	(3) Any activity	(4) SN size
Parental health shock	-0.841	-0.092**	-0.019	0.034
	(0.638)	(0.036)	(0.014)	(0.047)
Mean	66.75	9.05	0.48	3.01
R2	0.211	0.135	0.316	0.251
N	5,366	5,358	5,352	5,366

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort X country FE. Controls: age, married, education, number of children and grandchildren, size of household, siblings, female siblings, coresidence, income quartile, wealth quartile, rural residence, and the baseline outcome $Y_{i,0}$. All other covariates are measured at $t-1$.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: OLS vs. 2SLS Estimates, all countries – SC (quality), SN Satisfaction, Activities, and SN Size

	Panel A: OLS			
	(1) SC (quality)	(2) SN Satisfaction	(3) Any activity	(4) SN size
Intense caregiving to parent	1.152	-0.00216	0.0164	0.0210
	(1.005)	(0.0516)	(0.0201)	(0.0677)
	Panel B: 2SLS			
Intense caregiving to parent	-15.37	-1.683**	-0.352	0.626
	(12.03)	(0.734)	(0.266)	(0.855)
F-test (IV)	28.64	28.60	27.90	28.81
Hausman p-value	0.15	0.01	0.15	0.47
Observations	5366	5358	5352	5366

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort X country FE. Controls: age, married, education, number of children and grandchildren, size of household, siblings, female siblings, coresidence, income quartile, wealth quartile, rural residence, and the baseline outcome $Y_{i,0}$. All other covariates are measured at $t-1$.
The F-test for the excluded instrument is the Kleibergen-Paap rk F-statistic. The Hausman p-value tests the endogeneity of intensive caregiving. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

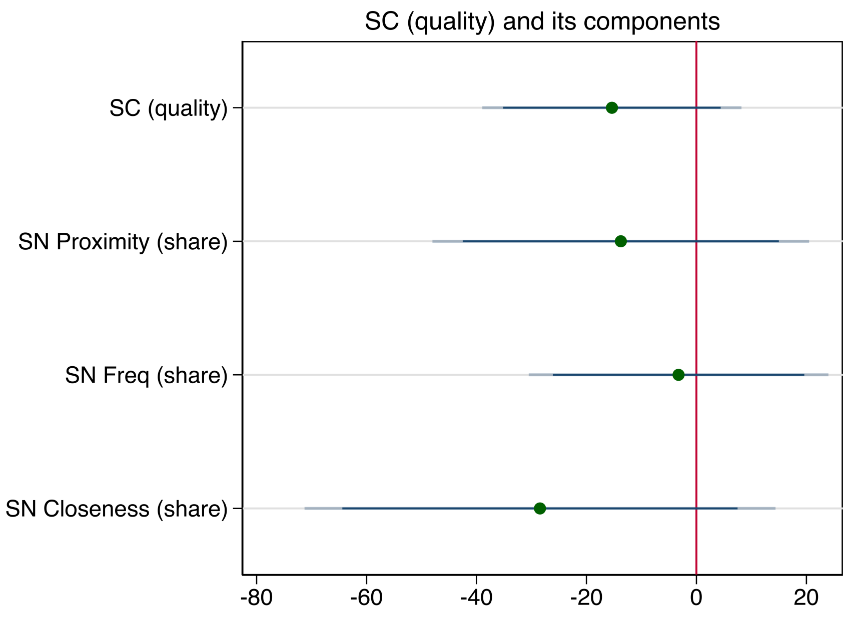

Fig. A.1. 2SLS estimates: SC (quality) and components, all countries

Notes: 2SLS estimates, all countries. Inner bars show 90% CI; outer lines show 95% CI. Components are the share of confidants who are geographically close, in frequent contact, and emotionally close.

Table A.6: First Stage estimates- Sample of countries with higher share of women informal caregivers- Dep. variable: Intense caregiving

	(1)	(2)	(3)	(4)
Parental health shock	0.148***	0.148***	0.148***	0.150***
	(0.028)	(0.028)	(0.028)	(0.025)
Basic controls	No	Yes	Yes	Yes
Siblings	No	No	Yes	Yes
Coresidence	No	No	No	Yes
F-Stat	28.5	28.8	28.9	35.8
R2	0.11	0.12	0.13	0.21
Observations	1317	1317	1317	1317

Notes: Standard errors in parentheses, clustered at individual level.
All regressions include country, year, cohort FE, and cohort X country FE. Dependent variable is a dummy for daily caregiving to mother. Basic controls: age, married, education, number of children and grandchildren, size of household. All covariates are lagged.
Countries selected include: Spain, Italy, Greece, Poland, Portugal, Croatia, Bulgaria, Malta, and Slovakia.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: 2SLS estimates: effect of intense caregiving on time allocation (robustness)

Panel A: All countries														
	Chores	Pers.care	Childcare	Help par.	Help part.	Help oth.	Leisure	Leis.part.	Admin	Paid work	Vol.work	Commute	Nap	Sleep
Intense caregiving to parent	-7.4	-0.2	-62.6	244.0*	9.7	-37.0	154.6	32.2	-22.6	-206.5	29.7	-18.5	6.4	-33.6
	(73.7)	(29.7)	(85.4)	(133.3)	(51.0)	(32.5)	(107.3)	(83.1)	(24.0)	(132.5)	(24.1)	(22.4)	(38.9)	(55.7)
F-stat	12.96	12.75	13.08	4.91	12.90	13.03	13.05	12.73	13.03	13.11	13.17	12.90	12.99	12.54
N	3,279	3,278	3,283	3,283	3,285	3,284	3,282	3,284	3,283	3,288	3,287	3,284	3,285	3,282
Panel B: High-IC countries														
Intense caregiving to parent	56.8	-5.6	-2.2	160.4**	-22.6	-8.0	2.8	10.3	-16.5	-152.0*	5.1	-22.4	20.4	-30.6
	(49.9)	(15.6)	(39.9)	(73.6)	(24.3)	(10.7)	(54.1)	(44.4)	(13.5)	(83.5)	(12.3)	(14.4)	(31.8)	(33.5)
F-stat	16.61	16.97	16.90	10.43	17.82	17.19	17.21	17.15	17.40	17.30	17.24	17.26	17.29	16.61
N	973	974	976	978	977	977	977	977	976	978	978	977	976	975
Panel C: All countries – controlling for lagged employment														
Intense caregiving to parent	-11.8	-1.4	-63.7	241.6*	9.4	-36.9	148.4	29.2	-21.6	-192.8	29.8	-16.3	3.8	-34.2
	(74.0)	(29.7)	(85.9)	(132.3)	(50.6)	(32.5)	(107.1)	(83.3)	(24.0)	(130.6)	(24.3)	(22.1)	(38.8)	(55.8)
F-stat	12.91	12.68	13.02	4.91	12.85	12.96	12.96	12.66	12.97	13.06	13.12	12.86	12.93	12.47
N	3,273	3,272	3,277	3,277	3,279	3,278	3,276	3,278	3,277	3,282	3,281	3,278	3,279	3,276
Panel D: High-IC countries – controlling for lagged employment														
Intense caregiving to parent	57.1	-5.1	-2.7	159.9**	-22.8	-8.1	1.8	11.0	-16.2	-149.8*	5.0	-21.9	20.8	-31.0
	(49.9)	(15.6)	(39.9)	(73.4)	(24.3)	(10.7)	(54.1)	(44.3)	(13.4)	(83.1)	(12.4)	(14.2)	(31.7)	(33.5)
F-stat	16.63	16.99	16.92	10.45	17.83	17.21	17.22	17.17	17.42	17.32	17.26	17.29	17.30	16.63
N	972	973	975	977	976	976	976	976	975	977	977	976	975	974

Note: 2SLS estimates of the effect of intense caregiving on time allocation (minutes). Instrument: indicator for parent in poor health at t, given the baseline non-poor restriction. Controls: education, marital status, number of children and grandchildren, household size, siblings, female siblings, coresidence, and the baseline outcome $Y_{i,0}$. All other covariates are measured at $t-1$. Panels C–D additionally control for lagged employment status. All regressions include country, year, cohort FE and country x cohort trend FE. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: Note: 2SLS estimates of the effect of intense caregiving on time allocation (minutes).
Instrument: parental health shock. Controls: education, marital status, number of children and grandchildren, household size, siblings, female siblings, coresidence, baseline outcome $Y_{i,0}$.
Panels C–D additionally control for lagged employment status.
All regressions include country, year, cohort FE and country x cohort trend FE. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

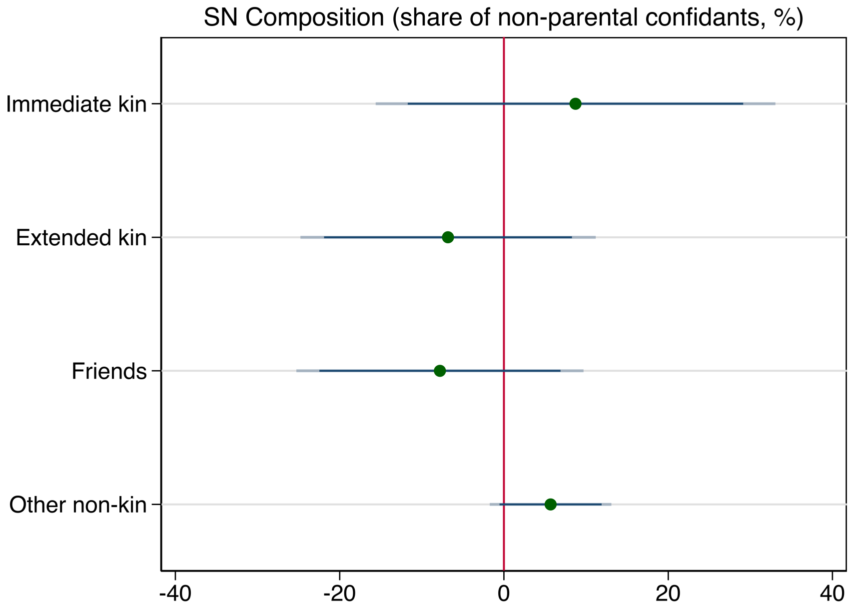

Fig. A.2. 2SLS estimates: network composition (shares), High-IC countries

Notes: 2SLS estimates for High-IC countries. Inner bars show 90% CI; outer lines show 95% CI. Outcomes are the *share* of the respondent's non-parental confidants who fall in each of four mutually exclusive relationship types: immediate kin (spouse and children), extended kin (siblings, in-laws and other relatives, excluding parents), friends, and other non-kin (colleagues, neighbours, and other contacts). Baseline shares (High-IC mean): immediate kin ≈ 73%, extended kin ≈ 13%, friends ≈ 12%, other non-kin ≈ 1.5%. Shares are used rather than counts to neutralise any composition effect driven mechanically by network size.

Table A.8: Control-set sensitivity of the main 2SLS estimates, High-IC countries

	(1) Y_t-1	(2) Y_0 (main)	(3) No SES lags	(4) No lag	(5) No co-res	(6) Y_t-1	(7) Y_0 (main)	(8) No SES lags	(9) No lag	(10) No co-res	(11) Y_t-1	(12) Y_0 (main)	(13) No SES lags	(14) No lag	(15) No co-res
care_intense_parent	-17.431*	-20.494**	-23.971**	-23.665**	-20.383**	-1.109**	-1.112**	-0.981*	-1.256**	-1.097**	-0.261	-0.244	-0.222	-0.287	-0.243
	(8.989)	(9.362)	(9.601)	(9.704)	(9.344)	(0.525)	(0.526)	(0.511)	(0.565)	(0.524)	(0.163)	(0.175)	(0.172)	(0.179)	(0.173)
F (K-P)	34.72	35.54	36.13	35.76	30.28	35.49	36.08	36.56	35.58	30.96	34.35	34.46	34.93	34.45	29.61
N	1317	1317	1390	1317	1317	1314	1315	1388	1315	1315	1315	1315	1388	1315	1315

Notes: 2SLS estimates of the effect of intense caregiving on the three primary outcomes in the High-IC sample. Columns 1–5 vary the outcome-control choice for SC (quality); columns 6–10 for SN satisfaction; columns 11–15 for any activity. Within each block, column 1 uses the contemporaneous lag $Y_{i,t-1}$; column 2 uses the baseline outcome $Y_{i,0}$ (the paper's main specification); column 3 drops the SES lags (income, wealth, rural); column 4 drops the outcome control entirely; column 5 drops co-residence. Fixed effects, cluster structure, and remaining covariates are as in Table 4. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Corroborating Evidence: Details

Descriptive evidence: associational activities and survival

We examine the long-run consequences of associational engagement for women aged 50 to 70 in Wave 4, following them until Wave 9. Unlike the main IV analysis, we do not restrict the sample to women with a surviving parent. We concentrate on the two outcomes most strongly affected by caregiving in the causal analysis: the number of associational activities and satisfaction with activities. To mitigate reverse causality, we measure these at baseline and then examine subsequent mortality hazards. Figure A.3 reports Kaplan–Meier survival curves by baseline participation category.

Women who do not participate in any activity at baseline experience substantially lower survival than those who take part in one, and the advantage is even larger for two or more activities. A log-rank test rejects equality of survivor functions at the 1% level ($\chi^2(2) = 121.1$, $p < 0.001$). Figure A.3(b) shows a similar gradient by satisfaction with activities. The separation of curves becomes more pronounced after age 70, consistent with early differences in social engagement generating differential mortality over time.

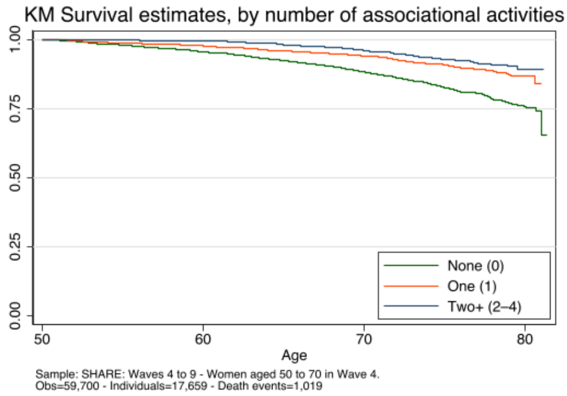

(a) By number of associational activities

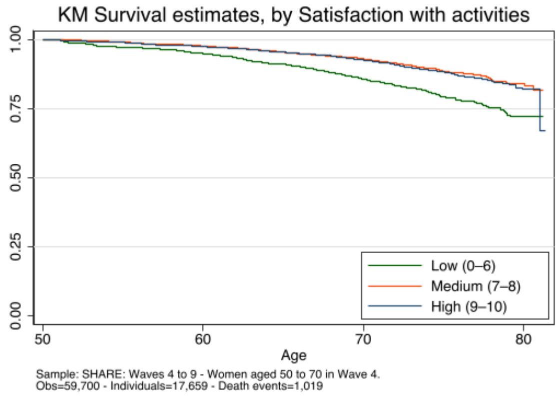

(b) By satisfaction with activities

Fig. A.3. Kaplan–Meier survival curves by associational activities at baseline (Wave 4).

Cox models with baseline health controls

To assess whether these associations survive adjustment for health, we estimate Cox proportional hazard models following Fawaz and Mira (2023), using the number of associational activities at Wave 4 as the main predictor. All models include frailty (Fried five-item phenotype),[9] self-assessed health, number of chronic diseases, functional limitations (ADLs and IADLs), cognitive recall, depressive symptoms (EURO-D), and socio-demographic controls (education, employment, income and wealth quartiles). Country and wave fixed effects are included through-

[9]Following Fried et al. (2001), the frailty index assigns one point for each of five criteria: (i) weakness (grip strength below the 20th percentile, stratified by gender and BMI); (ii) exhaustion (self-reported frequent fatigue); (iii) shrinking (involuntary weight loss or reduced appetite); (iv) slowness (difficulty walking 100 metres or climbing stairs without resting); and (v) low physical activity (infrequent moderate-energy activities such as walking or gardening). The index ranges from 0 to 5; scores of 3 or above correspond to clinical frailty.

out. To reduce reverse causality, the sample is restricted to respondents with at least 24 months of follow-up.

$$h_i(t \mid X_i) \;=\; h_0(t)\,\exp\!\Big(\beta_1 \cdot \text{Activities}_{i0} \;+\; \gamma'\text{Health}_{i0} \;+\; \delta' Z_{i0}\Big)$$

We report hazard ratios; values below one indicate lower mortality risk.

Table A.9: Cox models: Effect of Social Connectedness at Wave 4 on Mortality up to Wave 9.

	(1)	(2)	(3)	(4)	(5)	(6)
Participation in associational act.	0.699***	0.793***	0.880**		0.896*	0.782***
	(0.037)	(0.044)	(0.049)		(0.052)	(0.073)
Frailty index			1.320***	1.328***	1.308***	1.305***
			(0.052)	(0.052)	(0.055)	(0.055)
Number of chronic diseases			1.034	1.031	1.047*	1.047*
			(0.026)	(0.026)	(0.028)	(0.028)
Number of limitations			1.074***	1.073***	1.061**	1.061**
			(0.026)	(0.026)	(0.028)	(0.028)
Depression score (EuroD)			0.973	0.972	0.972	0.972
			(0.017)	(0.017)	(0.018)	(0.018)
Cognitive recall test (higher is good)			0.908***	0.904***	0.901***	0.902***
			(0.022)	(0.021)	(0.023)	(0.023)
Basic	yes	yes	yes	yes	yes	yes
Socio-Demo	no	yes	yes	yes	yes	yes
Health	no	no	yes	yes	yes	yes
Follow-up >24m	no	no	no	no	>24 mo	>24 mo
Binary	no	no	no	no	no	>0
Observations	52032	52032	52032	52032	47601	47601
Individuals	15267	15267	15267	15267	14740	14740

Exponentiated coefficients; Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Main explanatory variable is the Number of Associational Activities (0-4). Basic controls: age, age squared. Socio Demo controls: education categories, whether employed, whether has at least one child, income and wealth quartiles, living in a house (vs building), rural (versus urban). All regressions include wave and country-specific fixed effects. All covariates are fixed at Wave 4.

Each additional activity is associated with a 30% lower mortality hazard in the basic specification, falling to around 10–12% once the full battery of health controls is included – still significant. The association is concentrated at the extensive margin: women with at least one activity have markedly lower hazards than those with none, even with the follow-up restriction (column 6 of Table A.9).

Horse race: which network dimensions predict survival?

Table A.10 runs a horse race across all eight social network dimensions under the same full specification. The mortality-relevant dimension is participation in any associational activity (HR = 0.782, $p < 0.01$): baseline participants have hazard rates about 22% lower than non-participants, conditional on the full battery of health and socio-economic controls. Subjective network quality (HR = 0.998, insignificant) and network satisfaction (HR = 1.040, insignificant) show no association with survival. Network size shows a marginal protective association (HR = 0.955, $p < 0.10$); the remaining dimensions are indistinguishable from one. The mortality gradient runs through the behavioural margin of engagement, not through subjective quality or the structural size and frequency of ties.

Table A.10: Horse race: Cox proportional hazard models by social network dimension

	(1) SN Sat.	(2) SN Quality	(3) Closeness	(4) Any Activity	(5) Size	(6) Freq.	(7) Proximity	(8) SC Index
SN satisfaction	1.040							
	(0.028)							
SN Quality Index		0.998						
		(0.002)						
SN Closeness (share)			0.999					
			(0.001)					
Any associational activity				0.782***				
				(0.073)				
SN size (excl. parents)					0.955*			
					(0.025)			
SN Freq (share)						0.999		
						(0.001)		
SN Proximity (share)							1.000	
							(0.001)	
SC Index								0.934
								(0.054)
Observations	47,509	40,147	47,086	47,601	47,569	40,147	47,086	40,147
Individuals	14,710	12,365	14,584	14,740	14,732	12,365	14,584	12,365

Exponentiated coefficients. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: Hazard ratios from Cox proportional hazard models. Baseline covariates fixed at Wave 4 for women aged 50–70. Each column includes a single social network dimension as predictor. Full controls in all columns: frailty index, self-assessed health, number of chronic diseases, functional limitations (ADL), depressive symptoms (EURO-D), cognitive recall, education, employment, income and wealth quartiles, number of children, rurality, housing type, country and wave fixed effects. Sample restricted to individuals with more than 24 months of follow-up.