TIAA Institute

Research Dialogue | Issue no. 236
February 2026

The protective effects of a healthy spouse: Medicare as the family member of last resort

Abstract

We use novel Medicare data that link spouses to examine how one spouse's sudden incapacitation affects their partner's need for formal care. A spouse's health shock causes their partner to be 18% more likely to visit a skilled nursing facility. That pattern reflects both a change in health and a shift from informal care to formal care. After one spouse is incapacitated, the other spouse becomes less sensitive to the price of formal care. We explore the implications for optimal health insurance contracts, showing these within-household spillovers imply that the optimal health insurance contract should provide more generous coverage to those whose spouses are incapacitated relative to those whose spouses are available to provide care.

Itzik Fadlon
University of California
San Diego and NBER

Tal Gross
Boston University
and NBER

Alex Hoagland
University of Toronto

Tim Layton
University of Virginia
and NBER

1. Introduction

As populations age, long-term care becomes a central challenge. Projections for the United States and other high-income countries point to rapid growth in the share of adults who need help with the basic activities of daily living (Favreault and Johnson, 2021; OECD, 2024). Meeting that need has become a first-order issue for public budgets worldwide.

Long-term care is different from most other forms of health care because there exists a comparable alternative at home. Unlike, for example, surgery or diagnostic tests, help with the activities of daily living can be supplied informally by one's family. Long-term care thus comes in two forms: formal and informal. Formal care consists of nursing home stays, skilled nursing services, and paid home health aides. Informal care consists of the unpaid help of a spouse or child at home.

Little is known about informal care as it's difficult to measure and even more difficult to value. When providing care, family members don't clock in, track their hours, or log their effort. As a result, informal care is typically unobserved and unmeasured. This has led to a wide range of estimates of the value of informal care, from $86 billion (Gruber and McGarry, 2023) to more than $500 billion (Chari et al., 2015). The existence of informal care complicates standard public insurance design. The public sector must decide not only how much to pay for formal care but also when to pay, given that a private alternative—informal care—exists. Across the world, approaches vary. The Dutch system, for example, explicitly reduces a married beneficiary's eligibility for publicly funded care when a healthy spouse is present (Ilinca et al., 2017). As a result, the Dutch effectively raise the price of formal care when informal care is available. The long-term-care system in Germany does something close to the opposite—charging higher premiums to adults without children because they're more likely to need formal care (Rhee et al., 2015). America's Medicare program, in contrast, offers identical coverage to all beneficiaries regardless of access to informal care. Those contradictory rules signal genuine uncertainty regarding how public coverage of long-term care should take the availability of informal care into account.

This paper addresses that question in two steps. First, using novel data from the U.S. Medicare program that identifies couples, we provide causal evidence on the value of informal care from a healthy spouse. Using event-study regressions, we document that when one spouse suffers a sudden heart attack or stroke, their partner's probability of entering a skilled nursing facility (SNF) increases by 18%. The sudden loss of a caregiver thus shifts the demand curve for formal care. We show these effects are larger when the spouse who experiences the heart attack or stroke (the "shock spouse") is incapacitated for a longer period of time. We also show stronger effects when the other spouse (the "outcome spouse") is sicker and has more care needs. We decompose the increase in SNF stays into marginal stays due to declining health of the outcome spouse versus substitution away from informal care toward formal care, conditional on health. We find that around 90% of the overall effect is substitution. In addition, we study a particular feature of Medicare's reimbursement for formal care: a sharp rise in out-of-pocket costs after 20 days in a SNF. Beneficiaries with an able spouse at home are far more likely to exit on day 20, just before copayments increase, than those whose spouse was just incapacitated. That pattern suggests that the loss of informal care leads to a more price-inelastic demand for formal care—one that persists for some time after the initial health shock.

Second, we interpret those estimates with a simple insurance-design framework. We use our estimates of price sensitivity to show that those whose spouses are incapacitated get much more surplus from SNF stays than those who have healthy spouses at home. We then use our framework to show that because households without an able spouse face a steeper marginal-utility loss from leaving formal care and thus a lower demand-response to coverage, optimal coverage is more generous for them. Our estimated model implies that optimally tailoring coverage in this way would cut moral-hazard-driven deadweight loss by roughly two-thirds while holding government spending constant.

Our work contributes to three areas of inquiry. First and foremost, we contribute to research on elderly and long-term care in the United States, particularly informal care. Gruber and McGarry (2023) estimate that formal care expenditures totaled $280 billion in 2018, which were financed mostly by the public sector, and that informal care plays a critical role in providing for the elderly, showing that at every age and disability level they receive informal care more frequently than formal care. Studies have also found informal care may delay entry into formal care (Bergeot and Tenand, 2023; Bonsang, 2009; Charles and Sevak, 2005) and significantly affects the labor supply and earnings of caregivers, having effects on households that go beyond health (Maestas et al., 2024). In addition, Massner and Wikström (2024) document that lowering out-of-pocket costs for formal elderly care in Sweden not only improves seniors' health outcomes, but also improves adult children's labor-market outcomes. We contribute to this literature by providing, to the best of our knowledge, the first analysis of intra-family interactions in the demand for elderly care in the United States using administrative data at the population level.

Second, this paper contributes to a broad academic literature on the economics of risk sharing within the household. Since the pioneering work of Becker (1974), there has been significant progress in understanding the dynamics of risk-sharing within families, particularly how health shocks to one member of the family can affect the others.

This body of research has explored joint learning about health risks (Hoagland, 2025), the interdependence of family members' risks (Fadlon and Nielsen, 2019a), and the implications for public policy (Fadlon and Nielsen, 2019b; Gross et al., 2024). Additional work has highlighted the role of economic and health shocks on household labor supply and family stability (Fadlon and Nielsen, 2021; Hodor, 2021; Arteaga et al., 2024; Fontes et al., 2024). We contribute to this literature by providing novel evidence on the value of a healthy spouse via the intra-household interactions between formal and informal care among elderly American couples.

Third, this paper uncovers health care dynamics that can inform how health insurance contracts ought to be structured. When one person experiences a health shock, the risks their spouse faces change, ultimately affecting the willingness to pay for health care (Bauer et al., 2025). Since Zeckhauser (1970), we have known that the optimal contract provides more generous coverage when price sensitivity is lower. Our results reveal (as far as we know, for the first time) that price sensitivity is lower when one's spouse is incapacitated and unavailable to provide care, implying that optimal coverage should be higher for this group. And yet, the Medicare program offers individual health insurance plans that make no adjustment for the rest of the household. Many employer-provided plans also feature individual-specific rather than household-specific out-of-pocket limits or deductibles, despite evidence of household-specific cost-sharing features shaping consumption decisions (Anderson et al., 2024). In addition, other annuities—such as Social Security—take into account survivor benefits in their payment structure, while Medicare does not (Cottle Hunt and Caliendo, 2022; Brown and Poterba, 2000). This paper provides new evidence that can help assess the optimality of such arrangements. Our work is therefore related to previous work estimating optimal levels of social insurance (Baily, 1978; Chetty, 2006; Chetty and Finkelstein, 2013), where we provide novel findings with respect to the design of health care policies, which are large and growing.

The remainder of the paper proceeds as follows. Section 2 describes the data we analyze and the empirical framework we use. Section 3 presents empirical estimates in two pieces. First, it presents event-study estimates of the effects of a shock spouse's health shock on their partner's health and health care utilization. Second, the section assesses how a shock spouse's health shock affects outcome spouse's responses to the out-of-pocket price of formal care. Section 4 explores the implications of those estimates for policy and estimates the implied valuation of informal care. Finally, Section 5 concludes.

2. Empirical framework

2.1 Data

We study Medicare claims data that cover all types of care, including the utilization of skilled nursing facilities. We rely on a novel dataset of couples in Medicare using a custom crosswalk from the Centers for Medicare and Medicaid Services (CMS). The dataset includes beneficiaries' addresses as tracked by CMS. We then group beneficiaries who live at the same address and identify spouses by selecting residences with a maximum of two Medicare recipients of the opposite sex living at the same address. This allows us to study Medicare-covered health care among American couples with large-scale administrative data. Gross et al. (2024) validate this identification of spouses.

Our outcome measures are drawn from fee-for-service claims for traditional Medicare enrollees from 2010 to 2017. We use information on 100% of hospital inpatient stays and qualifying SNF stays using the Medicare Provider Analysis and Review (MedPAR) file, as well as beneficiary summary information and annual measures of utilization for each enrollee. In addition, we observe detailed outpatient, carrier claims, and Part D (pharmaceutical) utilization for a 20% sample of enrollees.

In our analysis, we require households to have full Medicare eligibility throughout the window of observation, and we assign spouses in the year prior to the index health event. Finally, following the work of Dobkin et al. (2018), we focus on precipitating health shocks for those who have not been admitted to a hospital in the two years preceding their index admission.

2.2 Research design

A macabre thought experiment that would identify the impacts of a spouse's health shock would randomly assign such shocks to couples. We take a quasi-experimental approach that aims to approximate that thought experiment. We restrict the analysis to couples who have experienced a health event—specifically, a heart attack or stroke—and identifying the treatment effect using variation in the timing of when the event was realized. We follow the work of Fadlon and Nielsen (2019a, 2021) and construct counterfactual outcomes for affected families using families who experience the same event but in the future.

We construct a treatment group, composed of spouses in households who experience a health shock in a particular calendar month, and a matched control group composed of spouses in households who experience a similar event, but 12 months later. We assign a placebo event to control households to match the calendar month of the event experienced by the treatment group. We then recover the treatment effect with a dynamic difference-in-differences estimator, identifying the shock's impact from the change in the differences in outcomes across the two groups over time. We can estimate the effects of the shocks for 12 months because the control group becomes "treated" 12 months later. The analysis is not subject to potential challenges involved in having units that switch in and out of experimental arms as posed by recent work (De Chaisemartin and d'Haultfoeuille, 2024). Similarly, the results are also robust to using estimators that explicitly model heterogeneous treatment effects, specifically the local-projections difference-in-differences estimator developed by Dube et al. (2023).

Table 1 presents summary statistics for the analysis sample. For this table and the subsequent analysis, we refer to the spouse who experiences the initial health event as the "shock spouse," and their partner as the "outcome spouse." Shock spouses are slightly older and more likely to have a chronic condition. Those differences are unsurprising—the older and less healthy spouse is the one more likely to experience the first health shock. Beyond those differences, the shock and outcome spouses are relatively similar.

We estimate the following event-study regression:

$$y_{i,t} = \alpha_i + \beta treat_i + \sum_{n\neq -1} \gamma_n \times I_n + \sum_{n\neq -1} \delta_n \times I_n \times treat_i + \lambda X_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where $y_{i,t}$ denotes an outcome for household i at time t, $treat_i$ denotes an indicator for whether a household belongs to the treatment group, and I_n are indicators for time relative to the assigned event period in weeks. The key parameters of interest are δ_n, which estimate the period n treatment effects ($n > 0$) relative to the pre-period $n = -1$. $X_{i,t}$ is a vector of controls that includes household fixed effects and calendar-time fixed effects.

TABLE 1. SUMMARY STATISTICS

	Shock spouses	Outcome spouses
Panel A. Demographics		
Age	76.95	75.75
Female	0.38	0.64
White	0.89	0.89
Black	0.05	0.05
Hispanic	0.01	0.01
FFS status	0.67	0.66
HMO status	0.32	0.32
Dual eligibility status	0.07	0.07
Panel B. Healthcare utilization		
Has a chronic condition	0.46	0.41
Predicted risk of SNF visit	0.00	0.02
Conditional # of SNF stays	1.11	1.35
Any inpatient admission	0.00	0.10
Conditional # of inpatient admissions	1.14	1.52
Total spending	$5,211	$8,339

Notes: This table presents summary statistics for the analytical sample N = 479,247. Shock spouses are those who experienced a first heart attack or stroke as discussed above; outcome spouses are their partners whose outcomes we study. Outcomes in Panel B are measured in the year prior to the index event. Total spending indicates the sum of Medicare and beneficiary annual payments.

This empirical strategy rests on the standard parallel-trends identifying assumption. It assumes that, absent the realization of the health shock, the outcomes of the treatment and control groups would evolve in parallel. To test the validity of this assumption, we study the evolution of the treatment and the control groups' outcomes in the periods prior to the event. In order to calculate average effects over several periods, we estimate the following difference-in-differences specification:

(2) $$y_{i,t} = \alpha_i + \beta treat_i + \gamma post_{i,t} + \delta \times treat_i \times post_{i,t} + \lambda X_{i,t} + \varepsilon_{i,t},$$

where $post_{i,t}$ is an indicator for whether the observation belongs to a period after the shock occurred, and δ measures the average treatment effect.

3. Empirical results

This section presents the paper's empirical results in two main steps. We first present reduced-form, event-study estimates of the effect of a spouse's health shock on their partner. Second, we explore an alternative estimation strategy, a regression-discontinuity design, based on the change in copayments after day 20 of a SNF stay.

3.1 Event-study analyses

Utilization of formal care. Figure 1 examines how an outcome spouse's risk of visiting a SNF evolves after the shock spouse's health event. This figure and the subsequent event-study figures plot the δ_n coefficients from estimates of equation (1). The horizontal axis plots time relative to the shock spouse's event in months, and the vertical axis plots the estimated coefficients. Average baseline levels in the six months prior to the event are reported in the subtitle of the plot, and we normalize the outcome variables such that plotted coefficients are presented in percent relative to baseline.

FIGURE 1. EFFECT OF A SPOUSE'S MAJOR CARDIOVASCULAR EVENT ON RISK OF A SNF VISIT

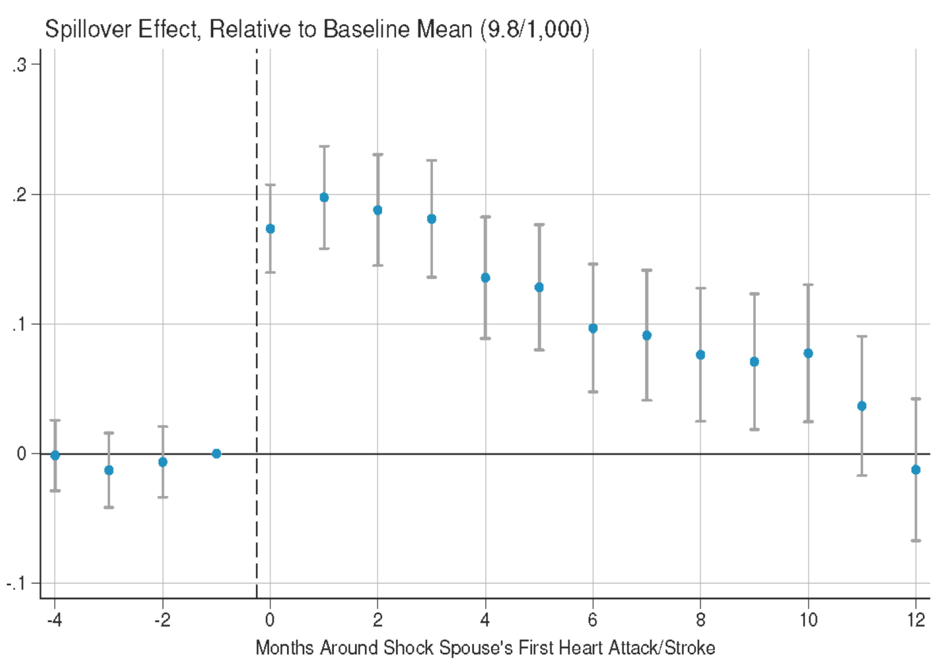

Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator for whether the outcome spouse visited a SNF. We rescale all coefficients such that they indicate the change relative to the initial baseline risk of hospitalization or SNF stay. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. The estimation includes calendar month fixed effects and person-specific fixed effects.

The figure suggests no systemic differentials across groups in the periods prior to the event, supporting the research design. We then observe an increase in the risk of an outcome spouse's visit to a SNF immediately following the index health shock. The risk of a SNF visit increases by 17.7% after their partner's cardiac event. The first column of Table 2 reports the effects for the month of the shock using equation (2).

A key factor driving potential heterogeneity underlying Figure 1 is the nature of the index health shock. As a first approach to explore that heterogeneity, we stratify the sample by whether the shock spouse died as a result of their health event. Outcome spouses who lose their partner experience a fundamental shift in the availability of informal care even after initial adjustments are made. And thus, that group offers a useful comparison in order to assess the long-term impacts of the initial event. Figure 2 presents the results of such a stratification. Outcome spouses who lose their partners face a more persistent and larger increase in the risk of visiting a SNF.

Figures 1 and 2 suggest large and sudden increases in the risk of a visit to a SNF once a person's partner experiences a health shock. Both figures suggest the effects are somewhat transitory, with a return to baseline after a year. This suggests that rather than permanently substituting formal care for informal care, outcome spouses may rely intermittently on formal care, particularly in the presence of potential health effects. We quantify these health and substitution effects in more detail below.

Taken together, these results suggest that one spouse's health shock increases their partner's reliance on formal care. When shock spouses do not recover, these effects are larger and somewhat more persistent.

FIGURE 2. HETEROGENEOUS EFFECTS OF SPOUSAL SHOCK, BY SHOCK FATALITY

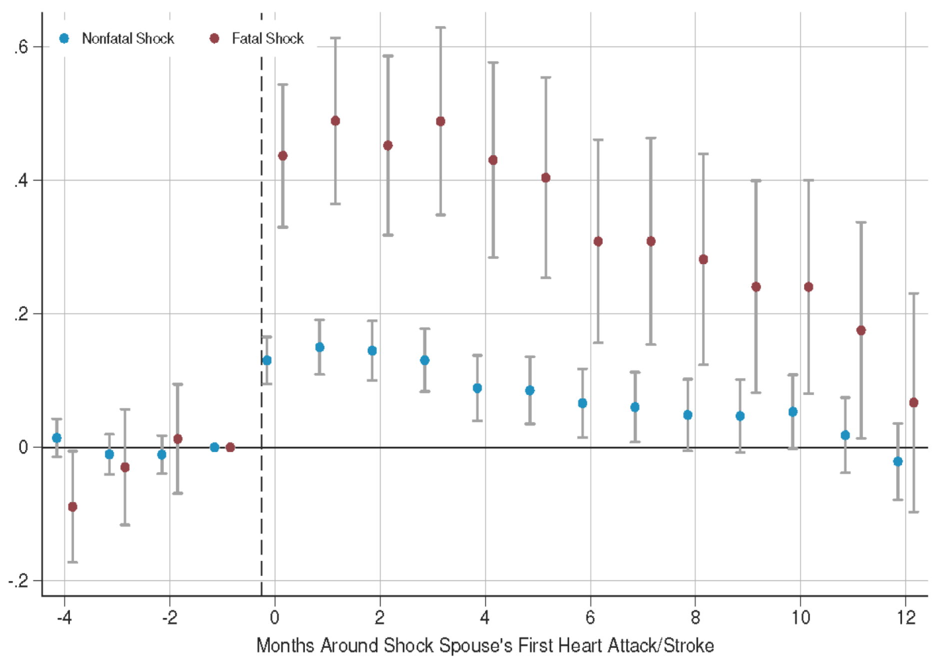

Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. Groups presented indicate whether a shock spouse passed away during the major cardiovascular event, measured within one year of discharge. The outcome of interest is an indicator for whether the outcome spouse visited a SNF. We rescale all coefficients such that they indicate the change relative to the initial baseline risk of hospitalization or SNF stay. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. The estimation includes calendar month fixed effects and person-specific fixed effects.

Health effect. We next evaluate the potential mechanisms driving the increase in formal care. We first assess the presence of a health effect. Motivated by prior work (for example, Arteaga et al., 2024; Bünnings et al., 2021), we test the hypothesis that adverse health shocks to one spouse can lead to a worsening of the other spouse's health and thus a greater need for care. Health effects may arise when those with preexisting conditions experience the short-term disruption associated with a spouse's health shock. That short-term disruption, in turn, may stem from emotional distress, a deterioration in sleep quality, or the transition to a new type of care provision at home.

Table 2 studies how a shock spouse's event affects the health of their partner across a range of possible outcomes that proxy for health: the risk of inpatient stays overall, inpatient stays specifically for falls, emergency department (ED) encounters, and a joint indicator variable for being hospitalized and then visiting a SNF. In each case, we report the treatment effect immediately following the event using equation (2).

Overall, outcome spouses exhibit an increased risk of all encounters, though the effects are more transient than their increased risk of a visit to a SNF. The effects are most pronounced in the month following the index event and persist for, at most, four months. The table suggests an increased risk of hospitalization of 8.9% in the first month and a much larger increased risk of a hospitalization for a fall: 78.5%. The risk of an ED visit increases by 23.3% in the month of their partner's health shock. In the month following their spouse's incapacitation, outcome spouses are 19.4% more likely to be hospitalized and then consume formal care at a SNF.

Table 2 offers an additional stratification in order to assess whether the increase in SNF visits is driven by the outcome spouse's diminished health. The second row of the table reports estimates of equation (2) when the sample is limited to health shocks in which the shock spouse passed away, similar to Figure 2. The third row of the table reports estimates when the index health shock was nonfatal. We observe larger increases in utilization for outcome spouses after fatal index events versus nonfatal index events, consistent with previous work (Fadlon and Nielsen, 2019a). The only exception here is for falls, which peak in month 0 and are more common among nonfatal spousal shocks.

TABLE 2. IMPACT OF SPOUSES' MAJOR CARDIOVASCULAR EVENTS ON PARTICULAR HEALTH CARE ENCOUNTERS

	(1) SNF	(2) Hospitalization	(3) Fall	(4) ED visits	(5) Hospitalization and SNF
Treatment effect,	0.177***	0.089***	0.785***	0.233***	0.194***
Month 0	(0.0176)	(0.0110)	(0.1999)	(0.0361)	(0.0248)
Fatal events only,	0.413***	0.243***	1.515	0.588**	0.594***
Month 0	(0.1146)	(0.0653)	(1.2240)	(0.2008)	(0.1579)
Nonfatal events only,	0.169***	0.084***	0.765***	0.221***	0.180***
Month 0	(0.0178)	(0.0111)	(0.2025)	(0.0368)	(0.0251)
Baseline rate/1,000	10.1	32.8	0.06	31.5	6.8
N	9,983,960	9,983,960	9,983,960	1,797,650	9,983,960

Notes: This table presents pooled difference-in-differences coefficients estimating the effect of a shock spouse's first heart attack or stroke on the outcome spouse's health outcomes (indicated in each column). "Hospitalization & SNF" indicates a spouse both was hospitalized and visited a SNF in the same month. ED visits are measured as total number of visits using the 20% sample of Medicare beneficiaries, hence the reduced sample size. All other outcomes are binary and use the 100% sample. Treatment effects are estimated in month 0, capturing the effect for the first four weeks post-event. We rescale all coefficients such that they indicate the change relative to the initial baseline risk of diagnosis in each category and cluster standard errors at the household level. Regressions include calendar month fixed effects and person-specific fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Substitution effect. We next examine evidence for a potential substitution effect. Outcome spouses with existing needs for care may switch from informal care to formal care once their partner is incapacitated. We provide two exercises to test for such an effect. First, we stratify households based on the outcome spouse's underlying need for care prior to the shock spouse's event. Second, we stratify households based on the degree of incapacitation of the shock spouse.

Figure 3 plots estimates of equation (1) for visits to SNFs separately by whether the outcome spouse suffered from a chronic condition before the shock spouse's event. We focus on chronic conditions most likely to require care from a spouse. In that sense, chronic conditions act as a proxy for a preexisting, underlying need for care. Outcome spouses with no underlying conditions exhibit a smaller—but still statistically significant—increase in SNF visits. By contrast, outcome spouses with chronic conditions exhibit a much larger increase in the risk of visit to a SNF. The effect for outcome spouses with chronic conditions is roughly four times as large as that for those with no conditions.

Next, we stratify by the severity of the initial health shock. We do so in order to capture the degree to which the shock spouse may still be able to provide informal care. We divide households based on whether the shock spouse was discharged home or to a medical facility, either a rehabilitation facility or a SNF. We limit this comparison to nonfatal index shocks because Table 2, above, demonstrates that spousal death is much more likely to induce health effects. Figure 4 suggests that the more severe the spouse's incapacitation, the more likely outcome spouses are to utilize formal care following their partner's event. Outcome spouses exhibit a nearly threefold increase over the pre-event baseline mean when the shock spouse is also discharged to formal care, compared to much smaller increases following a home discharge.

Together, Figures 3 and 4 indicate a substitution effect. We observe an increased risk of SNF stays when outcome spouses are less healthy and require care ex ante and also when the initial event leads to greater incapacitation of the shock spouse ex post.

FIGURE 3. EFFECT ON RISK OF A SNF VISIT BY OUTCOME SPOUSE'S PREEXISTING CHRONIC CONDITIONS

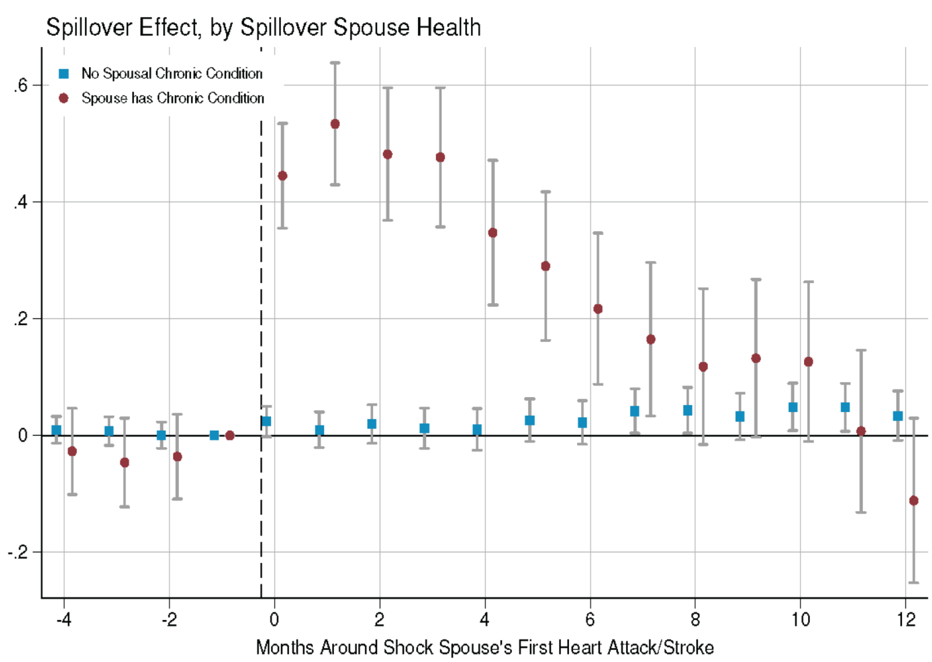

Notes: This figure plots estimates of the δ_n coefficients from equation (1), tracking the months since a shock spouse's first heart attack or stroke. The outcome is an indicator for whether the outcome spouse visited a SNF. We stratify the sample based on whether the outcome spouse had a documented chronic condition one year prior to the index spouse's event. Chronic conditions include ADRD, cancer (colorectal, lung, or prostate), chronic kidney disease, COPD, hip fractures, and strokes. We rescale coefficients to indicate changes relative to the initial baseline risk in each category. Error bars plot 95% confidence intervals based on standard errors clustered at the household level.

FIGURE 4. EFFECT ON RISK OF A SNF VISIT BY SEVERITY OF INDEX HEALTH EVENT

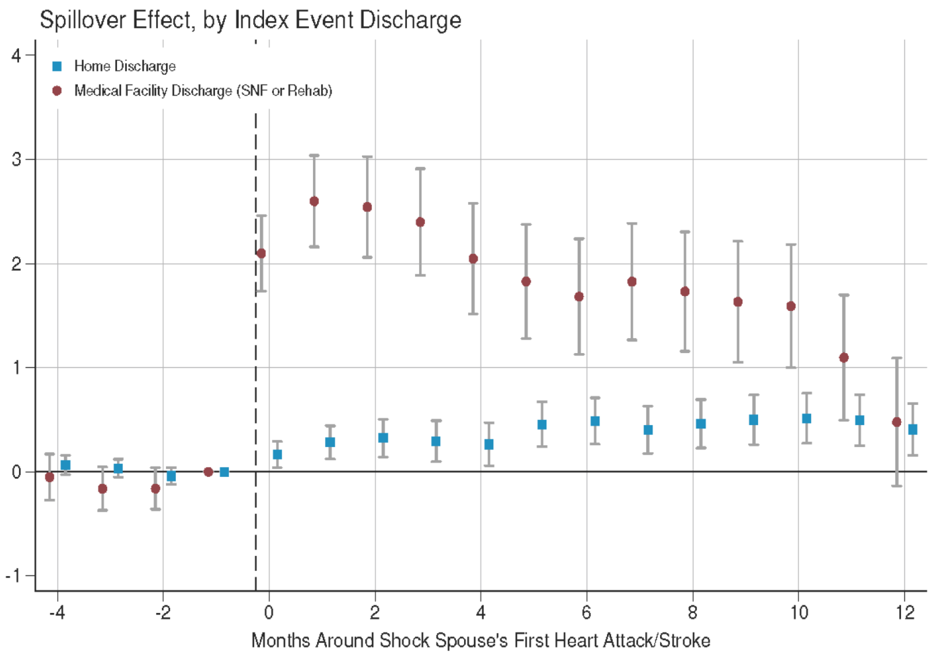

Notes: This figure plots estimates of the δ_n coefficients from equation (1), tracking the months since a shock spouse's first heart attack or stroke. Here, analysis is restricted to nonfatal index events only. We also exclude households whose shock spouses experienced a hospitalization in the 12 weeks prior to the index event to avoid bias from anticipatory SNF stays from preceding hospitalizations. The outcome of interest is an indicator for whether the outcome spouse visited a SNF during that week, where effects are split based on the discharge status of the index event. We rescale all coefficients such that they indicate the change relative to the initial baseline risk of diagnosis in each category. The error bars plot 95% confidence intervals based on standard errors clustered at the household level.

Comparison of health and substitution effects. The results above provide evidence in support of both a health effect and a substitution effect. We next explore a decomposition in order to characterize how much of outcome spouses' increased SNF visits is driven by substitution into formal care versus a deterioration in their health.

For all hospitalizations, we predict the probability that the patient will be discharged to a SNF. We then evaluate how those predicted probabilities evolve around their partner's health shock. That is, we generate a "predicted SNF stay" variable equal to zero for anyone who didn't have a hospitalization in the given time period and, for those who were hospitalized, is equal to the probability of a hospitalization with similar observable characteristics ending in discharge to a SNF. The prediction captures how we would expect outcome spouse's SNF use to change after the shock spouse's health shock based solely on preexisting characteristics. We then compare the actual change in the risk of a SNF visit after the shock spouse's event to the predicted change.

We rely on the following procedure in order to predict visits to SNFs. First, we randomly segment all hospitalizations into a 90% estimation sample and a 10% validation sample. We then run a Lasso regression of a binary indicator variable for any SNF discharge on the patient's age, sex, race, and the diagnosis-related groups (DRGs) for the hospitalization (Ellis et al., 2022). We rely on a Lasso regression given the large dimensionality of our covariate space (Abadie and Kasy, 2019).

FIGURE 5. CHANGE IN OUTCOME SPOUSE'S PREDICTED PROBABILITY OF SNF DISCHARGE

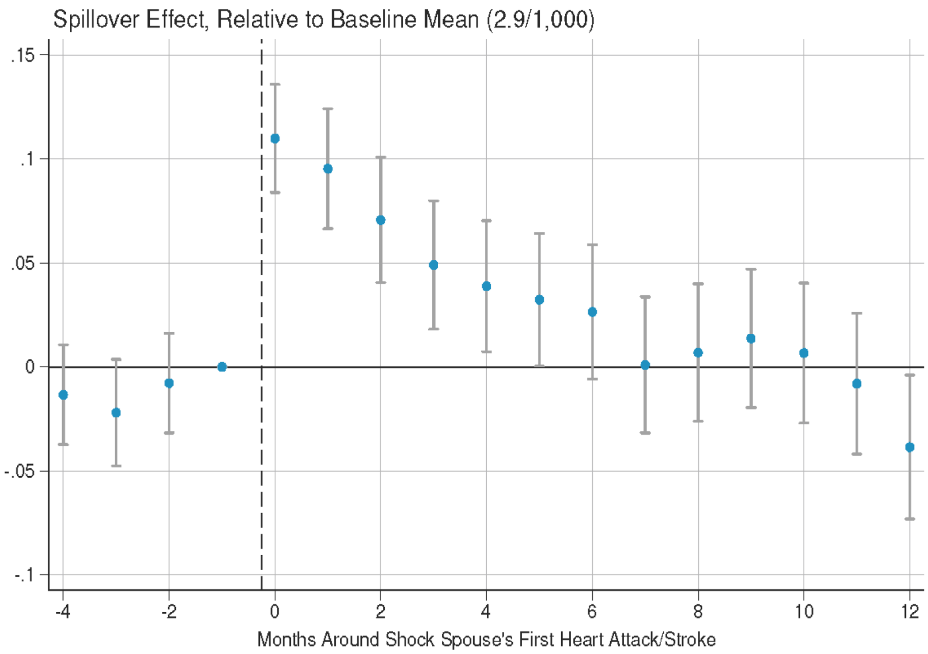

Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a spouse's first heart attack or stroke. The outcome of interest is the predicted probability of a SNF discharge based on hospitalization and patient characteristics using LASSO regression. We rescale all coefficients such that they indicate the change relative to the pretreatment average predicted probabilities. The error bars plot 95% confidence intervals based on standard errors that have been clustered on each Medicare household. The regression includes calendar month fixed effects and person-specific fixed effects.

Figure 5 presents estimates of equation (1) when the outcome of interest is the outcome spouse's predicted probability of visiting a SNF. In the month following the index shock, an outcome spouse's predicted probability of a SNF stay increases by 12%, up from a baseline risk of 0.35%. This estimated increase (0.35 stays per 1,000) constitutes roughly 18% of the 1.92 stays per 1,000 estimated overall. We would thus expect around 18% of the marginal SNF stays to have occurred solely due to the deterioration of the outcome spouse's health. The remaining 82% of SNF stays may arise solely from the substitution of formal care for informal care. In addition, in the long run, health effects dissipate, indicating that the bulk of observed effects are attributable to substitution.

That exercise, of course, assumes the only pathway from worsening health to a SNF is via hospitalization. We thus view this exercise as suggestive—though not conclusive—evidence that a majority of the rise in SNF stays we catalog is driven by a substitution from informal to formal care. More generally, it's not possible to disentangle the substitution and health effects in this setting. Instead, the goal is to offer one estimate of the degree to which observable health deterioration can account for the shift to formal care.

New diagnoses. The results above document an abrupt transition in the lives of outcome spouses, especially a transition into formal care. Those transitions may prompt a reassessment of the outcome spouse's health needs, including the detection of previously undiagnosed conditions. Such new diagnoses are more likely for medical conditions that develop gradually over time, making them unrelated to the immediate health event. New diagnoses are also more likely when the conditions are harder to detect due to the absence of clear objective tests and measures.

Two main classes of diseases among the Medicare population, which often require the type of elderly care that we consider, are Alzheimer's Disease and Related Dementias (ADRD) and cancer. Both ADRD and cancer are serious conditions that develop over time, so we wouldn't expect a shock spouse's health event to immediately result in these conditions. In addition, the two conditions substantially differ in the degree to which they can be directly and accurately diagnosed. Diagnosis of ADRD is notoriously difficult and requires multiple mental and physical tests. Cancer, in contrast, can be accurately diagnosed via biopsy (Chandra et al., 2023).

Table 3 reports estimates of equation (2) when the outcomes of interest are new diagnoses of ADRD and cancer. The likelihood of an outcome spouse being diagnosed with ADRD more than doubles immediately following a shock spouse's health shock and increases by an average of 43% over the first six months after the shock. By contrast, we observe no statistically significant effect on cancer diagnoses, despite these being nearly twice as prevalent at baseline. These patterns suggest that outcome spouses may be living with latent, undiagnosed conditions such as ADRD, whose detection is triggered by the changes in care following the shock spouse's event.

TABLE 3. IMPACT OF SPOUSES' MAJOR CARDIOVASCULAR EVENTS ON PARTNERS' NEW DIAGNOSES

	New ADRD diagnosis		New cancer diagnosis
	All	Excluding SNF diagnoses	All
Month of index event	0.904**	0.785***	-0.080
	(0.1805)	(0.1756)	(0.1071)
12 months since index event	0.262**	0.253**	0.056
	(0.0919)	(0.0901)	(0.0628)
Baseline rate per 1,000	1.08	1.05	3.08
N	1,785,930	1,785,930	1,564,790

Notes: This table presents pooled difference-in-differences coefficients estimating the effect of an outcome spouse's first heart attack or stroke on spousal diagnoses. The outcomes of interest are indicators for whether the focal spouse received a new diagnosis for ADRD or cancer. We exclude any new ADRD diagnosis arising from a SNF in column (2). Sample is restricted to beneficiaries consistently observed without an ADRD or cancer diagnosis during the first year of the analytical sample to correctly identify new diagnoses (which leads to the change in the sample size across columns). We rescale all coefficients such that they indicate the change relative to the initial baseline risk of a diagnosis in each category, and we cluster standard errors at the household level. Regressions include calendar month fixed effects and person-specific fixed effects. *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$

The fiscal value of a healthy spouse. The event-study estimates above demonstrate increases in outcome spouses' health care utilization: SNF admissions, hospitalizations, and ED visits. Those results naturally lead to the question of how much the shock spouse's event increases their partner's overall health care expenditures.

Table 4 presents estimates of regressions that track the impact of a shock spouse's event on overall Medicare-covered spending for the outcome spouse. In order to study spending, we rely on the annual summary measures tracked for all fee-for-service Medicare recipients. In the year of a shock spouse's event, total Medicare payments increase by roughly 5% ($317) and by an average of 12% ($773) for fatal index events. The table suggests increases in health care expenditure across categories that are large in both absolute and relative terms. For instance, spending on SNFs increases by $112 in the year of the event, roughly a 33% increase over the baseline mean. We observe similarly large increases in the use of home health and hospice care, as well as across various measures of inpatient hospitalizations. Those effects are more than doubled when the shock spouse passes away.

Some of the estimates in Table 4 are negative, indicating that some types of care may crowd out other less urgent expenditures. For instance, we observe a 2.3% decrease in hospital outpatient care and an 8.0% decrease in elective surgeries. We also observe declines in spending on prescription medication (2.4%). The decline in prescription medicine may be mechanical: Patients receive medications directly in a SNF, and those medications are unreported in claims data.

One can view the estimates in Table 4 as describing the spillover effects of an index health shock on the overall Medicare budget. In aggregate, this burden is substantial. Consider that 65.7 million Americans are covered by Medicare, 58% of whom are married or living with a partner. The incidence rate of heart attacks and strokes is roughly 5% annually, 15% of which are fatal within 30 days. The estimates in Table 4 thus translate into a national spillover effect of $734.3 million annually.

TABLE 4. EFFECT OF SPOUSE'S MAJOR CARDIOVASCULAR EVENT ON ANNUAL MEDICARE SPENDING

Spending measure	Pretreatment average	(1) All events	(2) Fatal events
Panel A. Total spending			
	\$6,433	\$253***	\$584**
		(36.92)	(236.11)
Panel B. Long-term care			
SNF	\$338	\$95***	\$308***
		(8.43)	(61.67)
Home health	\$283	\$24***	\$30
		(3.64)	(24.14)
Hospice	\$125	\$49***	\$69*
		(5.58)	(37.70)
Panel C. Hospital and surgical care			
Acute inpatient	\$1,457	\$123***	\$236**
		(15.73)	(102.03)
Other inpatient	\$177	\$20***	\$34
		(6.59)	(40.83)
Hospital outpatient	\$985	-\$17**	-\$14
		(7.48)	(40.71)
Ambulatory surgical center	\$76	-\$6***	-\$6
		(1.14)	(6.21)
Panel D. Other care			
Physician payments	\$337	-\$10***	\$2
		(0.75)	(4.65)
Evaluation and management	\$243	\$25***	\$49***
		(2.18)	(13.98)
Part D spending	\$896	-\$19***	-\$53**
		(4.84)	(25.99)

Notes: This table plots estimates of pooled post-treatment effects tracking the year following an index event's first heart attack or stroke. The outcome of interest is annual Medicare payments per beneficiary, across the spending categories indicated in each row. Regressions include calendar month fixed effects and person-specific fixed effects. Column (1) reports the estimate for all events, while column (2) reports the estimate only for fatal index events. Estimates are scaled by the time remaining in the year following the index event, and standard errors are clustered at the household level.

3.2 Sensitivity to the price of formal care

This section studies how a spouse's incapacitation affects their partner's price elasticity of demand for formal care. We rely on a regression-discontinuity design that exploits a sudden change in cost-sharing under the Medicare program.

When Medicare recipients visit a SNF after a hospitalization, they initially face zero out-of-pocket costs for the visit. Then, after the 20th day of their stay, they face copayments equal to 20% of the Medicare reimbursement. We study responses to these discontinuous changes in a regression discontinuity (RD) framework, and we recover the difference in the RD parameter across the shock spouse's health shock via a difference-in-discontinuities design (Grembi et al., 2016).

Figure 6 illustrates the variation we exploit. We plot mean residuals from a regression of discharge from the SNF on day-of-week-specific fixed effects and a linear trend for days in the SNF so far. We then calculate deviations in length of stay relative to this trend around the discontinuous change in cost-sharing at day 20. First and foremost, the figure demonstrates that Medicare recipients are much less likely to stay in the SNF after cost-sharing increases on day 20. More importantly, Medicare recipients with a healthy spouse are more likely to leave the SNF after day 20 than those whose shock spouse has just experienced a major health event.

In order to measure how shock spouses' health crises affect responses to cost-sharing, we estimate an RD design. Consider the following specification:

(3) $$y_{i,t} = \delta_0 + \delta_1 D_{i,t} + S_{i,t}(\gamma_0 + \gamma_1 D_{i,t}) + T_{i,t}[\alpha_0 + \alpha_1 D_{i,t} + S_{i,t}(\beta_0 + \beta_1 D_{i,t})] + \varepsilon_{i,t},$$

where $y_{i,t}$ represents an indicator for whether outcome spouse i resides in a SNF on day t, $D_{i,t}$ represents the recentered running variable of SNF days relative to the 21-day cutoff, $S_{i,t}$ indicates a post-cutoff dummy, and $T_{i,t}$ indicates a post-treatment dummy of a shock spouse's health shock. We estimate equation (3) for SNF stays of duration within a certain bandwidth of the cutoff and use a triangular kernel to weight observations more heavily if they are closer to the cutoff. The parameter β_0 indicates the difference in responses across the experimental groups.

We estimate equation (3) using observations within four months of both the treated and placebo index events and adjust for Medicare couple and time-of-event fixed effects. Table 5 presents the results. Prior to the shock spouse's event, the discontinuous change in Medicare payment for SNF stays induces a significant negative effect on the length of an outcome spouse's stay, decreasing the probability they'll remain in the SNF on the first day they face copayments by 1.8 percentage points. In contrast, the outcome spouse's response to copayments is reduced to 1.0 percentage points when they've been admitted following the shock spouse's event. Comparing the differences in the estimated responses, we find households affected by a shock are 0.8 percentage points more likely to remain in care across this threshold, absorbing roughly 45% of the pretreatment price responses. The results further suggest that these differences are almost entirely by outcome spouses with chronic conditions (see Figure 3).

FIGURE 6. RESPONSES TO COST-SHARING FOR SNF STAYS BY HEALTH EVENT

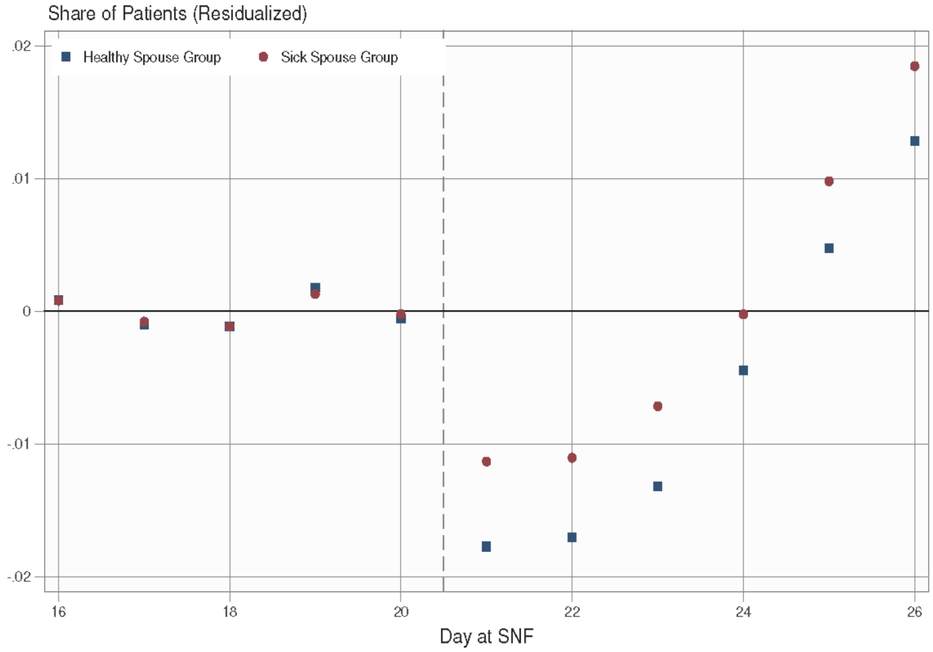

Notes: This figure plots the intuition and identifying variation for the regression discontinuity design (equation 3). We plot coefficients for the residual average share of patients in a SNF on a given day, adjusting for the day of the week and the number of days covered by Medicare previously. The vertical line indicates the day after which 20% coinsurance is applied to each additional SNF day.

TABLE 5. EFFECTS OF SHOCK SPOUSE HEALTH EVENT ON OUTCOME SPOUSE'S PRICE SENSITIVITY FOR SNF STAY DECISION, BY OUTCOME SPOUSE CHRONIC CONDITION STATUS

	(1) Pooled	(2) No chronic condition	(3) Has chronic condition
τ_{pre}	-0.018***	-0.019***	-0.018***
	(0.0030)	(0.0045)	(0.0038)
τ_{post}	-0.010***	-0.017***	-0.007**
	(0.0026)	(0.0041)	(0.0033)
	(0.0041)	-0.007**	$30
$\beta_{D\text{-}Disc}$	0.008**	0.001	0.011**
	(0.0040)	(0.0062)	(0.0051)
Couple FEs	✓	✓	✓
Time FEs	✓	✓	✓
Bandwidth	11.22	11.22	11.22
N	380,673	139,702	240,971

Notes: This table presents regression-discontinuity and differences-in-discontinuities estimators identifying the effect of ending Medicare coverage for SNF stays, which ends on day 21 for qualifying stays. The first two rows present the estimated effect of losing coverage on the probability an outcome spouse will remain in the SNF, stratified by before or after the index event. The third row presents the difference-in-discontinuities estimator as discussed in the text. SNF stays within four months of the treatment, and placebo events are included in the regression. Columns are stratified based on whether the outcome spouse has a chronic condition in claims (63% of the sample) or not in the year prior to the index event.

These results suggest that having a partner at home is an important determinant of the price elasticity of demand for formal care. In large part, this is likely driven by the availability of informal care, as discussed above. Additionally, however, healthy spouses may serve as advocates for those residing in SNFs, facilitating more rapid discharges for an individual to avoid paying unnecessary costs.

4. Implications

This section explores the implications of the empirical results. We first develop a theoretical model of the optimal cost-sharing for formal care. We then estimate the demand curve for formal care to explore alternative policies that would shape access to formal care.

4.1 Theoretical model

Consider a Medicare recipient who is already a resident of a SNF, having resided there for 20 days, and is deciding whether or not to stay for a 21st day, after copayments are required. The recipient has a healthy spouse in time period $t = 0$ and an incapacitated spouse in time period $t = 1$.

We index Medicare recipients by their valuation of formal care, $s_t \sim U(0,1)$. A recipient with $s_t = 0$ has the highest valuation for formal care at time t, and a recipient with $s_t = 1$ has the lowest valuation of formal care.

Medicare charges a price p for an additional night in the SNF after day 20. Given p, there is a marginal type of recipient s_0^* who has a healthy spouse and is indifferent between staying in the SNF and leaving it. Likewise, there is a marginal recipient with type s_1^* who has an incapacitated spouse and is similarly indifferent.

Define $V_t(s_t)$ as the incremental value of the 21st day in the SNF for a recipient of type s_t. A recipient and their household who remain an additional night in the SNF enjoy flow utility $W_t(s_t,p) = V_t(s_t) + U_t(Y_t - p)$, where U_t is the state-dependent household utility from nonmedical consumption and Y_t is disposable income. Disposable income may itself depend on family structure—let Y_0 and Y_1 denote the household's income when the spouse is healthy and incapacitated, respectively. The marginal type $s_t^*(p)$ is therefore determined by the indifference condition:

(4) $$V_t(s_t^*(p)) = U_t(Y_t) - U_t(Y_t - p),$$

so that households with $s_t < s_t^*(p)$ have a higher valuation of formal care and so purchase the 21st day and those with $s_t > s_t^*(p)$ leave the facility.

The indifference condition, equation (4), is a function of p. A first-order Taylor approximation of (4) around $p = 0$ yields:

(5) $$V_t(s_t^*(p)) \cong U_t'(Y_t) \times p.$$

This then provides the incremental value of an additional night in a SNF. Then define $v_t(s_t) \equiv V_t(s_t)/U_t'(Y_t)$ so that $v_t(s_t^*) = p$. This condition implies that, for the marginal household, the copayment of an additional night in a SNF is just equal to the dollar-valued benefit. We can then define the consumer surplus of all recipients who choose to stay in the facility as:

(6) $$\int_0^{s_t^*(P_t)} [v_t(s_t) - p] ds_t.$$

We next turn to the social planner's problem. Instead of a single price for a night in the SNF, p, suppose that the price can vary by family structure. In that sense, the social planner sets P_t rather than just p. Let μ_t be the population share in state t, with $\mu_0 + \mu_1 = 1$, c be the resource cost to Medicare of one day of SNF care, and R the program budget for the cohort. A social planner solves the following problem:

(7) $$\max \mu_0 \int_0^{s_0^*(P_0)} [V_0(s_0) + U_0(I_0)] ds_0 + \mu_1 \int_0^{s_1^*(P_1)} [V_1(s_1) + U_1(I_1)] ds_1$$

subject to the expected costs constraint:

$$\mu_0(c - P_0)\int_0^{s_0^*(P_0)} 1\, ds_0 + \mu_1(c - P_1)\int_0^{s_1^*(P_1)} 1\, ds_1 \leq R. \tag{8}$$

Taking $-P_t$ as the generosity variable and applying the envelope theorem yields the first-order condition:

$$\mu_t F_t\big(s_t^*(P_t)\big)\big(U_t'(I_t) - \lambda\big) = \lambda\mu_t(c - P_t)\left|\frac{dF_t\big(s_t^*(P_t)\big)}{dP_t}\right|, \tag{9}$$

where λ is the shadow price of public funds, and $I_t = Y_t - P_t$ is net income. This is the classic tradeoff between redistribution/insurance and efficiency. The left-hand side captures that fact that we want to transfer until marginal utility among the subgroup equals the economy-wide shadow price of government funds. The right-hand side captures the efficiency loss from fiscal externality (via effect on demand).

Next, consider a balanced-budget perturbation that lowers the sick-state price P_1 by $dP_t < 0$ and raises the healthy-state price P_0 by just enough to keep expected outlays fixed. Totally differentiating the constraint yields:

$$\frac{\partial P_0}{\partial P_1} = -\frac{\mu_1 F_1\big(s_1^*(P_1)\big)}{\mu_0 F_0\big(s_1^*(P_1)\big)}, \tag{10}$$

So that the impact on welfare is:

$$\frac{dSW}{d(-P_1)}\frac{1}{\mu_1 F_1\big(s_1^*(P_1)\big)} = \big(U_1'(I_1) - U_0'(I_0)\big) - \lambda\left(\frac{c - P_1}{P_1}|\varepsilon_1| - \frac{c - P_0}{P_0}|\varepsilon_0|\right), \tag{11}$$

where $\varepsilon_t \equiv \frac{P_t}{F_t\left(s_t^*(P_t)\right)}\frac{dF_t\left(s_t^*(P_t)\right)}{dP_t}$.

4.2 Implications for optimal copayments

The model above leads to two implications regarding optimal cost-sharing for the 21st day in a SNF. First, equal prices cannot be optimal. Evaluate (11) at a symmetric policy $P_0 = P_1$. Under the standard assumption that marginal utility is weakly higher when the spouse is sick, $U_1' \geq U_0'$, and that demand is *less* price-elastic when informal care is unavailable ($\varepsilon_1 < \varepsilon_0$), both terms on the right-hand side are positive. Hence $dSW/d(-P_1) > 0$. That is, a small shift of subsidy toward the incapacitated-spouse state raises welfare while leaving the budget unchanged. Equal copayments therefore cannot satisfy the social planner's first-order conditions whenever the empirical pattern of lower elasticities for sick households holds.

Second, the model leads to an optimal subsidy rule. Setting expression (11) to zero delivers the necessary condition for the optimal prices:

$$\big(U_1'(I_1) - U_0'(I_0)\big) = \lambda\left(\frac{c - P_1}{P_1}|\varepsilon_1| - \frac{c - P_0}{P_0}|\varepsilon_0|\right) = 0$$

Define $\rho_t \equiv \frac{c-P_t}{P_t}$, the inverse markup that measures the degree of subsidization with $\rho_t \to \infty$ as $P_t \to 0$. If health shocks are fully insured outside the copayment so that $U_1'(I_1)$ - $U_0'(I_0)$, then the optimal-pricing rule simplifies to

$$\frac{\rho_1}{\rho_0} = \frac{|\varepsilon_0|}{|\varepsilon_1|}. \tag{12}$$

Thus, the state with the lower price elasticity (the sick-spouse state) should receive the higher subsidy in proportion to the elasticity gap. Intuitively, when demand responds less to price, the efficiency cost of raising utilization through a subsidy is smaller, so the planner shifts resources toward that group until the marginal insurance benefit is equalized.

The RD estimates provide estimates of ε_0 and ε_1, where we find an optimal-pricing ratio of 1.80. In words, the degree of subsidization should be 80% higher for those with sick spouses than those with healthy spouses. This implies that for every \$100 Medicare pays for recipients with spouses, Medicare ought to pay \$180 for beneficiaries without healthy spouses.

5. Conclusions

We study couples enrolled in the Medicare program and find that when one spouse experiences a severe cardiac event, their partner's health care utilization increases by roughly 5%. Much of the increase is driven by a shift from informal care to formal care that must be covered by Medicare. This amounts to a fiscal externality of the index event. In aggregate, Medicare pays $734.3 million in care driven by that externality each year. More importantly, we estimate that people are willing to pay roughly four times more for formal care when their spouses become incapacitated.

These results demonstrate the importance of the public provision of formal elderly care when spouses cannot provide informal care. One policy response is a family-specific deductible. Such a deductible would effectively provide formal care at lower prices to households with an incapacitated spouse. As a result, family-specific deductibles would lower copayments for those whose demand curves for formal care have become more inelastic, and thus whose consumption is less driven by moral hazard.

Some countries do indeed adjust support for formal care based on the availability of family members who may provide informal care. The Dutch system, for example, explicitly limits eligibility for formal care based on whether an able spouse is present (Ilinca et al., 2017). Such an approach is consistent with the results in this paper. The Dutch system effectively raises prices for those with healthy spouses—those who exhibit an elastic demand for formal care—and lowers prices for those who lack them—those who exhibit an inelastic demand for formal care. Interestingly, other countries impose rules that go in the opposite direction. In 2005, the German long-term-care insurance system introduced a higher premium for childless adults (Rhee et al., 2015). Legislators justified such a policy on the grounds that those without children are more likely to need formal care. But the model and findings above suggests the opposite logic. Meanwhile, Medicare, by contrast, makes no adjustment for family structure in determining cost-sharing, eligibility, or coverage.

Finally, our findings also speak to the long-term secular trends in the aging of the population. Gerontologists sometimes refer to elderly patients who lack a family member to care for them as "elder orphans." As the population ages, elder orphans are predicted to become more prevalent (Roofeh et al., 2020). And, in turn, the increasing number of elder orphans may translate to greater need for more resources to be devoted to formal care.

References

Abadie, Alberto, and Maximilian Kasy. "Choosing Among Regularized Estimators in Empirical Economics: The Risk of Machine Learning." *Review of Economics and Statistics,* vol. 101, no. 5, 2019, pp. 743–762. https://doi.org/10.1162/rest_a_00812

Anderson, David M., Alex Hoagland, and Ed Zhu. "Medical Bill Shock and Imperfect Moral Hazard." *Journal of Public Economics*, vol. 236, 2024, 105152. https://doi.org/10.1016/j.jpubeco.2024.105152

Arteaga, Carolina, Natalia Vigezzi, and Pilar Garcia-Gomez. "In Sickness and in Health: The Broad Impact of Spousal Health Shocks," 2024. Working Paper.

Baily, Martin Neil. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics*, vol. 10, no. 3, 1978, pp. 379–402. https://doi.org/10.1016/0047-2727(78)90053-1

Bauer, Daniel, Darius Lakdawalla, and Julian Reif. "Health Risk and the Value of Life." *Journal of Public Economics*, vol. 245, 2025, 105346. https://doi.org/10.1016/j.jpubeco.2025.105346

Becker, Gary S. "A Theory of Marriage: Part II." *Journal of Political Economy*, vol. 82, no. 2, 1974, pp. S11–S26. https://www.jstor.org/stable/1829987

Bergeot, Julien, and Marianne Tenand. "Does Informal Care Delay Nursing Home Entry? Evidence from Dutch Linked Survey and Administrative Data." Journal of Health Economics, vol. 92, 2023, 102831. https://doi.org/10.1016/j.jhealeco.2023.102831

Bonsang, Eric. "Does Informal Care from Children to Their Elderly Parents Substitute for Formal Care in Europe?" *Journal of Health Economics*, vol. 28, no. 1, 2009, pp. 143–154. https://doi.org/10.1016/j.jhealeco.2008.09.002

Brown, Jeffrey R., and James M. Poterba. "Joint Life Annuities and Annuity Demand by Married Couples." *Journal of Risk and Insurance*, vol. 67, no. 4, 2000, pp. 527–553. https://doi.org/10.2307/253849

Bünnings, Christian, Lucas Hafner, Simon Reif, and Harald Tauchmann. "In Sickness and in Health? Health Shocks and Relationship Breakdown: Empirical Evidence from Germany." *Journal of Economic Behavior & Organization*, vol. 190, 2021, pp. 164–190. https://doi.org/10.1016/j.jebo.2021.07.029

Chaisemartin, Clément De, and Xavier d'Haultfoeuille. "Difference-in-Differences Estimators of Intertemporal Treatment Effects." *Review of Economics and Statistics*, 2024, pp. 1–45. https://doi.org/10.1162/rest_a_01414

Chandra, Amitabh, Courtney Coile, and Corina Mommaerts. "What Can Economics Say About Alzheimer's Disease?" *Journal of Economic Literature*, vol. 61, no. 2, 2023, pp. 428–470. https://doi.org/10.1257/jel.20211660

Chari, Amalavoyal V., John Engberg, Kristin N. Ray, and Ateev Mehrotra. "The Opportunity Costs of Informal Elder-Care in the United States: New Estimates from the American Time Use Survey." *Health Services Research*, vol. 50, no. 3, 2015, pp. 871–882. https://doi.org/10.1111/1475-6773.12238

Charles, Kerwin Kofi, and Purvi Sevak. "Can Family Caregiving Substitute for Nursing Home Care?" *Journal of Health Economics*, vol. 24, no. 6, 2005, pp. 1174–1190. https://doi.org/10.1016/j.jhealeco.2005.05.001

Chetty, Raj. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics*, vol. 90, no. 10–11, 2006, pp. 1879–1901. https://doi.org/10.1016/j.jpubeco.2006.01.004

---and Amy Finkelstein. "Social Insurance: Connecting Theory to Data." *Handbook of Public Economics*, vol. 5, 2013, pp. 111–193. https://doi.org/10.1016/B978-0-444-53759-1.00003-0

Davis, Matthew A., Chiang-Hua Chang, Sharon Simonton, and Julie P. W. Bynum. "Trends in US Medicare Decedents' Diagnosis of Dementia from 2004 to 2017." *JAMA Health Forum*, vol. 3, no. 4, 2022, e220346. https://doi.org/10.1001/jamahealthforum.2022.0346

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. "The Economic Consequences of Hospital Admissions." *American Economic Review*, vol. 108, no. 2, 2018, pp. 308–352. https://doi.org/10.1257/aer.20161038

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor. "A Local Projections Approach to Difference-in-Differences Event Studies." National Bureau of Economic Research, 2023, NBER Working Paper 31184.

Eliason, Paul J., Paul L. E. Grieco, Ryan C. McDevitt, and James W. Roberts. "Strategic Patient Discharge: The Case of Long-Term Care Hospitals." *American Economic Review*, vol. 108, no. 11, 2018, pp. 3232–3265. https://doi.org/10.1257/aer.20170092

Ellis, Randall P., Heather E. Hsu, Jeffrey J. Siracuse, et al. "Development and Assessment of a New Framework for Disease Surveillance, Prediction, and Risk Adjustment: The Diagnostic Items Classification System." *JAMA Health Forum*, 2022, e220276. https://doi.org/10.1001/jamahealthforum.2022.0276

Fadlon, Itzik, and Torben Heien Nielsen. "Family Health Behaviors." *American Economic Review*, vol. 109, no. 9, 2019, pp. 3162–3191. https://doi.org/10.1257/aer.20171993

---. "Household labor supply and the gains from social insurance." *Journal of Public Economics*, vol. 171, 2019, pp. 18–28. https://doi.org/10.1016/j.jpubeco.2018.01.010

---. "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records." *American Economic Journal: Applied Economics*, vol. 13, no. 3, 2021, pp. 1–30. https://doi.org/10.1257/app.20170604

Favreault, Melissa M., and Richard W. Johnson. "Projections of Risk of Needing Long-Term Services and Supports at Ages 65 and Older." Urban Institute, 2021. Prepared for the Office of the Assistant Secretary for Planning and Evaluation, U.S. Department of Health and Human Services.

Fontes, Luiz Felipe, Matias Mrejen, Beatriz Rache, and Rudi Rocha. "Economic Distress and Children's Mental Health: Evidence from the Brazilian High Risk Cohort Study for Mental Conditions." *The Economic Journal*, vol. 134, no. 660, 2024, pp. 1701–1718. https://ideas.repec.org/a/oup/econjl/v134y2024i660p1701-1718..html

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano. "Do Fiscal Rules Matter?" *American Economic Journal: Applied Economics*, vol. 8, no. 3, 2016, pp. 1–30. https://doi.org/10.1257/app.20150076

Gross, Tal, Tim Layton, Daniel Prinz, and Julia Yates. "Social Defaults and Plan Choice: The Case of Spousal Following." 25/28 Institute for Fiscal Studies, 2024, Working Paper.

Gruber, Jonathan, and Kathleen M. McGarry. "Long-Term Care in the United States." National Bureau of Economic Research, 2023, Working Paper 31881.

Hoagland, Alex. "An Ounce of Prevention or a Pound of Cure? The Value of Health Risk Information." 2025, Working Paper.

Hodor, Michal. "Family Health Spillovers: Evidence from the RAND Health Insurance Experiment." *Journal of Health Economics*, vol. 79, 2021, 102505. https://doi.org/10.1016/j.jhealeco.2021.102505

Hunt, Erin Cottle, and Frank N. Caliendo. "Social Security and Longevity Risk: An Analysis of Couples." *Journal of Public Economic Theory*, vol. 24, no. 3, 2022, pp. 547–579. https://doi.org/10.1111/jpet.12564

Ilinca, Stefania, Ricardo Rodrigues, and Andrea E. Schmidt. "Fairness and Eligibility to Long-Term Care: An Analysis of the Factors Driving Inequality and Inequity in the Use of Home Care for Older Europeans." *International Journal of Environmental Research and Public Health*, vol. 14, no. 10, 2017, 1224. https://doi.org/10.3390/ijerph14101224

Kang, Zi Yang, and Shoshana Vasserman. "Robustness Measures for Welfare Analysis." *American Economic Review*, vol. 115, no. 8, 2025, pp. 2449–2487. https://doi.org/10.1257/aer.20220673

Lage, Daniel E., Michael C. Jernigan, Yuchiao Chang, David C. Grabowski, John Hsu, Joshua P. Metlay, and Sachin J. Shah. "Living Alone and Discharge to Skilled Nursing Facility Care after Hospitalization in Older Adults." *Journal of the American Geriatrics Society*, vol. 66, no. 1, 2018, pp. 100–105. https://doi.org/10.1111/jgs.15150

Maestas, Nicole, Matt Messel, and Yulya Truskinovsky. "Caregiving and Labor Supply: New Evidence from Administrative Data," *Journal of Labor Economics*, vol. 42, S1, 2024, S183–S218. https://doi.org/10.1086/728810

Massner, Patrizia, and Jens Wikström. "Should Elder Care Be Subsidized? Theory and Evidence from Sweden." 2024, Job Market Paper.

McGarry, Brian E., David C. Grabowski, Lin Ding, and J. Michael McWilliams. "Outcomes After Shortened Skilled Nursing Facility Stays Suggest Potential for Improving Postacute Care Efficiency." *Health Affairs*, vol. 40, no. 5, 2021, pp. 745–753. https://doi.org/10.1377/hlthaff.2020.00649

OECD. "Is Care Affordable for Older People?" *OECD Health Policy Studies*, OECD Publishing, Paris, 2024. https://doi.org/10.1787/450ea778-en

Rhee, Jong Chul, Nutthaporn Done, and Gerard F. Anderson, "Considering long-term care insurance for middle-income countries: comparing South Korea with Japan and Germany," *Health Policy*, 2015, 119 (10), 1319-1329. https://doi.org/10.1016/j.healthpol.2015.06.001

Roofeh, Regina, Dylan M. Smith, and Sean A.P. Clouston. "Estimated Prevalence of Elder Orphans Using National Health and Aging Trends Study." *Journal of Aging and Health*, vol. 32, no. 10, 2020, pp. 1443–1449. https://doi.org/10.1177/0898264320932382

Zeckhauser, Richard. "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory*, vol. 2, no. 1, 1970, pp. 10–26. https://doi.org/10.1016/0022-0531(70)90010-4

About the authors

Itzik Fadlon is an Associate Professor of Economics at the University of California, San Diego (UCSD) and a Research Associate in the programs on the Economics of Health, Aging, and Public Economics at the National Bureau of Economic Research (NBER). He is also an Invited Researcher at J-PAL North America at MIT and an Associate Editor at the Journal of Health Economics. His primary fields of interest are public finance and health economics, with a particular interest in addressing central questions in the economics of aging and the optimal design of social policies. He received his PhD in economics from Harvard University in May 2015. Before joining UCSD in 2016 as an Assistant Professor, Itzik spent a year as a Postdoctoral Fellow in Disability Policy Research at the NBER funded by the Social Security Administration, and he spent the academic year of 2019-2020 as a Visiting Scholar in Aging and Health Research at the NBER funded by NIH/NIA.

Tal Gross is a professor in the Department of Markets, Public Policy & Law at Boston University. He is also a Research Associate at the National Bureau of Economic Research. His research focuses on health insurance and household finance. He received his bachelor's degree from The University of Chicago and his PhD from the Massachusetts Institute of Technology. He previously taught at the University of Miami School of Business Administration and the Columbia University Mailman School of Public Health.

Alex Hoagland is an Assistant Professor and early-career researcher in the University of Toronto's Institute of Health Policy, Management and Evaluation in the Dalla Lana School of Public Health, and Co-Director of the Canadian Centre for Health Economics. He received his PhD in Economics from Boston University in May 2022. His primary field is health economics, with additional research in industrial organization and applied econometrics. His research focuses on how information shapes equitable access to health care, health decision-making by patients and providers, and innovation adoption in healthcare.

Tim Layton is an Associate Professor of Public Policy and Economics in the Batten School of Leadership and Public Policy at the University of Virginia, a Research Associate at the National Bureau of Economic Research, Co-Editor of the Journal of Health Economics, and an affiliated researcher at J-PAL at MIT. His research focuses on the economics of health insurance markets, with a particular focus on markets and social health insurance programs for low-income households. He received his bachelor's degree from Brigham Young University and his PhD from Boston University. Before coming to UVA, he was the 30th Anniversary Associate Professor of Health Care Policy at Harvard Medical School.

About the TIAA Institute

The TIAA Institute helps advance the ways individuals and institutions plan for financial security and organizational effectiveness. The Institute conducts in-depth research, provides access to a network of thought leaders, and enables those it serves to anticipate trends, plan future strategies, and maximize opportunities for success.

To learn more, visit tiaainstitute.org.

5185091-A

TIAA Institute

(02/26)