NBER WORKING PAPER SERIES

THE PROTECTIVE EFFECTS OF A HEALTHY SPOUSE: MEDICARE AS THE FAMILY MEMBER OF LAST RESORT

Itzik Fadlon
Tal Gross
Alex Hoagland
Timothy Layton

Working Paper 35784
http://www.nber.org/papers/w35784

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

We are grateful to Elena Bassoli, Ashvin Gandhi, Jonathan Kolstad, Bruce Landon, Corina Mommaerts, and participants at the 2024 California Health Economics Conference, the 2025 NBER Health Economics Annual Program Meeting, the 2025 NBER-CEPRA Lugano Meeting on Aging and Health, the 2025 Collaborative Health Economics Seminar Series at George Washington University, and the 2025 NBER Aging Program Meeting Brigham Young University, Georgia State University, Simon Fraser University, University of Alabama - Birmingham, University of California - Irvine, University of Georgia, University of Miami, University of Toronto, and the 2026 SITE conference for helpful feedback. Prabidhik KC and Jimmy Yung provided superb research assistance. This project was supported by the National Institute on Aging (NIA P30 Project Grant #P30AG012810, via NBER) and the TIAA Institute/ Pension Research Council (UPenn Subaward 86688/00). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w35784

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Itzik Fadlon, Tal Gross, Alex Hoagland, and Timothy Layton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Protective Effects of a Healthy Spouse: Medicare as the Family Member of Last Resort
Itzik Fadlon, Tal Gross, Alex Hoagland, and Timothy Layton
NBER Working Paper No. 35784
September 2026
JEL No. I13, I18, J12

ABSTRACT

We study the value of informal care from a spouse and how access to such care shapes optimal insurance coverage for long-term care. Using Medicare claims and a novel spousal crosswalk, we examine how the sudden incapacitation of a spouse affects their partner's use of formal care. A spouse's health shock increases their partner's probability of a nursing-home stay by 30%. We document significant negative health effects for the partner, but provide evidence that their increase in nursing-home stays is largely driven by substitution from informal care to formal care. When their partners are incapacitated, Medicare recipients become less sensitive to the price of formal care, consistent with a higher valuation of that care and reduced moral hazard. We develop a framework to interpret these findings. The framework translates differences in the level and price elasticity of demand for formal care into a lower-bound estimate of the value of informal care provided by a spouse. It also suggests that optimal insurance for long-term care should be designed at the household level with more generous coverage for those who do not have access to informal care.

Itzik Fadlon
University of California, San Diego
Department of Economics
and NBER
fadlon@ucsd.edu

Alex Hoagland
University of Toronto
Dalla Lana School of Public Health
Institute of Health Policy, Management,
and Evaluation
alexander.hoagland@utoronto.ca

Tal Gross
Boston University
Questrom School of Business
and NBER
talgross@bu.edu

Timothy Layton
University of Virginia
and NBER
timothyjlayton@gmail.com

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 this growing need is a first-order issue for social welfare and public budgets worldwide.

Long-term care differs from most other forms of healthcare because family members can provide an alternative at home. Unlike, for instance, surgery or diagnostic tests, help with the activities of daily living can be supplied 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, typically covered by insurance. Informal care consists of the unpaid help of a spouse or adult children at home.

Informal care is important, but it is difficult to measure and even more difficult to value. Family members providing care do not clock in, track their hours, or log their effort.[1] As a result, informal care, like most other forms of home production, is typically unobserved and unmeasured. This has led to a wide range of estimates of the value of informal long-term care for older adults, from $86 billion (Gruber and McGarry, 2023) to over $500 billion (Chari et al., 2015).

Setting aside questions of measurement, the existence of informal care also complicates standard public insurance design. The public sector must decide not only how much to pay for formal care but also whether to even pay at all, given that there exists a private low-cost alternative: informal care. The presence of the private alternative increases the scope for moral hazard, whereby patients use expensive formal care when the government is paying, but opt to rely on informal care, which requires little direct financial expenditure, when facing the cost of formal care themselves. Such concerns are always present in social insurance design (Chetty, 2006; Zeckhauser, 1970), but they are exacerbated when informal alternatives are inexpensive and readily available.

Countries try to resolve these tradeoffs in different ways. The Dutch system, for example, reduces a beneficiary's eligibility for publicly funded care when they have a healthy spouse (Ilinca et al., 2017). As a result, the Dutch effectively raise the price of formal care when informal care is available. The German system does roughly the opposite: it charges higher premiums to adults without children (Nadash

[1] Some household surveys, such as the Health and Retirement Study, collect self-reported measures of informal care among surveyed households.

et al., 2018). By comparison, the US Medicare program offers identical coverage to all beneficiaries regardless of household structure or access to informal care. These contrasting approaches signal genuine uncertainty regarding how public coverage of long-term care should take the availability of informal care into account.

This paper makes three contributions in this setting. First, we offer an empirical approach to estimate the effects of spousal informal care. We implement this approach using novel data that identify spouses in Medicare claims. Second, we develop a framework to convert these effects into a lower-bound estimate of the value of informal care from a spouse. Third, we derive how long-term-care policies—such as those in Germany and the Netherlands—ought to account for the availability of informal care from a spouse.

We begin by estimating how the incapacitation of a spouse affects the level of demand for formal care. In doing so, we demonstrate that formal care acts as a substitute for informal care and quantify the burden that informal caregiving diverts from the formal healthcare system. We study couples in administrative Medicare records linked to a census of nursing-home stays. When one spouse—the "shock spouse"—experiences a heart attack or a stroke, the probability that their partner—the "outcome spouse"—enters a nursing home increases by an average of 30% over the following year.[2] The effects are larger when the shock spouse is incapacitated for a longer period, or when the outcome spouse is sicker and has a greater need for care. These results suggest that the loss of the primary caregiver is a likely mechanism for the increase in nursing-home visits, rather than some other direct effect of one's spouse experiencing a health shock. Further analyses similarly suggest that most of the increase in nursing-home visits reflects substitution from informal to formal care. Deterioration in the outcome spouse's health accounts for a small share of the increase.

We use an additional source of variation to explore the extent to which the loss of a spousal caregiver not only shifts the level of demand but also changes the price elasticity of demand for formal care. To do so, we leverage a feature of Medicare Part A coverage for skilled nursing facility (SNF) stays following a hospitalization: a sharp rise in patient out-of-pocket costs after 20 days in a SNF. Using a regression-discontinuity (RD) analysis, we show that beneficiaries with a healthy spouse at home are more likely to exit on day 20, just before copayments commence, than beneficiaries

[2]This estimate includes all nursing-home stays, whether paid by Medicare, Medicaid, commercial insurance, or by the beneficiary. We also observe an increase in entries to skilled nursing facilities, whose charges are covered by Medicare. The risk of such visits increases by an average of 20%.

whose spouse was recently incapacitated.[3] This pattern suggests that losing informal care makes demand for formal care less sensitive to price.

Next, we use those empirical results to estimate the value of informal care and study the design of optimal insurance contracts that adjust for the availability of informal care. A growing literature has documented the within-household spillover effects of spousal health shocks. Few studies focus on substitution between formal and informal care and thus few have highlighted the implications for the public financing of long-term care. We study these implications by developing a model of demand for formal care. The model converts the RD estimates into two indirect demand curves for formal care: one when a healthy spouse is present and one when that spouse has been incapacitated. We show that the area between those two demand curves offers a lower bound for willingness to pay for informal care from one's spouse. The lower bound is closer to the actual valuation of informal care when informal and formal care are closer substitutes. That insight is likely useful in other contexts whenever there exists an unobserved good for which there is no market—in this case, informal care—and an observed good for which prices vary—in this case, formal care.

Calibrating the model indicates that consumer surplus from an extra day of formal care is about $568 higher when a healthy spouse is unavailable as compared to when a healthy spouse is present. The model shows that this estimate offers a lower bound for the willingness to pay for an extra day of informal care. The estimate is around 1.6 times the typical cost of a day of nursing-home care and 2.6 times the typical cost of a day of paid in-home care.[4] It is also more than three times the median daily earnings of a full-time US worker, indicating substantial unmeasured value.[5]

We then use the model to derive a condition in the spirit of the work by Baily (1978) and Chetty (2006) that determines the optimal generosity of formal-care coverage for those with and without healthy spouses. The optimal generosity of insurance depends on the relative price elasticities of demand for formal care for those with and without healthy spouses.[6] We use the empirical estimates to calibrate the model and

[3]Exit rates are largely similar between the two groups for all days prior to the introduction of copayments.

[4]Both Gruber and McGarry (2023) and survey data indicate that the cost of nursing-home care ranges between $300 and $350 per day, and the cost of paid in-home care ranges between $185 and $215 per day.

[5]The Bureau of Labor Statistics reports median usual weekly earnings of $860 for full-time wage and salary workers in 2017, the last year of our sample (U.S. Bureau of Labor Statistics, 2018). We use 2017 rather than a current figure because the $568 estimate is denominated in dollars of the sample period.

[6]As in the Baily-Chetty model (Baily, 1978; Chetty, 2006), optimal generosity also depends on the

show that our estimated price elasticities of demand imply that the optimal subsidy is 80% larger for those with sick spouses than for those with healthy spouses.

Finally, we extend the analysis with a dynamic model that assesses welfare from day 20 onwards. The model extrapolates from day 20 to all other potential days in a SNF, allowing us to perform welfare analyses on a set of counterfactual Medicare benefit designs. Simulations suggest that shifting coverage generosity from beneficiaries with healthy spouses to beneficiaries whose spouses are incapacitated could increase social welfare by 13.4% over the status quo while holding government spending constant. These results demonstrate that ignoring the availability of a healthy spouse leads to suboptimal benefit design with meaningful welfare losses. According to the model, Medicare's current policy, which ignores the household entirely, under-insures those who have lost access to informal care and over-subsidizes those who can still find informal care at home.

This paper contributes to four areas of inquiry. First, we contribute to research on aging and long-term care in the US. Gruber and McGarry (2023) estimate that formal care expenditures totaled $230 billion in 2018, mostly financed by the public sector, and that informal care plays a critical role in providing for older adults. At every age and disability level, older adults receive informal care more frequently than formal care. Studies have also found that informal care may delay entry into formal care (Bergeot and Tenand, 2023; Bonsang, 2009; Charles and Sevak, 2005) and that it affects the labor supply and earnings of caregivers, with consequences for households that extend beyond health (Maestas et al., 2024; Massner and Wikström, 2024). We contribute to this literature by providing a revealed-preference method for assessing beneficiaries' valuation of informal care based on their formal-care choices. By linking spouses in population-level Medicare data, we study how shocks to one spouse's caregiving capacity change the other spouse's use of, and price sensitivity to, formal care. To the best of our knowledge, we provide the first analysis of intra-family interactions in the demand for elderly care in the US using population-level, administrative data.

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 can affect other family members. That research has explored joint learning about health risks (Oster et al., 2013;

relative marginal utilities of individuals in those two states. We make the standard assumption that households in the bad state are weakly under-insured so that their marginal utility of consumption is weakly higher. (The model also allows for the general case.)

Hoagland, 2025), the interdependence of family members' risks (Coile, 2004; Fadlon and Nielsen, 2019a; Hur and Lee, 2026), and the implications for public policy (Brown and Finkelstein, 2011; Fadlon and Nielsen, 2019b; Gross et al., 2025). Additional work has highlighted the effects of economic and health shocks on household labor supply and family stability (Fadlon and Nielsen, 2021; Hodor, 2021; Arteaga et al., 2025; Fontes et al., 2024). We contribute to this literature by providing evidence on the value of a healthy spouse via the intra-household interactions between formal and informal care among older US couples. We also introduce a general framework that uses these reduced-form estimates to address questions of optimal policy design at the household level.

Third, this paper documents dynamics that can inform the optimal design of health insurance contracts. When one person experiences a health shock, the care options available to their spouse change, ultimately affecting their willingness to pay for formal healthcare services covered by social health insurance programs. Optimal contracts provide more generous coverage when price sensitivity is lower and so moral hazard is less of a concern (Zeckhauser, 1970). This paper's results reveal that price sensitivity is lower when one's spouse is incapacitated, implying that optimal coverage is higher in that state. More broadly, optimal insurance design considers levels of utilization by the household rather than the individual. Yet, in many cases health insurance plans are individual and make no adjustment for the rest of the household. Medicare offers individual contracts with no consideration of spousal characteristics, and many employer-provided plans still feature individual-specific rather than household-specific out-of-pocket limits or deductibles, despite evidence that household-level cost sharing shapes consumption decisions (Anderson et al., 2024). In contrast, other social insurance programs—such as Social Security—explicitly incorporate survivor benefits into their benefit structure (Brown and Poterba, 2000; Cottle Hunt and Caliendo, 2022; Coyne et al., 2024). The model and estimates provide guidance on the potential benefits of health insurance plans mimicking this household-level structure.

Finally, this paper is related to studies that value non-market goods. A long-standing challenge has been the valuation of home production: informal care, childcare, cooking, and cleaning (Reid, 1934; Becker, 1965; Gronau, 1977).[7] Because those services are unpriced, they are typically valued by measuring either forgone wages or the price of a paid substitute. Those two methods, however, have long been known

[7]Similarly, researchers in environmental, urban, and real-estate economics seek to infer the value of non-market goods: clean air, open space, or public amenities (Rosen, 1974; Roback, 1982; Chay and Greenstone, 2005).

to be unreliable. Researchers cannot solely rely on forgone wages, because doing so conflates the value of the output with the value of time and imposes strong technological assumptions (Pollak and Wachter, 1975). This is especially relevant in our context, where care is provided by older spouses who are likely retired and thus not forgoing wages, yet their efforts are likely valuable to the care recipient. Researchers also cannot solely rely on the price of paid substitutes, because doing so assumes an equal-quality market substitute exists (Gronau, 1977; Abraham and Mackie, 2005). Getting these valuations right nonetheless matters: home production may rival measured market output in magnitude (Nordhaus and Tobin, 1972; Aguiar and Hurst, 2007). In this space, we provide a new, revealed-preference approach that recovers a lower bound on consumers' willingness to pay for an unobserved product. Our approach requires only variation in the good's availability and variation in the price of an observed substitute. This method therefore generalizes naturally to applications in other contexts.

The remainder of the paper proceeds as follows. Section 2 describes the data and research designs. Section 3 presents empirical estimates in two pieces. We first present event-study estimates for the effects of a spouse's health shock on one's own health and healthcare utilization, and we then assess how a spouse's health shock affects individuals' responses to the out-of-pocket price of formal care. Section 4 develops a model of demand for formal long-term care in the household setting and explores two implications of the model. First, the model converts estimates of the effects of the incapacitation of one's spouse into the valuation of informal spousal care. Second, the model demonstrates the implications of the estimates for policy. Section 5 concludes.

2 Empirical Framework

2.1 Data and Analysis Sample

We rely on three main datasets. First, we study Medicare claims. We observe 100% of hospital inpatient stays and qualifying SNF stays in the Medicare Provider Analysis and Review (MedPAR) files, as well as beneficiary summary information and annual measures of utilization. In addition, we observe detailed outpatient and carrier claims, along with prescription-drug utilization, for a 20% sample of enrollees.

Second, we rely on a crosswalk, created by analysts at the Centers for Medicare & Medicaid Services (CMS), that links spouses in Medicare. The crosswalk consists of beneficiaries' encrypted addresses as tracked by CMS. To create the dataset, analysts

encrypted each Medicare recipient's entire address field (number, street, apartment number, state, and ZIP Code). We then identify spouses by selecting encrypted addresses with a maximum of two Medicare recipients of the opposite sex and then linking beneficiaries who have the same encrypted address. Gross et al. (2025) validate this identification of spouses by comparing the Medicare data to data from the American Community Survey.

The third dataset we use is the Minimum Data Set (MDS). The MDS contains 100% of federally required assessments for residents of CMS-certified nursing facilities, allowing us to observe nursing-home stays and SNF stays regardless of payer. Whereas traditional Medicare claims data allow us to observe formal care covered by Medicare (i.e., SNF stays), the MDS captures demand for formal care that is financed privately or through Medicaid. Combined, these three datasets provide a comprehensive view of formal-care use among couples enrolled in fee-for-service Medicare from 2011 through 2017.[8]

We limit the sample to households that retain Medicare eligibility throughout the window of observation and the previous two calendar years. We assign individuals to couples based on observed cohabitation in the year prior to the index health event. We then identify index health events based on hospitalization claims for a heart attack or stroke. We term the individual who experienced the health event the "shock spouse," and call their linked partner the "outcome spouse." Following the work of Dobkin et al. (2018), we focus on health shocks occurring to shock spouses who have not been admitted to a hospital in the previous two years. We omit health shocks accompanied by external injuries because they may reflect common shocks, such as car accidents, that affect both spouses. In our preferred specifications, we retain only the first such qualifying event within each spousal pair; below, we highlight the robustness of our results to alternative approaches to constructing the sample.

Table 1 presents summary statistics for the analysis sample. 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.

[8]The crosswalk covers spouses from 2010 to 2017, allowing us to identify spouses based on the year prior to a health shock.

Table 1. Summary Statistics

	Shock Spouses	Outcome Spouses
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
B. Healthcare Utilization		
Has a Chronic Condition	0.46	0.41
# of SNF Stays, Conditional on Any	1.11	1.35
Any Inpatient Admission	0.00	0.10
# of Inpatient Admissions, Conditional on Any	1.14	1.52
Average Annual 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. "FFS:" fee-for-service; "HMO:" health maintenance organization.

2.2 Research Designs

In the empirical analysis below, we identify the effects of spousal incapacitation on the utilization of formal care and on the willingness to pay for it. The analysis involves two related designs. First, we use an event-study design to identify the dynamic effects of a spouse's incapacitation. This design compares outcome spouses to other outcome spouses who experience the same event but at a later date. Second, we use a difference-in-discontinuities design to exploit a discrete change in cost sharing under the Medicare program. This design identifies how a spouse's incapacitation affects price sensitivity for formal care.

Identifying the Effects of Spousal Incapacitation on Formal Care. A macabre thought experiment would identify the effects of a spouse's health shock by assigning such shocks at random. We take a quasi-experimental approach that aims to approximate that thought experiment. We restrict the analysis to couples that have

experienced a health event—specifically, a heart attack or stroke—and identify the treatment effect using variation in when the event occurred. We follow Fadlon and Nielsen (2019a, 2021) and construct counterfactual outcomes for affected families using families that experience the same type of event at a later date.

We construct a treatment group, composed of spouses in households that experience a health shock in a particular calendar month, and a matched control group composed of spouses in households that experience a similar event twelve months later. We assign control households a placebo event in the same calendar month as the treatment-group event. We then recover the treatment effect with a dynamic difference-in-differences estimator, identifying the shock's impact from the differential change in outcomes across the two groups over time. We can estimate the effects of the shocks for twelve months, because the control group becomes "treated" twelve months later.[9] By construction, the analysis is not subject to concerns regarding units that switch across experimental arms (De Chaisemartin and d'Haultfoeuille, 2026).

We estimate the following event-study regression:

$$y_{i,t} = \alpha_i + \theta_t + \sum_{n \neq -1,-5} \gamma_n \cdot I^n_{i,t} + \sum_{n \neq -1,-5} \delta_n \cdot I^n_{i,t} \cdot \text{treated}_i + \varepsilon_{i,t}, \tag{1}$$

where $y_{i,t}$ denotes an outcome for outcome spouse i during calendar month t, treated_i denotes an indicator for whether a household belongs to the treatment group, and $I^n_{i,t}$ denotes a set of indicators for months relative to the assigned event period.[10] Equation (1) includes calendar-time and person-specific fixed effects, θ_t and α_i, and fixed effects for the relative time indicators $I^n_{i,t}$, expressed with γ_n. The key parameters of interest are the δ_n coefficients, which estimate the period-n treatment effects relative to the pre-period $n = -1$. This empirical strategy rests on the standard parallel-trends identifying assumption. It assumes that, absent the health shock, the outcomes of the treatment and control groups would evolve in parallel. To assess the validity of this assumption, we study the evolution of the outcomes of the treatment

[9] In complementary analyses, we use a gap of more than twelve months between treatment and control events to identify longer-run effects. However, it comes with a potential tradeoff since it may introduce bias from less-comparable health events as the elapsed time between shocks lengthens. That said, all analyses consistently suggest no pre-existing, divergent trends, and show that effects are strongly persistent beyond the 1-year analysis window.

[10] Although control households are not treated during the analysis window, they are treated later. Therefore, we omit a second time period ($n = -5$, the period immediately preceding our pre-trend coefficients) in the baseline specification to separately identify the treatment effects from time trends across groups (Sun and Abraham, 2021). Appendix Figure A1 shows that the results are qualitatively unaffected by this choice.

and control groups in the periods prior to the event.

To calculate average effects over time, we estimate the following mean difference-in-differences specification:

$$y_{i,t} = \alpha_i + \theta_t + \gamma \cdot \text{post}_{i,t} + \delta \cdot \text{treated}_i \cdot \text{post}_{i,t} + \varepsilon_{i,t}, \tag{2}$$

where $\text{post}_{i,t}$ represents an indicator function equal to one when the observation belongs to a period after the shock occurred. Accordingly, δ measures the average treatment effect over the 1-year horizon.

A potential concern with this approach is that the shock spouse may have experienced a cardiac event as a result of the status of the outcome spouse. Consider an extreme case in which a husband is hospitalized every few months and his wife experiences a single, major cardiac event in the middle of the sample. In such a case, the wife would meet the criteria of a shock spouse, but her cardiac event could have been caused by one of her husband's many hospitalizations. We address this possibility in several ways. First, we test for anticipatory effects via event-study figures. Second, we exclude couples for whom the initial health shock had accompanying external injuries, the most common type of correlated shock. Third, in companion analyses below, we restrict the sample to couples in which neither spouse was hospitalized for at least three months prior to the index event. Finally, Appendix A offers a variety of alternative sample restrictions testing whether the main results are sensitive to a particular restriction. In particular, we analyze the extent to which the shock spouse's cardiac event causes cardiac events for the outcome spouse (Appendix Figure A2). We find no evidence for that pattern, which limits concerns regarding circularity.[11] All of these alternative approaches demonstrate the robustness of the main findings.

A final consideration involves the way in which mortality affects this sample. As documented below, the shock spouse's cardiac event leads to an increased risk that the outcome spouse dies. That effect then requires us to decide whether and how those who die are included in the analysis sample. In most of what follows, we limit the sample to a balanced panel consisting of outcome spouses who survived through the sample window. We do so because our primary interest is the potential substitution

[11] In addition, Appendix Figure A3 presents estimates of the main specification in which we do not require households to have a single shock spouse and a single outcome spouse and do not impose the requirement that the shock spouse's cardiac event be preceded by two years without hospitalizations. The estimates are similar to those presented in the main analysis. Fewer than 2% of households in the sample experience multiple health events that meet the criteria we use to define shock spouses. Removing the two-year no-hospitalization requirement only slightly attenuates the estimates.

between formal and informal care, and the main formal-care outcome of interest—a stay in a SNF or nursing home—cannot be observed after death. Accordingly, we apply the same balanced-panel restriction to control groups as we do to treatment groups so that the two groups are comparable. We also present estimates without conditioning on survival in order to estimate the overall effects on spending and utilization, rather than the specific formal-informal substitution that is our main focus (Appendix Figure A4).

Identifying the Effects on Formal Care Price Sensitivity. Second, we estimate how a spouse's incapacitation changes beneficiary sensitivity to the price of formal care. To do so, we exploit a discrete change in cost sharing under the Medicare program.[12]

When Medicare recipients visit a SNF after a qualifying hospitalization, they face zero cost sharing through day 20.[13] Then, after the 20$^{\text{th}}$ day of their stay, they discontinuously face copayments equal to one-eighth of the Medicare inpatient deductible, or approximately 20% of the Medicare reimbursement.[14] We study responses to that change via a regression discontinuity (RD) framework.

We use the same comparison across treatment and placebo control SNF stays as described above. That is, we estimate how outcome spouses respond to the change in the effective price they face at day 20 for SNF stays occurring within four months of their spouse's heart attack or stroke, and compare those price responses to those observed for similar SNF stays occurring within four months of the placebo control event. As above, we adjust our estimates for calendar-time and person-specific fixed effects.

To quantify how a spouse's health shock affects the response to cost sharing, we

[12]By necessity, then, we limit this analysis to Medicare-covered SNF stays.

[13]To qualify for no cost sharing (or any Medicare coverage, for that matter) for the first 20 days of a SNF stay, Medicare recipients must have been admitted to the hospital for at least three days (excluding the day of discharge) immediately prior to SNF admission; have been admitted to the SNF within 30 days of that inpatient stay; and have a documented requirement for daily skilled care.

[14]Medicare beneficiaries enrolled in a Medigap plan may not pay those increased copayments. Many Medigap plans cover the increased copayments in their entirety. Among beneficiaries entering SNFs for post-acute care, roughly 25% do not have any form of supplemental coverage, Medigap or Medicaid (McGarry et al., 2021). These beneficiaries drive the demand response we estimate. In Section 4, we use our estimates to construct demand curves, and we adjust the implied average price to account for the beneficiaries who face no copayments. Unfortunately, Medigap enrollment is not observed in any administrative data source, preventing us from assigning individual-specific prices.

use a difference-in-discontinuities design. Consider the following regression:

$$
\begin{aligned}
y_{i,t} = & \underbrace{\alpha_0 + \alpha_1(t-20)}_{\text{baseline trend}} \\
& + \underbrace{\gamma_0\,\mathbf{1}\{t>20\} + \gamma_1(t-20)\mathbf{1}\{t>20\}}_{\text{day-20 RD}} \\
& + \underbrace{\tau_0\,\mathbf{1}\{\text{treated}_i=1\} + \tau_1(t-20)\mathbf{1}\{\text{treated}_i=1\}}_{\text{spousal-shock baseline differences}} \\
& + \underbrace{\beta_1\,\mathbf{1}\{t>20\}\mathbf{1}\{\text{treated}_i=1\} + \beta_2(t-20)\mathbf{1}\{t>20\}\mathbf{1}\{\text{treated}_i=1\}}_{\text{day-20 difference-in-discontinuities}} \\
& + X_{i,t}'\Gamma + \varepsilon_{i,t},
\end{aligned}
\tag{3}
$$

where $y_{i,t}$ represents an indicator for whether outcome spouse i resides in a SNF on day t. We estimate equation (3) with a triangular kernel, which gives greater weight to observations closer to the day-20 cutoff. The parameter β_1 indicates the difference in responses across those with and without a healthy spouse. The vector $X_{i,t}$ collects calendar-time and person-specific fixed effects, as above.

This analysis identifies how the day-20 increase in Medicare cost sharing differentially affects the probability that beneficiaries remain in a SNF after a spouse's incapacitation. It relies on two identifying assumptions. First, it requires the same parallel-trends assumption that underlies the event-study analysis, described above. Second, it relies on the RD assumption that discharge probabilities in the absence of the change in cost sharing would evolve smoothly. Under those assumptions, the difference-in-discontinuities estimate recovers the causal effect of spousal incapacitation on the response to the day-20 increase in out-of-pocket costs. By itself, this estimate identifies a differential behavioral response to the price change. Section 4 imposes additional structure to show how these estimates can be interpreted as differences in price elasticity and how those elasticities inform the valuation of informal care and the design of optimal insurance.

A potential concern with this approach is that spousal incapacitation may change the composition of patients in the SNF, not just their response to cost sharing. Some of those with incapacitated spouses are marginal patients who enter a SNF solely because their spouse can no longer care for them. Marginal patients may have a lower willingness to pay for an additional day in the SNF, and so be more responsive to copayments. That composition effect alone would make those with sick spouses

more price sensitive—the opposite of what we find. It therefore works against the elasticity difference we estimate, and cannot account for it. Further, we find that prior to day 20, daily discharge rates are similar for those with healthy versus sick spouses, with discharge rates only diverging after copayments begin. This provides evidence that, at least on this observed characteristic, the two groups remain similar, despite the presence of additional SNF entrants in the sick-spouse group.

3 Empirical Results

This section presents the paper's empirical results in two steps. We first present event-study estimates of the effect of a spouse's health shock on their partner. Second, we compare the price sensitivity of beneficiaries with and without a healthy spouse.

3.1 Effects on Formal Care Utilization

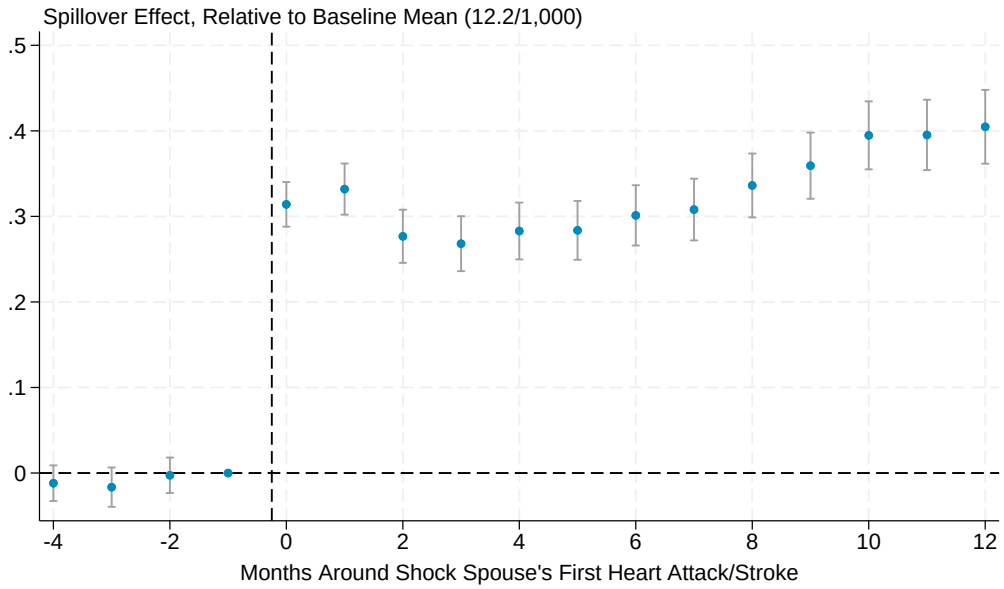
Figure 1 examines how an outcome spouse's risk of being in a nursing home evolves before and after their partner's health shock.[15] We normalize the coefficients by dividing by the baseline mean so they can be interpreted as percent effects relative to baseline. We observe no systematic differential changes across groups in the months prior to the shock spouse's cardiac event, supporting the research design. We then observe a sudden increase of more than 30% in the likelihood that an outcome spouse resides in a nursing home immediately after their partner's cardiac event. This risk remains elevated over the twelve months following the event, with an average increase of 40% by the end of the first year.

We further explore the dynamics of these effects by studying their immediacy and persistence. Appendix Figure A9 repeats the analysis for this main outcome at the weekly level: the effects materialize soon after the shock with no evidence of any effect prior to the shock, phasing in during weeks 0 and 1 and stabilizing at around 35% in week 2. Appendix Figure A10 explores longer-run effects using an alternative specification in which we match households experiencing health shocks three years apart rather than one year apart. That figure suggests lasting effects that grow larger three years after the event. As mentioned above, Appendix Figure A4 demonstrates that

[15] We study shock spouses of both genders together. Appendix Figure A5 presents results separately by gender. The effects are roughly similar regardless of the gender of the shock spouse, leading us to pool genders here. In the Appendix, we also show heterogeneity in our results by the age of both shock and outcome spouses (Appendix Figure A6), race and ethnicity (Appendix Figure A7), and the length of stay associated with the index heart attack or stroke (Appendix Figure A8).

effects persist when we do not condition on survival. Finally, Appendix Figure A11 takes a complementary approach that provides the "flow" counterpart to these "stock" dynamics: it plots the likelihood of *entry* into a nursing home or Medicare-covered SNF. We observe a sharp spike in entry—exceeding 100% of the baseline admission rate—during the month of a partner's health shock, and then sustained increases of around 20% in entry into nursing homes over the year following the shock. Taken together, these analyses show large, immediate, and persistent increases in the spouse's utilization of formal care.

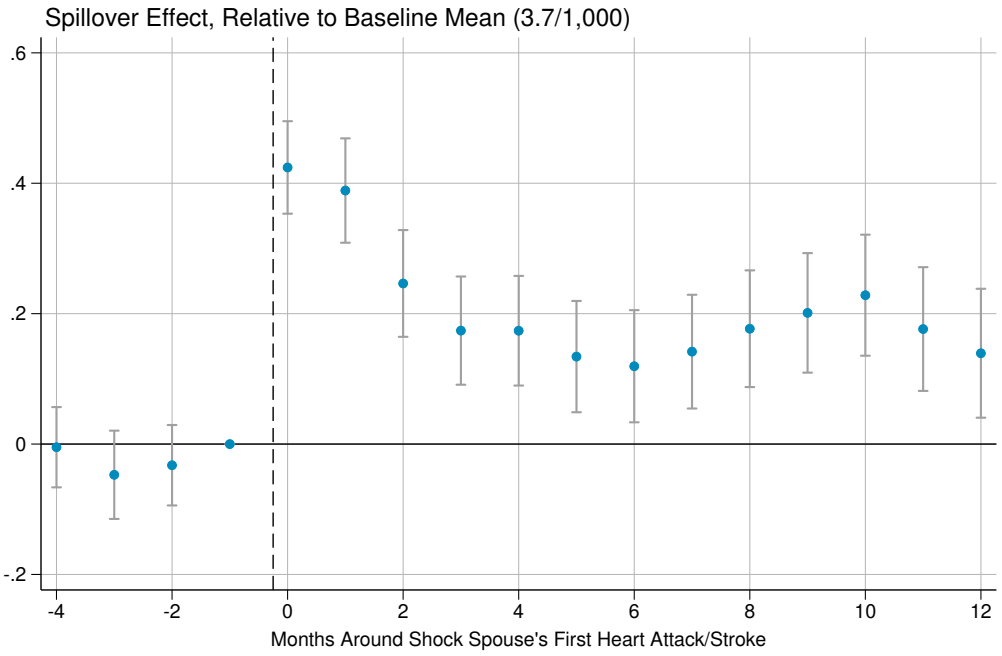
Figure 1. Effect of a Spouse's Major Cardiovascular Event on Risk of Any Nursing Home Visit



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 was visiting any nursing-home facility in a given period, including Medicare-covered SNF stays and all other nursing-home visits identified in the MDS. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stay or entry. 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.

Figure 1 presents estimates for any nursing-home visit: either a short-stay visit—a visit to a SNF—or a long-stay visit—a visit to a nursing home. SNF visits are of particular interest: they are covered by Medicare for up to 100 qualifying days, and account for $27.3 billion in 2023 Medicare spending (Centers for Medicare & Medicaid Services, 2026). For that reason, Figure 2 presents estimates solely for Medicare-covered SNF stays using MedPAR data.

Figure 2. Effect of a Spouse's Major Cardiovascular Event on the Probability of a SNF Stay



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 was visiting a SNF for a Medicare-covered stay in a given period. We rescale all coefficients such that they indicate the change relative to the baseline risk of a 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 pattern for SNF visits resembles the effects for all visits for the first 100 days following the shock spouse's event. Unlike the effect for all nursing-home stays, however, the SNF-only effects in Figure 2 decline over time, with the risk highest during the first two months. We observe a smaller average 20% increase beyond the 100-day mark, when full Medicare coverage for post-acute SNF care typically ends.[16] This divergence, combined with the patterns in Figure 1, suggests an accumulation of outcome spouses who exhaust Medicare-covered benefits and transition into longer-term nursing-home stays financed out-of-pocket, by Medicaid, or by private insurers.[17]

[16] Appendix Figure A11 highlights that SNF entry closely follows the spousal health shock, so that these two clocks overlap.

[17] For completeness, Appendix Figure A13 presents the corresponding event study for all nursing-home stays not covered by Medicare. As expected, the effects on these non-Medicare-covered stays increase over time.

3.2 Evidence on Substitution from Informal to Formal Care

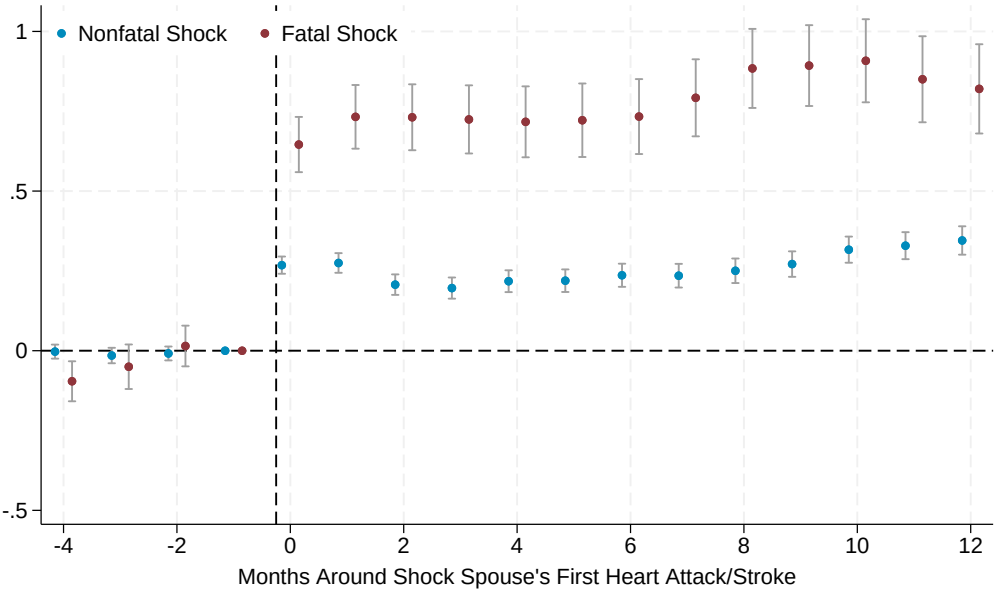
A key question is whether the results above are driven by the substitutability of informal and formal care. Two alternative explanations arise. First, the effects we estimate could be direct effects of the shock spouse's health shock itself on the outcome spouse, unrelated to the loss of care from the shock spouse. A spousal health shock can be a major short-term disruption, and, consistent with prior work, such a disruption could have real, immediate health effects for those with pre-existing conditions, stemming from emotional distress, deterioration in sleep quality, and disruptive effects of care transitions (Bünnings et al., 2021; Arteaga et al., 2025). Second, the loss of care from the shock spouse could cause a long-term deterioration in the health of the outcome spouse due to differences in the quality of care delivered by a spousal caregiver versus alternative sources. In both cases, outcome spouses could be more likely to visit nursing homes, but not as a matter of substituting formal care for informal spousal care.

We next differentiate between those alternative explanations in three steps. First, we explore heterogeneity in the effects of spousal health shocks that points toward the loss of spousal caregiving as the key mechanism. Second, we estimate the direct effects of a partner's cardiac event on the outcome spouse's health to assess whether there is evidence for this pathway. Finally, we develop a decomposition exercise, dividing marginal SNF stays into those that are likely to be caused by a deterioration in outcome spouse health versus those that are more likely to reflect substitution conditional on health.

Heterogeneity by Caregiver Availability and Care Needs. We first stratify the sample by whether or not the shock spouse died. The death of a spouse constitutes a permanent shift in the availability of informal care. This group thus offers a useful comparison for assessing long-term impacts: for nonfatal heart attacks or strokes, the shock spouse may be able to return to caregiving, but fatal spousal shocks preclude any return to informal spousal care. Figure 3 presents the results of that stratification. Outcome spouses who lose their partners face almost triple the risk of staying in a nursing home—over 75% on average. By contrast, those whose partners did not die within one year experience a roughly 25% increase in the risk of a nursing-home stay.

Appendix Figure A14 presents the same analysis separately for couples in which the shock spouse died in the first, second, third, or fourth quarters after their cardiac event. All groups exhibit statistically significant increases in the likelihood of

Figure 3. Heterogeneous Effects of Spousal Shock, by Shock Fatality



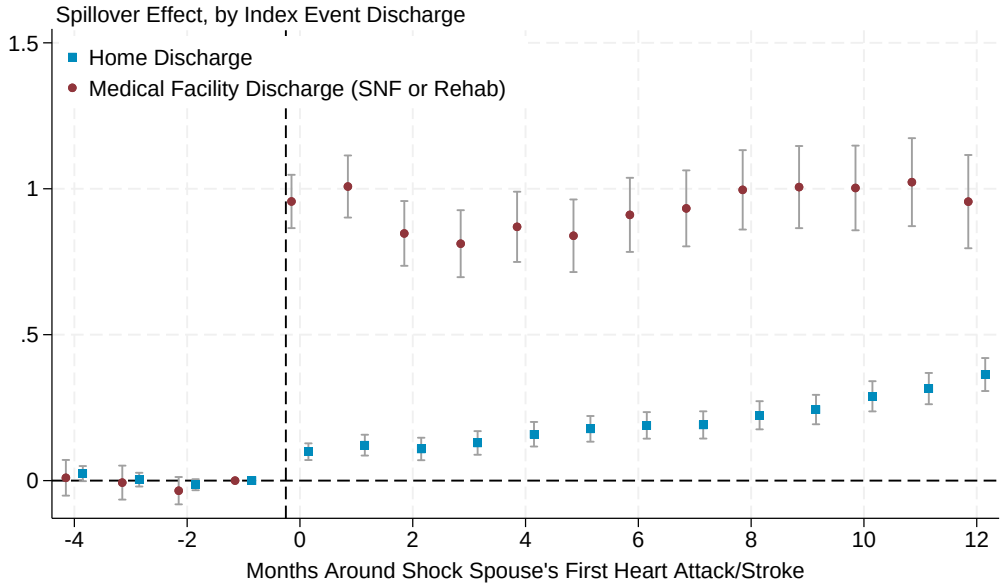
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The groups presented indicate whether a shock spouse passed away within 1 year of discharge following their major cardiovascular event. The outcome of interest is an indicator for whether the outcome spouse was visiting any nursing home in a given period. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stays. 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.

a nursing-home visit in the months after their partner's incapacitation, regardless of when the partner died. That is, outcome spouses whose partners die nine to twelve months later are just as likely to enter a nursing home in the month of the event as those whose partners die immediately. That pattern is important given prior work that has shown direct effects of spousal death (e.g., Fadlon et al. 2026): Appendix Figure A14 suggests instead that in this case, the key driver inducing outcome spouses to visit a nursing home is the loss of access to informal care rather than other effects of their partner's death.

Second, we stratify by the severity of the shock spouse's incapacitation among cases in which the shock spouse survived. We do so to capture the degree to which the shock spouse may still be able to provide informal care.[18] We divide households

[18]For this exercise, we limit the sample to couples whose outcome spouse was not hospitalized in the three months immediately preceding the shock spouse's cardiac event. We do so to ensure that we classify couples by the severity of the shock spouse's incapacitation and not by the status of the outcome spouse. The concern, otherwise, is that the outcome spouse's status could lead the shock

Figure 4. Effect on Risk of a Nursing Home 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. The analysis is restricted to nonfatal index events. We also exclude households whose outcome spouses experienced a hospitalization in the three months 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 any nursing home during a period. Estimates are stratified by the shock spouse's discharge destination. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stays in each category. The error bars plot 95% confidence intervals based on standard errors clustered at the household level.

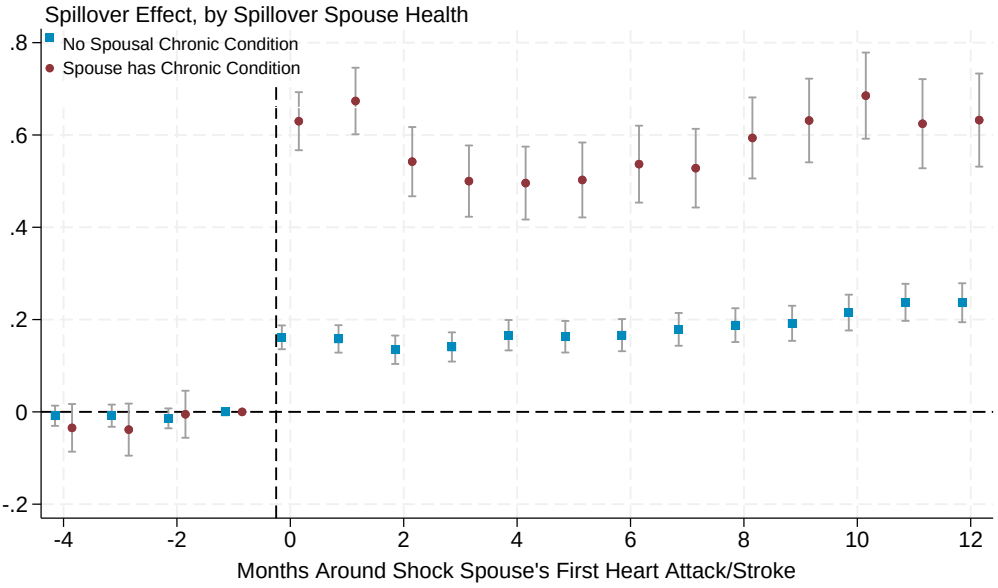
by whether the shock spouse was discharged home or to a medical facility, either a rehabilitation facility or a SNF. Based on that proxy, Figure 4 suggests that the more severe the spouse's incapacitation, the more likely outcome spouses are to utilize formal care. The risk of visiting a nursing home is, on average, five times as large for outcome spouses when their partners are discharged to a facility as when their partners are discharged home.[19]

spouse to be discharged to an institution rather than home.

[19]One concern is that transitions to formal care may reflect an outcome spouse's preference to co-locate with the shock spouse. Appendix Figure A12 decomposes the main effects into indicators for whether an outcome spouse resides in the *same* facility to which the shock spouse was discharged. In the short run, when the shock spouse is discharged to a SNF, outcome spouses who also enter a SNF enter the same facility as their partner at relatively high rates. Over time, however, shock spouses leave facilities and the increased use of formal care is mainly driven by couples where the outcome spouse is in a SNF while the shock spouse is not. Thus, it appears that if the shock spouse's event is severe enough to land them in a SNF, it is often the case that the spouses co-locate. This is not surprising, as spouses who both require institutional care are likely to opt to receive that care together. But in the long run, in most cases the shock spouse is not in a SNF, making the co-location hypothesis less relevant.

Finally, we stratify by the pre-existing health of the outcome spouse. If the increase in nursing-home visits is driven by the loss of spousal caregiving, then such effects should be concentrated among outcome spouses who require care. Figure 5 plots estimates of equation (1) for nursing-home visits separately by whether the outcome spouse had a chronic condition before their partner's health shock.[20] These conditions proxy for a pre-existing need for care from a spouse.

Figure 5. Effect on Risk of Any Nursing Home Visit by Outcome Spouse's Pre-existing 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 was visiting a nursing home in a given period. We stratify the sample based on whether the outcome spouse had a documented chronic condition during the year prior to the shock 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 baseline risk of a nursing-home stay in each category. Error bars plot 95% confidence intervals based on standard errors clustered at the household level.

Outcome spouses with no underlying conditions exhibit a smaller, but still statistically significant, increase in nursing-home visits: an average increase of 20%. By contrast, outcome spouses with chronic conditions exhibit a roughly threefold

[20]For this measure, we rely on the Medicare Chronic Conditions file, which records whether each Medicare beneficiary has any of 27 chronic conditions. We identify conditions that are more likely to require informal caregiving: Alzheimer's disease and related dementias (ADRD), cancer (colorectal, lung, or prostate), chronic kidney disease, chronic obstructive pulmonary disease, hip fractures, and strokes.

larger increase. Furthermore, Appendix Figure A15 documents heterogeneity in the estimated effects even across chronic conditions. For example, outcome spouses diagnosed with Alzheimer's disease and related dementias (ADRD) prior to their partner's shock are more than 200% more likely to use formal care after the shock, even though they have a higher pre-event utilization rate. This heterogeneity is consistent with the hypothesis that outcome spouses whose pre-existing needs for informal care are greater are more likely to transition to formal care after their partner's health shock.

Effects on the Outcome Spouse's Health. The heterogeneity above is consistent with the loss of a spousal caregiver, but it cannot distinguish two channels. Outcome spouses may substitute formal care for the informal care they have lost, at an unchanged level of health. Alternatively, their partner's cardiac event may have worsened their health, raising their need for care of any kind. We therefore estimate the health effects first, and then ask how much of the increase in SNF use they explain.

We first investigate a potential health effect by studying the mortality of the outcome spouse. To do so, we rely on a different sample: an unbalanced panel in which we require only that the outcome spouse be alive twelve months prior to the shock, when we identify the couple.[21] We then assess how mortality rates diverge for treatment and control outcome spouses after the shock spouses in the treatment group experience their cardiac event. Figure 6 plots those estimates. The figure presents an event-study analysis in which the outcome of interest is the death of the outcome spouse, stratified by whether the shock spouse died. Mortality rates increase by nearly 20% in the six months after the shock spouse's cardiac event, with effects roughly three times as large if the shock spouse died.

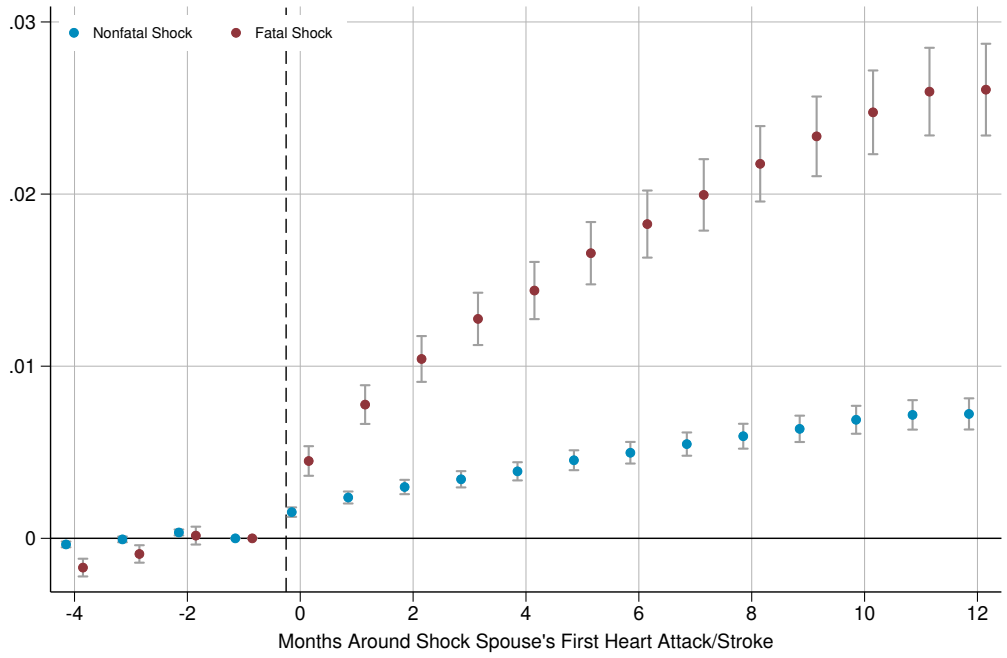
That persistent increase in mortality is partially a demonstration of "the widowhood effect."[22] For over a century, since at least the work of March (1912), researchers have documented higher mortality rates among those who have lost their spouses. More recently, Elwert and Christakis (2008) document that mortality rates rise especially in the first six months after the spouse's death.

In order to measure changes in health beyond mortality, consider Table 2. The table presents estimates of equation (2) in which the outcomes are a nursing-home visit,

[21] It is important to make this restriction several periods prior to the shock in order to be able to investigate the comparability of mortality trends in the treatment and control groups prior to the event.

[22] It is a partial demonstration, because we also observe small mortality effects when the shock spouse survives.

Figure 6. Effect of a Shock Spouse's Major Cardiovascular Event on Outcome Spouse's Mortality



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 has died, using mortality as an absorbing state. Coefficients are presented in levels (the baseline pre-treatment mean is 0.008). 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.

a SNF visit, any inpatient stay, inpatient stays for monitoring-sensitive conditions, a fall-related encounter, or an emergency department (ED) encounter. We define *monitoring-sensitive hospitalizations* as diagnoses associated with concerns about elder neglect, consistent with prior literature (Hoover and Polson, 2014; Centers for Disease Control and Prevention, 2024).[23] For each outcome, the table reports the treatment effect immediately following the event as well as the average treatment effect over the first twelve months using equation (2). All coefficients are normalized

[23]Monitoring-sensitive conditions are those whose prevention or early treatment depends on observation and caregiver responsiveness. They include falls as well as encounters coded with the following ICD-10 diagnosis codes: malnutrition (E43.*, E44.*, E46.*) and adult failure to thrive (R62.7); dehydration and fluid/electrolyte imbalance (E86.*, E87.*) and related intake symptoms (R63.*); pressure injuries (L89.*); other deprivation-related effects (T73.*); and untreated infection, including urinary tract infections (N39.0) and sepsis (A41.*). These conditions capture health events that are more likely to occur or go undetected following a disruption in spousal caregiving.

by each outcome's baseline mean.[24]

Table 2. Impact of a Spouse's Major Cardiovascular Event on Healthcare Utilization

	(1)	(2)	(3)	(4)	(5)	(6)
			Hospitalizations			
	Nursing Home	SNF	Any	Monitoring-Sensitive	Falls	ED Visits
Effect, Month 0	0.324***	0.436***	0.180***	0.216***	1.363***	0.227***
	(0.0136)	(0.0367)	(0.0194)	(0.0391)	(0.4012)	(0.0402)
Average Effect, Months 0–12	0.327***	0.237***	0.082***	0.086***	0.221	0.120***
	(0.0126)	(0.0267)	(0.0108)	(0.0255)	(0.1957)	(0.0203)
Baseline Rate per 1,000	12.2	3.7	15.3	4.2	0.06	26.3
N	16,789,284	16,789,284	16,789,284	16,789,284	16,789,284	3,357,857

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). ED visits are measured as the total number of visits. All other outcomes are binary. Note that ED visits are measured using the 20% sample of Medicare beneficiaries, hence the reduced sample size. Monitoring-sensitive conditions are defined as hospitalizations with specific diagnostic codes related to potential elder neglect, as discussed in the text. Treatment effects in the first row are estimated over the first four weeks post-event; average effects over the first year post-event are presented in the second row. We rescale all coefficients such that they indicate the change relative to the baseline risk or rate of each outcome and cluster standard errors at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects. Appendix Figure A16 presents corresponding event-study plots for columns (3) through (6). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Overall, outcome spouses exhibit an increased risk of all events. The effects are most pronounced in the month following the index event, but they persist over time. The table shows an 18.0% increase in hospitalization risk and a 21.6% increase in the risk of hospitalization for monitoring-sensitive conditions (which constitute roughly 27.5% of hospitalizations in this sample). We also observe a 22.7% increase in the number of ED visits. Average effects over the twelve months following the event are also statistically and clinically significant, including an 8.2% increase in hospitalization rates, an 8.6% increase in monitoring-sensitive hospitalizations, and a 12.0% increase in ED visits.

[24]For completeness, Appendix Figure A16 provides the analogous event-study figures.

Decomposing the Increase in SNF Use. The results above document a deterioration in outcome spouses' health after their partners' cardiac event, raising the possibility that the increase in SNF stays is actually due to those health effects rather than due to substitution of formal for informal care.[25] To that end, we next attempt to decompose the increase in SNF stays into stays driven by a deterioration in the outcome spouse's health and those driven by other causes.

To implement this decomposition, we focus on Medicare-covered SNF stays, noting that a SNF stay must follow a hospitalization. Medicare-covered SNF stays account for roughly 22% of the overall increase in nursing-home stays documented above. Though a minority of that increase, they represent a policy-relevant margin in this setting, because they are one of the few types of long-term care covered by Medicare—a point we return to in Section 4, when we consider the optimal design of public health insurance programs. The decomposition, however, applies only to these Medicare-covered stays and may not generalize to other nursing-home stays.

The decomposition rests on the following logic. The increase in SNF stays arises from three possible sources. First, a shock spouse's cardiac event could cause a hospitalization for the outcome spouse that would not have occurred otherwise, which then ends in a SNF stay. Second, the cardiac event could make an otherwise inframarginal hospitalization—one that would have occurred regardless—more severe, and so more likely to end in discharge to a SNF. Third, an inframarginal hospitalization could be unchanged in severity but still more likely to end in a SNF stay, because the outcome spouse who would have been discharged home now has no caregiver there. This last category most clearly reflects substitution, while the others could reflect health effects.[26]

Those three categories motivate the decomposition that we implement. As a benchmark, we first measure how much of the increase in SNF stays that we observe can be explained solely by the additional hospitalizations documented in Table 2 above. We apply the pre-existing probability of discharge to a SNF to those additional hospitalizations. That calculation captures SNF stays arising from marginal

[25]The increase in mortality motivates our focus on a balanced panel of couples for most analyses. Without imposing that restriction, we still see increases in nursing-home stays, but they fade over time (Appendix Figure A4). The fact that the effects do not fade for the balanced panel but do for the unbalanced panel indicates that the fading out is likely due to the competing risk of death.

[26]Marginal hospitalizations could also be interpreted as substitution, because inpatient care may stand in for informal spousal care, particularly for less severe hospitalizations. To be conservative, however, we assume that all marginal hospitalizations reflect health effects. Doing so understates the role of substitution.

hospitalizations—the first source described above—but does not allow the composition or severity of hospitalizations to change.

Next, we allow the shock to affect both the number and the composition of the outcome spouse's hospitalizations. For each hospitalization, we predict the probability of discharge to a Medicare-covered SNF using the details of the hospitalization (such as its primary diagnoses) and observable characteristics of the patient's health. Changes in predicted SNF use capture the portion of the increase attributable to additional hospitalizations or changes in their observable severity, the first and second sources described above. The difference between actual and predicted SNF use is the residual component, which we interpret as substitution from informal to formal care conditional on observed health, the third source. By construction, the predicted and residual components sum to the observed change in SNF use. The predicted component can nevertheless be negative if post-event hospitalizations are of types that are less likely to lead to SNF care than pre-event hospitalizations.[27] In that case, the decomposition will attribute nothing to observable health effects, with the entire observed increase attributed to the substitution component instead.

We study all hospitalizations in the data and predict the probability that the patient is discharged to a SNF. We train the model on a 90% sample and evaluate it on the remaining 10%. Once trained, we use the model to evaluate how predicted probabilities evolve around partners' health shocks. That is, we assign each individual-month a "predicted SNF stay" measure, which is equal to zero in months without a hospitalization. In months where a hospitalization occurred, the predicted SNF stay is estimated as the average SNF-discharge probability among all hospitalizations with similar observable characteristics. The prediction therefore captures how we would expect outcome spouses' SNF use to change after shock spouses' health shocks, but based solely on the outcome spouse's health during the precipitating hospitalization. The prediction captures both changes in predicted SNF stays due to marginal hospitalizations—those that occurred due to the shock spouse's cardiac event—and inframarginal hospitalizations—those that would have occurred either way but whose severity could be impacted by the absence of the spousal caregiver.

We estimate our prediction model using a lasso logistic regression of a binary indicator for any SNF discharge on patient age, sex, race, enrollment type (e.g., enrollment in a Medicare HMO or dual Medicaid eligibility), and the diagnosis-related

[27] For example, the spouse's shock could produce more frequent but milder hospitalizations than those observed before the event. Appendix Figure A17 illustrates this possibility.

group associated with the hospitalization, as well as year and month of admission fixed effects and all two-way interaction terms (Ellis et al., 2022). We employ the lasso approach to account for the high dimensionality of our covariate space (Abadie and Kasy, 2019). The penalty parameter is selected via three-fold cross-validation, whereby the sample is automatically partitioned into estimation and validation folds. Appendix Figure A18 presents a histogram of the predicted probabilities from the final model.[28] Reassuringly, those who are discharged to a SNF are assigned higher predicted values than those discharged elsewhere.

Table 3 compares estimates of equation (2) across actual and predicted probabilities of an outcome spouse visiting a SNF for post-acute care. The table presents those estimates in levels to highlight the decomposition of the total SNF effects into those captured by the two prediction models and residual variation. Column (2) provides the "naive" average-probability model (marginal hospitalizations only) as a benchmark, and column (3) provides our lasso predictions (marginal hospitalizations + changes in predicted discharge for inframarginal hospitalizations). In the month following the index shock, SNF stays for outcome spouses rise by an average of 1.58 stays per 1,000. Of this estimated increase, around one-third could be explained by an increase in marginal hospitalizations. However, when incorporating changes in the composition of hospitalizations, the model only predicts one-fourth (= 0.40/1.58) of the observed increase in SNF stays, so that 75% of the immediate increase cannot be explained by changes in hospitalization rates or type.[29] Over the full year post-event, the explanatory power of a health effect for the increase in SNF stays effectively vanishes.

This analysis provides an illustrative exercise for assessing health and substitution effects. It assumes that the only pathway from worsening health to a SNF is via hospitalization and it ascribes the residual to a substitution effect. One important takeaway of this analysis is that it offers an estimate of the contribution of observable health deterioration to the increase in SNF use, which we estimate to be relatively small. Overall, we view this exercise as suggestive evidence that the majority of the rise in SNF stays that we catalog is driven by substitution from informal to formal care. This exercise does not speak directly to the increase in non-Medicare-covered

[28]The estimates are similar when, instead, we rely on a simple prediction model based only on previous hospitalizations and SNF stays without a lasso model, or use other prediction algorithms such as ridge regression or elastic net.

[29]This implies that the mix shifted toward hospitalizations less likely to end in discharge to a SNF. This is consistent with marginal hospitalizations being for relatively minor events that typically do not result in post-acute care.

Table 3. Decomposing the Change in the Outcome Spouse's Probability of SNF Discharge

	(1)	(2)	(3)
		Predicted SNF Effects	
	Full Effect	Constant Average Effects	lasso
SNF Effect, Month 0	1.58***	0.50***	0.40***
	(0.135)	(0.055)	(0.054)
Average SNF Effect, Months 0–12	0.88***	0.23***	-0.10***
	(0.100)	(0.031)	(0.025)
N	16,789,284	16,789,284	16,789,284

Notes: This table presents estimates of the δ_0 coefficients from equation (1) and the δ coefficients from equation (2), difference-in-differences estimates that track an outcome spouse's SNF use in the months since a shock spouse's first heart attack or stroke. Column (1) shows the overall SNF effects in levels (rates per 1,000 individuals), following Figure 2. Column (2) indicates the predicted SNF effect that arises solely from changes in the volume of hospitalizations, assuming each hospitalization has the same average probability of discharge to a Medicare-covered SNF stay. Column (3) constructs the predicted probability of a SNF discharge based on hospitalization and patient characteristics using lasso regression. We report effects in levels here to illustrate the decomposition, and cluster standard errors at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects. Appendix Figure A19 presents corresponding event study plots for columns (2) and (3).

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

nursing-home stays, which we cannot decompose in the same way. We see no reason to expect substitution to matter less on that margin, but we cannot test it directly.

3.3 The Fiscal Value of a Healthy Spouse

The estimates above demonstrate increases in outcome spouses' healthcare utilization: SNF admissions, hospitalizations, and ED visits. Those results naturally lead to the question of how much the shock spouse's cardiac event increases their partner's overall healthcare expenditures.

Table 4 presents estimates of regressions that track the impact of a shock spouse's cardiac event on overall Medicare-covered spending for the outcome spouse. In order to study spending, we rely on the annual summary measures that are tracked for all fee-for-service Medicare recipients.[30] In the year of a shock spouse's event, total

[30]The regressions in Table 4 are at the annual level, and so involve only indicator variables for year relative to the shock spouse's major cardiac event. To appropriately estimate fiscal externalities taking mortality into account, here we relax the requirement that the outcome spouse survive at

Medicare payments increase by roughly 4% ($252), and by an average of 9% ($578) for fatal index events. The table suggests increases in healthcare expenditures across categories that are large in both absolute and relative terms. For instance, spending on SNFs increases by $95 in the year of the event, roughly a 28% 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 larger 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 1.7% decrease in hospital outpatient care, and a 7.9% decrease in elective surgeries. We also observe declines in spending on prescription medication (2.1%). The decline in prescription medicine may be mechanical: medications administered in a SNF are included in Medicare's Part A payment to the facility rather than billed as Part D claims.

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 people in the US are covered by Medicare, 58% of whom are married or living with a partner. The incidence rate of heart attacks and strokes among this group is roughly 5% annually and 15% of them are fatal within 30 days. The estimates in Table 4 thus translate into a national spillover effect of $558.7 million annually, attributable to heart attacks and strokes alone.[31]

3.4 Sensitivity to the Price of Formal Care

The results above show that the *level* of demand for formal long-term care is higher when one's spouse is incapacitated. What about the slope? This section studies how a spouse's incapacitation affects their partner's price elasticity of demand for formal care. That elasticity is a critical parameter for welfare analysis and the optimal level of insurance coverage, as originally pointed out by Zeckhauser (1970) and as we document in this setting. To estimate the price elasticity of demand, we rely on the difference-in-discontinuities design described above (equation (3)), which exploits the discrete change in cost sharing that patients face for Medicare-covered SNF stays

least twelve months after the shock spouse's health event.

[31] The corresponding increase in spillover spending for nonfatal heart attacks and strokes is $244 annually. Hence, we calculate the total spillover effects of the approximately 1.9 million heart attacks or strokes affecting married or partnered individuals on Medicare as the weighted average 1.9 million $\times$ ($0.15 \times \$578 + 0.85 \times \$244$) = $558.7 million.

Table 4. Effect of a Shock Spouse's Major Cardiovascular Event on Outcome Spouse's Annual Medicare Spending

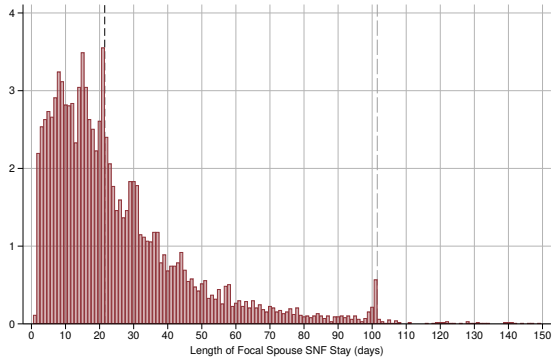
		Year of Index Event	
Spending Measure	Pre-Treatment Average	(1) All	(2) Fatal Events
A. Total Spending			
	\$6,438	\$252***	\$578**
		(36.93)	(236.14)
B. Long-term Care			
SNF	\$339	\$95***	\$303***
		(8.44)	(61.68)
Home Healthcare	\$283	\$24***	\$31
		(3.64)	(24.16)
Hospice	\$125	\$48***	\$65*
		(5.58)	(37.91)
C. Hospital & Surgical Care			
Acute Inpatient	\$1,458	\$123***	\$238**
		(15.73)	(102.12)
Other Inpatient	\$178	\$20***	\$34
		(6.59)	(40.82)
Hospital Outpatient	\$985	-\$17**	-\$14
		(7.48)	(40.70)
Ambulatory Surgical Center	\$76	-\$6***	-\$6
		(1.14)	(6.21)
D. Other Care			
Physician Payments	\$337	-\$10***	\$3
		(0.75)	(4.65)
Evaluation & Management	\$243	\$25***	\$50***
		(2.18)	(13.98)
Part D Spending	\$896	-\$19***	-\$53**
		(4.84)	(25.99)

Notes: This table reports pooled estimates of the effects in the year of a shock spouse's first heart attack or stroke (equation (2)). The outcome of interest is annual Medicare payments per beneficiary, across the spending categories indicated in each row. Regressions include calendar year 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 of the index event, and standard errors are clustered at the household level. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$

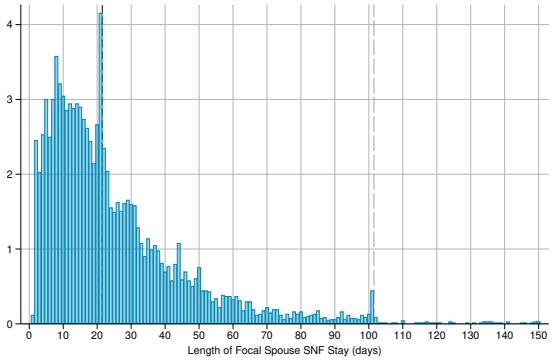
after the 20$^{\text{th}}$ day.

Figure 7. Responses to Cost Sharing for SNF Stays by Spousal Health-Shock Status

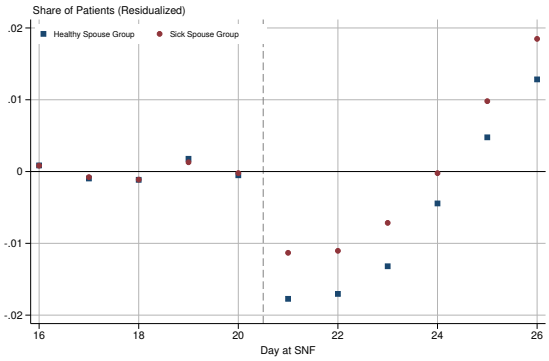
(a) Sick-Spouse Group



(b) Healthy-Spouse Group



(c) Regression Discontinuity, First Stage



Notes: This figure illustrates the identifying variation for the regression-discontinuity design (equation (3)). Panels (a) and (b) show the distribution of length of stay (LOS) for outcome spouses visiting SNFs within four months of the index event. Panel (a) provides the distribution among beneficiaries experiencing an actual spousal event (the "sick spouse" group), and panel (b) provides the corresponding distribution for those with a placebo spousal event (the "healthy spouse" group). Vertical lines indicate the days at which Medicare payment rates for SNF stays change discontinuously. In panel (c), we plot binscatter 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.

Figure 7 illustrates the variation we exploit. It plots the distribution of SNF lengths of stay across the treatment and control groups. Specifically, we compare the lengths of Medicare-covered SNF stays for the "sick-spouse" group, i.e., outcome spouses whose partner experienced a health shock in the last four months, with those for the "healthy-spouse" group, i.e., outcome spouses whose partner will experience a health shock in the future. Panel (a) plots the distribution for outcome spouses

in the sick-spouse treatment group, and panel (b) plots the distribution for outcome spouses in the healthy-spouse group. Both of those panels document "bunching" around day 20. The bunching appears more pronounced for outcome spouses with a healthy spouse, suggesting a higher level of price sensitivity for this group.

Panel (c) of Figure 7 plots mean residuals from a regression of a binary indicator for a SNF discharge on day-of-week-specific fixed effects and a linear trend for days in the SNF so far. Overall, the figure demonstrates that Medicare recipients are much less likely to stay in the SNF when cost sharing begins on day 21. More importantly, those with a healthy spouse are more likely to leave the SNF after day 20 than those whose partner has experienced a major incapacitating cardiac event in the last four months.[32]

Table 5 presents estimates formalizing this comparison using the difference-in-discontinuities design of equation (3). Prior to the shock spouse's cardiac event, the discontinuous change in copayments decreases the probability that outcome spouses remain in the SNF by 1.8 percentage points. In comparison, the response to copayments declines to just 1.0 percentage points when the outcome spouses have been admitted following their partner's cardiac event.[33] Those two estimates suggest that outcome spouses are 0.8 percentage points more likely to remain in a SNF once copayments commence, a difference that is 44% of the baseline control-group response.

Importantly, the difference in length of stay only emerges once copayments begin. Conditional on a SNF stay and adjusting for the day of the week, 44.47% of outcome spouses stay through day 20, compared with 44.56% of those in the control group, a statistically insignificant difference. That small pre-day-20 gap suggests that the two groups are similar. Because they are similar through day 20 and face the same price change on day 21, the ratio of their responses translates directly into a ratio

[32] As a complementary exercise, Appendix Figure A20 presents event-study estimates similar to those in Section 3.1, but with the outcome of interest being a stay in a SNF beyond 20 days. The figure suggests a persistent doubling in the likelihood that outcome spouses stay in SNFs beyond 20 days after their partners experience a major cardiovascular event. That implies that nearly a quarter of outcome spouses' SNF stays require payment beyond what Medicare covers. Similarly, Appendix Figure A21 documents that outcome spouses stay in SNFs for longer after their partners' cardiac events.

[33] This local-linear RD approach assumes that counterfactual discharge probabilities evolve smoothly over time. In a robustness exercise, we estimate an alternative specification adjusting for relative day-specific effects at baseline to accommodate discharge hazards in the control group that potentially vary by day. As we discuss below, conditional discharge hazards between the two groups are remarkably similar prior to day 21, as shown in Appendix Figure A22. The point estimate from this approach is -0.0188, with a standard error of 0.0032, so that the resulting estimates match those reported in Table 5 within 0.001.

Table 5. Effects of Spousal Incapacitation on SNF Price Sensitivity, by Outcome Spouse Chronic Condition Status

	(1) Pooled	(2) No Chronic Condition	(3) Has Chronic Condition
RD Coefficient, Healthy Spouse	-0.018***	-0.019***	-0.018***
	(0.0030)	(0.0045)	(0.0038)
RD Coefficient, Sick Spouse	-0.010***	-0.017***	-0.007**
	(0.0026)	(0.0041)	(0.0033)
Difference in Discontinuities	0.008**	0.001	0.011**
	(0.0040)	(0.0062)	(0.0051)
N	380,673	139,702	240,971

Notes: This table presents regression-discontinuity and differences-in-discontinuities estimators identifying the effect of ending full Medicare coverage for SNF stays, which ends on day 20 for qualifying stays (equation (3)). The first two rows present the estimated effect of losing full coverage on the probability an outcome spouse will remain in the SNF, stratified by group. The third row presents the difference-in-discontinuities estimator. 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. The estimation includes calendar month fixed effects and person-specific fixed effects. The MSE-optimal bandwidth is 11.22. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of elasticities: utilization of formal care is roughly 1.8 times as price elastic when a caregiver remains available at home.

Columns 2 and 3 of Table 5 stratify the price-sensitivity analysis across outcome spouses with and without chronic conditions. The results are almost entirely driven by outcome spouses with chronic conditions. These results bolster the interpretation that having a partner at home is an important determinant of the price elasticity of formal care. Further support for this interpretation is provided in the appendix with heterogeneity analyses along additional informative dimensions.[34]

[34]Appendix Table A1 further stratifies these findings by the type of cardiac event that the shock spouse experienced. Price sensitivity is lowest among beneficiaries whose spouses died following their cardiovascular event. Among this group, we observe no price responsiveness, a statistically significant difference compared to those whose spouse survived their cardiac event. Appendix Tables A2 and A3 further stratify the sample by gender and ADRD status. Price effects are larger for female outcome spouses and almost three times as large for outcome spouses with ADRD as for those without ADRD, although these estimates are less precise. Finally, Appendix Table A4 shows that these changes in demand elasticities occur throughout SNF stays over the first year after a partner's health shock.

4 Implications for Optimal Insurance

This section explores what the results above imply for policy in the following steps. Section 4.1 develops a model of the household's choice between formal and informal long-term care and shows how the gap in consumer surplus across the healthy- and sick-spouse states bounds the value of informal care from below. Section 4.2 solves the planner's problem and derives the optimal ratio of subsidies across the two states. Sections 4.3 through 4.5 then bring the empirical estimates to that model: they value an extra day of formal care on the 21st day of a SNF stay, extend the analysis to the full stay with a dynamic model of counterfactual benefit designs, and ask how such a contract could be implemented in practice.

4.1 Model Setup

Consider a Medicare beneficiary's decision to utilize formal care. Given the reduced-form estimates above, we first focus on the decision of a Medicare beneficiary who has resided in a SNF for 20 days and is deciding whether to stay for the 21st day.[35] Consider two states indexed by $h \in \{0, 1\}$. Recipients in state $h = 1$ have a healthy spouse present, and recipients in state $h = 0$ do not have that option. We order Medicare recipients by their state-dependent valuation of formal care. We denote this index by $s_h \sim U(0, 1)$, where a recipient with $s_h = 0$ has the highest valuation for formal care in state h, and a recipient with $s_h = 1$ has the lowest valuation of formal care in state h.[36]

Medicare charges the beneficiary a price p for an additional night in a SNF after day 20. Given p, there exists a marginal type of recipient s_1^* 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_0^* who has an incapacitated spouse and is similarly indifferent.

Define $V_h(s_h)$ as the incremental value of the 21st day in the SNF for a recipient of type s_h in state h. We assume that the household of a recipient who remains an additional night in the SNF enjoys flow utility $V_h(s_h) + U_h(Y_h - p)$, where $U_h(\cdot)$ is the state-dependent household utility from non-medical consumption and Y_h is disposable income. We allow the utility function $U_h(\cdot)$ to vary by state to account for potential

[35] Section 4.4 below imposes structure on the analysis to extrapolate to other margins.

[36] That notation appears "reversed," that is, the lowest value of s_h indicates the highest valuation of formal care. But that notation simplifies the integrals that follow and is consistent with the work of Einav and Finkelstein (2011) and Finkelstein et al. (2019).

changes due to the health effect of spousal incapacitation.[37] In addition, disposable income may itself depend on family circumstances, so we let Y_1 and Y_0 denote the household's income when the spouse is healthy and when the spouse is incapacitated.

The marginal type $s_h^*(p)$ is therefore determined by the indifference condition:

$$V_h\big(s_h^*(p)\big) = U_h\big(Y_h\big) - U_h\big(Y_h - p\big). \tag{4}$$

That is, households with $s_h < s_h^*(p)$ have a higher valuation of formal care and so purchase the 21$^{\text{st}}$ day in the SNF, and those with $s_h > s_h^*(p)$ leave the facility.

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

$$V_h\big(s_h^*(p)\big) \approx U_h'(Y_h) \cdot p.$$

Normalizing by marginal utility converts $V_h\big(s_h^*(p)\big)$ into the money-metric incremental value of an additional night in a SNF. Define this value as $v_h(s_h) \equiv V_h(s_h)/U_h'(Y_h)$. In turn, equation (4) leads to the condition $v_h(s_h^*) = p$. This expression implies that, for the marginal household, the copayment for an additional night in a SNF is just equal to the dollar-valued benefit. We can then define the consumer surplus from the extra day in the SNF for beneficiaries as:

$$CS_h(p) = \int_0^{s_h^*(p)} \big[v_h(s_h) - p\big]\, ds_h.$$

With this setup, the demand curve for an additional night in a SNF, denoted by $Q_h(p)$, can be expressed simply as $Q_h(p) = s_h^*(p)$. Note that $-\partial CS_h(p)/\partial p = s_h^*(p) = Q_h(p)$, so $Q_h(p) = s_h^*(p)$ is just the indirect demand curve and we can estimate consumer surplus using the area under that curve. Importantly, we can estimate the demand curve using the RD estimates described above.

In addition, this model allows us to estimate the value of *informal* care. It does so even though there is no market for informal care, and so we cannot infer value from a demand curve for informal care. The logic is as follows. Formal care and informal care

[37] Several papers study state dependence in the marginal utility of consumption (Viscusi and Evans, 1990; Finkelstein et al., 2009, 2013; Leive and Sun, 2025). State dependence in preferences may be central to household saving and the demand for long-term care insurance (Lockwood, 2018; Ameriks et al., 2020) and has direct welfare implications in the design of safety-net policies (Landais and Spinnewijn, 2021; Coyne et al., 2024). Note that while we allow $V_h(\cdot)$ to also be state-dependent throughout, this is not central to our results.

are substitutes, so when the beneficiary's spouse is healthy, the beneficiary holds an option: they can either pay for formal care or draw on care from their spouse. When the spouse is incapacitated, the option of informal care disappears. That reduction in the beneficiary's choice set causes their willingness to pay for formal care to increase. As long as the price of formal care does not change, the beneficiary's consumer surplus from formal care increases as well. We define the increased valuation of formal care in the absence of the option of spousal care as the "option value" of spousal care:

$$\text{OV}(p) \equiv CS_0(p) - CS_1(p),$$

where $CS_0(p)$ indicates the consumer surplus from formal care when a spouse is incapacitated and $CS_1(p)$ indicates the consumer surplus from formal care when a spouse is healthy.

Appendix B demonstrates that $\text{OV}(p)$ offers a lower bound for the overall value of informal spousal care. The intuition is as follows. The object $\text{OV}(p)$ combines differences in consumer surplus across two groups: those who visit the SNF regardless of whether their spouse is incapacitated and those who only visit a SNF when their spouse is incapacitated. One could refer to those two groups as "always users" and "switchers." The always users are on the margin between formal care and informal care when their spouse is healthy and on the margin between formal care and an outside option when the spouse is sick.[38] As a result, the consumer surplus we capture for them when their spouse is healthy measures the value of formal care relative to informal care. Similarly, when their spouse is incapacitated, the consumer surplus we measure for them captures the value of formal care relative to the outside option. The difference between those two is their valuation of informal spousal care.

Switchers, by contrast, choose informal spousal care over formal care when their spouse is healthy. By revealed preference, they value informal care above the outside option by at least as much as they value formal care relative to that same option. Because they do not use formal care in the healthy-spouse state, they contribute nothing to $CS_1(p)$. Because the outside option is their only alternative in the sick-spouse state, their contribution to $CS_0(p)$ is exactly their surplus from formal care. For switchers, then, $\text{OV}(p)$ equals the surplus they forgo when they rely on their

[38] Some "always users" may prefer the outside option to informal care and so never rely on informal care. But those people do not have their value of formal care affected by the presence of a healthy spouse, and thus their value of formal care differences out in the calculation of $\text{OV}(p)$. Their value of informal care is also irrelevant, as they never use informal care because they value it less than the (always available) outside option.

spouse, which, by revealed preference, is no greater than the value they place on informal care itself. In total, then, the area between the demand curves for formal care in the sick- and healthy-spouse states, $\text{OV}(p)$, provides a lower bound on the full value of informal care.

The bound omits two components of that value. First, for switchers, $\text{OV}(p)$ captures the formal-care surplus they forgo when they choose spousal care, but not any (necessarily positive) additional premium they place on care from their spouse. Second, $\text{OV}(p)$ omits the value of spousal care for beneficiaries who never choose formal care in either state. These beneficiaries may value spousal care more than the outside option, but their choices never enter either formal-care demand curve. Their valuation is therefore not revealed by $\text{OV}(p)$. This omitted component will be larger when the loss of spousal care leads beneficiaries to choose the outside option rather than formal care. Conversely, when formal and informal care are close substitutes, switchers' additional premium will be smaller and fewer beneficiaries are likely to remain outside the formal-care market in both states. In that case, the lower bound will be tighter.[39]

Importantly, the utility of this method for valuing informal care is not specific to this setting. Instead, it can be used to value any unobserved good. The necessary ingredients are (1) variation in the availability of the unobserved good and (2) variation in the price of an observed substitute. Possible applications include other forms of household production, such as childcare.

4.2 Social Planner's Problem: Optimal Copayments

Beyond valuing informal spousal care, the model is useful for deriving the optimal insurance contract. To that end, we next turn to the social planner's problem. We allow the social planner to set different prices (i.e., insurance coverage) based on family structure. Specifically, we consider prices that depend on the health of each person's spouse, p_h, rather than a single price, p. Let μ_h be the share of the population in state h, with $\mu_0 + \mu_1 = 1$. Further, let c be the constant resource cost to Medicare of one day in a SNF, and R be the program budget. Beneficiaries who remain in the SNF enjoy $V_h(s_h) + U_h(Y_h - p_h)$, while those who leave consume their full income and

[39] Appendix B formalizes this argument.

obtain $U_h(Y_h)$. The social planner solves:

$$\max_{p_0,p_1} W(p_0,p_1) = \sum_{h=0}^{1} \mu_h \left[\int_0^{s_h^*(p_h)} \Big(V_h(s_h) + U_h(I_h) \Big)\, ds_h \;+\; \big(1 - s_h^*(p_h)\big)\, U_h(Y_h) \right],$$

subject to the expected-cost constraint:

$$B = \sum_{h=0}^{1} \mu_h \,(c - p_h) \int_0^{s_h^*(p_h)} ds_h \leq R,$$

where we define $I_h = Y_h - p_h$ as net income. We also define the price elasticity of demand at the margin, $\varepsilon_h \equiv -\frac{p_h}{s_h^*(p_h)} \frac{ds_h^*(p_h)}{dp_h}$, and the inverse markup capturing the degree of public subsidization, $\rho_h \equiv (c - p_h)/p_h$. Note that, at one extreme, as the consumer-facing price goes to zero, $p_h \to 0$, the coverage term goes to infinity, $\rho_h \to \infty$. At the other extreme, as the price converges to the cost, $p_h \to c$, the coverage term goes to zero, $\rho_h \to 0$. We thus take $-p_h$ as the generosity variable.

To characterize long-term care coverage at the optimum, we consider a budget-neutral perturbation in generosity. Specifically, we consider a balanced-budget perturbation that lowers the sick-state price p_0 and raises the healthy-state price p_1 to keep expected outlays fixed. The total impact on social welfare of such a perturbation, which we denote by ΔW, is captured by:

$$\Delta W \equiv \frac{\partial W}{\partial(-p_0)} + \frac{\partial W}{\partial(-p_1)} \frac{dp_1}{dp_0}. \tag{5}$$

Taking behavioral responses in demand for formal care into account—the canonical fiscal externality as applied to our setting—budget-neutrality implies:

$$\frac{dp_1}{dp_0} \;=\; -\,\frac{\mu_0 \cdot s_0^*(p_0) \cdot (1 + \rho_0 \varepsilon_0)}{\mu_1 \cdot s_1^*(p_1) \cdot (1 + \rho_1 \varepsilon_1)}.$$

Next, differentiating welfare with respect to coverage generosity in each state, and using the indifference condition of the marginal type in equation (4), the envelope theorem yields:

$$\frac{\partial W}{\partial(-p_h)} = \mu_h \, s_h^*(p_h) \, U_h'(I_h), \qquad h \in \{0,1\}. \tag{6}$$

Put together, the total impact on social welfare of a budget-neutral perturbation in

generosity becomes:

$$\Delta W = \mu_0 \, s_0^*(p_0) \left[U_0'(I_0) - U_1'(I_1) \left(\frac{1 + \rho_0 \varepsilon_0}{1 + \rho_1 \varepsilon_1} \right) \right]. \tag{7}$$

To gain intuition, denote the Lagrange multiplier on the Medicare budget constraint by λ, and note that at an interior optimum the first-order condition with respect to $-p_h$ implies:

$$\mu_h s_h^*(p_h) \left[U_h'(I_h) - \lambda \left(1 + \rho_h \varepsilon_h \right) \right] = 0.$$

The optimality condition for an internal solution implied by equation (7) (that is, $\Delta W = 0$) can then be represented as follows:

$$\left[U_0'(I_0) - U_1'(I_1) \right] = \lambda \left[\rho_0 \varepsilon_0 - \rho_1 \varepsilon_1 \right]. \tag{8}$$

This captures the insurance-efficiency tradeoff in the socially optimal design of household-level long-term care insurance coverage, akin to the Baily-Chetty formula (Baily, 1978; Chetty, 2006). The left-hand side captures the relative value of a dollar to a household without a healthy spousal caregiver versus a household with a healthy spousal caregiver. The right-hand side captures the relative fiscal externality from more coverage for households without a healthy spousal caregiver versus households with a healthy spousal caregiver.

The welfare characterizations in equations (7) and (8) lead to two implications for the optimal cost sharing of a 21$^{\text{st}}$ day in a SNF. First, equal coverage cannot be optimal so long as either marginal utility or the price elasticities of demand differ across states.[40] To see that, evaluate equation (7) at a symmetric policy $p_0 = p_1$. In models of choice under risk, it is typically assumed that the consumer is not over-insured, so $U_0'(I_0) \geq U_1'(I_1)$. If demand is less price-elastic when informal care is unavailable ($\varepsilon_0 < \varepsilon_1$), equation (7) implies that $dSW/d(-p_0) > 0$. This implies that 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 so long as those with incapacitated spouses exhibit lower price elasticities of demand.

Second, equation (8) leads to a rule for optimal subsidies. If health shocks are fully insured outside the copayment so that $U_0'(I_0) = U_1'(I_1)$, then the optimal-pricing

[40] Of course, there is an edge case where differences in marginal utility perfectly offset differences in price elasticities and lead to equal coverage being optimal, but this is vanishingly unlikely in practice.

rule simplifies to:

$$\frac{\rho_0^*}{\rho_1^*} = \frac{\varepsilon_1}{\varepsilon_0}. \tag{9}$$

Thus, the state with the *lower* price elasticity—the sick-spouse state—should receive the *higher* subsidy in proportion to the ratio of elasticities. When demand responds less to price, the efficiency cost of raising utilization through a subsidy is smaller, so the social planner shifts resources toward that group.

The empirical results in Section 3 document that $\varepsilon_1 > \varepsilon_0$, and so equation (9) implies that the optimal contract should provide more coverage in the sick-spouse state. If $U_0'(I_0) \neq U_1'(I_1)$, the optimal subsidies are more complex, as copayments must take into account differences in marginal utility across states. However, it seems likely that marginal utility would be weakly *higher* in the sick-spouse state, and, in that case, the model would again suggest increased coverage in the absence of a healthy spouse.

4.3 Model Calibration

Next, we calibrate the model. We first derive demand curves for formal care and compute the option value of spousal care, $\mathrm{OV}(p)$, which Section 4.1 showed to be a lower bound on the full valuation of informal spousal care. We then determine the optimal subsidy ratio, ρ_0^*/ρ_1^*. Both calculations concern the 21$^{\text{st}}$ day in a SNF.

Demand Curves and the Value of Informal Care. In order to empirically estimate consumer surplus, we impose structure on the valuation function, $v_h(\cdot)$. We model $v_h(s_h)$ as linear in the type s_h:

$$v_h(s_h) = \alpha_h - \beta_h\, s_h, \qquad \alpha_h, \beta_h > 0, \tag{10}$$

where $\alpha_h = v_h(s_h = 0)$ is the willingness to pay of the highest-valuation household in state h, and β_h governs how quickly valuation declines across types. Imposing the indifference condition $v_h(s_h^*) = p$ and inverting equation (10) yields a demand curve that is linear in price,

$$Q_h(p) = s_h^*(p) = \frac{\alpha_h - p}{\beta_h}.$$

Two quantity-price pairs then identify the entire curve: for any two observed prices $p < p'$, $\beta_h = (p' - p)/[Q_h(p) - Q_h(p')]$ and $\alpha_h = p + \beta_h\, Q_h(p)$.

We calibrate demand curves reflecting differences in both the *level* of demand for

formal care and its price sensitivity (the *slope*), as our results above indicate that losing informal care from a spouse affects both. The regression-discontinuity design presented in Section 3 can provide those two quantity-price pairs.[41] It provides the share of beneficiaries in the SNF just before and just after copayments are imposed. The price of a 21st day in a SNF is $p \approx \$150$ for all Medicare recipients, but roughly three-quarters hold (full) supplemental coverage for out-of-pocket Medicare costs. So the average beneficiary faces $p \approx \$37.50$ on the 21st day after facing $p \approx \$0$ until day 20.[42]

On day 20, the share of beneficiaries with a healthy spouse in the SNF is $Q_1(0) = 0.225$ and the share of beneficiaries with a sick spouse in the SNF is $Q_0(0) = 0.432$. The RD coefficients in Table 5 provide the change in the probability of remaining in the SNF among the beneficiaries who reach day 20: 1.8 percentage points for those with healthy spouses and 1.0 percentage points for those with sick spouses. Scaling Q_1 and Q_0 down by those proportional responses provides the points (0.221, \$37.50) for those with healthy spouses and (0.428, \$37.50) for those with incapacitated spouses.[43]

Those numbers, in turn, lead to $\beta_1 \approx 9{,}300$, $\alpha_1 \approx \$2{,}083$ for the healthy-spouse group and $\beta_0 \approx 8{,}700$, $\alpha_0 \approx \$3{,}750$ for the sick-spouse group. These α_h, β_h are the parameters of the inverse-demand function $v_h(s_h) = \alpha_h - \beta_h s_h$ in (10), which returns the price at which a given marginal consumer is indifferent between remaining in and leaving the SNF on the 21st day and mirrors the condition $v_h(s_h^*) = p$ derived above. Appendix Figure A23 plots these inverse-demand curves.

[41] This step requires assuming that the entire utilization response to the discrete change in copayments operates through demand rather than supply. Supply-side responses are possible. If SNFs act as perfect agents for their patients, our consumer-surplus calculations remain valid. If instead SNFs act in their own interests, the demand curves conflate patient and provider responses, and the area beneath them is no longer consumer surplus. That could occur if, for example, patients often fail to pay the copayments they are charged, so that the revenue a SNF collects for the 21st day falls even as the posted price is unchanged. We assume away such responses on the grounds that they are likely second-order relative to the demand response to price.

[42] Medigap enrollment is not observed in any administrative data source, preventing us from assigning individual-specific prices according to Medigap enrollment status.

[43] Since the RD outcome is an indicator for SNF discharge, the RD coefficients reported in Table 5, β_h^{DD}, reflect the copayment-induced discontinuity in the daily discharge hazard. Since this hazard is a probability conditional on still being in the SNF, we apply it multiplicatively rather than additively to recover our demand curves. That is, $Q_h(37.50) = Q_h(0) \times (1 - \beta_h^{DD}/c_h)$, where c_h is the pre-copayment continuation rate; because daily discharge hazards around the threshold are small (Appendix Figure A22), $c_h \geq 0.92$ for both groups. In practice, this results in estimated points on the demand curve that are indistinguishable from applying the RD coefficients directly: $Q_h(0) \times (1 - \beta_h^{DD}/c_h) \approx Q_h(0) \times (1 - \beta_h^{DD})$ within 0.0002. Note that if we instead applied the estimates in Table 5 additively without correcting for the conditional outcome, we would recover demand curves with a larger difference in slopes.

The consumer surplus generated by the 21st day is the area between the inverse-demand curve and the price paid.[44] Using the notation above, we can derive consumer surplus for those in state h, $CS_h(p)$, as:

$$CS_h(p) = \int_0^{s_h^*(p)} [v_h(s_h) - p]\, ds_h = \tfrac{1}{2}\,(\alpha_h - p)\, s_h^*(p) = \frac{(\alpha_h - p)^2}{2\beta_h}.$$

Evaluated at $p \approx \$37.50$, this implies a consumer surplus of approximately \$226 per patient per day for the healthy-spouse group and \$794 for the sick-spouse group. This implies that the option value of spousal care at $p = 37.5$ is $\text{OV}(\$37.50) = CS_0(\$37.50) - CS_1(\$37.50) = \$568$.[45] In comparison, Medicare pays a median of \$556 per patient per day in a SNF (Schotland et al., 2023). Thus the estimates imply that the full value of a day of spousal care after a 20-day SNF stay is at least as large as the median Medicare payment for an additional day in the SNF.

Optimal Copayments. The model also provides a lens through which we can trace the implications of the empirical results for policymakers. In particular, the RD estimates provide estimates of ε_1 and ε_0, the price elasticities of demand for formal care among those with and without healthy spouses. Equation (9) then maps those price elasticities of demand to the optimal ratio of subsidies, and inversely, the optimal ratio of copayments.

As with the demand curves, a potential concern in converting the RD coefficients to optimal copayments is that we do not observe precise prices, because of the existence of private Medigap policies. That concern is less pressing here, however, because the main object of interest is not the elasticities themselves but a ratio of elasticities.[46] We therefore derive an optimal subsidy ratio of 1.80 using the RD estimates presented in Table 5; that is, the optimal degree of subsidization is 80% higher for those with sick spouses than for those with healthy spouses.

The calibrated model thus implies that optimal copayments are higher for beneficiaries with healthy, available caregivers than for those with incapacitated spouses.

[44] To see this equivalence, first plug in $v_h(s_h) = \alpha_h - \beta_h\, s_h$ for $v_h(s_h)$, solve the integral, and then plug in the condition $v_h(s_h^*) = p$.

[45] Similarly, at a price of 0, this option value is $\text{OV}(\$0) = \$808 - \$233 = \$575$.

[46] To see this, consider the price elasticity for group $h \in \{0, 1\}$ measuring the slope of the demand curve on day 21, given by: $\hat{\varepsilon}_h = -\frac{p_h}{Q_h^+} \cdot \frac{d\hat{s}_h}{dp_h} = -\frac{\Delta Q_h}{Q_h^+}$, where Q_h^+ indicates the probability of staying one more day in the SNF just after cost sharing is incurred (the right-hand intercept of the RD regression), and ΔQ_h is the RD coefficient. The ratio of elasticities therefore does not depend on the effective price change and is given by $\frac{\varepsilon_1}{\varepsilon_0} = \frac{\Delta Q_1}{\Delta Q_0} \cdot \frac{Q_0^+}{Q_1^+}$. Note that the ratio $\frac{Q_0^+}{Q_1^+} \approx 1$ in our context.

If an extra day in a SNF were to cost Medicare $1,000 and Medicare were to charge beneficiaries with healthy spouses a copayment of $200 (20%), then Medicare should charge beneficiaries with sick spouses a copayment of roughly $122 (or about 12.2%). Similarly, if Medicare were to charge beneficiaries with healthy spouses a copayment of $500 (50%), then Medicare should charge beneficiaries with sick spouses a copayment of roughly $357 (or about 35.7%).

4.4 Policy Counterfactuals

The calculations above focus on the 21st day in a SNF. We next extrapolate beyond the 21st day and explore how copayments affect beneficiaries with sick versus healthy spouses across the entire SNF visit. To do so, we rely on a dynamic model developed by Eliason et al. (2018). While this necessarily involves an extrapolation, the model allows us to estimate demand curves for each day of formal care, valuing both direct utility and the additional value of extending a SNF stay beyond the next day. Appendix C describes the model and estimation.[47]

We use the dynamic model to estimate SNF utilization and patient welfare in response to different coinsurance rates.[48] Figure 8 describes these estimates under alternative cost sharing policies. It reports how overall consumer surplus responds to different combinations of coverage generosity for beneficiaries with sick versus healthy spouses. On the horizontal axis, it plots the coverage generosity for recipients with sick spouses. On the vertical axis, it plots the generosity for recipients with healthy spouses.

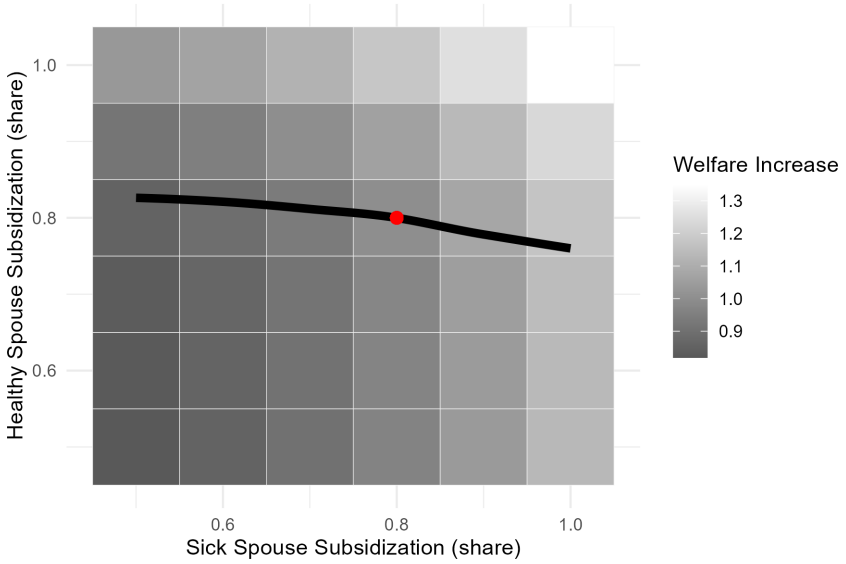
Currently, Medicare makes no distinction between those two groups. All recipients face copayments beginning on the 21st day of a SNF stay, which they either cover out-of-pocket or through Medigap. These copayments, although discrete, are approximately 20% of Medicare reimbursement. Accordingly, the figure includes a circular marker at $(0.80, 0.80)$ to approximate the current policy. The shading indicates the estimated change in total household welfare across both groups relative to this status quo, with lighter areas signifying larger welfare gains.

In addition, the figure plots a budget-neutral isoquant. The black line plots de-

[47] The model relies on a logit demand structure to estimate price elasticities across days in a SNF beyond day 21, informed by but independent of the RD estimates. The results are robust to other functional-form assumptions, including linear or constant-elasticity-of-substitution demand based only on the elasticities estimated in Section 3.4 (Kang and Vasserman, 2025).

[48] Here, we study coinsurance, rather than the fixed copayment schedule currently used by Medicare, as a continuous choice variable more in keeping with the social planner's problem above.

Figure 8. Welfare Gains from Subsidizing SNF Coverage by Spousal Incapacitation



Notes: This figure plots estimated changes in total household welfare using the dynamic model presented in Appendix C. Each point in the plot represents a different combination of government coverage generosity for SNF stays (from day 21 on), ranging from 0% to 100% coverage. For each point, the total household welfare is shown, relative to the status quo of $(80\%, 80\%)$ coverage (indicated by the circular marker). The black line indicates the set of policies which are budget-neutral relative to the status quo.

viations from the current policy that would involve different cost sharing policies for the two groups but would hold Medicare's budget fixed. The isoquant has a negative slope: reducing cost sharing for the sick-spouse group—moving rightward along the horizontal axis—requires increased cost sharing for the healthy-spouse group—moving downward along the vertical axis. The isoquant is also relatively flat. Large transfers to the sick-spouse group require only small reductions in coverage for the healthy-spouse group to hold the budget constant. For example, suppose that Medicare were to offer full insurance coverage to recipients with sick spouses. The figure suggests that such a policy change could be budget neutral solely by reducing coverage for those with healthy spouses from 80% to 76%.

The isoquant is relatively flat for two reasons. First, we observe limited demand response to the lower coinsurance rate for the sick-spouse group and a large demand response for the healthy-spouse group. Second, the healthy-spouse group is larger than the sick-spouse group.[49]

[49]In part, this may be due to our focus on households experiencing spousal heart attacks and strokes. This group is larger if we include other conditions; however, this group will likely remain small relative to the healthy-spouse group, and their demand response to lower coinsurance rates

Moving to the right along the isoquant in Figure 8 suggests a large improvement in welfare, 13.4% over the status quo, without increasing government expenditures. That increase in welfare requires a shift in the generosity of insurance coverage to those with incapacitated spouses.[50]

4.5 Practical Implementation: Family Deductibles

The analysis above assumes that the social planner can condition copayments for formal care on the availability of a healthy spouse, setting $p_0 \neq p_1$. In practice, setting those copayments may be difficult. Doing so would require the public sector to evaluate the ability of a spouse to provide informal care, which could entail many of the same difficulties as eligibility determinations in disability insurance.

More tractable alternatives, however, exist. Consider insurance contracts that involve a family-specific deductible or out-of-pocket maximum. Under such a contract, beneficiaries face copayments that are linked mechanically to the state of their spouse. A major cardiac event for one spouse could lead to spending that exceeds the deductible or out-of-pocket maximum and so directly lowers out-of-pocket prices for the other spouse. In that way, the same event that moves a shock spouse into a more vulnerable state also lowers cost sharing for the outcome spouse. That corresponds to the state-specific contract that the model above suggests would improve welfare relative to the status quo.

For that matter, the transfers provided by a family deductible may also be better targeted than those transfers provided by individual-specific deductibles. Consider a household that suffers two health shocks. Such a household would receive a larger transfer under a family deductible than under individual deductibles. That household likely has greater needs than other households, and so the family-specific deductible could improve social welfare in this way as well.

A family deductible is, of course, a blunt instrument. It would reset at the end of every year. In addition, a single threshold cannot fine-tune the ratio of subsidies across states of nature. However, even this imperfect instrument more closely approximates the optimal contract than the individual-specific status quo under the current

will continue to be small.

[50] Appendix Figure C2 offers a similar finding for the day cost sharing begins rather than the level of coverage generosity. It suggests that welfare could be improved by 4.7% over the status quo without increasing government expenditures by delaying cost sharing for those with incapacitated spouses. Delaying cost sharing for this group by a month, from 20 days to 50 days, would only require moving the onset of cost sharing earlier for those with healthy spouses by eight days, from day 20 to day 12.

Medicare program.

The broader lesson of the theoretical framework extends beyond any one instrument. The risks that beneficiaries face and the resources they have to absorb those risks are pooled within the household. When one spouse is incapacitated, the other's demand for formal care rises and becomes less price-sensitive. The findings therefore underscore that the natural unit of social insurance is not the individual but rather the household.

5 Conclusions

We study how the loss of informal spousal care changes demand for formal care among couples enrolled in Medicare. When one spouse experiences a major cardiac event, their partner becomes more likely to use formal care in a nursing home or Medicare-covered SNF. They also become less sensitive to the price of that formal care. These changes in both the level and slope of demand for formal care appear to be largely driven by substitution from informal to formal care once informal spousal care is no longer an option.

These results suggest that households place a high value on informal care from a spouse. Through the lens of optimal insurance, our results also have implications for the way in which public and private payers handle formal care. They suggest lower copayments for beneficiaries with an incapacitated spouse. Our simulations show meaningful welfare gains from shifting coverage toward this group. Family deductibles or out-of-pocket maximums offer one practical way to link coverage across spouses. Equal copayments for formal care under-subsidize those who have no access to informal care and over-subsidize those who do.

References

Abadie, Alberto and Maximilian Kasy, "Choosing Among Regularized Estimators in Empirical Economics: The Risk of Machine Learning," *Review of Economics and Statistics*, 2019, *101* (5), 743–762.

Abraham, Katharine G. and Christopher Mackie, *Beyond the Market: Designing Nonmarket Accounts for the United States*, Washington, DC: National Academies Press, 2005.

Aguiar, Mark and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *The Quarterly Journal of Economics*, 2007, *122* (3), 969–1006.

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti, "Long-Term-Care Utility and Late-in-Life Saving," *Journal of Political Economy*, 2020, *128* (6), 2375–2451.

Anderson, David M, Alex Hoagland, and Ed Zhu, "Medical Bill Shock and Imperfect Moral Hazard," *Journal of Public Economics*, 2024, *236*, 105152.

Arteaga, Carolina, Natalia Vigezzi, and Pilar Garcia-Gomez, "Health Spillovers: The Broad Impact of Spousal Health Shocks," 2025. NBER Working Paper #33994.

Baily, Martin Neil, "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 1978, *10* (3), 379–402.

Becker, Gary S., "A Theory of the Allocation of Time," *The Economic Journal*, 1965, *75* (299), 493–517.

Becker, Gary S, "A Theory of Marriage: Part II," *Journal of Political Economy*, 1974, *82* (2, Part 2), S11–S26.

Bergeot, Julien and Marianne Tenand, "Does Informal Care Delay Nursing Home Entry? Evidence from Dutch Linked Survey and Administrative Data," *Journal of Health Economics*, 2023, *92*, 102831.

Bonsang, Eric, "Does Informal Care from Children to Their Elderly Parents Substitute for Formal Care in Europe?," *Journal of Health Economics*, 2009, *28* (1), 143–154.

Brown, Jeffrey R and Amy Finkelstein, "Insuring Long-Term Care in the United States," *Journal of Economic Perspectives*, 2011, *25* (4), 119–142.

_ **and James M Poterba**, "Joint Life Annuities and Annuity Demand by Married Couples," *Journal of Risk and Insurance*, 2000, *67* (4), 527–553.

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*, 2021, *190*, 164–190.

Centers for Disease Control and Prevention, "About Abuse of Older Persons," https://www.cdc.gov/elder-abuse/about/index.html November 2024. Updated November 7, 2024. Accessed May 20, 2026.

Centers for Medicare & Medicaid Services, "CMS Program Statistics – Medicare Skilled Nursing Facility," 2026. Dataset last updated January 30, 2026. Catalog last checked February 1, 2026.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Difference-in-differences estimators of intertemporal treatment effects," *Review of Economics and Statistics*, 2026, *108* (4), 863–880.

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*, 2015, *50* (3), 871–882.

Charles, Kerwin Kofi and Purvi Sevak, "Can Family Caregiving Substitute for Nursing Home Care?," *Journal of Health Economics*, 2005, *24* (6), 1174–1190.

Chay, Kenneth Y. and Michael Greenstone, "Does Air Quality Matter? Evidence from the Housing Market," *Journal of Political Economy*, 2005, *113* (2), 376–424.

Chetty, Raj, "A General Formula for the Optimal Level of Social Insurance," *Journal of Public Economics*, 2006, *90* (10-11), 1879–1901.

Coile, Courtney, "Health Shocks and Couples' Labor Supply Decisions," 2004.

Coyne, David, Itzik Fadlon, Shanthi P Ramnath, and Patricia K Tong, "Household Labor Supply and the Value of Social Security Survivors Benefits," *American Economic Review*, 2024, *114* (5), 1248–1280.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo, "The Economic Consequences of Hospital Admissions," *American Economic Review*, 2018, *108* (2), 308–352.

Einav, Liran and Amy Finkelstein, "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, February 2011, *25* (1), 115–138.

Eliason, Paul J, Paul LE Grieco, Ryan C McDevitt, and James W Roberts, "Strategic Patient Discharge: The Case of Long-Term Care Hospitals," *American Economic Review*, 2018, *108* (11), 3232–3265.

Ellis, Randall P, Heather E Hsu, Jeffrey J Siracuse, Allan J Walkey, Karen E Lasser, Brian C Jacobson, Corinne Andriola, Alex Hoagland, Ying Liu, Chenlu Song 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, *3* (3).

Elwert, Felix and Nicholas A Christakis, "The Effect of Widowhood on Mortality by the Causes of Death of Both Spouses," *American Journal of Public Health*, 2008, *98* (11), 2092–2098.

Fadlon, Itzik and Torben Heien Nielsen, "Family Health Behaviors," *American Economic Review*, 2019, *109* (9), 3162–3191.

_ and _ , "Household Labor Supply and the Gains from Social Insurance," *Journal of Public Economics*, 2019, *171*, 18–28.

_ and _ , "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records," *American Economic Journal: Applied Economics*, July 2021, *13* (3), 1–30.

_ , Astrid Sophie Fugleholm, and Torben Heien Nielsen, "Survivors' mental health and the protective role of income stability," *American Economic Review: Insights*, 2026, *8* (2), 265–284.

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.

Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo, "Approaches to Estimating the Health State Dependence of the Utility Function," *American Economic Review*, 2009, *99* (2), 116–121.

_ , _ , and _ , "What Good Is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption," *Journal of the European Economic Association*, 2013, *11* (1), 221–258.

_ , Nathaniel Hendren, and Mark Shepard, "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," *American Economic Review*, 2019, *109* (4), 1530–67.

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*, 2024, *134* (660), 1701–1718.

Gronau, Reuben, "Leisure, Home Production, and Work—The Theory of the Allocation of Time Revisited," *Journal of Political Economy*, 1977, *85* (6), 1099–1123.

Gross, Tal, Tim Layton, Daniel Prinz, and Julia Yates, "Social Defaults and Plan Choice: The Case of Spousal Following," 2025. NBER Working Paper #34137.

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

Hoagland, Alex, "An ounce of prevention or a pound of cure? The value of health risk information," *Review of Economics and Statistics*, 2025, pp. 1–45.

Hodor, Michal, "Family Health Spillovers: Evidence from the RAND Health Insurance Experiment," *Journal of Health Economics*, September 2021, *79*, 102505.

Hoover, Robert M and Michol Polson, "Detecting Elder Abuse and Neglect: Assessment and Intervention," *American Family Physician*, 2014, *89* (6), 453–460.

Hunt, Erin Cottle and Frank N Caliendo, "Social Security and Longevity Risk: An Analysis of Couples," *Journal of Public Economic Theory*, 2022, *24* (3), 547–579.

Hur, Elizabeth and Ajin Lee, "The Economic Consequences of Nursing Home Admissions," 2026. Working paper.

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*, 2017, *14* (10), 1224.

Kang, Zi Yang and Shoshana Vasserman, "Robustness Measures for Welfare Analysis," *American Economic Review*, 2025, *115* (8), 2449–2487.

Landais, Camille and Johannes Spinnewijn, "The Value of Unemployment Insurance," *The Review of Economic Studies*, 2021, *88* (6), 3041–3085.

Leive, Adam and Ya Sun, "Fit to Consume: How Health Shapes Preferences for Consumption," 2025. TIAA Institute Research Paper.

Lockwood, Lee M., "Incidental Bequests and the Choice to Self-Insure Late-Life Risks," *American Economic Review*, 2018, *108* (9), 2513–2550.

Maestas, Nicole, Matt Messel, and Yulya Truskinovsky, "Caregiving and Labor Supply: New Evidence from Administrative Data," *Journal of Labor Economics*, 2024, *42* (S1), S183–S218.

March, Lucien, "Some Researches Concerning the Factors of Mortality," *Journal of the Royal Statistical Society*, 1912, *75* (5), 505–538.

Massner, Patrizia and Jens Wikström, "Should Elder Care Be Subsidized? Theory and Evidence from Sweden," November 23 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*, 2021, *40* (5), 745–753.

Nadash, Pamela, Pamela Doty, and Matthias von Schwanenflügel, "The German Long-Term Care Insurance Program: Evolution and Recent Developments," *The Gerontologist*, 2018, *58* (3), 588–597.

Nordhaus, William D. and James Tobin, "Is Growth Obsolete?," in "Economic Research: Retrospect and Prospect, Volume 5, Economic Growth," Cambridge, MA: National Bureau of Economic Research, 1972, pp. 1–80.

OECD, *Is Care Affordable for Older People?*, Paris: OECD Publishing, 2024.

Oster, Emily, Ira Shoulson, and E Ray Dorsey, "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease," *American Economic Review*, 2013, *103* (2), 804–830.

Pollak, Robert A. and Michael L. Wachter, "The Relevance of the Household Production Function and Its Implications for the Allocation of Time," *Journal of Political Economy*, 1975, *83* (2), 255–277.

Reid, Margaret G., *Economics of Household Production*, New York: John Wiley & Sons, 1934.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Schotland, Samuel, Rachel M Werner et al., "Medicare Payment Policy for Post-Acute Care in Nursing Homes," 2023. Penn LDI Research Brief.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

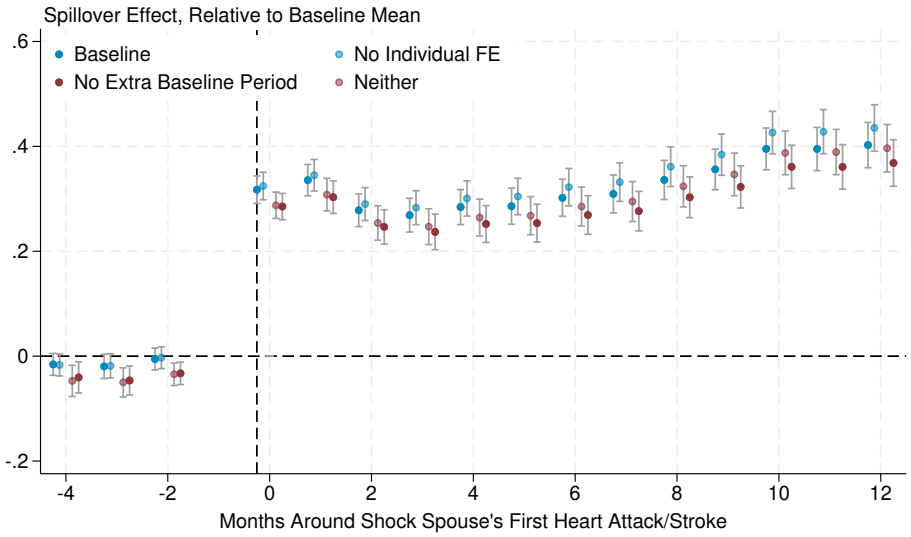
U.S. Bureau of Labor Statistics, "Median Weekly Earnings of Full-Time Wage and Salary Workers by Selected Characteristics," 2018. Current Population Survey, 2017 annual averages, Table 37.

Viscusi, W. Kip and William N. Evans, "Utility Functions That Depend on Health Status: Estimates and Economic Implications," *American Economic Review*, 1990, *80* (3), 353–374.

Zeckhauser, Richard, "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives," *Journal of Economic Theory*, 1970, *2* (1), 10–26.

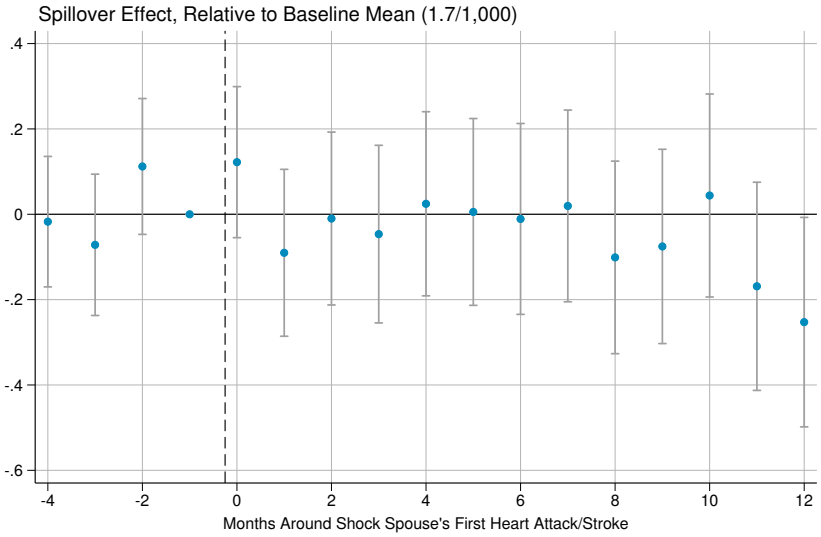
A Appendix: Additional Figures and Tables

Appendix Figure A1. Robustness Exercise: Alternative Specifications of Primary Effects



Notes: Compare to Figure 1. 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 was visiting a nursing home in a given period. Alternative models are shown as discussed in the text, including (a) omitting only one reference period instead of two, (b) excluding person-specific fixed effects from the regression, and (c) both. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stay, indicated at the top of the figure. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. All specifications include calendar month fixed effects.

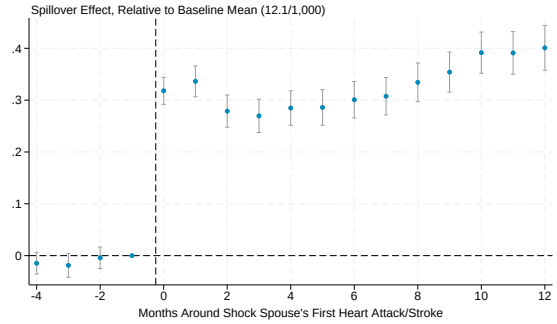
Appendix Figure A2. Effect of Shock Spouse's Index Event on Outcome Spouse's Heart Attacks and Strokes



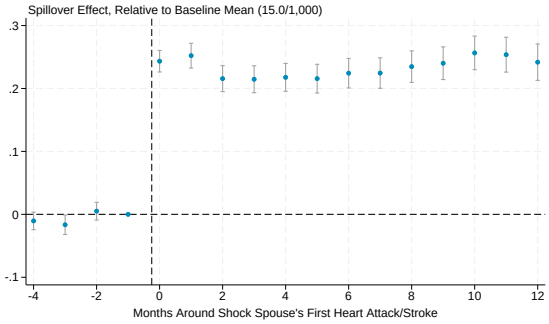
Notes: This figure plots dynamic difference-in-differences coefficients that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator for whether the outcome spouse also had a heart attack or stroke. Error bars plot 95% confidence intervals based on standard errors clustered at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects.

Appendix Figure A3. Robustness of Effects to Alternative Definitions of Index Events

(a) Allowing Multiple Health Shocks per Household



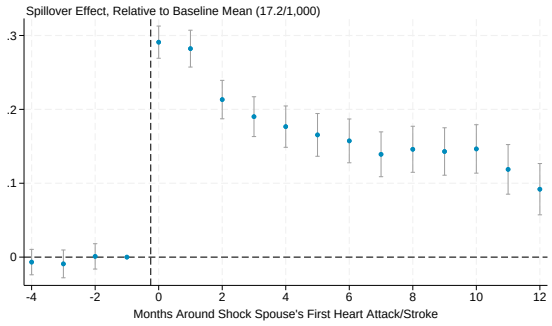
(b) No 2-Year Requirement



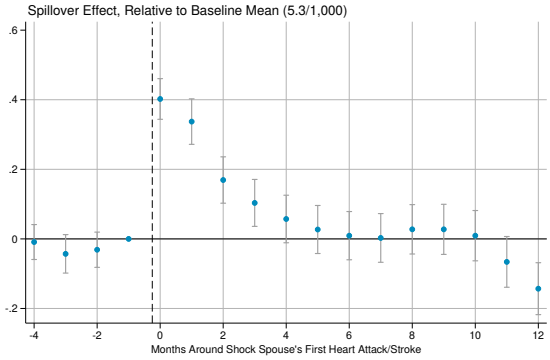
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator for whether the outcome spouse visited a nursing home. In panel (a), we include households where both spouses experienced a major cardiovascular event within our sample period (roughly 1.8% of all households). In panel (b), we relax the requirement that the shock spouse have no hospitalizations during the two years before the index event. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stays. 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. Compare to Figure 1.

Appendix Figure A4. Effect of a Spouse's Major Cardiovascular Event on Risk of a Nursing Home Visit: Unbalanced Panel

(a) Any Nursing Home Visit

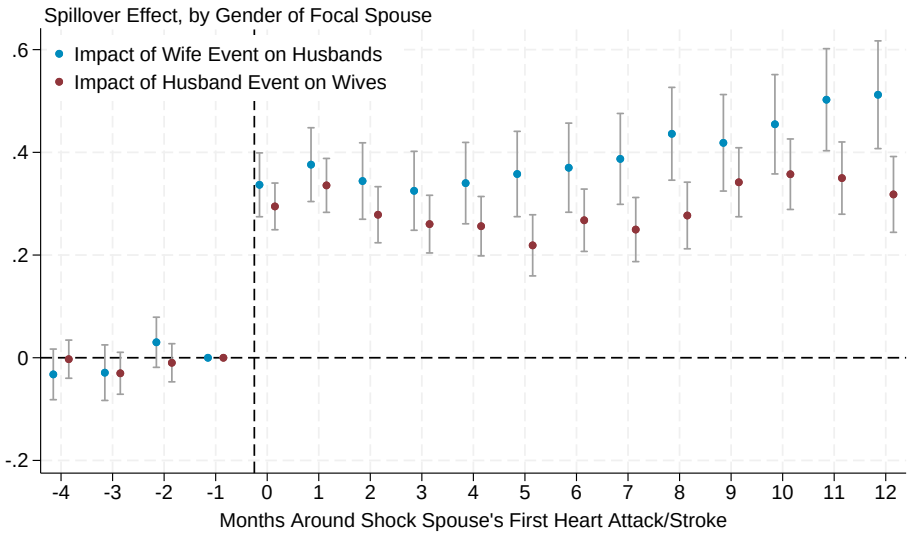


(b) Medicare-Covered SNF Stay



Notes: Compare to Figures 1 and 2. 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. Here, we do not require that outcome spouses survive for a full year following the health shock. We illustrate regression coefficients for an indicator for whether the outcome spouse was visiting (a) any nursing home (using the MDS) or (b) a Medicare-covered SNF (using the MedPAR data). We rescale all coefficients such that they indicate the change relative to the baseline risk of SNF stay or entry, indicated at the top of each panel. 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.

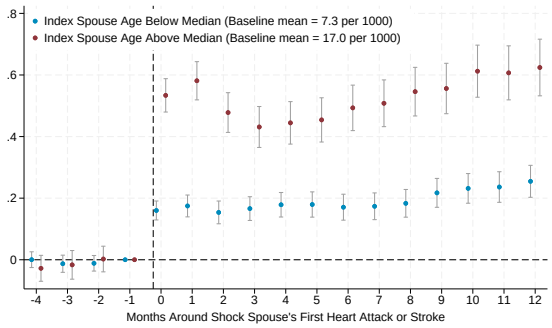
Appendix Figure A5. Spillover Effects by Gender: Visits to Nursing Homes



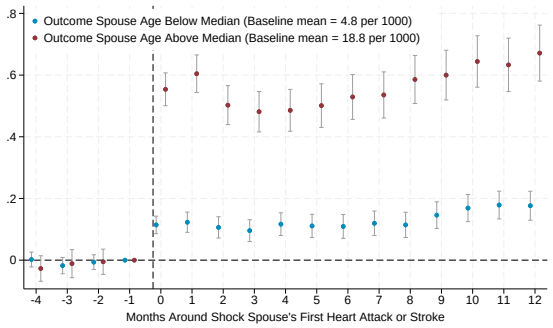
Notes: This figure plots dynamic difference-in-differences coefficients that track the months since the index event of a heart attack or stroke in the household, stratified by whether the husband or the wife experienced the index event. The outcome of interest is an indicator for whether the other partner visited a nursing home, and we rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stay. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects.

Appendix Figure A6. Nursing Home Effects, by Spouses' Ages

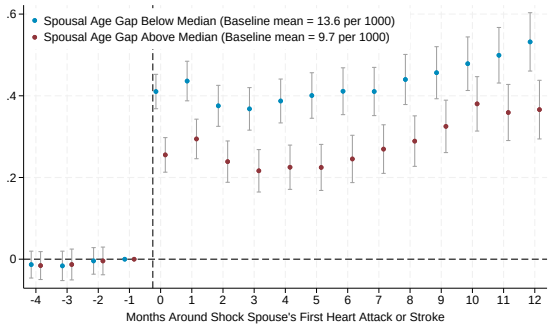
(a) Shock Spouse Age



(b) Outcome Spouse Age

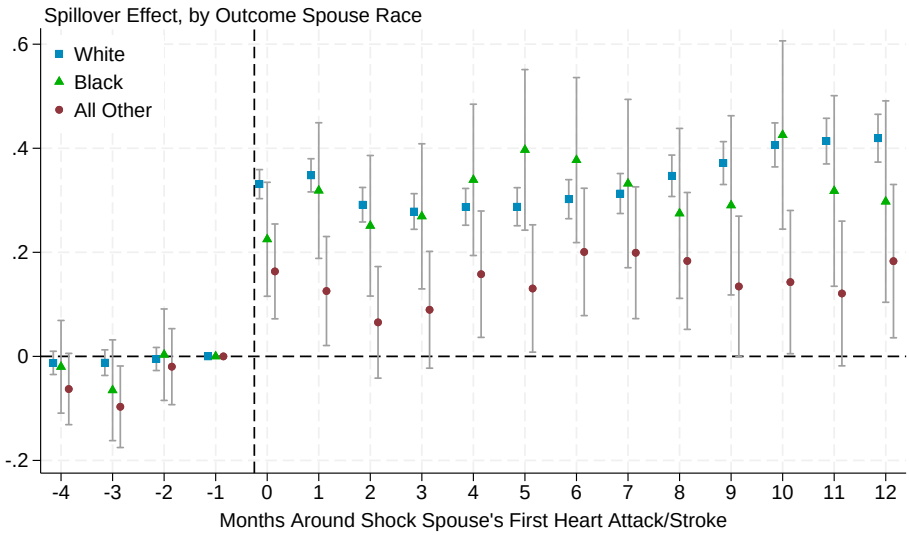


(c) Age Gap, Outcome and Shock Spouses



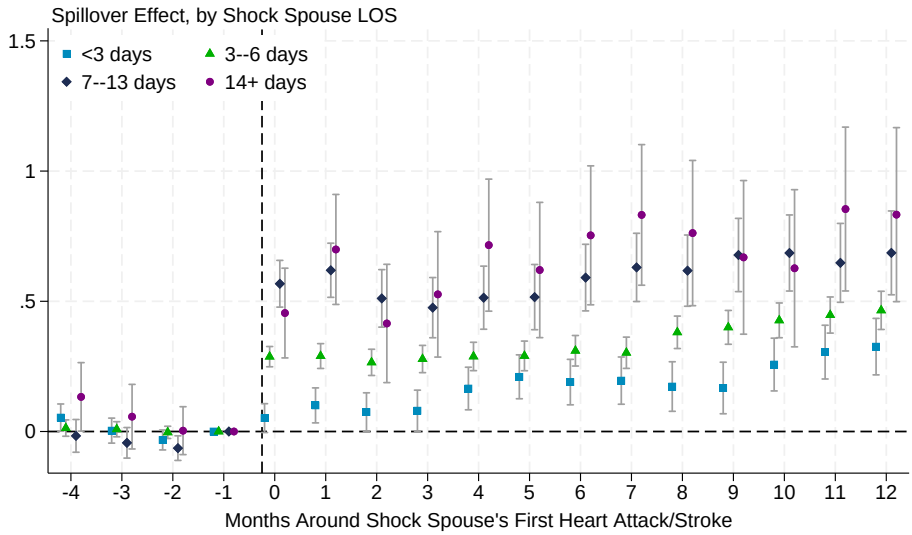
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator for whether the outcome spouse visited a nursing home. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home 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.

Appendix Figure A7. Nursing Home Effects, by Outcome Spouse's Race



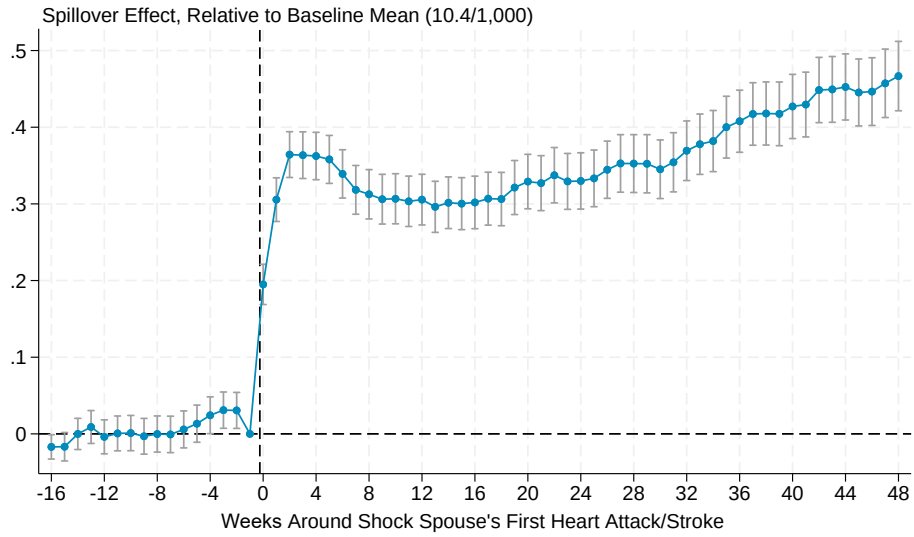
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator for whether the outcome spouse visited a nursing home. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home 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.

Appendix Figure A8. Nursing Home Effects, by Length of Stay for Initial Shock



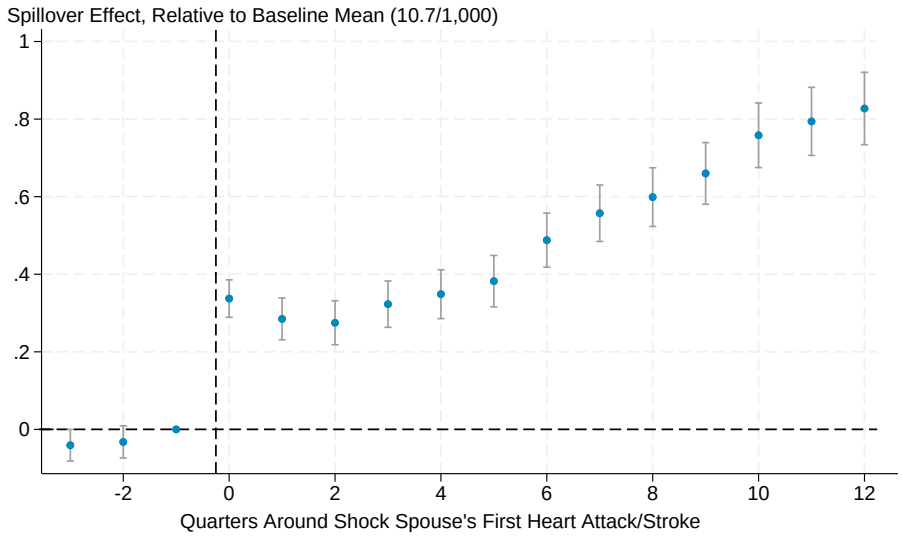
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. Groups are constructed based on the length of stay associated with the initial cardiovascular event affecting the shock spouse. The outcome of interest is an indicator for whether the outcome spouse visited any nursing home. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stays. 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.

Appendix Figure A9. Effect of a Spouse's Major Cardiovascular Event on Risk of a Nursing Home Visit: Weekly Results



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

Appendix Figure A10. Effect of a Spouse's Major Cardiovascular Event on Risk of a Nursing Home Visit: Longer-Run Results



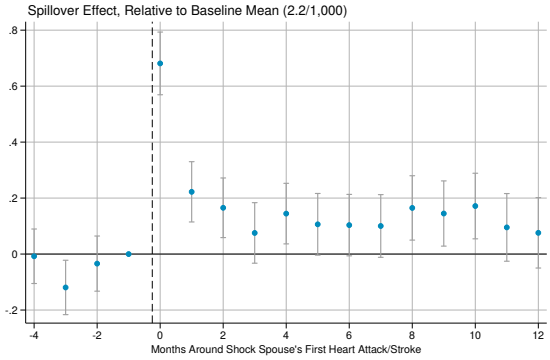
Notes: Compare to Figure 1. This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the quarters since a shock spouse's first heart attack or stroke; here, the control group is matched to facilitate estimating 3 years of post-treatment effects. The outcome of interest is an indicator for whether the outcome spouse was visiting a nursing home in a given period. We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stay, indicated at the top of the figure. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. The estimation includes calendar-time fixed effects and person-specific fixed effects.

Appendix Figure A11. Effect of a Spouse's Major Cardiovascular Event on Entry into Formal Care

(a) Any Nursing Home

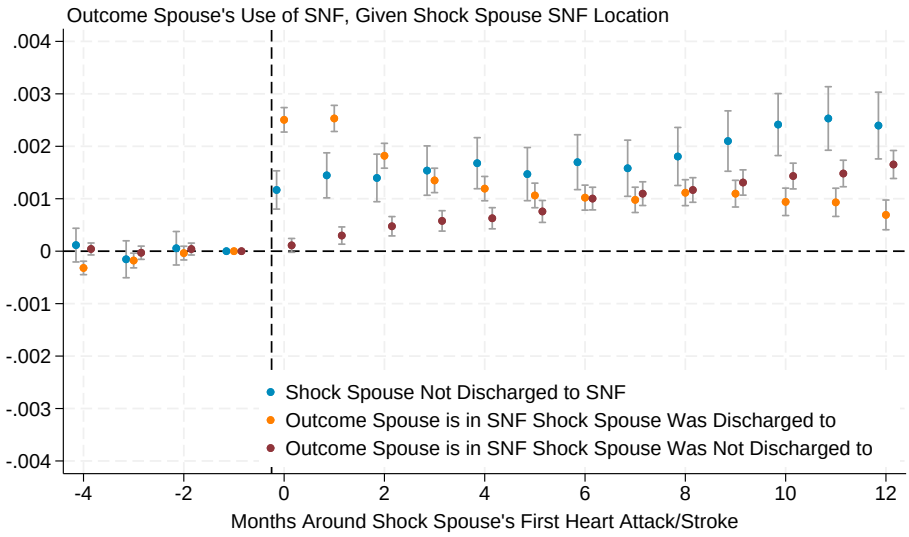


(b) Medicare-Covered SNF Stay



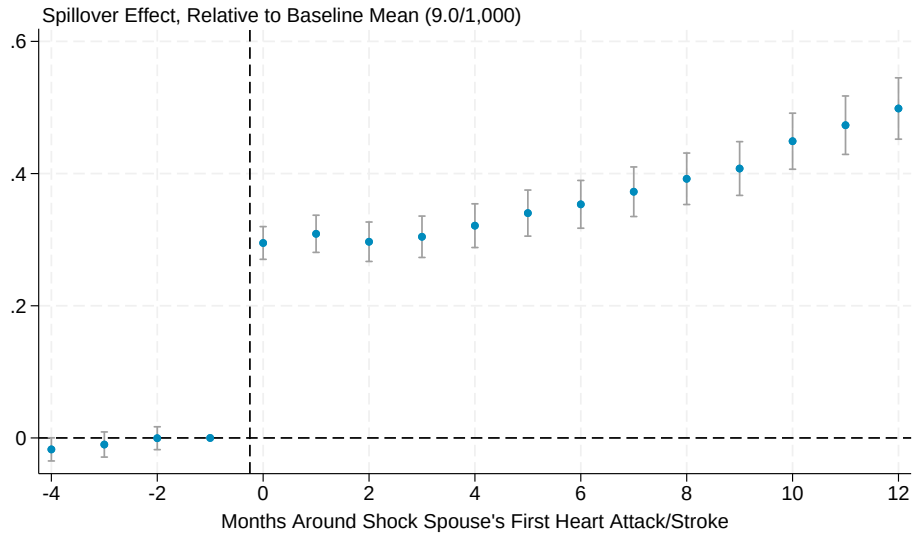
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest in both panels is an indicator for whether the outcome spouse began a new nursing-home stay (panel (a)) or a new Medicare-covered SNF stay (panel (b)) during the given period. Compare to Figures 1 and 2. We rescale all coefficients such that they indicate the change relative to the baseline risk of SNF stay or entry, indicated at the top of each panel. 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.

Appendix Figure A12. Nursing Home Effects Decomposed by Location of Shock Spouse



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 outcomes of interest are indicators for whether an outcome spouse is in a nursing home in a given month, stratified by whether the nursing home is the same as the one their partner (the shock spouse) is in or not. Nursing homes are identified using their organization NPI. Here, effects are presented in levels rather than being rescaled to assess the overall decomposition. 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.

Appendix Figure A13. Effect of a Spouse's Major Cardiovascular Event on Risk of Non-Medicare-covered Nursing Home Visits



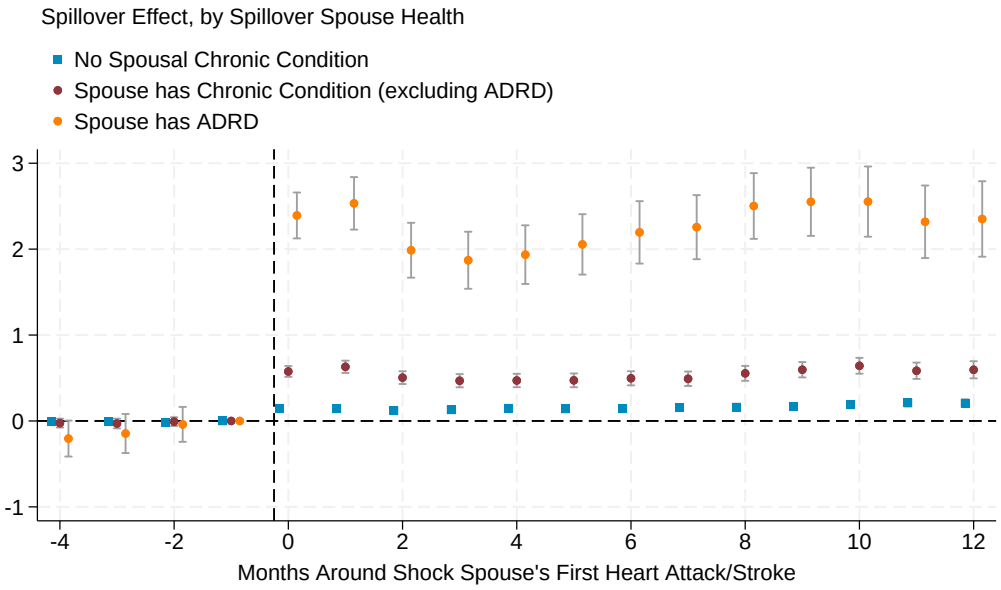
Notes: Compare to Figures 2 and 1. This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator for whether the outcome spouse was visiting any nursing-home facility identified in the MDS, excluding those covered by Medicare (in the MedPAR files). We rescale all coefficients such that they indicate the change relative to the baseline risk of nursing-home stays, indicated at the top of the figure. 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.

Appendix Figure A14. Nursing Home Effects for Fatal Shocks, by Time of Spousal Death



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, the sample is limited only to shock spouses who die within one year of their index event, and stratified by their time of death relative to the time of their initial event. The outcome of interest is a binary indicator for whether the outcome spouse was in a nursing home during a given month. We rescale all coefficients such that they indicate the change relative to the baseline mean. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects.

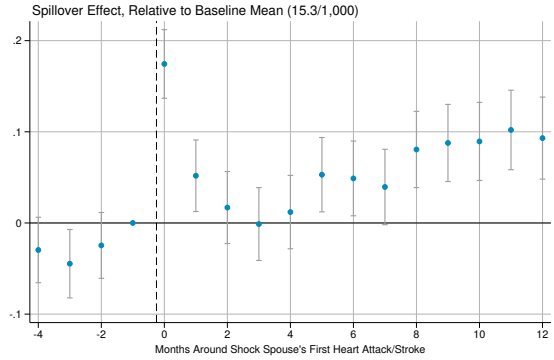
Appendix Figure A15. Effect on Risk of Any Nursing Home Visit by Outcome Spouse's Pre-existing ADRD Status and Other Chronic Conditions



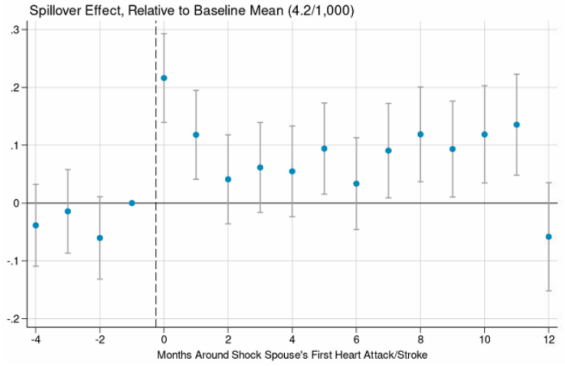
Notes: Compare to Figure 5. This figure plots estimates of the δ_n coefficients from equation (1), tracking the months since a shock spouse's first heart attack or stroke. The outcome is an indicator for whether the outcome spouse was visiting a nursing home in a given period. We stratify the sample based on whether the outcome spouse had diagnosed ADRD prior to their spouse's shock, another listed chronic condition but no ADRD, or neither. Chronic conditions include cancer (colorectal, lung, or prostate), chronic kidney disease, COPD, hip fractures, and strokes. We rescale coefficients to indicate changes relative to the baseline risk of a nursing-home stay in each category. Error bars plot 95% confidence intervals based on standard errors clustered at the household level.

Appendix Figure A16. Dynamic Treatment Effects for Health Outcomes

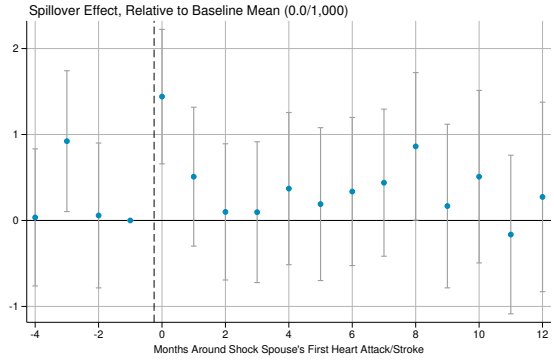
(a) Hospitalizations (Any)



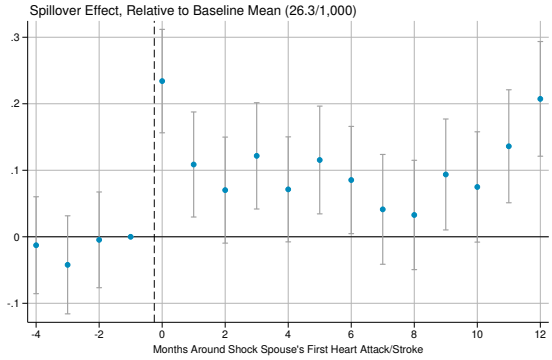
(b) Monitoring-Sensitive Hospitalizations



(c) Fall-Related Encounters

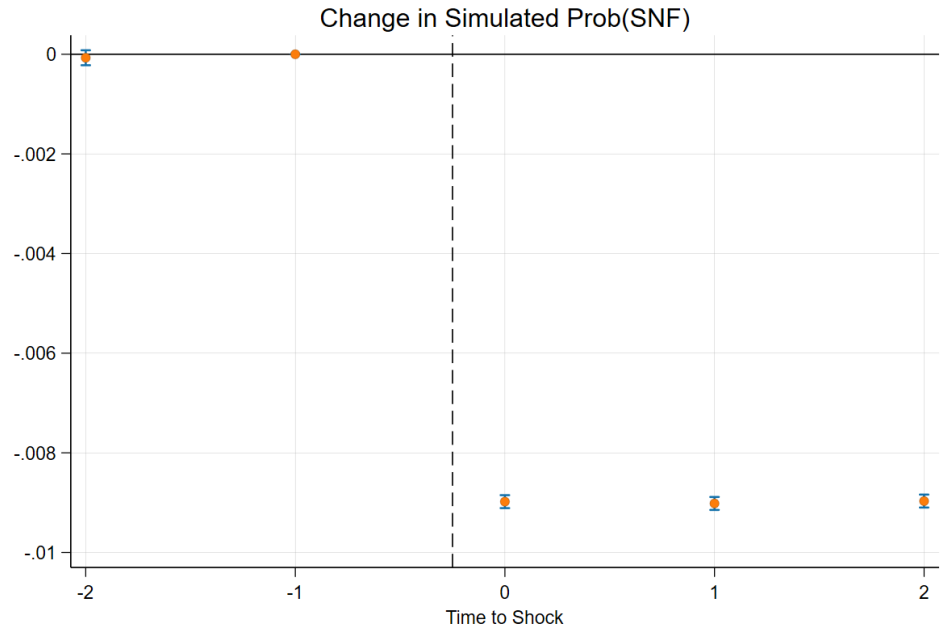


(d) ED Visits



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 outcomes of interest track those reported in Table 2. We rescale all coefficients such that they indicate the change relative to the baseline risk or rate. 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.

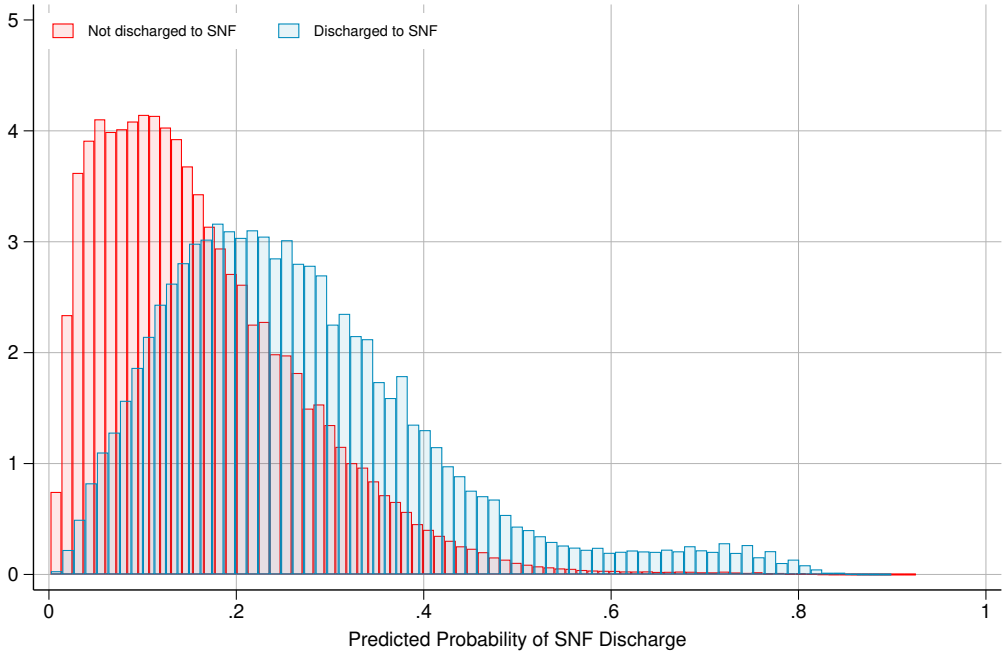
Appendix Figure A17. Simulation of Predicted SNF Probability Under Pure Health Effects



Notes: This figure presents results from a simulation designed to illustrate the behavior of the predicted-probability decomposition under a scenario where health effects contribute negatively to SNF discharge rates. We simulate outcomes for 1,000,000 outcome spouses who experience two types of hospitalizations in each period: severe hospitalizations, which lead to SNF discharge with probability 15%, and minor hospitalizations, which lead to SNF discharge with probability 1%. In the simulation, the treatment effect (following a spouse's hypothetical shock) increases only the probability of minor hospitalizations while holding the rate of severe hospitalizations fixed. The figure then plots the average rate at which an outcome spouse would be expected to be hospitalized and then discharged to a SNF in a given period, using the lasso-based prediction model.

The key insight illustrated by the figure is that, because any health effects induce only minor marginal hospitalizations, the overall predicted effect of a spousal shock on SNF discharge rates is negative as spouses switch from high-discharge to low-discharge hospitalizations. When spousal health shocks induce more frequent but milder hospitalizations relative to the pre-event baseline, the lasso decomposition can assign a negative predicted contribution to the health effect channel. Given an overall positive treatment effect on SNF discharges, this implies that the substitution channel accounts for more than 100% of the observed increase in SNF stays.

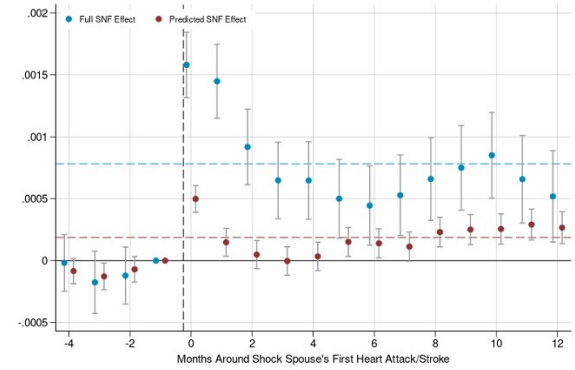
Appendix Figure A18. Predicted Probability of Discharge to a Medicare-Covered SNF: Validation



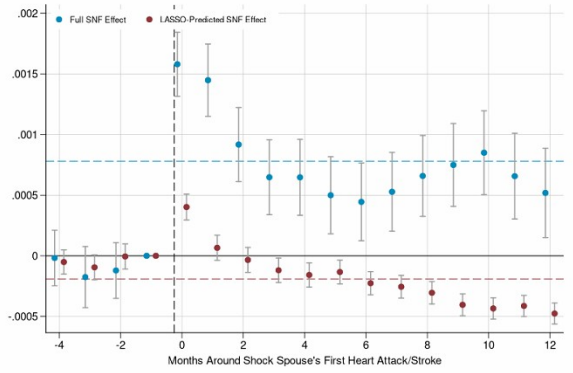
Notes: This figure plots the distribution of the predicted probability of a given hospitalization ending in a discharge to a SNF for 100% of inpatient hospitalizations between 2011 and 2017, for two groups: those that did not end with a discharge to a SNF (in red), and those that did (in blue). Predicted values are modeled using lasso logistic prediction with three-fold cross-validation with a binary outcome. Covariates include patient age, sex, race, dual Medicaid eligibility, Medicare enrollment type, and DRG indicators.

Appendix Figure A19. Change in Outcome Spouse's Predicted Probability of Medicare-Covered SNF Discharge

(a) Prediction Using Hospital Volume Only

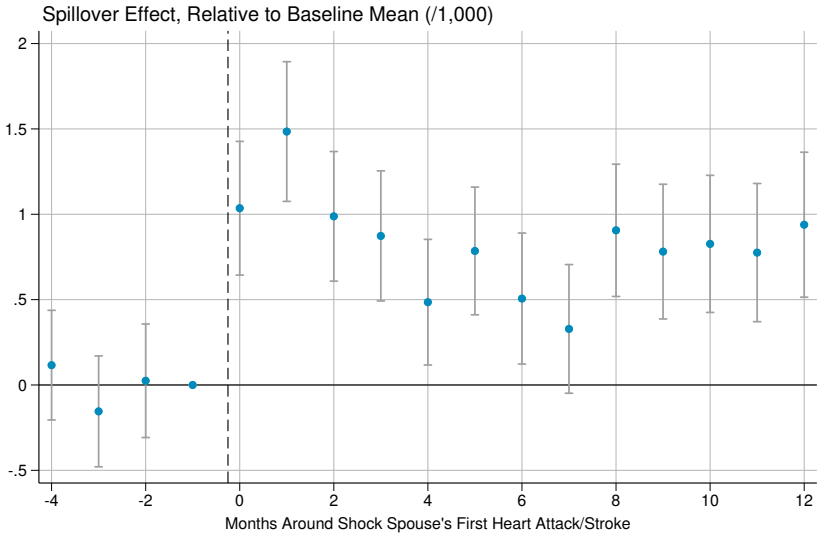


(b) Lasso Prediction



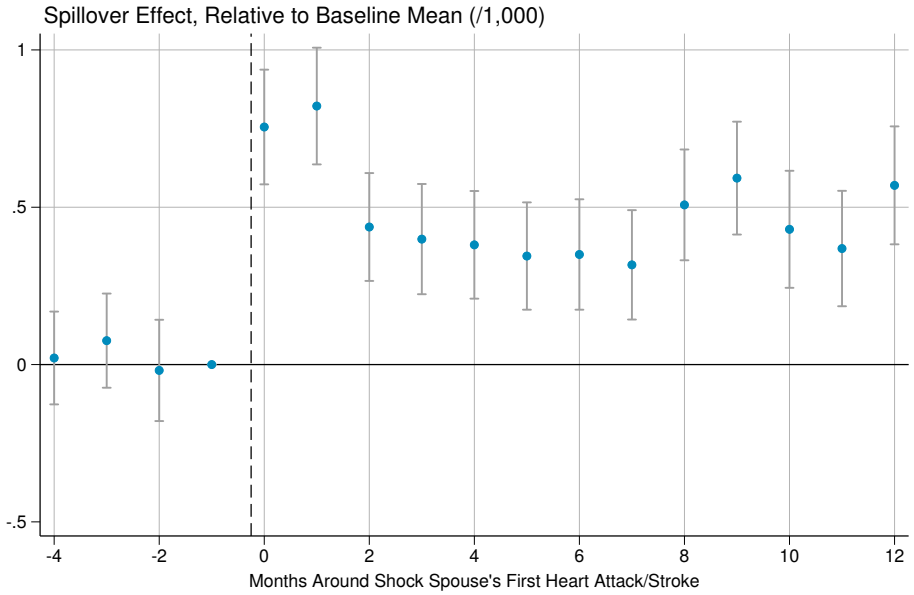
Notes: This figure plots event-study estimates for observed SNF use and the predicted components reported in Table 3. In each plot, the outcomes of interest are (1) an indicator for whether the outcome spouse was visiting a SNF in a given period and (2) the predicted probability of a SNF discharge. In the first panel, this prediction is based solely on hospital volumes, assuming each hospitalization has a fixed probability of SNF discharge. In the second panel, the predicted SNF discharge probability is based on hospitalization and patient characteristics using lasso regression. Coefficients are presented in levels here rather than rescaled (compare to Figure 2). The error bars plot 95% confidence intervals based on standard errors clustered at the household level. The regression includes calendar month fixed effects and person-specific fixed effects. Horizontal dashed lines indicate average post-treatment effects for each outcome.

Appendix Figure A20. Effect of Spousal Health Shock on the Probability of a SNF Stay Extending Beyond Full Medicare Coverage



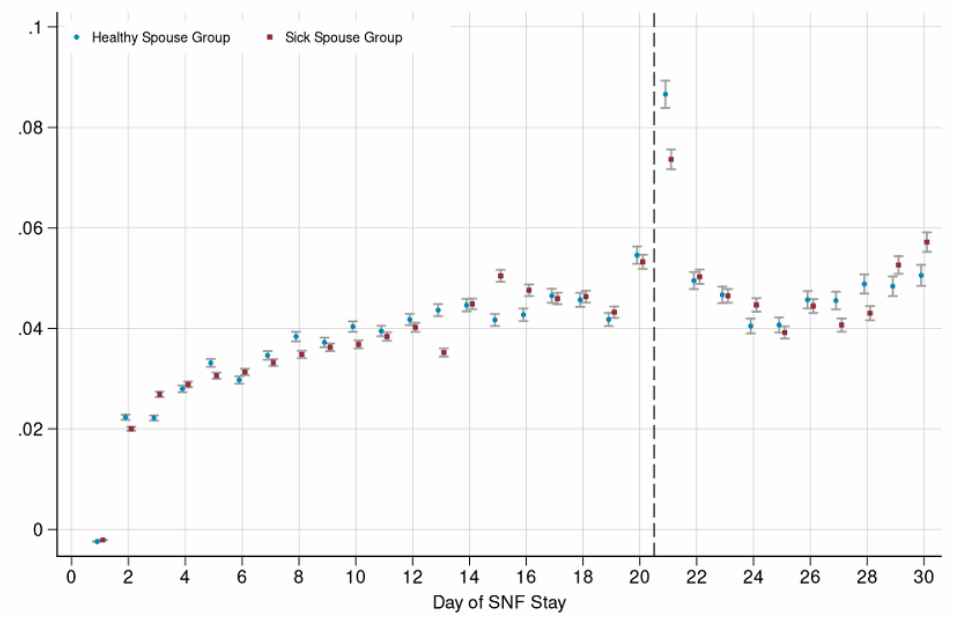
Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since a shock spouse's first heart attack or stroke. The outcome of interest is an indicator that equals one if at least part of that SNF stay was not fully covered by Medicare. All outcome spouses without SNF stays in both treatment and control groups are assigned values of 0 when no SNF admission occurred in a month. We rescale all coefficients such that they indicate the change relative to the initial unconditional mean values. Error bars plot 95% confidence intervals based on standard errors clustered at the household level.

Appendix Figure A21. Effect on Monthly Days in a SNF (Unconditional)



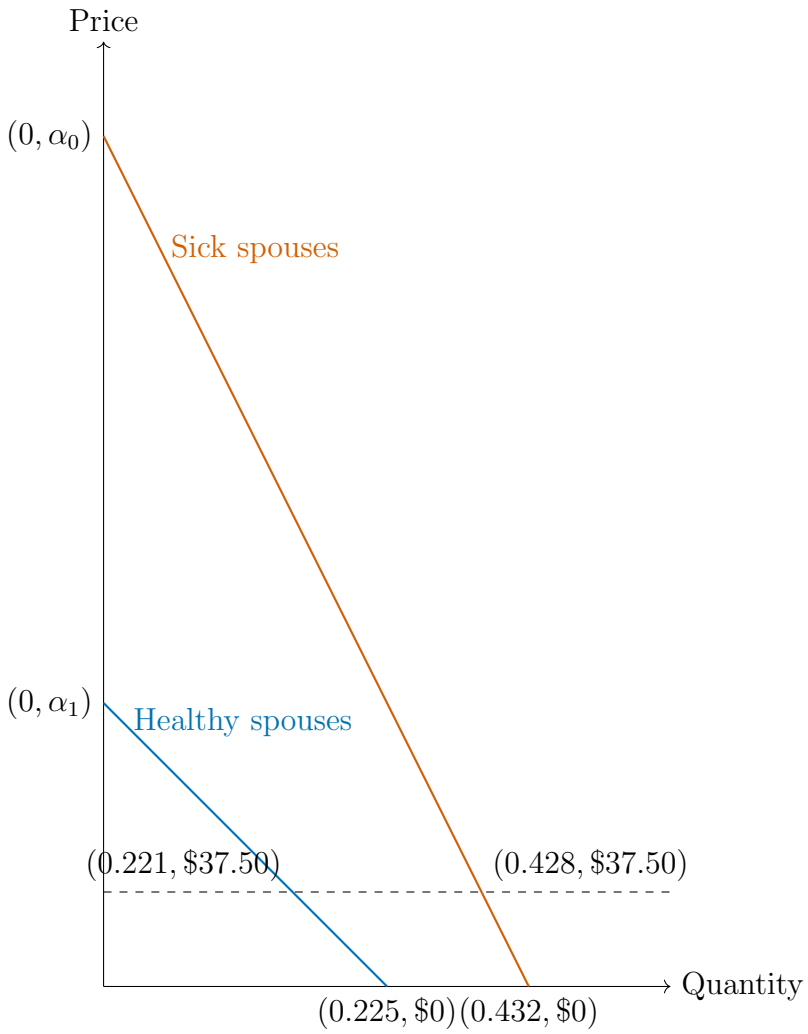
Notes: This figure plots estimates of the δ_n coefficients from equation (1), tracking the months since a shock spouse's first heart attack or stroke. The outcome of interest is a measure of days per month spent in a SNF across the entire sample, including months with zero SNF days in the sample. We rescale all coefficients such that they indicate the change relative to the baseline mean. The error bars plot 95% confidence intervals based on standard errors clustered at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects.

Appendix Figure A22. Residualized SNF Discharge Hazards, by Treatment Status and Length of Stay



Notes: This figure plots conditional discharge hazards for SNF stays based on whether the outcome spouse has a healthy spouse at home or not. Hazards are adjusted for the day of the week of admission and the number of days covered by Medicare previously (following Figure 7). The vertical dashed line indicates the day at which SNF patients begin to face cost sharing.

Appendix Figure A23. Estimating Demand for Formal Care across Groups



Notes: This figure demonstrates the intuition behind estimating linear inverse demand curves using variation in prices that Medicare recipients face for formal long-term care. See Section 4.3 for details. For each group, the quantity before cost sharing and the day-21 RD change provide two price-quantity points that identify a linear inverse-demand curve; these are derived from the RD estimates presented in Table 5. We use an expected daily out-of-pocket price for SNF care of $37.50 as discussed in Section 4.3.

Appendix Table A1. Effects of a Spouse's Major Cardiovascular Event on SNF Price Sensitivity, Stratified by Event Outcome

	(1)	(2)
A. RD, Healthy-Spouse Group		
Overall	-0.018***	-0.018***
	(0.0031)	(0.0017)
B. RD, Sick-Spouse Group		
Overall	-0.009**	-0.009***
	(0.0043)	(0.0024)
Fatal Event	0.003	0.003
	(0.0127)	(0.0062)
Event Ending in Discharge to SNF/Rehab	-0.011	-0.011**
	(0.0073)	(0.0041)
All Other Events	-0.010*	-0.010**
	(0.0058)	(0.0031)
Couple FEs		✓
Calendar Month FEs		✓
Bandwidth	11.22	11.22
N	380,673	380,673

Notes: Compare to Table 5. This table presents regression-discontinuity estimators identifying the effect of losing full Medicare coverage for SNF stays, which ends on day 21 for qualifying stays (equation (3)). Panel A presents the estimated effect of losing full coverage on the probability an outcome spouse will remain in the SNF prior to a spousal cardiovascular event. Panel B reports the post-event estimates stratified by event outcome. SNF stays within four months of the treatment and placebo events are included in the regression. The differences in Panel B between fatal and nonfatal events are statistically significant at the 99% confidence level. "FE" is an abbreviation for fixed effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A2. Effects of a Spouse's Major Cardiovascular Event on SNF Price Sensitivity, Stratified by Gender

	(1)	(2)	(3)
Shock Spouse:	Pooled	Female	Male
Outcome Spouse:	–	Male	Female
τ_{pre}	-0.018***	-0.015***	-0.021***
	(0.0030)	(0.0045)	(0.0039)
τ_{post}	-0.010***	-0.009**	-0.011***
	(0.0026)	(0.0040)	(0.0033)
$\beta_{\text{D-Disc}}$	0.008**	0.006	0.010*
	(0.0040)	(0.0061)	(0.0052)
Couple FEs	✓	✓	✓
Time FEs	✓	✓	✓
Bandwidth	11.22	11.22	11.22
N	380,673	164,358	216,315

Notes: Compare to Table 5. This table presents regression-discontinuity estimators identifying the effect of losing full Medicare coverage for SNF stays, which ends on day 20 for qualifying stays (equation (3)). In columns (2) and (3), we split the sample by the sex of the shock and outcome spouse. 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A3. Effects of Spousal Incapacitation on SNF Price Sensitivity, by Outcome Spouse's ADRD Status

	(1) Pooled	(2) No ADRD	(3) Has ADRD
τ_{pre}	-0.018***	-0.017***	-0.032***
	(0.0030)	(0.0031)	(0.0109)
τ_{post}	-0.010***	-0.010***	-0.010
	(0.0026)	(0.0027)	(0.0097)
$\beta_{\text{D-Disc}}$	0.008**	0.007*	0.022
	(0.0040)	(0.0041)	(0.0151)
Couple FEs	✓	✓	✓
Time FEs	✓	✓	✓
Bandwidth	11.22	11.22	11.22
N	380,673	356,983	23,690

Notes: Compare to Table 5. This table presents regression-discontinuity estimators identifying the effect of ending Medicare coverage for SNF stays, which ends on day 21 for qualifying stays (equation (3)). 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 an ADRD diagnosis in claims (6.2% of the sample) or not. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A4. Dynamic Effects on Outcome Spouse's Price Sensitivity for SNF Stay Decision

	(1) 0–90 days	(2) 91–180 days	(3) 181–270 days	(4) 271–365 days
τ_{pre}	-0.018***	-0.018***	-0.018***	-0.018***
	(0.0030)	(0.0030)	(0.0030)	(0.0030)
τ_{post}	-0.012***	-0.004	-0.015***	-0.011***
	(0.0029)	(0.0031)	(0.0031)	(0.0031)
$\beta_{\text{D-Disc}}$	0.006	0.015***	0.004	0.008*
	(0.0042)	(0.0043)	(0.0044)	(0.0043)
Couple FEs	✓	✓	✓	✓
Time FEs	✓	✓	✓	✓
N	329,452	294,009	291,916	289,386

Notes: Compare to Table 5. This table presents regression-discontinuity estimators identifying the effect of losing full Medicare coverage for SNF stays, which ends on day 20 for qualifying stays (equation (3)). The first two rows present the estimated effect of losing full 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. Each column indicates the duration of time for which SNF stays are included in the regression (relative to both treatment and placebo events). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B Appendix: Option Value versus Full Valuation

The main text provides intuition for the claim that the option value, $\text{OV}(p) \equiv CS_0(p) - CS_1(p)$, is a lower bound for the full valuation of spousal care. This appendix provides a formal derivation of that claim.

We embed the formal-care demand curves in a three-option choice model. Let $h = 1$ denote the healthy-spouse state and $h = 0$ the incapacitated-spouse state. Beneficiary i assigns money-metric values u_i^F and u_i^S to an additional day of formal care and to spousal care, respectively. We normalize the residual outside option to zero. For this valuation exercise, we impose a nested-choice-set assumption: the value of formal care and the outside option are fixed across states, and a healthy spouse changes the choice set only by making spousal care available. Individual indirect utility is therefore

$$\mathcal{U}_1^i(p) = \max\{u_i^F - p, u_i^S, 0\}, \qquad \mathcal{U}_0^i(p) = \max\{u_i^F - p, 0\}.$$

Population-average indirect utility is $\mathcal{U}_h(p) = \mathbb{E}[\mathcal{U}_h^i(p)]$.

Let $\bar{p}$ denote a common price high enough that formal care is chosen in neither spouse-health state. The purpose of $\bar{p}$ is to remove formal care from the choice set. At that price, a beneficiary with a healthy spouse can choose between spousal care and the outside option, whereas a beneficiary without a healthy spouse has only the outside option. The utility difference between the two states therefore equals the full value of access to spousal care:

$$FV_S \equiv \mathcal{U}_1(\bar{p}) - \mathcal{U}_0(\bar{p}) = \mathbb{E}\left[(u_i^S)_+\right].$$

Next define the value of access to spousal care at the prevailing price as

$$R(p) \equiv \mathcal{U}_1(p) - \mathcal{U}_0(p) \geq 0.$$

We call $R(p)$ the "unrevealed remainder." We will derive that

$$FV_S = \text{OV}(p) + R(p). \tag{11}$$

Equation (11) therefore implies that $\text{OV}(p) \leq FV_S$.

Let $Q_h(p)$ denote the share of beneficiaries who choose formal care in state h. To see how demand relates to indirect utility, consider a small price increase dp.

For a beneficiary who chooses formal care, the increase reduces utility by dp, because formal-care utility is $u_i^F - p$. It has no direct effect on beneficiaries who choose spousal care or the outside option. Moreover, beneficiaries who switch options following the price increase are indifferent at the point of switching, so their change in choice has no first-order effect on utility. The envelope theorem therefore gives

$$\frac{\partial \mathcal{U}_h^i(p)}{\partial p} = -\mathbf{1}\{i \text{ chooses formal care in state } h\}.$$

Averaging across beneficiaries yields

$$\frac{\partial \mathcal{U}_h(p)}{\partial p} = -\Pr(i \text{ chooses formal care in state } h) = -Q_h(p).$$

Integrating from the prevailing price p to the choke price $\bar{p}$ gives

$$\mathcal{U}_h(p) - \mathcal{U}_h(\bar{p}) = \int_p^{\bar{p}} Q_h(t)\, dt \equiv CS_h(p).$$

Thus, consumer surplus is the utility gained from making formal care available at price p, relative to the benchmark in which it is priced out.

Using this expression in the two spouse-health states, the option value can be written as

$$\begin{aligned} \mathrm{OV}(p) &= CS_0(p) - CS_1(p) \\ &= \big[\mathcal{U}_0(p) - \mathcal{U}_0(\bar{p})\big] - \big[\mathcal{U}_1(p) - \mathcal{U}_1(\bar{p})\big] \\ &= FV_S - R(p). \end{aligned}$$

Rearranging then yields (11).

The remainder has a simple economic interpretation. For beneficiaries who switch from formal care to spousal care when a healthy spouse is available, $R(p)$ contains the additional value they place on spousal care beyond the formal-care surplus they forgo. For beneficiaries who choose no formal care in either state, it contains their entire valuation of spousal care, which does not appear in either formal-care demand curve. These are the two components that make the option value a potentially strict lower bound.

C Appendix: Dynamic Model

Eliason et al. (2018) study length of stay in long-term-care hospitals. We adapt their model to study stays in a SNF. Assume that households have utility $u_t = \lambda_t + \alpha\, p_t$ for day t in a SNF. Here, λ_t represents non-monetary gains and p_t represents the monetary costs, with $\alpha < 0$. The monetary cost of staying in the SNF is very low for days $t \leq 20$ but increases suddenly on day $t = 21$ and remains high thereafter.

The household then faces a discrete-time decision problem. At the start of day t, the admitted patient in the household chooses either to remain in the SNF or to exit. The household experiences choice-specific idiosyncratic shocks $\xi_t = (\xi_{kt}, \xi_{dt})$, the components of which relate to staying or leaving the SNF, respectively.

Following the notation of Eliason et al. (2018), the household's dynamic problem is given by:

$$V_t(\xi_t) = u_t + \max\Big\{\xi_{kt} + \delta\, EV_{t+1},\ \xi_{dt}\Big\},$$

where EV_{t+1} is the expected continuation value of staying in the SNF at time $t+1$ and δ is the household's daily discount factor. Given the short time horizon, we set δ to 1.

We assume a finite time horizon and define a finite termination value of staying in the SNF on day 100, $\Omega \equiv EV_{100}$. We also assume that ξ_{kt} and ξ_{dt} are distributed via a standard type-I extreme-value distribution. Those assumptions lead to the relationship:

$$EV_t = u_{t+1} + \log(\exp(\delta EV_{t+1}) + 1).$$

Identifying choice probabilities requires estimating the following parameters: the λ_t parameters that represent the utility gain from each night in a SNF; α, the utility cost of the price of each night in a SNF; Ω, the termination value of the household's problem on the 100^{th} day; and the variance of ξ_t.[51] We model λ_t as a cubic polynomial in t, resulting in four parameters to estimate in $\vec{\lambda}$. This amounts to six parameters to estimate in total: the cubic polynomial in t for λ_t, α, and Ω. As is best practice, we normalize the variance of ξ_t as this would otherwise not be separately identified.

We estimate this model via the generalized method of moments. We divide the sample into two groups: households with a healthy spouse at home and those whose spouse has been incapacitated. We use the empirical distribution of length of stay in a SNF for each of the first 50 possible days of a stay. This provides 50 moments

[51] We assume that ξ_{kt} and ξ_{dt} have the same variance.

to identify the six parameters above, where each moment is the share of individuals leaving a SNF on a given day. To this, we add a 51^{st}: the difference in discontinuities estimate presented in Table 5 (β_1 in equation (3)). This moment reflects the difference in price sensitivity across groups.

Appendix Table C1 presents the estimated parameters for each group. In addition, Appendix Figure C1 illustrates how these parameters result in implied length-of-stay distributions matching the observed patterns in the data. Overall, the values for the model parameters give rise to distributions similar to those observed in the data, including key responses to price changes among both groups.

Appendix Table C1. Estimated Model Parameters

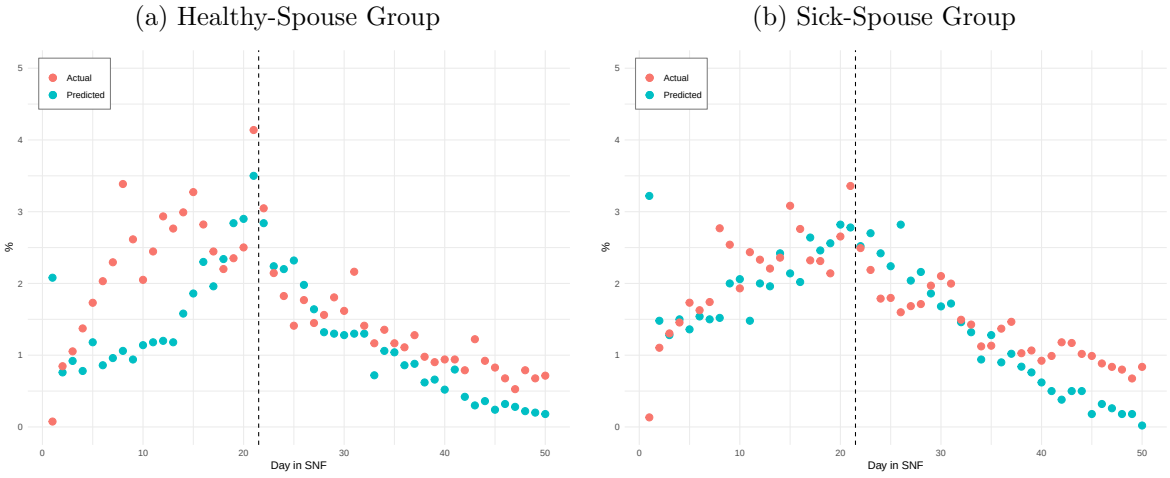
Parameter	**Description**	**Healthy-Spouse Group**	**Sick-Spouse Group**
λ_0	Utility from SNF, Intercept	-11.206	-5.680
λ_1	Utility from SNF, t	0.100	0.030
λ_2	Utility from SNF, t^2	-0.007	-0.005
λ_3	Utility from SNF, t^3	7.667×10^{-5}	5.945×10^{-5}
α	Price coefficient	-0.097	-0.009
Ω	Termination value	-2.928	-1.489

Notes: This table reports parameters of the dynamic model in Appendix C, estimated by the generalized method of moments (GMM). The model is estimated separately on those in the healthy spouse (control) and sick spouse (treatment) groups.

The model estimates suggest that individuals with healthy spouses at home are much more sensitive to changes in the price of formal care than those without an option for informal care. The price coefficient is -0.097 in the healthy-spouse group, more than 10 times as large in absolute value as the coefficient of -0.009 in the sick-spouse group. In addition, the model allows us to compare how overall demand for formal and informal care changes in the presence of alternative pricing and insurance schemes.

The main text describes one way to trace out the implications of this calibration via Figure 8 and Appendix Figure C2. Appendix Table C2 presents several alternative extrapolations. It considers two types of counterfactuals: changes in the household coinsurance rate and changes in the day cost sharing begins. For each scenario, we calculate the change in the overall use of the SNF, the average length of stay, the

Appendix Figure C1. Observed and Model-Implied SNF Length-of-Stay Distributions



Notes: This figure plots both the empirical and implied distributions for LOS among SNF stays in two groups. Implied distributions are constructed using the model described in Appendix C. In panel (a), distributions are compared for SNF occupants without an incapacitated spouse; in panel (b), we compare distributions for those whose partner has had a major cardiovascular event in the last four months.

average cost of a stay, and the average household consumer surplus for each group. We calculate the average ex ante expected cost of a stay taking into account the relatively high fraction of this population with Medigap coverage. Finally, consumer surplus is calculated as the estimated differences in average consumer utilities across states of the world, scaled by the relative price coefficient, α, in each group.

Those without a healthy spouse at home value the use of formal care roughly 8.5 times as much as those with healthy spouses. The sick-spouse group utilizes SNF care about 63% more than the healthy-spouse group in the four months following a spouse's health shock, and values that care at $179, on average, compared to $21 for the healthy-spouse group.[52]

Those simulations emphasize two main patterns. First, shifting how households are exposed to cost sharing changes utilization much more among the healthy-spouse group than among the sick-spouse group. That pattern is intuitive, given the higher price elasticity for those with a healthy spouse at home. Second, policies taking into account differences in moral hazard between the two groups can improve outcomes. Lowering coinsurance increases consumer welfare for both groups, but conditional

[52]These results are averaged over the entire population of each group. Conditional on a SNF stay, these valuations are $7,587 for the average SNF stay among the sick-spouse group, compared to $1,773 for the average SNF stay in the healthy-spouse group.

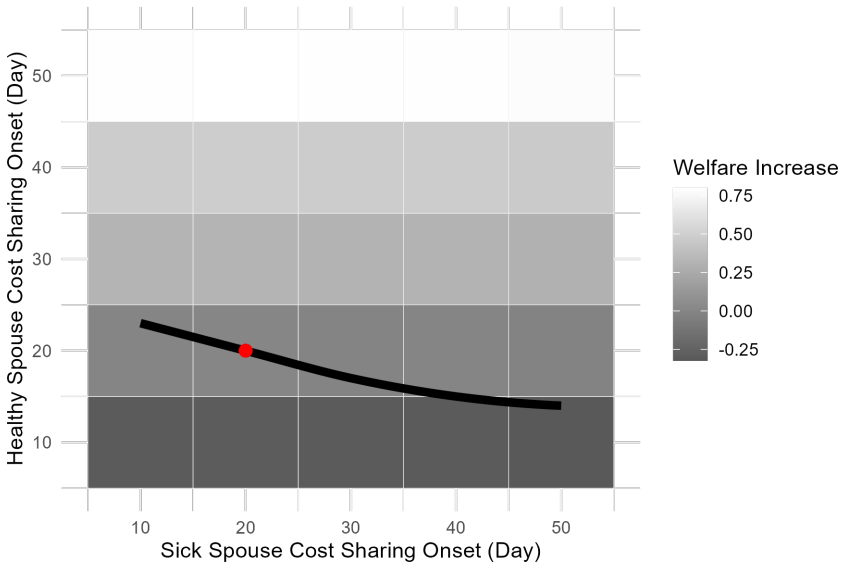
Appendix Table C2. Welfare Effects of Changing Prices for Formal Care

	Healthy-Spouse Group			Sick-Spouse Group		
	LOS	OOP	CS	LOS	OOP	CS
Base Case	0.52	\$13	\$21	0.85	\$19	\$179
A. Shifting patient cost sharing						
0% Coinsurance	1.03	–	\$57	1.10	–	\$216
50% Coinsurance	0.17	\$2	\$14	0.52	\$20	\$149
100% Coinsurance	0.16	\$1	\$13	0.35	\$6	\$137
B. Shifting day of cost-sharing onset						
Day 0	0.32	\$12	\$12	0.75	\$27	\$182
Day 7	0.46	\$15	\$20	0.81	\$25	\$197
Day 30	0.66	\$13	\$30	0.84	\$14	\$214
Day 60	0.96	\$13	\$41	0.86	\$8	\$223

Notes: This table presents estimated utilization and welfare measures based on the estimated model parameters following Section 4.1. The base case, presented in the first row, is when copayments begin on day 21. Each row presents results under a different counterfactual state of the world. Counterfactual welfare estimates are calculated separately for those in the healthy spouse (control) and sick spouse (treatment) groups. All columns are unconditional (on a SNF stay). Consumer surplus is measured as estimated utility normalized by group-specific price coefficients, α.

on entering the SNF, it leads to much longer stays in SNFs for those with a healthy spouse. We estimate that conditional LOS for this group nearly doubles, from 31 days to 61 days. By contrast, SNF patients with a sick spouse change their conditional length of stay by only 3.8%. The table thus suggests that eliminating cost sharing for those with a sick spouse changes behavior almost entirely along the extensive margin, increasing the likelihood of entering a SNF but with little impact on conditional stay length. Overall, the simulations suggest that tailoring cost sharing to whether a patient has a sick spouse at home could improve welfare for those with the greatest needs while limiting distortions in care use.

Appendix Figure C2. Welfare Gains from Subsidizing SNF Coverage by Spousal Incapacitation



Notes: This figure plots estimated changes in total household welfare using the dynamic model presented in Appendix C. Each point in the plot represents a different combination of government coverage generosity for SNF stays, indicated as the day that cost sharing starts for SNF stays. For each point, the total household welfare is shown, relative to the status quo of $(20, 20)$ (indicated by the circular marker). The black line indicates the set of policies which are budget-neutral relative to the status quo. Compare to Figure 8.