The Protective Effects of a Healthy Spouse: Medicare as the Family Member of Last Resort*

Itzik Fadlon† Tal Gross‡ Alex Hoagland§ Tim Layton¶

December 2024

Preliminary and incomplete
Please do not cite or circulate

Abstract

The aging of the US population has been a major source of concern, in part due to the high and rising costs of elderly care. We study the demand for elderly care among American couples, both formally (provided in nursing facilities) and informally (provided by family members at home). Using administrative Medicare claims data and a novel spousal linkage, we analyze how the health of one spouse affects the need for formal care of the other by constructing counterfactual to affected households using households that experience similar events later in time. We find that, immediately following an individual's incapacitation (for a heart attack or stroke), their spouses exhibit a 19% increase in the likelihood of entering a skilled nursing facility. With this effect being accompanied by sharp increases in hospitalizations and emergency department visits (on the order of 9% and 19%, respectively), we show that the increased utilization of skilled nursing facilities is driven by both direct health effects upon spousal incapacitation and changes in the use of formal care conditional on health. We offer revealed preference methods to evaluate households' valuation of formal and informal elderly care, which leverage a combination of our quasi-experimental variation with price variation of formal care, exploiting our finding that affected households are much less sensitive to the price of formal care which reveals their higher valuation. We discuss the implications of our results for optimal health insurance contracts, showing that the types of within-household spillovers that we uncover can provide justification for household-level deductibles and out-of-pocket maximums. Overall, our analysis provides new evidence on the interaction between formal and informal care and on their valuation among elderly American couples.

*We are grateful to Ashvin Gandhi, Jonathan Kolstad, Bruce Landon, and participants at the 2024 California Health Economics Conference for helpful feedback. Prabidhik KC provided superb research assistance. We acknowledge funding for this project from the National Institute on Aging.

†UC San Diego and NBER. fadlon@uscd.edu

‡Boston University and NBER. talgross@bu.edu

§University of Toronto. alexander.hoagland@utoronto.ca

¶University of Virginia and NBER. timothyjlayton@virginia.edu

1 Introduction

One of the challenges that an aging population presents is the growing need for elderly care (Johnson et al., 2021; Dall et al., 2013). When the elderly cannot care for themselves, then care must be provided, either formally or informally. Formal care for the elderly consists of nursing homes, skilled nursing facilities, home-based health aids, and other types of care often paid by Medicare or Medicaid. Informal elderly care consists of the time and resources that an older individual's family members invest to care for their needs.

The measurement of formal elderly care is straightforward—it can be directly observed in insurance claims and financial statements. Informal care, by contrast, is difficult to measure and even more difficult to value. When providing care, family members do not clock in, track their hours, or log their effort. As a result, informal care is typically unobserved and unmeasured.[1] Further, the opportunity cost of the time informal caregivers devote to caregiving activities is even more opaque and difficult to value, and differences in the quality and valuation of formal versus informal care are hard to assess. This has led to a strikingly wide range of estimates of the value of informal care, ranging from $86 billion (Gruber and McGarry, 2023) to over $500 billion (Chari et al., 2015).

In this paper, we study one of the most important sources of informal elderly care: care from a healthy spouse. In the US, care delivered by spouses accounts for roughly one-third of all informal care (Gruber and McGarry, 2023). Among married individuals, spouses provide the majority of informal care (Mommaerts, 2024). We identify the effects of receiving informal care from a healthy spouse and its valuation by leveraging plausibly exogenous changes in the availability of that care. Specifically, we identify Medicare beneficiaries who are hospitalized for a heart attack or stroke, and we then use event-study methods to recover the causal effects of that spouse's incapacitation on their partners. Our analysis relies on novel data that allow us to link across American couples covered by Medicare and study formal elderly care outcomes using administrative Medicare claims.

Event studies reveal large effects of the incapacitation of an individual (the "index spouse") on their spouse's utilization of formal care (the "focal spouse").[2] In the

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

[2] We focus on opposite-sex couples in this analysis to avoid problems of identifying same-sex spouses separately from aging siblings or people of the same sex who share an address.

months following an index spouse's health shock, partners become 19.1% more likely to be admitted to a skilled nursing facility (SNF).

To interpret this finding, we focus on two potential channels for the increase in SNF utilization: a health effect and a substitution effect. First, one spouse's health shock can directly lead to a worsening of their partner's health, which induces a need for formal care. Indeed, we find that affected spouses exhibit an 18.7% increase in emergency department visits and are 8.9% more likely to be hospitalized. Second, a substitution effect describes those who shift from receiving informal care to formal care because informal care, previously provided by their partner, is no longer available. Indeed, people with chronic conditions that require care are 3–4 times more likely to visit a SNF after their spouse's are incapacitated than those without such conditions. Similarly, we find that focal spouses are more likely to visit a SNF the more incapacitating their partners' health shock.

We then assess the welfare implications of those findings. Social welfare is composed of two components: the implications for the government's budget, that is, the "fiscal externality," and the household's valuation of formal and informal care. To assess the fiscal externality, we estimate the effects of index spouses' health shocks on total Medicare spending for focal spouses. The incapacitation of an index spouse leads to a 5% increase in spending for focal spouses in the year of the event, which increases to 12% when the event is fatal. Overall, healthy spouses lower US healthcare spending by $1.26 billion each year.

We then develop two complementary methods that assess households' "willingness to pay" for informal care. The two methods involve revealed-preference exercises that measure households' choices of formal care when the index spouse is incapacitated versus when the index spouse remains healthy. The first method evaluates how the marginal valuation of SNF days varies across states of the world. It relies on within-state first order conditions of SNF days and the fact the price schedule is invariant to the index spouse's incapacitation. For a plausible range of the relative risk aversion parameter, we find that an additional day of SNF is valued (relative to income) by 43% to 172% more after the event.

The second method exploits Medicare's discontinuous change in the out-of-pocket price of a night in a SNF. On the 20^{th} day of a SNF visit, Medicare recipients' out-of-pocket costs shift from 0% to 20% coinsurance. We find that individuals whose partners recently experienced a health shock are 64% less likely to leave the SNF upon facing co-payments, exhibiting a meaningfully more inelastic demand for formal

care. This pattern implies a higher willingness-to-pay for formal care in the absence of a healthy spouse, which in turn reveals the option value of the availability of spousal informal care. We use these estimates to plot demand curves for SNF care and evaluate consumer surplus. We find that affected households are willing to pay roughly 3.5 times more for formal care to the focal spouse.

Finally, we consider the implications for the design of optimal healthcare policies. The differential slopes of the demand curves across states of nature map directly to the magnitude of moral hazard in the consumption of formal care. The ratio of elasticities governs the relative deadweight loss from subsidies to formal elderly care across households with and without healthy spouses. The estimates accordingly imply a 64% reduction in efficiency loss among affected households. Medicare can provide higher levels of coverage for SNF care with lower deadweight loss due to moral hazard when a beneficiary's spouse is incapacitated versus when their spouse is healthy. This in turn provides strong support for the design of insurance contracts with family-level deductibles or a shared out-of-pocket maximum.

Our work contributes to three areas of inquiry. First and foremost, we contribute to research on elderly and long-term care in the US, particularly informal care. Gruber and McGarry (2023) estimate that formal care expenditures totaled $280 billion in 2018, which were financed mostly by the public sector, and that informal care plays a critical role in providing for the elderly, showing that at every age and disability level they receive informal care more frequently than formal care. Studies have also found that informal care may delay entry into formal care (Bergeot and Tenand, 2023; Bonsang, 2009; Charles and Sevak, 2005) and significantly affects the labor supply and earnings of caregivers, having effects on households that go beyond health (Maestas et al., 2024). We contribute to this literature by providing, to the best of our knowledge, the first analysis of intra-family interactions in the demand for elderly care in the US using administrative data at the population level. We provide novel analysis and estimates for the value of informal care by isolating the causal relationships in these interactions, evaluating the monetary value of the spousal relationship in elderly care, and assessing beneficiaries' valuation of formal care when their spouse becomes incapacitated.

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

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

Third, this paper uncovers healthcare dynamics that can inform how health insurance contracts ought to be structured. When one person is experiences a health shock, the risks that their spouse faces change. And yet, the Medicare program offers individual health insurance plans that make no adjustment for the rest of the household. Many employer-provided plans also feature individual-specific rather than household-specific out-of-pocket limits or deductibles. This paper provides new evidence that can help assess the optimality of such arrangements. Our work is therefore related to previous work estimating optimal levels of social insurance (Baily, 1978; Chetty, 2006; Chetty and Finkelstein, 2013), where we provide novel findings with respect to the design of healthcare policies which are large and growing.

The remainder of the paper proceeds as follows. Section 2 provides our empirical framework by describing the data we use and our identification strategy. Section 3 provides a conceptual framework that organizes our empirical analysis of the causal effects of spousal shocks. Section 4 provides the main empirical evidence of the effects of a index spouse's health shock on their partner's health and healthcare utilization. Section 5 provides our welfare analysis of households' valuation of informal care and the optimal design of healthcare policies. Section 6 concludes.

2 Empirical Framework

2.1 Data

We use Medicare claims data that cover all types of care, including primary and specialty care, inpatient and outpatient care, and formal care such as utilization of skilled nursing facilities. We have created a novel dataset of couples in Medicare using a custom crosswalk from the Centers for Medicare and Medicaid Services (CMS). The dataset includes beneficiaries' addresses as tracked by CMS. We then group beneficiaries who live at the same address and identify spouses by selecting residences with

a maximum of two Medicare recipients of the opposite sex living at the same address. This allows us to study Medicare-covered healthcare among American couples with large scale administrative data. Gross et al. (2024) validate this identification of spouses, specifically by comparing the gaps in ages for spouses identified with this crosswalk to spouses identified in the American Community Survey.

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

In our analysis, we require households to have full Medicare eligibility throughout the window of observation, and we assign spouses in the year prior to the index health event. We also exclude households where the index event has accompanying external injuries so as to avoid capturing joint events affecting both members of the household (for example, vehicle accidents). Finally, in keeping with previous work in order to focus on "new" events, we require the spouse who has the health shock to have not had a prior hospital admission in the two years preceding their index admission (Dobkin et al., 2018).

2.2 Research Design

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

We construct a treatment group, composed of spouses in households that experience the event in period τ, and a matched control group (based on event time alone), composed of spouses in households that also experience the event, but in period $\tau+\Delta$. We assign a placebo event to control households in period τ. We then recover the treatment effect by using a dynamic difference-in-differences estimator, identifying

the event's impact purely from the change in the differences in outcomes across the two groups over time. Accordingly, our analysis is not subject to potential challenges involved in having units that switch in and out of experimental arms as posed by recent work (De Chaisemartin and d'Haultfoeuille, 2024). Similarly, our results are also robust to using estimators that explicitly model heterogeneous treatment effects, specifically the local-projection difference-in-differences estimator developed by Dube et al. (2023).

Our choice of Δ is one year, so that we match households that experience health shocks 12 months apart. This implies that we can estimate the effects of shocks up to 12 months after the event since the control group becomes "treated" 12 months later. We normalize the time of observation such that event-time period n is measured with respect to the assigned event period. That is, $n \equiv$ calendar week $- \tau$, where $n = 0$ is when the assigned event occurs (the actual event for the treatment group and the placebo event for the control group). Table 1 presents summary statistics for our analysis sample. The treatment and matched control groups are closely similar in terms of patient demographics and healthcare utilization at the baseline year prior to the event (the actual event for the treated group and the placebo event for the matched control group).

Our estimating equation of the dynamic difference-in-differences estimator is of the following event-study type:

$$y_{i,t} = \alpha_i + \theta_n + \sum_{n \neq -1} \delta_n \times I_n \times \text{treat}_i + \lambda X_{i,t} + \varepsilon_{i,t}, \tag{1}$$

where $y_{i,t}$ denotes an outcome for household i at calendar time t, treat_i denotes an indicator for whether a household belongs to the treatment group, and I_n are indicators for months relative to the assigned event period. Equation (1) includes household fixed effects (α_i) and relative-time fixed effect (θ_n), and a vector of potential time-varying controls in $X_{i,t}$. The key parameters of interest are the δ_n coefficients, which estimate the period n treatment effects relative to the pre-period $n = -1$. Our empirical strategy rests on the standard parallel-trends identifying assumption. It assumes that, absent the realization of the health shock, the outcomes of the treatment and control groups would evolve in parallel. To test the validity of this assumption, we study the evolution of the treatment and the control groups' outcomes in the periods prior to the event. In order to calculate average effects over several

Table 1. Summary Statistics

	Index Spouses	Focal Spouses
Panel A: Demographics		
Age	76.39	75.72
Female	0.38	0.63
White	0.89	0.89
Black	0.06	0.06
Hispanic	0.01	0.01
FFS Status	0.68	0.67
HMO Status	0.31	0.31
Dual Eligibility Status	0.07	0.07
Panel B: Healthcare Utilization		
Has a Chronic Condition	0.45	0.42
Predicted Risk of SNF Visit	0.00	0.00
Any SNF Stay	0.00	0.02
Conditional # of SNF Stays	1.11	1.36
Any Inpatient Admission	0.00	0.11
Conditional # of Inpatient Admissions	1.18	1.54
Total Spending	$5,311	$8,430
Year	2012.54	2012.54
N	624,901	624,901

Notes: This table presents summary statistics for the analytical sample. Index spouses are those who experienced a first heart attack or stroke as discussed above; focal spouses are their partners whose outcomes we study. Averages with standard errors in parentheses are presented for the year prior to the true index event.

periods, we estimate a standard difference-in-differences specification of the following form:

$$y_{i,t} = \alpha_i + \theta \times \text{post} + \delta \times \text{treat}_i \times \text{post}_{i,t} + \lambda X_{i,t} + \varepsilon_{i,t}, \tag{2}$$

where $\text{post}_{i,t}$ is an indicator for whether the observation belongs to a period after the shock occurred, and δ is the estimate for the average treatment effect.

3 Conceptual Framework

This section presents a framework to conceptualize the effects of a spouse's health shock to guide the empirical analysis.

3.1 Setup

We consider two periods, $t \in \{0, 1\}$, where $t = 0$ represents time prior to the index spouse's incapacitating health event and $t = 1$ represents time after the event. In addition, the focal spouse has two possible states of a need for care, denoted by h_t, such that $h_t = 0$ indicates a healthy baseline state and $h_t = 1$ indicates a sick state with conditions that require care. Conditional on being sick ($h_t = 1$), the focal spouse's mode of receiving care is denoted by c_t, with $c_t = 0$ indicating informal care and $c_t = 1$ indicating formal care. The status $c_t = 0$ can be thought of as including a fallback option as well (e.g., self-care and care from other family members), and one can also extend the simple framework to model it directly. Overall, we need to trace the time period relative to the index spouse's incapacitation event ($t \in \{0, 1\}$), the focal spouse's state of need for care (h_t), and the mode in which care is delivered conditional on need (c_t).

The focal spouse transitions across states via the transition matrix in Appendix Table A1. Rows represent states in period 0 and columns represent states in period 1. Each cell provides the conditional probability of moving to the given column starting from a given row.

Capturing our main quasi-experiment of the index spouse's health event, in the transition to period 1 informal care is no longer available to a sick focal spouse. Hence, the transition probabilities for the second column are all 0, and previously sick focal spouses relying on either informal or formal care are assumed to move to formal care after their spouse's incapacitation. We model sickness as an absorbing state since our context is elderly care for conditions with long-term needs, which accordingly

restricts the transition probabilities in the first column. Note that by construction, $p_{0,0} + p_{0,(1,1)} = 1$.

3.2 Decomposing the Effects of Spousal Health Shocks

One key goal of our analysis is to identify the impact of spousal incapacitation on the demand for formal care. We denote this parameter by τ, where

$$\tau \equiv \Pr(c_1 = 1) - \Pr(c_0 = 1).$$

In terms of interpretation, τ can be thought of the elderly-care option value of a healthy spouse based on the causal effect of losing that option.

This then lends itself to a straightforward decomposition. Since we have that

$$\Pr(c_0 = 1) \;\; = \;\; \Pr(h_0 = 1, c_0 = 1),$$

and

$$\begin{aligned} \Pr(c_1 = 1) &= \Pr(h_1 = 1, c_1 = 1) \\ &= \Pr(h_0 = 0) \times p_{0,(1,1)} + \Pr(h_0 = 1, c_0 = 0) + \Pr(h_0 = 1, c_0 = 1), \end{aligned}$$

it follows that

$$\tau = \underbrace{\Pr(h_0 = 0) \times p_{0,(1,1)}}_{\text{Health Effect} \equiv \tau_H} + \underbrace{\Pr(h_0 = 1, c_0 = 0)}_{\text{Substitution Effect} \equiv \tau_S}. \tag{3}$$

That is, we can decompose the overall effect of the spouse's health shock on the focal spouse's formal care into two components. First, the shock can directly worsen the focal spouse's health which then leads to need for care, denoted by τ_H. Second, focal spouses who have a need for care at baseline may substitute to consuming formal care instead of the potential informal care from the spouse who became incapacitated, denoted by τ_S.

3.3 Effects on Social Welfare

To assess social welfare, we consider the differentials in the household's value function across states. We denote the value by $V(h_t, c_t)$, which is a function of the focal spouse's need for care (h_t) and their mode of receiving care conditional on need (c_t). We define $V_S \equiv V(1,0) - V(1,1)$ as the value of informal care versus formal care

for sick focal spouse who received informal care at baseline. We additionally define $V_H \equiv V(0,\cdot) - V(1,1)$ as the value of being healthy (relative to a counterfactual of being sick and receiving formal care). We further denote by μ the marginal value of public funds, which corresponds to the shadow value of government resources that we denote by R_t. With this formulation, we can represent social welfare by:

$$W_t = E_t(V) + \mu R_t,$$

where the expectation over the household's value is given by

$$E_t(V) = \Pr(h_t = 0) \times V(0) + \Pr(h_t = 1, c_t = 0) \times V(1,0) + \Pr(h_t = 1, c_t = 1) \times V(1,1).$$

We normalize the utility valuations by μ so that they are represented in their money-metric dollar equivalent, defining $v_S \equiv V_S/\mu$ and $v_H \equiv V_H/\mu$ accordingly. We denote by e_H the average expenditure per beneficiary who experiences a health condition, and we denote by e_F the average expenditure per beneficiary who utilizes formal care. Defining $w_t \equiv W_t/\mu$, we have that the money-metric change in welfare following the index event is given by:

$$\triangle w \equiv w_1 - w_0 = \underbrace{-v_H\tau_H - v_S\tau_S}_{\text{Household Welfare Effect}} + \underbrace{-(e_H + e_F)\tau_H - e_F\tau_S}_{\text{Fiscal Externality}= (R_1 - R_0)}.$$

This first component is the loss in household welfare from the health effect and from the imposed substitution away from informal care. The second component is the change in expenditure $(R_1 - R_0)$, which represents the fiscal externality on Medicare costs.

3.4 Parameter Estimation

Our analytical approach allows us to directly identify τ, the effect of a spousal health shock on the utilization of formal care. We will then provide different empirical exercises to gauge the presence and extent of the underlying channels: a health effect (τ_H) and a substitution effect (τ_S). We will study health effects by measuring changes in health conditions and proxies for poor health, such as hospitalizations and ED visits. We will study the substitution effect by examining differences in responses across focal spouses with pre-existing conditions as well as events with differing levels the index spouse's incapacitation. We will additionally offer a decomposition exercise

of τ into τ_H and τ_s based on predictions of the probability of a SNF stay using the focal spouse's demographics and baseline health conditions.

We can also directly estimate $(R_1 - R_0)$ by measuring the causal effect of the index event on total Medicare spending for focal spouses. This allows us to evaluate the fiscal externality effect of losing the option of a healthy spouse.

Finally, in Section 5.2, we develop our two revealed preference approaches to infer v_S. We will infer the average dollar valuation of informal care from spouses by excess valuation of formal care when informal care becomes absent. The first exercise will focus on marginal valuations of additional formal care and the second exercise will bring in an additional source of variation (in price) in conjunction with the quasi-experiment of the health event to infer total consumer surplus.

4 Empirical Evidence: Event-Study Analyses

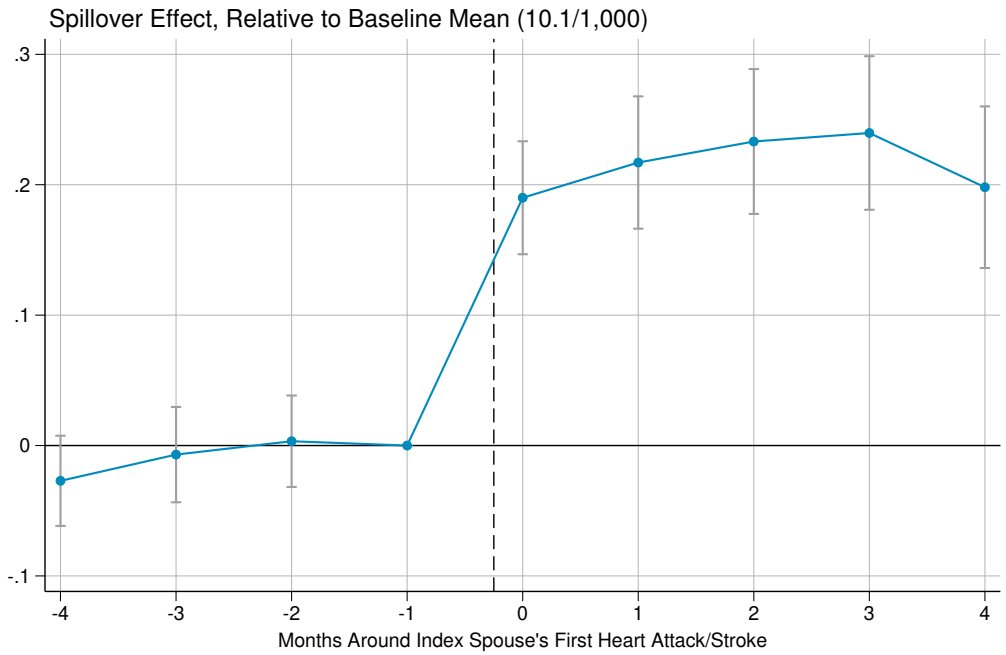
This section presents an empirical analysis of the effects of a spouse's health shock on their partner's healthcare utilization and health. We study index spouses of both genders together. Appendix B provides stratified results, showing little variation across the gender of the index spouse.

4.1 Utilization of Formal Care

We begin with Figure 1 that examines the risk of a focal spouse visiting a SNF after the index spouse's health event. This figure and the subsequent event-study figures plot the δ_n coefficients from estimates of equation (1). The x-axis plots time relative to the index event in month, and the y-axis plots the outcome. Average baseline levels in the 6 months prior to the event are reported in the subtitle of the plot, and the outcome variables are normalized such that plotted coefficients are presented in percent relative to baseline.

The figure first shows no systemic differentials across experimental groups in the periods prior to the event, in support of the research design. Then, the figure uncovers a sharp increase in the risk of a focal spouse's visit to a SNF immediately following the index health shock. This increased utilization remains persistent throughout the analysis horizon of 4 months. Column 1 of Table 2 reports the effects for the month of the shock using equation (2), showing an increase of 19.1% in the risk of a SNF visit. Appendix Figure A1 presents longer-run estimates: the observed increases persist for the entire year following the index event, reverting back to baseline only by month

Figure 1. Effect of a Spouse's Major Cardiovascular Event on Risk of a SNF Visit



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

12.

We now turn to focus on the potential mechanisms that may drive this effect.

4.2 Health Effect

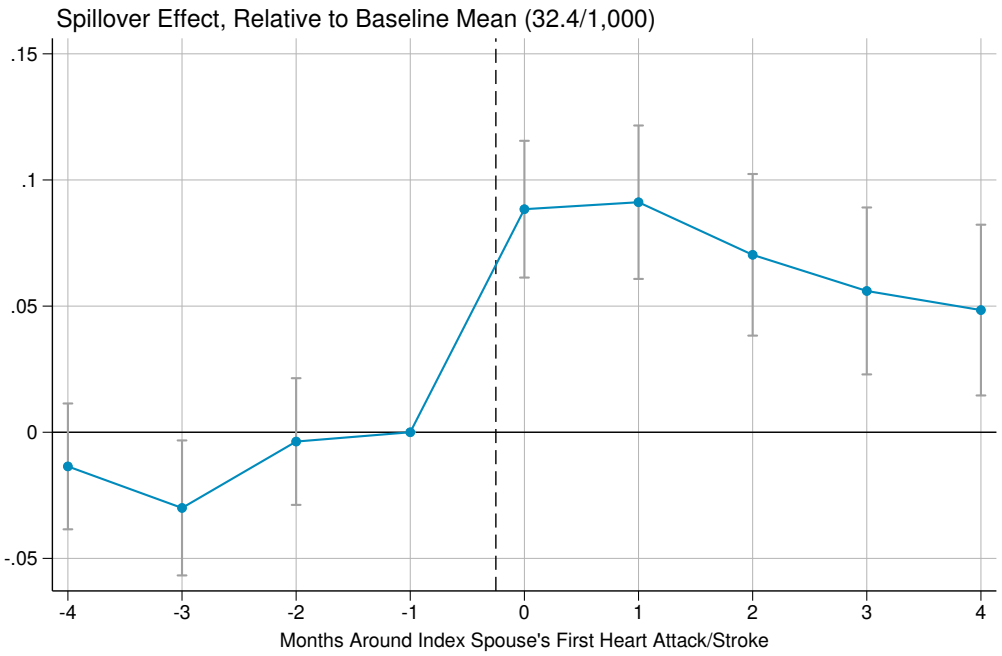
We first assess the presence of a health effect. That is, motivated by prior work (e.g., Arteaga et al. 2024; Bünnings et al. 2021), we test the hypothesis that adverse health shocks to one spouse can lead to a worsening of the other spouse's health and thus translate to a greater need for care.

In Figure 2 we study how an index event affects the focal spouse's probability of being hospitalized. We plots estimates of equation (1) when the outcome is an indicator for visiting a hospital in the given period for any reason. We find clear increases in the probability of a hospitalization immediately following the shock. The increased risk peaks in the immediate period and persists for the horizon of at least 4

months. Column 2 of Table 2 reports the treatment effect in the first month, showing an increased hospitalization risk of 8.9%. Appendix Figure A1 further provides the longer-run analysis over the 12 months horizon.

2.

Figure 2. Effect of a Spouse's Major Cardiovascular Event on Risk of Hospitalization



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

To focus on a more acute health proxy, column 3 reports estimates of a regression in which the outcome of interest is the focal spouse's Emergency Department (ED) encounters. The risk of an ED visit increases by 18.7% in the month of a spouse's health shock, providing further support for the presence of a health effect. Finally, column 4 provides a direct test for a health effect, as proxied by hospitalization, in driving the increased utilization of formal care (i.e., τ_H from the conceptual framework). It provides estimates of the causal effect on the joint event that the focal spouse is hospitalized and visits a SNF. In the month following the spouse's incapacitation, focal spouses are 15.6% more likely to be hospitalized and then consume

formal care at a SNF.

Table 2. Impact of Spouses' Major Cardiovascular Events on Partners' Risk of a Visit

	(1) SNF	(2) Hospitalization	(3) ED Visits	(4) Hospitalization & SNF
Treatment Effect, Month 0	0.191***	0.089***	0.187***	0.156***
	(0.0221)	(0.0139)	(0.0458)	(0.0307)
Fatal Events Only, Month 0	0.404***	0.187**	0.491**	0.215
	(0.1190)	(0.0702)	(0.2051)	(0.1650)
Nonfatal Events Only, Month 0	0.181***	0.085***	0.173***	0.153***
	(0.0225)	(0.0141)	(0.0470)	(0.0312)
Baseline Rate per 1,000	10.1	32.4	31.5	6.8
N	5,828,740	5,823,590	1,049,460	5,824,880

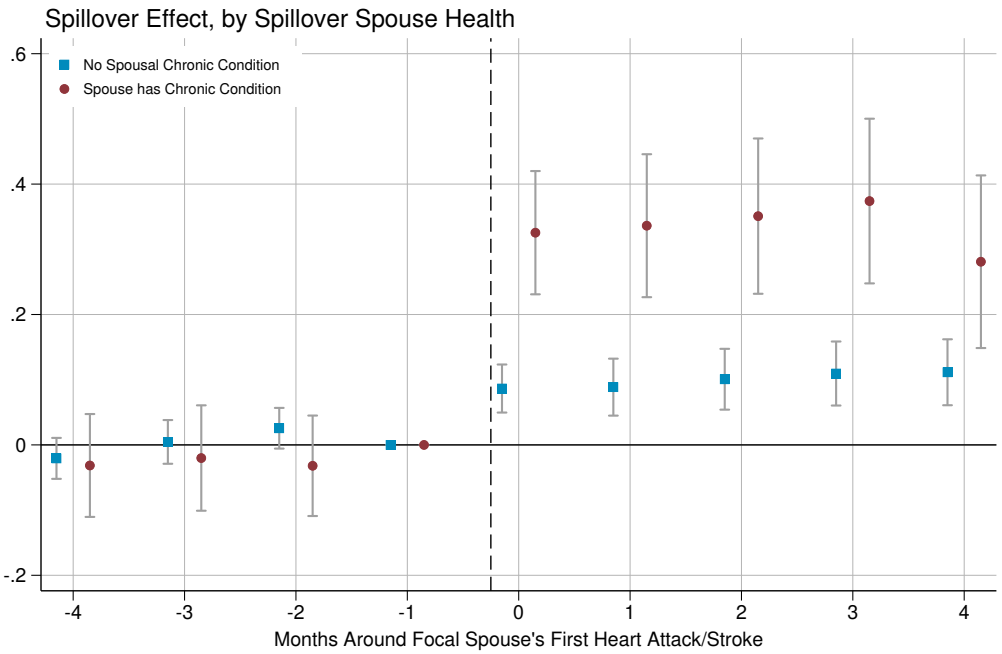
Notes: This table presents pooled difference-in-differences coefficients estimating the effect of a focal spouse's first heart attack or stroke on the spillover spouse's health outcomes (indicated in each column). "Hospitalization & SNF" indicates that a spouse both was hospitalized and visited a SNF in the same month. ED Visits are measured as 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. Treatment effects are estimated in month 0, capturing the effect for the first four weeks post-event. We re-scale all coefficient such that they indicate the change relative to the initial baseline risk of diagnosis in each category, and we cluster standard errors at the household level. Regressions include calendar-time fixed effects and person-specific fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Overall, these findings provide evidence in support of the hypothesis that part of the increased utilization of formal care after a spousal incapacitation is driven by changes in the health and care needs of the focal spouse. Importantly, health effects are more pronounced following fatal index events. Table 2 shows across all adverse health outcome that the treatment effects are larger among focal spouses when the index spouse's health event is fatal versus non-fatal. This is consistent with previous work showing that spousal death is more likely to lead to health effects (Fadlon and Nielsen, 2019a).

4.3 Substitution Effect

We next assess the presence of a substitution effect. Focal spouses with existing needs for care may switch from receiving informal care from their now incapacitated partners to utilizing formal care. We provide two exercises to test for such an effect. The first test relies on stratifying households based on the focal spouse's underlying needs for care prior to the index event. The second test relies on stratifying households based on the degree of incapacitation of the index spouse.

Figure 3. Effect on Risk of a SNF Visit by Focal Spouse's Pre-Existing Chronic Conditions

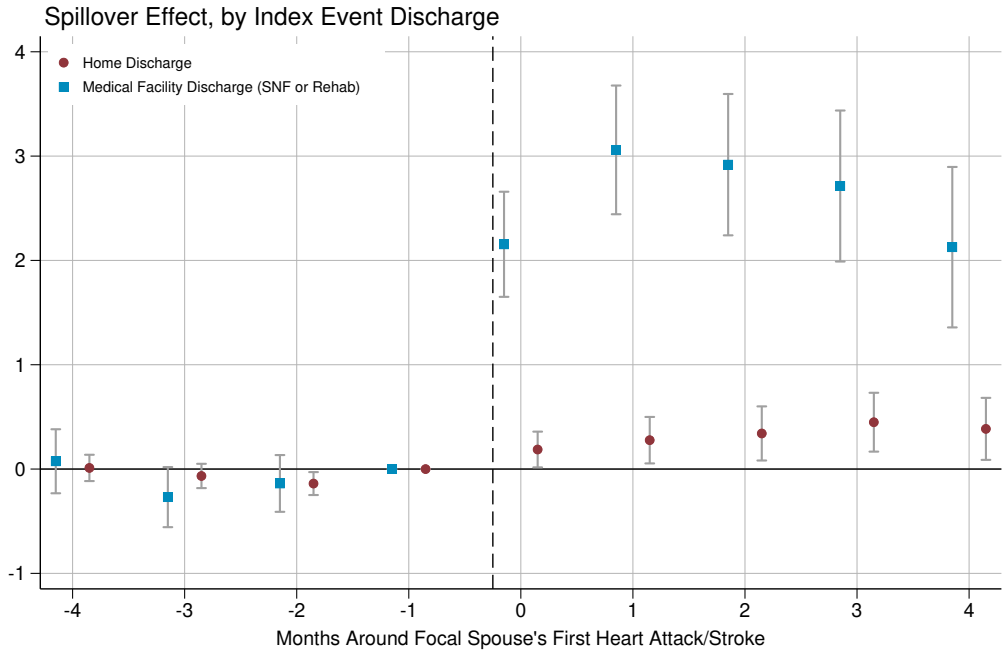


Notes: This figure plots estimates of the δ_n coefficients from equation (1), tracking the months since an index spouse's first heart attack or stroke. The outcome is an indicator for whether the focal (spillover) spouse visited a SNF. We stratify the sample based on whether the focal spouse had a documented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

Figure 3 plots estimates of equation (1) for visits to SNFs, separately by whether the focal spouse suffered from a chronic condition before the index event as a proxy for an underlying need for care. For this measure, we rely on Medicare's Chronic Conditions file, which catalogs for each Medicare beneficiary the presence of one of

27 chronic conditions. The figure clearly shows how the effects diverge: focal spouses with no underlying conditions display an attenuated but significant increase, whereas focal spouses with chronic conditions display a large effects that is roughly 3 times as large.

Figure 4. Effect on Risk of a SNF Visit by Severity of Index Health Event



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

We then stratify by the severity of the index health shock. This aims to capture the degree of incapacitation in terms of the potential to provide informal care. To do so, we divide households based on whether the index spouse was discharged home or to a medical facility (which could be either a rehabilitation facility or a SNF). We focus here on non-fatal health shocks, as we have shown that spousal death is much more likely to induce a health effects (in Table 2).[3] Figure 4 similarly shows a clear

[3]We also limit this sample to households whose focal spouse who did not recently experience a health event: those who have not been hospitalized in weeks -12 to -1. This restriction addresses potential contamination from reverse causality: the severity of an index spouse's health event could be worse due to a recent health event inflicted on the focal spouse.

pattern, where focal spouses are much more likely to utilize formal care following the index spouse's event when the incapacitation is more severe. The increases are up to a threefold increase over the pre-event baseline mean when the index spouse is also discharged to formal care, compared to much smaller increases following a home discharge.

Figures 3 and 4 together provide consistent support for a substitution effect. We observe an increased risk of SNF stays when focal spouses are less healthy and require care ex-ante and when the health event leads to greater incapacitation of the index spouse ex-post.

4.4 Effect Decomposition

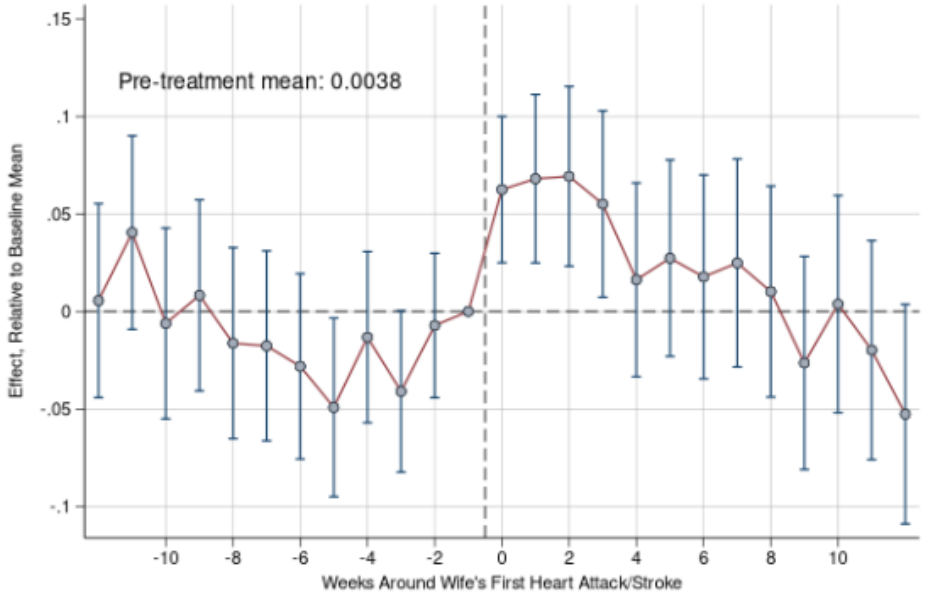
The findings provide evidence in support of both a health effect and a substitution effect in the increased utilization of formal care by focal spouses. In this subsection, we provide a decomposition exercise with the aim to characterize how much of the increased healthcare utilization is driven by substitution into formal care versus the deterioration in health.

We proceed as follows. For all hospitalizations, we predict the probability that an individual with a given set of hospitalization-specific observables will be discharged to a SNF, and we then evaluate how those predicted probabilities evolve around their partners' health shock. Essentially, we generate a "predicted SNF stay" variable equal to zero for anyone who did not have a hospitalization in the period and for those with a hospitalization equal to the probability of a hospitalization with similar observable characteristics (e.g. DRG) ending in discharge to SNF, where that probability comes from a predictive model that we describe below. Intuitively, this exercise captures how we would expect focal spouse SNF use to change after the index spouse's health shock solely from (1) the increase in hospitalizations and (2) how marginal hospitalizations differ from inframarginal hospitalizations in their observable characteristics that predict SNF use. A comparison of changes in this measure of focal spouse *predicted* SNF use around the index spouse's health event versus actual focal spouse SNF use around the event gives an approximation of how much of the actual change in SNF use stems from changes in health (that trigger hospitalizations and subsequent SNF use) versus changes in SNF use, conditional on health.[4]

[4]The decomposition should be accurate for capturing (1) assuming that the only path to a SNF due to a change in health is via a hospital and that hospitalizations have homogeneous probabilities of discharge to SNF. Our ability to capture (2) depends critically on the accuracy of the predictive model.

To estimate our predictive model for SNF stays based on hospitalization observable characteristics, we split the full set of inpatient hospitalizations into a 90% estimation sample and a 10% validation sample. We then predict a binary indicator for any SNF discharge using a Lasso model based on patient age, sex, race, Medicaid eligibility, Medicare enrollment type, and the Diagnostic Resource Groups (DRGs) for the hospitalization (Ellis et al., 2022), as well as relevant interaction terms. We rely on a Lasso regression given the large dimensionality of our covariate space (Abadie and Kasy, 2019). Appendix Figure A2 plots the predicted values on the 10% validation sample of hospitalizations. Those ending in SNF stays have considerably higher predicted values than those ending in another form of discharge.

Figure 5. Change in Focal Spouse's Predicted Probability of SNF Discharge



Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the weeks since a wife's first heart attack or stroke. The outcome of interest is the predicted probability of a SNF discharge based on hospitalization and patient characteristics using LASSO regression; see Section 4.2 and Appendix Figure A2 for more discusion. We re-scale all coefficient such that they indicated the change relative to the pre-treatment average predicted probabilities. The error bars plot 95-percent confidence intervals based on standard errors that have been clustered on each Medicare household. The regression includes calendar-time fixed effects and person-specific fixed effects.

Figure 5 evaluates how our measure of *predicted* SNF stays responds to an index spouse's health shock. In the month following the shock, a focal spouse's predicted probability of a SNF stay increases by roughly 6.5%, up from a baseline risk of 0.38%. We would thus expect 30% ($0.065/0.225 \approx 0.3$) of the marginal SNF stays to have

occurred solely due to the deterioration of the focal spouse's health. The remaining 70% of SNF stays may potentially arise solely from the substitution of formal care for informal care, again with the caveat that this assumes that the only pathway from worsening health to a SNF is via a hospitalization.

4.5 New Diagnoses

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

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

Table 3 reports estimates of equation (2) when the outcomes of interest are new diagnoses of ADRD and cancer. We find that the likelihood of a focal spouse being diagnosed with ADRD increases by nearly 40% immediately following an index spouse's health shock, which remains at the same level on average for six months following the shock. By contrast, we observe no statistically significant effect on cancer diagnoses, despite these being nearly twice as prevalent at baseline. These patterns suggest that focal spouses may be living with latent, undiagnosed conditions such as ADRD, whose detection is triggered by the changes in care following the index event. ADRD symptoms typically involve a process of cognitive decline and behavioral changes that may not be immediately linked to a medical condition. In contrast, cancer typically has more clear-cut physical symptoms and the diagnostic tools (such as biopsies and imaging) provide definitive results quickly.

Table 3. Impact of Spouses' Major Cardiovascular Events on Partners' New Diagnoses

	New ADRD Diagnosis	New Cancer Diagnosis
Month 0	0.399***	-0.054
	(0.1381)	(0.0963)
Months 0–6	0.391***	-0.091
	(0.0756)	(0.0556)
Baseline Rate per 1,000	1.6	3.2
N	2,055,820	1,830,930

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

5 Welfare Analysis

This section examines the welfare implications of our event-study findings. It assesses the fiscal externality, develops methods to recover household's valuation of different modes of care, and discusses the implications for the design of optimal health insurance contracts.

5.1 Government Fiscal Value of a Healthy Spouse

The event-study estimates demonstrate increases in focal spouses' healthcare utilization: SNF admissions, hospitalizations, and ED visits. Accordingly, we study how the effects of the index spouse's health shock translate into overall healthcare expenditures of the focal spouse, that is, the fiscal externality ($R_1 - R_0$).

Table 4 presents estimates of regressions that track the impact of index events on overall Medicare-covered spending for the focal spouse. In order to study spending, we rely on the annual summary measures that are tracked for all fee-for-service Medicare recipients in the data.[5] In the year of an index event, total Medicare spending

[5]The regressions in Table 4 are at the annual level, and so involve only indicator variables for year relative to the index event.

increases by roughly 5% ($317), and by an average of 12% ($773) for fatal index events. The table suggests increases in healthcare expenditure across categories that are large in both absolute and relative terms. For instance, spending on SNFs increases by $112 in the year of the event, roughly a 33% increase over the baseline mean. Those effects are more than doubled when the index spouse passes away following the shock. We observe similarly large increases in the use of home health and hospice care, as well as across various measures of inpatient hospitalizations.[6]

Table 4 summarizes the fiscal value of a healthy spouse for an affected household. To provide a back-of-the-envelope calculation of the aggregate associated fiscal burden among the Medicare population we note the following figures. The current Medicare population (as of 2023) is comprised of 65.7 million. The incidence rate of heart attacks and strokes is roughly 5% annually, 15% of which are fatal within 30 days. Using our estimates from Table 4, these translate into a national spillover effect of $1.26 billion annually.

5.2 Household Valuation of Care

We now turn to develop our two revealed preference approaches to infer measures of v_S, i.e., the average money-metric valuation of informal care. We assess the option value of a healthy spouse via households' "willingness to pay," based on revealed preference methods that leverage households' choices of formal care within and across the two states of the world. Section 3 modeled the transitions across states of nature. In this section we extend the model to analyze the within-state household choices as needed for the moment we want to identify. In subsection 5.2.1, we extend the framework to model the household decision of total SNF days, so we can analyze the value of a marginal SNF day for a given household. In subsection 5.2.2, we extend the framework to include heterogeneous households and model households' binary decision to utilize formal care in each day in the post event horizon, so we can analyze marginal entrants and market demand.

[6]The results suggest these increases may crowd out other, less urgent health expenditures, such as hospital outpatient care (2.3%), elective surgeries (8.0%), and outpatient payments to physicians (3.3%). We also find declines in spending on prescription medication (2.4%), but part of this may be mechanical to the extent that a patient receives medications directly while in a SNF which is not reported in the data.

Table 4. Effect of Spouse's Major Cardiovascular Event on Annual Spending

		Year of Index Event	
Spending Measure	Pre-Treatment Average	(1) All	(2) Fatal Events
Panel A: Total Spending			
	\$6,446	\$317***	\$773***
		(33.02)	(166.40)
Panel B: Long-term Care			
SNF	\$341	\$112***	\$267***
		(7.66)	(42.60)
Home Health	\$291	\$23***	\$61***
		(3.32)	(17.05)
Hospice	\$123	\$55***	\$77***
		(4.98)	(26.74)
Panel C: Hospital & Surgical Care			
Acute Inpatient	\$1,465	\$136***	\$225***
		(13.98)	(69.75)
Other Inpatient	\$178	\$33***	\$55*
		(5.90)	(29.14)
Hospital Outpatient	\$979	-\$17***	\$35
		(6.53)	(29.42)
Ambulatory Surgical Center	\$75	-\$6***	-\$8*
		(0.99)	(4.40)
Panel D: Other Care			
Physician Payments	\$333	-\$10***	\$5*
		(0.66)	(3.16)
Evaluation & Management	\$245	\$30***	\$53***
		(1.98)	(9.84)
Part D Spending	\$907	-\$20***	-\$42**
		(4.14)	(20.41)

Notes: This table plots estimates of pooled post-treatment effects tracking the year following an index event's first heart attack or stroke. Regressions include calendar-time fixed effects and person-specific fixed effects. Column (1) reports the estimate for all events, while column (2) reports the estimate only for fatal index events. Estimates are scaled by the time remaining in the year following the index event, and standard errors are clustered at the household level.

5.2.1 Marginal Rate of Substitution

The first method evaluates how the marginal valuation of SNF days relative to $1 varies across states of the world. It relies on within-state first order conditions of SNF days and the fact the price schedule is invariant to the index spouse's incapacitation.

We index the focal spouse's utility by the state of the index spouse's incapacitation. Incapacitation is modeled by t in the conceptual framework, where $t = 0$ corresponds to periods prior to the index event and $t = 1$ corresponds to periods after the index event. We let the focal spouse's utility, $U_t(x_t, m_t)$, take the following form:

$$U_t(x_t, m_t) = u_t(x_t) + m_t,$$

where x_t is the consumption of formal care and m_t is "money" so that the utility from care, $u_t(x_t)$, is relative to the within-state valuation of resources. A similar formulation can be obtained instead by writing the Lagrangian of the problem and normalizing it by the shadow value of the household's resources. In both cases we achieve a money-metric valuation of the utility from care. These valuations map to the household value functions that were presented in Section 3, where here we provide an explicit formulation for them.

Within-state optimality conditions imply that

$$u_t'(x_t) = P(x_t)$$

where $P(x_t)$ is the marginal price of a SNF day (capturing the non-linearity in pricing of SNF days). The optimality conditions within each state imply the the following holds across states:

$$u_0'(x_0) = u_1'(x_1).$$

That is, the marginal state-dependent utility from a SNF day is equated across states of nature, as the marginal price is similar across states. A first-order expansion (of the left-hand side) leads to the following welfare formula:

$$1 + \gamma(x_1)\frac{x_1 - x_0}{x_1} = \sigma(x_1),$$

where $\gamma(x_1) \equiv -\frac{u_0''(x_1)}{u_0'(x_1)}x_1$ represents relative risk aversion and can be calibrated from previous work to roughly vary from 1 to 4 (Schechter, 2007). The parameter $\sigma(x_1) \equiv \frac{u_1'(x_1)}{u_0'(x_1)}$ represents the state-dependence valuation parameter we want to assess. It

captures how the marginal (money-metric) valuation of a SNF day changes across states of nature. As such, it represents the ratio of the willingness to pay for care across the two states. In other words, it is the marginal rate of substitution of care across states of spousal ability to provide informal care.

This approach evaluates welfare based on changes in the curvature of the indirect utility from care. When an input into the household production function is missing in the constrained optimization, it would lead to a higher curvature of the indirect utility from any other input since there are fewer choice variables to smooth across marginal decisions. It has a similar effect as in consumption commitments, where there is a margin you cannot freely adjust on (Chetty and Szeidl, 2016). This approach relies on our quasi-experimental source of variation alone, but it is constrained to marginal assessments and requires calibration of risk aversion.

The term $\sigma(x_1) - 1$ represents the household valuation of a healthy spouse via the marginal value of formal care: it is the excess valuation of the marginal SNF day in the transition across the spousal incapacitation status. Based on our findings, we have that $\frac{x_S - x_H}{x_S} = 0.191$, and we provide values for a range of relative risk aversion parameters that vary from 1 to 4 in Appendix Table A2. We assess that, on average across the calibration, an additional day of SNF is valued (relative to income) by 48% more following the spousal incapacitation event.

5.2.2 Consumer Surplus

The second method evaluates the overall willingness to pay for formal care across the two states of the world. It evaluates the option value of spousal informal care by the excess valuation of formal care when this option vanishes. We explicitly model willingness to pay for formal care compared to the alternatives, where it is the set of alternatives that changes across states. We leverage price variation from the non-linear schedule of the cost of SNF days. We use the responses to the price variation to back out demand curves in each state and impute the excess valuation of formal care in a given day due to the spousal event. We then calculate the change in consumer surplus for each SNF day and aggregate over days to get the total option value of spousal informal care. Since we aggregate over households, we need to explicitly add heterogeneity in valuations into the model.

We let beneficiaries be indexed by $s_t \in [0, 1]$, where s_t orders consumers according to their willingness-to-pay for formal care (when $c_t = 1$) relative to informal care (when $c_t = 0$) if they have an underlying need (that is, $s_t = 1$). The subscript t again

characterizes the index spouse's incapacitation state, where $t = 0$ corresponds to periods prior to the index event and $t = 1$ corresponds to periods after the index event. We order consumers such that $s_t = 0$ is the consumer with highest willingness-to-pay and $s_t = 1$ is the consumer with lowest willingness-to-pay, which will simplify the formulas.[7] We denote the willingness-to-pay for formal care (relative to the available alternative) by $u_t(s_t)$ for an s_t-type consumer. $P(x_t)$ again denotes be the daily out-of-pocket price charged for formal care, which is similar across states of nature but can vary for based on the length of stay (x_t). In what follows, we drop the notation of x_t for notational convenience. Finally, let $s_t^*(P)$ represent the s_t-type consumer who is on the margin between purchasing and not purchasing formal care at price P, characterized by satisfying the condition $u_j(s_t^*(P)) = P$.

With this setup, we can write the consumer surplus for a given day as:

$$CS_t(P) = \int_0^{s_t^*(P)} \left[u_t(s_t) - P\right] f_t(s_t)\, ds_t$$

We note that since the probability density function is not conditional on takeup, it captures the option value of the alternative mode of care. Calculation of $CS_t(P)$ relies on the known fact that it is the integral of the inverse demand function since:

$$\frac{dCS_t}{dP} = \frac{ds_t^*(P)}{dP}\left[u_t(s_t^*(P)) - P\right] f_t(s_t^*(P)) - \int_0^{s_t^*(P)} f_t(s_t)\, ds_t = -F_t(s_t^*(P)).$$

We then infer the mean valuation of spousal informal care by

$$v_S = CS_1(0) - CS_0(0).$$

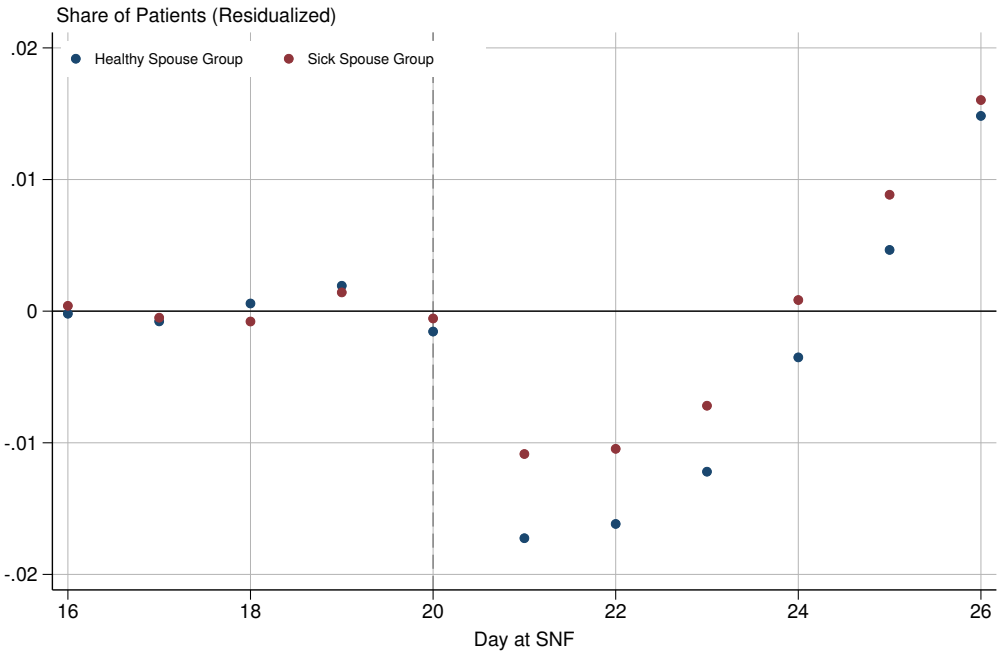
That it, we use the excess valuation of formal care in the absence of the option of spousal care to infer the option value of spousal care.

We recover the inverse demand curves using local price variation and a linear approximation. Specifically, we leverage the fact that Medicare coverage of SNF stays is limited to 100 days and patients face a discontinuous change in their coinsurance from 0% to 20% after 20 days. We study responses to these discontinuous changes in a regression discontinuity (RD) framework, and we recover the difference in the RD parameter across the index spouse's health shock in a difference-in-discontinuities design (Grembi et al., 2016).

[7] This allows us to work with the CDF instead of one minus the CDF in imputing demand.

In Figure 6, we illustrate the variation we exploit. First, we predict focal spouse's SNF stay decisions (measured in days), adjusting for day-of-week effects and assuming a linear trend based on decisions made before the 20-day threshold. We then calculate deviations in observed stay decisions relative to this trend around the discontinuous change in cost-sharing at day 20. We find visually clear evidence that focal spouses with a healthy index spouse appear to be meaningfully more price sensitive than those whose index spouse has just experienced a major health event.[8]

Figure 6. Responses to Cost-Sharing for SNF Stays By Health Event



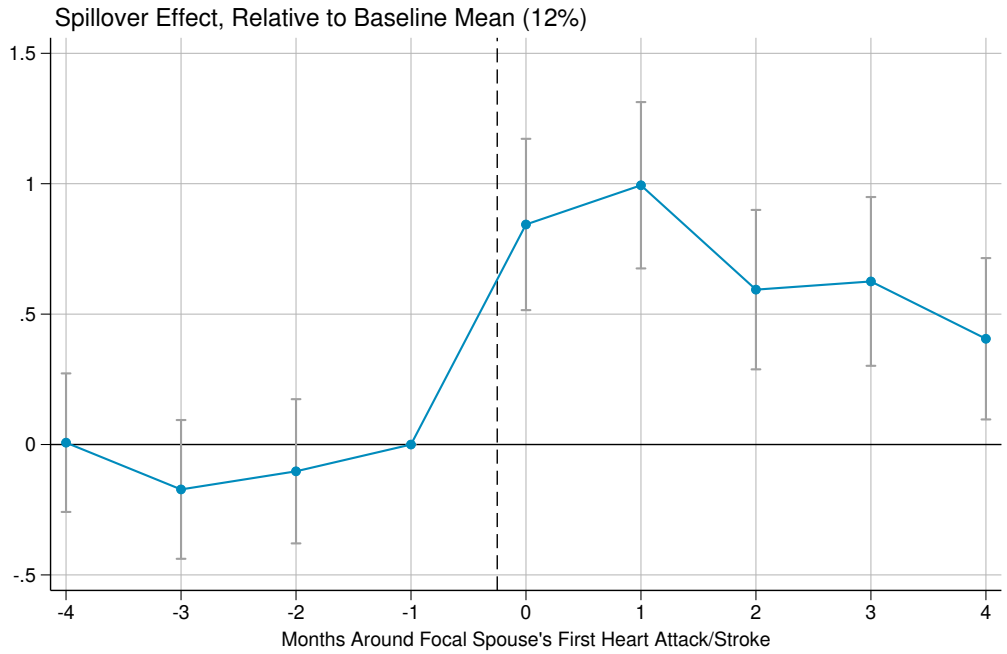
Notes: This figure plots binscatter coefficients for the average share of patients in a SNF on a given day. We plot residual means in this figure after adjusting for the day of the week and a measure of days covered by Medicare previously. Averages for the group with healthy index spouses (no health events) are shown in blue, while those whose index spouse has experienced an adverse health shock are shown in red. The vertical line indicates the day after which 20% coinsurance is applied to each additional day the focal spouse remains in the SNF. The selected bandwidth in the figure is chosen to match the optimal bandwidth used in estimating equation 4. See Appendix Figure A3 for a non-residualized histogram of SNF length of stay across groups.

In support of these patterns, we look at the likelihood that a focal spouse stays in a SNF for longer than Medicare fully covers. In Figure 7 we observe a doubling in the likelihood that focal spouses stayed in a SNF beyond 20 days in reponse to

[8]This bunching can also be seen directly when plotting the distribution of LOS conditional on SNF stays, as shown in Appendix Figure A3.

the spousal health event, implying that nearly one-quarter of focal spouse SNF stays following a index spouse's event required some patient cost-sharing. We also observe significant increases in the average length of stay for a SNF visit, consistent with our prior findings (Appendix Figure A4).

Figure 7. Effect of a Spouse's Major Cardiovascular Event on the Fraction of SNF Stays Not Fully Covered by Medicare



Notes: This figure plots estimates of the δ_n coefficients from equation (1), difference-in-differences estimates that track the months since an index 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 focal spouses without SNF stays in both treatment and control groups are assigned values of 0 when no SNF admission occurred in a month. We re-scale all coefficients such that they indicate the change relative to the initial unconditional mean values. Error bars plot 95-percent confidence intervals based on standard errors clustered on each Medicare household. The regression includes calendar-time fixed effects and person-specific fixed effects.

We estimate the difference in responses from Figure 6 using an RD design around the 20-day cutoff with the following specification:

$$y_{i,t} = \delta_0 + \delta_1 D^*_{i,t} + S_{i,t}(\gamma_0 + \gamma_1 D^*_{i,t}) + T_{i,t}\left[\alpha_0 + \alpha_1 D^*_{i,t} + S_{i,t}(\beta_0^{\text{D-Disc}} + \beta_1 D^*_{i,t})\right] + \varepsilon_{i,t}, \quad (4)$$

where $y_{i,t}$ is an indicator for whether focal spouse i is in a SNF on day t, $D^*_{i,t}$ is the re-centered running variable of SNF days relative to the 21-day cutoff, $S_{i,t}$ in-

dicates a post-cutoff dummy, and $T_{i,t}$ indicates a post-treatment dummy of a index spouse's health shock. Equation (4) is estimated only for SNF stays of duration within a certain bandwidth of the cutoff, $D^* \in [-h, h]$, and uses a triangular kernel to weight observations more heavily closer to the cutoff. The parameter β_0 indicates the difference in responses across the experimental groups.

We estimate equation (4) using observations within four months of both the treated and placebo index events and adjust for Medicare couple and time-of-event fixed effects. Table 5 presents the results. Prior to the index spouse's event, the discontinuous change in Medicare payment for SNF stays induces a significant negative effect on the length of a focal spouse's stay, decreasing the probability that they will remain in the SNF on the first day they will face copayments by 1.4 percentage points. In contrast, the focal spouse's response to copayments is reduced to 0.9 percentage points when they have been admitted following the index spouse's event. Comparing the differences in the estimated responses, we find that households affected by a shock are 0.5 percentage points more likely to remain in care across this threshold, absorbing roughly 36% of the pre-treatment price responses.

Table 5. Effects of Spouse Major Cardiovascular Event on SNF Price Sensitivity

	(1)	(2)
τ_{pre}	-0.014***	-0.014***
	(0.0031)	(0.0017)
τ_{post}	-0.009**	-0.009**
	(0.0043)	(0.0024)
$\beta_{\text{D-Disc}}$	0.005	0.005*
	(0.0053)	(0.0031)
Couple FEs		✓
Time FEs		✓
Bandwidth	11.22	11.22
N	783,840	783,840

Notes: This table presents regression-discontinuity and differences-in-discontinuities estimators identifying the effect of ending Medicare coverage for SNF stays, which ends on day 21 for qualifying stays. The first two rows present the estimated effect of losing coverage on the probability a focal spouse will remain in the SNF, stratified by before or after the index event. The third row presents the difference-in-discotinuities estimator as discussed in the text. SNF stays within 4 months of the treatment and placebo events are included in the regression.

Figure 8 illustrates the intuition of the approach. In our data, we can observe level differences in the demand for SNF care among focal spouses with spouses either affected or unaffected by health events. By identifying the fraction of focal spouses in a SNF at day 20 (the day before the price change) in each group, we identify the x-intercepts for each demand curve (points $(c, \$0)$ and $(a, \$0)$). The slope of the demand curve is then identified based on the estimated percentage change in demand for formal care between days 20 and 21 when cost-sharing for formal care increases discontinuously. The RD estimates allow us to identify the values b and d. Demand curves are then extrapolated using linear approximation based on these two points for each group.

From these estimated inverse demand curves, we estimate consumer surplus for each additional day of SNF coverage. Through day 20, we estimate that $CS_0(0) = \$32,558$ for the healthy spouse group, and $CS_1(0) = \$116,247$ for the sick spouse group. That is, focal spouses facing a reduction in the availability of informal care derive roughly 3.6 times the consumer surplus from receiving formal care at a given price. Appendix Figure A5 illustrates how these differences in consumer surplus accumulate over days in a SNF. By the end of day 100—the day at which consumers face full cost-sharing for a SNF stay—focal spouses with an incapacitated spouse have a total consumer surplus of $251,219, roughly 3.8 times larger than the estimated consumer surplus of the group with healthy spouses.

5.3 Design of Optimal Insurance Policies

Optimal health insurance involves a trade-off between the welfare benefits of risk protection and the deadweight loss from moral hazard (Zeckhauser, 1970). Increasing insurance coverage induces welfare gains from equalizing marginal utility across states of the world, but comes at the cost of the over-consumption of care.

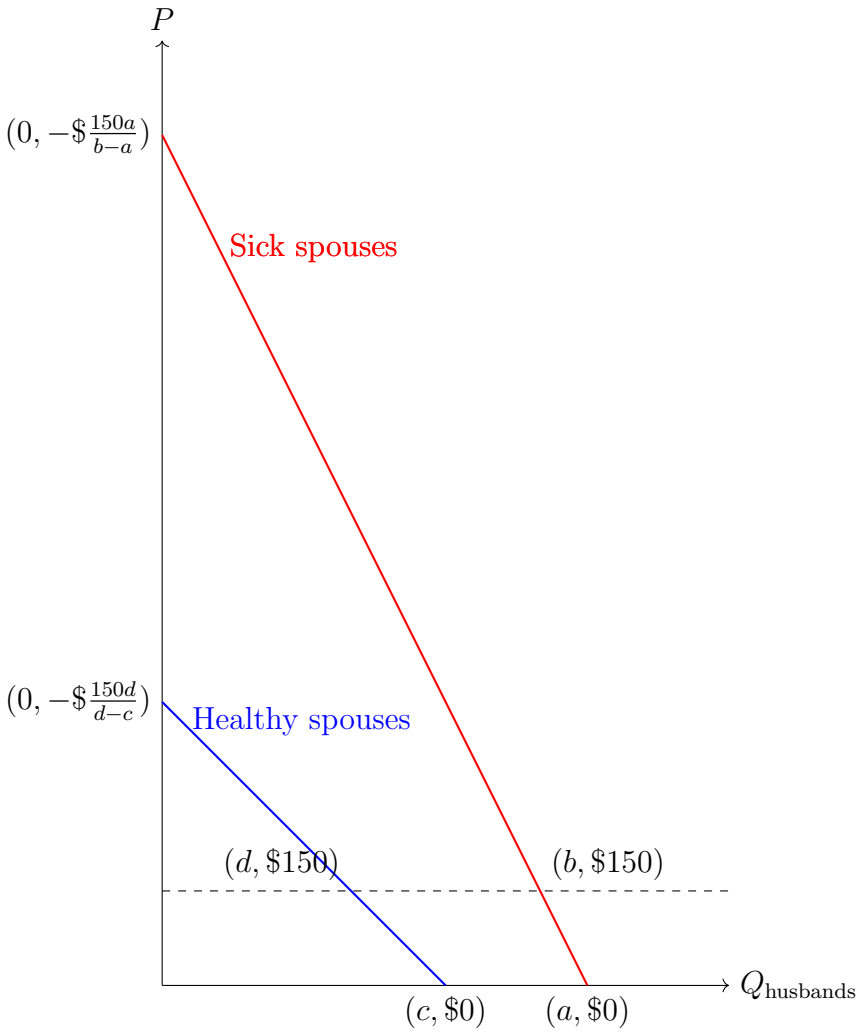
Total social welfare with respect to Medicare's formal care now takes the following explicit form:

$$W_t(P) = \int_0^{s_t^*(P)} \left[u_t(s_t) - P\right] f_t(s_t)\, ds_t + P \int_0^{s_t^*(P)} f_t(s_t)\, ds_t.$$

Accordingly, the welfare cost of a price change, the deadweight loss (DWL), is captured by

$$\frac{dW_t(P)}{dP} = P \frac{dF_t(s_t^*(P))}{dP}.$$

Figure 8. Estimating Demand for Formal Care across Groups



Notes: This figure shows the intuition for estimating linear inverse demand curves using variation in prices that husbands face for formal long-term care. Demand curve estimation relies on identifying differences in levels of demand (values a and c) and the relative change in this demand as a function of the RD estimates presented in Section 5.2. These changes identify b and d.

Importantly, this implies that the ratio of elasticities governs the relative DWL from subsidies to formal care across households with and without healthy spouses.

With the demand elasticities identified above, the estimated ratio is 36%. This implies that efficiency losses are reduced by 36% when subsidizing formal care for individuals without readily available informal care provided by a spouse. It captures that idea that people whose partners just experienced a shock have less access to informal care, and so exhibit a less elastic demand for formal care. As a result, generous coverage leads to less moral hazard among them.

This estimated difference in efficiency loss effectively provides a strong motivation to differential pricing for elderly care based on household structure. In particular, policymakers may consider family deductibles—which are more likely to bind if a spouse becomes incapacitated—as a means to charge different out-of-pocket costs to consumers with and without the option of informal care provided by a spouse or a family member.

6 Conclusion

We use Medicare claims data to study the elderly-care valuation of a healthy spouse among the Medicare population. We find strong evidence of a protective role of a healthy spouse, where negative health shocks adversely affecting spouses' health and their utilization of formal care. The large estimated effects of spousal incapacitation impose significant fiscal externalities, which we assess to be on order of $1.26 billion per year. More importantly, we estimate a high option value of informal care: beneficiaries are willing to pay roughly 4 times more for their own formal care when their spouses become incapacitated.

Our results illustrate the importance of government provision of formal elderly care when spouses are not able to provide informal care to tend their needs. One policy response to this dynamics would be a family-specific deductible. Such a deductible would effectively provide formal care at lower prices to households with an incapacitated spouse. As a result, the family-specific deductible would lower copayments for those whose demand curves for formal care have become more inelastic, and thus whose consumption is less prone to moral hazard inefficiencies. In fact, some countries do indeed adjust support for formal care based on the availability of family members who may provide informal care. The Dutch system, as one example, explicitly limits eligibility for formal care based on whether an able spouse is present

(Ilinca et al., 2017). Moreover, there are various organizations in the US that embed this type of pricing, such as the University of California system whose PPO plan (UC Care) is based on family deductibles and out-of-pocket maximums.

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

References

Abadie, Alberto and Maximilian Kasy, "Choosing among regularized estimators in empirical economics: The risk of machine learning," *Review of Economics and Statistics*, 2019, *101* (5), 743–762.

Arber, Sara and Jay Ginn, "Gender differences in informal caring," *Health & Social Care in the Community*, 1995, *3* (1), 19–31.

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

Baily, Martin Neil, "Some aspects of optimal unemployment insurance," *Journal of public Economics*, 1978, *10* (3), 379–402.

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.

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.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Difference-in-differences estimators of intertemporal treatment effects," *Review of Economics and Statistics*, 2024, pp. 1–45.

Chandra, Amitabh, Courtney Coile, and Corina Mommaerts, "What can economics say about Alzheimer's Disease?," *Journal of Economic Literature*, 2023, *61* (2), 428–470.

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.

Charles, Kerwin Kofi and Purvi Sevak, "Can family caregiving substitute for nursing home care?," *Journal of health economics*, 2005, *24* (6), 1174–1190.

Chetty, Raj, "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90* (10-11), 1879–1901.

_ **and Adam Szeidl**, "Consumption commitments and habit formation," *Econometrica*, 2016, *84* (2), 855–890.

_ **and Amy Finkelstein**, "Social insurance: Connecting theory to data," in "Handbook of public economics," Vol. 5, Elsevier, 2013, pp. 111–193.

Dall, Timothy M, Paul D Gallo, Ritasree Chakrabarti, Terry West, April P Semilla, and Michael V Storm, "An aging population and growing disease burden will require alarge and specialized health care workforce by 2025," *Health affairs*, 2013, *32* (11), 2013–2020.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo, "The economic consequences of hospital admissions," *American Economic Review*, 2018, *108* (2), 308–352.

Dube, Arindrajit, Daniele Girardi, Oscar Jorda, and Alan M Taylor, "A local projections approach to difference-in-differences event studies," Technical Report, National Bureau of Economic Research 2023.

Ellis, Randall P, HE Hsu, JJ Siracuse, AJ Walkey, K Lasser, B Jacobson, C Andriola, A Hoagland, Y Liu, C Song et al., "Diagnostic items: A new framework for disease surveillance, prediction, and risk adjustment," in "JAMA Health Forum" 2022.

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.

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.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano, "Do fiscal rules matter?," *American Economic Journal: Applied Economics*, 2016, pp. 1–30.

Gross, Tal, Tim Layton, Daniel Prinz, and Julia Yates, "Social Defaults and Plan Choice: The Case of Spousal Following," 2024.

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," May 2024. Working Paper.

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

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.

Johnson, Richard W, Melissa M Favreault, Judith Dey, William Marton, and Lauren Anderson, "Most Older Adults Are Likely to Need and Use Long-Term Services and Supports Issue Brief," *US Department of Health and Human Services: Office of the Assistant Secretary for Evaluation and Planning*, 2021.

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.

Mommaerts, Corina, "Long-Term Care Insurance and the Family," 2024. Unpublished.

Pillemer, Sarah, Jennifer Davis, and Geoffrey Tremont, "Gender effects on components of burden and depression among dementia caregivers," *Aging & mental health*, 2018, *22* (9), 1162–1167.

Roofeh, Regina, Dylan M Smith, and Sean AP Clouston, "Estimated prevalence of elder orphans using National Health and Aging Trends Study," *Journal of Aging and Health*, 2020, *32* (10), 1443–1449.

Schechter, Laura, "Risk aversion and expected-utility theory: A calibration exercise," *Journal of Risk and Uncertainty*, 2007, *35* (1), 67–76.

Tramonti, Francesco, Paolo Bongioanni, Rebecca Leotta, Irene Puppi, and Bruno Rossi, "Age, gender, kinship and caregiver burden in amyotrophic lateral sclerosis," *Psychology, health & medicine*, 2015, *20* (1), 41–46.

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 Tables and Figures

Appendix Table A1. Matrix of Transitions across States of Nature

	Healthy in $t=1$	Sick in $t=1$	
		Informal care	Formal care
	$(h_1=0)$	$(h_1=1, c_1=0)$	$(h_1=1, c_1=1)$
$(h_0=0)$	$p_{0,0}$	0	$p_{0,(1,1)}$
$(h_0=1, c_0=0)$	0	0	1
$(h_0=1, c_0=1)$	0	0	1

Notes: Table shows calibrated values for the state dependence of the marginal valuation of SNF days relative to \$1 (σ, as described in Section 5.2). Based on our data, we have that $\frac{x_S - x_H}{x_S} = 0.43$. We then calculate σ for a range of relative risk aversion parameters shown above. That is, each column of the table shows the marginal valuation of an additional day of SNF care relative to income, for a different relative risk aversion parameter γ.

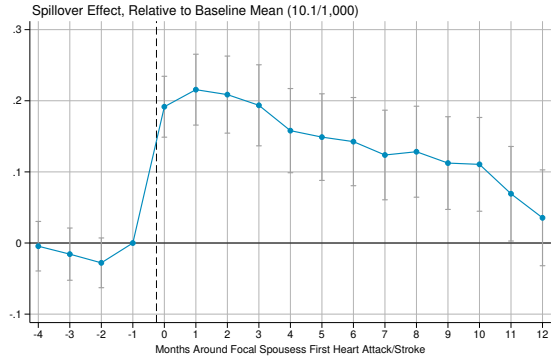
Appendix Table A2. Model Calibration

γ	1	2	3	4	
$\sigma - 1$	19.1%	38.2%	57.3%	76.4%	Mean: 47.75%

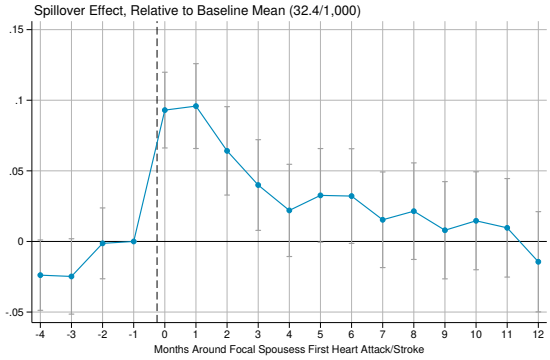
Notes: Table shows calibrated values for the state dependence in the marginal valuation of SNF days relative to \$1 ($\sigma(x_1) - 1$, as described in Section 5.2). Based on our data, we have that $\frac{x_1 - x_0}{x_1} = 0.191$. We then calculate $\sigma(x_1) - 1$ for a range of relative risk aversion parameters shown above. Each column of the table shows the excess marginal valuation of an additional day of SNF care relative to income for a different relative risk aversion parameter γ and the average across the range.

Appendix Figure A1. Long-run Spillover Effects After Initial Heart Attack/Stroke

(a) Visits to Skilled Nursing Facilities

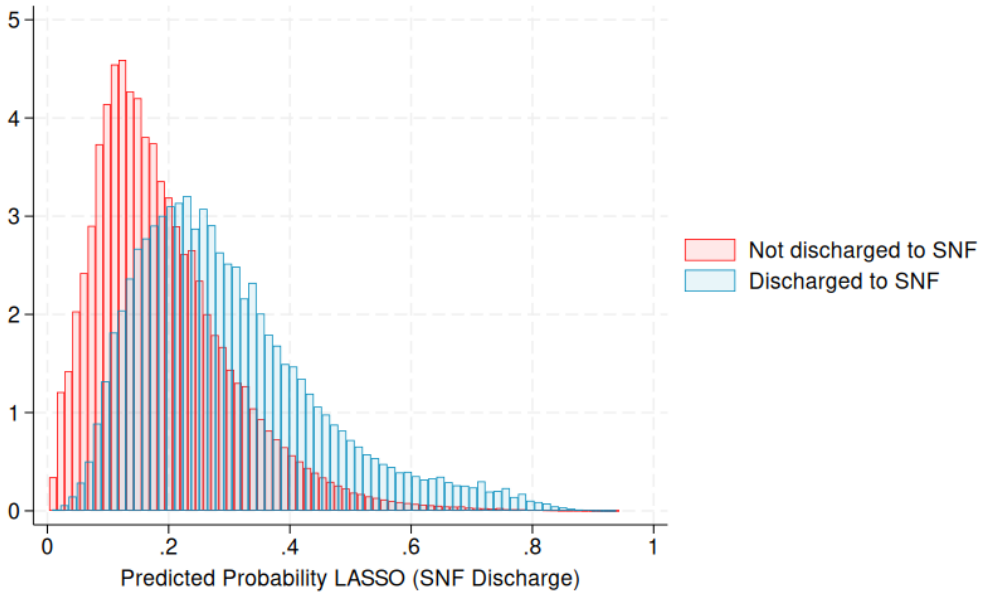


(b) Inpatient Hospitalizations



This figure plots dynamic difference-in-differences coefficients that track the months since a focal spouse's first heart attack or stroke. The outcome of interest is an indicator function equal to one if the spillover spouse was hospitalized or visited a SNF. Compare with Figures 1 and 2.

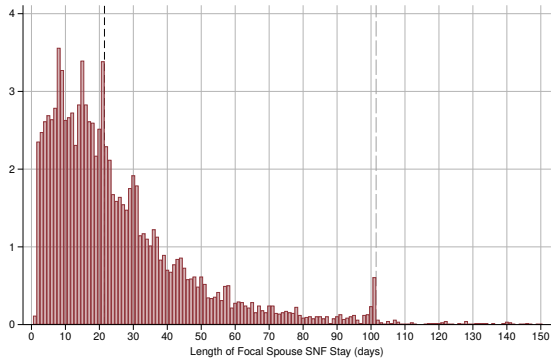
Appendix Figure A2. Predicted Probability of Discharge to SNF: Validation on 40% Sample



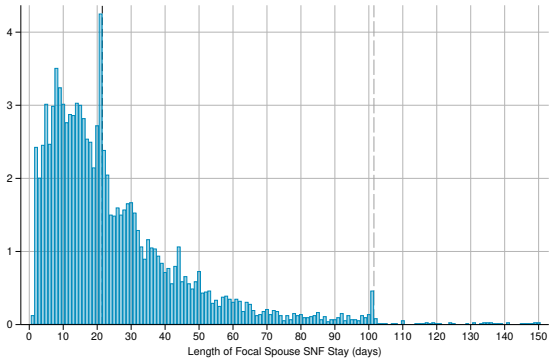
Notes: This figure plots the distribution of the predicted probability of a given hospitalization ending in a discharge to a SNF across a 40% validation sample of husband hospitalizations in the Medicare data, for two groups: those that did not end with a discharge to a SNF (in red), and those which did (in blue). Predicted values are modeled using LASSO prediction, with a binary predicted outcome for all hospital stays in the 100% Inpatient claims files from 2011 to 2017. Covariates include patient age, sex, race, dual Medicaid eligibility, Medicare enrollment type, and DRG indicators.

Appendix Figure A3. Focal Spouses' SNF LOS Following a Health Shock

(a) Sick Spouse Group

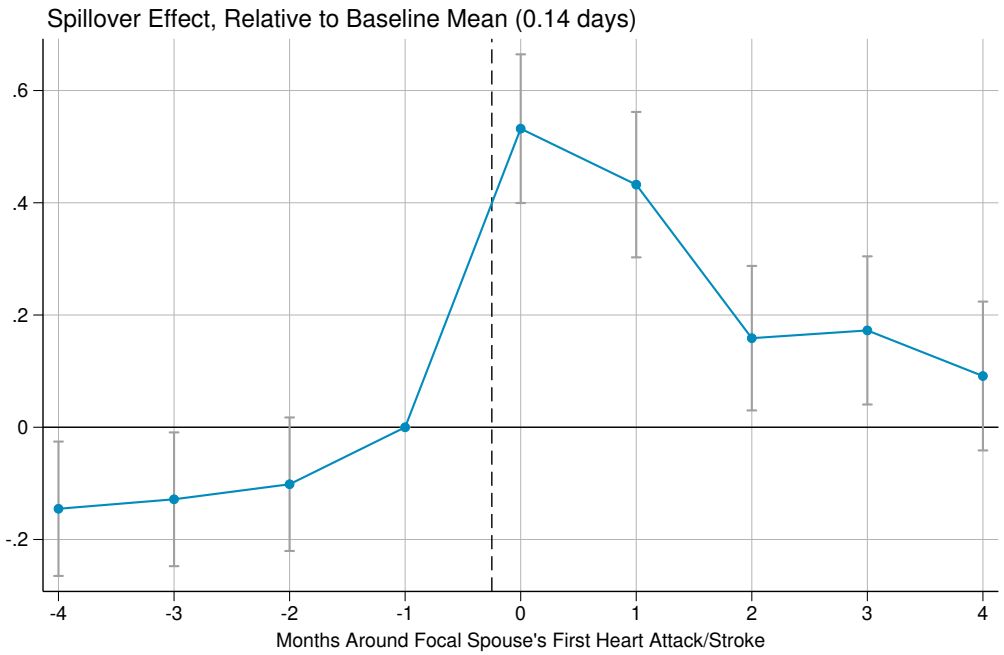


(b) Healthy Spouse Group



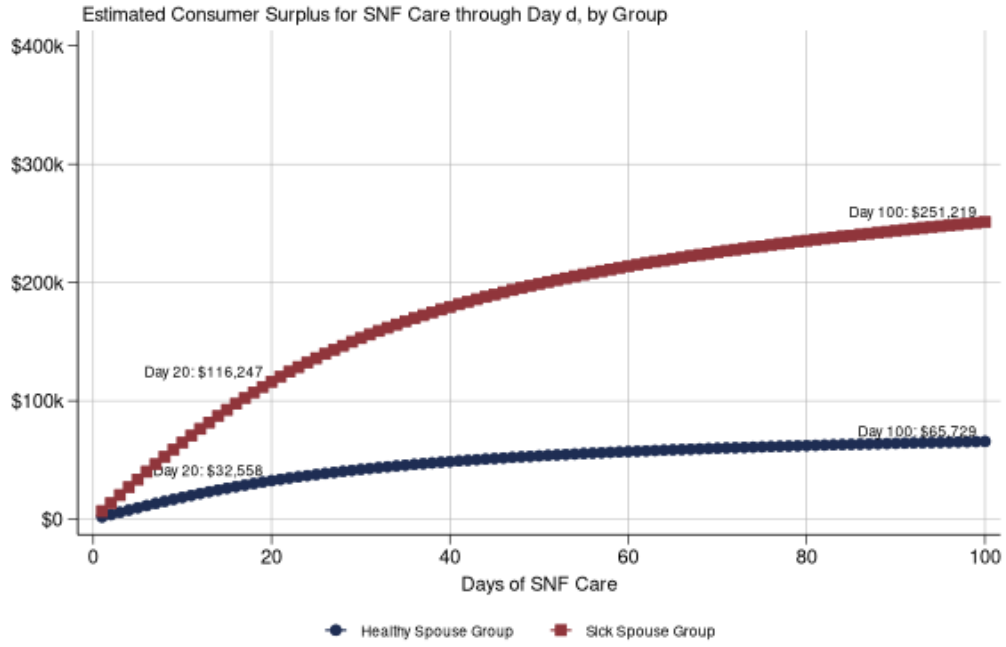
Notes: This figure plots histograms showing the distribution of length of stay (LOS) for focal spouses visiting SNFs around the time of their index spouse's initial heart attack or stroke. Panel A provides the distribution among beneficiary observations from the first 12 weeks following an actual spousal event (from the "sick spouse" group), and panel B provides the distribution among beneficiary observations from the first 12 weeks following a placebo spousal event (from the "healthy spouse" group). Vertical lines indicate the days at which Medicare payment rates for SNF stays change discontinuously.

Appendix Figure A4. Effect on Unconditional Average LOS for a SNF Stay



Notes: This figure plots estimates of the δ_n coefficients from equation (1), tracking the weeks since a index 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 0-day months. We re-scale all coefficients such that they indicate the change relative to the initial baseline average LOS. The error bars plot 95-percent confidence intervals based on standard errors clustered at the household level. Regression include calendar-time fixed effects and person-specific fixed effects.

Appendix Figure A5. Cumulative Consumer Surplus by Index Health Events

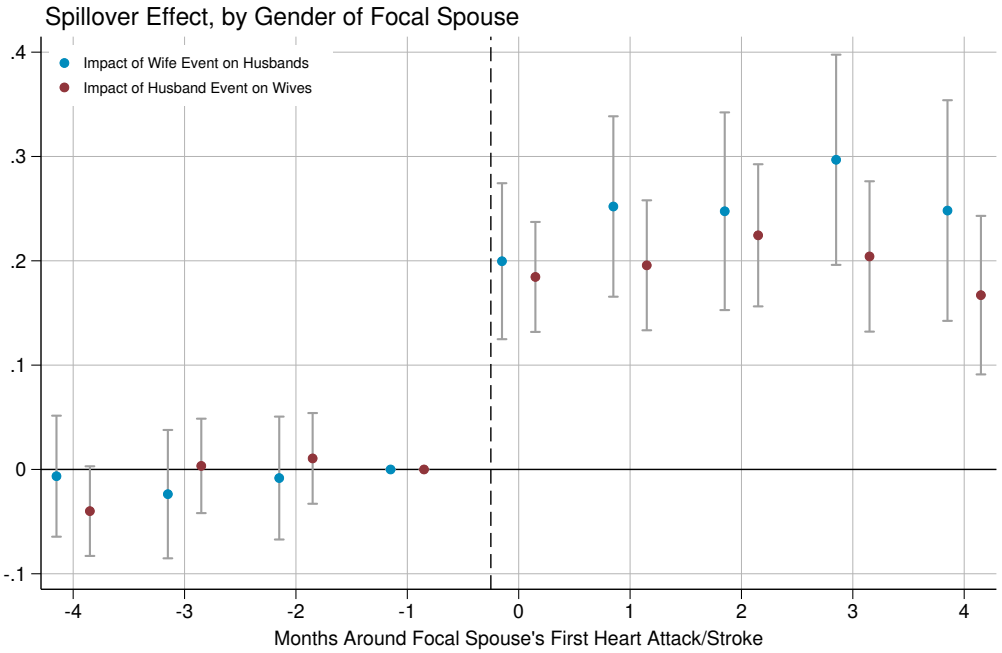


Notes: This figure plots estimated cumulative consumer surplus following the methodology of Section 5.2.2 and the framework presented in Section 3. The cumulative consumer surplus from a SNF stay for focal spouses with a healthy partner; the same values for those with a sick index spouse (for whom demand curves have shifted) are plotted in red.

B Appendix: Symmetry by Gender

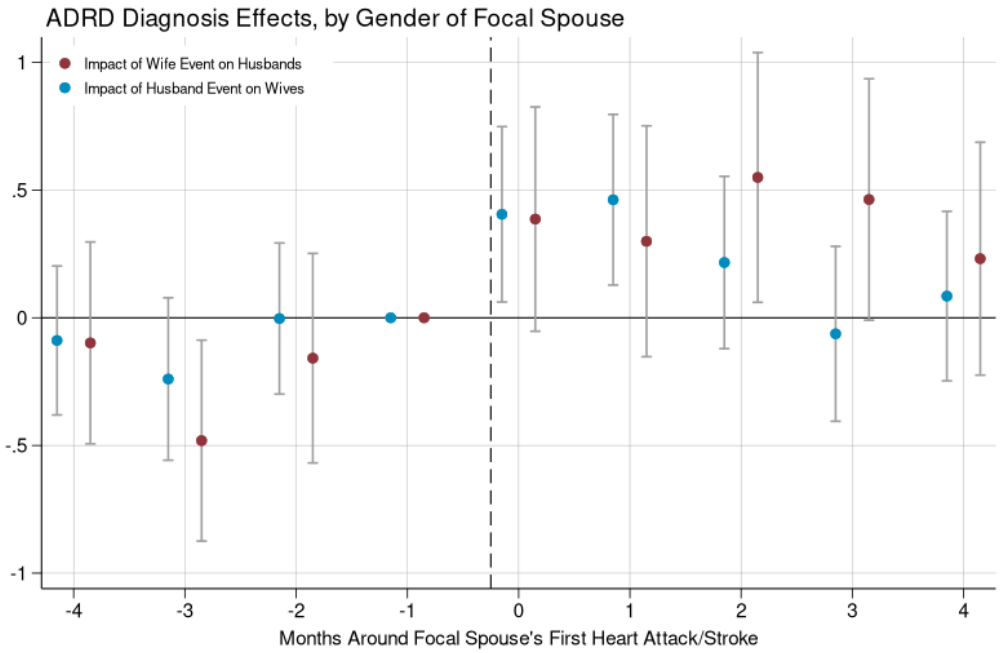
This appendix stratifies the paper's main results by gender. A priori, one might expect the main estimates to differ based on whether the index spouse is male or female. Previous studies have found differences in the burden of caretaking by gender (Arber and Ginn, 1995; Tramonti et al., 2015; Pillemer et al., 2018).

Appendix Figure B1. Spillover Effects by Gender: Visits to Skilled Nursing Facilities



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 event occurred to a husband or wife. The outcome of interest is an indicator for whether the other partner was hospitalized or visited a SNF, and we re-scale all coefficients such that they indicated the change relative to the initial baseline risk of hospitalization or SNF stay. The joint tests of different post-treatment effects across focal spouse sex had p-values of 0.23 for SNF stays and 0.13 for hospitalizations. The error bars plot 95-percent confidence intervals based on standard errors clustered at the household lever. Regressions include calendar-time fixed effects and person-specific fixed effects.

Appendix Figure B2. Spillover Effects by Gender: New ADRD Diagnoses



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 event occurred to a husband or wife. The outcome of interest is an indicator for whether the other partner obtained a new ADRD diagnosis, and we re-scale all coefficients such that they indicated the change relative to the initial baseline risk of hospitalization or SNF stay. The error bars plot 95-percent confidence intervals based on standard errors clustered at the household lever. Regressions include calendar-time fixed effects and person-specific fixed effects.

Appendix Table B1. Effects of Index Spouse Health Event on Focal Spouse's Price Sensitivity for SNF Stay Decision

	(1) All	(2) Focal Husbands	(3) Focal Wives
τ_{pre}	-0.014***	-0.012***	-0.015***
	(0.0017)	(0.0027)	(0.0023)
τ_{post}	-0.009***	-0.009**	-0.008***
	(0.0024)	(0.0037)	(0.0030)
$\beta_{\text{D-Disc}}$	0.005*	0.003	0.007*
	(0.0031)	(0.0048)	(0.0040)
Couple FEs	✓	✓	✓
Time FEs	✓	✓	✓
Bandwidth	11.22	11.22	11.22
N	783,840	330,694	453,146

Notes: This table presents regression-discontinuity and differences-in-discontinuities estimators identifying the effect of ending Medicare coverage for SNF stays, which ends on day 21 for qualifying stays. The first two rows present the estimated effect of losing coverage on the probability a focal spouse will remain in the SNF, stratified by before or after the index event. The third row presents the difference-in-discotinuities estimator as discussed in the text. SNF stays within 4 months of the treatment and placebo events are included in the regression.