

Aging 7/30/02
COILE 1:15 PM

Preliminary and Incomplete
Comments Welcome

Health Shocks and Couples' Labor Supply Decisions

Courtney Coile, Wellesley College

NBER Summer Institute Workshop on Aging
Cambridge, MA
July 30, 2002

Acknowledgements

The author gratefully acknowledges financial support from the Center for Retirement Research at Boston College and comments from Deborah Dwyer, Jonathan Gruber, and seminar participants at MIT and at the Retirement Research Consortium Conference. A preliminary version of this paper was circulated under the title "The Effect of Health Events and Health Insurance on Couples' Labor Supply Decisions."

For people of near-retirement age, unexpected negative health events are unfortunately quite common. For example, men and women participating in the Health and Retirement Study (HRS) were found to have a 5% chance of having a heart attack, stroke, or new cancer diagnosis, a 10% chance of being diagnosed with a new chronic illness, and a 3% chance of having an accidental injury over a two-year period. These negative health events lead individuals to decrease their labor supply, as has been shown by McClellan (1998) and others. However, the effect of each spouse's health events on the other spouse's labor supply has not been previously studied.

There are two reasons to be concerned about the spouse's response to health shocks. The first is that it has important financial implications for the family. The family has suffered a drop in lifetime income as a result of the negative health event; a decrease in spousal labor supply will exacerbate that loss, putting the family at greater risk for poverty in retirement, while an increase in labor supply may partially offset the loss. The second is the possibility that the availability of government benefits such as Disability Insurance may "crowd out" the spousal labor supply response. A considerable amount of crowd out would suggest that the existence of these benefits does not raise family income significantly, though presumably it raises family utility by allowing the healthy spouse to be at home with the sick spouse.

This paper explores the effect of negative health shocks, such as heart attacks or new diagnoses of chronic illnesses, on the labor supply of both the affected spouse and his or her partner. In so doing, the paper links two important strands of the retirement literature, the large literature on health and retirement and the small but growing literature modeling retirement in a family context. This paper may also be viewed as an extension of the literature on spousal labor

supply as insurance against negative events, which measures whether there is an “added worker effect” when one spouse becomes sick and whether it is crowded out by public insurance programs. This work uses the first five waves of the Health and Retirement Study (HRS), a recent, nationally representative survey of the young elderly with extensive information on health, labor force status, and demographics.

The use of health shocks as an alternative to self-reported health status is appealing due to the concern that health status may not be independent of labor force outcomes if people seek to rationalize their retirement status by claiming a health problem, a concern that has been validated in work by Bound (1991) and others. In theory, the onset of a negative health event may have several effects on the labor supply of the healthy spouse. The healthy spouse may increase labor supply, as the household has suffered a drop in lifetime income and the healthy spouse’s leisure time is likely a normal good. The healthy spouse may also increase labor supply in order to provide health insurance for the family. Alternatively, the healthy spouse may decrease labor supply if the health event has resulted in a change in home production opportunities or in the complementarity of leisure among the spouses. Thus the effect of the health event on the spouse’s labor supply is theoretically ambiguous.

This analysis examines the average response of spouses to their partners’ health shocks, as well as the possibility of different responses by couples depending on the importance of the drop in lifetime income, caregiving opportunities and complementarity of leisure, and the need to provide health insurance in their families. Particular attention is paid to whether men and women have similar responses to a spouse’s health shock, as Coile (1999) suggests that men and women respond differently to their spouse’s financial incentives for retirement due to a difference in the strength of complementarity of leisure.

The paper has several major findings. First, health shocks have an important effect on own retirement: a heart attack accompanied by a significant decrease in functioning is associated with a labor supply reduction of 1,030 hours for men and of 654 hours for women. Second, in the sample as a whole, health shocks have no significant effect on spouse's retirement for either men or women. Third, there is some evidence that couples respond to health shocks in the way we would expect, such as decreasing hours if the sick spouse has retiree health insurance or increasing hours if alternate care-givers are available. This suggests that the aggregate non-response to health shocks may be explained by offsetting responses by different groups. These offsetting responses are more often found for men, suggesting that men may respond more to their spouses' health shocks than women. Finally, there is evidence of substantial crowd out of spousal labor supply by Disability benefits.

The remainder of the paper is organized as follows. Section I describes the previous literature on the topic. Section II discusses the data and empirical strategy. Section III presents the empirical results, first for models incorporating only the effect of own health shocks, then for models incorporating the effects of the spouse's health shocks. Section IV discusses the implications of the findings and directions for future research.

I. Previous Literature

Three strands of literature are relevant for this project. The first is the literature on health and retirement. Within this literature, several methods have been used to capture the effect of health on retirement, the first of which is to include self-reported health status or self-reported work limitations in the retirement model. Studies using this method have typically found large

effects of health on retirement.¹ However, this method is subject to both measurement error, as individuals' subjective judgements of what constitutes poor health may vary substantially, and to endogeneity concerns, as self-reported measures may not be independent of labor force outcomes if people seek to rationalize their retirement status by claiming a health problem. If endogeneity is important, the estimates from these models may overstate the true importance of health.

As a result of these concerns, a second set of studies instead uses objective measures of health such as information on medical conditions or subsequent mortality.² While these measures may be only imperfectly correlated with working capacity, making estimates using them subject to measurement error, studies using them have found significant effects of health on retirement, though smaller than those estimated using self-reported health. A third set of papers instrument for self-reported measures using objective measures.³ Bound (1991) finds that self-reported and objective measures are both subject to the problems discussed above and shows that instrumenting for self-reported measures with objective measures may actually exacerbate endogeneity concerns. Thus the conclusion of this literature is that health has an important effect on retirement, but there is no perfect method for estimating the magnitude of the effect.

McClellan (1998) avoids the endogeneity problem by using shocks to health that occur between waves 1 and 2 of the Health and Retirement Study to estimate the effect of health on retirement. He finds that acute health events have significant effects on labor supply: an event such as a heart attack is associated with a labor supply reduction of 600 hours, or 1700 hours if accompanied by a serious loss in functioning. However, the study touches only very briefly on

¹ See, for example, Diamond and Hausman (1984), Hanoch and Honig (1983), Burtless (1986), or Gustman and Steinmeier (1986).

² See, for example, Anderson and Burkhauser (1985), Bazzoli (1985), and Chirikos and Nestel (1984).

³ See, for example, Stern (1989), Bound (1991), and Dwyer and Mitchell (1999).

the effect of one spouse's health shocks on the other spouse's labor supply, the subject of this analysis.

The second strand of literature that is relevant to this analysis is the literature on couples' retirement decisions. Several studies, including Hurd (1990), Maestas (1999), and Gustman and Steinmeier (2000), estimate structural models of joint retirement. These authors employ different models, but all find that complementarity of leisure is a key factor in explaining why so many husbands and wives retire simultaneously. A second set of studies, including Baker (1998), An et. al. (1999), Coile (1999), and Johnson and Favreault (2001), estimate reduced-form models exploring the cross effects of one spouse's characteristics on the other spouse's retirement decision. A feature common to both sets of studies, however, is that they either do not control for health or do so using self-reported health status, subjecting them to the critique that the resulting estimates of the effect of health are biased.⁴

Two studies on couples' labor supply deserve further mention. Coile (1999) examines the effect of one spouse's retirement incentives on the other spouse's labor supply and finds that men are more responsive to their wives' incentives due to a stronger complementarity of leisure effect. Gruber and Cullen (2000) points out that spousal labor supply can serve as insurance against a negative event (such as unemployment in their analysis, or a health shock here) and that the "added worker effect" may be crowded out by public insurance programs. This analysis builds on these studies by exploring whether the effect of health shocks on the spouse's labor supply is similar for men and for women and whether spousal labor supply responses are crowded out by the availability of government benefits such as Disability Insurance.

⁴ Johnson and Favreault (2001) control for both self-reported health status and an objective measure, the number of functional impairments, in the analysis. However, that study does not look at the change in functional impairments or in other health measures over time, as is done here, or explore whether the effect of health measures on the spouse's labor supply depends on factors such as the family's access to health insurance.

Finally, the literature on health insurance and retirement is also relevant for this analysis, as the need to provide insurance may influence the spouse's response to the health shock. Using the first two waves of the HRS, Blau and Gilleski (1997) find that the availability of employer-provided retiree health insurance increases labor force exit by 2 percentage points if the individual shares the cost of coverage with the firm and 6 percentage points if the firm pays the entire cost. Exploiting federally mandated continuation of coverage provisions through COBRA, Gruber and Madrian (1995) find that an additional year of available coverage increases the probability of retirement by 2.2 percentage points.

In brief, this paper extends the previous literature on health and retirement by estimating reduced-form models that measure the effect of one spouse's health shocks on the other spouse's labor supply. The empirical analysis also allows the response to the spouse's health shocks to differ depending on factors such as access to health insurance. The paper explores the extent to which spousal labor supply serves as insurance in the event of a negative health shock and tests for asymmetrical responses to health shocks by husbands and wives.

II. Data and Empirical Strategy

The data for the project is the Health and Retirement Study (HRS), a survey of persons aged 51-61 in 1992 and their spouses. The survey contains extensive information on health, labor supply, and demographic characteristics. Data for the first five waves of the study, 1992-2000, are used.⁵ The sample is constructed using the 4,617 married couples who are in the sample at wave 1 and are observed for at least two consecutive waves of the survey. The male and female samples consist of all person-year observations for waves 2-5 in which the individual

⁵ The wave 5 data is the preliminary release; all other waves are final release data.

is between the ages of 50 and 69 and was working at the previous wave.⁶ The final sample size is 9,699 male person-year observations and 8,135 female person-year observations.⁷

Following McClellan (1998), three types of health shocks are examined: acute health events (heart attack, stroke, new cancer), onset of new chronic illnesses (diabetes, lung disease, heart failure, and arthritis), and accidental injuries or falls.⁸ Another key health variable is the functional impairment index. The index is based on whether the individual reports any difficulty in performing a series of seventeen activities of daily living (ADL), such as walking, climbing stairs, lifting ten pounds, and getting out of bed; the index ranges from 0 (difficulty in no activities) to 1 (difficulty in all 17 activities).⁹ The analysis also makes use of the self-reported survival probabilities, as health shocks provide new information about mortality and this may influence labor supply decisions.

Table 1 shows the summary statistics for the male and female samples. Over a two-year period, the typical male decreases annual hours by 353 hours and has a 17% chance of exiting the labor force completely. Over a two-year period, 6% of men experience acute health events, 11% are diagnosed with a new chronic illness, and 3% are injured in an accident; 12% of men

⁶ Wave 1 observations cannot be used because the paper examines the effects of changes in health during the previous two years on labor supply, and it is not possible to determine whether health problems reported at wave 1 began within the past two years or at an earlier point in time.

⁷ The standard errors in the regressions below are corrected for repeated observations on the same individuals.

⁸ Measures are constructed so as to be as compatible across waves as possible. The two most significant changes are that asthma is included in the chronic illness measure only for waves 1-2 due to lack of data thereafter and that the question about accidents changes from recent accident or injury in waves 1-2 to fall-related injury in waves 3-5 (the fall-related injury question is also only asked of older respondents). Measures are also constructed to be similar to McClellan (1998); the most significant difference is the exclusion of back pain from chronic illness, due to implausibly large fluctuations in the number of new diagnoses in later waves.

⁹ The choices offered to respondents for the ADL questions change across waves:

- Wave 1: "Is [the activity] not at all difficult, a little difficult, somewhat difficult, very difficult, or something that you can't do at all?"
- Wave 2: "Do you have any difficulty [with the activity]?" followed by "Is that a little or a lot of difficulty?"
- Waves 3-5: "Do you have any difficulty [with the activity]?"

In creating the index, individuals reporting any level of difficulty are treated as impaired, as this is the only measure available for all waves. On average, individuals are 1-2% more disabled at each passing wave; however, individuals are 4% less disabled at wave 2 than wave 1, presumably due to the change in definition. To make the measure more comparable across time, I decrease each individual's wave 1 impairment index by one activity (roughly 6%).

have experienced an acute health event at some point in the past and 46% have been diagnosed with a chronic illness. The typical male reports difficulty performing 10% of the 17 activities of daily living and has experienced a 1.5% increase in the ADL index during the past two years; he rates his odds of living to age 75 as 66%.¹⁰ The average age of the male sample is 60. In the male sample, 24% of individuals have less than a high school education, 32% have high school only, 20% have some college, 11% have graduated from college, and 13% have graduate education. Median household net worth is \$145,200 and median liquid assets are \$4,700. The sample statistics for women are similar, except as noted. The typical female decreases hours by 275. Compared to men, women have a higher incidence of past chronic illness, 53%, and a higher ADL index, 13%. The typical female in the sample is younger (58) and has less education than the typical male.

The empirical strategy is to exploit exogenous shocks to health between waves of the survey to explore the effect of health on one's own and one's spouse's labor supply. Two dependent variables will be used: the change in hours (continuous) and exit from the labor force (dummy).¹¹ In addition to the health measures, the regressions will include a full set of age dummies, education dummies, industry and occupation dummies, net worth and liquid assets, and year dummies.

Table 2 displays trends in the incidence of health shocks and in the impairment index across waves of the survey. The fraction of men experiencing an acute event such as a heart

¹⁰ This average corresponds well with the actual probability of living to age 75 from the 1995 life tables, which is .678 for the average man in the sample. Interestingly, men greatly overestimate their probability of living to age 85 (sample average of .449 vs. .320 from life tables), while women greatly underestimate their probability of living to age 75 (sample average of .682 vs. .782 from life tables) but make an accurate predication regarding age 85. Using wave 1 of the HRS, Hurd and McGarry (1995) have a similar finding.

¹¹ Results are very similar using a dummy=1 if hours go to zero or a dummy=1 if the individual switches from reporting himself as not retired or partly retired to completely retired. Labor force re-entry is also an interesting phenomenon to explore, though it is sufficiently infrequent (affecting about 3% of the sample between each wave) so as to make such analysis infeasible.