Mitigating the Consequences of a Health Condition
The Family in the PSID

Michael Dalton *[1] and Daniel LaFave †[2]

[1]Duke University, Economics Department
[2]Colby College, Economics Department

March 29, 2013

Abstract

This study exploits longitudinal and genealogically linked data from the Panel Study of Income Dynamics to examine the relationship between health status, risk-sharing, and mechanisms of informal family insurance. In the context of well-developed theoretical models, we relate deteriorations in health due to specific ailments to losses in expenditure, labor market outcomes, and increased family assistance along a number of channels. These channels include drawing down wealth, increasing transfers, and taking in ill family members. The relationships are empirically robust to models exploiting the longitudinal structure of the data to look at changes within an individual over time. Our results suggest that family networks fill gaps left by formal insurance coverage. The results are informative for understanding patterns of resource allocation within families and how informal networks operate while facing adverse events.

*michael.dalton@duke.edu
†daniel.lafave@colby.edu

1 Introduction

Changes in health as an adult may have serious consequences for short and long-term economic outcomes. These combined impacts necessitate an established safety net for those individuals and families who are most affected. Previous literature has examined the extent that different channels of insurance are relied upon in the presence of shocks. However, while widely acknowledged as important, little is known empirically concerning the relationship between health deteriorations and within family responses. This is particularly true in developed country contexts. Using well-known models from the consumption smoothing and risk-sharing literature, we exploit longitudinal data from the Panel Study of Income Dynamics (PSID) to examine how individuals and families are exposed to health shocks and the mechanisms used to respond to poor health. The results emphasize the connectedness of families across generations while highlighting the exposure to health risk faced by families in a developed setting.

The framework for this study is founded in models of family interaction and informal insurance (e.g. Cochrane (1991); Townsend (1994)), as well as empirical evidence highlighting the role of family transfers in informal networks. We go beyond past work to not only test the existence of risk-sharing and insurance, but to examine particular mechanisms including labor supply responses, informal transfers, and co-residing with family members. Moving beyond examinations of expenditures allows us to illustrate multiple channels families may use to provide informal insurance. Literature in developing settings has documented a relationship between health shocks and non co-resident family transfers (e.g. Fafchamps and Lund (2003); Genoni (2012)), while evidence from developed countries has also highlighted the important role of intergenerational and intrafamily exchange (e.g. Cox (2003); Hotz et al. (2010)).

Our work relies on unique features of the PSID and contributes to the literature in a number of ways. First, by exploiting the split-off following rule of the survey, this paper emphasizes the importance of both coresident and non-coresident family as a main channel of insurance in response to deteriorating health. Second, we combine health and transfer responses with rich labor outcome data to determine how health conditions affect economic outcomes of both the individual and their family. This allows us to examine the extent that family transfers respond directly to health conditions relative to labor market impacts. Third, we incorporate responses regarding health insurance to show how informal forms of insurance and the family help fill in gaps that formal health insurance may miss. Finally, by using the unique genealogical structure of the data, we determine which specific relationships in the family provide

informal insurance against changes in health. We are also able to control for the extended family's health history to assess its impact on both the transmission of certain health conditions and a family's likelihood of different insurance responses.

Our results show that households and extended families are exposed to health risk and are unable to fully insure consumption or efficiently share the risk brought on by these shocks. This is true both for young and elderly respondents, as well as whether or not the respondent is employed. In the face of deteriorating health, we show that individuals reduce their labor supply and earn less in the labor market, and respond to these losses by drawing down assets. We then show that the family responds to these events by increasing labor supply, increasing transfers, and taking-in family members. The results paint a picture of partial insurance against health declines that is provided by family networks, findings particularly unique in the developed country setting.

The following section outlines a theoretical framework for studying consumption smoothing, risk allocation, and the responses of family members. We then discuss our empirical implementation and rich data before presenting results.

2 Theoretical Framework

The theoretical motivation behind this paper lies in models of insurance within households and informal networks. The development of both full insurance and risk sharing models are well known in the literature, including seminal work of Cochrane (1991), Townsend (1994), and Hayashi et al. (1996). We briefly discuss the two models we focus on - full insurance and efficient family risk sharing - below.

2.1 Full Insurance and Consumption Smoothing

In models of full insurance, households are able to smooth away risk in order to maintain consumption. Working from a point of expected utility maximization, what matters for determining consumption in a given period t is only permanent household resources, independent of idiosyncratic fluctuations in income (e.g. Deaton (1992)). While households face a number of states with potential consumption realizations, they have access to state contingent means of insurance to equate marginal utility across time. Income fluctuations or health shocks should have no effect on the realized change in consumption between periods.

While we begin with the model of full insurance as a starting point, we are particularly interested in examining how negative health events impact well-being along a number of dimensions. As a large literature has tested and rejected the full insurance

model, we use it as a baseline for a more detailed analysis of household and family responses to negative shocks.[1] Rather than focusing only on changes in consumption, we look to examine both the extent that households are unable to insure against deteriorating health as well as the mechanisms that they may use to attempt to achieve partial insurance.

This analysis follows in line with literatures examining whether households are able to fully insure consumption against deteriorating health (e.g. Gertler and Gruber (2002)), and the formal and informal mechanisms used as insurance (e.g. Genoni (2012)). The majority of this literature comes from developing countries, where the analysis is framed in terms of missing markets for insurance and health care. Our results show that even households with health insurance in the U.S. are unable to fully insure against this risk.

2.2 Efficient Risk Sharing

Another subset of work relaxes the prediction of full insurance that each household can completely insure themselves, and instead looks at insurance within groups. In models of static group risk sharing, there is an aggregate component of risk one's group or network is unable to insure against, but transfers smooth consumption within the group (e.g. Townsend (1994); Hayashi et al. (1996)). Each member of a group receives a pareto weight on their utility which ensures that the ratio of marginal utilities is constant regardless of the aggregate shock to the group.

Literature from developing contexts often considers the village as the insurance group (e.g. Townsend (1994)), or one's self-reported social network (e.g. Fafchamps and Lund, 2003). These papers often reject the model of efficient risk sharing, and cite a variety of frictions related to asymmetric information and limited commitment.[2]

We choose instead to focus on the genealogically linked family, as there are theoretical reasons why the family may be capable of sharing resources efficiently. In a series of influential papers, Chiappori and coauthors examine the sharing of resources within the household and fail to reject Pareto efficiency (e.g. Chiappori (1988); Chiappori (1992); Bourguignon et al. (1993)). If the family has superior ability to limit asymmetric information, or family members are agents interacting in a repeated game with no determined endpoint, then the extended family can be an efficient network for sharing risk and uncertainty. This is supported by recent work on kinship networks

[1] For rejections of tests of consumption growth being independent of fluctuations in income, see Cochrane (1991) and Hayashi et al. (1996) among other. The model has also been rejected in developing settings including Cote d'Ivoire (Deaton (1997)), India (Townsend (1994)), Ethiopia (Asfaw and Braun (2004)), Burkina Faso (Kazianga and Udry (2006)), and Thailand (Townsend (1995)).

[2] A subset of recent papers has highlighted the potential to refine tests of risk-sharing to incorporate heterogeneity in risk preferences, see Schulhofer-Wohl (2011) and Mazzocco and Saini (2012).

and risk reduction (Kinnan and Townsend (2012)), and extended families and child development (LaFave and Thomas (2013)), yet remains an open question on whether the family is an efficient unit (Coate et al. (2013)).

The subtle difference between the two models motivates an empirical strategy focusing on two different types of variation. We describe these models in the following section.

3 Empirical Framework

3.1 Full Insurance

As outlined in the previous section, the starting point of our analysis is to examine the extent that deteriorations in health are reflected in consumption patterns. Our primary specification relates our outcomes of interest to health status while controlling for time varying controls and individual fixed effects:

$$Y_{ht} = \beta\theta_{ht} + \delta X_{ht} + \mu_h + \epsilon_{ht} \tag{1}$$

where Y_{ht} is the outcome of interest, beginning with non-health expenditure, for household head h in time t. Health status of the household head is measured by θ, while time varying variables related to the individual and household are controlled in X. These include time fixed effects and a variety of demographic characteristics of the head of household, spouse of the head, other members of the household, and non-coresident family. We include polynomials on head and spouse's age and education, household income, household wealth, controls for head and spouse working status, polynomials for household size and composition, and a dummy for whether the head is covered by health insurance.

A primary concern with identification is that individuals that are more prone to experiencing a health condition are permanently different in other unobserved ways than others who may not experience health deteriorations. One example would be if poorer health is related to permanent income, which will affect consumption choices throughout life in a way that is separate from the actual health condition. One way to control for permanent characteristics is to incorporate the panel nature of the data and include individual level fixed effects, μ_h, where the individual is the head of household h. This will control for cases where reporting severe limitations due to a health condition are correlated with permanent, unobserved characteristics of the individual.

The parameter of interest, β, is then identified by examining how changes in health

status are related to changes in our outcomes of interest. In tests of insurance, if households are fully insured, fluctuations in health status should have no relationship to non-medical expenditure, and β will be zero. We will also use the model to examine additional outcomes including the specific channels which families rely upon for insurance, and the responses of non-coresident family members to negative health events.

3.2 Efficient Risk Sharing within the Family

As described in the theory section, we distinguish between full insurance of the household, where consumption does not respond to idiosyncratic health events, and efficient risk sharing with a family unit, where consumption may co-move with aggregate family expenditure, but is distributed consistently within the family unit. Even if we see evidence that extended families do share resources with one another, there are many potential reasons why households may fail to fully share risk. Issues of commitment or asymmetric information are two potential reasons why a family may fall short of fully insuring each other's idiosyncratic risks.

Using the PSID, we can test the efficient sharing of risk amongst genetically related households. The intuition behind the test is to control for aggregate resources of the network for the year, and determine whether the idiosyncratic shock specific to a member of the network affects their own consumption. While a health shock can impact both the consumption of the household and extended family, efficient sharing of risk implies that after accounting for the loss of consumption for everyone in the family due to the shock, individual shocks to household h should have no impact on household h's consumption.

The empirical model for household head h, in family unit f, and wave t is the following:

$$Y_{hft} = \beta\theta_{ht} + \delta X_{hft} + \psi_{ft} + \mu_h + \epsilon_{hft} \tag{2}$$

where family-year fixed effects, ψ_{ft}, are included on top of the controls in the previous model. This allows us to control for family wide consumption in a given year, and to examine the consumption response to a health shock relative to the consumption of other members of the family in the same year. By construction, time varying family fixed effects also control for time-constant characteristics shared across all family members capturing common, additive components of ability, permanent income, genetic background, and preferences.

Identification of β comes from comparisons in expenditure between family members

in a given year while continuing to incorporate individual fixed effects specific to the head of household in order to focus on deviations from the average consumption of the household. In this case, β equals zero if the extended family shares risk efficiently.

4 Data

We use the 6 waves of interviews from 1999-2009 of the Panel Study of Income Dynamics (PSID). The PSID has many unique features including its genealogical construction and recent additions to the question series that make it a useful dataset for our study. Since beginning in 1968 with a nationally representative set of households, the PSID has used a unique following rule that continues to interview individuals from the original households even as they move out and form new households. The study also follows the offspring of any original 1968 PSID household members. This following rule continues with each subsequent generation, leading to a current day sample of approximately 8600 households composed of multiple generations within numerous genetically linked families. In addition to contemporaneous interviews with multiple households within a family, the longitudinal nature of the PSID allows us to analyze changes in health reported by the same individual over time in order to control for fixed, unobserved individual heterogeneity.

We focus on the 1999-2009 waves of the survey that contain new detailed questions on health, consumption, and a consistent wealth module. Prior to 1999, the only consistent health question in the PSID is general self-reported health status.[3] This question has been found to have a number of issues related to measurement as well as correlation with observed characteristics that lead to concerns about correlations with unobserved characteristics (e.g. Strauss and Thomas (1998), Strauss and Thomas (2007)). In response to the need for a more robust measure of health, researchers developed health status questions anchored in objective events. In 1999, the PSID began including questions about 12 specific acute, chronic, and psycho-social health conditions including a question about the degree of limitation to daily physical activities due to each condition.[4] The respondent can report a lot, somewhat, a little, not at all, or never diagnosed in terms of how much the condition limits daily activities. The survey also began asking the now more common Activities of Daily Living (ADL) questions about whether or not the head of household faces limitations on specific regular activities. This newly available longitudinal health data allows for new contributions to the health and risk sharing literature in the developed country

[3]The question asks "Would you say your health in general is excellent, very good, good, fair, or poor?"

[4]The 12 conditions asked about are stroke, high blood pressure, diabetes, arthritis, asthma, lung disease, cancer, heart attack, heart disease, emotional distress, memory loss, and learning disabilities.

context.[5]

Another new feature of the PSID data is a more complete set of questions related to expenditures. Prior to 1999, the survey focused only on food expenditures with some additional consumption questions related to housing. The PSID currently asks a larger set of questions allowing one to seperate housing, education, transportation, health and care in addition to the previously asked food questions. With this more complete consumption module, we can perform tests related to consumption smoothing that were either not possible or limited prior to 1999.

Another important feature for our study is that the PSID began asking the wealth module of questions each wave beginning in 1999. These questions allow us to incorporate measures of specific assets and debts in order to determine exactly what means are used to insure against deteriorating health. Combined with data on labor market outcomes and family transfers, we are able to illustrate a detailed picture of informal insurance within the family.

The summary statistics for some of the key variables are presented in Table 1. This table shows that there are differences on some observable characteristics between respondents and the degree of physical limitation that they report. Column 1 includes the full sample of individuals, while Columns 2 through 5 limit the sample to only those who have ever reported a health condition. Many of the differences are not surprising: respondents who report "A lot" of limitations are older, more likely to be retired, and slightly less educated. Since there are differences in observed characteristics, this emphasizes the need for an empirical strategy that deals with observed and unobserved differences between respondents. Our results presented below rely on changes across time for a given individual rather than comparisons across the groups in Table 1.

5 Results

We begin by testing whether households experience changes in consumption due to deteriorations in health status. This then motivates a closer examination of the effects of deteriorating health, the mechanisms used to insure against this risk, and the role of both intrahousehold and intergenerational family behavior.

[5]The PSID health questions we use in this study have been shown to accurately reflect patterns of health in the National Health Interview Survey (Andreski et al. (2009))

5.1 Do Households Smooth Consumption Against Health Risk?

Table 2 presents the results from equation 1 with non-health consumption for a given household in a given year as the dependent variable. Recall that household-head fixed effects are included to focus only on variation within a person over time. We use two different measures of consumption in order to account for the effect that household size may have (Wagstaff (2007)): log of expenditures for all non-health consumption categories (food, housing, education, transportation, and care) and log of per capita expenditures for all non-health consumption categories. The per capita measure allows household size to directly affect the dependent variable.

The regressors of interest are the degree to which the head of household reports having limitations on normal daily activities due to any of the 12 specific conditions. Table 2 examines the relationship between each level of limitation and consumption. One issue with this strategy is the concern of non-separability between experiencing a health condition and utility. For instance, if experiencing a health condition leads to drops in consumption related to a change in the utility function, then we may see consumption fall for reasons that are not due to the inability to smooth utility over the event of a negative health condition. By looking across different levels of limitation, we are actually comparing all individuals who experience one of the 12 specific health conditions. If changes in consumption were due only to some sort of state dependence or non-separability, then it is likely that we should see consumption responses be consistent no matter the degree of limitation. However, if consumption smoothing is an integral part of the story, the degree of limitation will matter since more severe limitations can lead to larger bills of health expenditures and larger effects on labor supply. This begins to seperate stories of insurance versus state dependent preferences.

The results in Table 2 tell a number of stories. First, the degree of physical limitation does indeed matter for consumption. Using either measure of household consumption, or whether or not individual fixed effects are included, reporting "somewhat" or "a lot" of physical limitations has a statistically significant negative effect on consumption. For example, relative to never having experienced any limitations, experiencing "a lot" is related to a 5.5% reduction in expenditures (Column 2). While all coefficients are negative, after accounting for individual fixed effects, having "no" or "little" current limitations due to a health condition does not have a significant impact on consumption. The magnitude of effects follow an intuitive pattern, with "a lot" of limitations being associated with larger reductions in consumption than only some limitations. While we cannot completely rule out non-separability between a health condition and preferences, by looking within the series of respondents'

interviews experiencing a health condition, these results do provide evidence that at least part of the story is about consumption smoothing being more difficult when experiencing a stronger health "shock". Non-separability could still be an issue if preferences are determined by the degree of limitation, as opposed to only being a function of experiencing a health condition.

Also important to note is that the magnitude of coefficients do change after controlling for individual level fixed effects, leading to a larger than 50% decrease in magnitude for many of the coefficients. Furthermore, the coefficient on "a little" limitations loses its statistical (and economic) significance after controlling for fixed effects. These results suggest that individuals experiencing physical limitations may have permanently different consumption patterns that could bias our coefficient estimates. For this reason, we focus on models including individual fixed effects throughout the rest of the paper. We also focus on instances when individuals report "a lot" of limitations.[6]

The next step is to better understand what might lead to difficulties insuring against health shocks. We start by examining some of the characteristics of the heads of household in order to see how health shocks may differentially affect subgroups of the sample. First, we split the sample into a group aged 50 or younger and a group 50 or older. The goal is to examine whether insurability of health shocks is only an issue at certain points in the life cycle. Since many health conditions are more likely to occur as one gets older, experiencing a change in health may be more anticipated for the older population. The younger group may have both temporary and permanent earnings negatively affected by a bad health shock, suggesting that it will be more difficult to insure against a health shock for the younger population. On the other hand, the younger group may be more adaptable and capable of dealing with an unforeseen health shock. Panel A of Table 3 examines these two groups. The results show that both age groups have consumption significantly impacted by experiencing a health shock, although the younger group has nearly twice as large of an effect.

It may be possible that this age distinction is simply picking up labor supply decisions. Table 2 provided suggestive evidence that non-separability of utility and health is not the sole explanation for the dip in expenditures; however, it may be the case of non-separability between labor supply and utility drive the results. If labor supply responds to health shocks, we may be picking up changes in consumption correlated with additional leisure. Panel B examines household heads that worked at all during the previous year compared to household heads that did not work. Interestingly, we see that both groups are significantly affected by the health shock.

[6]Continuing to distinguish between "a lot" and "some" limitations reveals results qualitatively similar to those presented here.

While the coefficients are statistically indistinguishable, the group that works is more affected by the health shock. The results in Panel B suggest that while health shocks may affect labor supply decisions, health shocks impact even those individuals who are not working.

While rejecting consumption smoothing in and of itself is not particularly novel, we are able to go further than past work and exploit the richness of the PSID to examine the responses to deteriorating health on a number of dimensions. Some reasons we may see non-health consumption drop in response to a health shock is that preferences may change, health expenditures can increase to deal with the health condition, or labor supply may decrease. While Table 2 and Table 3 provided evidence that preferences are not the only story at play, we look to examine the other two explanations in Table 4. In Column 1 of Table 4 we see that health expenditures increase by 530 dollars more than average when experiencing a severe physical limitation due to a health condition, with the coefficient significantly different from zero at the 1% level. Also matching with expectations, we see that the head of household reports working approximately 150 fewer hours than average over the year when experiencing severe physical limitations. Relatedly, the head of household earns about $1500 less in labor income over the year. All of these results provide further evidence that experiencing physical limitations can have a serious impact for an individual and household if not insured properly.

We now shift our focus to examining the way individuals, households, and families may respond to a member experiencing a negative health shock. Panel A of Table 5 examines the effect of severe physical limitations on wealth and the likelihood of entering a nursing home. Experiencing severe physical limitations has a very large and statistically significant effect on total wealth of the household. Since wealth also includes total debt obligations of the members of the household, this is likely evidence that respondents who experience these health conditions are turning towards credit or debt in order to deal with the costs associated with major medical bills. This is consistent with results in developing settings showing depletions of more liquid assets in the face of negative shocks (Frankenberg et al. (2003), Fafchamps and Lund (2003)).

Column 2 shows that a respondent is approximately 2 percentage points more likely to enter a nursing home when having many physical limitations. This is a near doubling of the probability given the mean of 2.2%. On top of potential major medical bills, the respondent may also face new expenses related to assisted living services, which may not be captured in the questions asking about health related expenditures. This is another reason why we may see the respondent's spending down wealth where

we see a more modest increase in annual health expenditures.

Panel B further explores the detailed wealth module to examine specific components of wealth. We look at reports from the respondent about the household's money in savings and checking accounts, the amount of debt outside of mortgages, amount in stocks, amount in a retirement account, value of automobiles, and home equity. The results show severe physical limitations have a significant positive impact on debt, and negative impacts on auto value and home equity. These results suggest that some of the ways that people deal with the costs of health conditions is to take on additional home equity debt and to reduce their durable assets. The respondents may be selling off automobiles and either selling off homes or taking additional mortgages on the equity of a home in dealing with a decline in health.

5.2 What role does the Family Play?

Having established significant impacts of deteriorating health to the individual and their household, Table 6 examines how an individual may rely on their family both within and outside of the household to smooth consumption and utility using equation 2. Column 1 shows how the spouse's labor supply increases by approximately 40 hours in the face of partner's health deteriorating. Recall from Table 4 that respondents decreased their own labor supply by about 150 hours in response to a health shock. These results suggest that while the spouse does change labor patterns, only a little more than 25% of the hours lost due to a health condition of the respondent are made up.

Another potential informal insurance mechanism is transfers from non-coresident family members. Column 2 shows the respondent is 4 percentage points more likely than average to receive a transfer from a family member outside of the household upon experiencing severe physical limitations. The value of the received transfers also increases as shown in Column 3. Monetary transfers between households has been studied in a variety of contexts, but this result provides new evidence supporting the idea that transfers are used as a form of informal insurance against negative events. In particular, the family seems to be the source of these transfers.

A different form of insurance from the family can be the option to co-reside. The next column looks at the likelihood that the respondent will move into the household with a family member also interviewed by the PSID. The result suggests that the respondent is about 1 percentage point more likely than average to share residence with a family member after experiencing severe physical limitations. This is about half as large as the effect on entering a nursing or assisted living home (Table 5, Column 2). However, this is also statistically significant and large relative to the baseline mean,

which suggests this is another way that individuals can cope with a health shock. Moving in with a family member can help in a variety of ways. First, co-residing makes it possible for the family to look after and provide assistance to the respondent with deteriorating health. Second, co-residing can help share costs and reduce some of the financial burden brought on by the health shock. Furthermore, in the context of unemployment, other papers have found the phenomenon of moving back in with parents to be one form of insurance for young adult experiencing unemployment (McElroy (1985); Kaplan (2012)). These results provide new evidence in a different way in which adult family members may choose to co-reside. Hotz et al. (2010) examine the location decision for parents later in the life cycle, and the results in this paper add to the findings in that research.

5.3 Do Families Efficiently Share Risk?

Since there is evidence of non-co-resident, genetically linked households sharing resources, it is a reasonable question to ask whether households that are genetically linked are capable of sharing risk in an efficient manner. As outlined in Section 3, the test of risk sharing examines whether idiosyncratic health shocks to an individual within a family unit impact consumption after controlling for aggregate family resources in that period. If families share risk efficiently, only the aggregate resources will matter.

The results for the test of full risk sharing within the family from equation (2) with both family-year and individual fixed effects appear in Table 7. For both measures of consumption, we find evidence that the family does not fully share the risk of declines in health. Since the health decline specific to the head of household is related to approximately a statistically significant 3.5% reduction on household consumption after controlling for family resources, this is evidence that risk is not efficiently shared amongst genetically related family members. Despite the evidence from Table 6 on the smoothing responses of non-coresident family through transfers, we still see that families are not able to fully smooth away risk.

The previous specification treats the entire extended family interviewed in the PSID as the primary risk sharing network. However, it may be the case that risk sharing occurs at a more contained level between closely related individuals. For instance, Genicot and Ray (2003) provide evidence that endogenous networks may have bounded size in equilibrium in the case of non-cooperative risk-sharing. Even without bounds on the size of a network, households may only want be connected to one another when a sharing arrangement is reasonably enforceable. A bond with a parent is likely to be stronger than the bond with distant cousins. A stronger bond

may imply reduced asymmetric information and better commitment devices. For this reason, it is an interesting exercise to explore whether a tighter, more contained network may do a better job at sharing risk. This is left for ongoing work.

5.4 Comparing Channels of Informal Insurance

Having established that families are not able to fully insure against risk but do respond to negative health events of their members with partial insurance, the final table examines which of the potential insurance methods does the best at smoothing non-health expenditures in response to a negative health event. In a parsimonious model, we examine how having health insurance, moving in with a family member in the PSID, drawing down wealth assets, receiving a financial transfer from a family member, or having a spouse that works each counter negative effects of health deterioration. The variables of interest are the interactions between each of these possible insurance channels with a dummy for having severe physical limitations. The goal is to see how well each type of insurance mitigates the drop in consumption related to experiencing a negative health event. The models continue to incorporate individual level fixed effects as in equation (1).

Columns 1 and 2 of Table 8 present these results. Having a spouse working and moving in with a family member have the largest positive impacts on consumption when faced with a negative health shock. Moving in with a family member comes closest to fully mitigating the consumption loss associated with experiencing the health event. In fact, moving in with a family member actually more than balances the effect of severe limitations for log of per capita expenditures (-0.14 + 0.15 in Column 2). Part of the reason moving in with family may do a good job at smoothing consumption is related to the result in Column 3, where health expenditures is the dependent variable. Moving in with a family member has a very large and significant negative impact on health expenditures, suggesting that this can reduce some of the costs associated with going to hospitals or receiving treatment. If a family member can provide care for an individual who has severe physical limitations, there is potentially less necessity for a longer hospital stay or other medical services. Having this means of informal insurance seems to be an important element for individuals coping with severe physical limitations.

6 Robustness and Future Work

There are several potential robustness checks that can strengthen our results. First, along with the health limitations questions used throughout, the PSID also asks a

series of questions similar to the typical ADL (Activities of Daily Living) battery. The questions ask about very specific daily activities - dressing oneself, getting in or out of a chair, eating, walking, bathing, getting outside, and using the toilet - and whether these activities cause problems for the head of household. While these questions ask about similar concepts to the physical limitations questions used in the previous results, the ADL questions are not tied to any specific health condition. Additionally, the ADL questions are used in a variety of data sets, with claims that ADLs are a reliable measure of physical functioning that help to distinguish serious health problems. For this reason, it will be a good check that the results are consistent using this different measure for physical limitations.

Second, since many of the outcomes we are looking at either are binary (whether or not household received transfer from the family, whether they moved in with a family member, etc.) or have clustering at zero for continuous variables (health expenditures, hours worked), it makes sense to apply a limited dependent variable model to check if the results hold under alternative specifications. While a limited dependent variable model has some benefits over a basic linear model, there are also some strong additional assumptions necessary for the models to be valid. However, if the results presented thus far are robust and valid, we would expect to see similar estimates using a logit or tobit model.

Third, a primary concern with our identification strategy would be if the changes in physical limitations we measure were endogenously expected events, where the individual could mitigate the effects of the change beforehand. While this would bias our results towards the null which we clearly reject, the possibility still deserves consideration. One way to focus on more unanticipated health shocks is to look for large jumps in physical limitation status from one interview to the next i.e. restricting our identification to be from transitions from no limitations to a lot. Large jumps may better represent an unanticipated change in health status. Another way to isolate unanticipated changes is to use demographic and consumption decisions at $t-1$ to predict health status at wave t. We then can use the unexplained error term as the unanticipated portion of their health status. The panel dimension of the survey can also be exploited to test whether health in $t+1$ effects consumption in wave t.

Lastly, due to the PSID's sample design, we can use the health information from interviews with family members of the respondent. Since many health conditions have some genetic component, this information is likely a strong predictor for the respondent's health. For instance, if the parents of a respondent have a history of high blood pressure or experience a negative health event themselves, we can use that

information as a predictor for whether the respondent has severe physical limitations due to high blood pressure. For the same reasons that physicians rely on family health history, this same information can be used as part of a first stage estimation.

The concern about this strategy is the exogeneity of using family health history as a predictor for personal health; individuals may undertake mitigating behaviors such as taking medication, altering diet, or exercise regimes knowing their own health history. Health conditions are also not only a function of genetic inheritance, but a complex series of interactions between genetic influences and environmental factors. While asthma can be inherited, perhaps the reason we see both a parent and child experience asthma is due to their location choice where there is significant amount of smog or pollution. These risk factors for health conditions are extremely difficult to parse out as the medical literature has shown. However, it is an interesting starting point and unique application of family health history data as instruments for current health conditions. We will use this estimation strategy as a robustness check for the results we have reported thus far.

Lastly, we also need to take into the account the potential for mortality playing a role in the results. If people without adequate support are more likely to die in response to a negative shock in health, then we may have population estimates that are upward biased. One issue is that the PSID does not report a respondent having passed away in the public use data. We can further examine how health conditions may lead one to attrit from the sample. This is left for future work.

7 Conclusion

This paper contributes new results to a variety of literatures relating health outcomes, labor market experiences, formal and informal insurance, and family economics. The results are informative for understanding patterns of resource allocation within families, as well as how informal networks operate during adverse events. We present results on how specific health conditions affect labor market outcomes, and shed light on the interplay between formal health insurance and the family's informal transfers.

Our results suggest that households are unable to fully insure against deteriorating health, but undertake a number of partial insurance mechanism in the face of health shocks. These include compensating labor supply changes from family members, monetary transfers, and residency decisions. Yet despite this evidence, genetically linked families do not efficiently share risk within their networks.

These results offer a number of possible avenues for further examination, and highlight the importance of family networks in a developed country context.

8 Tables

Table 1: Summary Statistics

	Mean of Variables of Interest by Reported Physical Limitation Due to Health Condition				
		Respondents with Health Condition			
	Full Sample	**No Current Limitation**	**Little Limitation**	**Somewhat Limited**	**A lot of Limitation**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Age	45.71	51.88	53.75	56.98	60.58
Years of Education	13.10	13.09	12.47	12.09	11.51
Married (%)	56.80	55.74	46.33	44.59	36.25
Household Size	2.75	2.58	2.48	2.41	2.31
Male (%)	46.53	46.02	40.84	38.59	40.47
Retired (%)	13.05	20.81	25.46	31.59	35.33
Non-Health Expenditures	28864	29180	24786	22270	17644
Health Expenditures	2401	2866	2889	2773	3272
Received Transfer from Family (%)	9.14	8.37	9.95	9.77	11.72
Value of Family Transfer Received (Non-Zero Average)	2402	2574	2314	2139	2300
Number of Individual-Year Observations	37308	13232	5769	4340	3054

Table 2: Changes in Consumption in Response to Physical Limitations Due to Health Condition

Degree of Limitation on Household Head's Daily Physical Activities due to Health Condition:	*Dependent Variable: Log of Household […]*			
	Non-Health Expenditures		**Per Capita Non-Health Expenditures**	
	(1)	**(2)**	**(3)**	**(4)**
A lot	-0.10***	-0.055***	-0.10***	-0.050***
	(0.016)	(0.015)	(0.016)	(0.016)
Somewhat	-0.058***	-0.023**	-0.052***	-0.022**
	(0.012)	(0.011)	(0.012)	(0.011)
Little	-0.030***	0.0032	-0.029***	-0.00080
	(0.010)	(0.0087)	(0.010)	(0.0086)
No Current Limitations	0.011	-0.0036	0.011	-0.0034
	(0.0084)	(0.0069)	(0.0084)	(0.0070)
Never Any Limitations	-	-	-	-
Mean of Dependent Variable	10.00	10.00	9.14	9.14
Individual Fixed Effects	N	Y	N	Y

Note: All standard errors are clustered at the Family level. Included controls described in text.
37308 individual-year observations; 8212 extended families
*** p<0.01, ** p<0.05, * p<0.1

Table 3: Heterogeneity in Consumption Response to Physical Limitation

		Dependent Variable: Log of Household […]	
		Non-Health Expenditures	**Per Capita Non-Health Expenditures**
		(1)	**(2)**
Panel A			
"A lot" of physical limitation for household head interacted with..	Household head Younger than 50	-0.079***	-0.077***
		(0.026)	(0.026)
	Household head Older than 50	-0.037**	-0.030*
		(0.018)	(0.018)
Panel B			
"A lot" of physical limitation for household head interacted with..	Household head Works	-0.055**	-0.053**
		(0.023)	(0.023)
	Household head Does not Work	-0.050***	-0.043**
		(0.019)	(0.019)
	Mean of Dep. Var.	10.00	9.14
	Individual Fixed Effects	Y	Y

Note: All standard errors are clustered at the Family level. Included controls described in text.
37308 individual-year observations; 8212 extended families
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Difficulties from Physical Limitations

	Dependent Variable:		
	Household Head's […]		
	Health Expenditure	**Labor Income**	**Hours Worked Past Year**
	(1)	**(2)**	**(3)**
"A lot" of limitations for household head	530***	-1,529***	-145***
	(179)	(542)	(17.2)
Mean of Dep. Var.	2400	28667	1484
Individual Fixed Effects	Y	Y	Y

Note: Standard errors are clustered at the Family level. Included controls described in text.

37308 individual-year observations; 8212 extended families

*** p<0.01, ** p<0.05, * p<0.1

Table 5: Responses for Dealing with Physical Limitations

Panel A

	Household Wealth	Live in Nursing Home
	(1)	(2)
"A lot" of limitations for household head	-62,041***	0.018***
	(21,735)	(0.0059)
Mean of Dep Var.	242404	0.022
Individual FE	Y	Y

Panel B - Specific Household Assets

	Dependent Variable - Value of […]					
	Savings	**Debt**	**Stocks**	**Automobiles**	**IRAs**	**Home Equity**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
"A lot" of limitations for household head	-3,492	1,179*	-1,442	-697**	1,186	-11,750***
	(2,493)	(708)	(11,422)	(328)	(1,878)	(1,982)
Mean of Dep Var.	18847	8447	35859	13132	26947	68528
Individual FE	Y	Y	Y	Y	Y	Y

Note: All standard errors are clustered at the Family level. Included controls described in text.
37308 individual-year observations; 8212 extended families
*** p<0.01, ** p<0.05, * p<0.1

Table 6: The Family as Insurance Against Physical Limitations

	Dependent Variable			
	Hours Worked by Spouse of Head for the Year	**Y/N Did Head Receive Transfer from Family?**	**Amount of Transfer Received from Family**	**Y/N Did Respondent Move into House with a Family Member?**
	(1)	**(2)**	**(3)**	**(4)**
"A lot" of limitations for household head	42.7***	0.040***	102**	0.0092**
	(12.1)	(0.0082)	(47.9)	(0.0045)
Mean of Dep. Var.	878	0.091	220	0.037
Individual Fixed Effects	Y	Y	Y	Y

Note: All standard errors are clustered at the Family level. Included household controls described in text.

37308 individual-year observations; 8212 extended families
*** p<0.01, ** p<0.05, * p<0.1

Table 7: Does the Family Fully Share Risk?

	Dependent Variable: Log of Household […]	
	Non-Health Expenditures	**Per Capita Non-Health Expenditures**
	(1)	**(2)**
"A lot" of limitations for household head	-0.033*	-0.036**
	(0.018)	(0.018)
Individual Fixed Effects?	Y	Y
Family-Year Fixed Effects?	Y	Y

Note: All standard errors are clustered at the Family level. Included controls described in text.

37308 individual-year observations; 8212 extended families

*** p<0.01, ** p<0.05, * p<0.1

Table 8: Comparison of Insurance Channels in Smoothing Consumption

	Dependent Variable: Annual Household […]		
	Log Expenditure	**Log Per Capita Expenditures**	**Health Expenditures**
	(1)	**(2)**	**(3)**
Dummy for "A lot" of physical	-0.13***	-0.14***	453*
limitation for head of household	(0.048)	(0.050)	(272)
I(Spouse of Head Works)			
Interaction with Physical Limitation	0.080***	0.074***	-176
	(0.027)	(0.027)	(271)
Main Effect	-0.022	-0.032*	-111
	(0.018)	(0.018)	(131)
I(Received Transfer from Family)			
Interaction with Physical Limitation	-0.023	-0.027	-144
	(0.043)	(0.044)	(352)
Main Effect	-0.032**	-0.036**	-118
	(0.014)	(0.014)	(73.7)
I(Wealth Drawn Down from Previous Year)			
Interaction with Physical Limitation	-0.016	-0.017	-128
	(0.020)	(0.021)	(290)
Main Effect	-0.013***	-0.016***	-1.06
	(0.0051)	(0.0051)	(41.5)
I(Moved in with Family)			
Interaction with Physical Limitation	0.11**	0.15***	-1,222***
	(0.051)	(0.050)	(364)
Main Effect	-0.024	-0.060**	14.9
	(0.024)	(0.023)	(289)
I(Have Health Insurance)			
Interaction with Physical Limitation	0.066	0.075	318
	(0.049)	(0.051)	(280)
Main Effect	0.027**	0.026**	528***
	(0.012)	(0.012)	(67.5)
Individual/Household Fixed Effects?	Y	Y	Y

Note: All standard errors are clustered at the Family level. Included household controls described in text.
37308 individual-year observations; 8212 extended families
*** p<0.01, ** p<0.05, * p<0.1

References

Andreski, P., McGonagle, K., and Schoeni, R. (2009), "An analysis of the quality of the health data in the Panel Study of Income Dynamics," *Ann Arbor: University of Michigan*, pp. 09–02.

Asfaw, A. and Braun, J. (2004), "Is Consumption Insured against Illness? Evidence on Vulnerability of Households to Health Shocks in Rural Ethiopia," *Economic Development and Cultural Change*, 53, 115–129.

Bourguignon, F., Browning, M., Chiappori, P.-A., and Lechene, V. (1993), "Intra Household Allocation of Consumption: A Model and Some Evidence from French Data," *Annals of Economics and Statistics / Annales d'conomie et de Statistique*, pp. 137–156.

Browning, M. and Crossley, T. F. (2009), "Shocks, Stocks, and Socks: Smoothing Consumption over a Temporary Income Loss," *Journal of the European Economic Association*, 7, 1169–1192.

Chiappori, P.-A. (1988), "Rational household labor supply," *Econometrica: Journal of the Econometric Society*, pp. 63–90.

Chiappori, P.-A. (1992), "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100, 437–467.

Coate, P. (2012), "Parental Influence on Labor Market Outcomes of Young Workers," Working Paper.

Coate, P., Dalton, M., Hotz, V. J., and Thomas, D. (2013), "The Family that Shares is the Family that Cares: Are Extended Families Efficient in their Sharing?"," .

Cochrane, J. H. (1991), "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99, 957–976.

Cox, D. (2003), "Private transfers within the family: mothers, fathers, sons and daughters," *Death and dollars: The role of gifts and bequests in America*, pp. 168–210.

De Weerdt, J. and Dercon, S. (2006), "Risk-sharing networks and insurance against illness," *Journal of Development Economics*, 81, 337–356.

Deaton, A. (1992), *Understanding consumption*, Clarendon Press Oxford.

Deaton, A. (1997), *The analysis of household surveys: a microeconomic approach to development policy*, Johns Hopkins University Press.

Dercon, S. and Krishnan, P. (2000), "In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia," *Journal of Political Economy*, 108, 688–727.

Dercon, S., De Weerdt, J., Bold, T., and Pankhurst, A. (2006), "Group-based funeral insurance in Ethiopia and Tanzania," *World Development*, 34, 685–703.

Dynarski, S., Gruber, J., Moffitt, R. A., and Burtless, G. (1997), "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity*, 1997, 229–303.

Fafchamps, M. and Lund, S. (2003), "Risk-sharing networks in rural Philippines," *Journal of Development Economics*, 71, 261–287.

Frankenberg, E., Smith, J. P., and Thomas, D. (2003), "Economic Shocks, Wealth, and Welfare," *The Journal of Human Resources*, 38, 280–321.

Frankenberg, E., Gillespie, T., Preston, S., Sikoki, B., and Thomas, D. (2011), "Mortality, The Family and The Indian Ocean Tsunami*," *The Economic Journal*, 121, F162–F182.

Genicot, G. and Ray, D. (2003), "Group Formation in Risk-Sharing Arrangements," *Review of Economic Studies*, 70, 87–113.

Genoni, M. E. (2012), "Health shocks and consumption smoothing: Evidence from Indonesia," *Economic Development and Cultural Change*, 60, 475–506.

Gertler, P. and Gruber, J. (2002), "Insuring Consumption Against Illness," *The American Economic Review*, 92, 51–70.

Hayashi, F., Altonji, J., and Kotlikoff, L. (1996), "Risk-Sharing between and within Families," *Econometrica*, 64, 261–294.

Hotz, V. J., McGarry, K., and Wiemers, E. (2010), "Living Arrangements of Mothers and Their Adult Children Over the Life Course," *Unpublished Manuscript*.

Kaplan, G. (2012), "Moving back home: Insurance against labor market risk," *Journal of Political Economy*, 120, 446–512.

Kazianga, H. and Udry, C. (2006), "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso," *Journal of Development Economics*, 79, 413–446.

Kinnan, C. and Townsend, R. (2012), "Kinship and Financial Networks, Formal Financial Access, and Risk Reduction," *American Economic Review*, 102, 289–93.

LaFave, D. and Thomas, D. (January 2013), "Extended Families and Child Well-Being," Working Paper.

Mazzocco, M. and Saini, S. (2012), "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, 102, 428–68.

McElroy, M. (1985), "The joint determination of household membership and market work: The case of young men," *Journal of Labor Economics*, pp. 293–316.

McGarry, K. and Schoeni, R. (1995), "Transfer behavior in the health and retirement study: Measurement and the redistribution of resources within the family," *Journal of Human Resources*, pp. 184–226.

Schulhofer-Wohl, S. (2011), "Heterogeneity and Tests of Risk Sharing," *Journal of Political Economy*, 119, 925–958.

Smith, J. P. (1999), "Healthy Bodies and Thick Wallets: The Dual Relation between Health and Economic Status," *Journal of Economic Perspectives*, 13, 145–166.

Strauss, J. and Thomas, D. (1998), "Health, nutrition, and economic development," *Journal of Economic Literature*, pp. 766–817.

Strauss, J. and Thomas, D. (2007), *Chapter 54 Health over the Life Course*, vol. Volume 4, pp. 3375–3474, Elsevier.

Townsend, R. M. (1994), "Risk and Insurance in Village India," *Econometrica*, 62, 539–591.

Townsend, R. M. (1995), "Consumption insurance: An evaluation of risk-bearing systems in low-income economies," *The Journal of Economic Perspectives*, 9, 83–102.

Wagstaff, A. (2007), "The economic consequences of health shocks: Evidence from Vietnam," *Journal of Health Economics*, 26, 82–100.