

Insuring Consumption Against Illness

Paul Gertler

University of California at Berkeley and NBER

Jonathan Gruber

MIT and NBER

January 3, 2001

One of the most sizable and least predictable shocks to economic opportunities in developing countries is major illness, both in terms of medical care expenditures and lost income from reduced labor supply and productivity. As a result, families may not be able to insure their consumption over periods of illness. In this paper, we investigate the extent to which families are able to insure consumption against major illness using a unique panel data set from Indonesia that combines excellent measures of health status with consumption information. We focus on the effect of large exogenous changes in physical functioning. We find that there are significant economic costs associated with these illnesses, albeit more from income loss than from medical expenditures. We also find that there is very imperfect insurance of consumption over these episodes of illness, and that in particular families are less able to insure the costs of rare and severe illnesses and more able to insure the costs of more frequent and less severe illnesses. Indeed, we find that families are able to insure less than 40 percent of the income loss from illnesses that severely limit physical function. These estimates suggest that the large public subsidies for medical care typical of most developing countries may improve welfare by providing consumption insurance.

We are grateful to two anonymous referees, Anne Case, Don Cox, Angus Deaton, Alain de Janvry, Jeff Hammer, Levis Kochar, Jack Molyneaux, Jonathan Morduch, Christina Paxson, Mark Rosenzweig, Elizabeth Sadoulat, Ken Wolpin, and from seminar participants at Harvard, University of Oslo, University of Pennsylvania, Princeton, UC Berkeley, University of Washington, and the NBER. Nga Vuong and My Vuong provided excellent research assistance. Gertler gratefully acknowledges financial support from the National Institute of Aging. Gruber gratefully acknowledges financial support from the National Institute of Aging and the National Science Foundation. Both thank the Indonesian Government for supporting the data collection.

One of the most sizable and least predictable shocks to the economic opportunities of families is major illness. There are two important economic costs associated with illness: the cost of the medical care used to diagnose and treat the illness, and the loss in income associated with reduced labor supply and productivity. The size and unpredictability of both of these costs suggests that families may not be able to insure their consumption over periods of major illness, especially in developing countries where few individuals are covered by formal health and disability insurance (World Bank, 1993 and 1995a). The possibility that there is less than full consumption insurance suggests a potentially large loss in welfare from an illness shock to the household's resources. Recognizing this, many developing countries have or are considering social insurance to help insure the economic costs of illness.

While families with sick members in developing countries are not able to access formal insurance markets, they do rely on private informal coping mechanisms such as drawing on savings, selling assets, transfers from their families and social support networks, and borrowing from local credit markets.¹ In this case, there may be relatively little welfare gain from social insurance, as such insurance would serve largely to “crowd out” these other sources of insuring.

This conclusion appears to be supported by the existing empirical evidence in developing countries. Townsend (1994) finds that the percentage of the year that an adult male is sick has no impact on consumption. Kochar (1995) models wage income and informal borrowing as a function of illness in the family, as measured by a member of the family experiencing a loss of work due to illness. She finds that illness to the male lowers wage income and increases informal borrowing during peak periods in the agricultural cycle, but that there are no effects during slack periods and no effects of female illnesses. These studies appear to indicate that families living in low-income countries are able to insure illness shocks fairly well.²

A key limitation of past work, however, is that the measures of health employed may reflect only small,

¹ See Morduch (1995 and 1999) and Townsend (1995) for reviews of the consumption insuring literature.

² In contrast, Cochrane (1991) finds that consumption in the U.S. is sensitive to major illness, defined as being ill for more than 100 days.

and even potentially anticipated, changes in health status, not the kind of large unexpected major illnesses that may be difficult to insure. Even if families are able to insure illness shocks on average, they may be able to more effectively insure the frequent small illness shocks as opposed to the large rare shocks. This would imply that it is important to use measures of more severe illness in fully assessing the ability of households to insure consumption.

If families are able to insure the costs of minor illness, but not major illness, then optimal social insurance interventions should use their limited revenues to insure catastrophic events. This would suggest that limited user fees do not reduce the insurance value of public health care systems. It also would suggest that copayments and deductibles are in order for social insurance plans, as opposed to the current opposite trend in developing countries of first dollar coverage with low capped benefits (Gertler, 1998 and Gertler and Solon, 1999).

To address these issues we estimate the ability of families in Indonesia to insure consumption over periods of major illness using a unique panel data set that contains excellent measures of health status combined with data on consumption. Specifically, we investigate the effect of major illness on earnings and medical care expenditures, and the extent to which these costs of illness are financed out of consumption. We address the potential shortcomings of previous work in this area by using measures of individuals' physical abilities to perform activities of daily living (ADLs). ADLs have been proven reliable and valid measures of physical functioning ability in both developed and developing countries, and distinguish the type of serious exogenous health problems that are likely to be correlated with changes in labor market and consumption opportunities (Stewart et al., 1990; Strauss et al., 1993).

Our analysis proceeds in five steps. First, in Part I, we present the institutional setting and an overview of our rich data source. In Part II, we describe the risk of ill-health using measures that capture increasing degrees of severity. In Part III, we document that severe illness has dramatic implications for family resources by reducing the labor supply and earnings of household heads, and (to a much lesser extent) by increasing medical spending. In Part IV, we specify and implement a reduced-form model of families' abilities to insure consumption

against illness. We find that while families are able to fully insure minor illness measured by the types of indicators such as those used in previous work, they are not able to insure illnesses that limit their ability to physically perform activities of daily living. Moreover, our results do not appear to be driven by state-dependence? i.e. changes in tastes associated with illness.

Fifth, we then combine the results from Parts III and IV to assess the magnitude of consumption insuring; for each rupiah (Rp. 2000 ~ \$1) of income loss due to illness, how much does consumption fall? We find that the ability of families to insure falls dramatically with the severity of the illness shock. Indeed, families are able to insure less than 40% of the income loss from illnesses that are associated with a very severe loss in functioning. Finally, Part VI concludes by considering the welfare implications of our findings.

Part I: Institutional Setting and Data Source

The Setting

Indonesia is the fourth most populous country in the world with tremendous cultural and economic diversity. Until recently, economic growth has been impressive with an average real per annual capita growth rate of 3.9% over the last 15 years. Despite this, per capita incomes were still low even before the recent crisis, at \$US 880 per year in 1996 (Asian Development Bank, 1997). Indonesia had also seen remarkable improvements in health status (World Bank, 1993). Between 1960 and 1990 life expectancy at birth increased by 24 percent to 59 years and child mortality decreased 68 percent to 111 per thousand.

Indonesia has invested heavily to develop a comprehensive government-operated health care delivery system that individuals are able to access by paying a modest user fee. In 1991, there were at least one health center and several sub-centers in each of Indonesia's 3400 sub-districts. A network of government-operated hospitals at the district, provincial and central levels backs up this large primary care system. Despite this, Indonesia's health care expenditures remain low relative to those of its neighbors (World Bank, 1993). In 1990, annual expenditures on health care from both public and private sources were only about \$12 per person, which

amounts to about 2 percent of GDP.

Few individuals in Indonesia are covered by health insurance other than the implicit insurance provided through the almost free public health care system. On average, user fees at public facilities amount to 5% of costs (World Bank, 1995b). While the public health care system provides extensive primary care services, its hospital care is more limited. Moreover, many individuals opt to pay out of pocket for higher quality private sector services as over half of all utilization is provided by the private sector (Gertler and Molyneaux, 1996). About 10 percent of the population is covered by health insurance provided to civil servants. However, this insurance only covers utilization at public facilities. Therefore, the benefit to the individual is to only cover the small user fees. An additional four percent of the population is covered by health insurance offered through employers, but this insurance typically has capped benefits, minimizing absenteeism from minor illnesses but not paying the costs of major illness (Dow and Gertler, 1997). Similarly, there is limited disability insurance because there is no government program, over two-thirds of workers are self-employed, and few firms provide extensive sick leave.

Data Source and Sample

The data used in our analyses, collected as part of the Indonesian Resource Mobilization Study (IRMS), come from a panel survey of households designed to evaluate an experimental increase in user fees charged at public medical care facilities in two of Indonesia's 27 provinces. The two study provinces are West Nusa Tenggara (NTB), which is comprised of the two Islands just east of Bali, and East Kalimantan (KalTim), which is located on the east coast of the Island of Borneo. Together they account for about six million residents. KalTim has the third highest per capita income among all 27 provinces, while NTB is ranked twenty-second.

The data were collected in 1991 and 1993, allowing us to examine health, income, and consumption changes over a two-year period.³ The data are from a stratified random sample of households, consisting of

³ See Molyneaux et al. (1995) for a complete description of survey methods and results, and Dow et al. (1996) for an analysis of potential attrition bias in the data.

about 20 households per village (enumeration area) from 216 villages. The data were collected for each household at the same point in the year in both waves, so that we condition out seasonality effects in our fixed-effects models. The response rate in the first round of the survey was relatively high at 91%, and attrition between the first and second rounds was low at 7%. The complete sample is a panel of 4,889 households. Our sample of 3933 consists of all households whose who were in the survey in both rounds, whose first round head was in the sample in the second round, and who have non-missing and non-proxy data on the health measures described below.

The IRMS questionnaire was developed based on detailed roster, labor supply, consumption and health modules from existing surveys and were augmented when found to be incomplete. The survey team used focus groups and extensive pilot testing to ensure that the questions fit into the cultural context. In order to minimize measurement error, every adult in the household was interviewed directly as opposed to interviewing one individual and using proxy responses for the rest. We discuss the measurement of specific variables below.

Part II: The Risk of Illness

The key to our analysis is that we have unusually good measures of the change in the health status of household members. We explore the effects of two types of health measures: self-reported illness symptoms and limitations in the physical ability to perform activities of daily living (ADLs). Self-reported illness symptoms are similar to the measures used by the previous literature. We measure the presence of an illness symptom by a dummy for whether the individual reports any symptom (ill), and a dummy for whether they report a symptom that has lasted more than one month (chronically ill). This measure aggregates the 10 categories of self-reported specific symptoms (e.g. fever, respiratory congestion, etc.) for adults.

There are three important problems with these measures. First, as noted above, these symptoms may not represent the type of major health changes that impinge on consumption decisions. Second, as highlighted by both Strauss and Thomas (1996) and Bound (1991), using these self-reported symptoms may overstate the effect

of health status on labor supply because individuals who have left their jobs for other reasons may justify this decision by reporting a deterioration in health. To the extent that the labor force transition was planned, families may have already accounted for it in their consumption decisions at period $t-1$, so that there is no effect on the change in consumption, resulting in spurious evidence of consumption insuring. Third, there is substantial evidence that wealthier and more educated individuals have different definitions of illness as these types of individuals are more likely to report having an illness symptom in the last month (e.g. Sindelar and Thomas, 1991; Schultz and Tansel, 1997).

As an alternative we use a measure that assesses an individual's physical ability to perform activities of daily living (ADLs). These physical functioning measures are based on individuals' self-ratings of ability to engage in specific activities, not based on general assessments of illness symptoms, which are more likely to be endogenous to labor supply decisions. Initially developed for studying levels of disability among the elderly, these measures are used increasingly to study the health status of all adults. Physical functioning measures have been tested extensively for reliability (consistency between tests and interviewers) and validity (consistency between individual assessments of different skills). In the United States and South East Asia, they have been found to be reliable and valid self-assessments with a high degree of internal consistency (Andrews et al. 1986; Guralnik et al. 1989; Ju and Jones 1989; Strauss, et al., 1993; Ware, Davies-Avery, and Brook 1980). They are routinely used in studies of labor supply in the US (e.g. Bound, 1991; Bound et al., 1995; Stern 1989), and are the key measures of health status in the new Health and Retirement Survey (Wallace and Herzog, 1995). In addition, in contrast to self-reported symptoms, these measures tend to be negatively correlated with income and education in both US and low-income samples (e.g. Strauss et al., 1993; Kington and Smith, 1996; Gertler and Zeitlin, 1996).⁴

The specific ADL questions in the IRMS survey were adapted from the RAND 32 Question Short Form, which was developed using clinical outcome data and has been found to provide reliable and valid self-

⁴ADLs have been used in a number of studies of the relationship between health and labor market outcomes. See Strauss and Thomas (1996) and Bound (1991) for reviews of the developing country and US applications, respectively.

assessments with a high degree of internal consistency (Ware, Davies-Avery, and Brook 1980). The IRMS team used focus groups and extensive pilot testing to ensure that questions were modified to fit into the local cultural context. In order to minimize measurement error, every adult in the household was interviewed directly and proxy responses were not accepted.

Activities of daily living (ADLs) are divided into two categories. Intermediate ADLs consist of ability to carry a heavy load for 20 meters; sweep the floor or yard; walk for 5 kilometers; take water from a well; and bend, kneel, or stoop. Basic ADLs consist of ability to bathe yourself; feed yourself; clothe yourself; stand from sitting in a chair; go to the toilet; and rise from sitting on the floor. A limitation in any of these activities, particularly basic ADLs, clearly represents a major change in health status.

The responses to these questions on the survey are coded either as can do it easily (a value of 1), can do it with difficulty (2), and unable to do it (3). The responses to these questions were then combined in accordance with the following algorithm developed for the RAND Medical Outcome Study (Stewart et al., 1990):

$$Health = \frac{Score - Min\ Score}{Max\ Score - Min\ Score}$$

The ADL index takes on a value of 1 if the individual can perform all ADLs without difficulty, and zero if the individual cannot perform any ADLs. Our central model uses an overall index of ADL limitations. We also present results below for a disaggregation of this ADL index into both its intermediate and basic components.

The means and standard deviations of the health outcome measures for the head of household are presented in Table 1. The left-hand panel shows the means for period 1 levels, while the right hand panel shows them for changes from period 1 to period 2. In period 1, a large proportion of adults, 30 percent, reported some ADL limitation. In addition, there is substantial change in health status over time. Between 1991 and 1993, 11% of the sample reported a deterioration of their ADLs, and 23% reported an improvement. On average, the size of the deteriorations and improvements were comparable. Decomposing our overall ADL index into its basic and

intermediate components, we see that the basic ADL index is very close to one, reflecting the rare nature of these very serious health changes. Indeed, less than 4% of our sample in period 1 has a basic ADL limitation. Intermediate ADLs are more common with almost 30% of the sample having one in period 1.

Despite their severity, changes in ADLs do not appear to be permanent on average. There are actually more improvements than deteriorations in intermediate ADLs, and the figures are roughly comparable for basic ADLs. This reflects an important difference in the interpretation of ADLs, and in particular basic limitations, in developed and developing country contexts. In wealthier and more developed countries such as the US, limitations in the ability to feed oneself, bathe and toilet indicate a severe incapacitation that would make one close to bed-ridden and may reflect long-term disability. However, in a developing country setting such as Indonesia, performing basic physical activities requires more ability than in developed countries. For example, bathing in Indonesia generally requires going to the river and bathing using a sarong (large tubular like fabric) to maintain modesty. This requires much more effort and coordination than bathing in one's house. Also, toileting requires the use of eastern as opposed to western toilets, which are often located outside the main living quarters. Hence, basic indicators capture less severe limitations in developing country settings better than in developed country settings. As a result, it is not surprising that as many people recover from basic limitations as develop them, suggesting that we are indeed measuring severe temporary changes in health as opposed to permanent deterioration. This is confirmed by the fact that we find similar impacts when we restrict our estimates to downward movements in ADLs only.

Turning to the self-reported symptoms, more than half the sample reported an illness symptom last month in the first survey round. This raises questions about the usefulness of this indicator for investigating consumption insuring as its huge frequency indicates that it is picking up many minor health problems that do not need expensive medical care or affect labor supply. However, a much smaller share of individuals report chronic symptoms lasting more than one month. While there is some reduction in symptoms across these two years, there is a very large increase in chronic symptoms that may be expected to some extent as this cohort ages.

Part III: Illness and the Reduction in Family Resources

A prerequisite for there to be an effect of illness on consumption through imperfect consumption insurance is that there must be a sizable cost of illness. In this section, we quantify the cost of illness to the head of household in terms of reduced labor supply, lost earnings and increased medical spending. We choose the head of household because he or she is the major source of household income. We will consider the health of other household members later.

Model Specification

We estimate labor supply, earnings and medical care spending equations using the following fixed effects specification:

$$\Delta L_{ij} = \alpha_j \Delta h_{ij} + \sum_k \beta_k X_{ijk} + \epsilon_{ij} \quad (1)$$

which is a regression of the change in labor supply or earnings or medical care spending ΔL_{ij} for individual i in community j against community fixed effects (α_j), the change in health Δh_{ij} , a series of demographic controls (X_{ijk}), and a random error (ϵ_{ij}).

Equation (1) regresses first differenced labor outcomes and medical care spending against the change in health. We include a full set of community dummies to control for aggregate determinants of labor supply or medical spending.⁵ We also include demographic controls to capture other secular trends in the labor supply of household heads: the head's sex, age, education, and marital status, the wife's age and education, and the change in log family size. To measure a change in the indicator variables for symptoms, we define a variable which is 0 if there is no change, 1 if the person moves from ill to healthy, and -1 if the person moves from healthy to ill. The

⁵ Communities for our purposes as defined as IRMS "enumeration areas", which are village sampling clusters.

change for ADLs is simply the change in the ADL index value.

The model is a fixed effects specification, and as such, controls for unobserved heterogeneity. In particular, the first differencing sweeps out correlation from omitted unobserved individual characteristics (such as preferences and health endowments) that confound identifying the effect of illness on labor market outcomes. However, there may be unobserved correlates of changes in income and changes in health outcomes that confound identification. We control for one major source of spurious correlation, shocks to the local community economy such as weather that affect both changes in permanent income and changes in health, by including a set of community fixed effects.

A related concern is idiosyncratic changes in household income that feedback into health; for example, job loss that results in a deterioration of health (perhaps through mental depression). However, our empirical results suggest that this alternative explanation does not account for our findings. In particular, we find that larger health shocks are associated with bigger income losses and larger consumption losses. Therefore, if our results reflect effects of labor supply on health, this feedback mechanism would have to operate more strongly the larger the negative income shock. This means, for example, that the effect on health from a job loss would be bigger for high wage individuals than low wage individuals. This type of feedback seems to us to be unlikely. In addition, our findings below that the effects of illness on consumption are differentially strong for families who are not well insured are inconsistent with an idiosyncratic feedback explanation.

Data and Measurement

Labor supply is measured in two ways: as the change in hours worked and as a dummy for participation in the labor force. Earnings and wages are only reported in the IRMS data for the one-third of heads who work in the formal labor market. Earnings from family farms and enterprises were not collected. We therefore impute wages to all workers based on market rates. This imputation proceeds by first taking an average of hourly market wages by province (NTB or KalTim), age (<25, 25-49, 50+), education (the four categories denoted at the bottom

of Table 1), and sex. This cell-specific average wage is then matched to all persons in the cell, regardless of whether or not they worked in the market.⁶ This imputed hourly wage is then multiplied by hours per week to get weekly earnings, and by 4.3 to get monthly earnings, in order to match our monthly consumption figures.

Our measure of earnings raises a potential problem in estimating the impact of illness on income. Specifically, illness can only affect earning through labor supply and not through the returns to labor. It is quite possible that illness reduces the returns to labor, especially for those working on family farms and enterprises, and in self-employment. Since we do not have data on earnings from these sectors, we are unable to account for this possibility, implying that we may have a downward biased estimate of the impact of illness on earnings.⁷ However, the extent of this bias is likely to be lower for more severe illnesses, especially those that reduce labor supply to zero.

The means and standard deviations of the labor supply and earnings variables are reported in Table 1. Earnings are measured in real per capita terms in order to match our consumption specification below.⁸ Among heads, average hours of work are almost 37 and over 80% of heads work in period 1. There is a slight reduction in labor supply between rounds, but only a very modest change in average earnings.

Spending on medical care is measured as the product of reported medical utilization and prices from the sites at which medical care was delivered, following Gertler and Molyneaux (1996). Descriptive statistics are reported in Table 1. Spending on medical care is quite low, averaging less than 1% of non-medical consumption. This reflects both low levels of utilization and the extensive subsidization of medical care costs by the public sector. Even conditional on having some positive spending, mean spending on medical care is only about 2% of non-medical consumption, although roughly 5% of sample households spend more than 10% of non-medical

⁶ The valuation of non-market work at the market wage is only appropriate if labor markets clear. This assumption is supported for Indonesia by Pitt and Rosenzweig (1986) and Benjamin (1992).

⁷ We did examine the effect of illness on wages for the sample of individuals working in the formal wage sector. This analysis produced a negative, but substantively and statistically insignificant, correlation between health and wages.

⁸ All figures are reported in 1991 urban NTB rupiah.

consumption.

Results

Table 2 reports the full regression specification for our first measure of labor supply—change in hours worked. For symptoms, there is a negative effect of becoming ill on hours of work, but neither coefficient is significant. The result suggests that having chronic symptoms is associated with a 1 hour per week reduction in labor supply.

The next column shows the results for ADL changes. Here, illness is represented by a reduction in the index, so that a positive coefficient indicates that illness reduces labor supply. There is a sizable and significant effect of ADL changes. The coefficient implies that moving from completely healthy (index=1) to completely sick (index=0) would lower hours of work by almost 31 hours per week, which is a fall of 84 percent of baseline hours worked. In other words, if the head moved from able to perform all of the ADLs to unable to perform one ADL, his hours of work per week would fall by 2.8 hours (7.6 percent of baseline hours).⁹ The control variables are generally insignificant, except that the most educated heads are found to work more hours.

Table 3 presents the coefficients of interest for other measures of labor supply. The first row replicates the findings from Table 2. The next row shows the results for change in labor force participation. The finding parallels that of Table 2: positive effects of symptoms (becoming ill raises non-participation), and negative effects of ADL changes (improved physical functioning lowers non-participation). In this case, the effects are significant for all of our health status measures, although much stronger for ADLs. Indeed, moving from being able to perform all of the ADLs to being able to perform none implies a 74 percent likelihood of becoming a non-participant, while experiencing an illness symptom or chronic illness increases the likelihood of not participating by about 3 percent.¹⁰

⁹ Note that this evaluation can be done for all of the ADL coefficients by simply multiplying the coefficient by 0.0909, which is the change in the ADL index arising from a movement of being able to do all ADLs to being completely unable to do one ADL.

¹⁰ These results are consistent with Pitt and Rosenzweig (1986) who find significant effects of having an illness symptom in the head's labor supply and farm profits using a 1980 Indonesian data set.

The third row shows the effect on imputed earnings, expressed in 10,000 Rupiah units. Surprisingly, the effect of chronic symptoms on earnings is actually lower than for non-chronic symptoms, despite a larger effect on hours worked. This implies that the population for which chronic symptoms are associated with reduced work is a relatively low (predicted) wage population. While our imputation of wages by demographic characteristics limits the generality of this result, it is consistent with the notion that individuals who are marginally attached to the labor force justify their exit from the labor force by reporting a chronic symptom.

For ADL change, the coefficient is once again much stronger. It implies that moving from completely able to perform ADLs to completely unable to perform ADLs would lower earnings by R.20,170. This is roughly as large as baseline mean earnings, suggesting that such a shift would (unsurprisingly) leave the head with little earnings. Moving from completely able to perform all ADLs to being unable to perform one ADL lowers earnings by R.1834 or about 10% of baseline earnings.

The fourth row of Table 3 shows the effects of illness on medical spending. There are significant effects in the expected direction for all three measures (having symptoms = more spending, lower ADLs = more spending). But these effects are trivially small relative to the effects on earnings. This is not surprising since publicly provided care is heavily subsidized (i.e. user fees are well below the cost of care).

Effects of Head's Health on Other Family Members' Labor Supply

While the primary focus of this paper is to measure the effect of major illness on consumption and the loss in welfare from not being able to insure consumption, there could be effects on the labor supply (leisure) of other household members. Specifically, one of the major forms of informal insurance is for other family members to increase their labor supply when the head falls ill.

We repeat the above analyses using the change in earnings of other family members as the dependent variables and report the results in the final two rows of Table 3. The next to last row shows the effect of the change in the head's health on the earnings of all other family members. The results suggest that the labor supply

of other family members does not respond to a health shock to the head when health is measured as a change in ADLs or chronic illness symptoms. However, when the health shock to the head is measured by the change in illness symptoms, the labor supply of other family members significantly increases. The overall effect of the head experiencing an illness symptom on total household earnings is small and not significantly different from zero.

The last row of the table summarizes the findings for the head and others by using total family earnings. While health shocks to the head measured by illness and chronic illness symptoms do not significantly affect total household earnings, health shocks measured by ADLs do significantly reduce total household earnings. Consistent with little labor substitution, the coefficient on the head's ADLs in the total family income model is very close to the coefficient in the head's earnings model.

These results have several implications. First, neither illness nor chronic illness symptoms have any effect on total household earnings and, therefore, are unlikely to have any effect on consumption. Second, since the head's ADLs did not affect the labor supply of other family members, households are using methods other than labor substitution to insure consumption if at all.¹¹ Finally, since the head's ADLs have the same coefficient in the head's earnings model and in the total household earnings model, it will not matter whether we use the change in household earnings due to the change in the head's ADLs or the change in total family earnings in the consumption models.

Part IV: Health and Consumption ? Reduced Form Results

In the previous sections we demonstrated that major illnesses as measured by changes in basic ADLs are associated with large financial costs to households. In this section, we test whether households are insuring consumption against these unexpected costs of illness and investigate whether consumption falls because of state dependent utility or because families are unable to insure consumption. The answer to this second question is critical for a welfare interpretation of the main results reported in the following sections.

¹¹ This also implies that other family members' welfare is not being reduced through less leisure.

Empirical Specification

Our empirical specification is based on the theory of full insurance, which casts consumption insurance in terms of interhousehold risk sharing (e.g. Cochrane, 1991; Deaton, 1994; and Townsend, 1994). In practice, however, our empirical specification follows the previous developing country literature in examining consumption insurance, either through mutual insurance with others or through self-insurance (i.e. savings). Indeed, our empirical tests do not distinguish between these two channels for consumption insurance. The key empirical insight of the theory of full insurance is that the growth in each household's consumption will not depend on changes in household resources that are uncorrelated with shifts in preferences once the growth in community resources has been taken into account. Therefore, we can test whether families are able to insure consumption against illness by estimating the following equation:

$$\ln\left(\frac{C_{ij}}{n_{ij}}\right) = \alpha_j + \beta h_{ij} + \sum_k \gamma_k X_{ijk} + \epsilon_{ij} \quad (2)$$

which is a regression of the growth in log per capita (non-medical care) consumption for household i in community j against community fixed effects (α_j), the change in health (h_{ij}), a series of demographic controls (X_{ijk}), and a random error (ϵ_{ij}). If there is full insurance of illness, then there will be no effect of the change in health on the change in consumption, implying $\beta = 0$.¹²

A major assumption of the insurance interpretation, however, is that the utility function is separable in consumption and health, and in consumption and leisure. As a result, the marginal utility of consumption does not depend on the state of health directly, nor indirectly through induced changes in leisure. If this is not true, then even with full insurance the growth of consumption will vary with the state of health. That is, in the formulation above, h_{ij} will be correlated with omitted preferences and thereby with the error term, biasing the estimated

¹² We have 213 community dummies, with an average of 18 observations per community. These community dummies are jointly significant in our key consumption equation. If they are excluded from the basic equation, then the coefficient on health is reduced by roughly 20%, although it remains marginally significant.

coefficient γ in equation (2). An important feature of our empirical strategy is to test for such "state dependence" in consumption behavior. Indeed, this is major purpose of this section. As we discuss in detail later, we find no evidence that state dependence explains our empirical results.

Data and Measurement

The dependent variable for our analysis is the change in the log of monthly non-medical consumption per capita. The means for consumption expenditures are shown in Table 1. Like earnings, consumption is reported in real terms by deflating for price differences across locales and over time. Consumption was measured using a 30-minute 101-item questionnaire covering a comprehensive list of food and non-food consumption items. The questionnaire was developed based on Indonesia's 2-hour 300 plus item consumption questionnaire used in their budget-expenditure survey collected every three years as part of the 60,000 household National Socio-Economic Survey of Households (SUSENAS). Extensive pilot testing was used to aggregate the 300 plus item list to 101 items. The questionnaire collected both expenditures and the value of home produced consumption for each of the 101 items. The survey allows families to report their consumption in their own time units in order to reduce measurement error. The IRMS log per capita consumption distribution closely matched the same distribution generated from the longer SUSENAS.

The independent variables include the change in health measured by illness symptoms, chronic illness symptoms, and the ADL indices. In addition, we control for preference shifts associated with changes in family size or structure by including the change in log family size and a series of measures of the change in the share of the family that is male and female family members in age groups 0-5, 6-17, 18-49, and 50 plus. As noted above, we also control for other potential taste shifters that might be correlated with illness: the head's sex, age, education, and marital status, and the wife's age and education.

Results

Our estimates of equation (2) are presented in Table 4. For illness symptoms, we cannot reject the

hypothesis of full insurance. The coefficients on both measures are insignificant. Indeed, they have the wrong signs, indicating that illness is associated with higher levels of consumption, not lower.

In contrast, when we use the ADL index, we strongly reject full insurance. Changes in the ADL index have a significant and sizeable effect in the expected direction; negative increments to health are associated with reductions in consumption. Moving from being able to perform all ADLs to being able to perform none of them would lower consumption by almost 20%. A move from completely able to unable to perform one ADL would lower consumption by 1.8%.

The control variables show the expected pattern of effects. Consumption growth rates are higher for male heads, for older heads (although the effect increases with age at a diminishing rate), and for more educated heads. Per capita log consumption changes fall with the change in log family size, indicating some economies of scale in consumption. There is no clear pattern to the (unreported) coefficients on the changes in demographic shares, which are mostly insignificant.

These findings suggest two conclusions. First, traditional measures of illness change using illness symptoms, while weakly associated with labor supply changes, are not associated with consumption changes. There are two possible interpretations for this finding. On the one hand, there may be full consumption insurance against relatively minor health changes. Alternatively, as suggested by Bound (1991) and Thomas and Strauss (1996) these types of self-reported symptom measures may in fact be endogenous to labor supply decision-making.

Second, the more severe illnesses measured by ADL changes are very strongly associated with consumption changes. These latter types of illness changes appear to represent shocks to a family's opportunity set that cannot be insured.

State Dependence?

An alternative interpretation of our results is state dependence. That is to say that the types of serious

illnesses that drive our consumption findings may be associated with changes in underlying preferences of the household. It seems unlikely that state dependence could account for the very large family consumption effects that we find, given that we are measuring illness to the head only, and the average family size in our data is almost five. For example, if consumption is distributed equally across family members, a movement in intermediate ADLs from 1 to 0 would have to lower the head's consumption by roughly 100% to explain our result.

Of course, due to differential economies of scale, the head may account for more than 20% of family consumption, so that this discussion cannot refute state dependence as an explanation for our findings. Moreover, the illness of the head may be correlated with illness to other family members, so that the large percentage effect on family consumption reflects family-wide changes in tastes. This latter hypothesis is directly testable, by including in our model changes in the ADLs of other adult members of the household. In fact, when we do so, we find no effect of others' ADLs on consumption, nor any change in the coefficient on the head's ADLs. This is to be expected as there is a very low correlation between the illnesses of other family members and the illness of the head (correlation = 0.10), and the other family members individually contribute much less to family income. The impact of other family members' health shocks is discussed in more detail below. While the a priori facts are suggestive, in this section we propose three more rigorous tests to demonstrate that state dependence is not driving our findings.

Workers vs. Non-Workers

The assumption underlying our discussion is that illness affects consumption by lowering the earning potential of workers. On the other hand, state dependence through illness should affect workers and non-workers equally. Thus, a test of our view versus the alternative is to consider the effects of illness on workers relative to non-workers. There will be some effect on non-workers through medical spending, but as we demonstrated in Table 3 the effects of ADLs on medical spending are trivial relative to their effects on earnings. We define workers as those working at least 20 hours per week in the first survey year, and non-workers as those who

worked less than 20 hours¹³. For this test, we consider only downward movements in ADLs, since these changes will have no budgetary impact on non-workers but will lower income of workers; for upward movements, both groups could see increases in income, and indeed the coefficient on upwards movements are similar across the two samples.

We examine the separate effects on these two samples in the first row of Table 5. For workers, the impact of ADL changes are much larger than for the full sample (the coefficient on downward movements only for the full sample is 0.21). We find that moving from being able to perform all ADLs to being able to perform none lowers consumption by about 37%. But there is no impact at all on non-workers; the coefficient is actually wrong-signed and very close to zero.

This test provides some confirmation that the dependence of tastes for consumption on underlying health is not driving our results. Both workers and non-workers suffer from illness, but only workers experience the shock to their budget constraint implied by the illness. However, a problem with this test is that while it rules out state dependence through changes in health, it does not rule out state dependence through health-induced changes in leisure. That is, for the working sample but not for the non-working control group, hours of work change dramatically with illness. This could lead to lower consumption, for example, through lower work-related consumption purchases (ie. bus fare or new clothes). Once again, it seems unlikely that complementarities between work and consumption could explain the very sizeable effects that we see for total family consumption. But this suggests the value of additional tests.

Self-Insurance

Our second test for distinguishing state dependence is to assess how our effect varies with the ability of families to self-insure. Families that are better able to self-insure should see a smaller effect of illness on

¹³ The median age of non-workers is 46 compared to 42 for workers, and three-quarters of non-workers are less than age 60.

consumption. However, there is no reason why state dependence should be smaller for these well self-insured families. Thus, if the effect of illness is much larger for poorly self-insured families, it suggests that these effects are operating through the budget constraint, and not through state dependence. Our test therefore consists of including an interaction of illness with the indicator for ability to self-insure in equation 2. A negative interaction suggests that having self-insurance mitigates the effect of health shocks, which would be consistent with our hypothesis of imperfect insurance but inconsistent with state dependence.¹⁴

Our measure of self-insurance for this test is the labor supply of other family members. For families whose other members provide a large share of income in the first period, the effects of reduction in the labor supply of the head should have much less severe implications for family consumption opportunities. This means that as the share of the other members earnings gets larger, the net percentage change in family income from a reduction in the earnings of the head gets smaller. On the other hand, there is no reason to believe that state dependence operates more strongly in these types of families because there is no reason that changes in tastes for consumption when ill should depend on the share of earnings that is derived from the other members in the base period.

We investigate this proposition in the second column of Table 5. We measure ability to self-insure through family labor supply as the ratio of others' earnings to the earnings of the head in period 1¹⁵. A value of one indicates that others earn as much as the head in period 1, so that the family has substantial self-insurance against illness to the head. We then interact this measure with the change in the ADL index. We estimate this model for workers only, both because we have just demonstrated that shocks to health matter only for workers, and because we need a sample of workers to have an earnings denominator for the head.

In fact, there is a sizable and significant negative interaction, which is consistent with our interpretation of the ADL change as measuring a shock to the budget constraint. That is, the impact of health shocks is mitigated

¹⁴ Similar considerations of differential consumption insuring by the ability to self-insure are discussed in Morduch (1990).

through self-insurance, as measured by the share of family earnings from other members. Indeed, this interaction is roughly one-half as large as the coefficient on the ADL change itself, which is sensible. This indicates that if other family members earn the same amount as the head, then the illness-related fall in consumption is about half as large.

Other Family Members' Health

Our third test for distinguishing state dependence extends the results just presented to assess whether the health of other family members affects consumption in proportion to the amount they contribute to family income. If there is no state dependence, then the health of other adult family members should affect consumption only in proportion to their contribution to family earnings. On the other hand, if there is state dependence, so that our findings reflect changing tastes for consumption when ill, then the health of other family members should affect consumption regardless of the amount they work.

We test this hypothesis by adding to the regression just presented (a) the change in health of other family members and (b) this health change interacted with other family members' share of earnings in period one. Thus, as above, we should continue to find that health changes for the head matter less when other family members have a larger share of earnings. But now we should also find that health changes for other family members matter more when they have a larger share of earnings. That is, as other family members earn more in the base period, it should mitigate the consumption impacts of illness to the head (as we showed above to be true), but also exacerbate the consumption impacts of illness to others.

We investigate this proposition in the third row of Table 5. We find that the other family member ADL index is insignificant, while the interaction with their share of earnings is positive and significant. This suggests that the health of other family members only affects consumption in proportion to their contribution to income. This is clearly inconsistent with state dependence.

We also find that the interaction of the head's ADL index with the share of income contributed by other

¹⁵ We also tried using period 1 level earnings instead of the earnings ratio. The results were almost identical.

family members is negative and significant, suggesting that the effect of the head's health on consumption diminishes with the family's ability to self-insure. More specifically, the implied effects of the head's ADLs and other family members ADLs have the same effect if they contribute equal shares to earnings. When both contribute equally to earnings the ratio of others' earnings to head's earnings is 1. The implied effect of a head's ADL shock is $0.507 - 0.229 = 0.278$, and the implied effect of a shock to others is $-0.056 + 0.343 = 0.287$. This strongly suggests that health is just operating through the budget constraint, which is inconsistent with state dependence.

The fact that health shocks weighted by relative earnings affect consumption in the same way also means that our results generalize beyond illness to the head. Indeed, our estimates of the reduction in consumption from the head experiencing an illness are the same as the effect of an illness to any family member weighted by their earnings contribution. This means that our estimates of the family's ability to smooth the costs are robust across family members.

Thus, to summarize the results of the second and third panels, we find that health shocks to both the head and to others appear to matter in proportion to their contribution to family earnings. Taken together with the results for workers and non-workers, these findings strongly refute the contention of state dependence. In addition, they also serve to address concerns over reverse causality, with shocks to consumption opportunities driving both consumption and health (idiosyncratic health feedback). This alternative explanation is inconsistent with our differential effect on insured/uninsured heads. There is no reason why this reverse causality should operate differentially for one group and not another.

Part V: How Incomplete is Insurance?

We view our results so far as a convincing demonstration that major illness reduces consumption because households are unable to insure the economic costs of illness and not because of state dependence. In this section we estimate the extent to which households are able to insure consumption. This magnitude is critical for

assessing the importance of our findings for welfare and for considering their policy implications. We measure the extent to which households are not able to insure consumption against illness as the share of the costs of illness that are financed out of consumption. To do so, we estimate a model of the effect of changes in (net of medical spending) income on the growth of non-medical care consumption:

$$\ln\left(\frac{C_{ij}}{n_{ij}}\right) = \alpha_j + \beta Y_{ij} + \sum_k \gamma_k X_{ijk} + \epsilon_{ij}, \quad (3)$$

where Y_{ij} is earnings minus medical care expenditures. Then the share of the costs of illness that are financed out of reduced consumption is simply β/C_{ij} .¹⁶

The results of estimating equation (3) by OLS are shown in the first row of Table 6. We show only the coefficient of interest - that on change in income - from regressions that include all of the regressors shown in Table 4. Income is expressed in units of Rp. 10,000. We find that there is a significant, but very small, relationship between income changes and consumption changes. A Rp. 10,000 increase in income is estimated to increase consumption by only 0.7%, or Rp. 255. As we show in the second column, this implies that for each rupiah that income falls, consumption falls by only Rp. 0.03. This is a trivial change, which would indicate very close to full consumption insuring.

However, there are two potential problems with estimating equation (3) by OLS, both of which would bias towards a finding of consumption insuring. The first, as noted by Morduch (1990), is that the growth in income is correlated with the error term through the production process. Risk adverse families may choose the variation in income so that consumption can be insured using available mechanisms. The second is measurement error in the growth of income, particularly since we have imputed earnings in our data.

In order to solve these problems, we employ an instrumental variables approach, which uses the change in the illness variables to instrument for the change in income. This instrument is valid given three conditions: (i)

¹⁶ We use the level of income, instead of the log, since roughly one-quarter of cases where there is a change in the ADL index have zero earnings in one period, and we do not want to exclude these cases.

that the utility function is not state dependent, (ii) that there is no feedback from changes in consumption to changes in health, and (iii) that measurement error in health changes is uncorrelated with measurement error in income changes. The first two conditions are supported by the results in Table 5 and the last one seems reasonable. In this case, the regression in equation (3) allows us to assess whether the major changes in income due to illness are insured differently from average income changes.

Once we instrument income by the change in the head's ADLs in the second row of Table 6, the coefficient rises dramatically and becomes significant. The estimate indicates that for every Rp.10,000. of income lost due to illness, there is a fall in consumption of 9.6%, or Rp. 3490. That is, for each rupiah that income falls, consumption falls by Rp. 0.35 (as shown in the second column). This suggests that families are able to insure only 65% of the loss in income from ADL changes on average.

In the next row, we instrument for income using the ADL index of the head and the ADL index of all other adults in the household weighted by the share of income contributed to the household in period one. The coefficient is positive and significant. The coefficient is close and not significantly different from the same coefficient estimated just using the head's ADLs as the instrument. This is consistent with no state dependence since an illness shock affects consumption the same regardless of who becomes ill, holding the contribution to household income constant. As is clear from the small change in the coefficient, we easily pass an overidentification test of these joint instruments; the test statistic is 1.57, which is well below the critical value of 3.64 for a chi-squared with one degree of freedom.

Heterogeneity

The results thus far have aggregated all of the available ADL information into one index. While this provides a convenient summary measure, it masks underlying heterogeneity in the response of consumption to different types of health changes. In particular, households may be better able to insure the modest income losses associated with minor health changes than they are the larger income losses associated with major health changes.

This would arise if individuals have available limited self-insurance (e.g. savings) that can only cover small income losses, or if consumption insurance (e.g. transfers and loans from extended family) were available in small but not large amounts. This suggests that the model in equation 3 should be non-linear in change in income, with the effect on the change in consumption rising with the change in income.

We explore this issue using a local average treatment effects approach (Angrist and Imbens, 1994) by disaggregating our ADL index into changes in intermediate and changes in basic ADLs. As noted earlier, the latter set of limitations are much more serious; while 29% of our sample has some intermediate ADL limitation, only 3.6% has some basic ADL limitation. The coefficients in the last two rows of table 6 are estimates of the effect of income variation on consumption, where the instruments are not the overall ADL index but rather the separate indices for basic and intermediate ADLs. The IV coefficients in each case represent the impact of income variation induced by less and more severe changes in illness. In other words, they provide an assessment of whether there is a better ability to insure less severe health changes (intermediate ADL changes in the second last row) than more severe health changes (basic ADL changes in the last row).

Our results suggest that the impacts of basic ADL limitations on consumption are much more sizeable. We find that each Rp. 10,000 of income loss due to intermediate ADL changes causes consumption to drop by 7.9%, but that each Rp. 10,000 of income loss due to basic ADL changes causes consumption to drop by 17.1%. That is, we estimate that families can insure 71% of the income loss associated with intermediate ADL changes, but only 38% of the loss associated with basic ADL changes. This is a striking difference, and it suggests that there is definite heterogeneity in the ability of families to insure illness shocks.

Part VI: Conclusions and Policy Implications

Using reliable and valid measures of ill health that distinguish varying degrees of severity, we find that Indonesian households are not able to fully insure consumption against the economic costs of illness. We estimate that 35 percent of the costs of serious illness are not insured by other sources available to households. We also

find that the more severe the illness, the less households are able to insure. Households are able to fully insure the economic costs of illnesses that do not affect physical functioning, insure 71 percent of the costs resulting from illnesses that moderately limit an individual's ability to function physically, but only 38 percent of the costs from illnesses that severely limit physical functioning.

Our findings imply that there are non-trivial costs to the Indonesian economy from incomplete insurance of even these very extreme health events. We can measure the fall in consumption associated with illness as the aggregate reduction in family resources from illness, times the extent to which that reduction in resources is reflected in lower consumption. The former is derived from the probability distribution of illness (ADL deviations from perfect health) times the loss in earnings plus medical care expenditures associated with each level of illness. The latter simply comes from our estimates in Table 6, where we weight separately by their coefficients intermediate and basic health changes to reflect the heterogeneity in smoothing.

Doing so, we find that illness is associated with a fall in consumption of 0.84 percent of baseline. This is a non-trivial effect given the very low frequency of very serious health limitations. Moreover, this understates the total welfare cost of illness for at least two reasons. First, there are additional welfare costs from the uninsured variability in consumption, beyond the reduction in the level of consumption. Second, there are costs to others of the resources used to smooth consumption when family members become ill. For example, there is some cost to family and friends from private transfers of resources to the ill household head.

Our analysis has also suggested an additional rationale for subsidized medical care prices in developing countries: consumption insurance. While our results indicate that families are able to insure the costs of small frequent illnesses, they are unable to insure the costs of rare major illness. This suggests that there may be an important welfare cost to raising user fees at public hospitals in order to shift subsidies to primary and preventive cares (which are affordable). Governments considering raising hospital user fees must consider how to insure the medical care costs of large illnesses, for example through caps on fees for inpatient hospital stays, or by developing prepayment schemes in conjunction with reducing subsidies.

These findings also have implications for another major form of social insurance in developing countries, which is to finance medical care publicly through payroll taxes, but allow beneficiaries to purchase medical care from private providers (Gertler, 1998). Since low-income countries have limited abilities to tax, the resources available for SI are severely constrained. This results in a tradeoff between catastrophic coverage, with a high deductible but uncapped coverage, and first dollar coverage, with coverage of all expenses from the first dollar but a low cap on total covered expenditures. Surprisingly, many low-income countries have adopted the latter strategy, providing the minimum benefits for all illnesses rather than full insurance for rare high-cost illnesses (Gertler and Solon, 1998). The rationale for this choice is the concern that the lower income groups may not be able to “afford” the deductible, and therefore would not benefit from the insurance. However, if families are able to insure small health shocks, then first-dollar capped-benefits provides little increased insurance but rather simply crowds out private informal insurance. Our findings suggest that these countries would be much better served by moving towards catastrophic -type models of insurance coverage.

References

- Andrews, G., A. Esterman, A. Braunack-Mayer, and C. Rungie (1986). Aging in the Western Pacific. Manila: World Health Organization.
- Asian Development Bank (1997). Asian Development Outlook 1996 and 1997. Manila: Asian Development Bank.
- Benjamin, Dwayne (1992). "Household Composition, Labor Markets, and Labor Demand: Testing for Separation on Agricultural Household Models," Econometrica, 60, 287-322.
- Bound J., (1991). "Self-Reported versus Objective Measures of Health in Retirement Models," Journal of Human Resources, 26, 106-38.
- Bound J., M. Schoenbaum, and T. Waidman (1995). "Race and Education Differences in Disability Status and Labor Force Participation," Journal of Human Resources 30, S227-S267.
- Cochrane, John (1991). "A Simple Test of Consumption Insurance," Journal of Political Economy, 99, 957-76.
- Cornia, G., R. Jolly and F. Stewart (1987). Adjustment With a Human Face. Oxford: Clarendon Press.
- Deaton, Angus (1992a). "Saving and Income Insuring in Cote d'Ivoire," Journal of African Economics, 1, 1-24.
- Deaton, Angus (1992b). Understanding Consumption. Oxford: Clarendon Press.
- Deaton, Angus (1994). The Analysis of Household Surveys: Microeconomic Analysis for Development Policy. Monograph, World Bank, Washington DC.
- Dow, William, and Paul Gertler (1997). "Private Health Insurance and Public Expenditures in Indonesia". Mimeo, RAND, Santa Monica, CA.
- Dow, William, Paul Gertler, Robert Scheoni and Duncan Thomas (1996). "Health Care Prices, Health and Labor Outcomes: Experimental Evidence". Mimeo, RAND, Santa Monica, CA.
- Engen, Eric (1993). "A Stochastic Life-Cycle Model with Mortality Risk: Estimation with Panel Data." Mimeo, UCLA.
- Gertler, Paul (1998) "On The Road to Social Health Insurance: the Asian Experience," World Development, Vol.26, No. 4, pp.717-732.
- Gertler, Paul, and Jack Molyneaux (1996). "The Effect of Medical Care Prices on Utilization and Health Outcomes: Experimental Results". Mimeo, RAND, Santa Monica, CA.
- Gertler Paul and Orville, Solon (1999) "Who Benefits From Social Health Insurance," mimeo, University of California, Berkeley
- Gertler, Paul and Jenifer Zeitlin (1996). "The Returns to Childhood Investments in Terms of Health Later in Life," forthcoming in T. Hu and H. Chee (ed.s) Health and Health Care in Asia, Oxford University Press.

- Guralnik, J., L. Branch, S. Cummings and J. Curb (1989). "Physical Performance Measures in Aging Research," Journal of Gerontology, 44, 141-46.
- Hayashi, Fumio, Joseph Altonji and Laurence Kotlikoff (1996). "Risk Sharing Between and Within Families," Econometrica, 64, 261-294.
- Jiminez, E. (1996). "Human and Physical Infrastructure," in J. Behrman and T.N. Srinivasan (eds.) Handbook of Development Economics. Amsterdam: North Holland.
- Ju, A. and G. Jones (1989). Aging in ASEAN and its Socio-economic Consequences. Singapore: Institute of Southeast Asian Studies.
- Kington R. and J. Smith (1996) "Socio-Economic Correlates of Adult Health," Demography, forthcoming.
- Kochar, Anjini (1995). "Explaining Household Vulnerability to Idiosyncratic Income Shocks," American Economic Review, 85, 159-164.
- Mace, Barbara (1991). "Full Insurance in the Presence of Aggregate Uncertainty," Journal of Political Economy, 99, 928-56.
- Morduch, Jonathan (1990). "Risk, Production, and Saving: Theory and Evidence from Indian Households". Mimeo, Harvard University, 1990.
- Morduch, Jonathan (1995). "Income Insuring and Consumption Insuring," Journal of Economic Perspectives, 9 (Summer), 103-114.
- Morduch, Jonathan (1999). "Between the market and State: Can Informal Insurance Patch the Safety Net?" forthcoming in The World Bank Research Observer.
- Pitt, M. and M. Rosensweig (1986). "Agricultural Prices, Food Consumption, and the Health and Productivity of Farmers," in I. Sing, L. Squire and J. Strauss (eds.) Agricultural Household Models: Extensions and Applications. Baltimore: Johns Hopkins Press.
- Ready, Sanjy (1996). "A Critical Review of User Charges for Basic Services," Policy Review No. 136, UNICEF, New York.
- Rothschild, Michael, and Joseph E. Stiglitz (1976). "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," Quarterly Journal of Economics, 90, 629-650.
- Schultz, T.P. and A. Tansel (1997). "Wage and labor supply effects of illness in Cote d'Ivoire and Ghana: Instrumental variable estimates for days disabled," Journal of Development Economics, forthcoming.
- Sindelar, J. and D. Thomas (1991). "Measurement of child health: Maternal response bias," Economic Growth Center Discussion Paper #663, Yale University.
- Stern, S. (1989). "Measuring the Effect of Disability on Labor Force Participation," Journal of Human Resources, 24, 361-95.

- Stewart, A., J. Ware, R. Brook, and A. Davies-Avery et al. (1990). Measurement of Adult Health Status: Physical Health in Terms Functional Status. Cambridge: Harvard University Press.
- Strauss, John, and Duncan Thomas (1996). "Human Resources: Empirical Modeling of Household and Family Decisions," in J. Behrman and T.N. Srinivasan (eds.) Handbook of Development Economics. Amsterdam: North Holland.
- Strauss, J., P. Gertler, O. Rahman, and K. Fox (1993). "Gender and Life-Cycle Differentials in the Patterns and Determinants of Adult Health," Journal Of Human Resources, 28, 791-837.
- Townsend, Robert (1994). "Risk and Insurance in Village India," Econometrica, 62, 539-92.
- Townsend, Robert (1995). "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies," Journal of Economic Perspectives, 9 (Summer), 83-102.
- Wallace, R. and A. Herzog (1995). "An Overview of Health Status Measures in the Health and Retirement Survey," Journal of Human Resources, 30, S84-S107.
- Ware, J., A. Davies-Avery, and R. Brook (1980). Conceptualization and Measurement of Health Status for Adults in the Health Insurance Study: Vol. IV, Analysis of relationships Among Health Status Measures. R-1987/6-HEW. Santa Monica, CA. RAND.
- World Bank (1987). Financing Health Services in Developing Countries. World Bank Policy Study, Washington DC.
- World Bank (1993). World Development Report: Investing in Health. Oxford: Oxford University Press.
- World Bank (1995a). Averting the Old Age Crisis. Oxford: Oxford University Press.
- World Bank (1995b). Indonesia: Public Expenditure, Prices and the Poor. Washington DC: World Bank.
- Zeldes, Stephen P. (1989). "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, 305-346.

Table 1: Means and Standard Deviations

	Period 1 (1991)	Change (1993-1991)
<i>Head's Health Status Measures</i>		
ADL Index	0.966 (0.082)	0.005 (0.088)
Some ADL Limitation (=1)	0.300	-0.130
Any ADL Deterioration (=1)	---	0.111
Mean ADL Deterioration Conditional on Some Deterioration		-0.016 (0.064)
Any ADL Improvement (=1)	---	0.227
Mean ADL Improvement Conditional on Some Improvement		0.021 (0.055)
Illness Symptom (=1)	0.601	-0.015
Chronic Illness Symptom (=1)	0.144	0.161
Intermediate ADL Index	0.933 (0.146)	0.013 (0.148)
Basic ADL Index	0.993 (0.045)	-0.002 (0.064)
<i>Socio-Economic Variables</i>		
Non-Medical Consumption Per Capita	36,350 (31,868)	7,229 (62,265)
Medical Spending/Capita	335 (1,020)	68 (1,412)
Head's Earnings/Capita	17,573 (19,256)	742 (19,746)
Head's Hours of Work	36.6 (24.6)	-4.72 (27.40)
Head Not Working (=1)	0.157	0.024
Head is Male (=1)	0.86	---
Head is Married (=1)	0.79	---
Spouse's Age Conditional on Head Married	35.8 (11.6)	---
Family Size	4.80 (2.13)	-0.53 (0.97)
Head Has No Education (=1)	0.33	---
Head Has 1-5 Years Education	0.32	---
Head Had 6 Years Education	0.19	---
Head Has 7+ Years Education	0.16	---

Notes: Tabulated by authors from IRMS data. Standard Deviations are in parentheses for continuous variables. N=3933.

Table 2: Effect of Change in Head's Health on Change in Head's Hours Worked

Change in Head's Illness Symptoms	-0.52 (0.80)		
Change in Head's Chronic Symptoms		-1.04 (1.02)	
Change in Head's ADL Index			30.85 (5.08)
Head is Male (=1)	3.14 (2.28)	3.23 (2.32)	3.21 (2.30)
Age of Head	-0.26 (0.21)	-0.26 (0.21)	-0.31 (0.21)
Head's Age Squared/100	0.22 (0.20)	0.22 (0.20)	0.26 (0.19)
Head Has No Education (=1)	2.78 (1.62)	2.66 (1.62)	2.82 (1.62)
Head Has 1-5 Years Education (=1)	2.69 (1.47)	2.64 (1.46)	2.76 (1.46)
Head Has 6 Years Education (=1)	4.71 (1.59)	4.59 (1.59)	4.67 (1.59)
Head is Single (=1)	0.39 (3.27)	0.47 (3.27)	0.05 (3.26)
Wife's Age	-0.08 (0.07)	-0.08 (0.07)	-0.08 (0.06)
Change in Log (Family Size)	-1.49 (2.28)	-1.39 (2.28)	-1.51 (2.27)

Notes: Standard errors in parentheses. Estimates are from models such as (2) in text. N=3933. Included in the estimation, but not reported in the table, are community fixed effects.

Table 3: Effect of Change in Head's Health on
Change in Labor Supply, Earnings and Medical Spending

	Change in Head's Illness Symptoms	Change in Head's Chronic Symptoms	Change in Head's ADL Index
Change in Head's Hours Working	-0.52 (0.80)	-1.04 (1.02)	30.9 (5.08)
Change in Head's Labor Force Participation	-0.031 (0.012)	-0.029 (0.014)	0.738 (0.080)
Change in Head's Imputed Earnings (in Rp. 10,000)	-0.128 (0.060)	-0.060 (0.076)	2.02 (0.35)
Change in Head's Medical Care Spending (in Rp. 10,000)	0.022 (0.004)	0.015 (0.006)	-0.118 (0.026)
Change in Imputed Earnings of Other Household Members (in Rp. 10,000)	0.062 (0.032)	-0.029 (0.038)	0.256 (0.220)
Change in Total Imputed Household Earnings (in Rp. 10,000)	0.068 (0.084)	-0.030 (0.101)	2.19 (0.57)

Notes: Standard errors in parentheses. Each coefficient in the table is from a separate regression model. The dependent variables are listed in the first column, and the independent variables of interest are listed in the first row. The coefficient reported is that on health change in a regression that includes all covariates shown in Table 2 and reported in the footnote to Table 2. N=3933.

Table 4: Effect of Change in Head's Health on Change in Non-Medical Consumption

Change in Head's Illness Symptoms	0.003 (0.012)		
Change in Head's Chronic Symptoms		0.012 (0.015)	
Change in Head's ADL Index			0.195 (0.083)
Head is Male (=1)	0.022 (0.037)	0.021 (0.037)	0.023 (0.038)
Age of Head	0.002 (0.003)	0.002 (0.003)	0.002 (0.003)
Head's Age Squared/100	-0.003 (0.003)	-0.003 (0.003)	-0.003 (0.003)
Head Has No Education (=1)	0.003 (0.027)	0.006 (0.026)	0.002 (0.026)
Head Has 1-5 Years Education (=1)	-0.005 (0.024)	-0.005 (0.023)	-0.003 (0.023)
Head 6 Years Education (=1)	0.035 (0.026)	0.036 (0.026)	0.033 (0.024)
Head is Single (=1)	0.077 (0.053)	0.079 (0.052)	0.076 (0.053)
Wife's Age	0.0003 (0.001)	0.0002 (0.001)	0.0003 (0.001)
Change in Log (Family Size)	-0.476 (0.037)	-0.476 (0.037)	-0.475 (0.037)

Notes: Standard errors in parentheses. Estimates are from models such as (8) in text. Coefficients on change in share of family members in 8 age/sex groups and community fixed effects not reported. N=3933.

Table 5: Tests for State Dependence

	Workers vs. Non-Workers	Other Family Earnings Share Interaction (N=2817)	Other Family Health and Earnings Interactions (N=2817)
Change in Head's ADL Index for Workers (N=2817)	0.369 (0.153)		
Change in Head's ADL Index for Non- Workers (N=1116)	-0.005 (0.207)		
Change in Head's ADL Index		0.507 (0.148)	0.507 (0.150)
Other Family Members' Period 1 Earnings ? Head's Period 1 Earnings		0.004 (0.007)	0.001 (0.007)
Interaction of Change in Head's ADL Index and Other/Head Earnings Ratio		-0.201 (0.085)	-0.229 (0.086)
Change in Other Family Members' ADL Index			-0.036 (0.131)
Interaction of Change in Other Members' ADL Index and Other/Head Earnings Ratio			0.343 (0.129)

Notes: The dependent variable is the change in household per capita consumption. The models include all controls shown in Table 4 and noted in footnote to that table. Standard errors are in parentheses. The first column compares the effect of a change in the head's ADL index on family consumption estimated separately for heads working and not working in the first period. The second and third columns each represent a separately estimated model.

Table 6: Estimating Magnitude of Consumption Insurance

Estimation Method	Instruments for Δ in Imputed Income Minus Medical Expenditures	Δ	$\frac{\Delta}{C_{t+1}}$
OLS		0.007 (0.004)	0.03
IV	Head's Overall ADL Index	0.096 (0.044)	0.35
IV	Head's and Other Family Members' Overall ADL Indices	0.107 (0.037)	0.39
IV	Head's Intermediate ADL Index Only	0.079 (0.039)	0.29
IV	Head's Basic ADL Index Only	0.171 (0.085)	0.62

Notes: Standard errors in parentheses. Each regression includes controls shown in Table 4 and noted in footnote to that table. Coefficient in column (1) is that on change in imputed earnings of head minus change in medical spending from regressions of the form of (9), where instruments are listed at left. Figures in column (2) are implied effects of imputed income changes on consumption changes.