INTER-AMERICAN DEVELOPMENT BANK

CONFERENCE ON SOCIAL PROTECTION AND POVERTY

Insuring the Economic Costs of Illness

Paul Gertler

University of California at Berkeley and NBER

February 5, 1999

Organized by the Poverty and Inequality Advisory Unit
of the Sustainable Development Department

Washington D.C.

Abstract

One of the most sizable and least predictable shocks to the economic opportunities of families in developing countries 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 smooth 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). While families with sick members in developed countries are able to access formal insurance markets, families in low income countries must rely on informal mechanisms such as drawing on savings, selling assets, transfers from their family and social support networks, and borrowing from local credit markets. The possibility that there is less than full consumption smoothing through these mechanisms suggests a potentially large loss in welfare from this shock to the household's resources.

This paper has been prepared for the Conference on Social Protection and Poverty. Its findings, interpretations and conclusions reflect those of the author and should not be attributed to the Inter-American Development Bank. I am grateful to John Giles for useful comments and advice.

I. INTRODUCTION

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 smooth 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 smoothing through these mechanisms suggests a potentially large loss in welfare from this shock to the household's resources. Recognizing this potential loss in welfare and the failure of private health insurance markets, many developing countries have or are considering social insurance to help smooth the economic costs of illness.

In this paper, we investigate the potential welfare gain from social insurance. In particular, we pay attention to the possibility that social insurance may crowd out or replace private informal insurance. 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 family and social support networks, and borrowing from local credit markets[1]. Estimates of the welfare gain from social insurance must net out the crowd out effect.

The extent to which families are able to smooth consumption over periods of illness is also important for the design of social insurance programs. In particular, as we discuss in Section II, the extent to which families are able to smooth small health shocks has implications for user fees and the ability of social insurance to finance the rarer expensive health shocks.

To determine the extent to which households are able to employ private informal insurance mechanisms to smooth health shocks, we investigate the ability of families in Indonesia and China to smooth consumption over periods of major illness. To do so, we use unique panel data, which contains excellent measures of health status combined with data on consumption.

While there is a growing literature on consumption smoothing in developing countries, little explicit attention has been paid to smoothing of health shocks[2.] Townsend (1994) includes in his regression analysis the "percentage of the year sick", and finds no effect on consumption changes. Kochar (1995) models wage income and informal borrowing as a function illness, 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 smooth illness shocks fairly well.[3] A key limitation of past work, however, is that the measures of health employed may reflect only small, and even potentially anticipated, changes in health status, not the kind of large unexpected major illnesses that may be difficult to smooth. Gertler and Gruber (1997) overcome these problems by using measures of individuals' physical abilities to perform activities of daily living (ADLs). ADLs have been proven reliable and valid

---

[1] Morduch (1999) provides a useful summary of this literature

[2] See Morduch (1995) and Townsend (1995) for reviews of the consumption smoothing literature.

[3] 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.

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. They find that while households are able to smooth 70 percent of the costs of moderate health shocks, they can only smooth about 40 percent of the costs of serious illness.

The paper is organized as follows. In the next section, we provide a policy framework. The policy framework argues that optimal user fee policy and the design of social insurance benefits depends on the extent to which households are able to smooth small and large health shocks. Section III describes welfare measurement using a certainty equivalence framework and identifies how crowd-out affects the welfare measurement. In the section, we also describe the empirical specification, which we use to estimates the degree of crowd-out. Section IV and V provides the analysis for Indonesia and China, respectively. We find that in both countries households are able to smooth insure small health shocks, but not the costs of major illness. We also find evidence that households are willing to pay substantial amounts to reduce the variation in consumption from the cost of the major illness, suggesting a role for government in providing social insurance. Finally, we end with policy conclusions.

II. POLICY FRAMEWORK

The classic reason for most governments to intervene in health care markets is the inherent uncertainty in health status (Arrow 1963). No one knows what tomorrow will bring. Seemingly healthy individuals can be struck by cancer, injured in accidents, or experience severe diarrhea. The uncertainty is compounded the longer one looks into the future and the less one knows about one's current health. While the costs associated with most illnesses are small, the costs associated with rare serious illness can be quite large. When a serious illness hits unpredictably, the economic consequences for the family can be significant. Given aversion to risk, families would prefer to have predictable health costs smoothed over time, rather than incurring the costs unexpectedly all at once. Insurance smoothes the *ex ante* cost of illness over time and individuals, thereby making the costs predictable.

Despite the demand, families have difficulty purchasing insurance from private sources because of two problems created by information asymmetries: adverse selection and cream skimming. Adverse selection arises because insurers are not able to observe differences in health status among different people. Individuals are born with different genetic makeup, have different life course experiences with respect to exposure to environmental contagion and accidents, and investments in their health (e.g. diet, exercise, etc.). Insurers are not able to observe individuals propensity to become ill and so they are forced to offered insurance based on the population's average medical care expenditures. For high propensity individuals (bad risks) , the terms of the contract are quite good. However, the terms are quite bad for good risks (low propensity individuals). The incentive is for the good risks to drop out of the market, leaving the bad to insure among themselves. This substantially drives up the price of insurance, making it unaffordable to large segments of the population. Cream skimming occurs when insurers can observe poor health. Competing on their ability to select good risks leads insurers to avoid individuals with pre-existing conditions such as cancer or AIDS, who "certain" bad risks with predictably high medical care expenditures. In addition, since insurance contracts are written for discrete periods of time (e.g. one year), individuals who develop serious illness are prohibited from renewing their insurance.

Private insurance market failure suggests the potential for welfare gains from government intervention. The main problem with insurance market failure is that it is voluntary so that good risks can choose not to buy insurance and bad risks can be denied coverage. Many countries solve this problem through mandatory social insurance (SI). SI takes two main forms.

The first is universal publicly financed and delivered health care modeled after the British National Health

System. In these NHS-like systems governments provide medical care through public facilities which are accessed by paying at most a nominal user fee. These systems are financed through some combination of general tax revenues and payroll taxes. There is a heated policy debate about raising user fees at public health care facilities. Governments have or are actively considering raising user charges at public facilities as a means of financing improvements in the health sector (e.g. World Bank, 1987; Jimenez, 1994). Vocal opponents are concerned that increased fees will adversely affect the poor's access to medical care and, consequently, their health outcomes (e.g. Cornia, Jolly and Stewart, 1987; Ready, 1996). This debate, however, has ignored the possible role of public subsidies as insurance. Subsidies reduce risk by spreading the medical costs of uncertain illness across healthy and sick times; taxes incurred when healthy finance medical care purchased when sick. As a result, raising user fees in a world of imperfect consumption insurance has an important welfare cost: higher user fees "tax families while they are down", imposing higher costs at exactly the point where the marginal utility of consumption is highest.

The second approach is to finance medical care through mandatory payroll taxes, and allow beneficiaries to purchase medical care from both private and public providers. The movement towards this model has been motivated out of the desire to reduce financial pressure on government budgets (Gertler, 1998). SI is seen as a way to shift a portion of the public burden of delivering and financing health care to the private sector. SI reduces the out-of-pocket prices at the time of purchase of higher quality private care relative to lower quality public care. As a result, SI provides an incentive to choose the private sector over the public sector, thereby lowering the demand for publicly delivered services. Since SI is financed through additional off budget earmarked taxes, it also relieves pressure from the general budget.

Low-income countries, however, have limited abilities to tax. Therefore, the resources available for SI are severely constrained which greatly diminishes their ability to provide insurance. In such poor environments, there is a strict budget constraint on SI benefits. SI plans thus face the trade-off between providing a large number of individuals with a small benefit or a small number of people with a large benefit. This means that a very large deductible (and possibly a large copayment) would be required to provide uncapped benefits for the rare large financial risks—e.g. those associated with rare catastrophic illnesses such as cancer. The high deductible would ensure that benefits are available for expensive catastrophic illnesses and are not used up on less expensive higher-probability events (e.g. influenza). Because of the budget constraint, lowering the deductible would require capping benefits. In the limit, a zero deductible implies the lowest possible benefit cap and least effective insurance against catastrophic illness.

All the same, this is exactly what many countries have done. Gertler and Solon (1998) show that many low-income countries have adopted first-dollar coverage and therefore have placed the lowest cap possible on benefits. In essence, they have chosen to provide the minimum benefits for all illnesses rather than full insurance for rare high-cost illnesses.

A number of powerful political interest groups see a first-dollar capped benefit design as in their self-interest. The most obvious interest group is the collection of international donors and other political groups worried about the poor and who want to use the health care system as a means for redistribution (Besley and Gouveia, 1994). First dollar coverage ensures universal access to medical care regardless of income. It alleviates the widespread concern that that even small out-of-pocket costs may deter utilization, especially among the very poor (Cornea, Jolly and Stewart, 1987)[4]. Politicians also support first-dollar coverage. Since the median voter is poor in most

---

[4] Indeed, a good portion of public intervention in health care markets is justified on the tenet of universal access to basic minimum medical care, regardless of income, as embodied in the populist slogan "health for all by the year 2000." Most countries recognize that poor individuals may not be able to afford health care and therefore subsidize the participation of poor individuals in universal insurance schemes. This explains, in part, why most countries set up large universal public health care delivery systems that charged at most nominal user fees (Jimenez, 1987; World Bank, 1987).

of these countries, first-dollar coverage puts money into more voters' pockets. Employers also have strong financial incentives to support capped-benefit first-dollar coverage. Typically, social insurance premiums are co-financed by employers. In many countries, large employers historically have provided workers with limited health benefits as a means of reducing absenteeism (Gertler and Sturm, 1997). However, quick treatment of minor illnesses reduces absenteeism more than the treatment of catastrophic illnesses. Employers capped benefits since it was cheaper to fire severely ill individuals who had little chance of returning to work. For similar reasons, employers benefit more from the capped first-dollar coverage that is more likely to reduce absenteeism, than from catastrophic coverage with high deductibles that is less likely to affect workforce productivity.

The solution to both the user fee debate and the debate over first dollar coverage depends in part on the extent to which families are able to smooth the costs of small health shocks. If yes, then there is little welfare loss to charging small user fees or copays and deductibles.

III. Welfare Measurement

In this section, we consider measuring the welfare gains from insuring the economic costs of illness. We focus solely on the welfare gains from more complete insurance and abstract from other important potential welfare gains such as improvements in health status and gains in social welfare from redistribution. We begin by discussing the measurement of the welfare loss from not being able to insure consumption. We then turn to estimating the gains from formal disability insurance and medical subsidies.

One measure of the welfare cost of not being able to fully insure the costs of illness is the amount that households are willing to pay to eliminate consumption variability due to illness. This measures households ex ante valuation of insurance that would fill the gap in existing insurance markets for the income loss due to illness, arising either through reduced earnings or increased medical expenditures.

We calculate the willingness to pay in a certainty equivalence framework. Let C^* be consumption when healthy and $L(Hi)$ be the economic cost of illness with severity H_i which occurs with probability p_i. Then the welfare loss from uncertain illness is the amount, W, such that the welfare from getting $C^* - W$ with certainty is equal to the

$$U\left(\frac{C^* - W}{C^*}\right) = E\left[U\left(\frac{C^* - L(H_i)}{C^*}\right)\right] \tag{1}$$

expected welfare when the loss is uncertain:

However, the willingness to pay formula overstates the welfare gain from social insurance because it ignores that fact that households will smooth some of the shocks through private informal mechanisms. So, if we let *?* represent the share of the loss that cannot be smoothed, then (1) can be written as

$$U\left(\frac{C^* - W}{C^*}\right) = E\left[U\left(\frac{C^* - \gamma L(H_i)}{C^*}\right)\right] \tag{2}$$

Assuming a constant relative risk aversion form for the utility function, where *?* is the coefficient of relative risk aversion, (10) can be rewritten as:

$$\frac{\left(\frac{C^* - W}{C^*}\right)^{1-r}}{1-r} = \Sigma_j \, p_j * \frac{\left(\frac{C^* - gL_j}{C^*}\right)^{1-r}}{1-r} \qquad (3)$$

where there are *j* discrete health states. Rearranging terms, we can express the certainty equivalent as a

$$\frac{W}{C^*} = 1 - \left(\Sigma_j \, p_j * \left(\frac{C^* - gL_j}{C^*}\right)^{1-r}\right)^{\frac{1}{1-r}} \qquad (4)$$

percentage of consumption when healthy:

W/C* measures the value of insurance that fully smoothes consumption across illness states, as a percentage of baseline consumption. This measure is a lower bound of the willingness to pay for insurance, however, since it is calculated based on the variation in consumption due to illness after families have already used informal mechanisms to smooth some of the costs of illness. These smoothing activities themselves have costs that are not reflected in the calculation. For example, there is some cost to family and friends from private transfers of resources to the ill household head; similarly, if consumption smoothing is occurring through increased labor supply by family members, the value of the reduced leisure to those family members is not reflected here.

The key to estimating (4) is *?*, the share of the costs of illness that are not able to insure. We estimate it in the context of the theory of full insurance, as discussed in Cochrane (1991), Deaton (1992a), and Townsend (1994). This theory posits that households will fully share the risk of idiosyncratic shocks so that the growth in household consumption will not depend on changes in household resources once the change in aggregate community resources has been taken into account.

The theory of full insurance implies that we can test whether families are able insure consumption against illness by estimating the following equation[5,6]:

$$\Delta \ln\left(\frac{C_{ij}}{n_{ij}}\right) = a\Delta \ln(C_j) + m\Delta \ln(n_{ij}) + d\Delta y_{ij} + e_{ij} \qquad (5)$$

[5] The theoretical model developed above is cast in terms of consumption insurance through interhousehold risk sharing. In practice, our empirical analysis follows the previous developing country literature in examining consumption smoothing, either through insurance from others or through self-insurance (ie. savings). Our tests do not distinguish between these two channels for consumption smoothing.

[6] A major assumption of the full insurance interpretation 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,

where C^{it} is household per capita expenditures in village j, C_j is per capita consumption in village j, y_{ij} is household income per capita net of medical care expenditures, and n_{ij} is household size. This is a regression of the growth in per capita (non-medical care) consumption against the change in net income, controlling for the growth in community resources by including the change in community level consumption. In addition, we control for preference shifts associated with changes in family 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 ages 0-5, 6-17, 18-49, and 50 plus. 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. Since the dependent variable is a natural log, the share of the costs of illness that are financed out of reduced consumption is

$$? = \delta \times Cij.$$ [7]

IV. INDONESIA

A. The Institutional Setting

Indonesia is the fourth most populous country in the world with tremendous cultural and economic diversity. Though until recently economic growth has been impressive with an average real annual per capita growth rate of 3.9% over the last 15 years, per capita incomes are still low, at $US 880 per year in 1996 (Asian Development bank, 1997). Since the Asian economic crisis, real income has fallen substantially and a large number of households have fallen into poverty. Indonesia had also seen remarkable improvements in health status (World Bank, 1993). The infant mortality rate fell by about 35 percent from 1965 to 1990. 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. By 1991, there was at least one primary health center and several subcenters in each of Indonesia's 3400 subdistricts, and a large network of hospitals at the district, provincial and central levels backing up the primary care system. Despite this large public health care system, Indonesia's expenditures on health care remain low relative to those of its neighbors (World Bank, 1993). In 1990, annual expenditures on health care from both public and private sources are only about $12 per person which amounts to roughly 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 average 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 are covered by health insurance provided to civil servants. However, this insurance only covers utilization at public facilities and, therefore, the benefit to the individual is coverage of only low 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).

---

then even with full insurance the growth of consumption will vary with the state of health. Gertler and Gruber (1997) uniformly reject state dependence using the same data.

[7] 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.

Similarly, there is limited disability insurance as there is no government program, over two-thirds of workers are self-employed, and few firms provide extensive sick leave.

B. Data.

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 Tengarah (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 ranks twenty-second. The data were collected in 1991 and 1993, allowing us to examine health, income, and consumption changes over a two year period. 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 differences models.

Our sample consists of all household who were in the survey in both rounds, whose heads were at least 18 years old in the second round, and who have non-missing data on the health measures described below. Gertler and Molyneaux (1996) and Dow et al. (1997) discuss attrition from this sample, and conclude that it does not cause significant sample selection problems.

C. 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 assessments of 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).[8]

Activities of daily living 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 as "can do it easily" (a value of 1), "can do it with

[8]ADLs 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.

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 (Stewert et al., 1990):

$$Health = \left(\frac{Score - \text{Min } Score}{\text{Max } Score - \text{Min } Score} \right)$$

so that the ADL index takes on a value of 1 if the individual can perform all ADLs without difficulty, and zero is 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 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 proportional of adults, 29 percent, reported some ADL limitation. In addition, there is substantial change in health status over time. Between 1991 and 1993, over 33% of the sample reported changes (either upwards or downwards) in ADL limitations.

Table 1: Descriptive Statistics for The Health Measure (N=3,933)

Period 1 Levels		Changes	
ADL Index	0.966	Change in ADL Index	0.005
	(0.082)		(0.088)
Any ADL Limitations	0.29	Change in Any ADLs	0.337

Despite their severity, changes in ADLs do not appear to be permanent on average: there are roughly as many upward movements as downward movements in the ADL index. This reflects an important difference in the interpretation of ADLs, and in particular basic limitations, in developed and developing country contexts. In wealthier 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 ones house. Also, toileting requires the use of eastern as opposed to western toilets, which are many times located in outside the main living quarters. Hence, basic indicators capture less severe limitations in developing country settings than in developed 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 very similar impacts when we restrict our estimates to downward movements in ADLs only.

D. The Cost of Illness

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 in terms of reduced labor supply, lost earnings and increased medical spending.

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

$$\Delta L_{ij} = b\Delta h_{ij} + u_j + e_{ij} \tag{6}$$

where $?L_{ij}$ is the change in labor supply (or earnings, or medical care spending) for individual we living in community j, $?h_{ij}$ is the change in health for that individual, and $?_j$ is a community-level error component. Equation (1) therefore regresses first differenced labor outcomes and medical care spending against the change in health and aggregate determinants of labor supply (or medical spending). We include a full set of community dummies to control for these aggregate determinants.[9] 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 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, it 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 income and 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 permanent income and health, by including a set of community fixed effects.[10]

Earnings and wages are only reported in the IRMS data for the one-third of heads who work in the market. We therefore impute wages to all workers based on these 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 1b), 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.[11] 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.

The means and standard deviations of the earnings and medical care expenditures variables are reported in Table 2. 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 for both are reported in Table 2. 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. Both are measured in real per capita terms, in order to match our consumption specification

---

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

[10] A related source of 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). But our pattern of results suggests 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.

[11] 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).

below.[12]

Table 2: Descriptive Statistics for Dependent Variables (N=3,933)

Variable	Mean	Standard Deviation
Head's Earnings Per Capita	17,573	19,253
Household Non-Medical Consumption Per Capita	36,351	36,351
Household Medical Spending Per Capita	335	335

Table 3 reports the coefficient of interest from the earnings and medical care spending regressions expressed in 10,000 Rupiah units. The coefficient on the ADL change is significantly different from zero and implies that moving from completely able to perform ADLs to completely unable to perform ADLs would lower earnings by 20,170 Rp. 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 1834 Rp., or about 10% of baseline earnings.

Table 3 also shows the effects of illness on medical spending, once again in units of 10,000 Rupiah. There is a significant effect in the expected direction, but the magnitude is trivially small relative to the effect on earnings. This is not surprising since publicly provided medical care is heavily subsidized (i.e. user fees are well below the cost of care).

Table 3: Effect of Illness on Head's Earnings and Medical Spending*

	Change in Earnings	Change in Medical Spending
Change in ADLs	2.02 (0.35)	-0.118 (0.026)

* These are coefficients and standard errors from a regression of the change in earnings or change in medical spending on change in ADLs controlling for the age and sex of the head, marital status, spouse's age, head's education, change in the log of household size, and village fixed effects. See Gertler and Gruber (1998) for full regression results.

E. Consumption Smoothing

Equation (5) cannot be estimated by OLS. There are two potential problems. The first, as noted by Morduch (1990), is that the growth in income is correlated with the error term through the production process; risk averse families may choose the variation in income so that consumption can be smoothed using available mechanisms. The second is measurement error in the growth of income, particularly since we have imputed earnings in our data.

[12] All figures are reported in 1991 urban NTB rupiah.

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. We report the coefficient of interest in the first row of Table 4. The estimate indicates that for every 10,000 Rp. of income lost due to illness, there is a fall in consumption of 9.6%, or 3490 Rp. That is, for each rupiah that income falls, consumption falls by Rp. 0.35 (as shown in the second column). This suggests that, on average, families are able to smooth only 65% of the loss in income from ADL changes.

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. More specifically, it assumes that the consumption response is linear across all health shocks. There may be reason to believe that individuals are more able to smooth the modest income loss associated with minor health changes than they are the larger income loss associated with major health changes; this would be consistent with our much larger findings for ADLs than for other illness measures. This could arise, for example, if individuals have available limited self-insurance that can cover small income losses, or if there are other forms of consumption insurance available (e.g. loans from extended family) in small but not large amounts.

Table 4: Effect of Illness on Family Consumption*

	γ	γCt-1
Overall ADL Index	0.096 (0.042)	0.35
Intermediate ADL Index	0.079 (0.039)	0.29
Basic ADL Index	0.165 (0.095)	0.60

* These are coefficients and standard errors from a regression of the change in earnings or change in medical spending on change in ADLs controlling for the age and sex of the head, marital status, spouse's age, head's education, change in the log of household size, and village fixed effects. See Gertler and Gruber (1998) for full regression results.

We explore this issue by disaggregating our ADL index separately into changes in intermediate and basic ADLs, respectively. As noted earlier, the latter set of limitations are much more serious; while 24% of our sample has some intermediate ADL limitation, only 2% has some basic ADL limitation. We therefore vary the instrument set used to estimate Equation 2; that is, the coefficients in the table are estimates of the effect of income variation on consumption, but where the instruments are not just the overall ADL index, but also

the separate indices for basic and intermediate ADLs. This provides a natural basis for comparing the sensitivity of our key finding to the specification of the illness instrument.

Our results, reported in the second and third rows of Table 4, suggest that the impacts of basic ADL limitations on consumption are much more sizeable. We find that families can smooth 71% of the income loss associated with intermediate ADL changes, but only 40% 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 smooth illness shocks. Specifically, families are much less able to smooth low frequency large shocks.

F. Insurance Market Policies

The results thus far demonstrate that families are not able to smooth the economic costs arising from serious illness to the household head, and we have documented the extent to which smoothing is imperfect. This incompleteness in private insurance markets suggests the potential for welfare gains from government provision of insurance against income loss and medical illness. In this section, we consider the magnitudes of these welfare gains. We focus solely on the welfare gains from more complete insurance and abstract from other important potential welfare gains such as improvements in health status and gains in social welfare from redistribution.

Disability Insurance

In this section, we contrast the benefits and costs of formal disability insurance that fully smoothes consumption over the loss in earnings arising from illness. The gain to the household from such insurance is the expected value of the transfer from the insuring agency, plus the welfare gain from consumption smoothing. The cost to the government is the expected value of the transfer, plus (i) a markup for administrative costs, (ii) the cost of moral hazard through increased reported illness in response to the existence of this program, and (iii) any deadweight loss from financing these benefit payments. That is, the expected benefit payout is just a transfer from the government to households; the ultimate efficiency of disability insurance policy rests on a comparison of the welfare gains from consumption smoothing and the inefficiencies inherent in operating a disability insurance program. Measuring these inefficiencies is beyond the scope of our paper. But, by comparing the welfare gain from consumption smoothing to the expected benefits payout (the transfer), we can offer a sense of how large these costs would have to be in percentage terms to offset the consumption smoothing benefits of disability insurance.

We begin by estimating the expected payout from a disability insurance policy that fully replaces the earnings loss to those who become ill by our ADL metric. We measure the loss in earnings from an illness, L_j, by using the estimates of equation (2) reported in the third row of Table 3 to predict the loss in earnings from downward movements in the ADL indexes.[13] We measure the probability of experiencing the loss, pj, using the observed frequency distribution of downward movements in the ADL indexes. The expected loss in earnings, then, is the sum of the p_j times L_j. Our estimate for the expected insurance payout for earnings losses is 0.89% of baseline consumption. This estimate is quite small, despite the large impact on earnings from major changes in ADLs, because of the low frequency of these major changes in the data.

It is important to note that the cross-sectional frequency distribution is not the theoretically appropriate set of probabilities p_j to use for this exercise; we should be using the longitudinal probabilities for each individual in our data set, not the cross-sectional distribution of risk. Unfortunately, however, we only have two observations on each individual, so we must rely on this cross-sectional distribution. This will lead us to understate somewhat the welfare loss from illness with concave utility, since we are using the average risk rather than the underlying

[13] It is inappropriate to incorporate the consumption increases from upward movements in health, which simply reflect recovery from earlier downward shifts. As we show in the Appendix, our results for the consumption smoothing effects of changes in ADL index are very similar if we just use downward shifts to identify our estimates.

distribution. Moreover, we are assuming that individuals know the true underlying distribution of risk; if individuals overstate/understate the probabilities of serious illness, they may value more highly/less highly having insurance against a bad health state.

We next use equation (12) to estimate the private willingness to pay to eliminate the variation in consumption do to the income loss from serious illness. We have already discussed how we measure p_j and L_j. We use the coefficients reported in the second column of Table 7 as estimates of ?. Since we do not have a direct estimate of the coefficient of relative risk aversion, we have evaluated (12) for a range of values of *?*, and found it to be fairly insensitive; we therefore use an estimate of 3, which is in the range of previous studies using individual micro-data (Zeldes, 1989; Engen, 1993).

The results of this exercise are presented in Table 5. In the first column, we present the expected loss as a percent of baseline consumption. These results indicate that unconditional expected costs of illness are quite low (.01 percent) as a share of permanent income, even when the costs conditional on illness are quite large as a share of income. The reason the unconditional costs are low is that major illness is a rare event. These results suggest that premiums for disability insurance could be quite low as a share of income. The second column shows the welfare loss from illness, which is (as noted above) a lower bound on the willingness to pay for insurance, as a percentage of ex-ante consumption. This is 0.36% of baseline consumption. Despite the large economic cost of these rare illnesses, this is once again quite small due to the infrequency of illness. But, while these welfare losses are small as a share of baseline consumption, they are fairly large relative to expected insurance payouts, as we show in the third column of Table 5. The gain is over 40% as large as the amount of insurance payout.

As noted in Table 5, the consumption effects of illness are larger for workers than for non-workers. If a disability insurance program was devised as in most developed countries, where individuals qualified for benefits through work, then the relevant question is the value of this insurance program to workers. This is shown in the second row of Table 9. Here, the willingness to pay is 0.5% of consumption, which amounts to over half of expected payouts.

Table 5: Welfare Gains from Insurance

	Expected Loss ÷ Baseline Consumption	Willingness to Pay ÷ Baseline Consumption	Willingness to Pay ÷ Expected Loss
Insuring Earnings Loss			
All	0.009 %	0.36 %	40.4 %
Workers	0.010 %	0.50 %	50.5 %
Insuring Medical Spending			
All	0.003 %	0.14 %	45.2 %
Workers	0.003%	0.16	51.6 %

These findings of large welfare gains from insurance, relative to expected payouts, suggest the potential for welfare improvements from government insurance provision. Only if the deadweight loss of government provision, through administrative costs, moral hazard, and the marginal cost of public funds, amount to more than 40-50% of expected payouts, will there be no welfare improvement from formal disability insurance.

The Insurance Value of Public Medical Care Subsidies

In this section we estimate the insurance value of these public medical care subsidies. First, we estimate what medical care expenditures would be in a world with no subsidies. User fees at public facilities are estimated to be about 10 percent the cost of providing care (World Bank, 1995). Using the same data, Gertler and Molyneaux (1996) estimated the price elasticity of demand for medical to be -0.4. Therefore, if prices were increased to the full costs of care - a ten fold increase - this would raise medical spending by 600%. The unsubsidized medical care expenditures for each change in ADLs is measured by the subsidized expenditure predicted from Table 3, increased by a factor of 6.

Using our estimates of p_j, we calculate the expected unsubsidized medical care expenditures arising from ADL changes. Doing so, we find that the expected increase in medical spending arising from ADL changes is about 0.32% of consumption. Even unsubsidized, expected medical care expenditures are only about one-third of the expected loss in earnings from illness.

The willingness to pay to eliminate the variation in consumption due to medical care expenditures, in a world with no formal disability insurance, is reported in the second panel of Table 5. The willingness to pay computation follows that above, but uses the unsubsidized value for medical care expenditures in addition to the earnings loss, raising the welfare cost of imperfect consumption smoothing. We then take the difference between the willingness to pay to insure total income loss, including unsubsidized medical spending, and the willingness to pay to insure the earnings loss only.

Not surprisingly, as we show in the second column of Table 5, the welfare gain as a percent of consumption is small, both absolutely and relative to the welfare gain from insuring earnings loss. But, as a percentage of expected medical care expenditures, the welfare gain is roughly similar to that of earnings loss, suggesting that unless the inefficiencies of subsidizing medical care exceed 40-50% of the cost of provision, there will be welfare gains from these subsidies. Relative to insuring earnings loss through disability, however, the inefficiencies from subsidizing medical care may be particularly large. This is especially true when one recognizes that raising user fees will mostly tax high frequency illnesses that individuals are able to smooth well; medical spending may be a poor instrument for discriminating more serious health shocks (consider the extreme case of instant death from an accident). Thus, while suggestive, our results do not prove that government subsidies of medical care expenditures are efficient, at least on consumption smoothing grounds.

V. CHINA

A. The Setting

During the 1980s and '90s, China astonished the world with annual GDP growth rates averaging over 8 percent in most years. These numbers mask the tremendous differences in economic growth across China's regions, as rural areas of the hinterland witnessed much slower growth over this period. Much of the same story can be told with respect to both access and quality of health care during the reform period. Variation in access can be quite striking -- in 1993, for example, the number of hospital beds per thousand was 5.50 for Shanghai and 6.24 for Beijing, but only 1.1 per thousand in China's rural areas (SSB, 1994; MPH, 1994).

Improved access to health care services actually predated the reform era. China's rural residents followed in the wake of a 1965 State Council directive that called for health reform and more emphasis on rural services, and led to the training of barefoot doctors and extension of basic health services down to the village (West, 1997). Dramatic improvements in health outcomes followed. The under-five mortality rate, for example, fell from 144 per thousand live births in 1965 to 60 in 1980, and further to 44 per thousand by 1990 (World Bank, 1996).[14] Since 1980, the basic institutions of the rural health care system developed in the 1960s have remained preserved within the post-reform township and village institutions.

The upper tier of the rural health care system run by the County Public Health Bureau is comprised of one or more general hospitals, an epidemic prevention station, and a maternal and child health station. Township health centers (*weishengyuan*) form the middle tier of the rural health care system and generally include a small hospital, an outpatient clinic, and an office for maternal and child health care. Finally, the lowest tier of the rural health system is the village health office (*weishengsuo*). Part-time rural health aides typically treat only common medical problems and refer serious illnesses to the township health center. In addition to family planning services, rural health aides are responsible for immunization treatment and provide general health education (West, 1997).

On the eve of reform in the late '70s, costs of health care for rural residents were quite low, and collective insurance coverage was quite broad. Dating to the 1950s, one of three forms of insurance were available to rural residents. Government health insurance (*gongfei yiliao*) available to state cadres, teachers, and township government workers was funded through cost-reimbursement. Labor health insurance (*laodong baoxian*) was provided to any rural residents employed in state-owned enterprises. Finally, cooperative health insurance (*hezuo yiliao*) was set up by villages and managed in coordination with township health centers (Wei, 1995; World Bank, 1996). By the early 1980s, the majority of China's rural population had access to some form of risk pooling arrangement that guaranteed access to care. Various estimates place the uninsured population at between 10 and 29 percent in 1981 (West, 1991, World Bank 1984). With introduction of the household responsibility system in agriculture and the dismantling of the communal welfare fund that supported cooperative health insurance, however, the uninsured population to exploded to almost 80 percent of rural households by 1993 (Wei, 1995). The villages used in the analysis of this paper reflect this trend. As of 1997, residents in only one quarter of them participated in any type of health insurance plan.

At present, the Ministry of Public Health is experimenting with different types of community financing programs for rural residents. Most small scale experiments provide a form of coinsurance requiring residents to pay part of the cost of health services at the point of service (World Bank, 1996). Whether health insurance should be compulsory or not has been a matter of debate within policy circles. While a 1994 policy announcement emphasized that participation in community organized coinsurance programs should be voluntary at the individual level (*Peoples Daily*, July 2, 1994), there have been reports of experiments with compulsory participation as well.

B. The Data

The analyses of health shocks in this paper use household survey data provided by the Survey Department of the Research Center on the Rural Economy (RCRE) at the Ministry of Agriculture in Beijing. Annual household surveys from 104 villages of Shanxi, Jiangsu, Anhui, Shandong, Henan and Shaanxi Provinces and the Shanghai Suburbs spanning the period from 1986 to 1997 are used to analyze both the impact of health shocks on household consumption, and evaluate potential demand for health insurance in rural China.

RCRE has collected data from a panel of households since 1986, the survey was not conducted in 1992 and 1994 due to funding difficulties. Households are asked a range of questions regarding income from on-farm activities

[14] The under-five mortalty rate is defined as the number of children who die between birth and their fifth birthday.

and off-farm employment, household consumption, land use, asset ownership, savings, formal and informal access to and provision of credit, and transfers from both village members and friends and family outside the village. Values of non-marketed grain that show up in income, consumption and grain balance sections of the survey are adjusted to reflect market prices following a procedure outlined in Chen and Ravallion (1996). County agricultural research offices that collect expenditure, income and labor allocation information from households on a monthly basis monitor the household survey. A staff person from the office works with households to clear up inconsistencies in the survey.

In the analysis below, we use changes in expenditures on health care to identify the occurrence and severity of a health shock. While small expenditures generally reflect routine preventative care administered at village or township health offices, large expenditures reflect catastrophic shocks that require visits to more expensive fee-for-service hospitals at the Township or County level.

C. Health Care Costs

Until the mid-1980s the rural health care system was financed largely through transfers from the central government. Since the mid-80s, however, there has been increasing pressure on local governments to contribute matching funds for the construction of health centers and offices. With respect to the provision of services, a switch from salaried compensation from local government budgets to fee for service financing exacerbated the effect of deteriorating mechanisms for risk-pooling at the village level, and contributed to the growth in household out of pocket expenses. Nationally, patient fees represented roughly 14 percent of total expenditures on health care in rural areas in 1980, but by 1992 patient fees accounted for 34 percent of total expenditures on health care (West, 1997). Data from the RCRE household survey also reflects this trend. From the table below, we see that in constant 1986 Yuan, the share of households spending more than 100 yuan on health expenditures doubled between 1987 and 1997. For the households in this sample, an expenditure of 100 Yuan is roughly eight percent of total household expenditures for the year, making this a sizable health shock.

As in other countries around the world, the data from this survey indicates that most expenditures on health care are devoted to treating the catastrophic illnesses faced by a small proportion of the population. The Lorenz curve below plots the cumulative share of the population versus cumulative share of health expenditures. It is evident from this graph that roughly 20 percent of the population in this sample account for 75 percent of the expenditures on health care. West (1997) has notes that in 1993 the household portion of payment for major surgery is roughly 150 Yuan, implying that expenditures of over 100 Yuan signify and major shock.

Table 6: Increase in Household Expenditures on
(Values in 1986 RMB Yuan)

Year	1987	1989	1991	1993	1995	1997
Average Household Consumption of Non-Durable Goods	1360	1280	1221	1221	1433	1342
Average Household Healthcare Expenditure	31	32	37	45	54	60
Share of Household Health Expenditures Over 100 RMB Yuan	0.09	0.11	0.11	0.14	0.17	0.18

Source: RCRE Rural Household Survey. Non-Durable Goods include food, clothing, fuel, and services.

Lorenz Curve of Health Care Expenditures

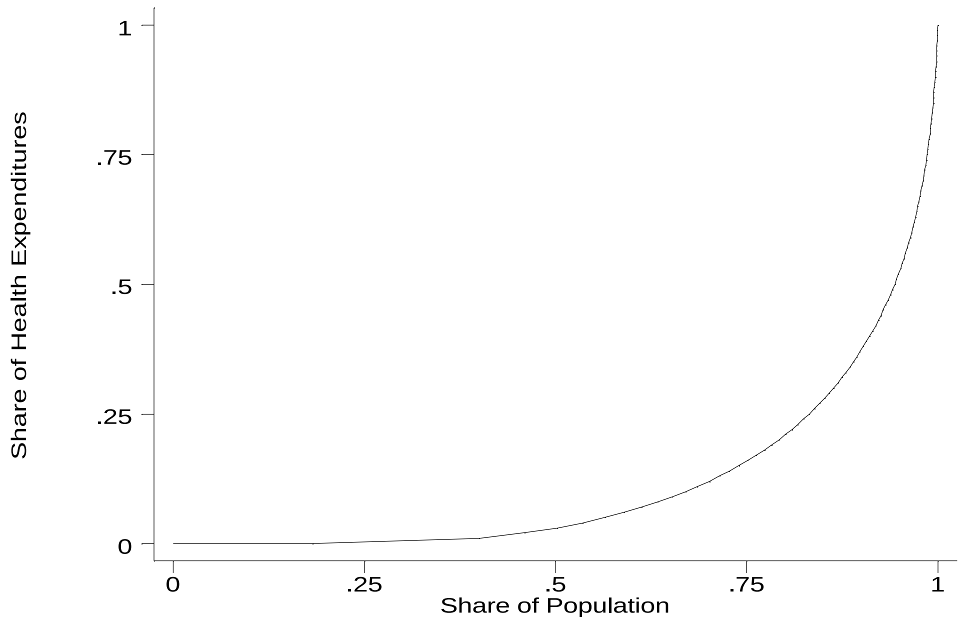

Source: RCRE Household Surveys, 1987-1997. Ministry of Agriculture, Beijing.

While households might be able to cover small health shocks relatively easily, catastrophic shocks have become a more serious problem with introduction of "fee for service" payment schemes in China's rural hospitals. In order to cover fees for serious illnesses, households must deplete savings if they exist or rely on informal loans from family members. Financing medical care to cover the serious illness of a family member can have serious long-term consequences for households in rural China. As the health policy community in Beijing strives to develop a new rural insurance system, the clear task they face is to find a way to insure against the economic impact of catastrophic shocks.

D. Consumption Smoothing and The Welfare Loss From Lack of Insurance

In the case of of Indonesia we observed the illness shocks directly. In China we only observe medical care expenditures. Moreover, since we do not observe income, we are not able estimate the impact of illness on income nor are we able to directly estimate (5). Instead, we estimate (5) replacing the change in income with the change in medical care expenditures. The coefficient times baseline consumption, then, is interpreted as how much consumption falls for each Yuan of medical care spending. However, the reason consumption falls might be either due to medical care spending or associated reductions in income from lost work due to illness. In the case of China we cannot separate the two.

By using the change in medical care expenditures in place of the change in income in (5) we avoid the measurement error problem. However, there could be feedback from the change in consumption to the change in medical care spending. Families with higher income spend more on medical care. The first order effect is solved by using a difference model which controls for heterogeneity in permanent income. Then by including village-time fixed effects, we controls for community level changes in permanent income. We also control for changes in human capital at the household level. Therefore, the remaining sources of concern include changes in household level idiosyncratic shocks to permanent income. While we cannot completely rule out feedback from changes in consumption to medical care expenditures, the degree to which it is likely to affect the results is limited. Moreover, feedback from consumption to medical care expenditures should be positive, while the effect of medical expenditures on consumption should be negative. Therefore, to the extent there is feedback, it biases the coefficient towards zero. This suggests that results are biased towards finding insurance. Hence, our estimates of the welfare loss from not being able to insure the economic costs of illness are lower bound estimates.

In Table 7 we report the results for two specifications of equation (5). The first specifies the change in medical care expenditures linearly and the second specifies them as piecewise linear splines with knots at 50 and 250. The splines allow nonlinear responses to the health shock. This specification tests the hypothesis that households are able to smooth small shocks, but less able to smooth big shocks[15].

The results overwhelmingly indicate that households are not able to fully insure medical care expenditures. In the linear specification, the coefficient on medical care expenditures indicates that 94 percent of medical care expenditures come out of household consumption or that households can smooth only about 6 percent of medical care costs. The spline specification, however, indicates that households are fully able to smooth small health shocks under 50 Yuan, but not able to smooth larger shocks above 50 Yuan. In fact, consumption falls by 1.06 Yuan for each Yuan of medical spending between for spending levels between 50 and 249, and consumption falls by 1.11 Yuan for each Yuan of medical spending for levels above 250. The results above 50 Yuan suggest could be interpreted as households not being able to insure all medical care costs above 50 Yuan or that there are other costs associated with large expenditures such as lost income, transportation and food costs associated with inpatient stays in hospitals.

Using the spline model, we calculated the willingness to pay to eliminate the variation in consumption due to medical care expenditures. We estimate that households are on average willing to pay 1.87 percent of household consumption. This is approximately six times average medical care spending.

VI. CONCLUSIONS

Using data from Indonesia and China, we find that while households are able to insure the economic costs of small health shocks, they are not able to insure the costs of major illness. The results indicate that there are large welfare gains from reducing the variation in consumption from the economic costs of major illness. However, households are unable to purchase insurance from the private sector due to market failures. This suggests a role for government to improve welfare through social insurance for income loss due to disability and medical care

[15] Note that about a 27.9 percent of the household have no medical care expenditures in a given year, 53.3 percent have less than 50 Yuan, 16.3 percent have expenditures between 50 and 250, and 2.5 percent have shocks greater than 250.

expenditures. The fact that households are able to insure against the costs of small illnesses suggests that employing user fees to finance the expansion of publicly delivered health care services will have little impact on welfare. Moreover, this result implies that first dollar capped social insurance benefit structures provide little welfare gain. Indeed, such benefit structures just crowd private informal insurance. Rather the big welfare gains are from insuring the rare large illnesses.

Table 7
Smoothing the Consumption Effects of Shocks to Non-Durable Consumption in Rural China
(Standard Errors in Parentheses)

Dependent Variable: Δln(Non-Durable Consumption)	Model 1	Model 2	Mean (Standard Deviation)
Lag Household Prime Laborers	-0.031 (0.001)	-0.031 (0.001)	2.522 (1.138)
Lag Household Dependents	-0.005 (0.001)	-0.005 (0.001)	1.852 (1.158)
Lag Male Share of Prime Laborers	0.026 (0.008)	0.026 (0.008)	0.516 (0.216)
Lag Land Per Capita	0.002 (0.003)	0.002 (0.003)	1.434 (1.012)
Lag Share of Laborers with Elementary Education	0.006 (0.006)	0.007 (0.006)	0.352 (0.329)
Lag Share of Laborers with Lower Middle School Education	0.015 (0.007)	0.015 (0.007)	0.323 (0.325)
Lag Share of Laborers with Upper Middle School Education	0.010 (0.010)	0.010 (0.009)	0.071 (0.188)
Lag Share of Laborers with Special Skill	0.007 (0.008)	0.007 (0.008)	0.102 (0.205)
Change in Health Expenditures	-0.00071 (0.00002)		36.530 (80.110)
Small Changes in Health Expenditures (Less than 50 1986 Yuan RMB)		0.00009 (0.00008)	
Medium Changes in Health Expenditures (50 to 249 1986 Yuan RMB)		-0.00080 0.00005	
Large Changes in Health Expenditures (Greater than 250 1986 Yuan RMB)		-0.00092 (0.00003)	
Number of Observations	45263	45263	45263
R^2	0.051	0.053	
*Village*Year Fixed Effects, F(780, 44473)*	9.567	9.626	

Note: Average Non-Durable Consumption for was 1321 RMB Yuan Per Household (in 1986 Yuan), with a standard deviation of 705 RMB Yuan.

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.

Attanasio, Orazio, and Steven Davis (1996). "Relative Wage Movements and the Distribution of Consumption," Journal of Political Economy, 104, 1227-1262.

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 Retirment 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.

Chen, Shaohua and Martin Ravallion (1996), ``Data in Transition: Assessing Rural Living Standards in Southern China," *China Economic Review*, Volume 7, Spring 1996.

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: Claredon Press.

Deaton, Angus (1992a). "Saving and Income Smoothing in Cote d'Ivoire," Journal of African Economics, 1, 1-24.

Deaton, Angus (1992b). Understanding Consumption. Oxford: Claredon Press.

Deaton, Angus (1994). The Analysis of Household Surveys: Microeconometric 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.

Dynarski, Susan and Jonathan Gruber (1997). "Earnings Instability and Family Living Standards," Brookings Papers on Economic Activity, forthcoming.

Engen, Eric (1993). "A Stochastic Life-Cycle Model with Mortality Risk: Estimation with Panel Data." Mimeo, UCLA.

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 P. and J Zeitlin (1996) "The Returns to Childhood Investments in Terms of Health Later in Life," mimeo, RAND, Santa Monica.

Gruber, Jonathan (1996). "Cash Welfare as a Consumption Smoothing Mechanism for Single Mothers". NBER Working Paper #5738.

Gruber, Jonathan (1997). "The Consumption Smoothing Benefits of Unemployment Insurance," American Economic Review, forthcoming.

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. Srinavasan (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.

Mace, Barbara (1991). "Full Insurance in the Presence of Aggregate Uncertainty," Journal of Political Economy, 99, 928-56.

Ministry of Public Health (MPH) (1994). *Chinese Health Statistical Digest, 1993*.

Morduch, Jonathan (1990). "Risk, Production, and Saving: Theory and Evidence from Indian Households". Mimeo, Harvard University, 1990.

Morduch, Jonathan (1995). "Income Smoothing and Consumption Smoothing," Journal of Economic Perspectives, 9 (Summer), 103-114.

Nelson, Julie (1993). "On Testing for Full Insurance Using Consumer Expenditure Survey Data," Journal of Political Economy, 102, 384-94.

Paxson, Christina (1992). "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," American Economic Review, 82, 15-33.

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.

State Statistical Bureau (SSB) (1994). *Zhongguo Tongji Nianjian 1994* (A Statistical Survey of China). Beijing: *Zhongguo Tongji Chubanshe*.

Stern, S. (1989). "Measuring the Effect of Disability on Labor Force Participation," Journal of Human Resources, 24, 361-95.

Stewert, 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. Srinavasan (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.

Wei, Ying (1995). "An Introduction to Health Financing Patterns in China," National Health Economics Institute (Beijing). Paper presented at a Health Care Conference in Beijing, May 2-5, 1995.

West, Lorraine (1997). "Provision of Public Services in the PRC," in **Financing Local Government in the People's Republic of China**, Christine P. W. Wong (ed.), Asian Development Bank, Oxford University Press, 1997.

World Bank (1984). "China. The Health Sector. A World Bank Country Study." Washington, DC, World Bank.

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.

World Bank (1996). "China: Issues and Options in Health Financing." Washington, DC, World Bank.

Zeldes, Stephen P. (1989). "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, 305-346.