Social Insurance in Developing Economies*

Raj Chetty, UC-Berkeley and NBER

Adam Looney, Federal Reserve Board of Governors

May 2005

Abstract

This paper assesses the potential welfare gain from introducing social safety nets in developing economies. Using panel surveys of households in Indonesia and the United States, we find that food consumption falls by approximately ten percent when individuals become unemployed in both countries. This finding is surprising given that the U.S. has an extensive social safety net while Indonesia has virtually none. Prior studies have interpreted such results as evidence that social insurance is of limited value in developing economies. We show that this conclusion is incorrect if the consumption path is smooth because individuals are highly risk averse. Exploratory tests suggest that Indonesian households are indeed quite risk averse because of subsistence constraints. These results imply that social safety nets may be valuable in low-income economies despite the smoothness of consumption.

Keywords: liquidity constraints, consumption smoothing, unemployment

*E-mails: chetty@econ.berkeley.edu, looney@frb.gov. This paper has been prepared for the Sixteenth NBER East Asian Seminar in Economics to be held in the Philippines. We have benefited from comments by Ted Miguel. James Sly and Bhuvan Jain provided excellent research assistance.

1 Introduction

Social safety nets in developing countries are far smaller than in developed economies. In 1996, the average expenditure on social insurance as a fraction of GDP in countries with below-median per capita income was 6.8 percent; the corresponding figure in above-median countries was 18.5 percent.[1] In the rapidly growing developing economies of South and East Asia, social insurance may be viewed as an unnecessary precaution that could potentially hamper growth without yielding substantial welfare gains. However, income shocks are prevalent in these economies. For example, at least 15 percent of households in the Indonesian Family Life Survey, a representative sample of 83% of the Indonesian economy, report some type of income shock in a given year. Recent large-scale shocks in this region such as the financial crises and the Asian tsunami further underscore the point that rapidly growing economies are not immune to large fluctuations. Anecdotal evidence suggests that inadequate insurance programs can lead to extreme suffering, at least for certain subgroups. Hence, studying the welfare consequences and optimal design of social insurance in developing economies is a particularly important issue from a public finance perspective. This paper takes a step in this direction by identifying the key factors that influence the welfare gain from social insurance programs in low-income economies.

Social insurance is only necessary if private insurance markets are inadequate. A straightforward and intuitive method of testing for full private market insurance frequently implemented in the development literature is to examine consumption fluctuations associated with shocks. As a starting point in our analysis, we compare the effects of unemployment on consumption in the U.S. and Indonesia. We use two large panel datasets that contain consumption and labor force data for each of these countries – the Panel Study of Income Dynamics (PSID) and the Indonesian Family Life Survey (IFLS). We compare the growth rate in consumption for agents who remain employed and agents who report job loss in the

[1] Source: International Labour Organization (2000). See section 2 for further details on this data.

two panels. The mean and median consumption drop associated with unemployment in both economies is roughly 10 percent . The similarity in the consumption drop is remarkable given that Indonesia has no formal UI system whereas the United States insures 50 percent of the pre-unemployment wage for most individuals. This finding appears to suggest that Indonesians are not underinsured relative to Americans despite their lack of a formal social safety net. Consistent with these findings, Morduch (1995) remarks that, "The emerging consensus of the empirical literature [on consumption-smoothing in developing economies] is that holes in effective [consumption] insurance exist.... But, in general, the holes are a good deal smaller than many had assumed.... The results have clear policy implications. If ... markets and alternative mechanisms do indeed provide reasonably good insurance and credit, publicly provided financial services and social security could crowd out private efforts with limited net gain to society."

To investigate the validity of this inference and estimate the welfare gains to social insurance formally, we draw upon some results from the public finance literature. The welfare gain from social insurance (ignoring efficiency costs caused by distortions in behavior) is determined by the product of the percentage consumption drop caused by the shock ($\frac{\Delta c}{c}$) with the coefficient of relative risk aversion (γ) in a general class of stochastic dynamic models (Baily 1978; Chetty 2005). Therefore, in order to understand whether a social safety net could be useful, one must determine the reason that $\frac{\Delta c}{c}$ is small. If it is small because agents have good private insurance, then social insurance may indeed be unnecessary. But $\frac{\Delta c}{c}$ could also be small because γ is large. An agent who is very risk averse experiences a sharp decline in utility when he cuts consumption, and will therefore undertake costly efforts to maintain consumption when hit by a shock. Since the welfare cost of a drop in consumption is proportionate to $\gamma\frac{\Delta c}{c}$, small consumption fluctuations may belie large welfare changes. Hence, if $\frac{\Delta c}{c}$ is small because of high risk aversion, social insurance could still yield large welfare gains.

To see why this point could be important in practice, note that much of the Indonesian

population appears to be close to a subsistence level of consumption. The average household in the Indonesian sample devotes nearly 70 percent of its budget on food, compared to 20 percent in the United States. It is plausible that Indonesians may be very reluctant to cut consumption when they face a shock. As a result, they may take costly actions such as removing children from school precisely in order to prevent large consumption fluctuations. But these behaviors have costs that could be reduced or avoided if a social safety net were available. It is therefore critical to investigate the degree of risk aversion to understand the cost of income fluctuations, especially in a low-income setting where γ could be quite large.

We make inferences about risk aversion in Indonesia using two simple, exploratory methods in this paper. First, we examine the effect of shocks on staple consumption goods such as rice, which presumably would be reduced only in the most dire circumstances. Many households do in fact consume significantly fewer staples during unemployment shocks. Second, we examine the methods households use to mitigate the income loss associated with unemployment. Strikingly, parents appear to reduce expenditures on children's education substantially during idiosyncratic unemployment spells. To the extent that these reductions permanently diminish children's educational attainment, the welfare costs of transitory unemployment shocks could be particularly large and long-lived. The fact that families go to such lengths to maintain food consumption further underscores the point that the marginal utility of consumption must rise quickly as consumption falls. In addition, more than 30 percent of households report raising labor supply to maintain their income stream. This high degree of responsiveness is further evidence that consumption-smoothing is costly for Indonesian households (i.e., it is not accomplished simply by depleting savings or borrowing).

Many of the empirical results in this paper are not new findings. Similar results on consumption, labor supply, and changes in educational attainment have been established in other contexts by existing studies (see section 3 for a review). Our objective here is to show how this type of descriptive evidence can be used in a tractable but general optimal social insurance model to answer the key policy question, "What are the welfare consequences of

social insurance?" We reach two conclusions in this regard. First, evidence that consumption fluctuations are small does not necessarily imply that existing insurance is "adequate" in developing economies. In fact, the converse may be true: Consumption may be smooth precisely because the welfare costs of consumption fluctuations are very high. Second, studies that identify other responses to shocks (such as reduced human capital investment or choice of less risky crops with lower expected yield) provide important information about risk aversion and thereby inform the policy question. These two observations, coupled with the empirical findings of this and previous papers, suggest that programs that reduce consumption fluctuations may provide greater welfare gains than suggested by prior work.

Of course, the costs of social insurance – e.g., reduced employment or opportunity costs such as forgone infrastructure or health investments – may also be very large. Several studies have found that social insurance programs have large distortionary costs in the United States (see Krueger and Meyer 2002 for a survey), and these costs could be even larger in developing countries. Therefore, one cannot conclude from the results here that introducing a large safety net will raise aggregate welfare. Hence, the most important lesson of this study is perhaps that further research on social insurance programs in developing economies could be very useful given their potential benefits.

The remainder of the paper proceeds as follows. The next section briefly describes existing social safety nets around the world. Section 3 compares the effects of unemployment on consumption in the United States and Indonesia empirically. Section 4 outlines a simple model of consumption responses to shocks to analyze the welfare gains from social insurance. Section 5 discusses tests of whether risk aversion can explain the smoothness of consumption in Indonesia. Section 6 offers concluding remarks.

2 Social Safety Nets in Developing Countries

The size of the formal government-provided social safety net is substantially smaller in developing countries than in developed economies. According to statistics collected by the International Labour Organization (2000) for 91 countries in 1996, the average GDP share of social insurance – defined as total expenditures on social security, disability insurance, unemployment insurance, insurance against work-related injuries, and government provided health insurance – was 12.5%, with a range spanning 0.7% to 34.7%. Figure 1a plots the fraction of GDP devoted to social insurance programs against GDP per capita for these countries (with log scales). There is a striking positive correlation between these two variables. As shown specification 1 of Table 1, a 1 percent increase in GDP per capita is associated with a 0.7 percent increase in the GDP share of social insurance in this cross-section. Perhaps more interestingly, the share of social insurance in government expenditure is also significantly higher in richer countries (Figure 1b and specification 2 in Table 1). Wealthier countries not only have higher government expenditure but also devote a larger fraction of that expenditure to social insurance.

Notably, the rapidly growing East Asian economies are on average 1.3 log units below the trend line plotted in Figure 1. In other words, they devote about 10 percentage points less of GDP to social insurance than other countries of similar income. East Asian economies devote on average 4.9% of their GDP to social insurance, compared with 16.5% in the U.S. and 22% in Europe. The positive relationship between GDP per capita and social safety nets is evident even among the small subsample of East Asian economies, with Indonesia having the lowest income and expenditure on social insurance and Japan having the highest of both.

These statistics understate the size of the social safety net in developing countries because they ignore other forms of in-kind and charity assistance, such as minimum food grants and NGO aid. However, these types of programs are generally quite limited in size (Gough et. al.

2004) and moreover have two features that considerably limit their scope relative to western social safety nets. First, they are often means-tested and so may not provide consumption smoothing benefits to a majority of the population. Second, aid tends to flow toward large-scale catastrophes (such as the recent tsunami), with much fewer funds available for the smaller but more numerous idiosyncratic shocks like unemployment or disability. Hence, although the welfarist approach to social insurance in developing countries is a potential consumption smoothing device in some instances, it is substantially different from a formal social insurance system with well defined taxes and event-conditioned benefits.

There are many reasons that developing countries might choose not to implement such social safety nets. The most plausible reason is perhaps that financing such systems is infeasible given limitations on the government's ability to extract revenue from standard sources (see Gordon and Li (2005, this volume) for a compelling argument on how such constraints shape government policies). While it is interesting and important to understand the political economy of social insurance in developing countries, the purpose of this study is to assess the normative value of such a program if it could be implemented. As the recent introduction of a formal unemployment insurance system in Korea suggests, some of these countries are reaching a point where such systems are feasible, making this normative question of substantial practical relevance.

3 Consumption-Smoothing in the U.S. vs Indonesia

Social insurance can only be beneficial if private insurance markets are incomplete or inadequate. The natural first step in determining whether there is a role for social insurance is therefore to test whether private insurance markets are adequate for agents to smooth consumption over shocks. The standard method of testing for full consumption insurance, originally implemented by Cochrane (1991) using U.S. data and Townsend (1993) using data on Indian farmers, is to directly examine the effect of idiosyncratic shocks such as job loss,

health changes, or weather shocks on consumption. Under the assumption that utility is additively separable over consumption and leisure, a drop in consumption associated with these shocks is evidence that insurance markets are incomplete. More recently, in the public finance literature, Gruber (1997) and Browning and Crossley (2001) have implemented tests of full insurance that do not rely on additive separability by examining whether the size of consumption drops during unemployment spells is related to the amount of government-provided unemployment insurance. Their estimates show that with full unemployment insurance, consumption would not fall at all during job loss, implying that most or all of the consumption fluctuations identified in prior studies are indeed attributable to incomplete insurance rather than complementarity between consumption and leisure.

Following this literature, we begin our comparison of the welfare gains of social insurance in developing versus developed economies by establishing consistent measures of consumption drops for a developing and developed economy for a specific shock. The shock we focus on is unemployment, since it is a well-defined and common event in both types of economies. We focus on the United States as the developed economy, primarily because of our familiarity with the institutions and the availability of the longitudinal PSID data there. We focus on Indonesia as the developing economy because it has high-quality panel data with a design very similar to the PSID. Indonesia also has minimal social insurance, making it an ideal laboratory in which to investigate the response of families to idiosyncratic shocks in a low-income economy without any social safety net.

Our methods and empirical results are borrowed from and consistent with a large body of prior work. Most relevant are studies that examine responses of Indonesian households to shocks. The general consensus of these papers on Indonesia and of the literature on developing countries more generally is that transitory shocks seldom translate into significant fluctuations in consumption because households have developed a variety of coping mechanisms. These mechanisms include spending down household wealth and assets or borrowing (Frankenberg, Thomas & Beegle 2003; Frankenberg, Smith, and Thomas, 2003),

increasing family labor supply (Beegle, Frankenberg, Thomas 2000; Cameron and Worswick 2003), and reducing investments in children's health and education (Frankenberg, Thomas, Beegle 1999; Thomas, Beegle, Frankenberg, Sikoki, Strauss, Teruel 2004). The smoothness of consumption has been taken to imply that economic shocks are not costly and that the scope for publicly provided social insurance for transitory shocks is small (Morduch 1995; Cameron and Worswick 2003).[2]

Our goal here is not to attempt to improve on the work of these prior studies, but rather to examine the normative conclusions that should be drawn from their empirical findings. To do so, we begin by establishing comparable measures of the consumption response to unemployment in Indonesia and the U.S.

3.1 Data

We use two household-level panel datasets in this study. The first is the Panel Study of Income Dynamics (PSID), which tracks approximately 8,000 households and their children over more than 30 years in the United States. We use an extract of the PSID that contains consistently defined annual data between 1980 and 1993. The second is the Indonesian Family Life Survey (IFLS), which follows roughly 7,500 households over a span of 7 years, with interviews in 1993, 1997, and 2000.

To examine the impact of unemployment shocks, we focus on households for which longitudinal data exists and with household heads who were employed at the time of the immediately preceding interview. Hence, we include only households where the head was employed one year before the current interview in the PSID, and three or four years before in the IFLS. We discuss below how the lack of annual data in the IFLS could affect the comparison between the datasets.

[2]Studies which examine large, persistent health shocks in Indonesia (Gertler and Gruber (2002), Gertler, Levine, and Moretti (2001)) do find large consumption drops. However, Gruber and Gertler observe that their results offer "little insight into consumption smoothing of more likely and less costly risks" that are our primary focus here.

Table 2 provides summary statistics on these households. Inflation in Indonesia was high over this time period, largely due to the 1998 financial crisis. The price level rose an average of 91 percent in the 3-4 year periods between interviews. In comparison, average annual inflation in the U.S. was 5 percent over our sample. The IFLS statistics reported in Table 2 are deflated using an aggregate CPI series from the Asian Development Bank and are converted to year 2000 dollars using the US/Rupiah exchange rate as of January 2000. The PSID statistics are deflated using the standard CPI series from the BEA. Note that real food consumption growth rates are small in both samples. In our empirical analysis, we use nominal growth rates for transparency, since inflation rates are thought to differ significantly across goods and regions in Indonesia around the financial crisis. Naturally, nominal growth rates are much higher in the IFLS than in the PSID.

The most striking differences between the samples are in economic characteristics. PSID household heads earn on average $31,828 per year and PSID households consume $7,255 of food per year ($2,687 per person). In contrast, IFLS households report average total incomes of $1,484, and consume approximately $926 in food each year ($162 per person). Note that this figure includes food purchased and food produced (important given the large number of farmers in the data). Unemployment shocks appear more frequently in the IFLS data; approximately 8 percent of heads of household become unemployed between interview waves while 4 percent become unemployed between years in the PSID. The median household holds total assets of $2,692, but much of this wealth is in the form of illiquid assets such as farms.

Because of data constraints we define unemployment spells slightly differently in the two samples. In the PSID, a household head is defined to be unemployed if he or she is not working and searching for a job at the time of the interview. Replicating this measure in the IFLS is not always possible because weekly employment data (module TK) for the 1997 interview has not yet been publicly released. Instead, we use a question corresponding to employment status during the last 12 months. In 1993 and 2000, when both weekly

and annual employment statistics are available, these measures are highly correlated and we find that the effects of unemployment on consumption are very similar regardless of which variable is used. We use the annual employment variable to maximize the sample size and to avoid focusing only on changes in outcomes over seven years, as required if we dropped 1997 interview information.

A concern with our definition of unemployment in the IFLS data is that the IFLS annual employment variable provides little detail on employment status, so that we cannot always differentiate involuntary unemployment from endogenous transitions out of the labor force such as retirement. Papers by Frankenberg, Thomas, and Beegle (cited above) are able to address this issue better as they have access to the unpublished data. The results we report below are very similar to their results. In addition, when we restrict the sample to cases where we do know whether the individual is still in the labor force, we obtain similar point estimates. We are therefore fairly confident that this data issue is not a significant source of bias.

3.2 Results

We begin our analysis with a simple comparison of growth rates of food consumption in the U.S. and Indonesia. Define the growth rate of food consumption for household i from year t to year t' as

$$g_{it} = \log c_{it} - \log c_{it'}$$

where c_{it} denotes household i's food consumption in period t. Ideally, the gap between t and t' would be small, but in Indonesia data is available only every 3-4 years while in the U.S. data is annual. In the baseline analysis, we attempt to get as close a measure to the true drop as possible in each dataset, by examining the growth rate from t to $t+1$ in the U.S. and t to $t+3$ or $t+4$ (as data permits) in Indonesia.

Our basic identification strategy is to divide our sample of employed heads-of-household

in the pre-period year t' into two groups: Job losers, who reported being unemployed at the time of the survey in year t, and job keepers, who reported still having a job. We then compare various moments of the distr-ibution of growth rates for these two groups to estimate the effect of unemployment on consumption. The key identification assumption that must hold for this method to give a consistent estimate of the causal effect of unemployment on consumption growth is that the "treatment" group of job losers and the "control" group of job-keepers have identical consumption growth rates absent the shock. This identification assumption may be questionable given that individuals prone to job loss are generally lower skill types, and therefore may have relatively lower rates of trend wage and consumption growth in a society with increasing income inequality. In this case, the simple differences below will overstate the true consumption drop caused by unemployment. We implement some tests to address this concern below.

We first demonstrate the effect of unemployment on consumption using the long time series available in the PSID data with an "event study" method in Figure 2. This figure is constructed by redefining as year 0 the year of job loss for the set of household heads who lost their jobs once during the PSID sample. We then plot real average annual consumption growth rates (more precisely, change in real log household consumption) against year relative to year of job loss (e.g. -3 is 3 years before job loss). The figure shows that consumption grows at a real rate of roughly 2-4 percent per year before time 0, and then drops by nearly 10 percent in the year of job loss. Consumption then recovers gradually over the next few years back to its original level. This graph confirms that unemployment causes a sharp, temporary decline in consumption for the typical household in the United States, consistent with the results of Cochrane (1991) and Gruber (1997). Unfortunately, a similar graph showing a long pre-event and post-event period cannot be drawn for Indonesia, since there are at most three observations per household in the IFLS. We are therefore forced to compare single observations on growth rates in consumption from time -1 to time 0 across job losers and job keepers to identify the effect of unemployment in the IFLS. We adopt a similar strategy

in the PSID for purposes of comparability.

We begin our comparison of Indonesia and the U.S. with a nonparametric, graphical analysis of the effect of job loss on consumption. We estimate kernel densities for the distribution of nominal growth rates by employment status in each country. Following the convention in the consumption growth literature (see e.g. Zeldes 1989 or Gruber 1997), we trim outliers (the lower and upper 2 percent of the reported distribution), though our results are insensitive to this restriction. The kernel densities are estimated using an "optimal" bandwith chosen to minimize the asymptotic mean squared error of the fitted distribution.

Figure 3a plots the density of growth rates for job losers (red curve) and keepers (blue) in Indonesia. It is clear that unemployment leads to a left-shift in the distribution, indicating that households are unable to fully smooth consumption over this transitory shock. The medians of each distribution are depicted by vertical lines of corresponding color. The median nominal growth rate of food consumption for job keepers in the sample is 67 percent (due to the high rate of inflation in Indonesia over this period), in comparison with a growth rate of 56 percent for job losers. Hence, at the median, unemployment appears to reduce food consumption by approximately 11 percent.

Figure 3b plots analogous curves for the United States. Again, it is clear that agents are not fully insured, consistent with the results of Gruber (1997). Of greater interest here is the comparison of these distributions to their analogs in Indonesia. The distribution of growth rates reported by Indonesian households has variance twice as high as that in the U.S., which could be either because of measurement error or because outcomes in developing countries tend to be more stochastic. Despite this general difference in the distributions, the within-sample difference between job losers and job keepers is strikingly similar. In the U.S., the median nominal growth rate for job keepers is approximately 8.5 percent, compared to -1.5 percent for job losers. Hence, job loss appears to reduce food consumption by approximately 10 percent in the United States, only 1 percent different from the Indonesian value. Other

quantiles of the distribution shifts are also quite similar across the two economies.[3]

We now examine whether robustness of this conclusion to controls using a more structured regression analysis. We estimate specifications of the following form:

$$g_i = \alpha + \beta \text{unemp}_i + \theta X_i + \varepsilon_i \tag{1}$$

where $\text{unemp}_i = 1$ if the agent reports unemployment at time t', $\text{unemp}_i = 0$ if the agent is employed at time t', and X_i denotes a vector of covariates. The key coefficient β equals the effect of job loss on the consumption growth rate.

Table 3 reports several estimates of (1) for Indonesia and the U.S. The first specification is estimated with OLS using no controls except year dummies. Consistent with the graphical results, unemployment is estimated to reduce consumption by about 9 percent in the U.S. and 10 percent in Indonesia. The second specification introduces several controls: age, gender, marital status, education, and region dummies (to control for differential inflation patterns). The coefficient estimates on the unemployment dummy are essentially unchanged. These results show that after controlling for observable heterogeneity in trend growth rates across job losers and job keepers, consumption drops remain quite similar in the two countries.

The third specification tests the "common trends" identification assumption more directly by restricting the sample to individuals who lost jobs at some point within the panel. In this specification, the counterfactual for the job losers in year t' are individuals who lost their jobs at some other point in the dataset. The advantage of this specification in terms of identification is that growth rates in consumption for job losers are compared to what is arguably a better "control" group. The problem of unobservable differences between job losers and keepers is mitigated in the restricted sample by identifying purely from variation in the date of job loss rather than whether or not job loss occurred. As shown in the last

[3]The estimated consumption drops become *larger* in the PSID if we use changes from t to $t+3$ (as in the IFLS). Hence, using a comparable strategy across the two datasets only further reinforces the point that consumption is as smooth during shocks in Indonesia as it is in the U.S.

two columns of the table, this smaller sample yields estimates that are generally similar to the original results, supporting the claim that the causal effect of the unemployment shock on consumption is being identified.

We also conducted a series of other robustness checks and sensitivity analyses that are not reported in the table. Quantile regressions generally yield estimates very similar to the OLS results. Different trimming criteria for outliers, such as 1 percent or 5 percent also yield similar results. Broader measures of consumption also follow a similar pattern. Gruber (1998) augments the results from the PSID with broader measures of consumption from the Consumer Expenditure Survey and finds that the decline in total consumption mirrors that of food consumption. We find a similar decline in total consumption in the IFLS sample as well (not reported).

To summarize, the evidence from the IFLS and the PSID suggests fairly strongly that idiosyncratic unemployment shocks lead to consumption fluctuations of similar magnitude in the U.S. and Indonesia. Even if there are differences across the data in measurement error and collection procedures that affect the comparability of the estimates, one can at least rule out the hypothesis that the drop in Indonesia is far larger than in the U.S given the evidence above. This similarity is surprising given that the U.S. has a large UI system whereas Indonesia has very little formal social insurance (Figure 1).

On the surface, these results suggest that families in developing economies have "adequate" insurance because they are able maintain a reasonably smooth consumption path when faced with shocks, as originally suggested by Townsend's (1993) classic study of Indian farmers. The apparent normative implication is that social insurance would offer relatively modest welfare gains in the Indonesia (and perhaps other developing economies). The next section examines the validity of this conclusion more closely.

4 A Normative Framework

Though many papers have studied consumption-smoothing in developing economies, none have explicitly investigated the potential welfare gains of implementing a social insurance program. To do so, consider a general dynamic lifecycle model where agents maximize expected lifetime utility and face employment shocks based on an arbitrary stochastic process. Agents have utility over consumption and N other choice variables and face M arbitrary constraints (e.g. borrowing constraints) in making their decisions. Suppose that a constant social insurance benefit of b is provided in the unemployed state, and this benefit is fully financed by taxation in the employed state. Chetty (2005), building on the canonical analysis of Baily (1978), derives a formula for the optimal level of social insurance under some weak regularity conditions in this environment. Assuming that third-order terms of utility are small (i.e. ignoring precautionary savings motives), the formula is:

$$\frac{\Delta c}{c}(b^*)\gamma \approx \varepsilon_{d,b} \tag{2}$$

where $\varepsilon_{d,b}$ is the elasticity of unemployment durations with respect to benefits, which captures the moral hazard cost of benefit provision due to behavioral response; $\frac{\Delta c}{c}(b)$ is the average observed drop in consumption as a function of social insurance benefits b, which quantifies the consumption-smoothing benefits; and $\gamma = -\frac{u_{cc}}{u_c}c$ is the coefficient of relative risk aversion, which reflects the value of having a smoother consumption path. The left hand side of this formula quantifies the marginal change in expected utility from increasing the benefit level by \$1, which arises from a smoother consumption path. The right hand side quantifies the marginal efficiency cost, which arises from distorted prices. At the optimal benefit level, the marginal costs and benefits are equated.

The important aspect of (2) for our purposes here is that the marginal welfare gain from raising b by \$1 depends on the *product* of γ and $\frac{\Delta c}{c}$, not just the latter term. This is especially relevant because γ and $\frac{\Delta c}{c}$ are inversely related – it is precisely in situations where γ is high

that agents will try to keep $\frac{\Delta c}{c}$ small. To see this point and its normative consequences, consider a stylized static example where agents have no ex-ante savings. Suppose the agent has CRRA utility over food consumption

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

Let the disutility of supplying $\$c$ of consumption be given by a linear function

$$\psi(c) = \theta c$$

A shock such as unemployment or injury can be modeled in this framework as an increase in θ, which makes earning money more difficult. When the worker is employed, θ captures the disutility of regular labor. When he loses his job, θ rises because financing $\$c$ of consumption now requires undertaking more costly activities like job search, increased spousal labor supply, or reduced human capital or health investments in children. To simplify matters, suppose that there are two states (employed and unemployed), with $\theta_u > \theta_e = 1$. With this normalization, θ_u can be interpreted as how much more difficult it is to earn money in the bad state than the good state. For example, $\theta_u = 2$ implies that the disutility of generating consumption is doubled.

In this setting, the worker chooses consumption in each state by solving

$$\max_c \frac{c^{1-\gamma}}{1-\gamma} - \theta c$$

Hence

$$c^*(\theta) = \theta^{-1/\gamma}$$

The consumption drop from the employed to the unemployed state is therefore

$$\frac{\Delta c}{c} = \frac{c_e - c_u}{c_e} = 1 - \frac{c_u}{c_e} = 1 - \left(\frac{1}{\theta_u}\right)^{1/\gamma}$$

This expression shows that $\frac{\Delta c}{c}$ is decreasing in γ and increasing in θ_u. Intuitively, high γ makes consumption reductions particularly costly, and the agent therefore exerts greater effort in the unemployed state to maintain consumption close to the pre-employment level. Similarly, high θ_u makes earning income while unemployed particularly costly, making it preferable to tolerate a larger consumption drop. These comparative statics indicate that the $\frac{\Delta c}{c}$ observed in the data could be small for two independent reasons: (1) θ_u is low, i.e. agents are able to easily and inexpensively smooth consumption by borrowing or through informal insurance mechanisms or (2) γ is high, i.e. agents are very risk averse to fluctuations and work hard to have a small consumption drop even though θ might be high. In case 1, the marginal welfare gain from social insurance $\gamma\frac{\Delta c}{c}$ is likely to be small. In contrast, in case 2, the gain from social insurance could be quite large even if $\frac{\Delta c}{c}$ is small because γ may be very high.

Table 4 illustrates this point quantitatively by showing the implied consumption drop and welfare gain for a range of γ and θ_u. Part A of the Table shows that a consumption drop of $\frac{\Delta c}{c} \approx$ 10-13 percent as estimated in the data can be generated by a variety of combinations of γ and θ_u, indicated in bold on the diagonal of the table. Part B shows that the welfare implications implied by the different combinations above can vary widely. With high γ and θ_u, the marginal gain in expected utility from the provision of an extra dollar of social insurance can be three times as large as the gain with low γ and θ_u.

To understand this point intuitively, consider two different descriptions of the economy in Indonesia, both of which could generate a relatively small consumption drop of 10 percent. In the first scenario (low γ, low θ), agents have access to credit markets and networks that allow them to smooth consumption easily when hit by a shock. In this case, a mandated

social insurance program would simply crowd out existing private market arrangements, without much of a welfare gain. In the second scenario (high γ, high θ), private market insurance arrangements are very poor. However, households are close to a subsistence level of consumption while employed, and are therefore very reluctant to cut consumption further when they are hit by a shock. They therefore use costly, high θ methods to avoid a substantial consumption drop, such as taking children out of school. In this case, the provision of social insurance could yield large welfare gains, because it reduces the need to rely on costly consumption-smoothing mechanisms when hit by shocks.

This analysis shows that estimates of the consumption fluctuations associated with shocks are inadequate to compute the welfare gains from social insurance. It is equally important to estimate the degree of risk aversion of households over food gambles to answer this question. The next section shows how existing evidence on other behavioral responses to shocks can be used to infer γ, and also provides some suggestive evidence on this parameter from the IFLS.

5 Suggestive Evidence on Risk Aversion

A natural way to infer risk aversion would be to give people real or hypothetical gambles in an experimental setting. as in Binswanger's (1981) pioneering study of risk aversion in rural India. Unfortunately, a growing body of evidence suggests that this method of preference elicitation leads to misleading and often contradictory estimates of risk aversion, partly because of framing effects and the limited size of gambles feasible in experiments (see e.g. Rabin 1999 or Starmer 2000 for a recent synthesis). If a social planner wants to maximize an agent's expected utility, he should be especially interested in the *experienced* marginal utility of food at different levels of food consumption, irrespective of whether those marginal utilities are the basis of decisions. Hence, it is helpful to have some measures of γ that are not completely reliant on the axioms of expected utility holding in experiments. We now

provide a few suggestive pieces of evidence about γ in Indonesia along these lines.

The simplest indicators that risk aversion could be quite high in Indonesia are summary statistics about expenditures. The average household in the IFLS reports annual expenditures of \$1,400, with nearly 70 percent of that sum devoted to food. In contrast, in the United States, the average household spends \$7,255 per year, only 20 percent of which is devoted to food. Part of the difference is obviously due to different price levels across the countries. Nonetheless, it is clear that a large share of Indonesian families' expenditures are devoted to basic necessities, suggesting that any reduction in expenditure could be quite costly for them (γ is high).

We attempt to infer the utility cost of a reduction in expenditures in two ways. First, we attempt to establish direct evidence that γ is high by examining changes in consumption of staples such as rice, which presumably would fall only if agents were unable to reduce consumption on "luxuries" which have lower marginal utility. We implement empirical specifications analogous to (1) to test whether the consumption of staple foods (including rice, corn, cassava, and flour) falls in households experiencing unemployment shocks relative to households that do not experience such shocks. The regression sample specifications and trimming procedures are analogous to those described above for total food consumption.

We begin with an OLS regression on the full sample. The estimate in column (1) indicates that mean consumption of staple foods falls by 6 percent during unemployment spells; however the estimate is not statistically significant. As one might expect, the magnitude of this decline is smaller than the drop total food consumption (see Table 3) and total consumption (not shown) because households are presumably more willing to cut back on "luxuries" than "necessities." A kernel density plot (not shown) for growth in staples consumption by job status analogous to Figure 3a reveals a clear downward shift in consumption of staples for job losers who experience the most negative growth rates, but little shift for those who fared better. This is consistent with the claim that only the worst off reduce consumption of staples. This suggests that even though the change in the mean growth rate may not

be statistically significant, other moments could reveal a more robust response. Column (2) confirms this point by showing that median staples growth rate is 10 percent lower for job losers relative to keepers. This estimate is highly statistically significant. Column (3) shows that the mean drop in staple consumption is 12 percent among households without any farmers, who might have less capacity to store crops. In sum, these results indicate that many households do reduce consumption of the most basic and important sources of nutrition when the household head loses his job. These findings are consistent with those of Beegle, Frankenberg, and Thomas (2000) and Frankenberg, Smith, and Thomas (2003), who study the effects of the 1998 Asian Financial Crisis on consumption using an augmented IFLS sample.

Our second approach is to establish that households use costly (high θ) methods to smooth consumption. Based on the model in the previous section, evidence that households resort to costly smoothing methods is evidence that γ is high. One such costly method is reducing educational expenditures on children. Intuitively, if families find it worthwhile to pull children out of school to deal with a temporary idiosyncratic shock, they must be very reluctant to cut consumption on food and other goods, and therefore must have high γ.

The first three specifications in Table 6 report the effect of unemployment shocks on educational investment. In these regressions, we restrict the sample to households with children under 24 years of age who reported educational expenses at the time of the previous interview. Specifications (1) and (2) examine extensive-margin (participation) effects by using a dummy for positive household educational expenditure as the dependent variable. The results reported in column (1) imply that families experiencing unemployment were 13 percentage points more likely to stop spending on education entirely (presumably by withdrawing their children from school). This is a large reduction relative to the sample mean of 77 percent participation in education in this group. Controlling for household characteristics reduces the estimated magnitude of this response slightly, but does not alter the conclusion that unemployment shocks significantly reduce the likelihood a household will

spend on education. Column (3) examines the intensive margin by changing the dependent variable to the log change in education expenditures (with 2% trimming as above). Median educational expenditure falls by 12 percent in households experiencing unemployment. Average educational spending (not shown) falls by less than this 12 percent, largely because richer households do not appear to reduce expenditures as much as poorer households, for reasons similar to the staples results. Figure 4 shows the distributional shift on the intensive margin, confirming the regression results visually.

Taken together, these results suggest that many households reduce spending on education to mitigate the income loss during an unemployment shock. A concern with the interpretation of these results is reverse causality. One might worry that families with children who finish school are those where the parent stops working, generating the observed correlation. However, it is reassuring that Frankenberg, Thomas, Beegle (1999) and Thomas et. al. (2004).have documented similar patterns in educational expenditure among households affected by the Asian Financial Crisis. These studies address the causality and identification concerns much more carefully, suggesting that shocks cause reductions in education.

A second behavioral response, which perhaps has a lower cost than reduced human capital accumulation but is nonetheless much more costly (higher θ) than borrowing, is augmented labor supply by other members of the household. If the loss of income due to unemployment imposes significant hardship, other household members should be expected to try to re-coup that lost income though increased market work.

Columns (4)-(6) of Table 6 provide evidence of labor supply responses. On the extensive margin, Column (4) shows that other household members are 17 percentage points more likely to work for wages when the head of household becomes unemployed. Controlling for other household characteristics does not significantly affect this conclusion. Column (6) examines the income earned by other family members on the intensive margin with a specification analogous to (3) for educational expenditures. The point estimate suggests that income earned by other household members increases by between 11 percent in house-

holds where the head becomes unemployed. Figure 5 corroborates this result visually. These results suggest that unemployment shocks increase the labor supply of other family members along a variety of margins. Part of these effects may again be due to reverse causality. But other studies (Beegle, Frankenberg, Thomas (2000), Cameron and Worswick (2003), Frankenberg, Smith, and Thomas (2003)) report similar responses in terms of labor market participation, second jobs, and additional hours of work among household members using richer data where causality is clearer and much more effort is devoted to identification concerns.

Interestingly, the magnitude of the "added worker effect" estimated here contrasts sharply with corresponding patterns in the United States. Cullen and Gruber (2000) observe that there is no change in labor supply of secondary earners at the mean when household heads lose their jobs in the U.S.. This difference suggests that the social safety net and financial system in America relieves families of the burden of replacing lost income through the costly methods used by Indonesian households.

In summary, the simple tests implemented here suggest that the small consumption drops in Indonesia are a consequence of high risk aversion rather than cheap consumption smoothing technologies. Based on the preceding theoretical analysis, this result implies that the welfare gains from social insurance (ignoring efficiency costs due to behavioral distortions) could be large despite the small consumption drop observed in the data. These gains would arise because households would not be forced to pull children out of school or send additional members into the workforce to maintain consumption in the short run.

It is important to stress that the methods used here to infer risk aversion are only illustrative and could be improved upon considerably in subsequent work. In addition, the behavioral responses examined here are only two examples among many possibilities. Examining the costs of other consumption-smoothing methods used by households in the context of a model similar to that above would be very useful.

6 Conclusion

An extensive literature in public finance has examined the costs and benefits of social safety nets in developed economies such as the United States and Europe. An equally large literature in development has tested for full consumption insurance by estimating the effect of agricultural and health shocks on consumption in developing economies and examining the mechanisms used to smooth consumption. Few studies, however, have applied the analytic tools of the public finance literature to assess the potential gains from formal social safety nets in developing economies. This paper has shown how existing evidence from the development literature can be used in a simple but general normative framework to evaluate the benefits of social insurance.

Our main conclusion from this analysis is that the welfare gains from government provision of social insurance cannot be quantified simply by estimating effect of shocks on consumption smoothing. It is important to determine *why* the path of consumption is relatively smooth in the presence of shocks by examining other behavioral responses. This research agenda is especially relevant for South and East Asian economies as they reach a phase of development where implementation of a formal social safety net is feasible.

References

Baily, Martin, "Some Aspects of Optimal Unemployment Insurance," Journal of Public Economics 10 (1978), 379-402.

Beegle K, Frankenberg E, Thomas D. 2000. "Labor Market Transitions of Men and Women During an Economic Crisis: Evidence from Indonesia." Rand Labor and Population Working Paper 00-11.

Binswanger, Hans P., "Attitudes Toward Risk: Theoretical Implications of an Experiment in RuralIndia," Economic Journal, 91, December 1981, 867-890

Browning, Martin and Timothy Crossley, "Unemployment Insurance Levels and Consumption Changes," Journal of Public Economics 80 (2001), 1-23.

Cameron, L. and C. Worswick 2003. "The Labor Market as a Smoothing Device: Labor Supply Responses to Crop Loss", Review of Development Economics,7(2), 327-341.

Chetty, Raj, "A General Formula for the Optimal Level of Social Insurance," UC-Berkeley mimeo, 2005.

Cochrane, John H, "A Simple Test of Consumption Insurance." Journal of Political Economy 99 (1991), 957-76.

Crossley, Timothy and Hamish Low, "When Might Unemployment Insurance Matter? Credit Constraints and the Cost of Saving," 2004.

Cullen, Julie and Jonathan Gruber, "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" Journal of Labor Economics 18 (2000), 546-572.

Frankenberg E, Thomas D, Beegle K. 1999. "The Real Costs of Indonesia's Economic Crisis: Preliminary Findings from the Indonesia Family Life Surveys." DRU-2064-NIA/NICHD. Santa Monica, CA: RAND.

Frankenberg, E, Smith, J, and Thomas, D. 2003. Economic shocks, wealth and welfare, Journal of Human Resources, 38.2: 280-321.

Gertler, P. and J. Gruber, "Insuring Consumption Against Illness," The American Economic Review, March (2002), vol. 92, no. 1, pp. 51-70(20).

Gertler, P., E. Moretti, and D. Levine "Do Microfinance Programs Help Families Insure Consumption Against Illness?" Mimeo. (2001)

Gough, Ian, et. al., *Insecurity and Welfare Regimes in Asia, Africa and Latin America : Social Policy in Development Contexts*, Cambridge University Press, 2004.

Gordon, Roger and Wei Li, "Tax Structures in Developing Countries: Many Puzzles and a Possible Explanation," UC San Diego mimeo, 2005.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," American Economic Review 87 (1997), 192-205.

Gruber, Jonathan, "Unemployment Insurance, Consumption Smoothing, and Private Insurance: Evidence from the PSID and CEX," Research in Employment Policy 1 (1998), 3-31.

International Labour Organization, *World Labour Report 2000: Income Security and Social Protection in a Changing World*, 2000.

Krueger, Alan and Bruce Meyer, "Labor Supply Effects of Social Insurance," NBER Working Paper 9014, 2002.

Morduch, Jonathan, 1995. "Income Smoothing and Consumption Smoothing." Journal of Economic Perspectives, 9(3) , 103–114.

Rabin, Matthew, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," Econometrica, 68 (2000), 1281-1292.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. Journal of Economic Literature. 38, 332-382.

Thomas D, Beegle K, Frankenberg E, Sikoki B, Strauss J, Teruel G. 2004. "Education in a Crisis." Journal of Development Economics, 74(1): 53-85.

Townsend, Robert, ""Risk and Insurance in Village India", Econometrica, 62(3), May 1994, 539-591.

Zeldes, Stephen, "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, University of Chicago Press, vol. 97(2), pages 305-46, 1989.

Table 1

RELATIONSHIP BETWEEN SOCIAL INSURANCE AND PER CAPITA GDP

	(1) SI as % of GDP vs GDP	(2) SI as % of Gov't Exp vs GDP	(3) Continent dummies	(4) East Asian Countries
Dependent Variable:	log SI % of GDP	log SI % of Gov't Exp	log SI % of GDP	log SI % of GDP
log GDP Per Capita	**0.630**	**0.351**	**0.357**	**0.674**
	(0.070)**	(0.064)**	(0.069)**	(0.062)**
Constant	-3.376	0.267		-3.673
	(0.626)**	(0.589)		(0.550)**
East asia indicator				**-1.318**
				(0.250)**
Continent dummies	No	No	Yes	No
Observations	89	64	89	89

NOTE--Social Insurance statistics are from ILO (2000); GDP statistics are from the Penn World tables.
Social insurance is defined as sum of expenditures on ocial security, disability insurance, unemployment insurance, insurance against work-related injuries, and government provided health insurance.
East Asian countries in the sample are Indonesia, Japan, Korea, Malaysia,Thailand, and Singapore.
GDP is measures in 1996 US dollars.

Table 2
SUMMARY STATISTICS FOR IFLS AND PSID

	Mean	Median	Standard Deviation
IFLS (Indonesia)			
Currently Unemployed	8%	0	27%
Age of Head	48	46	25
Married	83%	1	37%
College	6%	0	24%
Number of people in household	5.7	5.0	2.5
Household food consumption (annual)	$926	$703	$1,065
Real food consumption growth rate	4%	3%	61%
Inflation rate	91%	132%	42%
Staples consumption	$191	$144	$247
Total consumption	$1,604	$1,073	$2,047
Wage income of head	$580	$308	$1,056
Other family members earn income	58%	1	49%
Total household income	$1,484	$811	$3,569
Total household assets	$7,525	$2,692	$17,189
No household member is a farmer	58%	1	49%
Education expenditure	$144	$49	$344
Positive education expenditure	77%	1	42%
Number of obs: 12,236; Number of households: 7,197			
PSID (United States)			
Currently Unemployed	4%	0	21%
Age of Head	38	36	12
Married	65%	1	48%
College	40%	0	49%
Number of people in household	2.7	3.0	1.4
Household food consumption (annual)	$7,255	$6,303	$4,646
Real food consumption growth rate	2%	3%	56%
Inflation rate	5%	4%	2%
Wage Income of Head	$31,828	$27,285	$30,267
Number of obs: 70,889; Number of households: 11,685			

NOTE-Sample includes all households who remain in panel for two or more years where head is employed in previous observation (1 year lag in PSID, 3 or 4 years in IFLS).
All monetary values are annual figures in real 2000 US dollars.
Education expenditure data in IFLS are for households with children under 24 years old
Inflation rate data is from Asian Development Bank for Indonesia and BEA for United States.
Summary statistics for inflation rates and food consumption growth rates are changes between observations for households (1 year in PSID, 3 or 4 years in IFLS).

Table 3

EFFECT OF UNEMPLOYMENT ON FOOD CONSUMPTION: INDONESIA VS UNITED STATES

Dependent variable: Food cons. growth rate (change in log household food consumption)

	(1)		(2)		(3)	
	No controls		With controls		Only Those Unemployed Exactly Once	
	US	Indonesia	US	Indonesia	US	Indonesia
Unemployed dummy	**-0.087**	**-0.097**	**-0.106**	**-0.078**	**-0.095**	**-0.098**
	(0.006)***	(0.027)***	(0.010)***	(0.022)***	(0.017)***	(0.038)**
People in household			0.01	-0.005	0.012	-0.004
			(0.002)***	(0.002)**	(0.005)**	(0.007)
Age			-0.001	0.000	0.000	0.001
			(0.000)***	(0.000)	(0.001)	(0.000)
Married			0.033	0.057	0.032	0.02
			(0.007)***	(0.027)**	(0.018)*	(0.06)
Sex			-0.012	-0.007	0.006	-0.035
			(0.007)*	(0.014)	(0.017)	(0.03)
School			0.000	-0.005	0.000	0.005
			(0.000)	(0.008)	(0.001)	(0.025)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Province/state dummies	No	No	Yes	Yes	Yes	Yes
Observations	50769	11284	50763	11284	7894	1231

NOTE-Sample includes all households who remain in panel for two or more years where head is employed in previous observation. Observations with nominal food consumption growth rates in bottom 2% and top 2% of distribution are discarded to trim outliers. Dependent variable in all specifications is $\log(c_t)$-$\log(c_{t-1})$ where t-1 refers to the previous observation (1 year lag in PSID, 3 or 4 years in IFLS. Unemployed dummy is 1 if head of household is not working at time of interview; else 0.

Table 4

CALIBRATIONS OF CONSUMPTION DROP AND WELFARE GAINS OF SOCIAL INSURANCE

A. Consumption Drop ($\Delta c/c$)

		Coefficient of relative risk aversion (γ)				
		1	2	3	4	5
Disutility of effort in unemp. state (θ_u)	1	0.00	0.00	0.00	0.00	0.00
	1.25	0.20	**0.11**	0.07	0.05	0.04
	1.5	0.33	0.18	**0.13**	0.10	0.08
	1.75	0.43	0.24	0.17	**0.13**	0.11
	2	0.50	0.29	0.21	0.16	**0.13**

B. Marginal Welfare Gain ($\gamma\Delta c/c$)

		Coefficient of relative risk aversion (γ)				
		1	2	3	4	5
Disutility of effort in unemp. state (θ_u)	1	0.00	0.00	0.00	0.00	0.00
	1.25	0.20	**0.21**	0.22	0.22	0.22
	1.5	0.33	0.37	**0.38**	0.39	0.39
	1.75	0.43	0.49	0.51	**0.52**	0.53
	2	0.50	0.59	0.62	0.64	**0.65**

NOTE-Panel A shows the implied consumption drop without social insurance for a given combination of risk aversion and disutility of effort to earn income while unemployed for the stylized model in section 4. The table shows that many combinations of risk aversion and disutility of effort can generate consumption drops similar to those observed in the data (in bold on diagonal). Panel B shows the marginal welfare gains of social insurance for each combination of parameters. Welfare gains are rising on the diagonal even though the consumption drop is constant.

Table 5

EFFECT OF UNEMPLOYMENT ON CONSUMPTION OF STAPLES

Dependent variable: Staples cons. growth rate (change in log staples consumption)			
	(1) OLS	(2) Median Reg.	(3) OLS, No Farmers
Unemployed dummy	**-0.060**	**-0.100**	**-0.119**
	(0.039)	(0.035)***	(0.048)**
People in household	-0.009	-0.005	-0.013
	(0.004)**	-0.004	(0.006)**
Age	0.000	0.000	0.001
	(0.000)	(0.000)	(0.001)
Married	0.129	0.147	0.060
	(0.047)***	(0.043)***	(0.068)
Sex	0.042	0.048	0.037
	(0.024)*	(0.022)**	(0.033)
School	0.052	0.042	0.080
	(0.014)***	(0.013)***	(0.020)***
Year dummies	Yes	Yes	Yes
Province dummies	Yes	Yes	Yes
Observations	9,466	9,466	5,205

NOTE-Sample includes all IFLS households who remain in panel for two or more years where head is employed in previous observation. Observations with nominal staples consumption growth rates in bottom 2% and top 2% of distribution are discarded to trim outliers. Dependent variable in all specifications is $\log(c_t)$-$\log(c_{t-1})$ where t-1 refers to the previous obs. Unemployed dummy is 1 if head of household is not working at time of interview; 0 otherwise. Median regression is a quantile regression at the 50th percentile. No farmers specification excludes all households with one or more individual working on a farm.

Table 6

OTHER RESPONSES TO UNEMPLOYMENT: EVIDENCE OF RISK AVERSION

	Educational Expenditures			Other family members' labor supply		
	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive margin		Intensive margin	Extensive margin		Intensive margin
	No controls	With controls	Median Reg	No controls	With controls	Median Reg
Dependent Variable:	Education dummy		log Δ ed exp	Participation dummy		log Δ other fam inc
Unemployed dummy	**-0.13**	**-0.09**	**-0.12**	**0.17**	**0.15**	**0.11**
	(0.02)***	(0.02)***	(0.07)	(0.02)***	(0.02)***	(0.07)*
People Per Household		0.01	-0.03		0.06	0.02
		(0.00)***	(0.01)***		(0.00)***	(0.01)***
Age		0.00	-0.01		0.00	0.00
		(0.00)***	(0.00)***		(0.00)***	(0.00)
Married		0.13	-0.03		0.07	0.28
		(0.02)***	(0.09)		(0.03)***	(0.09)***
Sex		0.01	-0.05		0.04	0.21
		(0.01)	(0.04)		(0.01)***	(0.04)***
School		0.06	0.05		0.00	-0.04
		(0.01)***	(0.02)**		(0.01)	(0.02)*
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Province dummies	No	Yes	Yes	No	Yes	Yes
Observations	7,700	7,457	6,156	6,778	6,407	3,478

NOTE-Sample includes all IFLS households who remain in panel for two or more years where head is employed in previous observation. Dependent variable in (1) and (2) is an indicator for whether household reported positive education expenditures. Only households with positive education expenditures in previous year are included in (1) and (2). Dependent variable in (3) is log change in education expenditures; sample includes all households reporting positive education expenditures in both previous year and current year. In (3), outliers are trimmed at upper and lower 2% as in Table 3. Dependent variable in (4) and (5) is an indicator for whether any household member besides the head is earning income in current year. Only households where no other member besides head was working in prior year included in (4) and (5). Dependent variable in (6) is log change in other family members' income, with 2% trimming analogous to that in Table 3.
Sample in (6) includes households reporting positive non-head income in both previous year and current year. Unemployed dummy is 1 if head of household is not working at time of interview; 0 otherwise. Median regression is a quantile regression at the 50th percentile.

Figure 1a

Social Insurance vs. GDP Per Capita in 1996

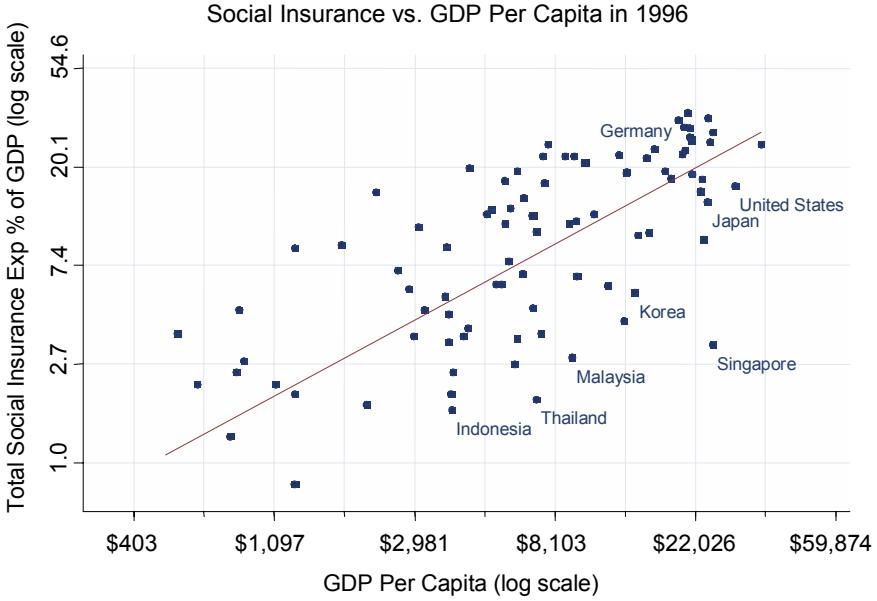

Figure 1b

Govt. Budget Share of Social Insurance vs. GDP Per Capita

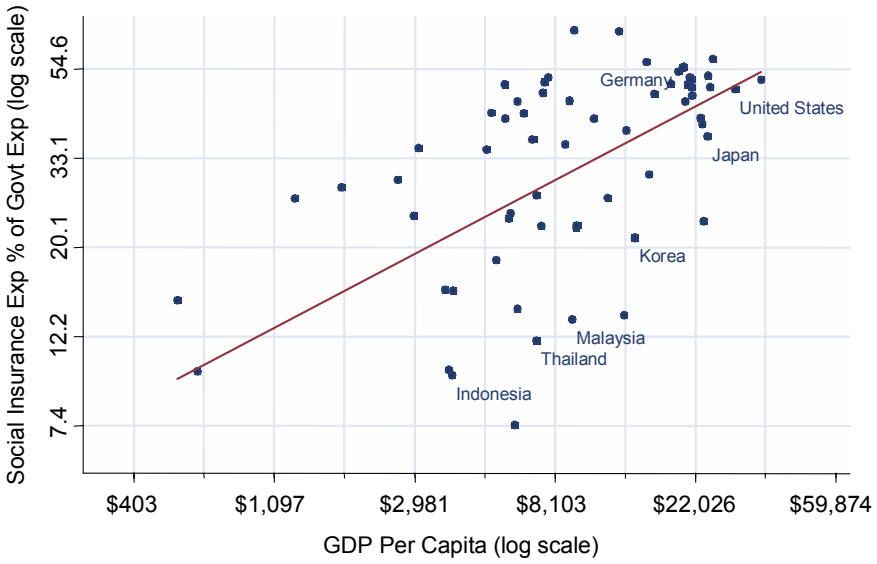

NOTE–Social Insurance statistics are from ILO (2000); GDP statistics are from the Penn World tables. GDP is measured in 1996 US dollars. Panel A shows relationship between social insurance share of GDP and GDP per capita. Panel B shows relationship between social insurance share of government budget and GDP per capita.

Figure 2

NOTE–Sample consists of all household heads who experienced exactly one unemployment spell between 1980 and 1993 in the PSID. Annual growth rates of food consumption are computed as change in log of real (CPI deflated) food consumption from year $t-1$ to year t. Year of job loss is normalized at 0 and all other years are defined as difference relative to that year.

Figure 3a

Effect of Unemployment on Food Consumption in Indonesia

Source: IFLS 1993-2000

Figure 3b

Effect of Unemployment on Food Consumption in the US

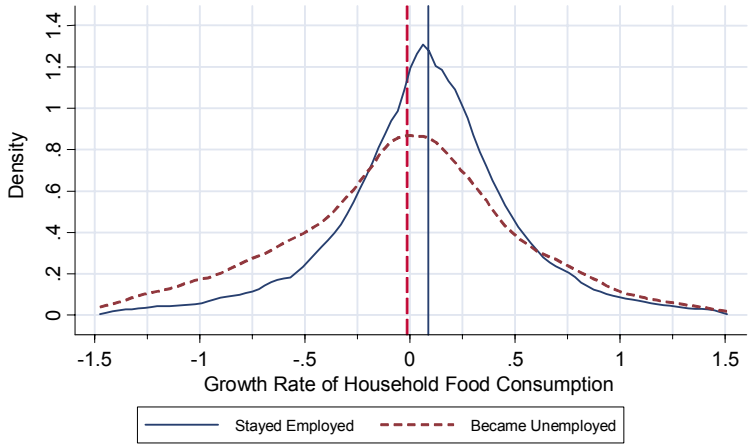

Source: PSID 1980-1993

NOTE-In each figure, vertical lines denote median for density of corresponding color. Sample consists of all household heads in IFLS or PSID who reported being employed at the time of previous interview. "Stayed Employed" group includes household heads who remain employed in interview t. "Became Unemployed" group includes household heads who are not working at time of inteview t. Growth rate of household food consumption is defined as nominal difference in log food consumption in inteview t and interview $t-1$. Gap between interviews is one year in PSID and 3 or 4 years in IFLS. Observations with growth rates in top 2% or bottom 2% of unconditional food growth distribution in each dataset are discarded to trim outliers. Kernel densities are estimated using an optimal bandwith procedure.

Figure 4

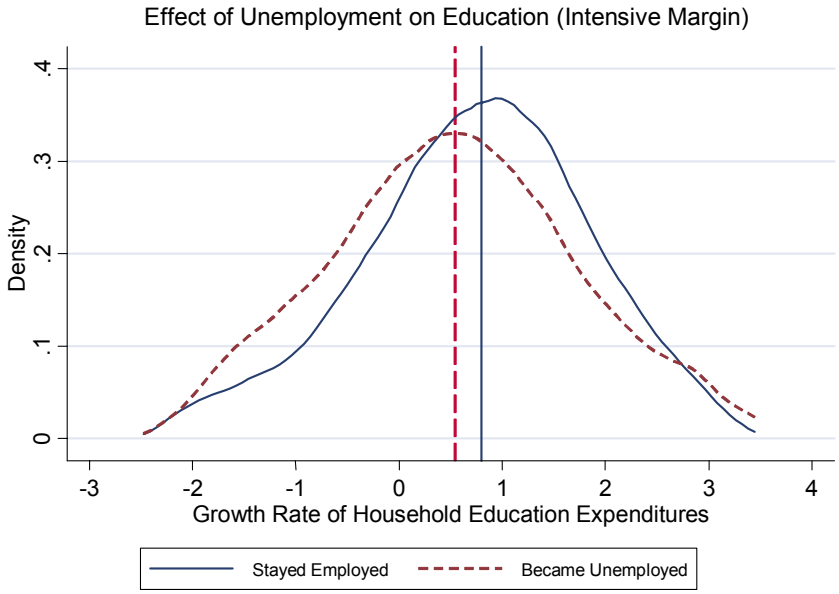

Source: IFLS 1993-2000

NOTE-Vertical lines denote median for density of corresponding color. Sample consists of all household heads in IFLS who reported being employed at the time of previous interview and who report positive educational expenditures in both previous interview and current interview. "Stayed Employed" group includes household heads who remain employed in interview t. "Became Unemployed" group includes household heads who are not working at time of inteview t. Growth rate of is defined as nominal difference in log educational expenditure in inteview t and interview $t-1$. Gap between interviews is 3 or 4 years in IFLS. Observations with growth rates in top 2% or bottom 2% of unconditional educational expenditure growth distribution are discarded to trim outliers. Kernel densities are estimated using an optimal bandwith procedure. See Table 6 for corresponding results on extensive margin.

Figure 5

Effect of Unemployment on Others' Labor Supply (Intensive Margin)

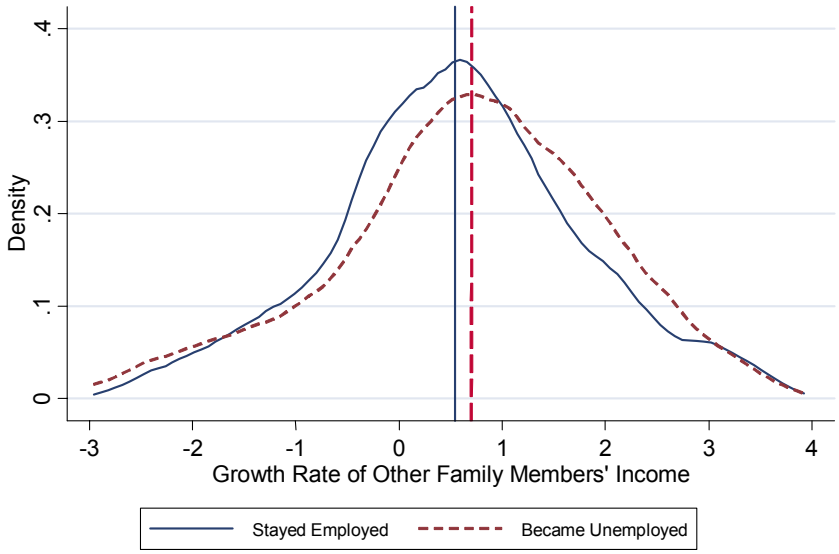

Source: IFLS 1993-2000

NOTE-Vertical lines denote median for density of corresponding color. Sample consists of all household heads in IFLS who reported being employed at the time of previous interview and who report positive income from other family members in both previous interview and current interview. "Stayed Employed" group includes household heads who remain employed in interview t. "Became Unemployed" group includes household heads who are not working at time of inteview t. Growth rate of is defined as nominal difference in log of other family members' income in inteview t and interview $t-1$. Gap between interviews is 3 or 4 years in IFLS. Observations with growth rates in top 2% or bottom 2% of unconditional other-income growth distribution are discarded to trim outliers. Kernel densities are estimated using an optimal bandwith procedure. See Table 6 for corresponding results on extensive margin.