RESEARCH SEMINAR IN INTERNATIONAL ECONOMICS

School of Public Policy
The University of Michigan
Ann Arbor, Michigan 48109-1220

Discussion Paper No. 482

The Distributional Impacts of Indonesia's Financial Crisis on Household Welfare: A "Rapid Response" Methodology

Jed Friedman
RAND Corporation

James Levinsohn
University of Michigan

September, 2001

Recent RSIE Discussion Papers are available on the World Wide Web at:
http://www.spp.umich.edu/rsie/workingpapers/wp.html

The Distributional Impacts of Indonesia's Financial Crisis on Household Welfare: A "Rapid Response" Methodology

Jed Friedman
RAND Corporation
Santa Monica, CA 90407

James Levinsohn
Ford School of Public Policy and Department of Economics
University of Michigan
Ann Arbor, MI 48109
and
NBER

Abstract: Analyzing the distributional impacts of economic crises is important and, unfortunately, an ever more pressing need. If policymakers are to intervene to help those most adversely impacted, then policymakers need to identify those who have been most harmed and the magnitude of that harm. Furthermore, policy responses to economic crises typically must be timely. In this paper, we develop a simple methodology to fill the order and we've applied our methodology to analyze the impact of the Indonesian economic crisis on household welfare there. Using only pre-crisis household information, we estimate the compensating variation for Indonesian households following the 1997 Asian currency crisis and then explore the results with flexible non-parametric methods. We find that virtually every household was severely impacted, although it was the urban poor that fared the worst. The ability of poor rural households to produce food mitigated the worst consequences of the high inflation. The distributional consequences are the same whether we allow households to substitute towards relatively cheaper goods or not. However the geographic location of the household mattered even within urban or rural areas and household income categories. Additionally, households with young children may have suffered disproportionately adverse effects.

The authors are grateful to the World Bank for partially funding this research. Alan Winters and Will Martin provided helpful comments. Special thanks is due to Angus Deaton who essentially suggested the methods used in this paper.

I. Introduction

The collapse of the Indonesia Rupiah during the 1997 Asian currency crisis precipitated a 12% decline in Indonesia's GDP the following year as well as rampant inflation. In an 18-month span, food prices nearly tripled and prices for other goods also rose substantially. The degree to which Indonesian households were vulnerable to these changes depended on a mix of factors including the types of goods the household consumed, which goods' prices rose the fastest, and the degree to which changes in income were able to buffer households from the brunt of the price shocks. In this study we focus on the first two factors, namely household consumption choices and goods price changes, in order to explore how the price changes impacted households across both the income distribution and the different regions of Indonesia. We develop and apply a simple methodology that requires only the sort of data that is often readily available immediately after an economic crisis. Although this paper focuses on the Indonesian experience, our methodology is intended to be applicable to a wide range of situations and countries. By so doing, we hope to present a meaningful yet straightforward methodology that can be adopted to analyze the distributional consequences of financial crises and inflation anywhere in the world.

A careful and definitive investigation of the impacts of the Indonesian currency crisis, and potential differential impacts across levels-of-living, ideally requires detailed income and expenditure information for a large number of households, both before and after the crisis. This ideal data does not exist. Some sources of data approach this ideal, in particular the Indonesian Family Life Survey (IFLS). Consecutive waves of the IFLS (namely waves 2 and 2+) have gathered pre- and post- crisis information for a panel of 2500 households. Various studies using this data have investigated impacts of the crisis on consumption, employment, and education (see Frankenberg et al. 1999, Smith et al. forthcoming, Thomas et al. 2001). While this data presents an impressive depth of information for households, unfortunately it is limited to a relatively small sample of households in a minority of provinces and thus cannot speak to the breadth of the crisis across the sweep of Indonesian geography or income distribution. Approximately a year and a half after the release of the initial IFLS based reports (and three years after the onset of the crisis), studies employing nationally representative post-crisis (non-panel) household information have begun to appear, an example of which is Suryahadi et al. (2000).

Our approach will be distinct from the above studies in that we will only utilize household data collected before the onset of the crisis. We then match this consumption data with information on commodity price changes brought on by the crisis in order to calculate simple measures of compensating variation - the amount of money sufficient to compensate households

following price changes and enable a return to pre-crisis levels of utility. We calculate this compensating variation with a variety of methods and compare and contrast the strengths and weaknesses of each approach. Our analysis employs data with sufficient degrees of freedom to allow an exploration of differences in compensating variation across the spectrum of household income and location.

Because the analysis presented here only requires pre-crisis household consumption information and price change data, our approach is applicable to many other settings. This is because numerous countries now conduct periodic household consumption surveys and even more collect more-or-less current price data used for computing price indices. An important benefit of these methods is the relative immediacy of the findings. Post-crisis household surveys can yield valuable and even definitive information however this data is only available after the substantial lag needed for data collection and processing. In the face of rapid economic change and social disruption, the information needs of policy makers are immediate. We hope to suggest a "rapid response" method that can be implemented at the onset of a financial crisis well before this post-crisis household data can be collected and disseminated.

The next section discusses the details of the data sources. This is then followed by an explanation of the methods of analysis in Section III. Section IV presents the basic results while Section V extends the results and investigates the robustness of the findings. The final section (Section VI) offers some concluding thoughts.

II. Data

In this study we match household level data on consumption with provincial level information on commodity price changes. The consumption data derive from the 1996 National Socio-Economic Survey, known by the Indonesian acronym SUSENAS. Indonesia regularly conducts this extensive household consumption survey that typically covers fifty to sixty thousand households. These surveys are conducted every three years and the 1996 wave, which surveyed 61,965 households, was the most recent survey before the onset of the crisis. While these surveys are large, they are not panels. That is, there is no systematic effort to track the same households over time. However they do cover the entire geographic range of the country and contain very detailed consumption data on a total of 306 food and non-food goods. SUSENAS also records whether food goods were purchased in the market or produced by the household. If food is self-produced, SUSENAS imputes a value of this consumption at prevailing local prices. SUSENAS also imputes a value for owned housing.

It is important to note that SUSENAS does not contain information on prices. Rather the data enable the computation of unit values, which are defined as expenditure for a particular good divided by quantity consumed. These unit values may differ across households that in fact face identical prices due to differences in the quality of consumption. For example, while all households in a village may face the same prices for high quality and low quality rice, the unit values recorded for a household that bought mostly high quality rice will be higher than the unit values recorded for the household that bought mostly low quality rice. These higher unit values simply reflect the higher mean quality of total rice purchases. This type of data can be (and in fact have been) used to estimate demand elasticities exploiting the spatial variation in the data using methods developed by Deaton (1988, 1990, 1997). We note this, as the unit value data will be utilized in subsequent sections.

We also have recent price data that has been supplied by the Indonesian Central Statistical Office (the Biro Pusat Statistik, or BPS). The price data contain monthly price observations for 44 cities throughout the country over the period January 1997 to October 1998. This time period, which begins before the advent of the crisis, spans the steep devaluation of the Rupiah and subsequent (and temporary) stabilization at the new higher rate. We employ a single price change measure- the percent change in price from January 1997 to October 1998. By adopting such a long time period from before the onset of rapid inflation until after the inflation has largely abated, we hope to capture a robust measure of the price changes associated with the crisis.

The price data supply information for both aggregate goods, such as food or housing, as well as for individual goods such as cassava or petrol. There are approximately 700 goods with observed prices in the data. However, the types of goods observed vary by city, perhaps reflecting taste/consumption heterogeneity throughout the country. On average, a particular city has price information for about 350 goods. Jakarta has as many as 440 goods listed while some small cities only have price information for 300 goods.

Each of the 27 Indonesian provinces is represented by at least one city in the price data. In order to match households from the SUSENAS data to as local a price change as possible, we calculate province specific price changes from the city-level data. For those provinces that have only one provincial city in the price data, we take those price changes as representative of the entire province. For those provinces with more than one city in the price data, we calculate an average provincial price change using city-specific 1996 population weights.

The accuracy of this extrapolation of city price data to an entire province will surely vary with the size and characteristics of the province considered. For example, Jakarta, the national

capital, is also its own province and the observed price changes will fairly accurately represent the price changes faced by residents throughout the province. On the other hand, the price changes for Irian Jaya, a vast mountainous province, are based on price changes observed in the provincial capital Jayapura. Price changes in the provincial capital may not be a completely accurate proxy for price changes in remote rural areas. Indeed a recent paper suggests that post-crisis inflation in rural areas may have been 5% higher than in urban areas (Frankenberg et al. 1999). We frequently report separate results for rural and urban households and the fact that price data were collected in cities should be kept in mind as those results are reviewed.[1]

We match the price change data with the consumption data at the most disaggregate level possible in order to calculate the measures of compensating variation, which will be detailed in the next section. There are 219 products and product aggregates that appear in both the SUSENAS and our price data. We attempt to match goods across the two data sets at the lowest level of aggregation possible. For the case of food (both raw and prepared) we were able to match 155 individual goods between the two data sets. In the case of non-food items we matched 64 different goods, both individual goods such as firewood and kerosene, as well as aggregate goods such as toiletries and men's clothing.

For certain groups of goods the price data are more disaggregated than the consumption data reported in SUSENAS. In order to link the new price data with the existing consumption data, we use the prices for those commodities that appear in both the price data set and in SUSENAS. In some cases we also aggregate commodities in the price data to match a product category in the SUSENAS data. The match between the price data and the consumption data is good, but not perfect. We find that we have detailed price data for most, but not all, of the goods that comprise a household's total expenditure. On average, expenditures on matched goods account for 79 percent of a household's total expenditure- a little greater for poor households and a little less for wealthy ones.

For use in subsequent analysis, we calculate the budget shares of each of the 219 items based on the reported monthly expenditures for each item. For durable goods and other non-food

[1] We have looked at how unit values in urban and rural areas have changed across different Susenas survey years (1984, 1987, 1990, 1993, and 1996) to investigate potential differences in the time trend of urban and rural prices. We do this by fixing a specific food basket and then pricing this basket separately for both urban and rural areas in each survey year. Instead of actual prices, however, we use the mean national urban or rural unit values as our price measure. In essence, we've generated separate price indices for urban and rural areas. The time trends of these indices are virtually identical. For example the urban index increased by 187% over the 1984 – 1996 period while the rural index increased by 182%. Each three year change in urban areas is even closer in magnitude to its rural counterpart. This pre-crisis co-movement of urban and rural unit values suggests that urban and rural prices may behave quite similarly following the crisis and thus our extension of urban price changes to rural areas may not introduce significant bias.

items, we use the monthly average of annual expenditure, and not the expenditures in the month preceding the survey, in order to more accurately measure monthly expenditures for durables that are infrequently purchased. Table 1 gives an overview of the consumption data by reporting budget shares for selected composite goods. These goods are not chosen from among the 219 items but rather are composite aggregates constructed only for the expositional purposes of Table 1. Even the rice good in the first row of Table 1 is an aggregate of three different varieties. To highlight the heterogeneity in consumption patterns, we report mean budget shares for the entire sample as well as for the top and bottom decile of household expenditures. Clearly, rice is the single most important commodity, as measured by the budget share, for the majority of Indonesians. Households in the bottom expenditure decile devote more than a quarter of all outlays to rice, while for the mean household a still substantial 16% of total expenditures goes towards rice. The next most important aggregate consumption category encompasses housing and utilities, especially so for the top expenditure decile where 22% of spending goes towards those ends.

Alongside the budget shares, Table 1 also reports the average price increase for each product aggregate. This is accomplished by calculating the household specific price increase of the composite goods using household expenditure shares to weight the price increases of each constituent individual good. We then average these household-specific price increases over all households. By any measure, the inflationary impacts of the crisis were large. The all-important rice price increased by an average of almost 200%, and the prices for many foodstuffs increased by more than 100%. Non-food prices did not rise nearly as rapidly, with the housing and utilities price increasing the least - only 24% on average. Listed next to the mean price increases are the standard deviations of the price increases for the aggregate goods. Due to the constructed nature of the reported price changes, variations in price change will arise due both to geographic variation in price changes as well as household variation in consumption of individual goods. For rice, a relatively homogenous good, all of the variation in the rice price increase is geographical and a standard deviation of 30% shows how varied the price increases actually were.[2] If the price changes for rice were distributed in a roughly normal fashion then fully one-third of households experienced an increase in the rice price outside the interval (165%, 225%). Other reported price changes combine variation in household consumption choice with regional variation in price changes and as such the standard deviations of these price changes tend to be larger.

[2] Although there are three rice varieties from the Susenas consumption data, the BPS price data supplies only one price change for all rice varieties.

Given the wide dispersion of price changes, both within and across product aggregates, what a household consumes and where a household lives will go a long way towards determining the particular impacts of the crisis for each household. The next section discusses how we measure these household specific consequences.

III. Methodology

To consider the impacts of the price increases on household welfare, we will look at changes in consumer surplus brought about by the change in prices. We start with a minimum expenditure function C(u,p) which, given existing prices p, relates the minimum cost needed to attain utility level u [see Deaton and Muellbauer (1980), chapter 2, for a discussion of the general properties of cost functions]. A first-order Taylor expansion of the minimum expenditure function with respect to price will yield an approximation of the income required to compensate the household after a price change and to restore that household to the pre-change utility level. Hence this expression will approximate the compensating variation. Noting that the partial derivative of the minimum expenditure function with respect to price yields quantities consumed, we derive this simple expression:

$$\Delta C \approx q\Delta p \tag{1}$$

where q is a 1 x n vector of consumption goods quantities, Δp a 1 x n vector of price changes, and n the number of consumption goods in the total demand system. We note that this first approximation of compensating variation requires information only on pre-crisis consumption quantities and on price changes; neither pre-crisis price levels nor, more importantly, post-crisis consumption choices are needed.

It is straightforward to reformulate (1) in terms of budget shares, w, and proportionate price changes with the following expression:

$$\Delta \ln C^h \approx \sum_{i=1}^{n} w_i^h \Delta \ln p_i^h \tag{2}$$

where i subscripts the individual goods in the commodity system and h refers to the household. The budget share w is simply the household cost of good i divided by pre-crisis total household expenditures. Made clear in (2) is the simple fact that any differential distributional impact of the

price changes must derive both from the presence of large relative price changes and large differences in the budget shares across households. Table 1 shows this exact combination of factors existed in Indonesia following the crisis.

In general, the costs of attaining pre-crisis utility levels will increase less rapidly than (2) may suggest since households have the ability to substitute away from goods whose prices have disproportionately risen. Hence this compensating variation measure provides a maximum bound on the impact of the crisis, since this index does not take into account the substitution toward relatively less costly products that will take place. Given the large relative price changes following the crisis this substitution surely occurred to some extent and thus (2) may not be an entirely accurate approximation. Returning to the minimum expenditure function, a second-order Taylor expansion of the minimum expenditure function does permit substitution behavior:

$$\Delta C \approx q\Delta p + \frac{1}{2}\Delta p^T s \Delta p \tag{3}$$

In (3), q and Δp are quantity and price change vectors as before and s is the n x n matrix of compensated derivatives of demand. Similar to (2), we can reformulate the above expression in terms of budget shares and proportional price changes as:

$$\Delta \ln C^h \approx \sum_{i=1}^{n} w_i^h \Delta \ln p_i^h + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} c_{ij} \Delta \ln p_i^h \Delta \ln p_j^h \tag{4}$$

where the expression c_{ij} contains the Slutsky derivatives s_{ij} and is defined by the expression:

$$c_{ij} = \frac{p_i s_{ij} p_j}{C^h}$$

With some simple algebraic manipulation we can show the c_{ij} term above to be equivalent to $w_i\varepsilon_{ij}$, where ε_{ij} is defined as the compensated price elasticity of good i with respect to price change j. Thus we can restate (4) as:

$$\Delta \ln C^h \approx \sum_{i=1}^{n} w_i^h \Delta \ln p_i^h + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} w_i^h \varepsilon_{ij} \Delta \ln p_i^h \Delta \ln p_j^h \tag{5}$$

It is the two formulations of compensating variation given in (2) and (5) that we will use to explore the possible differential impacts of the Indonesian currency crisis. The only additional pieces of information required in (5) and not found in (2) are the ε_{ij} terms. Thus an approximation to the compensating variation that also wishes to account for potential household substitution behavior requires estimates of a complete set of price elasticities in addition to the pre-crisis consumption quantities and post-crisis price changes. Exactly how these elasticities are estimated depends on the types of data used in the analysis. Our task appears difficult since we have no information on household consumption changes over time nor do we have information on price levels. Instead of prices we have data on unit values. In a series of papers Deaton (1988, 1990, 1997) presents an approach to elasticity estimation using only a single cross-section of household information. Crucial to this approach is the twin recognition that prices for equivalent goods can greatly vary across space in a lesser developed country setting and that household survey information is often gathered in clusters in order to reduce survey costs. Given these insights as well as certain assumptions on how households choose the quality of goods purchased, the clustered nature of this data can be exploited to purge the unit value data of quality components. The cross-spatial variation in these purged unit values can then be utilized to identify own-price or cross-price elasticities. This is the method adapted here to estimate the ε_{ij} terms.

We now summarize this method in a bit more detail before moving on. Deaton suggests adopting the following econometric specifications for the log quantity (lnq) and log unit value (lnv) of a particular good:

$$\ln q_{hc} = \alpha^0 + \beta^0 \ln x_{hc} + \gamma^0 z_{hc} + \varepsilon_p \ln \pi_c + f_c + u^0_{hc}$$
$$\ln v_{hc} = \alpha^1 + \beta^1 \ln x_{hc} + \gamma^1 z_{hc} + \psi \ln \pi_c + u^1_{hc}$$

where h and c index household and cluster respectively, x represents total household expenditures, z household demographic characteristics, and π the (unobserved) price of the good. The quantity equation contains a cluster fixed effect, f_c, and the coefficient of interest is ε_p, the price elasticity. The simplified process described here only concerns the estimation of own-price elasticities- cross-price terms can be added through a relatively straightforward extension. The final estimate of ε_p derives from two main steps. In the first step, the within cluster variation of household income and characteristics is utilized to estimate β and γ (since prices are constant within clusters, these parameters can be consistently estimated). The estimated coefficients are then employed to generate the two variables:

$$\hat{y}^{o}_{hc} = \ln q_{hc} - \hat{\beta}^{0} \ln x_{hc} - \hat{\gamma}^{0} z_{hc}$$
$$\hat{y}^{1}_{hc} = \ln v_{hc} - \hat{\beta}^{1} \ln x_{hc} - \hat{\gamma}^{1} z_{hc}$$

The next step is to calculate the cluster-level averages of y^0 and y^1. Then a "regression" of the cluster averaged y^0 on the cluster averaged y^1 will yield an estimate of the ratio of ε_p to ψ:

$$\frac{\varepsilon_p}{\psi} = \frac{\text{cov}(\hat{y}^{0}_{c} \hat{y}^{1}_{c})}{\text{var}(\hat{y}^{1}_{c})}$$

Finally, combining the above expression with an estimate of ψ (which itself is identified from previously estimated coefficients in an expression (not shown) determined by the model of household quality choice) enables the researcher to calculate the price elasticity estimate.[3]

If one wishes to estimate a demand system of our dimensions, some product aggregation is necessary. There are simply not enough degrees of freedom in the SUSENAS data to estimate a demand system for 219 products complete with the all-important cross-price elasticities. The types of goods for which we can estimate price elasticities are also limited by the fact that SUSENAS reports unit values solely for food goods. Hence we reduced the dimensions of the problem through aggregation and decided to estimate elasticities for 22 composite goods - 21 aggregate food goods and a residual non-food consumption category. A subsequent table in the next section lists each of these aggregate goods.

Beyond estimating the ε_{ij}'s, another issue concerns the services provided by owner-occupied housing and self-produced agriculture. Many households, especially in rural areas, own their own home. Although the price of housing has increased, these households are, in an absolute sense, perhaps not better off (they are still living in the same house). However these households are better off relative to those who do not own their own home. We choose to account for these services provided by owner occupied housing by treating the imputed rental value for these homes as a negative expenditure. Many households, mostly rural, also produce some of their own food. Households that consume self-produced foodstuffs are also potential net exporters of agricultural products. As the price of food rose, the value of their production also increased.

[3] This brief discussion ignores the identification of ψ as well as the important role of measurement error corrections in the original series of papers. We refer readers to those works for a much more extensive presentation and discussion.

Clearly, if the household were a net exporter of food, the household would benefit from the price increase. To the extent that a household produced some of its own food, such production would mute the impact of price increases relative to a household that purchased food in the market. Thus we account for self-produced agricultural products by treating the imputed value of self-produced food as a negative expenditure.[4,5]

Once the budget share and price change data have been matched, and the price elasticities estimated, we then calculate our two measures of compensating variation for each household. So that we can explore in a flexible manner how (2) and (5) vary across the levels-of-living, our principal approach will be non-parametric. Specifically, we use locally weighted least squares to estimate the compensating variation at each point in the income distribution (see Fan, 1992, for an introduction of this method). Local observations were weighted with a biweight kernel and, after experimentation, we choose to adopt a bandwidth of 0.4 units of the independent variable (which will be log per capita monthly household expenditures).[6]

To assess a household's level-of-living, we divide the data two ways. The first is merely by per capita household expenditure, and the second is a binary poor/non-poor measure dependent on whether the household's per capita expenditure exceeds or falls below a predetermined poverty line. Each measure is discussed in turn.

Perhaps the most standard approach to measuring the level-of-living in a developing country setting is to use some estimate of household expenditures. In this view, the level of household consumption constitutes the lion's share of total household utility and total consumption is most easily proxied by the household's actual expenditures. Expenditure levels are generally viewed as a better measure of welfare than income since the ability to smooth consumption in the presence of income shocks suggests that expenditures rather than income more closely tracks actual welfare.[7]

---

[4] Note that this approach will understate the effects of the price increases to the extent that we do not observe nor adjust for price increases of intermediate inputs used in agricultural production.

[5] There is a long standing debate over whether shadow prices in rural households engaged in agricultural production equate market prices for agricultural inputs such as labor or land. To the extent that these shadow prices may diverge from market prices, the "valuation" for self-produced food, based on market prices, will not be entirely accurate. Benjamin (1992) presents evidence from rural Java that hosuehold shadow prices for agricultural inputs such as labor are not significantly different from market prices.

[6] The Stata code for all of the analysis will be made available at www.econ.lsa.umich.edu/~jamesL.

[7] Chaudhuri and Ravallion (1994) investigate the competing merits of using these two welfare indicators and find little difference when the goal is to distinguish poor from non-poor households. This paper remains within the standard literature and uses household expenditures as a main measure of household welfare.

In addition to this continuous measure of level-of-living, an alternative binary poverty measure is also adopted. A household is deemed poor if its per capita expenditure falls below a predetermined poverty line. The poverty lines used here are calculated from the 1996 SUSENAS using a "cost of basic needs" approach to poverty determination as set forth in Ravallion (1994), Bidani and Ravallion (1993), and Ravallion and Bidani (1994). The details of the particular method used here are presented in Friedman (2001), but the general approach is summarized as follows: a nutritionally adequate food bundle (with nutritional guidelines stipulated by WHO, 1985) that reflects the actual consumption choices of Indonesian households is determined and then priced. The total cost of this bundle is scaled upwards by an econometrically estimated factor that represents the cost of essential non-food goods. Thus this final value, which we take as the poverty line, proxies the total cost of essential food and non-food consumption needs. Due to important differences in relative prices between urban and rural areas, poverty lines are computed separately for each region. For the 1996 SUSENAS, this method translates into a poverty line of 36,956 Rupiahs per person per month in urban areas and 32,521 Rupiahs in rural areas. These values yield poverty headcounts of 9.3% in urban areas and 24.9% in rural areas.

IV. Results

The impacts of the crisis were not uniform. Instead household consumption choices, sources of income, and location mattered greatly in determining the specific impact. The diversity of impacts was due both to wide geographical variation in price changes as well as wide variation in household structure and consumption. An earlier paper (Levinsohn et al., 1999) explores this heterogeneity in much detail. Our concern here is solely focused on the relative differences in the compensating variation (cv) measures across the income distribution. This relative difference is exhibited quite clearly in Table 2, which reports summary mean values of equation (2) by decile of household expenditure as well as poor/non-poor status. For all households we see that the compensating variation has an inverted u-shape, with the lowest decile having an average cv of 73% of initial household expenditures, rising to a 85% of household expenditures for those in the sixth, seventh, and eighth deciles, while falling back to 77% for households in the top decile. From this perspective, it was the Indonesian households in the middle of the distribution that were most adversely impacted by the price changes. Indeed poor households would need to earn less income (as a proportion of initial expenditures) than non-poor households- 77% versus 82%- in order to return to original consumption levels.

However we see in the next two columns that this story obscures important differences between households in rural and urban areas. When separating the sample in this manner, the distributions of the impacts tell opposing stories. For urban areas, households in the lower deciles need the greatest relative amount of new income to return to pre-crisis consumption levels and, indeed, this amount declines monotonically as household expenditures increase. For rural areas, lower income households need the least relative compensation and then this proportion increases monotonically with expenditures. Table 2 suggests that it is the urban poor who are the most adversely affected by the crisis needing, on average, 109% of their pre-crisis income in order to reach pre-crisis utility levels. The rural poor, on the other hand, require the least amount, only 70% of their pre-crisis income. In general, urban households, composed mainly of households that do not grow their own food, fare the worst under the price changes.

This is the same story captured in Figure 1, which depicts the entire distribution of the compensating variation measure as estimated by locally weighted least squares. The figure also includes the urban and rural poverty lines for reference, as well as bootstrapped 95% confidence interval for each regression line.[8] The urban regression line declines almost completely monotonically from its peak at the bottom of the income distribution to its trough at the top of the distribution. In contrast, the rural regression line rises from its low at the bottom of the distribution and then flattens out for households beyond the top third of the distribution. After this point in the expenditure distribution there are virtually no differences in the cv measure, and no statistically significant difference, between urban and rural households. However the large differences between poor urban and poor rural households are indeed significant at conventional levels and as Figure 1 shows quite clearly, the urban poor were the most adversely affected while the rural poor were perhaps the least impacted.

The results in Table 2 and Figure 1 derive from equation (2) and likely overstate the true compensating variation since (2) does not allow for the substitution behavior that surely occurred to some degree. As already discussed, the addition of the second order terms in (5) may give a better approximation to the true cv since it does include substitution terms. These elasticities were identified by the spatial variation of consumption choices and unit values in the 1996 SUSENAS following the methods discussed earlier. Before moving on to estimates of equation (5), Table 3 presents these estimated price elasticities for the 22 composite good demand system (21 food goods such as rice or meat and the residual non-food category). The own price elasticities for each composite good are located on the diagonals in the price matrix and they are negative for

[8] The bootstrapped standard errors were estimated with 50 draws (with replacement) from the total sample and took into account the clustered nature of the underlying survey data.

almost every good. The estimated own price elasticity of rice is estimated to be -.48, exactly equal to that found by Case (1991) using earlier SUSENAS data and different methods of estimation. The three goods (preserved meat, prepared beverages, and alcohol) that are estimated to have positive own-price elasticities are goods that have substantially fewer positive consumption values than the other goods, in other words they are not widely consumed and as such are not likely to be precisely estimated. The cross price elasticities are generally smaller in magnitude than the own price elasticities and, of course, vary in sign depending on whether the data suggests a particular pair of goods to be substitutes or complements.

With this matrix of own and cross price elasticities we re-estimate the compensating variation using (5) and then contrast the results with those we found with (2). The comparisons, again estimated with locally weighted least squares and shown in Figure 2, are separated by urban and rural household location. As is readily apparent, the qualitative conclusions drawn with equation (2) also hold with results that now allow for substitution behavior. Across urban areas, the compensating variation declines as household expenditures increase, again suggesting that poor urban households are affected the most severely by the price changes. Similarly, poor rural households appear to fare the best, with little difference between wealthier urban and rural households.

However, the differences in the levels estimated for (2) and (5) are quite pronounced. The cv measures that allow for substitution are substantially less than those that do not, and this holds true at all expenditure levels and for both urban and rural households. Indeed, as a rule of thumb, the estimates of (5) are roughly half as much as the estimates of (2), with the difference being greatest for lower income urban households. Thus (5) suggests that the overall impacts of the crisis weren't nearly as severe as found with (2).

Without further information it is difficult to know which of the results from (2) and (5) are closer to the truth. We know that (2) surely overstates the impacts of the price changes since it restricts households to consume goods in the same proportions as they have done before the large relative price changes of the crisis. However we have reason to believe that (5) as currently estimated may dramatically understate the true compensating variation. If so then the true cv lies somewhere between the two regression lines for (2) and (5). We believe the results for equation (5) may overstate the true degree of substitution because the reduction in food consumption implied by the ε matrix and the price changes results in very low caloric intakes, much lower than would actually be exhibited (and indeed has been suggested by measured changes in the body-mass index in Frankenberg et al., 1999). Essentially the problem lies with the estimated point elasticities themselves. We believe these estimates may not be accurate for two important

reasons. One, (5) specifies the compensated price elasticity of demand while the estimation methods employed yield uncompensated price elasticities, which also involve the income effects from a price change. To the extent that the uncompensated price elasticities are larger than the compensated price elasticities, our measured effects in Figure 2 will be too large. Another and perhaps even greater problem is due the fact that the estimated elasticities are essentially local approximations based on consumer behavior at the observed prices. Hence SUSENAS might give fairly good estimates of how households respond to a price change on the order of five or ten percent. When the price changes under consideration are on the order of 100 to 300 percent, the answer is essentially dictated by choice of functional forms. This is troubling for most any parametric approach to the estimation of demand elasticities. In essence, we are forced to make out-of-sample predictions for every household and the further the real price changes are from the range of prices (or unit values) in SUSENAS, the more important our particular choice of functional form.[9]

We still present results with the cross-price elasticities since, in principal, they are an important refinement over (2). Noting the difficulties of accurately accounting for substitution behavior given only one cross-section of households and given price changes of the magnitude found in Indonesia in 1998, we don't wish to claim that the true post-crisis cvs are those estimated from equation (5). We do find it reassuring that the distributional consequences implied by (5) are the same as those implied by (2), however, and present results from both specifications. The combined results from both (2) and (5) may be of greater use to policy makers than either (2) or (5) alone.

V. Extensions

Having presented the basic results for equations (2) and (5), we now turn to four extensions that explore the robustness of the findings. These extensions: i) Explore how the results may differ if we ignore the services provided by owned housing and self-produced food; ii) Explore potential differences if, instead of the 219 highly disaggregate goods, we use a fewer number of more highly aggregated goods; iii) Explore the degree of spatial variation in the cv measures, and iv) Explore how the cv measures may be influenced by household size and demographic composition.

[9] We did re-estimate the cross-price elasticity matrix with a variant of Stone's demand system (see Deaton and Muellbauer, 1980, chapter 3) that does attempt to *approximate* compensated price elasticities. The results were largely similar to those presented in Figure 2, suggesting that most of the overstatement is due to the unavoidable out-of-sample predictions.

The first extension investigates differences in our findings if we do not account for the services provided by owned housing and self-produced food. Figure 3 presents this scenario, separate for urban and rural households, by presenting the non-parametric regression lines for the compensating variation given in (2) with and without valuing self-produced food and owned housing as negative expenditures. It is quite apparent that ignoring household self-production dramatically changes the results, especially for rural households. For households in urban areas, where only a minority of households produce some of their own food, the qualitative results are the same whether or not we value self-production- poor urban households are affected substantially more than wealthy ones. However without self-production and owned-housing, the regression line is shifted upwards in an almost parallel fashion, so that the levels of compensating variation are now approximately 15% greater than before.

For rural areas, however, ignoring self-production results in attributing the greatest adverse consequences to the rural poor as opposed to the rural wealthy, a complete reversal of the findings in Figure 1. The levels of cv for the rural poor also increase dramatically, almost doubling to approximately 130% of initial expenditures from the 70% reported in Table 2. The levels also rise for the rural wealthy but by a much smaller proportion. Clearly the ability of rural households, especially lower income rural households, to produce their own food served to buffer those households from the worst effects of the crisis. Urban households to a large degree could not share in this benefit.

We are also interested in exploring how the degree of aggregation affected the results. Remember that we attempted to match consumption and price changes at as low a level of aggregation as possible to more fully allow for heterogeneity in both consumption choices and price changes. The motivation, however, for looking at a more aggregate index stems from the fact that the disaggregated index only accounts for 79 percent of household expenditures on average. It is possible that we excluded important unobserved goods and this exclusion can either exacerbate or mitigate the measured welfare effects, depending on the relative price changes of those excluded goods. Concerned about this potential bias, we compute another compensating variation measure based on 19 aggregate commodities instead of the original 219. These aggregates include 15 food categories, such as cereals or meat, and four non-food categories such as housing or clothing. A benefit of this aggregate measure is that it covers 97 percent of the individual household's expenditures (this coverage is virtually the same for rural or urban households and across the income distribution).

The results shown in Figure 4 suggest that little is changed if we base the compensating variation measures on the more aggregate consumption goods. Indeed the regression lines

representing the aggregate and disaggregate measures are virtually identical for both urban and rural households. The analysis based on aggregated data is essentially unaffected by aggregation bias, at least in this case where our disaggregate measures include many important consumption goods. We find this reassuring on two fronts. First, Figure 4 implies that our main results are not biased by any "missing" consumption. Second, not every household survey from the rest of the world records consumption at such a disaggregate level as found in SUSENAS. However Figure 4 suggests that similar analysis conducted with these other surveys may suffer little detriment from aggregation bias as long as the basic consumption categories are covered in the data.

All of the preceding analysis has ignored cross-spatial variation in the compensating variation measures except by distinguishing urban from rural households. However Indonesia's population is spread out over 27 provinces on thousands of islands. Many of the studies previously cited concerning post-crisis household changes have shown that different areas of the country were impacted differently by the crisis due to geographic variation in both price changes and sources of income. Our findings are no different. When we calculate the mean province level values of our main cv measure the geographical diversity is readily apparent. For example, households in urban East Nusa Tenggara, a collection of islands east of Bali and Lombok, needed an additional 53% of pre-crisis expenditures in order to maintain consumption, while households in urban Southeast Sulawesi needed 124%. Although in every province rural households faced a smaller cv than their urban counterparts, the regional variation among rural households is equally dramatic. In Bengkulu, a Sumatran province, the cv for rural households averaged 105% while the figure for Irian Jaya was only 30%.

Besides these summary measures, we've also estimated the non-parametric regression lines separately for each province, ordered the estimated cvs at each point in the expenditure distribution, and then plotted the 10^{th} and 90^{th} percentile of the provincial specific cvs, along with the median. The resulting figure, Figure 5, presents some measure of the geographic variance of the impacts while controlling for per capita household expenditures. It is quite apparent that the effects of the crisis depended not only on the location of the household in the national expenditure distribution but also on the location of the household in space. For urban households, the 90^{th} percentile is roughly twice that of the 10^{th} percentile although this ratio is even greater for poorer households. Among rural households the spread between the 90^{th} and 10^{th} percentiles is even greater than that for urban households. Clearly, even within rural and urban areas, the household location is an important determinant in the overall impact of the crisis.

Until this point our principal measure of household welfare has simply been the household's per person expenditure level. Although a common measure, it imposes certain

restrictions on how welfare may or may not vary across observable demographic information such as household size or age and gender composition. Specifically, this measure does not recognize the possibility of scale economies at the household level nor does it recognize that consumption needs of individual household members may vary across gender or the lifecycle. Larger households, especially those with a greater number of working age adults, may be better off than smaller households at equivalent income levels since purchases of household public goods are shared among a greater number of household members. A consequence of this may be proportionally greater household expenditures for food (an important household private good) as public goods such as housing are more easily afforded. In addition to household size, the demographic composition of the household is likely to affect household consumption choices to the extent that consumption needs vary across the lifecycle or across gender. For example, households with children will almost surely spend more on education than otherwise equivalent households without children. Of course any differences in household consumption due to demographic influences will affect our cv measures.

We explore these issues in our final extension with some simple OLS regressions of the main cv measure on household size and demographic composition, as well as some relevant covariates including per capita household expenditures. These regressions, estimated separately for urban and rural households, are presented in Table 4. An earlier finding of this paper is also apparent in Table 4 in the estimated coefficients for household expenditures; the positive coefficient for rural households indicates the crisis impact increases with income levels in rural areas while the opposite story is indicated by the urban household coefficient. Turning to the question of household size, larger households are associated with higher cvs, especially households in rural areas. The potential reasons for this result are numerous but, surprisingly, higher foodshares resulting from the larger household sizes is not one of these explanations. If anything, food shares are negatively related to household size (results not shown), especially in urban areas, once we control for per capita household expenditures. This finding may be somewhat surprising in light of the discussion above but it is largely consistent with the multi-country results reported in Deaton and Paxson (1998). For whatever reasons, larger rural households tend to consume more of goods whose prices have disproportionately risen. The finding for urban households is the same, although not as pronounced. Indeed once we control for household demographic composition the impact of household size on the cv measure disappears for urban households.

The second columns in both the urban and rural panels of Table 4 report the results from a regression of the cv measure on the proportion of household members falling into eight age and

gender categories: young (under 5 years) boys and girls, children and adolescents (5-14 years), adult men and women, and male and female elders (ages 60 and over). The excluded reference category is the proportion of adult men in the household. The results do indeed suggest that consumption patterns differ by age and, to a lesser extent, by the gender composition of the household members. Urban and, even more, rural households with a large proportion of young children face a significantly higher cv measure. Households with young children tend to spend more on food, especially rice, and since the prices of these commodities rose the fastest, these households disproportionately suffered. Conversely households with a higher proportion of adult women and, especially, elderly women (and elderly men in urban areas) tend to face lower cvs, in part reflecting the relatively lower food needs of these groups. Thus in addition to urban/rural status, provincial location, and overall income, other important factors that mediate the crisis impact on the household level include household size (in rural areas) and household composition.

VI. Conclusions

Analyzing the distributional impacts of economic crises is important and, unfortunately, an ever more pressing need. If policymakers are to intervene to help those most adversely impacted, then policymakers need to identify those who have been most harmed and the magnitude of that harm. Furthermore, policy responses to economic crises typically must be timely. In this paper, we've developed a simple methodology to fill the order and we've applied our methodology to analyze the impact of the Indonesian economic crisis on household welfare there. In particular, we estimated the compensating variation for Indonesian households following the 1997 Asian currency crisis. We found that virtually every household was severely impacted, although it was the urban poor that fared the worst. The ability of poor rural households to produce food mitigated the worst consequences of the high inflation. We found that the distributional consequences were the same whether we allow households to substitute towards relatively cheaper goods or not. Furthermore, these findings were not biased by any missing consumption, however we computed very different results, even opposite in some sense, if we ignored the relative benefits of self-production or owned housing. Finally, even within urban or rural areas, the geographic location of the household greatly mattered and households with young children suffered disproportionately adverse effects.

While our methodology is simple and uses more or less readily available data, it is not perfect. Two limitations in particular need to be kept in mind. First, it is easy to forget that the Indonesian economic crisis was not the only change in the economic environment over this

period. Concurrent with the crisis, some areas of Indonesia were hard hit by forest fires and others by drought. These and other disasters impact prices so that not all the price changes we observe in the data are due solely to the economic crisis. Put another way, prices would have changed some even absent the crisis. Our analysis speaks to the net effect of the many concurrent economic changes Indonesian households faced. We do not make any attempt to decompose what portion of the actual price changes are due to the financial crisis.

Second, all of the presented analysis concerns nominal changes. In terms of the real impacts of the crisis we are mute. We have estimated compensating variation, but we have no household information on actual changes in income or wages. This missing piece prevents us from discussing the real costs of the crisis. However any attempt to comprehensively measure the real costs to households requires time and energy intensive data collection. As a result, the ultimate results may be available long after policymakers have responded to the crisis with new or modified social policies. Because informational needs are relatively immediate, the simple measures presented should be useful. Exactly how these measures predict actual outcomes remains a topic of ongoing research. The necessary data are just now becoming available.

References

Benjamin, Dwayne. 1992. "Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models". Econometrica, 60(2): 287-322.

Bidani, Benu and Martin Ravallion. 1993. "A Regional Profile of Poverty in Indonesia". Bulletin of Indonesian Economic Statistics, 29(3): 37-68.

Case, Anne C. "Spatial Patterns in Household Demand". 1991. Econometrica, 59(4): 953-966.

Chaudhuri, Shubham and Martin Ravallion. 1994. "How Well Do Static Indicators Identify the Chronically Poor?". Journal of Public Economics, 53(3): 367-394.

Deaton, Angus. 1988. "Quality, Quantity, and Spatial Variation of Price". American Economic Review, 78(3): 418-431.

Deaton, Angus. 1990. "Price Elasticities from Survey Data: Extensions and Indonesian Results". Journal of Econometrics, 44(3): 281-309.

Deaton, Angus. 1997. The Analysis of Household Surveys. Johns Hopkins University Press, Baltimore, MD.

Deaton, Angus and John Muellbauer. 1980. Economics and Consumer Behavior. Cambridge University Press, Cambridge, United Kingdom.

Deaton, Angus and Christina Paxson. 1998. "Economies of Scale, Household Size, and the Demand for Food". Journal of Political Economy, 106(5): 897-930.

Fan, Jianqing. 1992. "Design-adaptive Nonparametric Regression". Journal of the American Statistical Association, 87(4): 998-1004.

Frankenberg, Elizabeth, Duncan Thomas, and Kathleen Beegle. 1999. "The Real Costs of Indonesia's Economic Crisis: Preliminary Findings from the Indonesia Family Life Surveys". RAND Working Paper No. 99-04, Santa Monica, CA.

Friedman, Jed. 2001. "Determinants of Poverty Change in Indonesia, 1984-1996". Manuscript, University of Michigan.

Levinsohn, James, Steven Berry, and Jed Friedman. 1999. "Impacts of the Indonesian Economic Crisis: Price Changes and the Poor". NBER Working Paper No. 7194, Cambridge, MA.

Ravallion, Martin. 1994. Poverty Comparisons. Fundamentals in Pure and Applied Economics, Vol. 56. Harwood Academic Press, Chur, Switzerland.

Ravallion, Martin and Benu Bidani. 1994. "How Robust is a Poverty Profile?" The World Bank Economic Review, 8(1): 75-102.

Smith, James, Duncan Thomas, Elizabeth Frankenberg, Kathleen Beegle, and Graciela Teruel. Forthcoming. “Wages, Employment, and Economic Shocks: Evidence from Indonesia”. Journal of Population Economics.

Suryahadi, Asep, Sudarno Sumarto, Yusuf Suharso, and Lant Pritchett. 2000. “The Evolution of Poverty during the Crisis in Indonesia, 1996 to 1999”. Policy Research Working Paper No. 2435. The World Bank, Washington D.C.

Thomas, Duncan, Kathleen Beegle, Elizabeth Frankenberg, Bondan Sikoki, John Strauss, and Graciela Teruel. 2001. “Education in a Crisis”. Manuscript, Department of Economics,UCLA.

WHO/FAO/UNU Expert Panel. 1985. Energy and Protein Requirements. WHO Technical Report #724, Geneva.

Table 1. Budget shares and price changes for selected aggregate goods

Product aggregate	Mean budget shares			Price changes	
	Bottom decile	All households	Top decile	Mean price increase	Standard deviation
Rice	0.269	0.164	0.048	195.2%	29.2
Other cereals & tubers	0.030	0.010	0.003	137.5%	101.8
Fish	0.033	0.040	0.032	89.1%	67.4
Meat	0.008	0.025	0.040	97.0%	49.3
Dairy & eggs	0.015	0.027	0.031	117.1%	31.9
Vegetables	0.034	0.032	0.020	200.3%	129.5
Pulses, tofu, & tempeh	0.025	0.023	0.012	95.2%	76.0
Fruit	0.016	0.021	0.027	103.7%	61.3
Oils	0.040	0.030	0.015	122.0%	74.8
Sugar, coffee, and tea	0.041	0.034	0.019	142.9%	28.3
Prepared food and beverages	0.025	0.047	0.058	81.4%	51.7
Alcohol, tobacco, and betel	0.039	0.049	0.031	93.9%	43.8
Housing, fuel, lighting, and water	0.146	0.162	0.223	23.8%	10.9
Health	0.010	0.014	0.021	50.7%	32.9
Education	0.013	0.021	0.037	55.3%	31.9
Clothing	0.044	0.045	0.041	84.4%	25.2
Durable goods	0.013	0.034	0.075	114.3%	34.3

Note: Price increases are from January 1997 through October 1998. Mean price increases are computed as the average across all households reporting positive consumption for a given good. Mean budget shares are reported for the entire sample, as well as separately for the top and bottom expenditure decile.

Source: Authors' calculations from 1996 SUSENAS and BPS Price Data.

Table 2. Compensating variation by expenditure decile and poor/non-poor status

Expenditure decile	Compensating variation		
	All	Urban	Rural
1	0.73	1.08	0.67
2	0.79	1.03	0.73
3	0.82	1.00	0.74
4	0.83	0.96	0.77
5	0.84	0.93	0.77
6	0.85	0.92	0.78
7	0.85	0.89	0.78
8	0.85	0.84	0.79
9	0.84	0.81	0.79
10	0.77	0.70	0.81
Poor	0.77	1.09	0.70
Non-poor	0.82	0.90	0.78
All households	0.82	0.91	0.76

Note: Compensating variation measured as a proportion of 1996 household expenditures

Source: Authors' calculations from 1996 SUSENAS and BPS Price Data.

Table 3. Estimated price elasticities for aggregate food goods and residual consumption

Product	Rice	Other cereals	Tubers	Fresh fish	Preserved fish	Fresh meat	Preserved meat	Eggs	Dairy	Green vegetables	Other vegetables
Rice	-0.479	0.082	-0.032	-0.029	-0.038	0.098	-0.016	-0.008	-0.009	-0.018	0.003
Other cereals	2.762	-5.046	-0.413	-0.074	0.387	-0.200	-0.134	-0.300	-0.014	0.048	-0.667
Tubers	2.521	-0.127	-0.590	0.233	0.205	-0.672	0.087	-0.531	-0.167	-0.919	-0.147
Fresh fish	-0.383	0.027	0.217	-0.996	0.026	0.169	0.219	-0.087	-0.012	0.026	0.071
Preserved fish	-0.533	-0.295	-0.059	0.373	-0.686	0.013	-0.015	-0.022	0.138	-0.103	0.091
Fresh meat	0.042	0.073	-0.046	0.056	0.118	-0.616	-0.004	-0.134	0.109	-0.135	-0.173
Preserved meat	-0.224	0.318	0.127	0.256	0.254	-0.418	0.955	-0.281	-0.260	-0.215	-0.106
Eggs	-0.458	0.128	0.013	-0.006	-0.080	0.084	-0.080	-0.985	-0.028	0.113	0.031
Dairy	-0.194	0.121	0.097	-0.072	-0.083	-0.216	0.548	0.040	-0.133	0.077	-0.002
Green vegetables	-0.384	0.097	0.189	-0.202	-0.041	-0.067	0.136	0.014	-0.023	-0.789	0.057
Other vegetables	-0.465	-0.005	-0.042	0.125	0.017	-0.115	0.074	0.034	-0.004	0.002	-0.840
Pulses	-0.406	0.367	-0.001	-0.153	-0.064	0.266	-0.271	-0.248	-0.474	-0.014	0.169
Tofu & tempeh	-0.104	0.077	0.010	0.102	-0.033	-0.111	-0.159	0.160	-0.025	0.017	0.052
Fruit	-0.181	-0.144	-0.141	0.098	-0.006	-0.253	0.044	-0.147	-0.110	-0.021	0.003
Oils	-0.238	-0.012	0.027	-0.143	-0.003	-0.136	-0.019	-0.004	0.007	-0.009	0.022
Beverage additives	-0.173	0.059	0.044	-0.167	0.013	0.001	-0.111	-0.047	-0.106	0.064	-0.036
Spices	-0.210	-0.018	0.104	-0.072	-0.007	0.000	-0.034	-0.057	-0.107	0.032	0.023
Other food	0.140	-0.056	0.069	-0.027	0.004	-0.238	0.098	0.112	0.013	0.029	-0.010
Prepared food	0.020	0.243	0.055	0.092	-0.006	-0.037	-0.037	0.060	-0.093	0.042	-0.033
Prepared beverages	-0.429	0.026	-0.083	0.246	0.005	0.034	0.259	-0.191	-0.203	0.146	-0.227
Alcohol	-2.806	-0.681	0.161	0.859	0.265	-2.175	2.039	0.506	-0.012	-0.770	0.447
Tobacco and betel	-0.441	0.053	0.001	-0.104	-0.037	0.151	-0.182	0.001	-0.025	0.033	0.030
Other consumption	0.010	0.017	0.008	-0.010	-0.003	0.019	-0.003	-0.013	-0.008	-0.002	0.002

Source: Authors' calculations from 1996 SUSENAS.

Table 3. (cont.)

Product	Pulses	Tofu & tempeh	Fruit	Oils	Beverage additives	Spices	Other food	Prepared food	Prepared beverages	Alcohol	Tobacco and betel	Other consumption
Rice	0.001	0.023	-0.053	0.058	0.032	-0.042	-0.037	-0.026	0.032	0.036	-0.138	0.274
Other cereals	0.002	0.262	-0.195	0.289	-0.239	0.086	-0.210	-0.292	0.430	0.649	-0.276	2.684
Tubers	0.010	0.467	-0.142	-0.590	-0.387	-0.196	0.621	0.225	0.660	-0.292	0.052	-0.822
Fresh fish	0.042	-0.038	0.165	-0.012	-0.104	0.135	-0.003	-0.025	0.185	0.128	0.304	-1.065
Preserved fish	-0.186	0.524	-0.198	0.011	-0.014	-0.093	0.058	-0.043	-0.143	0.164	-0.603	1.066
Fresh meat	0.092	-0.091	-0.023	-0.182	0.052	0.065	0.113	0.163	-0.075	-0.047	-0.063	-1.091
Preserved meat	0.259	-0.212	-0.094	-0.276	0.325	0.007	0.372	0.269	1.018	0.330	0.540	-5.594
Eggs	0.071	-0.052	0.027	0.019	-0.155	0.131	0.033	0.007	-0.019	-0.135	0.015	0.390
Dairy	0.122	-0.097	0.132	-0.150	0.216	0.283	-0.145	0.124	-0.296	-0.037	0.441	-2.696
Green vegetables	0.060	-0.099	0.033	0.110	0.106	0.138	0.004	0.064	-0.023	0.257	-0.086	-0.088
Other vegetables	0.041	-0.001	-0.114	-0.134	0.031	0.061	-0.095	-0.028	0.156	0.242	-0.096	0.432
Pulses	-0.772	-0.136	0.120	-0.135	-0.378	0.216	0.192	-0.048	-0.149	-0.281	0.038	0.936
Tofu & tempeh	-0.057	-0.965	0.035	0.032	0.059	0.292	-0.194	0.022	-0.056	-0.106	-0.069	0.606
Fruit	0.115	-0.006	-0.831	-0.111	-0.094	0.074	0.005	0.030	-0.060	-0.069	0.016	0.478
Oils	0.008	0.058	-0.064	-1.003	-0.039	0.038	-0.017	-0.031	0.144	0.019	0.036	0.789
Beverage additives	0.056	-0.029	-0.082	0.109	-0.625	0.017	0.029	-0.013	0.132	0.101	0.019	0.055
Spices	0.049	-0.073	-0.016	0.032	-0.033	-0.305	0.027	-0.062	-0.005	-0.080	-0.114	0.248
Other food	0.091	-0.059	0.027	0.019	0.039	0.177	-1.161	0.064	-0.008	0.114	0.139	-0.763
Prepared food	-0.078	-0.090	0.209	-0.044	-0.060	0.062	0.089	-0.775	-0.339	-0.135	0.124	-0.401
Prepared beverages	-0.040	-0.206	0.025	-0.162	0.720	0.137	-0.231	0.310	1.912	-0.517	-0.296	-3.857
Alcohol	0.143	0.204	-0.311	-0.545	-0.602	0.668	-0.569	1.183	1.501	6.106	1.226	-9.039
Tobacco and betel	-0.070	0.061	-0.017	0.070	0.067	-0.025	0.026	-0.147	-0.260	0.023	-0.876	0.664
Other consumption	0.007	0.005	-0.001	0.002	0.013	0.029	-0.003	0.002	-0.001	0.020	-0.011	-0.482

Source: Authors' calculations from 1996 SUSENAS.

Table 4. CV Regressions with household demographic controls

Independent variables	Rural households		Urban households	
ln(Household PCE)	0.0919	0.0948	-0.1709	-0.1720
	0.0144	0.0147	0.0075	0.0076
ln(Household size)	0.1034	0.0722	0.0293	0.0013
	0.0105	0.0127	0.0067	0.0068
Proportion of household:				
male, 0-4 years old	--	0.1362	--	0.0817
		0.0401		0.0211
female, 0-4 years	--	0.1224	--	0.0758
		0.0404		0.0235
male, 5-14 years	--	0.0511	--	0.0289
		0.0300		0.0158
female, 5-14 years	--	0.0186	--	-0.0106
		0.0307		0.0165
male, 15-59 years	--	--	--	--
female, 15-59 years	--	-0.0516	--	-0.0293
		0.0276		0.0139
male, 60 years or more	--	-0.0226	--	-0.0791
		0.0352		0.0303
female, 60 years or more	--	-0.1441	--	-0.1857
		0.0322		0.0276
R^2	0.1117	0.1131	0.2620	0.2677
Unweighted N	37493		24472	

Note: OLS regressions include age, gender, and education of household head as well as province dummies. Standard errors, reported below the estimated coefficients, are corrected for dependence within survey clusters.

Source: Authors' calculations from 1996 SUSENAS and BPS Price Data.

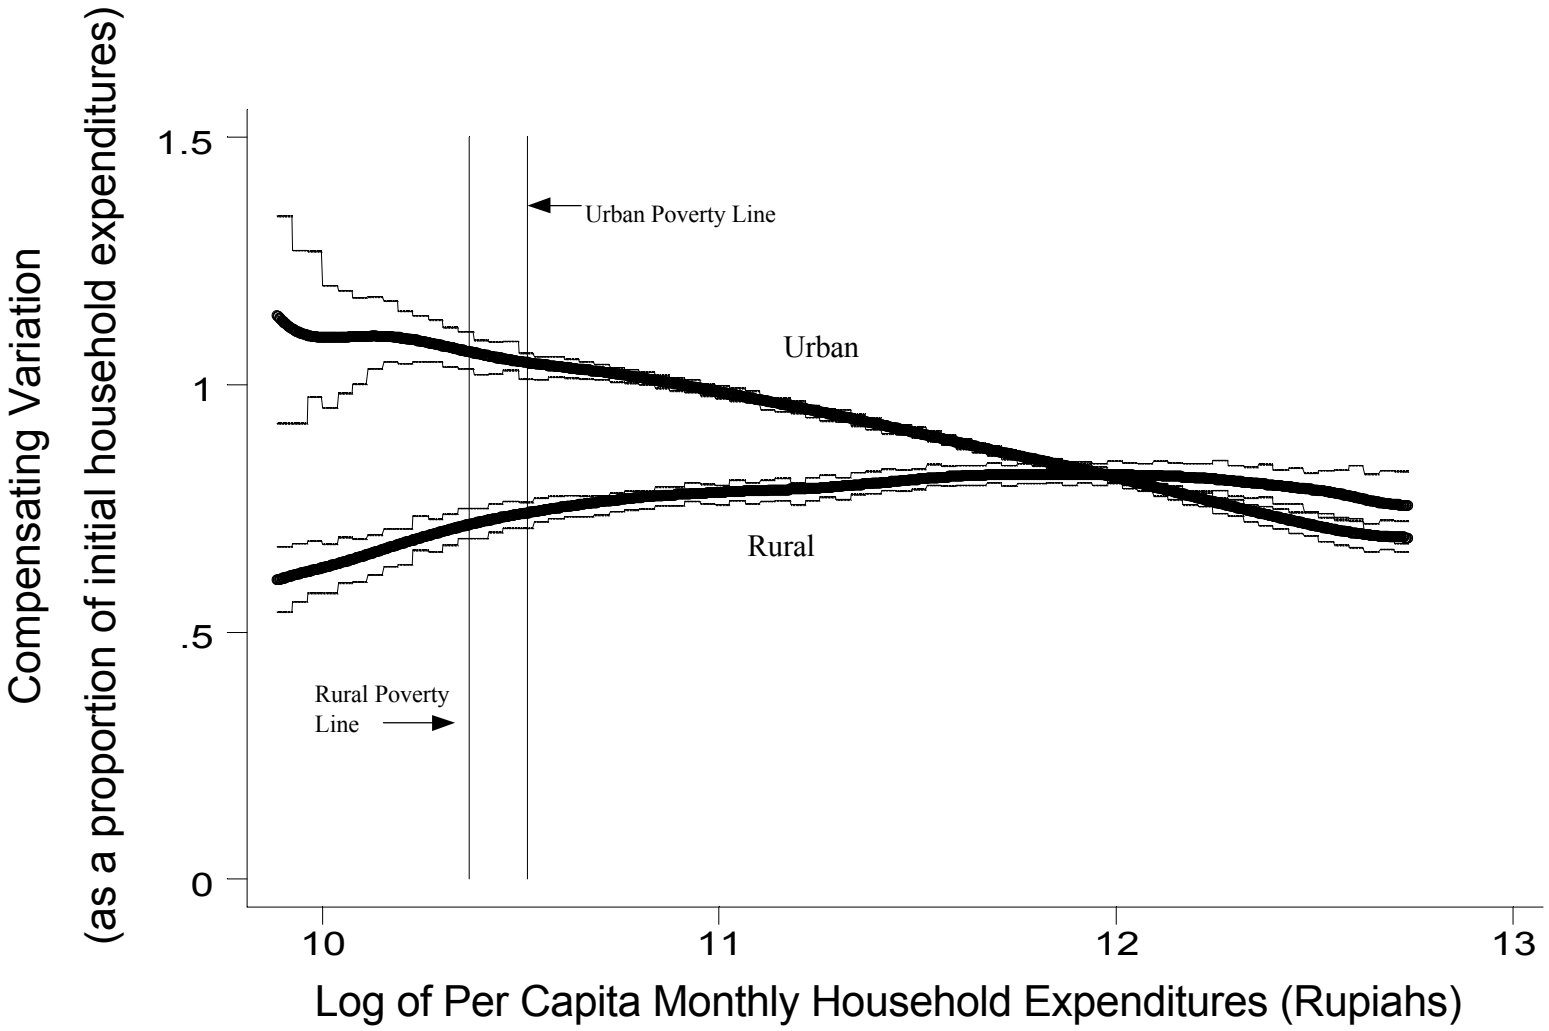


Figure 1. Compensating Variation, with 95% Confidence Intervals

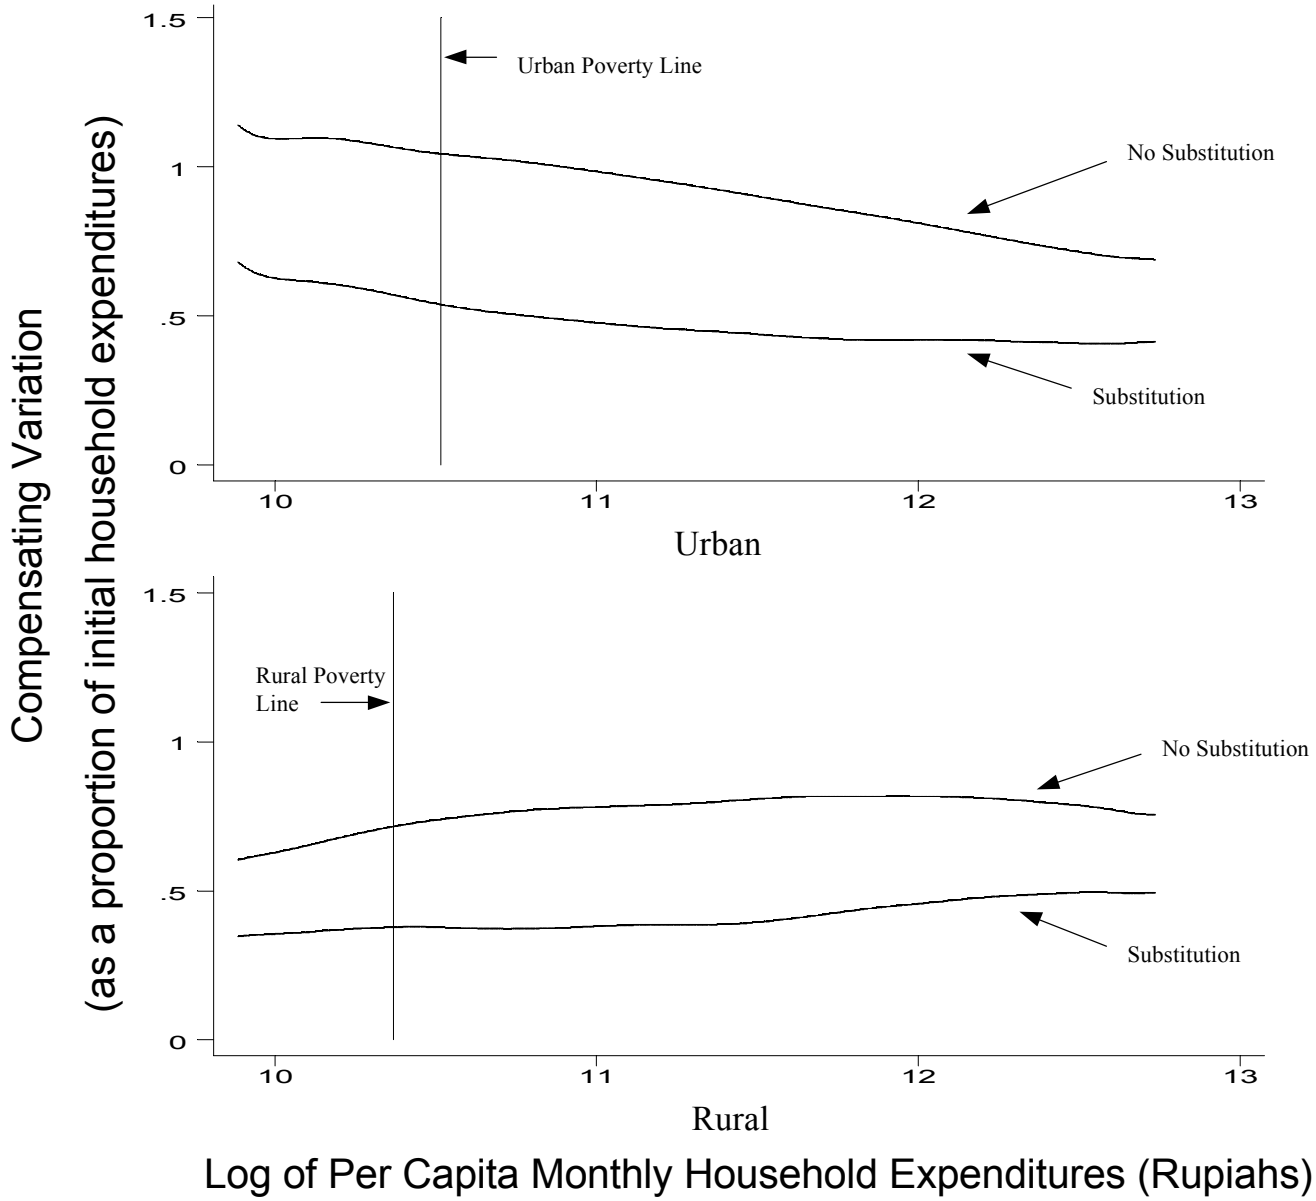


Figure 2. Compensating Variation, with and without Substitution Effects

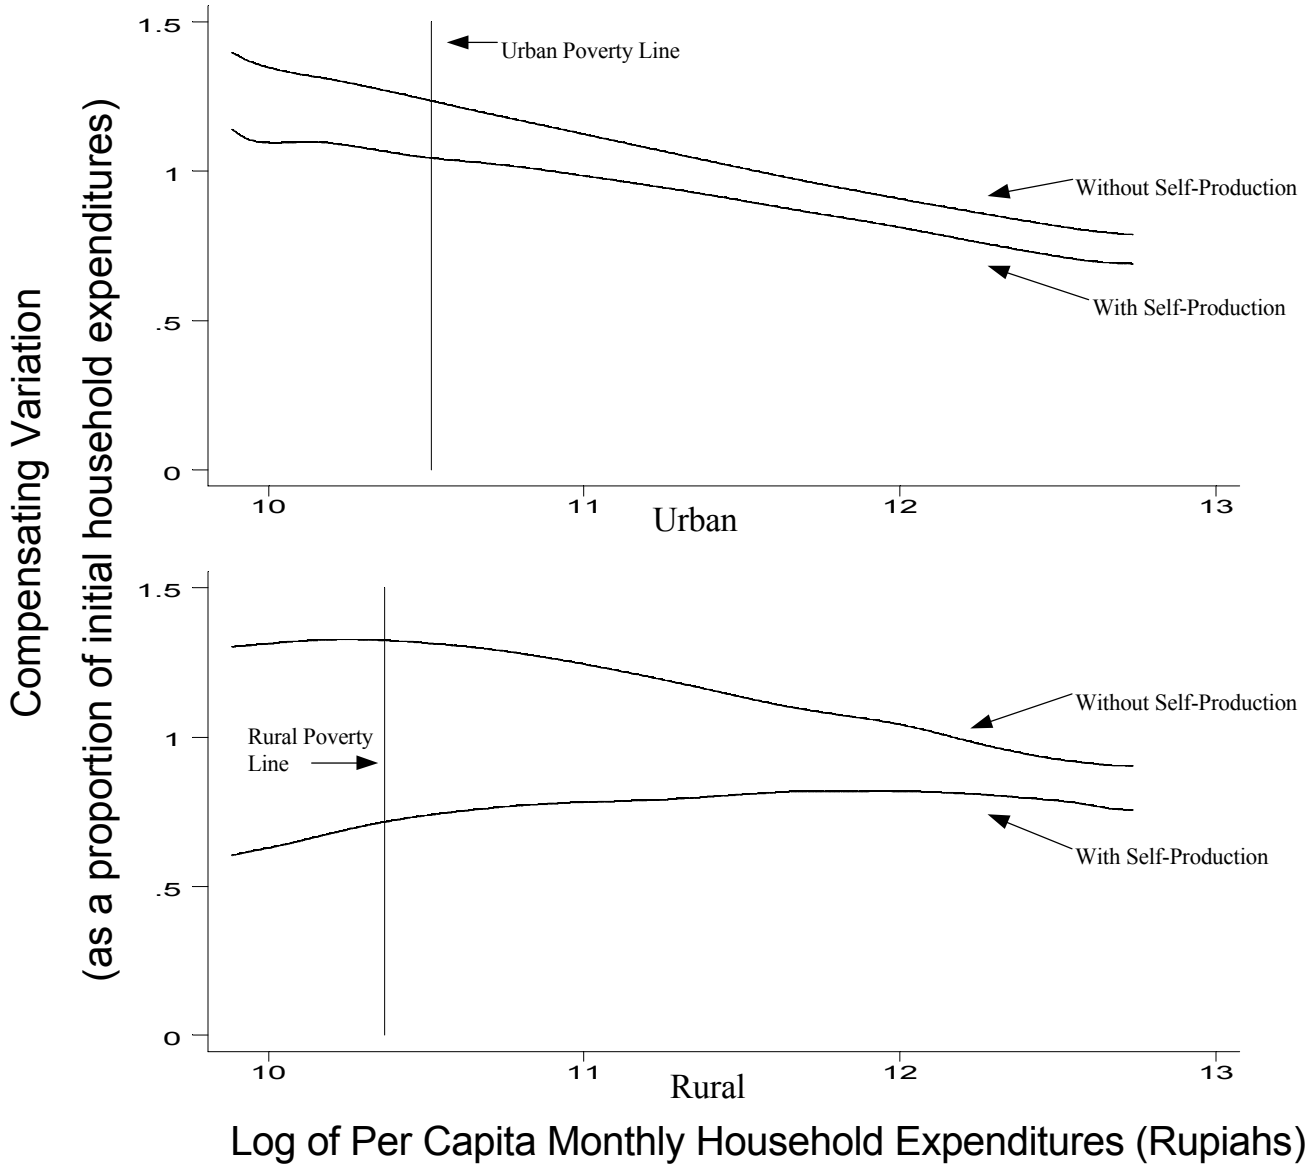


Figure 3. Compensating Variation, with and without Self-Production

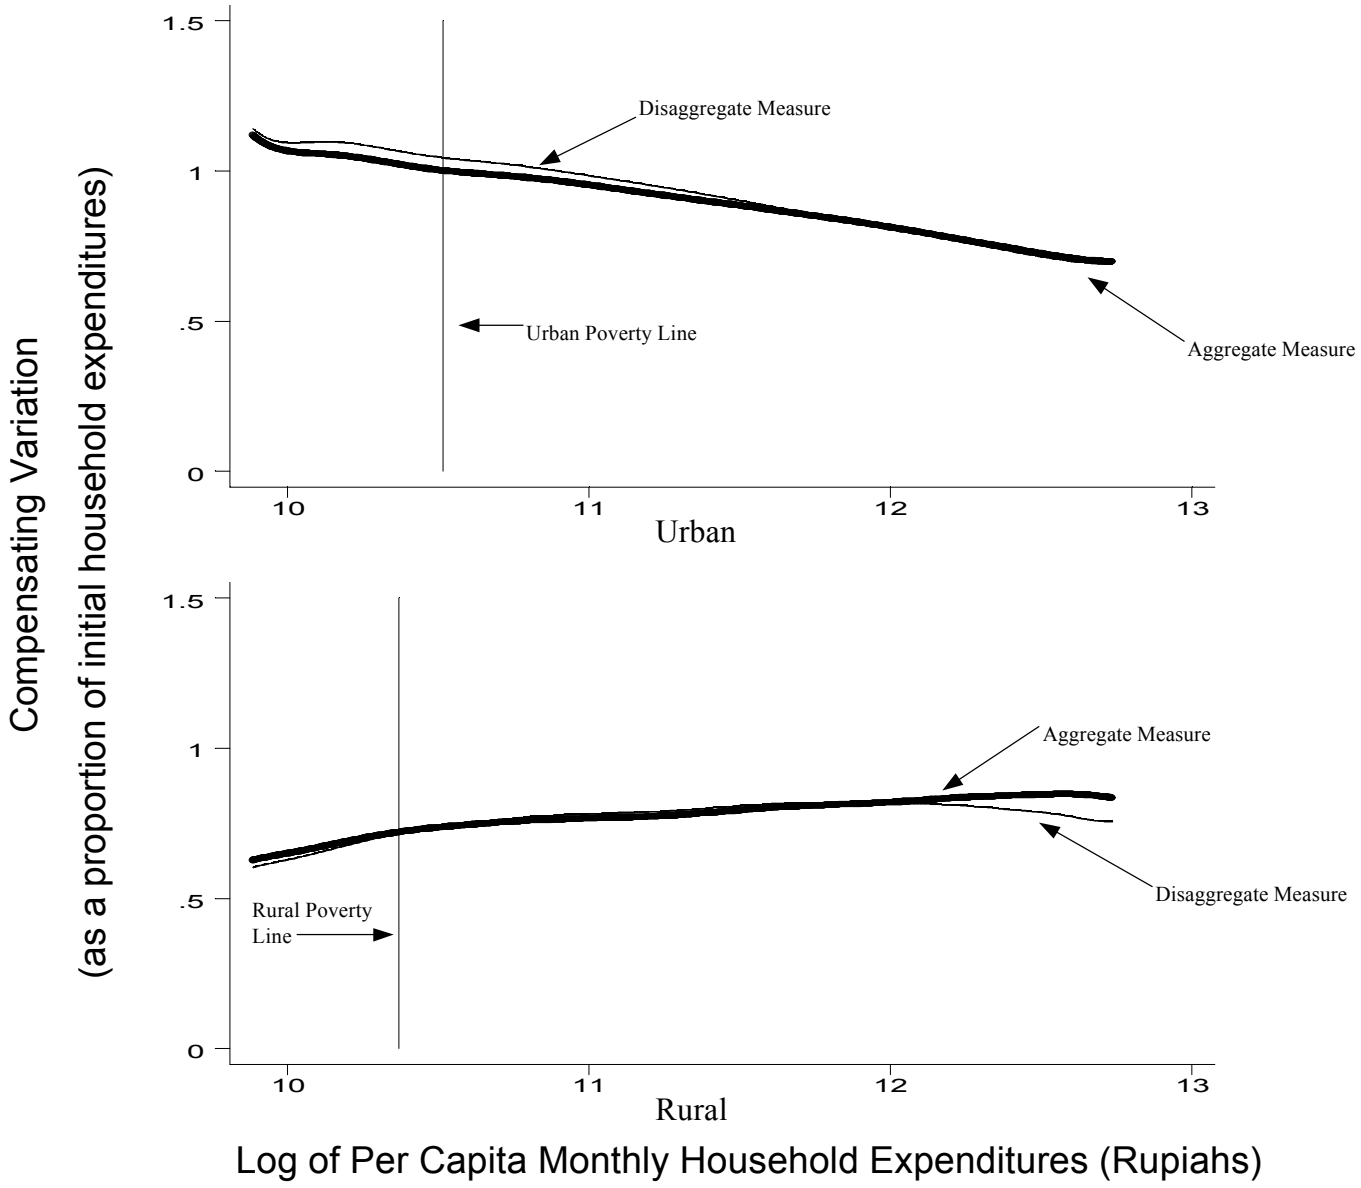


Figure 4. Compensating Variation, Aggregate and Disaggregate Measures

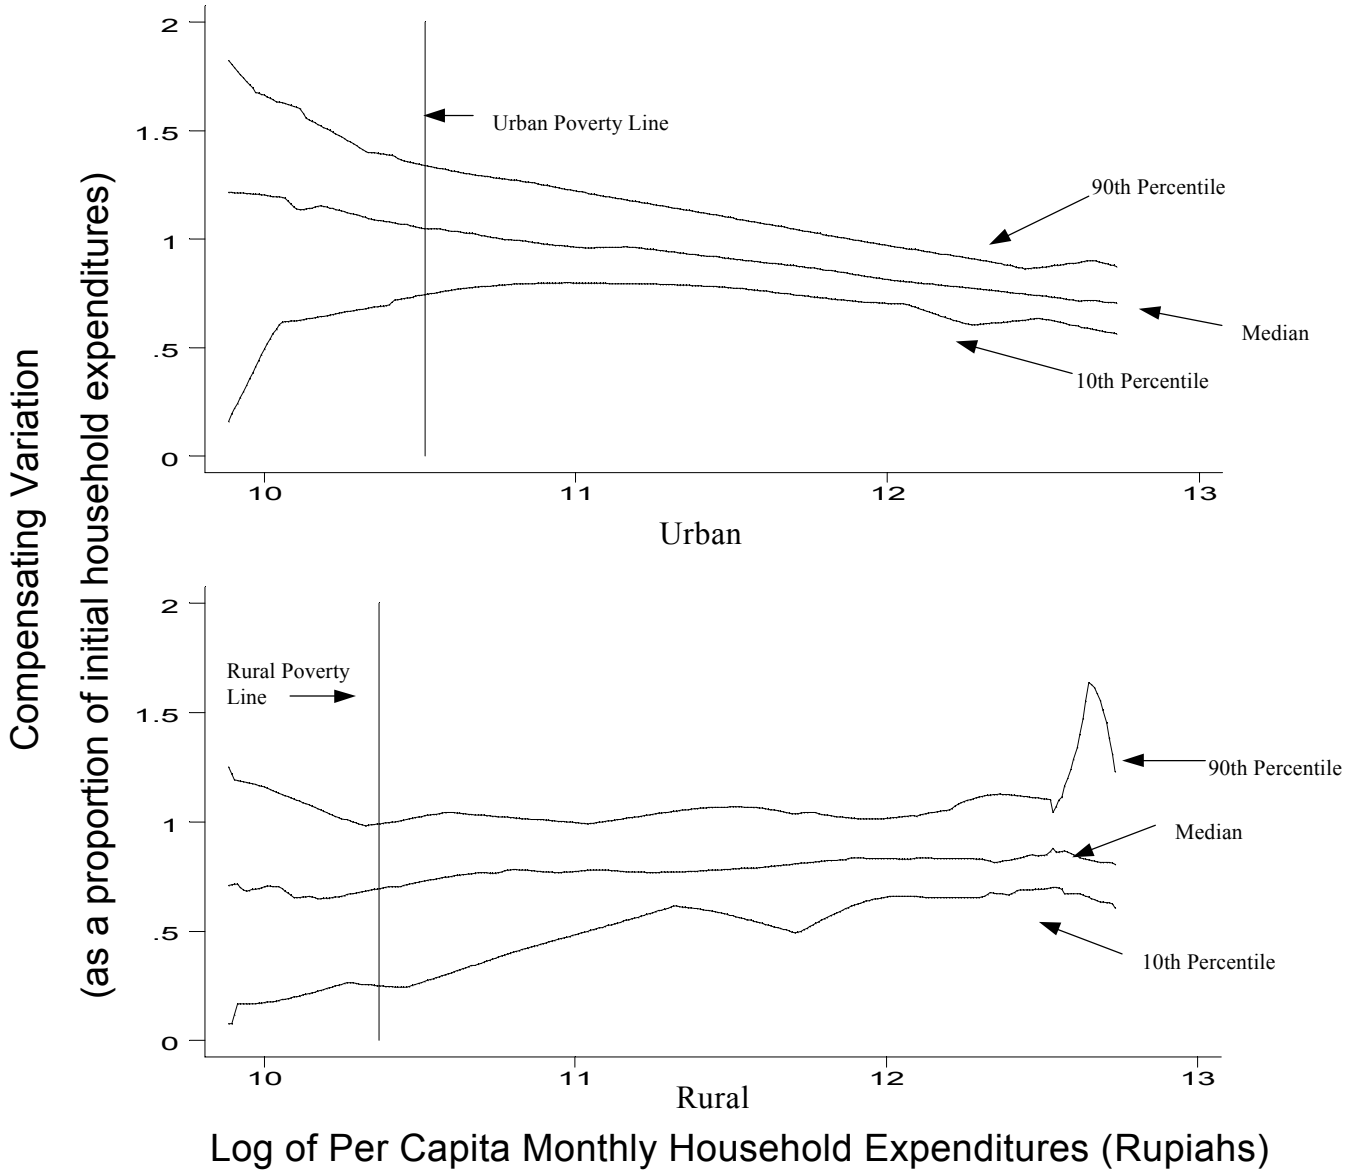


Figure 5. Dispersion of Compensating Variation Across Provinces, 10th, 50th, and 90th Percentiles (with Self-Production)