RESEARCH SEMINAR IN INTERNATIONAL ECONOMICS

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

Discussion Paper No. 446

Impacts of the Indonesian Economic Crisis: Price Changes and the Poor

James Levinsohn
University of Michigan

Steven Berry
Yale University

Jed Friedman
University of Michigan

June 22, 1999

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

Impacts of the Indonesian Economic Crisis: Price Changes and the Poor

by

James Levinsohn
University of Michigan
National Bureau of Economic Research
Steven Berry
Yale University
National Bureau of Economic Research
Jed Friedman
University of Michigan

Current version: June 22, 1999

Address. Berry: Department of Economics, 37 Hillhouse Ave., Yale University, New Haven, CT 06520. steveb@econ.yale.edu; Friedman: Department of Economics, University of Michigan, Ann Arbor, MI 48109. jedf@umich.edu; Levinsohn: Department of Economics and School of Public Policy, University of Michigan, Ann Arbor, MI 48109. jamesL@umich.edu

Impacts of the Indonesian Economic Crisis: Price Changes and the Poor

James Levinsohn
University of Michigan
National Bureau of Economic Research
Steven Berry
Yale University
National Bureau of Economic Research
Jed Friedman
University of Michigan

1. Introduction

In July 1997, following the decline of the Thai Bhat, the Indonesian Rupiah fell dramatically (or so it seemed at the time.) Since that initial decline of the Rupiah, the Indonesian economy has undergone tremendous change. The Rupiah has been subject to large swings, prices of some goods have risen substantially, and billions of dollars have been loaned by international lending organizations. These are not subtle changes. In this paper, we make a first-pass attempt at providing early estimates of the impact of the Indonesian economic crisis on Indonesia's poor.

While some might argue that the very poor are so impoverished that they are essentially insulated from swings in the international economy, it is more frequently argued that the very poor are among the most vulnerable to such swings. This is especially probable for the urban poor. Furthermore, in countries with little or no social insurance, any impacts of price changes on the very poor are unlikely to be muted by government policies in the way that they might be in richer countries.

These issues matter. From a broad humanitarian view, the magnitude of the price changes and the size of the impacted population argue that there is value to simply understanding what has happened. From a more narrow political view, the political economy of price changes may well depend in crucial ways on who bears the brunt of price increases. From the viewpoint of organizations such as the IMF which offer policy advice and (sometimes) loan conditionality, understanding how that advice might impact the poor is important. Finally, from a ridiculously narrow academic perspective, there is not an abundance of research on possible links between the international economy and the very poor.

We are grateful to Anderson Ichwan and Edwin Pranadjaja for invaluable assistance in translating from Indonesian to English. We are grateful to Alan Winters for helpful suggestions and encouragement. The paper has benefitted from suggestions from Mark Gersovitz and Gene Grossman. Thanks to the World Bank and the Swedish Ministry for Foreign Affairs for research support. The views expressed in this paper are those of the authors and do not necessarily represent those of the World Bank.

2. Some Background

We begin by setting the stage. The changes the Indonesian economy has undergone are dramatic. The purpose of this section is to very briefly review some of those changes. As background, Table 1 provides some information on recent changes in prices and exchange rates. From December 1996 until July 1997, the Rupiah traded in a narrow range of around 2400 to the U.S. dollar. The consumer price index (CPI) provided by the Bank of Indonesia shows stable prices for each of four aggregates-- food, housing, clothing, and health. In July 1997, the Thai Bhat nose-dived and the Rupiah followed suit. In Table 1, this appears in the August 1997 entry where the Rupiah is reported at 3035 to the dollar. Although this was a sudden depreciation on the order of 20 percent, prices rose only with a lag. Throughout the remainder of 1997, the Rupiah continued to depreciate against the dollar and (except in November) against the Yen. The food CPI rose from 105 to 120-- a noticeable increase but not an overwhelming one. The CPI for housing, clothing, and health care rose yet more modestly. On the economic policy front, the IMF approved a $10 billion loan while the World Bank pledged $4.5 billion for a three year program.

It was not until 1998 that matters became considerably more problematic. On January 8, sometimes referred to as ''Black Thursday,'' the Rupiah began a free fall and news accounts reported panic-like food purchasing. The exchange rate fell at one point in January to above 16,000 Rupiah per dollar and the CPI for food jumped almost as much in January as it had the previous six months combined. The CPI for clothing jumped even more. As international pressure to drop a proposed currency board increased, and aid was deferred, uncertainty mounted. For the first four months of 1998, prices continued to rise as documented in the last four columns of Table 1. In May 1998, riots spread and over a thousand people were reported killed. The World Bank postponed two loans totaling over a billion dollars, and the Bank and Fund as well as many embassies evacuated non-essential staff. On May 21, President Suharto resigned. The Rupiah traded at around 11,000 immediately after the resignation. The Bank of Indonesia reported the largest monthly Rupiah to dollar rate in June 1998-- 14,900. Thereafter the Rupiah began a gradual appreciation (albeit from an astoundingly low level.) The CPI reported rising prices through September 1998. The CPI for food reached 261 (relative to a level of about 100 in January 1997) while the CPIs for housing, clothing, and health hit 156, 225, and 204 respectively. (Throughout this period the CPI for housing was relatively more stable-- perhaps reflecting the somewhat non-traded nature of housing.) Although peaceful protests turned violent in Jakarta in mid-November 1998, order was quickly restored.

It would of course be a tremendous oversimplification to attribute these changes to the international economy-- or to any other single cause. Price levels and exchange rates are endogenously determined. Our goal is to analyze the impact of the changes surveyed in Table 1, but we do not attempt to analyze the root cause(s) of the macroeconomic changes. We realize, for example, that it is (barely) conceivable that

purely domestic inflation suddenly ran rampant leading to the Rupiah's depreciation, and in this (unlikely) scenario, the price changes in Table 1 would have little to do with the international economy. Given most accounts of the East Asian crisis and the contagious behavior of other East Asian exchange rates and price levels, it seems plausible that there was indeed an international element to the changes surveyed in Table 1.

Our goals include a more disaggregated analysis of the impacts of the price changes. The aggregated nature of the figures in Table 1 hide potentially important heterogeneity. The first type of heterogeneity concerns heterogeneity within commodity groups. For example, ''Food'' contains hundreds of items and it is possible that the price behavior of those items consumed by the non-poor is quite different than the price behavior of food items consumed principally by the poor. The second type of heterogeneity is geographical. Indonesia is a geographically dispersed country where simple arbitrage may be costly due to transport costs. This suggests that there may be significant price variation within a narrowly defined product class across geographic areas. What happens to prices in especially poor areas may be quite different from what happens to prices in the wealthier areas. The third type of heterogeneity is across consumers. Our focus is *not* on the representative consumer, but rather we care about the consumption patterns of the very poor. Examining aggregate consumption patterns may be quite misleading in this context.

3. Data Concerns and Constraints

There are many ways one could estimate how the large changes in prices in Indonesia over the last one and a half years have impacted the poor. In the end, the methods used will depend quite crucially on the available data. With this in mind, we briefly outline the data that are, and are not, available. We begin with the unattainable ideal. In the best case, one would have detailed consumption data that spanned the period before and after the financial crises of 1997-98 for thousands (or tens of thousands) of households. The time series variation would allow the researcher to examine how consumption patterns changed when faced with the large price changes. The large household survey would give the researcher enough households so that a focus on the very poor would still allow a sufficient number of observations. It would also be important to have detailed price data on a disaggregated set of commodities. This data would need to cover the most recent two years. Even this data, ideal and unattainable as it is, would pose significant econometric issues due to the nature of the questions posed. This is because what we want to know is how households in a particular part of the income distribution behaved in response to price changes and even the most sophisticated demand systems typically estimate a utility-consistent demand structure for a *representative* consumer. While, with infinite data, one could estimate a demand structure for just a particular decile of the income or wealth distribution, this would be massively inefficient. (A topic of future research is the estimation of a utility consistent demand system that explicitly accommodates the heterogeneity inherent in studying how the poor respond to price changes.)

In fact, the data described above simply do not exist.[1] The good news, though, is that reality is less far-removed from this ideal than is usually the case. Indonesian data sources are in fact quite good. Indonesia conducts an extensive household consumption survey (SUSENAS) covering on the order of 50,000 households. Most recently, these surveys have been conducted in 1981, 1984, 1987, 1990, and 1993.[2] While the surveys are large, they are not panels. That is, there is no systematic effort to track the same households over time. These surveys cover a wide geographic range of the country and contain very detailed consumption data.[3] The data do not contain prices, though. Rather, the data contain unit values which are defined as expenditure divided by quantity. These unit values may differ across households that in fact face identical prices due to differences in the quality of the households' consumption. (i.e. 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.) This type of data can be (and in fact has been) used to estimate demand elasticities exploiting the spatial variation in the data using methods developed by Deaton (1988).

We base our analysis on consumption data from the1993 SUSENAS-- the most recent wave available to us. The 1993 SUSENAS surveyed 65,600 households throughout the entire country. We have reduced our sample to the 58,100 households that have sufficient consumption and household information for the the analysis which follows. To the extent that consumption patterns change over time, we are concerned about the accuracy of using 1993 consumption data to measure behavior in 1997. We investigate this by examining expenditure patterns as they evolved over the course of prior waves of the SUSENAS. We found some definite trends. In particular, the proportion of expenditure on food decreases slightly but steadily across each SUSENAS. This is probably due to rising real incomes. These trends may have persisted until 1997. To the extent that our consumption baskets are calculated with 1993 and not 1997 data, our measured impacts of the crisis will diverge from the actual impacts. However one of our primary concerns is to highlight the heterogeneous effects of the crisis among households. The relative consumption baskets (among rich

[1] A special wave of the Indonesia Family Life Survey was conducted in late 1998 to investigate the immediate effects of the crisis. This data set, a true panel of households, can compare household consumption in late 1998 to a corresponding period one year earlier. Frankenberg, et al. (1999) summarize the initial findings. The study surveys 1900 households in 7 provinces and thus does not provide the geographic coverage or sample size suitable for our purposes.

[2] A survey was also conducted in 1996, but we have not been provided with those data yet.

[3] For 203 individual food items, the survey recorded the quantity and value consumed by the household in the last week. For 89 individual and aggregate non-food items, the survey recorded annual expenditures as well as expenditures in the month preceding the survey. For those households that consumed their own self-produced food, the survey imputed the value of that food. For those households that owned housing, SUSENAS imputed a monthly rental payment.

and poor households, or rural and urban households) did not change as much as the absolute consumption baskets over the 1993-97 period and as such the bias along this dimensions is likely to be slight.

We also have very recent price data that has been supplied by the Badan Pusat Statistik (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 stabilization at the new higher rate. We employ a single price change measure--the percent change in prices from January 1997 to October 1998. By adopting such a long time period from before the onset of rapid inflation until after the inflation had largely abated, we hope to capture a robust measure of the price changes brought on by the crisis.

The price data supplies price information for both aggregate goods, such as food or housing, as well as individual goods such as cassava or petrol. There are approximately 700 goods with observed prices in the data. However, the types of goods observed varies by city, perhaps reflecting taste and consumption heterogeneity throughout the country. On average, a particular city has price information on 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 whole 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 may 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 study suggests that overall inflation in rural areas is approximately five percent higher than in urban areas (Frankenberg, et al., 1999).[4] We frequently report separate results for the urban and rural poor, and the fact that the price data were collected in the cities should be kept in mind as those results are reviewed.

[4] The same study also presents some evidence that the BPS price data may understate inflation by as much as 15 percent. To the extent that this is true, the impact of the crisis is even greater than measured here.

For certain groups of goods the price data are more disaggregated than the consumption data reported in the SUSENAS. In order to link the new price data with the existing consumption data, we use the prices for those commodities which appear in both the price data set and in the SUSENAS. In some cases, we also aggregate commodities in the price data to match a product category in the SUSENAS data.[5] 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 accounts for 75 percent of a household's total expenditure. We return to this point later.

4. Methodologies

Given our data sources, the usual approach to investigating how the Indonesian poor were impacted by the recent crisis would be to do the following. First, one would estimate a demand system-- ideally one based on an underlying utility-consistent framework. The SUSENAS surveys would provide the data for such a demand system. Based on the estimated elasticities from that demand system, one would then estimate the welfare impact of the price changes that occurred recently in Indonesia. Special emphasis would be placed on how the poor were impacted by the crisis. It turns out that there are some very severe problems with this approach given the data and the policy goals. In order to better motivate what we *do* do, we first highlight the problems with the approach outlined above.

Estimating demand elasticities from the SUSENAS is not an especially satisfying endeavor. The SUSENAS is a cross-sectional survey of households. While we do have multiple waves, there is no panel, or time-series, nature to the data. As noted above, the SUSENAS contains data on expenditures and on quantities consumed, but not on prices. Expenditures divided by quantities give unit values and, as outlined in Deaton (1988), there is a misguided temptation to use these unit values as prices. As noted earlier, a naive swap between unit values and prices is wrong because unit values reflect the quality of the product as well as the market price. Deaton shows that under the appropriate separability conditions, one can exploit the spatial nature of the data to back out the true price elasticities. The idea is that within a geographic unit, say a village, the prices will be the same, although they are unobserved by the econometrician. Unit values, though, will differ across households within the village. This within variation allows the econometrician to identify the quality effect-- incomes vary and the observed unit values vary, but by assumption underlying prices are the same. The variation across villages, controlling for village fixed effects, allows one to then back out the true price elasticities, since the real price variation occurs only through the spatial dimension.

[5] In these cases, we take simple averages of the products that comprise a single product in the SUSENAS data.

All of this leads to a multi-step estimation algorithm developed by Deaton (1988). The estimator employed deals quite carefully with the errors-in-variables issues that the use of unit values raises.

So what's the problem? This methodology is probably the best available, but it has some real drawbacks. From an economic perspective, it is troubling that the resulting demand elasticities are not consistent with an underlying utility framework. If at the end of the day one wants to compute a welfare measure such as compensating or equivalent variation, one needs to work with a framework that allows the researcher to identify the primitives of the underlying utility function. From an econometric perspective, it is problematic that the methodology does not deal with the endogeneity of product quality. Consumers choose the quality as well as quantity of the products bought and this induces the usual simultaneity concerns. These issues, though, are perhaps just academic quibbles. The bigger problems arise due to the policy application at hand. Recall, we are concerned with better understanding how the price changes impacted the poor. There are at least three reasons that the methodology is ill-suited to adequately addressing this concern. First, the estimated elasticities are essentially local approximations based on consumer behavior at the observed prices. Hence the SUSENAS might give pretty good estimates of how households respond to a price change on the order of five percent. When the price changes under consideration are instead on the order to 100 to 300 percent, the answer is essentially dictated by the choice of functional forms. This is troubling for most any parametric approach to the estimation of demand elasticities. Second, the underlying framework is one of a representative consumer. Our concern, though, is with anything but the representative consumer. Rather, we are especially focused on the very poor. Fitzgerald wrote that the rich are different. So, we suspect, are the poor. A demand system that explicitly considers consumer heterogeneity is called for, but is not currently available. Finally, it is not feasible to estimate a complete demand system at a highly disaggregated product level. There are simply too many products. The obvious solution is to aggregate products, but this aggregation hides very important variation in consumption patterns and price changes. Alternatively, one can estimate own price elasticities (but not cross-price elasticities) for many disaggregated products.

We have done such an exercise with the SUSENAS data. Employing a simple OLS framework, and controlling for some observed household characteristics, we have estimated own price elasticities for individual food items. We do not attempt to correct for the quality effects discussed above. The elasticities, identified by the cross-sectional variation in unit values and quantities, yield the expected negative coefficients and are quite precisely estimated. For example we estimated the own-price elasticity for rice to be -0.43 with a standard error of .02, while the same estimate for ground coffee yields a coefficient of -.84 with a standard error of .01. Most of the point estimates for the 193 food items fall between -.3 and -.8. Only a handful of estimates exceed -1, perhaps indicating relatively inelastic demand even at the most disaggregate level. When the analysis includes fixed effects for each district (kabupaten), the point

estimates, still precisely estimated, tend to be a bit larger in absolute value but still very few exceed an estimate of -1.[6] Of course these estimated own price elasticities, like most parametric approaches, are subject to the some of the problems mentioned above.

Our principal approach in this paper is non-parametric. As with the econometric approach outlined above, we will need to assume that the 1993 SUSENAS survey provides a reasonably accurate picture of consumption patterns before the crisis. We then use the price changes that actually occurred to predict who the price changes would have impacted. This approach has both advantages and disadvantages. On the upside, it does not rely on functional forms and we can more easily explore the three types of heterogeneity listed above. On the down-side, it essentially ignores the possibility of substituting away from relatively more expensive goods. As such, our method will provide an upper bound on the predicted impacts of the price changes on the poor.

The best approach is to combine the heterogeneity highlighted with the non-parametric approach with the structural economic relationships estimated by the econometric approach. We'll do this, and this exercise has convinced us of the need to do this, but it is a longer-term project.

5. Heterogeneity

Our methods are motivated by our desire to capture the heterogeneity in prices, products, and consumers. We begin our analysis by simply documenting the extent of this heterogeneity. This serves two functions. First, it illustrates the importance of using methods that do not aggregate across the dimensions of heterogeneity. Second, it highlights exactly which sorts of heterogeneity are most important, and this will inform our analysis of the price changes.

Price heterogeneity across regions

We begin by analyzing how prices for narrowly defined products vary across Indonesia over the course of the financial crisis. The raw data that are used for this exercise are monthly prices for about 700 products that are collected on a city basis by the BPS. These data are used to then create the official consumer price indices for the entire country. Monthly prices for so many products in very many cities constitute a rather unwieldly data set. We have aggregated the data in 3 dimensions. In terms of the time-series dimension, we simply computed the price change for each product for the period spanning January 1997 to October 1998. Hence the 22 monthly price changes were reduced to one price change that spanned from before the crisis to the most recent data. This simplification is not without costs, for the reduced data set is no longer able to

[6] A positive correlation between unobserved quality and price might also bias our estimates toward zero.

address questions about the timing of price changes across provinces. It may have taken more time for price increases to have occurred in the more distant provinces, and this sort of information is no longer retrievable with the reduced data set. In the geographic dimension, we have aggregated to create price series for each of 27 provinces as explained above. In the product dimension, for some analysis we have collapsed the 700 or so products into approximately 180 products or aggregates that we are able to match with goods in the consumption data (SUSENAS).

The price data are reported in *levels*, but we focus our analysis in this paper on *changes.* There is little doubt that some places are more expensive to live in than others. Our interest, though, is whether the financial crisis had a differential impact on different regions of Indonesia. Hence, price changes seem the appropriate focus.

The notion that the overall impact of the financial crisis may have had geographically differential impacts finds some empirical support in ongoing work by Poppele, Pritchett and Sumarto at the World Bank in their working paper ‘‘Social Impacts of the Indonesian Crisis: New Data and Policy Implications (1999)". Relying on data sources different than those used in this paper, Poppele et. al. found that the geographic impact of the crisis on poverty was quite uneven. We return to these results in section 6 where we evaluate the impact of the crisis on the poor.

The geographic pattern of price increases differs according the the specificity of the products considered. At the most general level, the price index encompassing all goods does not show much regional variation. An unweighted average of the general price index for each province shows that prices increased an average of 92.5 percent from January 1997 to October 1998. The general price index on a province-by-province basis ranged from an increase of 70 percent in Nusa Tenggara Timor (NTT) to an increase of 119 percent in East Java. As a base line, the standard error of the series of provincial general price indices is about 11 percent. Figure 1 shows the empirical distribution of the provincial general price index increases. As noted above, it varies from between 70 percent and 120 percent and most provinces are in the 80 to 100 percent range. Given the different consumption patterns across provinces and the geographic separation of many provinces, this does not seem like very much heterogeneity. But this is deceiving.

There are 184 products and product aggregates that appear in both the SUSENAS and our price data. We have computed the change in the price index for all of them. The standard error of the change in the price index, as one looks across provinces, is greater than 11 percent (that of the general price index) in over 170 of them. There are some extreme examples, but one that is more representative and is especially important is the geographic dispersion of the price increase for rice. Figure 2 shows the empirical distribution of the percentage changes in the price of rice. It varies from around 110 percent (in South Sumatra) to around 280 percent (in South and Central Kalimantan.) The fact that the price increases of individual products

show much more geographic variation than that of the overall price index means that the price increases of products covary negatively across provinces. Loosely speaking, when the price of one product goes up a lot in a province, the price of another product goes up by less so that the increase in the general price index is not that different across provinces.[7]

The substantial geographic variation of price increases following the financial crisis has economic implications. Suppose the poor consume a similar basket of goods regardless of where they live. In such a case, the economic impact of the crisis on the poor may vary substantially across regions. For example, if the poor always devote a substantial share of their budget to rice, the poor would have been much more adversely impacted in South and Central Kalimantan than in Sumatra. Alternatively, if the poor consume very different baskets of goods in different regions, spatial price variation may in fact be coupled with a fairly uniform impact of the crisis on the poor. Further, if the poor are not evenly distributed across the provinces (and they are not), the geographic variation in prices has an additional impact which can serve to either mute or exacerbate the impact of the crisis on the poor.[8]

Product heterogeneity across product aggregates

The previous sub-section documented the spatial variation of prices. The general price index did not vary that much across provinces, but the prices of individual goods did. This finding has implications for product aggregation. If one wishes to estimate a demand system, some product aggregation is necessary. It is simply too hard to estimate a demand system for 184 (much less the original 700!) products complete with the all-important cross-price elasticities. One common practice is to aggregate products into groups such as ''food,'' ''housing,'' ''clothing,'' and the like. One can then estimate a demand system using the aggregated products. This is a relatively attractive option when the products that underlie the aggregate have price changes that are somewhat uniform. That is simply not the case in the Indonesian data. In this section, we document this finding and explain some economic implications of product heterogeneity across product aggregates. Like the spatial heterogeneity documented in the previous section, this type of heterogeneity also informs the methodology we use to investigate the impact of the crisis on the poor.

[7] An alternative explanation, which we have investigated and rejected, is somewhat more complicated. There are about 700 products that comprise the overall price index. Not all of these appear in the SUSENAS consumption data. It could have been the case that the products that contribute to the general price index but do not appear in our our consumption data contribute to the dampening of the variance of the general price index. This would happen if the excluded products had price increases that negatively covaried with the included products. We have gone back and investigated this possibility using all 700 prices and while there is some negative covariance between the price increases of included and excluded products, it is modest and does not explain the dampened variance of the province-level general price index.

[8] The spatial variation in price changes might in principal help to econometrically identify demand elasticities, but this would require concurrent (and unavailable) data on household expenditures.

The price data have seven aggregate commodities which in turn sum to the general price index. These aggregates are foodstuffs, prepared food, housing, clothing, health services, education and recreation, and transportation. Each of these is comprised of many individual products. The degree of disaggregation varies. There are 262 individual items under ''foodstuffs'' while there only about 40 or 50 for Health Services and for Transportation.

In order to abstract from heterogeneity across provinces and focus on the heterogeneity at the product level, we first collapse the data set and consider only the average price increase for each product when the average is taken across provinces. Hence, we compute the average increase in the price of the aggregate ''foodstuff'' as well as the average price increase in each of the 262 goods that comprise that aggregate. This removes the spatial dimension of the data. Figure 3 graphically illustrates the heterogeneity of the price increases of the products that comprise the aggregate for foodstuffs. In Figure 3, one notes that while one or two products had either price decreases greater than 50 percent or price increases greater than 400 percent, most products had price increases in the 0 to 200 percent range. The results of this aggregation across provinces for all product categories are provided in Table 2.

Table 2 lists, for each of the seven aggregates, the number of individual products, the average price increase when the average is taken across all the products that comprise the aggregate, the standard deviation of the price increases, as well as the minimum and maximum price increase. (One should keep in mind that these standard deviations do not account for the regional variation in price increases, only the variation of the average price increases.) For example, there are 262 products that comprise the aggregate ''Foodstuffs.'' Of these 262 products, one had an average price decrease of about 68 percent (a leafy vegetable that defies English translation) while one had an average increase of over 600 percent (red onions). Of all foodstuffs, the average price increase was 114 percent and the standard deviation of the price increase was about 80 percent. There was, in sum, tremendous variation in the average price changes of individual food items. This pattern holds for all of the aggregate commodities.

Once we have abstracted from spatial price variation, we have seen that how much prices increase depends on the degree of aggregation with which we define a product. This too has economic implications. Consider foodstuffs as an example. If poor households consume a different basket of specific food items than do the non-poor, the poor may be quite differentially impacted by the crisis. Perhaps the food items whose prices most skyrocketed were imported luxury items while the price rise for basic stables was more modest. Using an aggregate for foods will hide this important source of heterogeneity. This reasoning suggests that one should examine the impact of the crisis at the most disaggregated level. There is, though, a line of reasoning that works in the opposite direction. The methodology we use to investigate the impact of the crisis on the poor essentially assumes that there is no substitutability among goods (this is discussed

in some detail below). While the assumption of perfectly inelastic demands is clearly not correct, it is less incorrect as goods are more broadly defined. For these competing reasons, we analyze the impact of the crisis at different levels of product aggregation.

Heterogeneity across households

The above two sections have documented the heterogeneity of prices across provinces and within product aggregates. The purpose of this section is to illustrate the heterogeneity of households in the sample. One can either do this correctly, and write the ensuing book, or be too brief while giving a glimpse into relevant dimensions of household heterogeneity. Our choice will be obvious.

Table 3 quantifies how a handful of household characteristics vary across the population, both overall and by income groups. The first column gives the (weighted)[9] means of per-capita household income, expenditure, whether the head of the household had completed secondary school, the size of the household, the budget share of food in total expenditure, whether the household was rural, and the age of the household head. Means are reported for three separate deciles in the income distribution, as well as the overall sample. The sample is made up of the 58,100 households (from SUSENAS) included in the subsequent analysis.

Table 3 indicates that income is quite unequally distributed as the average income at the top decile is almost 12 times that of the bottom decile. Expenditure is less unequally distributed. Only about 25 percent of the very poor household heads have graduated secondary school while almost 75 percent of those in the top decile have done so. Richer households are smaller. (We have defined household size as the number of adults plus one half times the number of children.) About 90 percent of the households in the bottom decile are rural while about 70 percent of those in the top decile are urban. Households in the bottom decile devote about 85 percent of their income to food while those in the top devote only a bit more than one-third that share. As noted in Table 1, the CPI for food rose by more than the CPI for other categories, and this alone suggests that at this very aggregated level, the poor may have been more adversely impacted by the financial crisis.

As important as the averages across deciles reported in Table 3 are the standard errors of these averages. Even within households in the poorest decile, there is tremendous variation in the income share devoted to food consumption, the household size, the age of the head of the household, and whether the household head has completed secondary school. The very poor are themselves a quite heterogeneous group.

The poorest households don't just spend a larger share of their budget on food than middle and high income households but, as mentioned earlier, they also purchase a very different basket of products. Even

[9] When computing means, we use the sampling weights reported by SUSENAS.

within the category of food, poor households typically buy different items than do wealthier households. This is apparent in Table 4, which presents the mean expenditure shares for the overall sample as well as for those households in the top and bottom per-capita household income deciles. As expected, poor households spend a greater share of total expenditures on food than rich households (68 percent for those in the bottom decile compared with 47 percent in the top decile).[10] Even within food items, spending patterns vary by income level. The poor spend a far greater share on basic foodstuffs such as cereals and tubers (30 percent of all expenditures) than the wealthy (7 percent). Indeed expenditures on rice alone comprise a quarter of all expenditures for poor households compared with 6 percent for the wealthy. In contrast, the wealthy devote more than twice the expenditure share as the poor to meat, eggs and milk, and prepared food and beverages. Among non-food expenditures, the wealthy devote proportionately more resources to housing and education and are more reliant on electricity and gasoline (for transport), while the poor spend significantly higher proportions on kerosene and firewood. Because the prices of individual products do not all move together, the fact that richer and poorer households buy different products suggests that the financial crisis may have differentially impacted richer and poorer households in a complicated way. If one could simply multiply the poor's consumption basket by some scalar to get the rich's consumption basket, untangling the impact of the financial crisis on the poor would be simpler. But that is not the case.

6. Changes in the Cost-of-Living and the Impact of the Crisis on the Poor

The purpose of the previous section has been to establish that: i) Price changes varied a great deal across Indonesian provinces so that *where* a household lived may matter when evaluating the impact of the financial crisis; ii) Price changes varied a great deal depending on how one aggregates products so that the degree of disaggregation of product definition matters when evaluating the impact of the financial crisis; and iii) Households themselves are very heterogeneous so that a methodology investigating the impact of the financial crisis should accommodate this heterogeneity. With these concerns in mind, we now turn to measuring the impact of the crisis on the poor.

We measure the impact of the crisis on households (rich and poor) by computing household-level cost-of-living indices. Because we only have data on consumption patterns well before the crisis, we use these pre-crisis consumption baskets to compute what is essentially a Lespeyres cost-of-living index for each household. This index provides a *maximum* bound on the impact of the crisis, since the index does not take into account the substitution toward relatively less costly products that surely takes place (to some extent)

[10] Since we are now looking at food outlays as a share of total expenditures, and not income, the figures here will differ from those in Table 3.

after price increases. Denoting the price of good i faced by household j in time t by p_{ijt} and expenditure shares by q_{ijt}, the household cost-of-living index for household j is given by:

$$C_j = \frac{\sum_{i=1} p_{ij1} q_{ij0}}{\sum_{i=1} p_{ij0} q_{ij0}}$$

We compute 58,100 cost-of-living indices, or as many indices as there are households in our sample.

We actually compute three such household level indices. The first index that we compute matches the price changes of goods in the price data with the monthly expenditures of those same goods in the 1993 SUSENAS. For the monthly expenditure of food items, we simply convert the recorded weekly expenditures to monthly equivalents. For non-food items, we use the monthly average of annual expenditures, and not the expenditures in the month preceding the survey, in order to more accurately measure monthly expenditures for durables that are infrequently purchased. 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 132 different individual goods between the two data sets. In the case of non-food items, we matched 52 different goods, both individual goods such as firewood and kerosene, as well as aggregate goods such as toiletries or men's clothing. Hence, the i subscript in the Lespeyres formula above runs from 1 to 184. Through this matching, we were able to account for 75 percent of total household expenditures on average--a little greater for poor households and a little less for rich ones. This index is then an average of the observed price changes with each price change weighted by the household specific expenditure share for that good.

The second index is computed for the case in which we use 19 aggregate commodities instead of the original 184 which we matched between SUSENAS and the BPS price data. These aggregates include 15 food categories, such as cereals and meat, and 4 non-food categories such as housing and clothing. The motivation for this is twofold. First, recall the Lespeyres index, by construction, ignores substitutability across products. By defining products more broadly, as in the second index, we reduce the likely overstatement of the impact of the crisis. Put another way, when products are broadly defined, those aggregates are going to be less elastically demanded than the disaggregated products. The second motivation for this index stems from the fact that the disaggregated index only accounted for about 75 percent of households' expenditures. It is possible that for many households, the goods excluded in index 1 may either exacerbate or mitigate the measured welfare effects, depending on the relative price changes of those goods. The expenditures for these aggregates (i.e. meat, cereals, housing, etc.) are also supplied by the 1993 SUSENAS and the price changes for these aggregates are found in the price data. A benefit of this index is that it covers nearly 100 percent of the individual household's expenditures. Of course, by

attempting to compensate for the above potential biases, we may be introducing another bias-- aggregation bias-- which we have also previously discussed.

The third index that we compute accounts for the services provided by owner-occupied housing and for self-produced agriculture. Many households, especially in rural areas, own their 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 home. We account for the services provided by owner-occupied housing by treating the imputed rental value for these homes as a negative expenditure. Many households produce some of their own food. Over 90 percent of these households are classified as rural. Households that consume self-produced foodstuffs also tend to be net exporters of agricultural products.[11] As the price of food rose, the value of their production also increased. Clearly, if the household was 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. We modify the first index to account for self-produced agricultural products by treating the imputed value of self-produced food as a negative expenditure.[12] Note that this modified index 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. On the other hand, this index does not allow supply responses to the increased food prices and one would expect more self-production in the goods whose relative prices increased the most.

Non-Parametric Evidence

Figure 4 provides kernel density estimates of the first and third indices. The biggest difference between the two densities is the existence of households that are *better* off after the crisis due to the consumption of self-produced (and now more expensive) agricultural goods. While there are not very many of these households (less than ten percent of the sample), they are non-negligible. More generally, including household production muted the cost-of-living increases and this was especially true in the rural areas.

Table 5 summarizes how the cost-of-living index varied across per capita household income deciles. This table begins to address the impact of the crisis on the very poor.[13] Since the rapid price change resulting

[11] 54 percent (weighted) of households in the sample report operating an agricultural business. 69 percent (weighted) of those households report income from the sale of agricultural goods.

[12] 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 household shadow prices for agricultural inputs such as labor are not significantly different from market prices.

[13] We use income to measure ‘‘poor’’ and ‘‘non-poor’’ although we are well aware of the controversy around this definition. Our results were substantively the same when we looked at expenditure or food-related expenditure instead of income. See Chaudhuri and Ravallion (1994) for an investigation of the relative merits of these various poverty indicators.

from the crisis may have differentially impacted rural and urban areas, we report separate results for rural and urban households. When we do not take into account self-production and use disaggregated product definitions (Index 1), we find that the cost-of-living for the poorest urban households increased an average of 128 percent from the January 1997 through October 1998 period. The increase for urban households in the top income decile was 89 percent. Among rural households, where the overall cost-of-living increases were greater, the parallel figures are 136 percent for the poorest and 107 percent for the wealthiest households. For both rural and urban households, we find that the increase in the cost-of-living declines monotonically by income decile. Hence in both areas, the poorer a household was, the greater was the increase in its cost of living.

The middle columns of Table 5 (Index 2) use the more aggregated product definitions. The results are very similar to those using the disaggregated product definitions, although the differences across income deciles are mitigated, mostly because wealthier households in both areas now have somewhat higher cost-of-living increases. The overall similarities between the two indices are striking given that Index 2 employs broadly aggregate goods demanded less elastically than the individual goods of Index 1. Thus, while the Lespeyres index is, by construction, an overestimate of the true change in the cost-of-living, the bias may not be huge. Indeed, the averages of the the first and second indices are almost identical. It remains the case that the poorer households saw their cost-of-living increase by more than did wealthier households.

The final two columns of Table 5 now account for housing services and self-produced food. As suggested by the kernel density estimates in Figure 4, the increases in the cost-of-living are substantially muted relative to the figures for Index 1 in both areas, although this is especially true for rural households due to the tendency of rural households to engage in agricultural production. The average increase falls from 120 percent in index 1 to 84 percent. (In results not reported here, we find that it is indeed self-produced agriculture and not owner-occupied housing that accounts for most of the difference between the indices.) Interestingly, for rural households there is no longer a differential impact across the income spectrum. Indeed the cost-of-living increase is now slightly *greater* for the wealthier rural households suggesting that self-production has equalized the impact of the financial crisis across the income deciles in rural areas. The story is quite different for urban households, however, where there is still a clear monotonic decrease in the changes in cost-of-living. The cost-of-living for the poorest urban households increased 111 percent while the wealthiest households faced an increase of only 80 percent. It is important to note that the cost of living increases for the wealthy *urban* households are greater than those measured for the poor or median *rural* households. Thus this index suggests the impacts of the crisis have been greater for urban areas than rural, and greatest overall for the urban poor.

We turn now from variation across household income to the regional variation in the cost-of-living indices. Table 6 gives the change in the cost-of-living index for urban and rural households in each province. As an example, using Index 1 (disaggregated products and no correction for housing services or self-production), the cost of living for urban households in Aceh increased 102 percent while for rural households it increased 125 percent. For most provinces, Index 1 and Index 2 (with aggregated products) give similar results. Similar to Table 5, we find that when we do not account for housing services and self-production, rural households consistently faced greater increases in their cost-of-living than did their provincial urban counterparts. Our Index 1 results suggest that there was substantial regional variation in the the cost-of-living. In Irian Jaya, to the far east, the increase in urban cost-of-living measured 75 percent while in Southeast Sulawesi, the increase was 138 percent. The regional variation for rural households is equally dramatic. In Southeast Sulawesi, rural households faced a 161 percent cost-of-living increase while in Irian Jaya the increase was only 91 percent.

When we focus on an index that accounts for housing services and self-produced agriculture (Index 3), the regional variation remains, but the urban-rural comparison is changed. The depreciation of the Rupiah helped export-oriented provinces and the increased cost of food was offset by household production in rural areas. In every province, rural households faced a smaller increase in their cost-of-living index than did their urban counterparts. In some provinces, the differences are especially large. In East Timor and Irian Jaya, the cost-of-living for rural household only increased about 39 percent. The pattern of regional variation remains as other provinces had increases (for rural households) more than double that.

We view the particular results for some provinces with some caution as our results do not always coincide with the results reported by Poppele, Prichett and Sumarto (1999). For example, Poppele et. al. report that on Sulawesi, 70 percent of rural Kecamatans reported that things were better in August 1998 than they were a year prior. Although these measures are subjective, this figure is hard to reconcile with our cost-of-living changes. Overall, the Poppele et al. results suggest substantial regional variation (as do we) and that rural households fared relatively better (as do we once we account for household production.) These general findings are also echoed in Frankenberg, et al. (1999).

It might be interesting to investigate how the provincial and regional changes in the cost-of-living reported here vary with other provincial characteristics such as mean income or expenditure levels. Have wealthier or poorer regions of the country experienced higher cost-of-living increases? For rural areas, there is little correlation between mean per capita household income or expenditures and increases in mean cost-of-living. However for urban areas the cost-of-living changes are negatively and significantly correlated with provincial mean household income (and expenditures to a lesser extent). Thus urban areas with lower average household income experienced greater price changes than the more wealthy cities.

These findings hold true for any of the three indices. Another provincial characteristic more sensitive to the distribution of income within the province is the provincial population share categorized as poor. What might be the relation between regional variation in price changes and regional variation in poverty? We take as our poverty measure the population share deemed poor by Bidani and Ravallion (1993) from calculations based on the 1990 SUSENAS. Although these poverty indicators pertain to a period seven years before the currency crisis it is unlikely that the relative variation in regional poverty profiles would be much changed in the intervening years. We find no relation between the provincial cost-of-living increases and the provincial poverty indicators for urban areas. We also find no relation between cost-of-living changes and poverty indicators in rural areas when our cost-of-living is measured by Index 1 or Index 2.[14] However the cost-of-living changes as determined by Index 3 are negatively and significantly correlated with the share of rural provincial population deemed poor. Hence provinces with a greater proportion of poor experienced lower cost-of-living increases than the more well-off rural areas once adjustments for agricultural self-production were made.

Parametric Evidence

The results in Tables 5 and 6 suggest the crisis impacted the cost-of-living of the poor more than that of the rich, at least for indices 1 and 2, and urban households more than rural ones, after we account for owned housing and self-produced food. In order to investigate how the cost-of-living varies conditional on more than one household attribute, regression analysis is helpful. Our approach is simple and descriptive. It is without structural interpretation.

Regression results are summarized in Table 7.[15] The top half of Table 7 includes results using the cost-of-living index that does *not* account for housing services or self-produced food. We include three specifications and two estimation methods. All specifications are linear with the index being regressed on two to four explanatory variables. In each specification, we include the log of income and a dummy variable which takes a value of one if the household is rural. (We use the actual index instead of its log since the index when accounting for self-production may be negative.) In the most parsimonious specification, OLS yields a coefficient of -0.100 on log income and 0.163 on the rural dummy variable. Each are quite precisely estimated. The coefficient on log income has a natural interpretation. The negative sign on the coefficient indicates that the cost-of-living rises with declines in income. The poor are harmed most. A value of -0.100

[14] Poppele et al. (1999), using different poverty measures, also find no association across regions between the impacts of the crisis and pre-crisis levels of poverty.

[15] The regressions reported in Table 7 are not weighted by sampling weights, but we find that doing so makes little to no substantive difference in the results.

indicates that as income doubles (a 100 percent increase), the cost-of-living index falls by 10 points (0.10). The coefficient of -0.10 is large when considered in conjunction with the range of incomes. At the tenth percentile, household income is 75,802 while it is only 101,667 at the twentieth percentile. At the 80th and 90th percentile, income is 324,167 and 460,656. These large absolute differences translate into large differences in the cost-of-living. Since income is easily five times larger at the high end of the distribution than at the low end, the -0.100 coefficient corresponds to a cost-of-living index that is 50 points higher for poor households. This strikes us as a large disparity in the cost-of-living. The coefficient on ''rural'', still in the simplest OLS specification, is 0.163 indicating that rural households have a cost-of-living increase 16 points higher than their urban counterparts. Recall, though, that this result is for the index that does not account for self-production.

Our main focus is on how robust these results are to other specifications. Because we are not being guided by theory, the decision of which other regressors might be included in the regression is essentially arbitrary. Insofar as included regressors might co-vary with income or rural location, the coefficients on log income and the rural dummy might change. We include two additional regressors. One is the log of household size where children under 10 are counted as one half and adults as one. The other included variable (''Degree'') is a dummy variable set to one if the head of the household is a secondary school graduate. We find that the inclusion of household size increases (in absolute value) the coefficient on log income to -0.151. The coefficient on the rural dummy remains relatively stable. Household size itself conditionally covaries positively with the price index. Larger households face larger cost-of-living increases. Controlling also for the education of the head of the household has virtually no impact on the other coefficients.

There is good reason to believe that the residuals of the regression may be correlated by province. This would be consistent with the substantial regional variation that we found in earlier cuts of the data. Provinces seem to matter. For this reason, we estimated all specifications with a province fixed effects estimator. This effectively sweeps out any cross-province variation so the estimates instead capture only within province variation. We find that the inclusion of province fixed effects makes remarkably little difference to the estimates. All coefficients are about the same as with OLS. Put another way, the variation in the data that gave rise to the OLS estimates also exists at the province level.

In the bottom half of Table 7, we use the cost-of-living index that accounts for housing services and self-produced agriculture (Index 3). The results are broadly consistent with those in Tables 5 and 6. We focus first on the OLS estimates. In our most parsimonious specification, we find that the coefficient on log income is -0.046 while the coefficient on the rural dummy variable is -0.214. Hence if household income doubles, the cost of living decreases 4 points. As income at the tenth percentile is about one sixth that at the ninetieth percentile, these results indicate that the cost of living is about 25 points (0.25) higher for the very

poor. This is a large difference since the mean of the cost-of-living index is only 0.83. The OLS coefficient on log income is stable across specifications and is always precisely estimated. The coefficients on log income are always smaller with the index that includes self-production and this is consistent with the notion that self-production mutes the impact of the crisis on the poor. It remains the case, though, that the poor are more adversely impacted than the wealthy.

When we include housing services and self-production, the coefficient on the rural dummy variable becomes negative. Hence, when accounting for these influences, urban households faced a higher cost of living. The difference is on the order of 20 points which, again, is large given the mean of the index (.83). This coefficient is also stable across specifications. That the crisis impacted urban households more than rural ones is consistent with the preliminary results of Poppele et. al. (1999) as well as with Frankenberg et. al. (1999). We find that household size no longer seems to matter and that the coefficient on the education of the head of household becomes negative. The former effect is consistent with larger households having more housing services and self-produced agriculture. The later effect is consistent with higher education households engaging in less self-production. Finally, including provincial fixed effects mutes the impact of log income but it remains the case that the coefficient is precisely estimated and negative. Little else changes with the fixed effects.[16]

7. Conclusions and Caveats

Conclusions

The recent financial crisis in Indonesia has resulted in dramatic price increases. When we ask if these price increases have impacted the cost-of-living of poor households disproportionately hard, the answer is usually "Yes." Just how hard the poor have been hit, though, depends crucially on where the household lives, whether the household is in a rural or urban area, and just how the cost-of-living is computed. What is clear is that the notion that the very poor are so poor as to be insulated from international shocks is simply wrong. Rather, in the Indonesian case, the very poor appear the most vulnerable.

Our results emphasize the importance of heterogeneity when measuring the impact of the Indonesian economic crisis on households. We find that prices vary substantially across the disparate regions of Indonesia. Prices also vary across the types of goods considered. Households are also quite heterogeneous, even within income deciles, with respect to observable characteristics. On top of this variation, consumption

[16] We also experimented with an interaction term between log income and the rural dummy. In those (unreported) results, we find that $Priceindex = 2.22 - 0.10 * ln(y) - 1.46 * Rural + 0.10 * (ln(y) * Rural)$. All coefficients were precisely estimated. Hence, the negative relationship between income and the price index (Index 3) is present only for households in urban areas (as we might expect given the findings in Table 5).

patterns vary both regionally as well as by income class. For these reasons, we find it most helpful to think about *distributions* of responses and we have employed methods that, in most cases, do not rely on particular parametric assumptions.

By matching data on price changes with data on household consumption from a nationally representative Indonesian data source, we have calculated household specific cost-of-living increases. Since our measure-- a Lespeyres type index-- does not account for potential substitution among products, our figures provide an upward bound on the likely increase in the cost-of-living. We find a substantial increase in the mean cost-of-living, on the order of 130 percent, if we disregard the relative benefits of self-produced agriculture and owned housing. The measured increase is greater for poorer households and households in rural areas. There is a great deal of provincial variation in the measured cost-of-living increases although, as evidenced by the the fixed effects estimation results, there is as much variation within provinces as between them.

Our results also illustrate the role that agricultural self-production and owned housing played in dampening the impact of the crisis. When accounting for these benefits, the estimated mean cost-of-living falls to 84 percent and this cost is now lower for rural households. Of all households, the urban poor appear the most adversely impacted by the crisis. Their cost-of-living tended to rise the most and, being poor, these households are presumably among the least able to absorb these increases.

And Caveats

There are several reasons to view our results with caution. These include the absence of information about wages and incomes, potential problems with the price data that underlie our indices, the fact that we used 1993 consumption data to proxy 1997 consumption patterns, the biases inherent in a Lespeyres index approach, and the confounding influences of shocks other than that of the economic crisis. Each are discussed in turn.

Wages: This paper has analyzed variations in the changes in nominal prices during the Indonesian economic crisis. Of course we would also like to know what has happened with wages and income to better measure the real effects of the crisis. Unfortunately, our data contain no information on the changes in household income over the course of the crisis. However two alternative sources of data do have some information on wage changes. Data from the Badan Pusat Statistik (obtained from their web site at http://www.bps.go.id) reveal that nominal wages for many broad occupational classifications have increased throughout 1998. For example the reported increase in the mean wage from September 1997 through September 1998 for industrial workers stands at 26 percent. The median wage has also increased an almost identical 25 percent. Workers in the basic metal and metal working industries witnessed the highest wage increases, of about 40 percent, while wages in the paper and chemical industries increased less than

20 percent. There is also extensive regional variation in nominal wage increases. The largest wage gains reported were for workers in Sulawesi who experienced increases of 87 percent, while wages in Jakarta increased only 12 percent. It is apparent, however, that the nominal increase in wages was not nearly enough to offset the detrimental effects of the rapid price changes. Frankenberg et al. (1999) find significant erosion in the real wage, especially for workers in urban areas where the real wage has fallen 30 percent for men and 37 percent for women. The real wage has declined less in rural areas (18 percent for men and 19 percent for women) although overall wages are still significantly lower for rural workers.

While rising nominal wages will dampen the impact of rising prices, that only helps workers who actually earn the wages. Workers who instead become unemployed are hit doubly hard. BPS statistics indicate that unemployment rose from about four million workers in 1997 to over five million in 1998. On the other hand, the crisis has led to a slightly higher proportion of men, and a considerably higher proportion of women, currently working. The increased proportion working is largely due to unpaid family workers entering the labor force and somewhat mitigates the detrimental effects of the decline in real wages and rise in unemployment. This is apparent in the reported changes in household per capita expenditures, where the declines, although still significant, are not as large as the declines in wages. According to Frankenberg et al. (1999) mean per capita household expenditures have fallen 34 percent in urban areas and 13 percent in rural areas (although the median per capita expenditures have declined only 5 percent and 2 percent for urban and rural households respectively).

Inaccurate Price Data: Our price data comes from observations in urban areas. Due to the lack of information on rural prices, we extend these measured price changes to rural areas. However prices in rural areas, especially remote areas, may behave quite differently. Frankenberg et al. (1999) determine that overall inflation may be slightly higher (5 percent higher) in rural areas than urban. As well, at least for the 7 provinces for which they have some limited independent price data, Frankenberg et al. suggests that actual inflation may be as much as 15 percent higher than the BPS derived inflation estimates. This is another reason to view our results with caution.

1993 consumption data: We base our household expenditure shares on consumption data from the 1993 SUSENAS. As incomes rise, consumption patterns change. This is apparent if we review expenditure shares over the 1987 to 1993 period, where a smaller proportion of total expenditures is devoted to basic foodstuffs, such as rice, for all households throughout the period-specific income distribution. Up until the economic crisis this trend was likely to continue. To the extent that our consumption baskets are calculated with 1993 and not 1997 data, our measured impacts of the crisis will diverge from the actual impacts. However we are also concerned with the heterogeneous effects of the crisis among households, and as such,

the relative consumption baskets (among rich and poor households, or rural and urban households) are not likely to have changed as much as the absolute consumption baskets over the 1993-97 period.

The Lespeyres Index: We have examined the impact of the crisis with price data that both pre- and post-date the crisis, but we do not observe quantities corresponding the the higher prices. For households that do not engage in any self-production (which would include virtually all urban households), this means that our cost-of-living index is an upper bound on the true change in the cost of living. For households that do engage in agricultural self-production, the bias is lessened.

Not a controlled experiment: 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.

For all of these reasons (and surely more), one should view our results with some caution. On the other hand, the severity of the crisis and the sheer magnitude of the impacted population argue for presenting some evidence given the currently available data. That has been one aim of this paper. Another was to document the high degree of heterogeneity in the effects of the crisis across such dimensions as region, household location, and income.

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 Poverty Profile for Indonesia,'' *Bulletin of Indonesian Economic Studies*, **29(3)**, 37-68.

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 the Spatial Variation of Price,'' *American Economic Review*, **78(3)**, 418-430.

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 Labor and Population Program, Working Paper Series 99-04.

Poppele, Jessica, Sudarno Sumarto, and Lant Pritchett (1999) ''Social Impacts of the Indonesian Crisis: New Data and Policy Implications,'' A Background Note for the Consultative Group for Indonesia Update, January 1999, the World Bank.

TABLE 1
Some Background

	Rupiah Exchange Rates		CPI for			
	USD	100 Yen	Food	Housing	Clothing	Health
Dec. 1996	2,383.00	2,058.39	100.52	101.98	100.99	102.08
Jan. 1997	2,396.00	1,965.56	103.33	102.67	101.91	104.46
Feb.	2,406.00	2,000.63	105.99	102.90	102.43	105.32
Mar.	2,419.00	1,955.92	105.28	103.29	102.64	105.59
Apr.	2,433.00	1,921.19	105.24	103.99	102.62	107.56
May	2,440.00	2,095.15	105.30	104.82	102.71	107.69
Jun.	2,450.00	2,148.49	104.45	105.18	102.88	108.15
Jul.	2,599.00	2,210.83	105.93	105.82	102.80	108.41
Aug.	3,035.00	2,546.48	107.60	106.34	103.48	108.77
Sep.	3,275.00	2,715.56	109.59	107.58	104.56	109.21
Oct.	3,670.00	3,061.33	113.50	108.35	107.14	110.67
Nov.	3,648.00	2,867.48	117.25	106.82	107.01	112.27
Dec.	4,650.00	3,578.31	120.54	107.84	110.58	114.18
Jan. 1998	10,375.00	8,304.99	133.26	113.79	127.30	124.22
Feb.	8,750.00	6,895.21	157.79	123.28	145.14	148.98
Mar.	8,325.00	6,316.16	166.71	128.61	161.39	155.88
Apr.	7,970.00	6,034.46	176.56	131.56	168.39	164.12
May	10,525.00	7,580.14	183.42	136.99	176.01	168.06
Jun.	14,900.00	10,583.91	196.39	139.17	195.29	171.97
Jul.	13,000.00	9,048.21	220.27	146.93	219.23	186.41
Aug.	11,075.00	7,824.11	240.31	153.51	225.73	197.99
Sep.	10,700.00	7,921.25	261.00	155.92	225.22	204.49
Oct.	7,550.00	6,546.72	256.16	157.35	220.97	208.58
Nov.	7,300.00	5,903.77	255.70	158.11	215.99	210.71
Dec.	8,025.00	7,000.49	263.22	159.03	219.71	212.54
Jan. 1999	8,950.00	7,697.62	281.09	160.62	232.11	214.07

Source: Bank of Indonesia data available online at http://www.bi.go.id/ind/datastatistik/index.htm.

TABLE 2
Product Heterogeneity

Product Aggregate	Number of Individual Products	Average Price Increase	Std. Dev. of Price Increases	Minimum Price Increase	Maximum Price Increase
Foodstuffs	262	112.8%	80.5%	-68.3%	612%
Prepared Foods	72	78.4%	41.6%	0.04%	169%
Housing	105	107.7%	76.4%	0.4%	499%
Clothing	94	80.3%	46.4%	-0.04%	214%
Health Services	38	85.8%	51.2%	0.0%	263%
Transportation	48	77.3%	84.1%	-0.13%	482%
Education & Recreation	43	73.1%	71.5%	-9.70%	310%

Notes: Price increases are from January 1997 through October 1998. Average price increases are computed as the average across all provinces reporting price data for a given good.

TABLE 3
Household Heterogeneity

	Bottom Decile	Middle Decile	Top Decile	Overall
per-capita income	19241 (3916)	51959 (2411)	229097 (74424)	61596 (218335)
per-capita expenditure	21687 (11342)	46028 (10985)	136271 (91594)	49726 (41859)
Schooling	.2526 (.4345)	.5097 (.4999)	.7628 (.4253)	.4734 (.4993)
HH Size	4.3958 (1.6940)	3.7722 (1.650)	3.6142 (1.7225)	3.8911 (1.6911)
food share of income	.83483 (1.3818	.5569 (.1375)	.3233 (.1462)	.5824 (.5182)
rural	.9222 (.2678)	.6767 (.4677)	.3042 (.4601)	.6959 (.4600)
age of HH head	47.877 (13.714)	43.828 (13.98)	43.147 (13.467)	45.000 (13.910)
Number of HH's	5811	5811	5811	58,100

Notes: Deciles are by per-capita household income. The middle decile includes households with per-capita incomes between the 50th and 60th percentile. All means are weighted by population sampling weights. Household size is defined as number of adults plus one half the number of children under 10. Income and expenditure values are in current (1993) Rupiahs.

Source: 1993 SUSENAS.

TABLE 4
Expenditure Shares

Product	Bottom Decile	Mean	Top Decile
Food	68.1%	62.2%	46.9%
Cereals	27.6%	17.8%	6.9%
Rice	24.8%	16.7%	6.4%
Tubers	2.2%	1.1%	0.4%
Cassava	0.7%	0.4%	0.1%
Fish	4.6%	5.4%	4.4%
Meat	0.7%	2.2%	4.0%
Eggs and Milk	1.2%	2.5%	3.5%
Chicken Eggs	0.8%	1.3%	1.3%
Vegetables	7.3%	5.9%	3.7%
Legumes and Soy Products	2.8%	2.6%	1.6%
Fruit	1.9%	2.5%	3.0%
Oil and Animal Fat	3.8%	3.1%	1.8%
Beverages	4.2%	3.7%	2.4%
Sugar	2.7%	2.4%	1.5%
Seasonings	2.6%	2.3%	1.4%
Salt	0.3%	0.2%	0.1%
Ready Made Food and Beverages	4.2%	6.9%	8.7%
Tobacco and Beetle Leaf	4.5%	5.3%	3.9%
Filter clove cigarettes	1.1%	2.6%	2.7%
Non-Food	31.9%	37.8%	53.1%
Housing, Fuel, Lighting and Water	15.8%	17.5%	22.2%
Estimated monthly rent if owned	5.7%	7.5%	11.8%
Electricity	0.7%	1.4%	2.4%
Kerosene	2.7%	2.8%	1.9%
Firewood	5.3%	3.0%	0.5%
Health Care	0.9%	1.2%	1.7%
Education	1.4%	1.9%	3.2%
Gasoline (for transport)	0.0%	0.6%	2.1%
Clothing, Shoes, and Hats	6.4%	6.2%	5.5%
Durable Goods*	1.7%	2.7%	4.9%
Taxes and Insurance	0.6%	1.0%	2.1%

Notes: Source is 1993 SUSENAS. Durable Goods include items such as furniture, household utensils,

jewelry, and vehicles. Expenditure shares are given as a percentage of total household expenditures. Deciles are ranked by per-capita household income.

TABLE 5
Cost-of-Living Indices and Income Levels

Income Decile	Index 1		Index 2		Index 3	
	Urban	Rural	Urban	Rural	Urban	Rural
1	1.28	1.36	1.25	1.32	1.11	0.73
2	1.26	1.34	1.24	1.30	1.07	0.74
3	1.21	1.32	1.21	1.28	1.06	0.74
4	1.19	1.30	1.19	1.27	1.04	0.75
5	1.14	1.28	1.16	1.25	1.03	0.75
6	1.11	1.26	1.14	1.24	1.01	0.79
7	1.07	1.22	1.12	1.21	0.98	0.81
8	1.03	1.19	1.09	1.19	0.95	0.80
9	0.98	1.14	1.05	1.16	0.91	0.84
10	0.89	1.07	0.99	1.12	0.80	0.83
OVERALL AVERAGE	1.19		1.20		0.84	

Notes:

Price index 1 is computed across all disaggregated commodities and does not take into account either self-produced agriculture or owner-occupied housing. Price index 2 is computed across about 20 aggregated commodities and does not take into account either self-produced agriculture or owner-occupied housing. Price index 3 is computed across all disaggregated commodities and accounts for both owner-occupied housing and self-produced goods. Deciles are by per-capita household income.

TABLE 6
Regional Variation of Cost-of-Living Indices

Province		Index 1	Index 2	Index 3
Aceh	Urban	1.02	1.10	0.98
	Rural	1.25	1.25	0.63
North Sumatra	Urban	1.13	1.18	1.05
	Rural	1.39	1.36	0.87
West Sumatra	Urban	1.02	1.12	0.85
	Rural	1.26	1.26	0.64
Riau	Urban	0.94	1.04	0.90
	Rural	1.14	1.14	0.85
Jambi	Urban	1.00	0.99	0.95
	Rural	1.24	1.16	0.78
Bengkulu	Urban	1.18	1.29	1.10
	Rural	1.52	1.51	1.09
South Sumatra	Urban	1.06	1.16	0.96
	Rural	1.31	1.31	0.63
Lampung	Urban	0.98	1.11	0.85
	Rural	1.17	1.18	0.69
Jakarta	Urban	0.93	1.01	0.87
	Rural			
West Java	Urban	1.06	1.12	0.93
	Rural	1.21	1.23	0.67
Central Java	Urban	1.05	1.06	0.96
	Rural	1.18	1.16	0.78
Yogyakarta	Urban	1.15	1.14	0.92
	Rural	1.36	1.23	0.80
East Java	Urban	1.13	1.17	0.97
	Rural	1.33	1.29	0.98
Bali	Urban	1.08	1.10	0.97
	Rural	1.29	1.27	0.95
NTB	Urban	1.25	1.22	1.14
	Rural	1.50	1.41	1.01

Notes: Table is continued on the next page.

TABLE 6 (continued)
Regional Variation of Cost-of-Living Indices

Province		Index 1	Index 2	Index 3
NTT	Urban	0.84	0.85	0.74
	Rural	0.98	1.02	0.09
East Timor	Urban	0.92	1.06	0.83
	Rural	1.02	1.07	0.39
West Kalimantan	Urban	1.10	1.26	0.98
	Rural	1.49	1.53	0.76
Central Kalimantan	Urban	1.09	1.14	1.03
	Rural	1.43	1.37	0.71
South Kalimantan	Urban	1.13	1.12	1.07
	Rural	1.45	1.28	0.85
East Kalimantan	Urban	0.90	1.02	0.83
	Rural	1.14	1.18	0.80
North Sulawesi	Urban	1.05	1.04	0.98
	Rural	1.23	1.14	0.88
Central Sulawesi	Urban	1.21	1.23	1.10
	Rural	1.52	1.40	0.70
South Sulawesi	Urban	0.93	1.00	0.84
	Rural	1.09	1.13	0.42
Southeast Sulawesi	Urban	1.38	1.37	1.31
	Rural	1.61	1.54	0.98
Maluku	Urban	0.95	1.06	0.90
	Rural	1.19	1.22	0.61
Irian Jaya	Urban	0.75	0.81	0.67
	Rural	0.91	0.89	0.38

Notes:

Price index 1 is computed across all disaggregated commodities and does not take into account either self-produced agriculture or owner-occupied housing. Price index 2 is computed across about 20 aggregated commodities and does not take into account either self-produced agriculture or owner-occupied housing. Price index 3 is computed across all disaggregated commodities and accounts for both owner-occupied housing and self-produced goods.

TABLE 7 Cost-of-Living Regressions						
Dependant Variable: Index 1 (Index without Housing or Self-Production)						
Ind. Var.	OLS	Fixed Effects	OLS	Fixed Effects	OLS	Fixed Effects
ln(Income)	-.100	-.098	-.151	-.154	-.146	-.150
	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)
Rural	.163	.158	.141	.138	.138	.135
	(.002)	(.001)	(.002)	(.001)	(.002)	(.001)
ln(Size)			.143	.147	.142	.146
			(.002)	(.001)	(.002)	(.001)
Degree					-.023	-.019
					(.001)	(.001)
Dependant Variable: Index 3 (Index with Housing and Self-Production)						
ln(Income)	-.046	-.028	-.046	-.029	-.049	-.031
	(.003)	(.003)	(.003)	(.003)	(.003)	(.003)
Rural	-.214	-.201	-.214	-.201	-.212	-.199
	(.004)	(.004)	(.004)	(.004)	(.004)	(.004)
ln(Size)			-.000	.001	.000	.001
			(.005)	(.005)	(.005)	(.005)
Degree					.014	.012
					(.004)	(.004)

Notes: Regressions had approximately 58,000 observations.

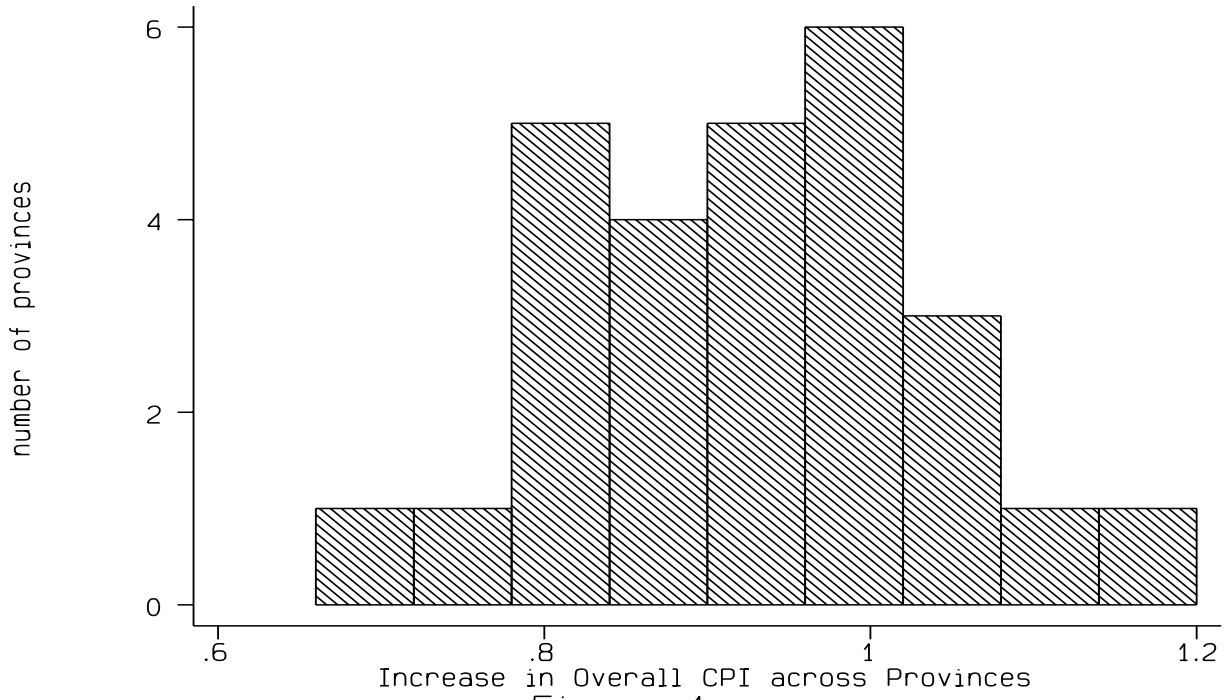

number of provinces
0
2
4
6
.6
.8
1
1.2
Increase in Overall CPI across Provinces

Figure 1

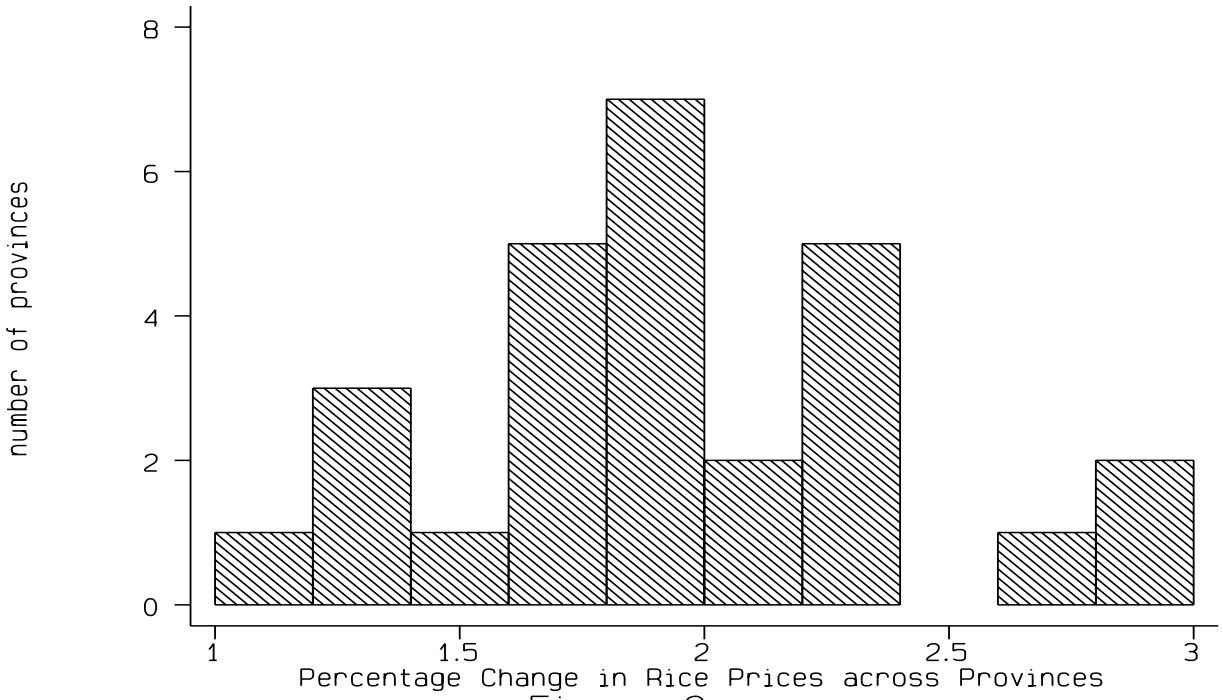

number of provinces
0
2
4
6
8
1
1.5
2
2.5
3
Percentage Change in Rice Prices across Provinces

Figure 2

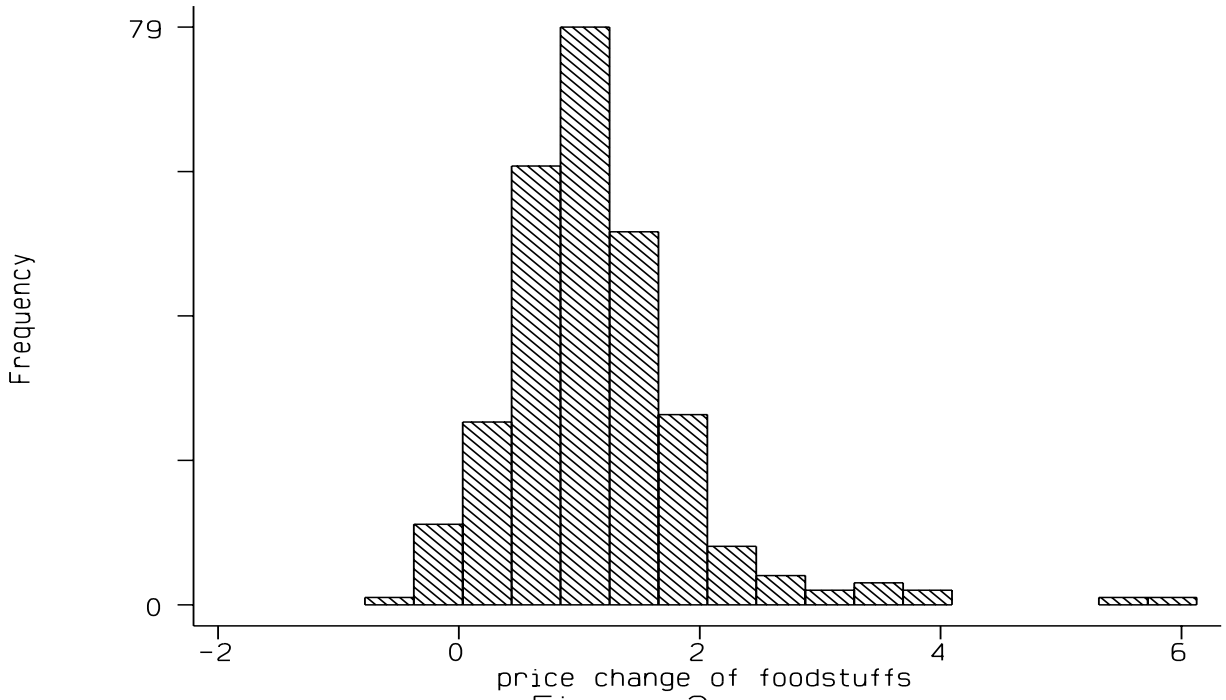


Figure 3

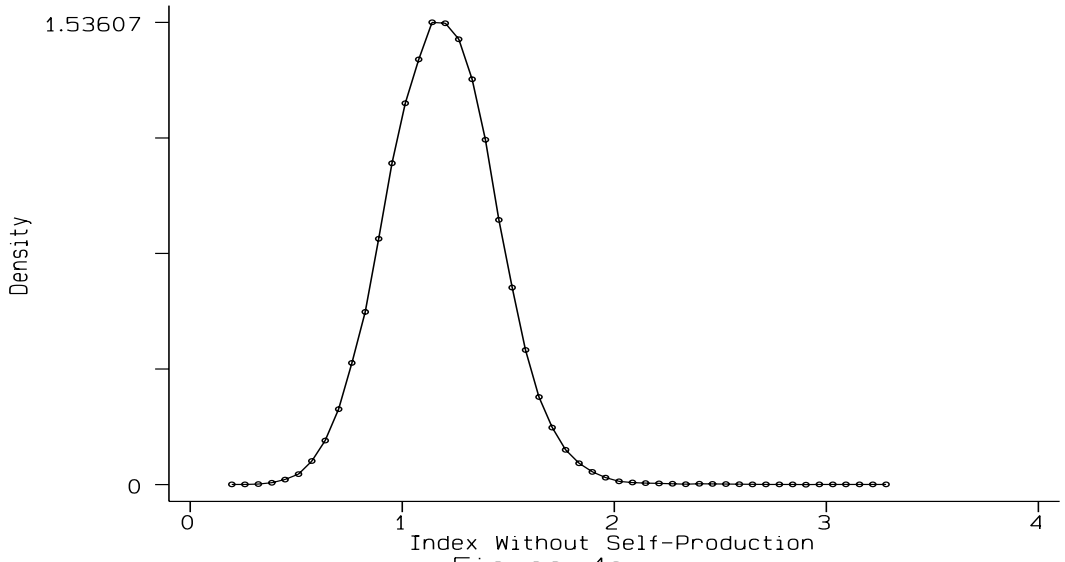


Figure 4a



Figure 4b