Distributional Analysis of U.S. Farm Household Income

Jeffrey Hopkins
USDA, Economic Research Service
1800 M Street, NW
Washington, D.C. 20036-5831
ph: 202-694-5584
fax: 202-694-5758
email:jhopkins@ers.usda.gov

Mitchell Morehart
USDA, Economic Research Service

July 18, 2000

Keywords: expenditures, farm safety net, household income, poverty, stochastic dominance, wealth

Selected Paper for the Annual Western Economic Association International Conference, Vancouver, B.C., June 29 - July 3, 2000. The views expressed are the authors' and do not reflect those of the U.S. Department of Agriculture.

1 Introduction

A shared goal of agricultural and rural development policy has been income maintenance. This goal has generally been taken to mean comparable incomes (in general or on average) between agricultural and non-agricultural households, although the definition has never been very precise. At the time the Federal Agricultural Improvement and Reform (FAIR) Act was enacted in 1996, the necessity of a program to address low incomes in agriculture was being seriously questioned, as median incomes for the farm population were similar to median incomes for the nonfarm population, while net worth for farm households far exceeded net worth for non-farm households (Tweeten, 1995). In 1998 and 1999, however, emergency assistance was enacted to supplement incomes and address economic hardship in farm households (Morehart and McElroy, 1999). An automated counter-cyclical policy, often called a safety net, has been a key component of recent policy initiatives designed to support farm household earnings (Harwood and Jagger, 1999).

Martin Feldstein recently argued that poverty is the primary issue of income distribution analysis, not changes in the incomes of the upper end of the distribution (Feldstein, 1998). We provide a distributional assessment of farm household welfare in 1998, with special attention to the lower end of the distribution. Our analysis will provide information that could be used to further define goals with respect to maintaining farm household welfare. Methods and data used in this study offer three principal advances over previous studies (Ahearn, Johnson, and Strickland, 1985; Allanson and Hubbard, 1999; Smith, Richardson, Anderson, and Knutson, 1999; El-Osta, G.A. Bernat, and Ahearn, 1995) that assess of farm household welfare. First, we use a nationally-representative sample to provide complete coverage of farm types and regions of the country, preserving the heterogeneity masked by

representative farm analysis. Second, our data provide the most complete coverage of all aspects of economic activity at the household level, including off-farm sources of income and wealth. Third, we examine all parts of the income distribution, particularly the lower tail, which will likely be the focus of targeted assistance, such as through a safety net.

A distributional analysis is necessary if we are to learn ways in which counter-cyclical mechanisms can be targeted towards those most affected by low returns. Even if post-FAIR farm policy resembles supply control more than it resembles a safety net, distributional issues will likely remain a part of overall rural development policy that addresses poverty and inequality.

2 Poverty and Inequality among Farm Households

Low returns in agriculture have traditionally used either the farm household or the farm business as the unit of analysis, but welfare implications extend from the farm household, not the farm business. Research on farm households tracks a rich vein of economic literature over that past 30 years on households and income poverty. This was a distinct departure from earlier welfare work by Pigou and Dalton on the inequality of incomes. The more modern literature addressed two distinct problems in the measurement of poverty: (1) setting a threshold that identifies the poor among the population, and (2) constructing an index of poverty from what we know about the poor. We take up the first problem in the empirical section of the paper. For the second problem, we implement a measure from the work of Sen (1976), Thon (1979), and Shorrocks (1995) that improves upon the traditional poverty measures that count the number of poor or estimate their poverty gap.

Poverty has traditionally been measured by the headcount. If we let z be the poverty

line for incomes (or some other measure of well-being), then the headcount ratio is the proportion of individuals or households that have income below z. The strength of the headcount is its simplicity, but it provides no information on the severity of poverty in the population. A second measure, the poverty gap, provides this information. The poverty gap is defined for a household with income x and poverty line z as $max(z - y, 0)$. In most empirical work, the poverty gap is normalized by z, resulting in a poverty gap ratio that takes on a value from zero (indicating no poverty) to one (indicating maximum deprivation). Poverty gaps can be averaged across a population, but then they reveal nothing about the inequality of the poor.

Sen critiqued the headcount and poverty gap measures for being blind to distribution among the poor, and proposed an index based on the headcount, poverty gap, and Gini index. The Gini measured inequality in the weighted poverty gap data, with the weights derived from each person's or household's rank in the ordering of the poor. Now known as the S index, it only satisfied weak versions of the monotonicity and transfer axioms. A modified-Sen index, known as the SST to signify the influence of Shorrocks (1995) and Thon (1979), satisfied the strong version of the transfer axiom (Shorrocks, 1995). Both the S and the SST index, are calculated as $(Headcount) * (Poverty\,Gap\,Ratio) * (1 + Gini(Poverty\,Gap\,Ratio)^{\gamma-1}$. When γ equals 2, the Gini formulation used here is that of the standard Gini coefficient where each observation, based on a social welfare function, is equally weighted. Large values of γ impose more weight on observations at the lower end of the distribution.

The difference between the S and SST poverty indices is that the S index measures the average poverty gap ratio and the Gini of the poverty gap ratios only among the poor, while the SST index considers the entire population when calculating the same measures (Xu and Osberg, 1999). The S and SST indices have both been widely embraced by researchers

because of their capacity to link absolute deprivation (through the headcount and poverty gap) with relative deprivation (through the Gini). However, because there are a multitude of indices (see Jenkins and Lambert (1997) for a recent accounting of different measures) from which to choose, measurement using one index may contradict measurement using another. The breadth of analysis offered by the SST index was broadened considerably by Shorrocks (1995) who expressed the indices as weighted areas beneath a Cumulative Poverty Gap (CPG) curve. The advantage brought by CPG curve analysis is the ability to compare the distributions themselves rather than comparing indices created to summarize the distribution.

3 Stochastic Dominance

From Sen onwards the criterion of dominance, which comes from stochastic dominance theory, has been widely viewed as a way around the problems associated with using index numbers to measure poverty. A welfare distribution is said to (stochastically) dominate another if and only if welfare, viewed through a welfare function, is greater. Likewise, in poverty dominance, a distribution is said to dominate another only if poverty, viewed through a poverty function, is less.

Consider two cumulative distributions of incomes F_A and F_B and let D_A^1 be such that $D_A^1(x) = F_A(x)$. More generally, let

$$D_(^s A)(x) = \int_0^x D_A^{(s-1)}(y)\, dy, \tag{1}$$

for any integer $s > 2$ and let D_B^s be defined analogously. It is useful to see that we can

also express $D^s(x)$ for any order s as (Davidson and Duclos, 1998):

$$D^s(x) = \frac{1}{(s-1)!} \int_0^x (x-y)^{s-1}\, dF(y). \tag{2}$$

First-order stochastic dominance of A by B up to poverty line z implies that $D_A^1(x) > D_B^1(x)$ for any poverty line $x \leq z$. This result is analogous to the situation where the headcount of individuals or households below the poverty line is always greater in A than in B for any poverty line below z. All stochastic dominance relations are nested, so that first order stochastic dominance assures second and higher order dominance as well. Second order stochastic dominance of A by B up to poverty line z implies that $D_A^2(x) > D_B^2(x)$ or that

$$\int_0^x (x-y)\, dF_A(y) \geq \int_0^x (x-y)\, dF_B(y), \tag{3}$$

for all $x \leq z$.

In words, equation (3) says that the poverty gap in A is greater than that in B for all poverty lines x less than or equal to z. Third order stochastic dominance is decided on the square of poverty gaps for all poverty lines x less than or equal to z. Higher order stochastic dominance is checked with larger powers of poverty gaps, (e.g. fourth order stochastic dominance requires testing for differences in poverty gaps raised to the third power).

Stochastic dominance addresses shortcomings of poverty index measures outlined above. First, stochastic dominance can be specified to be robust to wide classes of poverty indices upon which social welfare functions might be built. When first order stochastic dominance

holds, all persons with positive marginal utility prefer B to A. Second order stochastic dominance applies when all persons with positive and decreasing marginal utility share the same preference ranking. Third order stochastic dominance applies when all persons with positive, decreasing, and convex marginal utilities share the same preference ranking. Second, because the unit of analysis is the distribution itself, rather than a summary of a distribution, it is possible to differentiate two distributions that are ranked equally using an index. This is often the case when two distributions cross, and higher-ordered dominance is tested.

Figure 1 shows a pair of Cumulative Poverty Gap (CPG) curves, corresponding to 1995 and 1998 farm household income data. The curves were constructed by first ranking households by their incomes, then cumulating the probability-weighted poverty gaps, averaged across the entire population. The plot of these values at every p results in the CPG curve shown. The curve is drawn for a poverty line set at \$13,003, the official poverty line in 1998 for a household with three members. We first interpret the economic content using stochastic dominance, then switch to the poverty index approach.

The curves are drawn assuming that a common poverty line is used for both years, with incomes adjusted between the years according to the Consumer Price Index for Urban Consumers (CPI-U). Note that the poverty gap in 1995 is the same or greater than in 1998 over most of the population.

Because $CPG_{1995} \geq CPG_{1998}$ for all $p\epsilon[0,1]$, we say that CPG_{1998} second order dominates CPG_{1995}. This implies that for any social welfare function with positive and decreasing marginal utility, the 1995 distribution of incomes is inferior to the 1998 distribution of incomes. Because we have second order stochastic dominance, we have higher order

dominance as well.

Recall that CPG curves also have a poverty index interpretation. Figure 2 demonstrates how to extract information from the CPG curve for the SST poverty index. Figure 2 shows a normalized CPG (poverty gaps divided by the poverty line) rather than an absolute CPG, and both axes range from 0 to 1. The poverty *headcount* is reflected by the horizontal segment O, h. The poverty gap ratio is measured by the vertical segment $0, hi$. The Gini ratio is the distance from the diagonal line of perfect equality of poverty among the poor O, p, calculated by the ratio of area D to area C. The Gini for the poverty gaps of the entire population is adjusted by the area of the people who are not poor, equal to $1 - h$. Graphically, the SST poverty index is equivalent to the ratio of the sum of areas B, C, and D to area A. Moreover, the S index measure can be found by taking the ratio of the sum of areas C and D to area E (Osberg and Xu, 1999).

Referring back to the 1998 curve in Figure 1 we see that the headcount is 20 percent of the population, the average poverty gap for the entire population is \$1,824. The average poverty gap ratios, i.e. divided by z, is 0.14 (\$1,824/\$13,003). The Gini of poverty gap ratios for the entire population under the SST index is 0.86. The SST index of poverty for farm households, when the poverty line is equal to \$13,003, is equal to 0.052. For 1995, the headcount, poverty gap, and Gini measures were 0.22, 0.13, and 0.85, respectively, and the SST index for 1995 farm households was somewhat higher than in 1998, equal to 0.053. Our 1995 estimate for the farm population is smaller than an estimate for the entire U.S. in 1994 by Osberg and Xu, who report an SST index of 0.125. The two studies are not directly comparable, however, because we use a household rather than per-capita measure of poverty.

4 Household Poverty Indicators

Economic well-being, defined in terms of a household's command over goods and services, cannot be observed directly, proxies are used in practice. Common proxies include annual income, annual expenditures, and net worth. Note that the use of multiple proxies would have limited benefits for understanding household well-being if households held the same relative position in each distribution.

Income is the most popular measure of economic well-being in inequality studies. Income reflects one's potential control over economic resources, and can be used as an indicator of ability to sustain a level of consumption and therefore a standard of living. Expenditures, as an alternative to income, more directly reflect a household's consumption of goods and services. Any problems with the measure stem from the fact that it is consumption, not expenditure, levels that determine whether a household is in poverty. Net worth is the final alternative considered.

While the use of any single proxy will have advantages and disadvantages, data collected by USDA is inclusive enough to permit construction of all three proxies for well-being. See Table 1 for further detail on how the three measures of household welfare were constructed. Data used in the analysis come from the 1998 Agricultural Resource Management Survey (ARMS), an annual survey administered and maintained by the National Agricultural Statistics Service (NASS) and the Economic Research Service (ERS). The ARMS survey contains over 10,000 observations, stratified into 13 sales classes for each of the 48 contiguous states. The ARMS survey is also multi-phase, requiring the use of a complex weighting strategy in order to aggregate at the state, regional, or national level. Responses in ARMS are expanded according to the probability of being selected, so that each response represents

the respondent and other farm households that are like it.

Table 1. Computation of Household Welfare Measure with Poverty Line, 1998

Welfare Measure	Computation	Poverty Line
1. Income	Farm income (farm business income, operator and other household member on-farm wages, farm rental income, income from other farm businesses, commodities paid to household measures) + off-farm income (business income, wages and salaries, interest and dividends, Social Security and public assistance, other passive sources of income)	13,003
2. Expenses	Food and household supplies, household rent, mortgage, non-farm transportation, medical expenses, insurance, and retirement, other	10,750
3. Net worth	household share of farm assets and debt + non-farm assets and debt	66,271

Poverty thresholds were constructed for each of the three measures of well-being. The Bureau of the Census set the official poverty threshold in 1998 at $13,003 for a household with three members. Because the 1998 ARMS data does not include data on household size the average size reported in the 1995 ARMS was used. For the expenditure variable, a poverty threshold was set equal to half median household expenditures in the 1998 ARMS, approximately $10,750 on average, but the actual threshold (see Table 2, Panel B) was adjusted to reflect regional variation in prices. As well, a relative poverty threshold was used for the net worth measure using Survey of Consumer Finance data for 1995. Half the median net worth of households that operated sole proprietorships, an amount equal to $66,271 was used for the net worth poverty threshold.

While estimation of means and ratios using sample weights from the ARMS survey is straightforward, because of the complex sample design the classical asymptotic variance calculations can no longer be used, so resampling using replicates is used instead. For instance, the variance of the mean of a distribution is calculated using the mean from each replicate. Care is required in constructing the replicates, so that their design mimics the

design of the whole sample. A jackknife procedure is used to estimate a nearly unbiased estimate of mean squared error (Kott, 1999; Dubman, 2000). For an observation on, say, income x for household i with sample weight w_i, the estimated mean income for all farm households can be calculated as $\sum_{i=1}^{n} w_i x_i / \sum_{i=1}^{n} w_i$.

5 Results

CPG curves, as described above, are principally used for stochastic dominance analysis, although they also provide visual assessment of the headcount, poverty gap, and inequality of poverty that are frequently targeted by poverty alleviation programs. Figure 3 shows regional CPG curves for household income with the legend at the right ordering the regions by decreasing plateau values. Geographic region definitions are showed in Figure 4. The horizontal axis shows the percentile of the population, and the vertical axis measures the poverty gap in dollar units per household.

Only one consistent ordering is apparent in Figure 4, as the Northern Great Plains region shows a greater poverty gap at every p than every other CPG curve, implying that the region is pairwise dominated by every other region. Note that the Mississippi Portal region has a large aggregate poverty gap (indicated by the second-highest plateau. The frequent crossings from below indicate, however, that it's poverty gap is not greater at every point in the distribution, and therefore we cannot rank order Mississippi Portal poverty with other regions without further assumptions on the shape of marginal social welfare.

Table 2 explicitly show headcount, poverty gap, and inequality measures implicit in Figure 4. The first column shows the poverty threshold, the second column the poverty headcount, given the threshold. The next two columns give the poverty gap in dollars as

well as a ratio where the average poverty gap is divided by the poverty threshold. The inequality of poverty is shown in the next column, calculated as the S-Gini ratio (SST version) of the poverty gap ratios. Finally, the SST index is shown in the final column.

When poverty in household income is measured, a great deal of variation is observed among the regions in the headcount and poverty gap measures, with differences in the inequality of poverty being more muted. The Northern Great Plains region is seen in the table to have the highest headcount, at 30 percent of the population, and the Prairie Gateway has the lowest, as 14 percent of the households had incomes below $13,003. Poverty gaps ranged from a high of $2,848 in the Northern Great Plains region to $1,021 in the Basin and Range region. Inequality measures in general were highest where incidence and intensity were at their lowest, in the Prairie Gateway and Basin and Range regions. Similarly inequality in poverty was lowest in the Northern Great Plains, where incidence and intensity were the highest.

Up to here, analysis follows similar studies which do not assess sampling error, and reports distributions as though the distribution were known with certainty. One study that does consider the sampling error associated with surveys is Osberg and Xu (1999) who use the bootstrap method. In our study, a resampling method, the jackknife, was used to assess the sensitivity of the empirical distribution to the individual observations in the data.

Table 3 show the results of hypothesis tests of the differences in means among regions. Using the pooled-variance t-test with 14 degrees of freedom results in only 7 significant differences in poverty in household income (see Table 3) at the 95 percent confidence interval, showing that income poverty in the Northern Great Plains is greater than seven other regions (all but the Mississippi Portal). Northern Great Plains income poverty is greater

than the Mississippi Portal at the 90 percent confidence interval. No additional statistically different measures of poverty were found using either household expenditures or net worth as alternative poverty criterion.

Throughout this analysis, we have attempted to look at differences in poverty among regions. Other than the case where differences in household incomes are inferred for Northern Great Plains, overall poverty is more or less similarly distributed across regions. Rather than regional distributions, poverty might be better explained at the household level. Statistically significant differences were determined using a pooled-variance t-test, with jackknife variance estimates. Values annotated with a single star in the poor column are different from the value in the corresponding non-poor column at the 95 percent confidence level.

The first two rows of Table 4 show that none of the poverty measures results in a pool of poor making up more than a quarter of all households. The greatest pool size results from using the income measure, with 18 percent of farm households, followed by the expenditure measure with 15 percent, and the net worth poverty measure with 9 percent of farm households. Finally, the group *all* has only 17,990 members, roughly one percent of all farm households.

The second section of Table 4 presents average income, expenditure, and net worth for all groups, allowing us to show that targeting a single facet of economic power has often unexpected effects on other measures. For example, the group targeted for low incomes (column two in Table 4) is shown to also have lower expenditures but similar net worth to the non-poor, indicating that low incomes in 1998 affected households regardless of their wealth level. When expenditures, net worth, or all three measures are used, economic power of the poor group is uniformly lower than the economic power of the non-poor group,

regardless of which measure of power is used to classify the poor. Of course, the *all* measure targeting the three measures jointly maximizes the difference between the poor and non-poor groups.

6 Summary and Conclusions

We draw three main conclusions from this study, relating to the idea of how to best include a targeted safety net within a broad farm policy. First, even with the ARMS database, the most extensive database available on U.S. farm households, it was not possible to rank order most regions, even under the special case where we assumed away sampling error. Second, targeting by region or by crop may not be advisable. At least in 1998, poverty occurred to varying degrees in every region of the country and under every crop grown. Third, while many traditional government programs in place for farm households have either not been targeted at all or have been targeted by production levels, using economic criteria to target will have important consequences for what segment of the farm household population is reached. Evidence presented here indicates that farm households defined as poor under an income definition on average possess wealth levels that could allow economic smoothing through asset depletion or loans. An expenditure measure of poverty appears more able to exclude farm households that are able to smooth acute welfare losses. On the other hand, a net worth poverty definition targets younger farmers and the elderly, who did not successfully accumulate significant assets within their lifetimes. Using all measures simultaneously, on the other hand, results in a very small, but needy, group of households.

References

Mary Ahearn, Jim Johnson, and Roger Strickland. The distribution of income and wealth of farm operator households. *American Journal of Agricultural Economics*, 67(5):1087–1099, December 1985.

Paul Allanson and Lionel Hubbard. On the comparative evaluation of agricultural income distributions in the european union. *European Review of Agricultural Economics*, 26(1): 1–17, 1999.

Russel Davidson and Jean-Yves Duclos. Statistical inference for stochastic dominance and for the measurement of poverty and inequality. Working Paper 9805, Department of Economics, Queen's University, Kingston, Ontario, Canada, February 1998.

Robert W. Dubman. Variance estimation with usda's farm costs and returns surveys and agricultural resource management surveys. ERS Staff Paper AGES 00-01, Resource Economics Division, Economic Research Service, U.S. Department of Agriculture, 1800 M Street NW, Washington, DC 20036, April 2000.

Hisham El-Osta, Jr. G.A. Bernat, and Mary C. Ahearn. Regional differences in the contribution of off-farm work to income inequality. *Journal of Agricultural and Resource Economics*, 24:1–14, April 1995.

Martin Feldstein. Income inequality and poverty. NBER Working Paper 6770, 1998.

Joy Harwood and Craig Jagger. Agriculture's safety net: Looking back to look ahead. *Choices*, 14(4):54–60, 4th Quarter 1999.

Stephen P. Jenkins and Peter J. Lambert. Three $i's$ of poverty curves, with an analysis of uk poverty trends. *Oxford Economic Papers*, 49:317–327, 1997.

Phillip S. Kott. Some problems and solutions with a delete-a-group jackknife. National Agricultural Statistics Service, August 1999.

Mitchell Morehart and Robert McElroy. Net farm income to decline but remain near 1990-98 average. *Agricultural Outlook*, (AGO-258):11–14, January-February 1999.

Lars Osberg and Kuan Xu. International comparisons of poverty intensity: Index decomposition and bootstrap inference. *Journal of Human Resources*, 34(4), 1999.

Amartya Sen. Poverty: An ordinal approach to measurement. *Econometrica*, 44(2):219–231, March 1976.

Anthony F. Shorrocks. Revisiting the sen poverty index. *Econometrica*, 63(5):1225–1230, September 1995.

Edward G. Smith, James W. Richardson, David P. Anderson, and Ronald D. Knutson. Farm level comparison of the fair act to the 1990 farm bill. *AFPC Policy Working Paper 99-5*, (5):36, June 1999.

D. Thon. On measuring poverty. *Review of Income and Wealth*, 25:429–439, 1979.

Luther G. Tweeten. The twelve best reasons for commodity programs: Why none stands scrutiny. *Choices*, 10(2):4–7, April-June 1995.

Kuan Xu and Lars Osberg. An anatomy of the sen and sen-shorrocks-thon indices: Mulitplicative decomposability and its subgroup decomposition. June 1999.

Table 2. Components and overall index of poverty based on farm household income, 1998

	POVERTY LINE	HEADCOUNT	POVERTY GAP	INEQUALITY	INDEX
Region	dollars	rate	dollars	ratio	S-Gini	
Heartland	13,003	0.16	1,440	0.11	0.88	0.03
Northern Crescent	13,003	0.18	1,519	0.12	0.87	0.04
Northern Great Plains	13,003	0.30	2,848	0.22	0.78	0.12
Prairie Gateway	13,003	0.14	1,171	0.09	0.90	0.02
Eastern Uplands	13,003	0.20	1,368	0.11	0.87	0.04
Southern Seaboard	13,003	0.17	1,342	0.10	0.88	0.03
Fruitful Rim	13,003	0.17	1,488	0.11	0.88	0.04
Basin and Range	13,003	0.15	1,021	0.08	0.91	0.02
Mississippi Portal	13,003	0.24	1,625	0.12	0.86	0.06

Source: Agricultural Resource Management Study, USDA.

Table 3. T-statistics from inter-regional comparisons of SST Poverty Index, based on farm household income, 1998

	Heartland	N. Crescent	N. Great Plains	Prairie Gateway	E. Uplands	S.Seaboard	Fruitful Rim	Basin Range	Miss. Portal
Heartland		-0.73	-2.84	0.87	-0.53	-0.01	-0.28	0.80	-1.27
N. Crescent			-2.57	1.54	0.03	0.72	0.30	1.39	-0.85
N. Great Plains				3.13	2.49	2.84	2.62	3.09	1.80
Prairie Gateway					-1.17	-0.87	-0.96	0.07	-1.73
E. Uplands						0.52	0.22	1.12	-0.80
S.Seaboard							-0.28	0.81	-1.26
Fruitful Rim								0.92	-0.99
Basin Range									-1.68
Miss. Portal									

Table 4, Demographic Profile of non-poor and poor farm households under alternative economic measures, 1998

	Income		Expenditure		Net Worth		All	
	non-poor	poor	non-poor	poor	non-poor	poor	non-poor	poor
Absolute and Relative Pool Size								
Number of farms	1,663,521	358,891	1,721,708	300,705	1,842,536	179,877	2,004,422	17,990
Percent of farms	82	18	85	15	91	9	99	1
Economic power over goods and services					*dollars per household*			
Household Income	74,585	4,844*	68,468	26,373*	65,169	31,882*	62,725	4,655*
Household Expenditures	29,761	19,733*	31,579	7,381*	28,548	22,180*	28,178	6,115*
Household Net Worth	490,373	508,994	523,056	325,469*	538,685	32,651*	497,824	31,658*

* indicates differences between poor and non-poor 95 percent confidence intervals

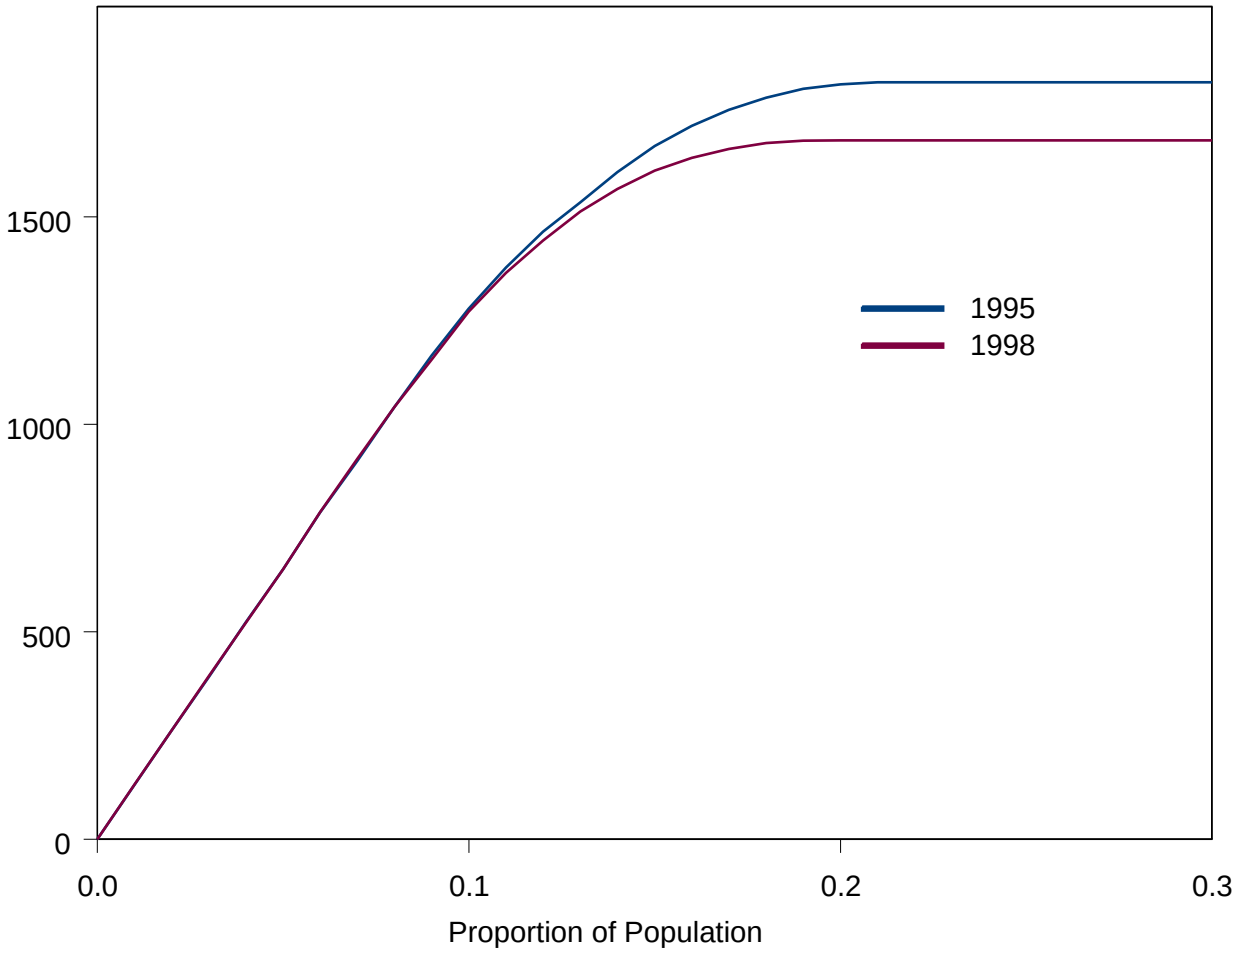

Figure 1: Cumulative Poverty Gap curves, U.S. Farm Household Income, 1995 and 1998

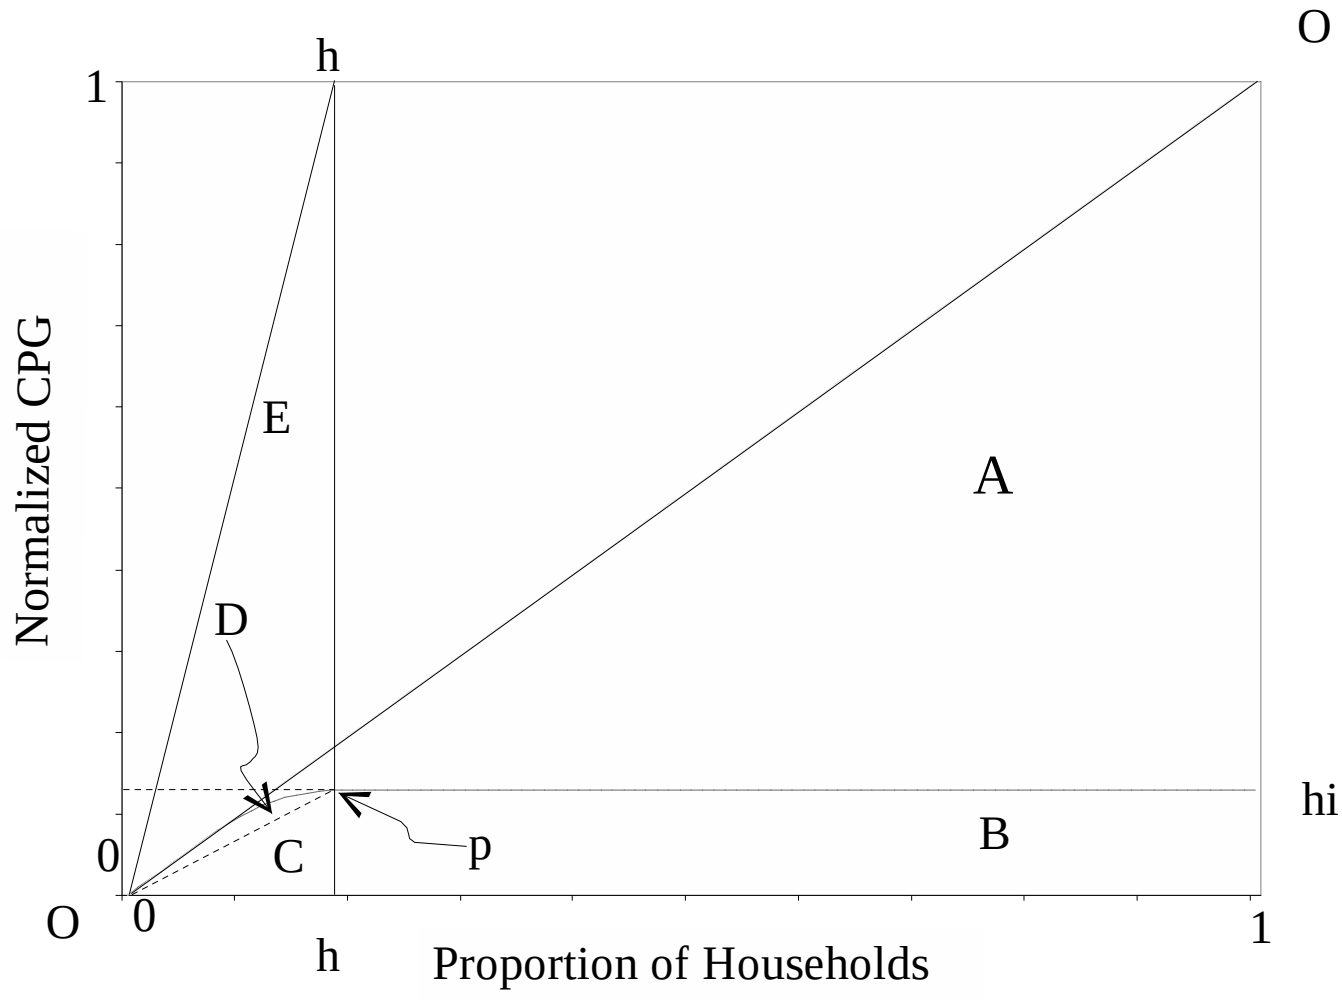

Figure 2: Cumulative Poverty Gap Curve

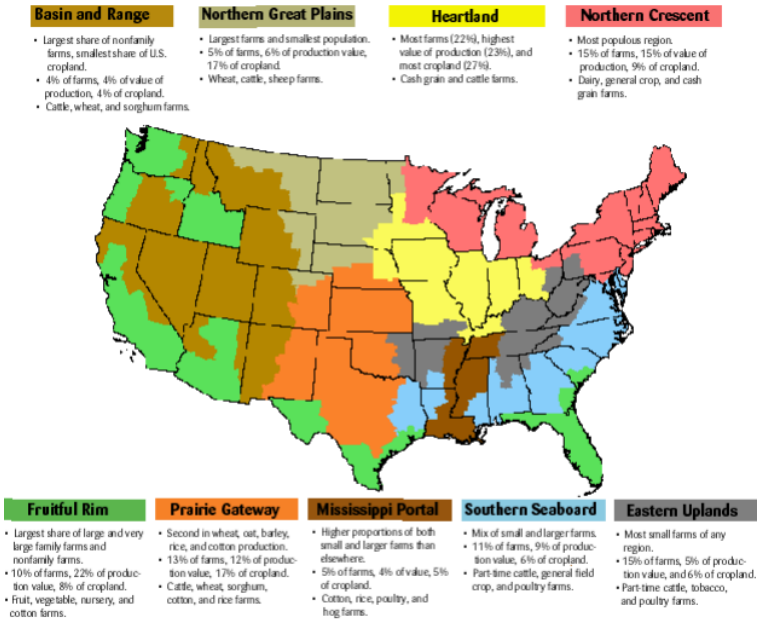

Figure 3: ERS Resource Regions

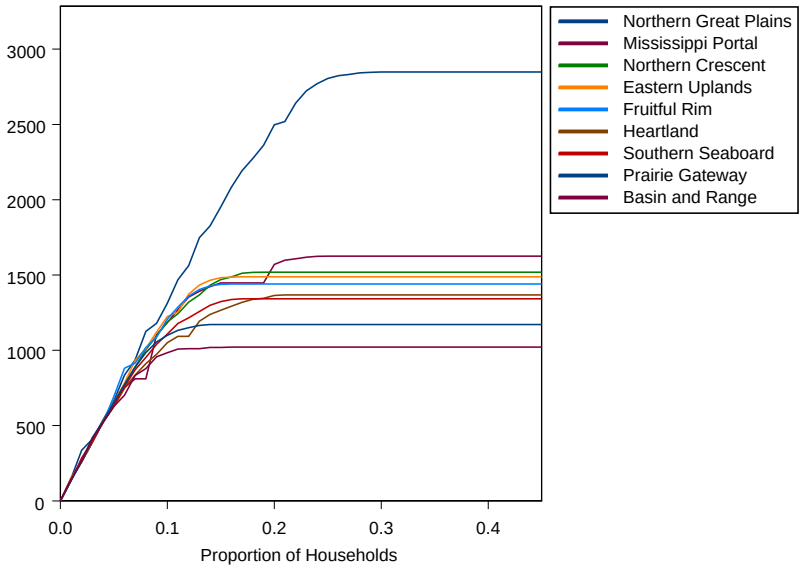

Figure 4: Cumulative Poverty Gap Curves for Household Income (top), by Region, 1998.