Household Welfare and the
Pricing of Cocoa and Coffee in Côte d'Ivoire:
Lessons from the Living Standards Surveys

Dwayne Benjamin and Angus Deaton

Cocoa and coffee are the most important crops in Côte d'Ivoire. Until recently, the difference between world and administered producer prices provided an important source of government revenue. As a result of a continued decline of world prices of both crops, however, the Ivoirien government was forced to cut producer prices in half. Because 40 percent of Ivoirien households grow either cocoa or coffee, this cut can be expected to have a considerable impact on the welfare level of these households. We use the 1985 Living Standards Measurement Survey to estimate the welfare effects of producer price changes for Ivoirien households, permitting an evaluation of the probable consequences of the recent price cut. Using nonparametric econometric techniques, we find that, although many households will suffer losses of income, the cuts will not have adverse distributional effects: cocoa and coffee farmers are scattered throughout the income distribution, but most are concentrated in the middle.

Cocoa and coffee are the most important crops in Côte d'Ivoire: cocoa and coffee orchards account for almost half of total agricultural land use, 40 percent of Ivoirien households are cocoa or coffee farmers (or both), 60 percent of export earnings are attributable to the two crops, and in some years as much as 40 percent of government revenue has come from the operating surpluses of the cocoa and coffee stabilization funds, that is, from cocoa and coffee taxes. The output of both crops is sold by farmers to licensed agents, who in turn sell to government marketing boards at prices set by the government. The choice of the official price is of central importance in determining government revenue and has been the source of some recent discomfort.

In the mid-1980s, many policy analysts recommended increasing the official price to bring farmer prices in line with world prices and to stimulate growth in the agricultural sector. However, during the late 1980s, the world prices of both

commodities measured in U.S. dollars declined steadily. To compound diffi-
culties, the domestic currency is tied to the French franc. The rise in the value of
the franc in relation to the U.S. dollar and the generally low level of world
commodity prices meant that by 1989 the government marketing boards were
receiving less for cocoa and coffee than the price paid to farmers. This turned a
revenue-generating tax into a subsidy. The consequent drains on government
revenue and foreign exchange are currently among the most difficult problems
facing the Ivoirien economy. To stabilize government revenue, the government in
1989 cut the prices of cocoa and coffee in half, resulting in the first nominal cut
in these prices to farmers in 25 years. Because of current trends in world cocoa
and coffee prices, producer prices may have to be cut again.

Two related policy questions arise. First, what price policy should the govern-
ment of Côte d'Ivoire follow in setting the farmer price in relation to the world
price? This, the conceptually more difficult question, is thoroughly addressed in
Trivedi and Akiyama (1992). Second, what are the likely effects of a change in
the level of farmer prices on production, welfare, and income inequality? This
article provides a statistical background for discussion of this question.

Ideally, we would like household and agricultural surveys taken before and
after the cuts in official prices. This would allow a precise analysis of the effects
of a price change on the welfare of households and on farm supply behavior.
However, such data do not exist. In their place we use the 1985 Living Standards
Measurement Survey (LSMS). Indeed, a major purpose of this article is to show
the usefulness of such data for analyzing price policy questions. Although these
data are significantly better than anything that existed previously, the lack of
price variation for cocoa and coffee make them inappropriate for studying pro-
ducer price-supply responsiveness. Instead, the focus is on the welfare implica-
tions of the price cuts: what happened to households after the prices of their
most important crops were halved?

Section I describes the institutional setting of cocoa and coffee pricing in Côte
d'Ivoire and places the study in the context of price policy analysis. It illustrates
the basic features of the current pricing policy and shows the magnitude of the
current fiscal problems. It also describes the distinguishing features of a price
policy and a price level and examines the informational requirements for eval-
uating policy reform. Section I outlines the analytic framework, which is the
theory of tax reform as outlined in Stern (1987) and applied in Deaton (1989a,
1989b). This theory leads to the commonsense conclusion that evaluating price
reform requires information on the supply characteristics of cocoa and coffee,
especially their price responsiveness, as well as information on the distribution
of price effects on household income.

Section II describes the LSMS data and explains how they meet the particular
needs of the analysis. It also explores some of the distinguishing features of
household income that anticipate the more sophisticated estimation results that
follow. Section III charts the distributional results of a change in crop prices for
the country as a whole and within the farming sector. Using nonparametric

techniques, we find evidence that the price changes that have taken place were unlikely to have had significant adverse distributional effects. These techniques are shown to be more revealing than the traditional parametric approach to distributional questions. Cocoa and coffee farmers, although scattered throughout the income distribution, are concentrated in the middle, so the price changes are not likely to increase inequality. However, this is not to dismiss the magnitude of the welfare loss for these farmers.

I. THE INSTITUTIONAL SETTING AND ANALYTIC FRAMEWORK

Cocoa and coffee are marketed through the government marketing board—CAISTAB (Caisse de Stabilisation des Prix des Produits Agricoles). CAISTAB pays farmers, through private agents, a preannounced price for their crops and sells the output on the world market. The difference between world and producer prices, net of marketing costs, is the CAISTAB surplus or the revenue from cocoa and coffee export taxes. Typically, this surplus has been an important part of government revenue. Through the early 1980s, cocoa and coffee tax revenue comprised 20 to 40 percent of government revenue. By 1988, world prices had dropped below producer prices and the CAISTAB surplus had become a deficit, with obvious consequences for the Ivoirien treasury. The most general statement of the policy question, then, is how producer prices should be set in relation to world prices. This should be distinguished from the more limited question of what the appropriate level of taxation should be, or in which direction the prices should be changed. The focus here is on the likely consequences of a tax change, but it is useful at the outset to explore the issues involved in determining a more general price policy.

Figures 1 and 2 show the relevant data for cocoa and coffee, respectively. In both figures, the upper series is the world price, converted at the official exchange rate and deflated by the domestic consumer price index. The world prices determine the potential revenue in constant Ivoirien currency that the government can obtain by selling a kilo of cocoa or coffee on the world market. Until 1987 the world price of robusta coffee (the type grown in Côte d'Ivoire) was about 25 percent higher than the world price of cocoa, although the prices have more recently converged. Both prices varied without obvious trend until the mid-1970s. After frosts in Brazil in 1975 had destroyed a third of the world coffee crop, prices rose very rapidly to reach their all-time high in the spring of 1977. The Ivoirien domestic equivalents of these world prices have behaved very differently. There has been considerable domestic price inflation, and the October 1991 price level was almost six times that of April 1963. However, there has been much less corresponding movement in the exchange rate between the French franc and the U.S. dollar. Because the CFA franc is tied to the French franc, as far as Côte d'Ivoire is concerned, there has been a progressive overvaluation of the currency. Consequently, although there has been much variation, there has been no trend in the domestic purchasing power of a kilo of coffee or cocoa sold on the world market.

Figure 1. *World and Official Farmgate Cocoa Prices in Côte d'Ivoire, 1963–93*
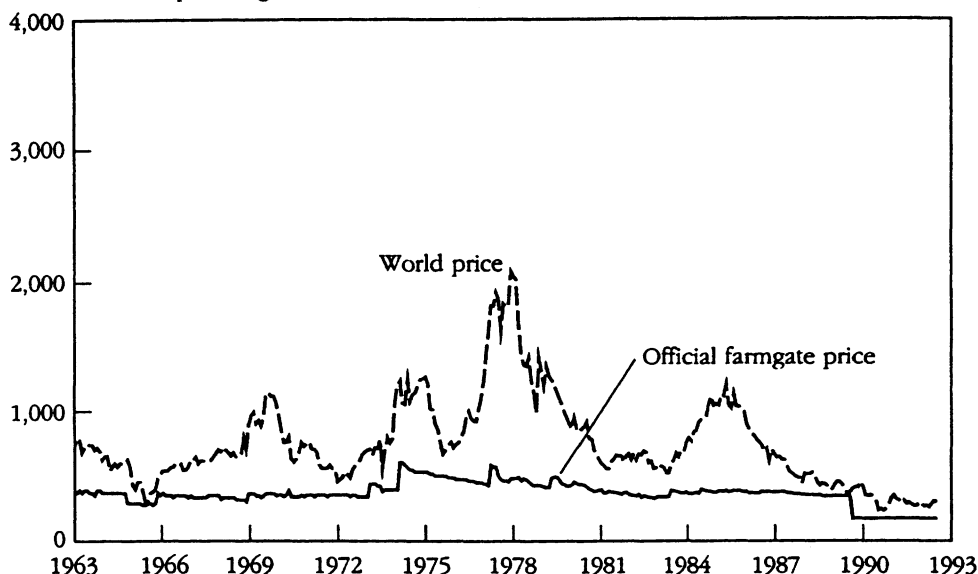
Real CFA francs per kilogram of beans

Note: Monthly price data are used. The world price is converted at the official exchange rate; all prices are deflated by the domestic consumer price index.
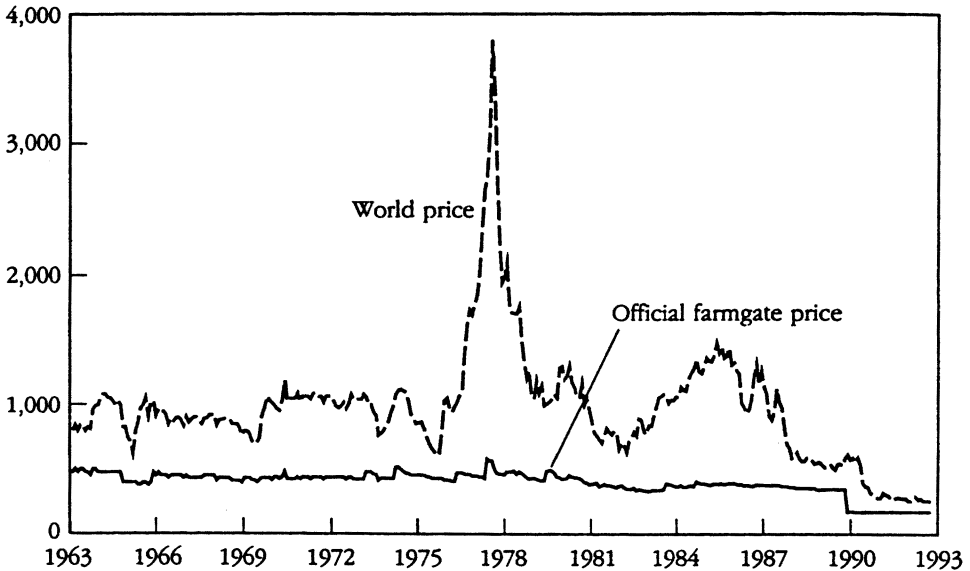Source: International Monetary Fund (various issues).

The lower series in figures 1 and 2 is the official farmgate price for the crop, again deflated by the price level. The official price series are close to being constant in real terms—until 1989. The occasional upward "blips" are characteristic of an administered series that is increased discretely in response to continuous price inflation. For 1963–92 the mean ratio of administered prices to world prices for cocoa is 54 percent; for coffee, 45 percent. The government has typically ignored the world price differential in favor of coffee and set the same price per kilo for both commodities; therefore, the ratio of domestic to world prices for coffee is lower than it is for cocoa.

Figures 1 and 2 suggest the kinds of issues that need to be examined when considering the consequences of price reforms. First, setting domestic prices equal to world prices is unlikely to be a viable option in the longer term. Had such a policy been followed in the past, the government would have been deprived of much of its revenue, and government expenditures based on those revenues could not have been sustained. Although higher cocoa and coffee prices would have meant higher incomes for farmers, we cannot claim that this by itself would have made the policy desirable without looking at the likely consequences for government expenditures. Reducing taxes may well reduce the size of dead-

Figure 2. *World and Official Farmgate Coffee Prices in Côte d'Ivoire, 1963–93*

Real CFA francs per kilogram of beans

Note: Monthly price data are used. The world price is converted at the official exchange rate; all prices are deflated by the domestic consumer price index.
Source: International Monetary Fund (various issues).

weight loss, and, if government expenditure is thought to be socially less valuable than private expenditure, a reallocation in favor of farmers will be desirable. But the argument is then about reducing government expenditures, not about the general desirability of having an untaxed agricultural sector.

Evaluation of the distributional consequences of the level of taxation in this context obviously depends on who is taxed and to whom the revenues from cocoa and coffee are diverted by the government. Some answers are provided by Gbetibouo and Delgado (1984). They identify three equity issues: interpersonal, regional, and sectoral. Interpersonal equity depends on the progressivity, or incidence, of the cocoa tax as well as on the distribution of government expenditure. The incidence issue is the focal point here. Sectoral and regional issues are closely linked in Côte d'Ivoire. There is one major city, Abidjan, and it receives most of the expenditure on urban infrastructure. The agricultural sector is divided between the cocoa- and coffee-growing Forest regions, and the poorer, drier Savannah in the north.

Gbetibouo and Delgado report that more than two-thirds of the CAISTAB surplus has historically been spent in sectors other than agriculture, especially urban infrastructure. Of the transfers to agriculture, only a negligible amount goes to the CAISTAB stabilization fund. Nevertheless, some expenditure does find its way

back to the cocoa and coffee farmers. Extension and research services, as well as input subsidies help encourage cocoa production. For farmers in the Forest region, there are also subsidies for diversifying into palm kernel and pineapple production. The Savannah region, which has very few cocoa or coffee farmers, also receives subsidies for rice, tobacco, and (especially) cotton. Furthermore, Gbetibouo and Delgado argue that the investment in rural infrastructure, such as roads, has fostered cocoa and coffee development and contributed to "absolute" equity, the overall level of opportunity in rural areas. This does not, however, change the fact that most of the coffee and cocoa surplus has left the sector.

A more subtle argument against the current pricing policy is that trade taxes are themselves typically inefficient and that the government could raise the same revenue with less distortion from alternative instruments. Certainly, the agricultural taxes, support prices, and subsidies each alter crop production from what would occur if the farmers faced world prices. In the case at hand, less coffee is probably grown than would be the case if cocoa and coffee were priced according to their world prices. But the desirable alternative taxes would typically be profits taxes or income taxes on cocoa and coffee producers, which are administratively more difficult to collect. In countries such as Côte d'Ivoire, where income taxes can be levied only on a very small fraction of the population and where only a limited range of goods can be taxed, government expenditure must be at least partially financed out of taxes on agricultural commodities, in this case cocoa and coffee. At the same time, many governments attempt to redistribute risk from farmers to the general population by a policy of stabilizing prices. Although there are doubts whether high-taxation stabilization policies can be sustained indefinitely, the two figures show that, until recently, the government had succeeded in stabilizing the farmgate levels of both prices, at least as they relate to world prices. The question that cannot be answered from current knowledge is whether the average level of cocoa and coffee taxation has been too high or too low. The general issues involved are well understood (see, for example, the excellent discussion by Newbery 1987).

Instead of determining the appropriate level of each producer price, one feasible reform would be to allow world prices to affect farmgate prices more directly. This could be accomplished by setting an explicit export tax, say of 50 percent, so that farmers would get much the same on average, but could respond directly to the signals provided by the world market. The rate of export tax could also be varied over time, with a high charge when world prices are high and a low or zero rate when prices are low. The argument favoring greater openness is the usual one—that farmers are then free to respond to the world prices that represent the true social opportunity costs of their crops. Economic efficiency and national income can be expected to increase at the expense of a different and perhaps less desirable distribution of risk. Also, if Côte d'Ivoire has any market power in world markets, it is in cocoa, not coffee. Standard optimal tariff arguments would therefore suggest a higher export tax on cocoa than on coffee—precisely the opposite of the current situation.

Yet, given the unsuccessful attempts of Côte d'Ivoire to manipulate the world price of cocoa in 1979 and again in 1988 by withholding its cocoa from the market, the monopoly power component of the optimal export tax should not be exaggerated.[1] The aspect of an explicit export tax policy that is less clear-cut is its effects on increasing the riskiness of cocoa and coffee farming. As figures 1 and 2 show, until 1989 the farmers' price had varied little compared with the world price. If domestic prices were to reflect even a fraction of the variability of world prices, planting decisions would be riskier, and risk-averse farmers would presumably plant less. Even so, the ability of rural smallholders to smooth income fluctuations should not be underestimated, and it is far from clear that the government itself can exercise the sort of control over its own expenditure that is required in the face of fluctuating export prices and tax revenues. Indeed, many of the current macroeconomic problems of Côte d'Ivoire are consequences of the government's having permanently increased public expenditure in response to the high and, in retrospect, transitory revenues of the late 1970s.

This discussion implies that evaluating the cocoa and coffee pricing policy decision requires (at least) two sets of measurements. The first set is own and cross price elasticities for cocoa, coffee, and other crops. Certainly farmers will adjust their investment and labor decisions in response to the changing relative price structure. In some areas, switching to other crops will be easier than in other areas. Although these magnitudes are very important, we have little to say about them. There are two problems with cocoa and coffee own price elasticities. First, in a cross-sectional data set such as the LSMS, there is virtually no price variation (the prices are administered). Furthermore, it would be difficult to measure long-run supply responses because both crops are grown on long-lived trees. Deaton and Benjamin (1988) explore some of the features of cocoa and coffee production that may be pertinent to agricultural policy, but estimation of price elasticities is not possible. Second, because both the absolute and relative prices of cocoa and coffee have been approximately constant (until recently), even time-series data would be difficult to use to identify supply elasticities. Unquestionably, one of the best ways to measure the supply response would have been to collect data before and after the price cut. This was not done. The second set of magnitudes useful for the price policy analysis is to measure the incidence, and the level of impact on household welfare, of changes in the level of cocoa and coffee producer prices.

A traditional approach to measuring the distributional consequences of the price cuts would be to compare Gini coefficients before and after the price cut. Summarizing the effects in a single measure ignores useful information that can be gleaned from the data. Furthermore, we are interested in the size of the welfare consequences of the price cuts as well as in the implications for income

---

1. See Gbetibouo and Delgado (1984) for a discussion of the 1979 episode and Hiltzik (1989) for the more recent one.

inequality. A more illuminating and convenient way to proceed is as follows. Household income (earnings), Y, can be written:

$$(1) \qquad\qquad Y = \pi(p; A) + E + \mu$$

where $\pi(p; A)$ is the (restricted) profit function depending on prices, p, and land holdings, A. E is household labor earnings, and μ is other income. Consider a "small" price change, dp. The change in farmer income is the derivative of the profit function with respect to prices multiplied by dp, or qdp, where q is the farm's output. Consider then the proportional change in consumption expenditures, x, that could be maintained out of the price change. If this is denoted $d\ln x$, we have

$$(2) \qquad\qquad d\ln x = \frac{qdp}{x} = \frac{pq}{x} d\ln p$$

so that the elasticity of the farm household's potential consumption with respect to price is the "benefit ratio," or ratio of sales to total consumption. High-ratio farmers earn more of their income from cocoa and coffee and are thus more exposed to changes in these prices. This benefit ratio is a convenient (and dimensionless) measure of the benefit (or cost) of a price change for each farm household. The distribution of the benefit ratio in relation to, say, household per capita expenditure (pcx) can indicate the distributional effects of price changes: who loses the most and who loses the least in response to a price cut.

The data requirements for this exercise are simple, although it will not always be easy to obtain data of the quality necessary for reliable estimation. Measures are needed for household income from cocoa and coffee sales, pq; total consumption, x; and the welfare measure, $\ln pcx$. We choose $\ln pcx$ as a welfare measure rather than an obvious alternative based on income for two reasons. First, consumption measures are likely to be better measured than income. Second, given the natural year-to-year variation in incomes, $\ln pcx$ will often be a superior measure of the "permanent" well-being or welfare level of households.

The above framework is easy to implement, and its simplicity highlights important caveats of interpretation. In reality, price changes are far from infinitesimal; procurement prices in Côte d'Ivoire have been cut by half. However, the benefit ratios are still likely to be a good guide to distributional effects. For small changes, substitution effects cancel out, but, when price changes are large, farmers will moderate the effects of a cut (say) by substituting out of cocoa and coffee. If the substitution effects are much the same for different farmers and do not vary by income level, the effects on income distribution will be correctly indicated by the benefit ratios. Of course, the estimated magnitude of the benefit ratios will overstate the welfare losses if farmers substitute out of cocoa and coffee. For a more detailed discussion of these issues, see Deaton (1989b). An additional consideration to bear in mind when analyzing cocoa and coffee supply, however, is that they are tree crops and that long-run adjustment of production may be slowed by rigidities in adjusting the stock of trees. As with the short-run supply elasticities, data are not available to answer these long-run questions,

but some of the first-order distributional consequences of the tree-crop nature of cocoa and coffee are examined.

II. THE LSMS DATA

For the analysis of questions of agricultural policy, the ideal data set would combine the precision of measurement of a traditional agricultural survey with the breadth of other information gathered in a typical household survey. Specifically, to implement our methodology, we need accurate information on household cocoa and coffee income and on per capita consumption. The LSMS samples only a relatively small number of households—1,600, of which almost half are urban—and sacrifices detail in any one area of household behavior to obtain the broader picture that can come only from collecting a great deal of interrelated data on all aspects of the household's economic activities. The survey is very useful for analyzing household incomes, and thus welfare, although less useful for learning about agricultural production.

Standard agricultural surveys collect information on quantities grown, land cultivated, yields by crop and variety, detailed use of inputs (including labor for each crop), and the relevant prices facing the farmer for inputs and a range of actual or potential outputs. For tree crops, it is also useful to know the characteristics of each stand, particularly the age structure and the varieties. Even in theory, only some of these data can be garnered from the LSMS. The main objective of the agricultural module of the survey was not to collect farm data as such but to determine the sources of farm income and the value of farm assets. To this end, questions were included on land, crops grown, age structure of tree crops, sharecropping, use of inputs, livestock, farm capital, and agricultural processing activities.[2] Respondents were asked to indicate current levels of stocks or a retrospect of flows over the previous year (or both). This section of the survey provided measures of household income based on cocoa and coffee, as well as on other crops. Income from cocoa and coffee corresponds to total farm revenue, less payments in kind made to *metayeurs*—Ivoirien sharecroppers. A particular strength of the LSMS is its measurement of household expenditures, including an attempt to measure the value of home-produced food. We were thus able to construct a comprehensive measure of household consumption, and therefore ln*pcx*. For more detail on the construction of household consumption, see Glewwe (1988, 1990) and Kozel (1990).

Before looking at estimation results, it is worth looking at some simple tables of means. Although sample averages hide detail, they do reveal the basic patterns that we will elaborate on with our estimation procedure. Table 1 describes the size distribution of farms in Côte d'Ivoire's main agricultural zones, and, specifically, the Savannah and Forests. Until recently, Côte d'Ivoire was

---

2. In Deaton and Benjamin (1988) we compare the LSMS estimates of cropping patterns and levels to results from other surveys and find that although the LSMS data are not perfect, they match up quite well.

Table 1. *Size Distribution of Farms in Côte d'Ivoire, 1985*

Size of farm (hectares)	Percentage of farms in size category			Percentage of overall cropped area in each farm-size category		
	Savannah	Forest	Overall	Savannah	Forest	Overall
Less than 0.99	2.8	2.7	2.7	0.0	0.0	0.0
1.0 to 1.99	4.6	4.1	4.3	0.6	0.3	0.3
2.0 to 4.99	37.3	13.5	21.0	14.9	2.8	5.1
5.0 to 9.99	33.0	25.1	27.6	28.9	11.8	14.9
10.0 to 19.9	17.4	34.5	29.1	31.6	32.8	32.6
20.0 to 49.9	4.5	17.3	13.3	20.0	34.1	31.5
More than 49.9	0.3	2.8	2.0	4.0	18.4	15.7
Average farm size	7.3	14.8	12.5			

Source: LSMS 1985 data base.

"land rich," and farmers could establish property rights as long as they used the land. The resulting distribution of land is quite equal, and most farmers are smallholders. For the Forest regions, where the cocoa and coffee are grown, 80 percent of the farms are under 20 hectares, and this represents about 50 percent of the total cropped area.

The most important determinant of cocoa and coffee output for a farmer is the state of the capital that produces the crops: the trees. Cocoa and coffee trees have well-pronounced age-productivity profiles. The LSMS data tell us the number of hectares of trees that are "too young," "fully mature," and "near the end of their productive life," with definitions of these categories left to the respondent. Table 2 shows that almost 40 percent of the cocoa stands were planted too recently to be productive; it takes about five years for a cocoa tree to reach maturity, and, when these trees come into production, output can be expected to increase significantly. The coffee orchard, by contrast, was old in 1985: only 20 percent of the trees were newly planted. If producer prices are set in favor of cocoa over coffee, it is likely to be reflected in a further decrease in coffee planting and a long-run decline in coffee output. Thus, there are two implications of the age structure for this study. First, the current benefit ratio is not a perfect measure of "exposure" in coffee or cocoa. A farmer with a lot of young trees may have a low benefit ratio but will probably be more affected by the price cuts as the trees mature. In this sense, there are dynamic aspects to the welfare question that are partially missed by a static approach. By contrast, the tree nature of the crops suggests that the short-run supply response is likely to be smaller than for an annual crop, such as corn. Thus, a farmer with only mature trees is likely to be affected in a similar manner for at least a few years, net of any short-run supply responses. A second important point from table 2 is the close correspondence between cocoa and coffee farming. Almost 80 percent of cocoa farmers grow coffee, and 67 percent of coffee farmers grow cocoa. This will have obvious implications for the role of cocoa and coffee income for farmers: cocoa and coffee farmers are mostly the same people.

Table 2. *Age Structure of Tree Stands and Percentage of Cocoa Farms Growing Coffee in Côte d'ivoire, 1985*

	Percentage of trees in cocoa stands by age structure			Percentage of cocoa farms growing coffee
	Too young	Fully mature	Near end	
East Forest	33	54	12	71
West Forest	47	50	3	91
Savannah	69	31	0	79
Total	39	52	9	78
	Percentage of trees in coffee stands by age structure			Percentage of coffee farms growing cocoa
	Too young	Fully mature	Near end	
East Forest	18	61	22	74
West Forest	17	76	7	73
Savannah	27	62	9	19
Total	18	67	15	67

Source: LSMS 1985 data base.

Although tables 1 and 2 illustrate important features of the distribution of agricultural capital, table 3 describes the relation between cocoa and coffee incomes and household incomes as a whole. Each column shows the mean annual income for a different type of household. The first column is for all Ivoirien households, or at least for all those represented in the survey, and the second column is for all farm households. The third and fourth columns give the household budgets for cocoa and coffee farmers, respectively. As already stated, the cocoa and coffee groups are not mutually exclusive; indeed, coffee farmers on average derive more net income from growing cocoa than from growing coffee. Starting near the bottom of the table, we see that farm households, that is, those reporting some sort of agricultural activity, have average incomes that are 73 percent of the national average.

Cocoa and coffee farmers do better than the average farm household: cocoa farmers have average incomes very close to the national average, and coffee farmers' average incomes are almost 88 percent of the national average. Recall that the northern, Savannah region, which has the lowest income levels in Côte d'Ivoire, grows essentially no cocoa and little coffee. The results are similar for total expenditure. Farm households have household expenditures equal to 71 percent of the national average, whereas cocoa and coffee are 88 and 79 percent, respectively. The welfare differences implied by per capita expenditure are more pronounced, reflecting the larger rural households. With this measure, farm households are only 58 percent as well off as the national average, and cocoa and coffee farmers, although better off than the average farm household, are, respectively, only 64 and 61 percent as well off as the national average.

From the first column of table 3, more than 30 percent of average household income derives directly from agricultural activities; cocoa alone accounts for 7.0

Table 3. *Average Household Income and Expenditure Data
for Various Groups in Côte d'Ivoire, 1985*

	All households	All farm households	Cocoa farmers	Coffee farmers
Cocoa				
Sales			502	
Less nonlabor inputs			9	
Less metayage (labor costs)			111	
Net cocoa income	110	166	381	266
Coffee				
Sales				228
Less nonlabor inputs				4
Less metayage (labor costs)				55
Net coffee income	56	85	145	168
Home-produced food	203	307	381	376
Net other agricultural income[a]	118	178	267	219
Total agricultural income	487	736	1,174	1,029
Nonagricultural income				
Wages	533	133	105	66
Self-employment	306	162	188	171
Other income[b]	236	115	112	101
Total nonagricultural income	1,074	410	404	338
Total income	1,562	1,146	1,579	1,368
Household expenditure	1,638	1,161	1,434	1,296
Per capita expenditure (*pcx*)	264	153	170	160
Sample size	1,559	1,033	450	522

a. Includes net other-crop, net livestock, transformational, and farm rental incomes, less labor costs not directly attributable to a particular crop.

b. Includes gifts and pensions, rental, transfers, imputed rent, and income from forced savings.

Source: LSMS 1985 data base.

percent and coffee for 3.6 percent. In addition, many income-generating activities in rural areas depend on agriculture. Less obviously, the prices received by farmers for cocoa and coffee in 1985 were well below world prices, and the figures take no account of the various publicly provided goods and services that were financed out of cocoa and coffee taxes. For farm households, agricultural income is 65 percent of total income, and cocoa and coffee combined account for about a third of average agricultural income. The most important source of income for farmers is home-produced food, the imputation of which accounts for 44 percent of agricultural income.[3] As noted also by Budd (1993), the high value of home-produced food will cushion all farm households from income shocks. At least on average, this places an upper bound on the sensitivity of household income to changes in crop prices.

Cocoa and coffee farmers are well diversified into other income-producing activities. Most farmers who grow one crop grow the other, and neither coffee

3. The value of home-produced food is provided by the survey respondents in answer to the question, "During the past 12 months, have members of your household eaten foods grown or raised by the household?" For the relevant recall period of the expenditure survey, they respond to "How much would it cost to buy the amount they eat each day?" There is no outside evidence that these data are overstated.

nor cocoa farmers derive more than half their agricultural incomes from the two crops. Again, home-produced food is an important source of income, as are various other crops. There is also a substantial amount of income generated from a range of other activities, including agricultural processing, non-agricultural business, wages, and rental income. The cost side of cocoa and coffee farming consists largely of the costs of *metayage*—payments in kind to sharecropperlike farm laborers. Expenditures on nonlabor inputs are very small because few farmers make any use of them.

For both crops, net income is 76 percent of gross sales, and both fetched the same price in 1985; however, the average yield per hectare of coffee trees is less than 60 percent of the average yield of cocoa trees, so cocoa is the much more profitable crop. Our sample is inadequate in documenting the position of the metayeurs; for example the data show 155 cocoa farmers who report using metayeurs, but only 15 households who claim to be metayeurs. We could partially allow for this by counting the cost of metayage as agricultural income, not cost. However, according to the little information that we have, the metayeurs, who are largely migrant workers, have low incomes, perhaps CFAF500,000, or less than a third of that of the average cocoa farmer. Here, then, is a group of low-income workers who would most likely be hurt by a significant decrease in crop prices but who are unfortunately much underrepresented in the data. This deficiency must be kept in mind when we look at the effects of price changes on income distribution.

III. Estimation of Distributional Consequences

As suggested earlier, perhaps the most straightforward way to evaluate the distributional consequences of a price change for cocoa and coffee would be to compare levels of inequality before and after the price change. Assuming there were no supply responses, simulation of the level after the price change simply involves recomputing farmer income with the appropriately changed prices. Such an experiment would allow us to determine the effects on inequality of a nonmarginal price change. The results of conducting this simple exercise are presented in table 4. Although the levels of these inequality measures are intrinsically interesting, they are extensively discussed elsewhere.[4] We are interested in the changes that result from changes in cocoa and coffee prices, as well as differences that may exist between farmers and the rest of the population.

The basic result appears to be that a price increase would decrease countrywide inequality but increase inequality among farmers; however, both changes are extremely small. A price decrease—the policy that actually occurred—slightly increases countrywide inequality and negligibly reduces inequality among farmers. Basically, lowering cocoa incomes shifts money away from farmers, increasing total inequality, but perhaps there is less of a relation between cocoa or

4. See, for example, Glewwe (1988, 1990) and Kozel (1990).

Table 4. *Gini Coefficients*

	All households	Farm households
Actual or "before" per capita consumption	0.455	0.365
Actual or "before" per capita income	0.536	0.475
Simulated per capita income, 50 percent increase in cocoa and coffee prices	0.529	0.482
Simulated per capita income, 50 percent decrease in cocoa and coffee prices	0.551	0.479

Source: Authors' calculations.

coffee farming and the level of income among farmers. Although the estimates that follow will not overturn these results, our approach provides in striking detail the reasons why these simple numbers yield the answers they do in this experiment. After all, why should cutting the value of the two most important crops have such a small effect on income inequality? Furthermore, the Gini coefficient is only a summary of inequality; it fails to indicate the magnitude of the welfare loss that these households experienced. To understand all that is hidden by these "experiments," it is necessary to explore the structure of income distribution in some detail.

There are two features of the income distribution that we wish to investigate. First, what is the distribution of the benefit ratio, and how does it relate to the level of household welfare; that is, who gets hit hardest by the price cut? Second, what is the underlying distribution of household welfare? Both questions can be explored using nonparametric statistical techniques. Average benefit ratios at different (log) per capita expenditure levels are given by nonparametric regressions (see, for example, Härdle 1990), and the joint and marginal densities of the two variables can be estimated using techniques in Silverman (1986). See Deaton (1989a, 1989b) and Budd (1993) for previous applications. The techniques are precisely described in these references.

The basic idea of the regression is that it estimates a weighted average of the benefit ratio around a particular value of ln*pcx*, where the weights are smaller for point observations more distant from the value of ln*pcx*. These means are computed over the entire domain of ln*pcx*, much like a moving average. More generally, if the regression function for a variable y, given x, is denoted $E(y|x) = m(x)$, then we estimate $m(x)$ over the n observations by

$$(3) \qquad \hat{m}(x) = \sum_{i=1}^{n} w_i(x, X_i)y_i$$

where $w_i(x, X_i)$ is the weighting function that depends on the kernel and bandwidth.[5] We use the Epanechnikov kernel, although the particular choice of kernel does not have much effect on the estimates.

5. For a detailed treatment of nonparametric regressions, including discussions of the choice of bandwidth and kernels, see Härdle (1990). See also Deaton (1989a), especially pp. 207–08, for a discussion of estimation issues in nonparametric regressions and density estimation.

The more important choice variable for the researcher is the bandwidth, which determines the degree of smoothness of the function to be estimated. The bandwidth affects the width of the interval over which the expected value of the dependent variable is calculated by determining when $w_i(x, X_i)$ becomes small. Wide bandwidths result in weights that allow data from a wider part of the domain around x to affect $\bar{m}(x)$. This causes too much smoothing and eliminates the detail of the empirical relation. Because they allow data further away from x to have more influence on $\bar{m}(x)$, wide bandwidths can result in biased estimators. By contrast, not enough smoothing allows noise in the data to obscure the patterns we seek to find. In the estimates that follow, the selected bandwidth sacrifices smoothness for less bias. The densities, denoted $f(x)$, are estimated in a manner similar to the regressions, only instead of the expected benefit ratio, we estimate the number of households at a point in the domain of lnpcx.[6]

The illustration of the results in the figures raises several presentation issues. First, like parametric estimators, nonparametric estimators are random variables with sampling distributions. Therefore, standard errors exist for these estimators.[7] To convey the precision of the estimates, 95 percent confidence bands can be computed and graphed with the estimated functions. A key ingredient in these pointwise confidence bands is an estimate of the variance at a point x:

$$(4) \qquad \hat{\sigma}^2(x) = \frac{1}{n} \sum_{i=1}^{n} w_i(x, X_i) [y_i - \bar{m}(x)]^2$$

that is, a weighted average of the squared difference between the actual and predicted y. As shown in Härdle (1990), a 95 percent confidence interval can be constructed at a point x as

$$(5) \qquad \bar{m}(x) \pm 1.96\hat{\sigma}(x) \frac{\sqrt{c_k}}{\sqrt{nh\hat{f}_h(x)}}$$

where c_k is the definite integral of the squared kernel ($c_k = 0.6$ for the Epanechnikov kernel) and $\hat{f}_h(x)$ is the estimated density at x, indexed by the bandwidth h. For clarity of presentation, we show only the standard errors for our main set of results, those for the benefit ratios. For the other estimates, the implied confidence bands are wide at the end points of the lnpcx domain because there are so few observations (low density) at these extremes.

The second presentation issue broadly concerns the scale of the graphs. The domain of lnpcx is quite wide, although the data are actually sparse at the end

6. See Silverman (1986) for a discussion of density estimation. Deaton (1989a: 207–08) also shows how kernel estimates of both univariate and bivariate densities are computed.

7. See Härdle (1990), especially pp. 98–102, for a discussion of pointwise confidence intervals. Härdle demonstrates the theoretical properties of the estimators in addition to outlining the computations. See also Budd (1993: 600–01) for a discussion of standard errors in nonparametric regressions.

points. As a result, the estimates at the end points are not well behaved, and, given the scale of the data, the plots can be misleading. To avoid presenting misleading estimates, even though the regressions are estimated over the entire sample, we show only the regression of the domain where there are "enough" observations. Basically, we plot only the regressions over the sample where the density is greater than 0.03; this means trimming 1 to 2 percent of the households, mostly those with especially high levels of lnpcx.[8]

Figure 3 presents the density of log per capita expenditure (lnpcx) for the entire country (the relevant one for national social welfare). These are not mutually exclusive densities. For all households the mode of lnpcx is approximately 4.4. For comparison, the log of average pcx for the country is closer to 5.6, indicating significant positive skew in the countrywide distribution of welfare. Indeed, one benefit of the nonparametric approach is that it reveals these features of the data. Sample means, especially with noisy or skewed data, over-summarize and may be midleading. For comparison, the densities of lnpcx for co-

Figure 3. *Distribution of Log per Capita Household Expenditure for All Households, Cocoa Farmers, and Coffee Farmers in Côte d'Ivoire, 1985*

Estimated density

Household per capita expenditure (natural log), ln(pcx)
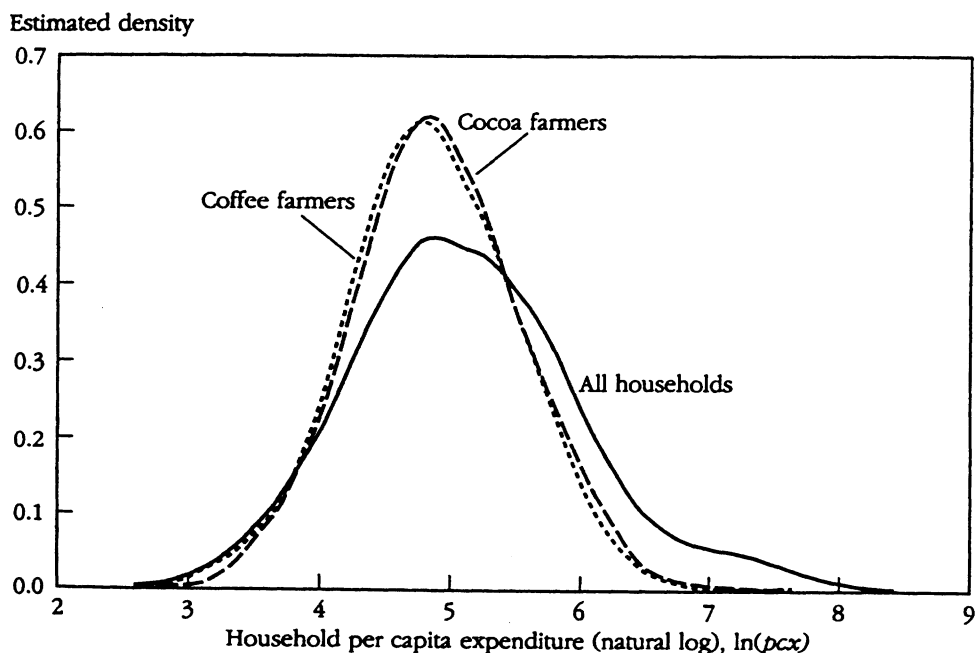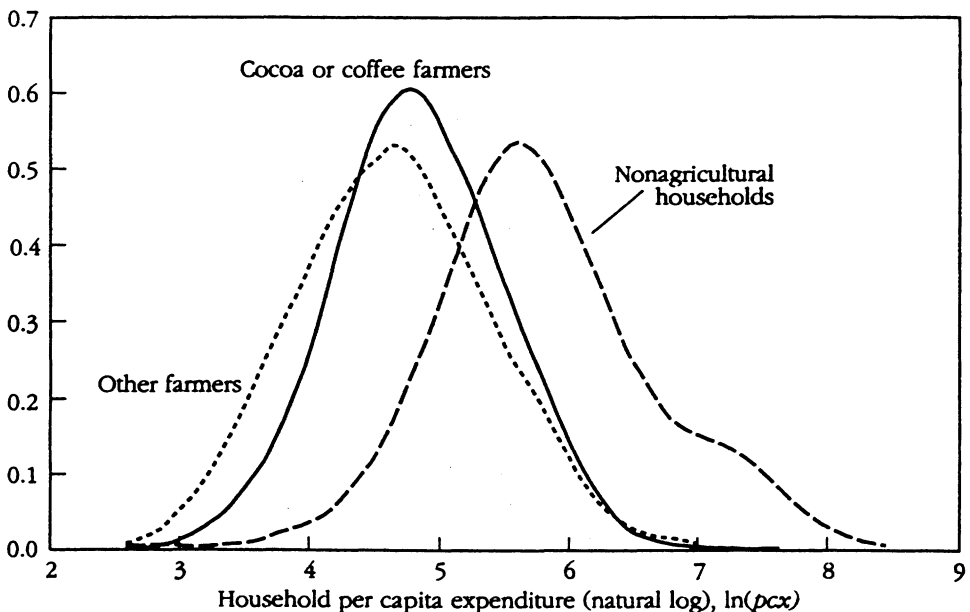
Note: Bandwidth = 0.5.

8. For regressions with all households, this means trimming estimates corresponding to 25 observations (1.6 percent); for farm households, this corresponds to 13 households, or 1.25 percent of farm households.

coa and coffee farmers are also presented in figure 3. Both groups have less disperse welfare than the country as a whole, with a mode not far from the national mode. The big difference in densities is the lack of an upper tail. Most of the rich households in Côte d'Ivoire are neither cocoa nor coffee farmers. Finally, the densities of cocoa and coffee farmers correspond closely. The only difference, as one might guess from table 3, is that the cocoa density is shifted slightly toward a position of higher mean welfare.

Because the densities in figure 3 have significant overlap of individuals, figure 4 breaks the country into three mutually exclusive groups to look at their distribution of welfare. The groups are cocoa or coffee farmers, other farmers, and nonagricultural households. The figure shows strikingly that the nonagricultural households have higher mean welfare, as measured by lnpcx. Only a small fraction of nonagricultural households have less than the median income of farm households, especially households that are not cocoa or coffee farmers. Also of importance, cocoa or coffee farmers appear slightly better off than other agricultural households. At least it appears that the very worse off households in rural areas are not cocoa or coffee farmers, because cocoa or coffee farmers are concentrated around the rural mode.[9] One feature in the figure appears different

Figure 4. *Distribution of Log per Capita Household Expenditure by Principal Source of Income in Côte d'Ivoire, 1985*

Estimated density

Household per capita expenditure (natural log), ln(pcx)

Note: Bandwidth = 0.5.

9. This is in line with the view of these farmers presented in the excellent article by Hecht (1983).

Figure 5. *The Proportion of All Households Engaged in Cocoa Farming, Coffee Farming, and Agriculture in Côte d'Ivoire, 1985*

Estimated probability

Household per capita expenditure (natural log), $\ln(pcx)$
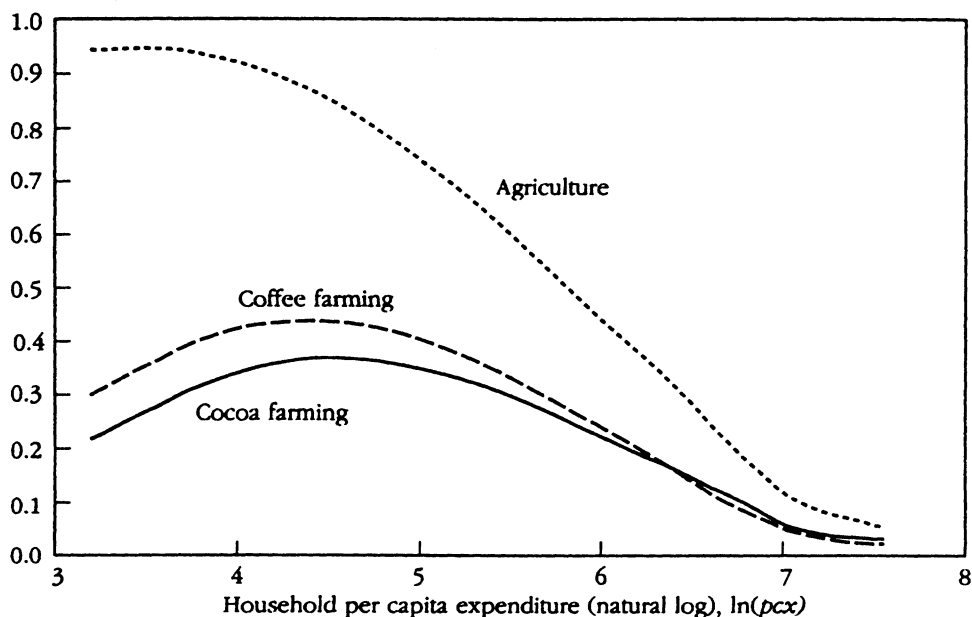
Note: Bandwidth = 1.0.

from the means in table 3: cocoa or coffee farmers are not as much better off than their rural neighbors as previously suggested.

A different way to map who does what and where in the income distribution is to estimate the proportion of households that engage in cocoa or coffee farming. This is a probability regression, where the dependent variable indicates whether the household engages in the relevant activity. Figure 5 examines these probabilities for the country as a whole, and figure 6 focuses on the rural sector. Whereas the densities in some sense show how many cocoa farmers are rich, these functions show how many of the rich households are cocoa farmers. First, for the country as a whole, the domain of $\ln pcx$ that accounts for most households is 3 to 8, and for agricultural households it is 3 to 6.5. One should thus discount regression results outside this part of the domain, because the estimated regressions are based on very few observations and are not reliable. The results are striking. If you are poor ($\ln pcx$ less than 4) you are more than 90 percent likely to be a farmer of some kind. Over this low range of welfare, you are somewhat less likely (20 or 30 percent likely) to be a cocoa or coffee farmer. The probability of being in agriculture declines rapidly as welfare rises. When $\ln pcx$ reaches 6.5, the probability of being a cocoa or coffee farmer is less than 20 percent, or less than 33 percent for agriculture in general. In figure 6, cocoa or

Figure 6. *The Proportion of Agricultural Households Engaged in Cocoa or Coffee Farming in Côte d'Ivoire, 1985*

Estimated probability

Household per capita expenditure (natural log), ln(*pcx*)
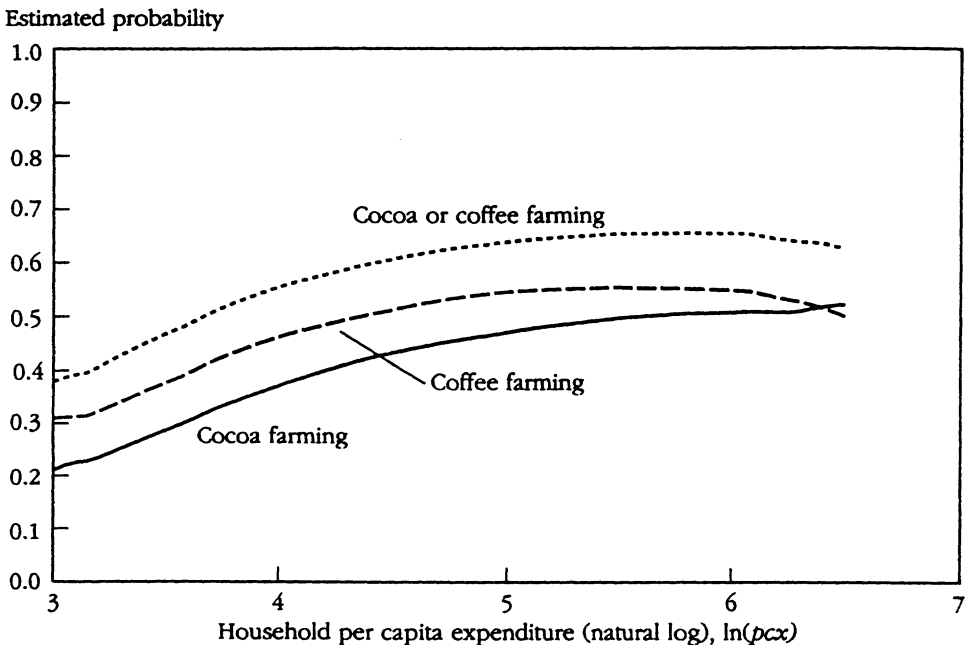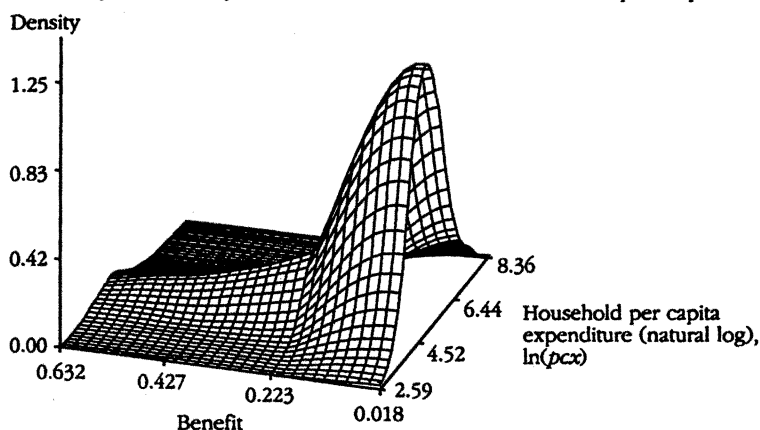
Note: Bandwidth = 1.0.

coffee farmers are not unrepresented in the poorest categories, but fewer than half of poor farmers grow either crop. The probability of being a cocoa or coffee farmer increases as ln*pcx* rises, but the probability eventually declines. There are clearly some rich farmers in rural areas who do not grow coffee or cocoa.

The more complicated sets of estimates are associated with the primary relation of interest in this section: that between the benefit ratio and ln*pcx*. We present the results of the estimation in figure 7, which has three panels. The first two panels show the estimate of the joint density of the benefit ratio and ln*pcx*. The densities are plotted directly in panel A, and the contours are plotted in panel B. The densities summarize all the relevant information on the statistical relations between the two variables. The average benefit ratio is estimated as a function of ln*pcx* (the regression of the benefit ratio on ln*pcx*) and presented in panel C.

The welfare effects of changing cocoa and coffee prices were estimated both separately and jointly. Most of the estimates revealed similar patterns. Because a joint cut in prices actually occurred, we present the results for a cut in both prices. In figure 7A the distribution is fairly symmetric about the median of ln*pcx*, with most of the mass distributed between the benefit ratio of 0 to 0.20. That is, few people in Côte d'Ivoire derive more than 20 percent of their income from cocoa and coffee. Of those who do, most are in the middle of the ln*pcx*
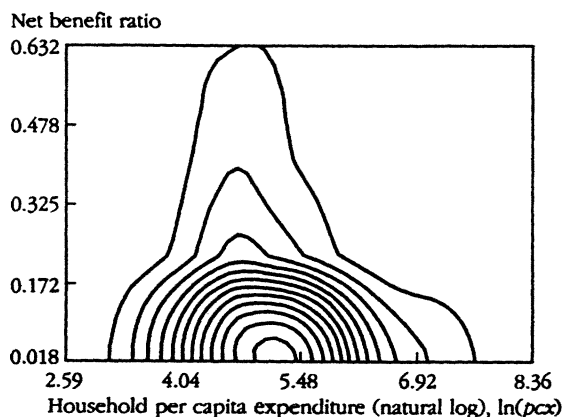
Figure 7. *The Cocoa and Coffee Benefit Ratio for All Households in Côte d'Ivoire, 1985*

A. Estimate of the Joint Density of the Benefit Ratio and Household per Capita Expenditure
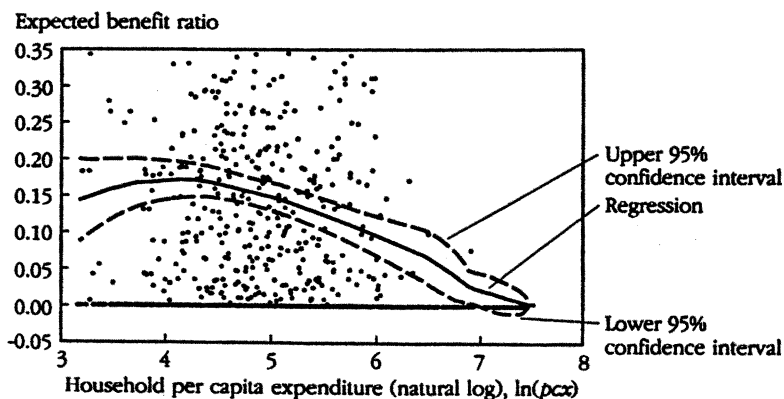
Density

1.25

0.83

0.42

0.00

8.36

6.44   Household per capita
expenditure (natural log),
4.52   ln(pcx)

2.59

0.632   0.427   0.223   0.018

Benefit

Note: Bandwidth = 0.7.

B. Estimates of the Contours of the Joint Density of the Benefit Ratio and Household per Capita Expenditure

Net benefit ratio

0.632

0.478

0.325

0.172

0.018

2.59   4.04   5.48   6.92   8.36

Household per capita expenditure (natural log), ln(pcx)

Note: Bandwidth = 0.7. The contour lines each represent increments of 0.1 units of density.

C. The Average Benefit Ratio Estimated as a Function of Household per Capita Expenditure

Expected benefit ratio

0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
-0.05

3   4   5   6   7   8

Household per capita expenditure (natural log), ln(pcx)

Upper 95%
confidence interval

Regression

Lower 95%
confidence interval

Note: Bandwidth = 1.0. The scatter points represent raw data.

distribution. This is clearer in figure 7B. The density for large-benefit farmers is seen to be highest before the mode of the distribution of ln*pcx* (4.7 compared with 5.0). The distribution is skewed toward higher-income households, so the mode is less than the mean.

The estimate of the regression in figure 7C shows the relation between benefit and welfare most clearly. In this panel we superimposed the estimated regression and confidence bands with a plot of the underlying data. To focus on the regression, we excluded observations with benefit ratios greater than 0.35 from the graph of the raw data. Many observations have a benefit ratio of 0. Nevertheless, the estimated regression has a clear pattern. Lower-to-middle-income households have expected benefit ratios of 13 to 17 percent. The maximum benefit occurs approximately at ln*pcx* 4.2, which is less than the mode of the ln*pcx* distribution from figure 3 (it is closer to 5.0). In other words, the biggest losers from a producer price cut are in the lower part of the overall welfare distribution. However, when the high benefit ratios are combined with the density of ln*pcx*, most of the loss for the country as a whole is concentrated on the low side of the middle of the distribution. The loss of welfare would be nonnegligible but would be borne by neither the richest nor poorest households in the country. Over most of the domain of ln*pcx*, the confidence intervals have a width of approximately 0.05, and we can certainly reject the hypothesis that the expected benefit ratio is zero. More important, the bands are sufficiently narrow that it seems reasonable to conclude that the true function has a negative slope when ln*pcx* is greater than 4.2.[10]

As a final exercise, we look at the agricultural sector in isolation, exploring the consequences for cocoa or coffee farmers of a change in both sets of prices. Obviously, from figures 8A and 8B, a greater fraction of households has high benefit ratios. Also the distribution of income is more symmetric (in the ln*pcx* dimension) than was the case for the entire country. The mass of high-benefit households is slightly tilted on the higher side of the mode. This is clearer in the regression function. In figure 8C, the expected benefit starts at around 0.15 (when ln*pcx* = 3) and rises monotonically to 0.23 (when ln*pcx* = 6) at the end of the relevant domain of ln*pcx*. Given the diversity of the sources of household income, farm households will be somewhat sheltered from the direct effects of the price cuts. Furthermore, changing the price level will not be distributionally neutral within the rural area. Households in all parts of the distribution will suffer, but the higher-income households will suffer the most.

We now have a clear picture as to why the Gini coefficients changed as they did in response to a simulated price cut. For the country as a whole, the mass of cocoa and coffee farmers, especially those with high benefit ratios (exposure to

10. At first glance, the confidence bands at ln*pcx* = 7.5 may seem suspiciously narrow. However, this results from the fact that the estimated function actually fits very well here: the predicted and actual benefit ratios are both zero, because there are virtually no cocoa or coffee farmers in this part of the distribution. Thus, the low sampling variance outweighs the low density, resulting in a very narrow confidence interval.
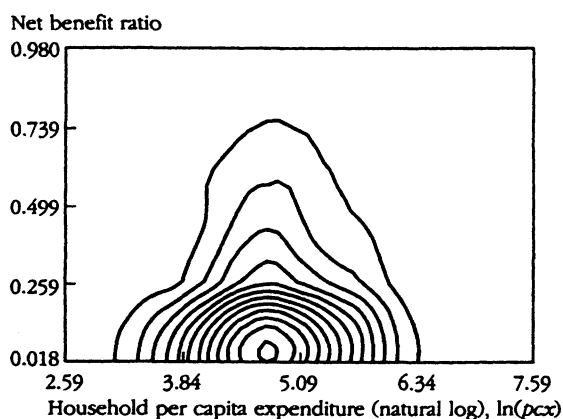
Figure 8. *The Cocoa and Coffee Benefit Ratio for Agricultural Households in Côte d'Ivoire, 1985*

A. Estimate of the Joint Density of the Benefit Ratio and Household per Capita Expenditure
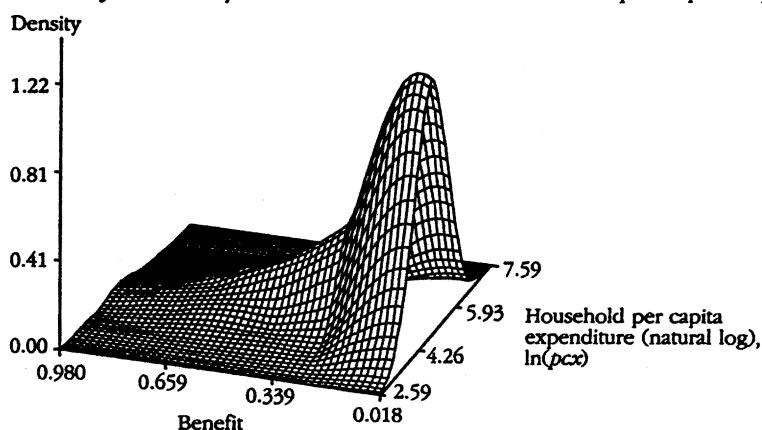
Note: Bandwidth = 0.7.

B. Estimates of the Contours of the Joint Density of the Benefit Ratio and Household per Capita Expenditure

Note: Bandwidth = 0.7. The contour lines each represent increments of 0.1 units of density.

C. The Average Benefit Ratio Estimated as a Function of Household per Capita Expenditure
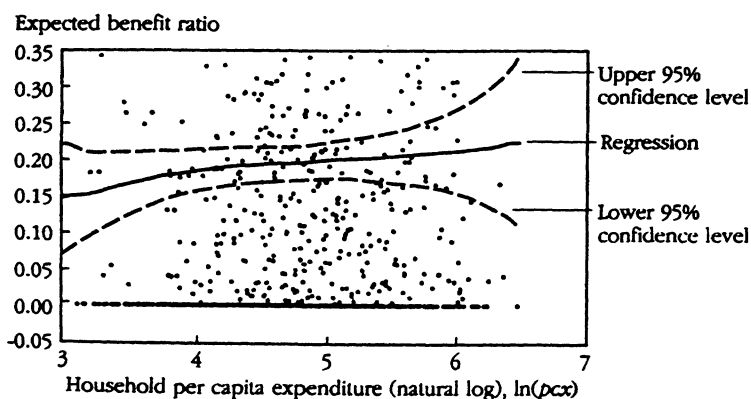
Note: Bandwidth = 1.0. The scatter points represent raw data.

cocoa and coffee prices), are slightly below the median of ln*pcx*. Thus, the price increase would decrease inequality, and a price decrease would increase inequality. In contrast, for agricultural households, although cocoa and coffee farmers are slightly better off, given the concentration of both farmers and the benefit ratio in the middle of the income distribution, the Gini changes only slightly, with price policy changes that hurt cocoa and coffee farmers slightly decreasing inequality.

It is worth reiterating some caveats. First, the metayeurs, who are underrepresented in the sample, are low-income households with a high benefit ratio. Second, we do not observe where government revenue flows, especially between rural and urban areas. If most government expenditure is in the urban areas, then tax and revenue effects will hurt the rural households even more. Third, the "general equilibrium" effects of the price cuts on the value of other sources of income are unpredictable. Probably, however, the entire rural area would see a decline in returns to these other activities. Thus, the level of benefits estimated here may be too low. There is no reason to believe, however, that this will vary across the ln*pcx* distribution. It is possible that low-income households, with a larger share of home-produced food, may be less affected by these secondary price changes.

Up to now the analysis has charted the distributional consequences, based on current benefit ratios, of a change in cocoa and coffee prices. But there are dynamic features to tree-crop supply, particularly features related to the age structure of the trees. We might be missing possible welfare effects by ignoring these dynamics. Farmers who have many young trees will have low benefit ratios (because their trees produce no output), so these farmers appear unscathed in the welfare analysis. Yet, these farmers will suffer "capital losses" and, although it is difficult to measure these potential losses with any precision, we can identify which households in the income distribution face the losses. Furthermore, if the crop supply elasticities are low, and the price cuts are permanent, then looking at the tree distribution suggests the longer-run consequences of the price cuts. As before, given the likely adjustment of farmers to the new prices, these welfare estimates will overstate the actual losses.

Figures 9 and 10 present regression estimates for the number of hectares of cocoa and coffee trees, respectively, held in each age category from a sample restricted to agricultural households. Not surprisingly, the holdings of both young and mature trees increase in size as one moves up the income distribution: richer households have more trees and are somewhat more exposed to "capital losses." Unlike the currently productive or young trees, however, there is very little relation between the number of old trees and ln*pcx*. With this pattern of tree ownership, the distributional results of figure 8C are likely to be maintained for some time. The current capital stock is distributed in a way that perpetuates the higher benefit ratios that were associated with the richer agricultural households.

The results are different for coffee trees. There is very little relation between the size of holdings of both young and old coffee trees and welfare, but mature tree holdings are positively associated. However, holdings of these trees also

Figure 9. *Household per Capita Expenditure and Ownership of Young, Mature, and Old Cocoa Trees in Côte d'Ivoire, 1985*

Hectares

Household per capita expenditure (natural log), ln(*pcx*)
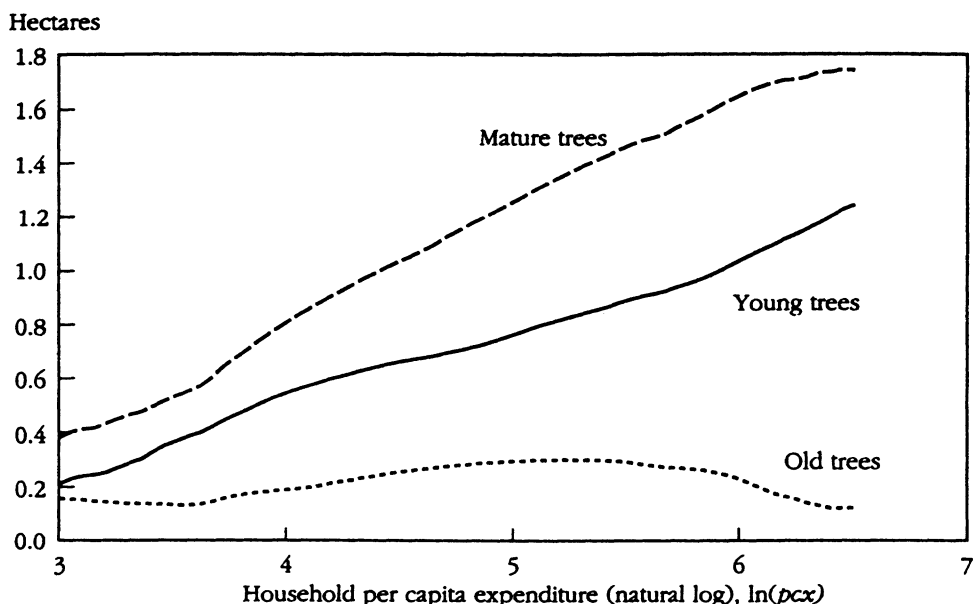
Note: Bandwidth = 1.0.

drop significantly before we reach the high end of the domain of ln*pcx*. Until coffee trees get too old, the positive association between the coffee benefit and ln*pcx* should be maintained. However, there is no relation between holdings of young trees and ln*pcx*, so there are not likely to be distributional consequences for this type of capital loss, nor should the positive relation between the coffee benefit and ln*pcx* be as strong as for cocoa. Of course, this result is at least partially explained by the fact that holdings of young coffee trees are very small: very few farmers, at any level of income, have planted coffee trees in the past 10 years. As already seen in table 2, the coffee orchard will continue to age, without the benefit of new trees coming into production. In summary, the richer rural households have the most to lose in terms of lost investment from a permanent decrease in cocoa prices, whereas for coffee, the potential losses are more evenly spread through the welfare distribution. The age structure of the trees thus reinforces our conclusions from the analysis of current benefit ratios.

IV. Conclusions

We used the 1985 LSMS of Côte d'Ivoire to provide data on cocoa and coffee farming and to examine the likely consequences on household welfare of changing the producer prices of cocoa and coffee, and more generally the conse-

Figure 10. *Household per Capita Expenditure and Ownership of Young, Mature, and Old Coffee Trees in Côte d'Ivoire, 1985*

Hectares

Household per capita expenditure (natural log), ln(*pcx*)
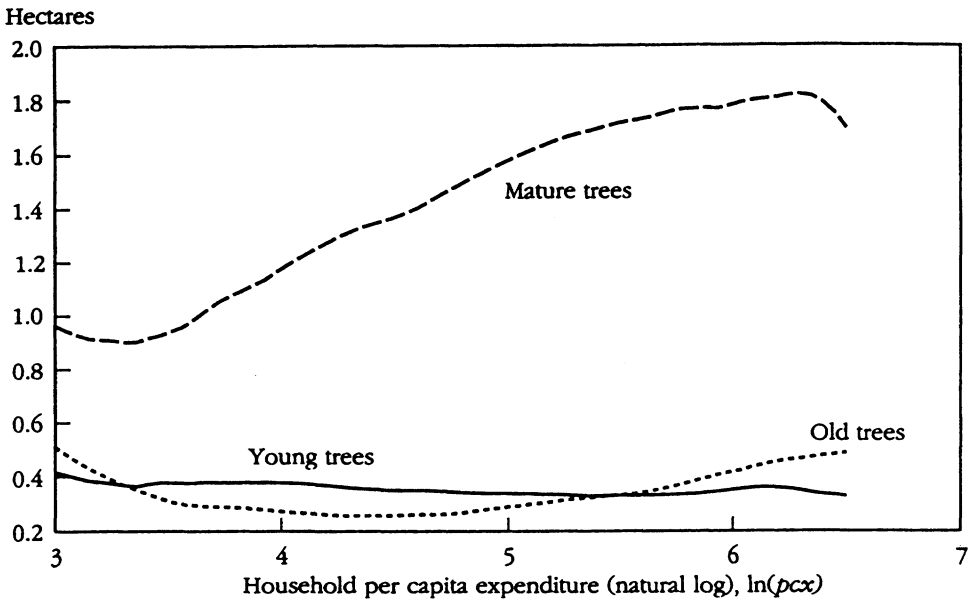
Note: Bandwidth = 1.0.

quences of price policy reform. Using the LSMS data to examine the contribution of cocoa and coffee to household incomes, we showed that cuts in cocoa and coffee prices that have taken place are unlikely to have had a dramatic effect on the distribution of income, essentially because cocoa and coffee farmers are well scattered through the population.

Our analysis does not permit any simple overall recommendation on pricing policy. However, we make three important points. First, it seems clear that distributional considerations should not play a large part in the discussions of pricing reform: cocoa and coffee farmers are typically smallholders and look very like other Ivoirien households. The exception here is the regional aspect; there is little or no cocoa and coffee farming in the Savannah, a region where many of Côte d'Ivoire's poorest farmers live. Second, general considerations would suggest that it is desirable for the government to derive a substantial amount of its revenue from taxing cocoa and coffee. Without good evidence on the long-run supply elasticities of the two crops, it is not possible to say whether welfare would be improved or worsened by decreasing the wedge between world and domestic prices. Third, the government has been successful in stabilizing the real value of farmgate prices for both crops. There are arguments both for and against such stabilization, but the main issues are, first, the continued feasibility of the policy at a time of low world prices and an overvalued currency; second, the ability of the government to manage its revenue in the face of revenue fluctuations; and, third, the determination of the relative cocoa and coffee prices.

REFERENCES

Budd, John. 1993. "Changing Food Prices and Rural Welfare: A Nonparametric Examination of the Côte d'Ivoire." *Economic Development and Cultural Change* 41 (3, April): 587–603.

Deaton, Angus. 1989a. "Household Survey Data and Pricing Policies in Developing Countries." *The World Bank Economic Review* 3(2):183–210.

———. 1989b. "Rice Prices and Income Distribution in Thailand: A Non-Parametric Analysis." *The Economic Journal* 99(supplement):S1–S37.

Deaton, Angus, and Dwayne Benjamin. 1988. *The Living Standards Survey and Price Policy Reform: A Study of Cocoa and Coffee Production in Côte d'Ivoire.* LSMS Working Paper 44. Washington D.C.: World Bank.

Gbetibouo, Mathurin, and Christopher Delgado. 1984. "Lessons and Constraints of Export Crop–Led Growth: Cocoa in Ivory Coast." In William Zartman and Christopher Delgado, *The Political Economy of Ivory Coast.* New York: Praeger.

Glewwe, Paul. 1988. *The Distribution of Welfare in Côte d'Ivoire in 1985.* LSMS Working Paper 29. Washington D.C.: World Bank.

———. 1990. *Investigating the Determinants of Household Welfare in Côte d'Ivoire.* LSMS Working Paper 71. Washington D.C.: World Bank.

Härdle, Wolfgang. 1990. *Applied Nonparametric Regression.* New York: Cambridge University Press.

Hecht, R. M. 1983. "Ivory Coast Economic 'Miracle': What Benefits for Peasant Farmers?" *Journal of Modern African Studies* 21(1, March):25–53.

Hiltzik, Michael. 1989. "Cocoa Turns Bitter for Ivory Coast: Felix Houphouet-Boigny's Boycott of World Market Fails." *Los Angeles Times* (July 12, 1989) section 1, p. 1.

International Monetary Fund. Various issues. *International Financial Statistics.* Washington, D.C.

Kozel, Valerie. 1990. *The Composition and Distribution of Welfare in Côte d'Ivoire.* LSMS Working Paper 68. Washington, D.C.: World Bank.

Newbery, D. M. G. 1987. "Agricultural Taxation: The Main Issues." In D. M. G. Newbery and N. H. Stern, *The Theory of Taxation for Developing Countries.* New York: Oxford University Press.

Silverman, B. W. 1986. *Density Estimation for Statistics and Data Analysis.* New York: Chapman and Hall.

Stern, N. H. 1987. "Aspects of the General Theory of Tax Reform." In D. M. G. Newbery and N. H. Stern, *The Theory of Taxation for Developing Countries.* New York: Oxford University Press.

Trivedi, P. K., and Takamasa Akiyama. 1992. "A Framework for Evaluating the Impact of Pricing Policies for Cocoa and Coffee in Côte d'Ivoire." *The World Bank Economic Review* 6(2):307–30.