Joseph H. Haslag
Senior Economist
Federal Reserve Bank of Dallas

Lori L. Taylor
Senior Economist
Federal Reserve Bank of Dallas

A Look at Long-Term Developments in the Distribution of Income

Strong economic growth in the United States during the last half of the 1980s did not translate into economic gains for all income groups. Poverty rates, for example, remained higher than those observed in the 1970s.[1] To paraphrase the most common findings, the rich got substantially richer during the 1980s, while the poor may have gotten poorer.

A trend toward greater income inequality can be cause for concern. Most Americans would not consider it desirable if, over time, all our society's resources became concentrated in the hands of a small group of individuals. On the other hand, few Americans would desire a perfectly equal distribution of income because income equality implies, among other things, that people who are college educated earn exactly the same income as people who are high school dropouts. If everyone earned the same income, there would be little incentive for people to work harder, become better educated, or find better, more efficient methods of production. Thus, the reasons underlying a trend toward greater income inequality are at least as important for policy analysis as the level of income inequality.

We set out to investigate how and why the distribution of income has changed over time. We find that the distribution of income has been becoming more unequal since the early 1950s, making what occurred in the 1980s a continuation of a longer-running trend. We also find that the distribution of income gains over the past dozen years is close to its historical average. Finally, we examine how rising income inequality relates to changes in the economy's demographic, business-cycle, and policy characteristics. We find that factors outside of direct policy control, such as the age and education profiles of the population, the gender composition of the labor force, and (mostly) inertia in income inequality, explain the lion's share of the forecast error variance. Policy variables, such as transfer payments and tax rates, account for only 15 percent of the variation in prediction errors.

What has happened to the distribution of income?

Our description of changes in the distribution of income proceeds in two parts. In the first part, we divide the population into five equal-sized subgroups, or quintiles, and examine each subgroup's gains from income growth.[2] In the second

We wish to thank Zsolt Besci, Stephen P.A. Brown, Christopher Carroll, and Keith R. Phillips for their helpful comments and suggestions, and Anne E. King and Adrienne C. Slack for their research assistance.

[1] *A change in how poverty rates were calculated means that poverty rates before 1975 are not comparable to those since 1975. Poverty rates stayed below 10 percent during the period 1975–80. During the 1980s, poverty rates climbed and then fell, staying above the 10-percent threshold. While economic growth appears to have roughly coincided with the declines in poverty rates, growth failed to lift enough people out of poverty to reduce the poverty rates below 10 percent.*

[2] *Another issue arises because we use tax returns as our data source. People do not have to file tax returns if their income levels are too low. Consequently, the sample we use is truncated in the sense that the lowest paid people are omitted.*

part, we look at developments in an aggregate measure of income inequality known as the Theil entropy index. The Theil index measures the degree of income inequality across the entire population in one number.

The distribution of income gains. Much of the recent attention to the issue of income inequality has focused not on the distribution of income but on how the gains from income growth were distributed across different income strata. According to work by Paul R. Krugman related in a memorandum from the Congressional Budget Office (1992), the top-paid 20 percent of the population received 94 percent of the gains in after-tax income between 1977 and 1989. In contrast, a 1992 U.S. Treasury report shows that people who were among the richest 1 percent of American taxpayers in 1979 received only 11.3 percent of the total gains in income during the 1980s (Sylvia Nasar 1992). In the *New York Times*, Nasar quotes Isabel V. Sawhill's finding that people who were in the top-paid 20 percent of the population in 1977 saw their incomes decline 11 percent over the next decade, while people who had climbed into the top category by 1986 had experienced, on average, a 65-percent increase in income.

To contribute to these discussions, we calculate summary statistics for the proportion of gains from income growth received by each population subgroup over the recent twelve-year period and compare those results with what occurred over the entire 1952–89 sample.[3] We find that nearly 60 percent of the gains in adjusted gross income during the period 1977–89 accrued to the top income quintile.

Analysts can reach such strikingly dissimilar conclusions because analyses of the distribution of income are very sensitive to the definitions of income, time horizon, and population with which the analyst works. For example, there are many possible definitions of income. One can examine the distribution of total income before taxes and transfers, total income after taxes but before transfers (such as Aid to Families with Dependent Children), total income after taxes and transfers, wage income, and still other variations. It is not too surprising that one finds different results when comparing, say, the income distribution of individuals with the income distribution of households, or the income distribution of taxpayers with the income distribution of all persons.

We focus on adjusted gross incomes rather than after-tax incomes because reliable data on both taxes and transfers are not available. We consider misleading any estimates of the income distribution after taxes but before transfers because they illustrate only part of the government's redistributive activities. In our opinion, one should either analyze the distribution of total factor incomes before taxes and transfers, which indicates the distribution of market-based claims on society's resources, or the distribution of income after both taxes and transfers, which indicates the distribution of purchasing power. Analyses of after-tax incomes that do not include information on transfers are neither fish nor fowl and are very problematic to interpret.

Analysts can also reach different conclusions from one another when they use different measurement techniques. In Krugman's analysis, the share of income gains accruing to population group i is represented as

$$(1) \qquad P_i = \Delta\mu_i \,/\, n\Delta\mu,$$

where μ_i is the change in average income for income group i, μ is the change in average income for the total population, and n is the number of equal-sized income groups. One can interpret equation 1 as the ratio of the weighted-average gain of a specified group to the average gain of the population as a whole. The equation highlights changes over time in the average income of a quintile.

The Council of Economic Advisers (CEA) analyzes growth in income by quintile for the 1992 *Economic Report of the President*. We com-

[3] *Annual adjusted-gross-income data for this study are obtained from various issues of* Statistics of Income, *published by the United States Internal Revenue Service (IRS). For our analysis, we divide the population into quintiles, as follows: the IRS organizes tax returns by adjusted gross income, ranking by groups from lowest paid to highest paid. By dividing the total number of returns by five, we obtain the number of returns for each quintile. Thus, between any two periods the change in adjusted gross income earned by each quintile necessarily equals the aggregate change in adjusted gross income.*

bine the CEA and Krugman approaches to represent the share of income gains earned by the *i*th group of the population as

$$(2) \qquad \hat{P}_i = \Delta Y_i \,/\, \Delta Y,$$

where Y_i is the change in income received by the *i*th group, and Y is the change in income for society as a whole. Equation 2, therefore, is the ratio of income gains received by the *i*th group to the income gains received by the population as a whole. Note that the Krugman and CEA-based measures yield identical results only when the population size is constant.

The following example illustrates how the interaction of population and income growth can affect the distribution of income gains as measured by the Krugman approach and the CEA-based approach. Suppose the economy has two workers with annual incomes of $30,000 and $20,000, respectively. The next year, these same workers earn $40,000 and $30,000, respectively, and two new workers obtain jobs and earn $20,000 each. Total income increased by $60,000. Average income among wage earners increased by $2,500 (from $25,000 to $27,500), while average income for the top half of the distribution increased by $5,000 (from $30,000 to $35,000). According to Krugman's original calculation, the top half of the distribution accounted for 100 percent of the gains from economic growth—$5,000/(2 × $2,500). Average income for the bottom half of the distribution did not change, indicating that the bottom half accounted for zero percent of the gains from economic growth using equation 1. Therefore, although each of the four participants in this hypothetical society earned more in the second year than they did in the first, the Krugman measure would indicate that all of the income gains accrued to the top half of the distribution.[4]

Using the CEA-based technique, the interpretation is somewhat different. The top half of the distribution received $40,000 more ($70,000 – $30,000), while the bottom half of the distribution received $20,000 more ($40,000 – $20,000). Thus, the top half accounted for 66.6 percent of the gains from economic growth using equation 2, while the bottom half received 33.3 percent of the gains.

The intuition behind the difference between Krugman's approach and the CEA-based approach is fairly straightforward. In Krugman's approach, income received by a particular subgroup of the population must grow at a rate faster than population growth. Otherwise, average (per capita) income for that subgroup would not rise, and Krugman's measure would indicate that they failed to share in the income gains. Thus, if the original workers in the example above earned $35,000 and $25,000, respectively, in the second year, and all other aspects of the example remained unchanged, then the average income of the top half of the distribution would remain at $30,000 [($35,000 + $25,000)/2], and the average income of the bottom half of the distribution would remain at $20,000. Krugman's measure would indicate that neither group has experienced any income gains.

By the CEA-based measure (equation 2), the condition for subgroups to share in income gains is that the sum of population growth and income growth for that particular subgroup exceeds zero. Given sufficient population growth, it is possible for the CEA-based measure to indicate that each quintile experienced income gains even when average income was falling for all quintiles. Thus, relative to Krugman's measure, the CEA method requires that a weaker condition is satisfied for any one subgroup to have a positive share in the distribution of income gains.

Both Krugman's and the CEA-based measures do not use longitudinal data. Neither of the statistics follows a particular group of people through time to trace how much of the aggregate gains are distributed to that group. Therefore, when there is substantial income mobility, these measures of the aggregate economy say little about the incomes received by specific individuals. However, they say much about the distribution of possible incomes and, therefore, about individual opportunities. (For a discussion of income mobility in the United States, see the box titled "Trends in Income Mobility.")

Table 1 reports the distribution of changes in income for several periods. Specifically, the table

[4] *Michael Boskin (1992) lays out this example in describing Krugman's distribution of income gains.*

Trends in Income Mobility

From year to year, people can and do move from one income class to another. For periods as long as a decade, the changes in people's income, especially for people in the lowest income group, are remarkable. Table B1 shows movements in the income distribution from 1979 to 1988.[1] The data indicate substantial income mobility, particularly among the lower income groups. Only 65 percent of the people who were in the top-paid 20 percent of the population in 1979 were still in the top-paid 20 percent in 1988. Two-thirds of the people in the middle income quintile changed classification over that ten-year period, while more than 85 percent of the people in the lowest income quintile changed income classifications. More than 17 percent of the people in the lowest income category in 1979 had climbed into the highest income category by 1988. Except for people in the highest income category (who, by definition, could not improve), those who changed income quintiles were more likely to move up than down.

One caveat to interpreting this evidence is that the study follows individuals who filed IRS returns in each of the ten years from 1979 through 1988. Accordingly, those who earned such low amounts that they did not have to file returns in any of the ten years were omitted from the sample. These people may well be permanently poor, making the upward mobility evidence less strong. Further, mobility out of the lowest income categories may be overstated because the low income groups in 1979 undoubtedly include students and part-time workers who became better compensated as they accumulated education and experience.

[1] See Joel Slemrod (1992) for evidence on the upward bias imparted to income inequality when looking at year-to-year income changes. Slemrod calculates the average income for each taxpayer over the seven-year period from 1979 to 1985. Slemrod refers to this approach as a time exposure. Compared with the snapshots of the income inequality over the same time period, the time-exposure Gini coefficient is roughly 7 percent lower, suggesting that income inequality declines somewhat as the time horizon lengthens. The findings are consistent with changes in income from period to period that are smoothed over when one uses income measured over several years, instead of capturing jumps in income that occur in any given year.

Table B1
Changes in Income Quintiles, 1979 and 1988

	Status in 1988				
Status in 1979	**Top-paid 20%**	**Next highest paid 20%**	**Middle 20%**	**Next lowest paid 20%**	**Lowest paid 20%**
Top 20%	64.7	20.3	9.4	4.4	1.1
Next highest paid 20%	35.4	37.5	14.8	9.3	3.1
Middle 20%	15.0	32.3	33.0	14.0	5.7
Next lowest paid 20%	11.1	19.5	29.6	29.0	10.9
Lowest paid 20%	17.7	25.3	25.0	20.7	14.2

Table 1
Percentage of Income Gains Distributed Among the Five Population Quintiles

Using the CEA-based method:

Period	Q1	Q2	Q3	Q4	Q5
1977–89	2.7	6.4	12.1	21.8	57.1
1980–85	3.3	9.8	11.0	23.6	52.4
1985–89	1.0	2.3	11.4	17.3	68.0

Using Krugman's method:

Period	Q1	Q2	Q3	Q4	Q5
1977–89	2.8	5.8	11.1	20.8	59.5
1980–85	3.3	10.0	10.5	23.3	53.0
1985–89	1.8	0.0	10.7	14.7	74.5

reports the changes in adjusted gross income received by each quintile for the periods 1977–89, 1980–85, and 1985–89.[5] The top half reports the distribution of gains in nominal, pre-tax, and transfer income using the CEA-based method (equation 2). The bottom half of the table reports the distribution of income gains from the same data using the Krugman method (equation 1).

Somewhat surprisingly, the results using the Krugman method are quite similar to those using the CEA-based method. The Krugman method indicates that a slightly higher percentage of income gains is going to the top quintile than indicated by the CEA-based method, but this difference does not change the implication that the top quintile reaped the majority of the income gains. Table 1 shows that over the period 1977–89, about 60 percent of the gains in factor income (income before taxes and transfers) went to the top-paid quintile.

Another question is how income gains are distributed across different subperiods. For example, were the 1980–85 or 1985–89 periods substantially different in terms of how income gains were distributed? The evidence presented in Table 1 suggests that the 1985–89 period saw gains going more to the highest paid quintiles and less to the lower-paid quintiles. For example, during the 1980–85 period the two lowest paid quintiles accounted for slightly more than 13 percent of the income gains, while during the 1985–89 period the same two quintiles accounted for less than 4 percent of the income gains. The share of income growth received by the middle-paid quintile was virtually unchanged from the 1980–85 period to the 1985–89 period, while the share of income growth received by the second highest paid quintile declined more than 25 percent. The declines in the first, second, and fourth quintiles were matched by the increases of the highest paid quintile. Using the CEA-based method, the top-paid quintile accounted for about

[5] *These statistics also represent the distribution of real income gains, assuming that each quintile has the same deflator.*

52 percent of the income gains in the first half of the 1980s, rising to about 68 percent of the gains in the second half.

By reporting the historical averages received by each quintile, one can see how these recent time periods compare with the entire sample. Using the CEA-based method, we calculate the share of income gains received by each quintile annually for the period 1952–89. As the evidence in Table 2 illustrates, on average the two lowest paid quintiles received about 4.5 percent of the real income gains over the 1952–89 sample. The highest paid quintile averaged about 58 percent of the income gains, while the middle-paid and second highest paid quintiles averaged almost 11 percent and 27 percent of the gains, respectively. The evidence, therefore, suggests that what happened during the 1977–89 period is not that different from what happened during the postwar period.

Note that the standard deviations are substantially different across the five quintiles. As Table 2 shows, the standard deviation is 5.2 percentage points for the lowest paid quintile. There is much greater variability in the highest paid (44.7), the second highest paid (25.0), and the second lowest paid (31.8) quintiles. Thus, the standard deviation for the lowest paid quintile is only about one-third the size of the standard deviation for the other quintiles. This evidence suggests that the lowest paid quintile receives a fairly steady proportion of the income gains across time, especially when compared with the proportions received by the other four quintiles.

In short, the IRS data suggest that the top-paid quintile did account for most of the gains in factor income during the period 1977–89. However, a substantially smaller proportion went to the top-paid quintile than was reported in Krugman's study of after-tax (but before transfer) incomes. The proportion of income gains accruing to the top-paid group increased during the latter half of the 1980s. However, the tendency toward increasing income inequality began in the 1950s, and the 1977–89 period appears to be well within the variability observed historically.

Table 2
Summary Statistics of the Proportion of Real Income Gains for each Quintile, 1952–89 (CEA-based method)

Quintile	Mean	Standard Deviation
1	2.9	5.2
2	1.5	31.8
3	10.9	13.8
4	26.7	25.0
5	58.0	44.7

The Theil index. At a more aggregate level, one can measure income inequality with the Theil entropy index:

$$(3) \qquad T(y,n) = (1/n)\sum_{j=1}^{n}(y_j/\mu)\ \ln(y_j/\mu),$$

where n is the number of equal-sized population groups, y denotes income for population group j, and μ is average income (that is, $\mu = \sum_{j=1}^{n} y_j/n$).[6]

The Theil index is defined so that it takes on values greater than or equal to zero and increases as income inequality increases. To illustrate this point, consider the limiting case in which income is evenly divided among all people in the economy. With $y_j/\mu = 1$ for $j = 1, 2, ..., n$, then $\ln(y_j/\mu) = 0$ in equation 3 and the Theil index equals zero. Another attractive feature of the Theil index is that a transfer of income from a high income person to a low income person will cause the Theil index to fall.

Figure 1 plots estimates of the Theil entropy measure over the period 1952–89. The plot indicates an upward trend in the Theil index since 1952.

[6] See Anthony Shorrocks (1980), John Bishop, John Formby, and Paul Thistle (1989), and Daniel J. Slottje (1989) for the set of properties that an income inequality measure possesses. This version of the Theil index differs from the population-weighted version used by Keith R. Phillips (1992).

Figure 1
Theil Index for Adjusted Gross Income, 1952–89

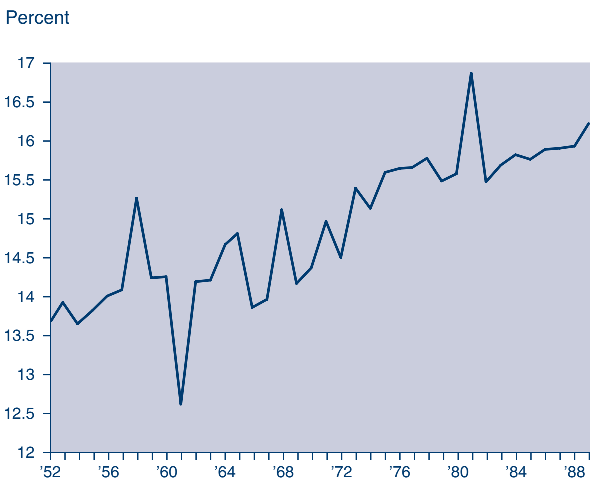

SOURCE OF PRIMARY DATA:
U.S. Internal Revenue Service, *Statistics of Income*.

Regressing the Theil index against time, one finds more concrete support for the notion that the Theil index has been increasing. The regression coefficient on time suggests that the Theil index has, on average, increased 0.4 percentage points each year. Figure 1 also shows a decline in volatility in the Theil index. Beginning around 1980, the Theil index appears to follow a less variable path compared with the swings observed in the pre-1980 sample. Thus, the two apparent inferences drawn from plotting the Theil index are that income inequality has been increasing over the past forty years and that volatility in the income inequality measure has decreased somewhat during the past ten years.[7]

Our conclusions about the time path of income inequality differ somewhat from other work in the field. For example, Barry Bluestone and Bennett Harrison (1988) find that the percentage of workers falling into the low-wage stratum follows a U-turn: the statistic falls in the 1950s and 1960s, reaches its trough in 1970, fluctuates in a narrow band around the trough value, and then increases beginning in 1979. The results in Figure 1 suggest that the trend toward greater income inequality may have started even earlier than Bluestone and Harrison identify. As such, our findings support those of Peter Henle and Paul Ryscavage (1980), who find that inequality in wages has been increasing since the late 1950s. A note of caution in comparing the results: we are using adjusted gross income from the IRS, while Bluestone and Harrison and Henle and Ryscavage are using wage data from the Current Population Survey. Hence, the results are not directly comparable; our analysis does not overturn Bluestone and Harrison's.

There is evidence on income inequality for broader measures of income. Using the Current Population Survey, Daniel J. Slottje (1989) calculates the Theil index over the period 1947–84. The income measure used here is before taxes but after transfers. According to Slottje, the Theil index falls to its lowest value in the late 1960s, rising thereafter for the remainder of the sample period. The implication is that a U-turn is present in comprehensive income measure employed in the Current Population Survey data. Again, the income concepts used in Slottje are not the same as ours because they include transfers and the data sets are different. Our finding, however, raises a question about when (and if) the U-turn occurred in a broader income measure, when one looks at income before taxes and transfers. It also suggests that the U-turn in the inequality of after-transfer income may reflect changes in the distribution of transfer policy rather than factor incomes.

What determines the distribution of income?

In this section, we turn from describing what happened to income inequality to examining why the distribution of income changed. The factors affecting income inequality that we consider fall into three categories: demographics, economic conditions, and fiscal policy. Our demographic data are

[7] *The Theil index invariably measures changes in the income distribution differently from other inequality measures because it weights transfers differently. To check the robustness of our inferences, we also calculated another aggregate measure of income inequality—the Gini coefficient. The results are not materially different whether one uses the Gini coefficient or the Theil entropy measure.*

the age and educational composition of the potential labor force, and the female share of the labor force.[8] Following Alan S. Blinder and Howard Y. Esaki (1978), we include some measure of business-cycle conditions and the inflation rate as variables that might explain variation in income inequality. Finally, we use the maximum marginal personal income tax rate and real per capita transfer payments to individuals (defined as the sum of federal, state, and local transfers) as the policy variables. The Theil index is our measure of income inequality.

To determine whether these characteristics can be used to predict changes in the Theil index over time, we estimate a VAR system of eight equations—one for each of our eight variables—using ordinary least squares regression. Each equation estimates contemporaneous values of the variable as a function of two lagged values of itself and two lagged values each of the other seven variables. All the variables except the growth rates for prices and gross domestic product are expressed as first differences (the change in value between period *t* and period *t*–1), and the variables that are bounded by zero and one (such as the percentage of women in the labor force) are logistically transformed.[9]

Table 3
Tests of Exclusions Restrictions in Theil Index Equation

Variable	*F* statistic	*p* value
AGE	.15	.86
ED. ATTAINMENT	1.61	.23
FEM SHARE	1.02	.38
TAX	1.67	.22
GDP	.87	.44
INFL	.78	.47
TRANS	.61	.55

Table 3 reports the results of exclusion restrictions for each variable in the equation in which the Theil index is the dependent variable. The tested hypothesis is that the coefficients on lagged values of the variable also are jointly equal to zero. The interpretation of the test results, then, is whether changes in the variable help to predict changes in the Theil index. The *F* statistics are small in each case, which is consistent with the notion that none of these variables, except past values of the Theil index itself, helps to predict changes in the Theil index. However, all of the variables together explain almost 50 percent of the variation in the Theil index. Correlations among some of the variables may be introducing multi-collinearity.[10]

Although a variable may have insignificant explanatory power in the VAR regressions, its compound influence over time may still be considerable. Table 4 reports how much of the two-, five-, and ten-step-ahead forecast error variances result from innovations in the variables.[11] The evidence strongly suggests that innovations in the Theil index account for most of the forecast error variance. Indeed, 54 percent of the ten-step-ahead forecast error variance results from innovations in the Theil. The analysis suggests that the Theil index displays

[8] *The age composition of the labor force is defined as the ratio of people between 16 and 25 to those between 16 and 65. The education composition of the potential labor force is measured by the percentage of the population over age 25 that has graduated from college.*

[9] *David Hendry and Jean–Francois Richard (1982) argue that a logistic transformation should be applied to any dependent variable defined over the [0,1] interval. Formally, the transformation redefines the variable as follows:* $ln(x_t/1-x_t)$.

[10] *The issue with multicollinearity is that close correlation between the explanatory variables will result in inflated standard errors. The upshot of this is that test statistics are downwardly biased when multicollinearity is present.*

[11] *We apply a Choleski decomposition with the following ordering for the recursive system:* AGE, TAX, ED. ATTAINMENT, THEIL, FEM SHARE, GDP, INFL, *and finally* TRANS. *See Christopher Sims (1980) for a more complete discussion of the Choleski decomposition applied to VARs.*

Table 4
Proportion of Forecast Errors for the Theil Index*

	Innovation to							
Step-Ahead	***AGE***	***FEM SHARE***	***ED. ATTAINMENT***	**Theil**	***TAX***	***GDP***	***INFL***	***TRANS***
Two	5.6	1.1	5.7	72.0	10.3	0.7	1.6	3.0
Five	5.1	3.0	16.3	54.9	8.9	4.3	2.0	5.6
Ten	5.5	3.4	16.3	53.6	8.9	4.3	2.2	5.8

*Proportions may not sum to 100 due to rounding error.

lots of persistence, accounting for movements in the Theil over time.[12]

The other factors explain the rest of the ten-step-ahead forecast error variance. Together, the age composition and educational attainment of the population, and the female share of the labor force, account for slightly more than 25 percent. Fiscal policy variables account for about 15 percent of the forecast error variance, while inflation and output growth account for 2.2 and 4.3 percent, respectively. The evidence, therefore, suggests that 85 percent of the variation in income inequality arises from factors outside direct policy control.

Sheldon Danziger, Robert Haveman, and Peter Gottschalk (1981) also examine the role that the U.S. transfer payment system has on income inequality. They use both welfare payments and Social Security as their definition of the transfer payment system. Comparing the distribution of total factor income with the distribution of total factor income plus transfer payments (but before taxes), Danziger, Haveman, and Gottschalk conclude that income inequality (as measured by the Gini coefficient) is 19 percent lower after transfers than it is before them. The Danziger, Haveman, and Gottschalk comparison excludes the complex general equilibrium effects that transfer payments have.

Using a narrower definition of transfer payments and assuming that a recursive model represents the structure, we find that transfer payments have much less explanatory power. In our analysis, transfer payments explain less than 6 percent of the ten-step-ahead forecast error. In addition, the impulse response functions suggest that increases in transfer payments are associated with *increases* in factor income inequality.

Summary and conclusions

In this article, we examine developments in the income distribution over almost four decades and the relative contributions of demographic, policy, and economic conditions toward explaining these movements. Our analysis indicates that recent developments in income inequality and the distribution of gains from income growth are not much different from historical norms. We find that the top-earning 20 percent of the population reaped a disproportionate share of the income gains during the 1980s. However, we find that the top income

[12] *Steven N. Durlauf (1991) examines the evolution of income inequality and finds that persistent income inequality can develop even with identical starting conditions. Education and neighborhood effects reinforce one another to stratify the economy, imparting substantial persistence in income inequality.*

group has been receiving similarly large shares of the income gains for the past forty years. Further, using the Theil index to measure income inequality, we find that income inequality increased over the period 1952–89.

In addition, we look at the relationship between various factors associated with changes in the distribution of income. In particular, we ask which factors explain movements in income inequality over time. We investigate demographic factors, economic conditions, and policy variables. The evidence reported in this article suggests that there is a great deal of persistence in income inequality. Most of the forecast error variance in the income inequality measure is explained by innovations in the inequality measure itself. Fiscal policy actions—measured as the maximum marginal tax rate and per capita transfers—account for only about 15 percent of the variation in the forecast errors.

Overall, the evidence presented in this article focuses on developments in income inequality *over time*. In doing so, the contribution is largely in describing how income inequality has evolved, excluding some factors that are widely believed to affect income inequality. The aim of future research is to formulate theories about what determines income inequality.

Appendix

Data Definitions

Variable	Source
GDP	Board of Governors of the Federal Reserve System FAME dataset
CPI	Bureau of Labor Statistics
AGE	U.S. Department of Commerce, Bureau of Census; population between 15 and 25 divided by the population over 16, less those over 65 (Citibase data)
FEM SHARE	U.S. Department of Commerce, Bureau of Census; total number of employed women divided by the total labor force (Citibase data)
TRANS	National income and product accounts, transfer payments to individuals paid by federal, state, and local governments (Citibase data). This variable is deflated with the fixed-weight GDP deflator.
TAX	Internal Revenue Service, *Statistics of Income*, various issues.
ED. ATTAINMENT	Current Population Survey, series P-60; percent of population over 25 with four or more years of college.

References

Bishop, John, John Formby, and Paul Thistle (1989), "Statistical Inference, Income Distributions and Social Welfare," in *Research on Income Inequality,* vol. 1, Daniel J. Slottje, ed. (Greenwich, Conn.: JAI Press, Inc.): 49–81.

Blinder, Alan S., and Howard Y. Esaki (1978), "Macroeconomic Activity and Income Distribution in the Postwar United States," *Review of Economics and Statistics* 60 (April): 604–07.

Bluestone, Barry, and Bennett Harrison (1988), "The Growth of Low-Wage Employment, 1963–86," *American Economic Review* 78 (May): 124–28.

Boskin, Michael (1992), "Letters to the Editor," *Wall Street Journal*, July 3, Southwest edition.

Congressional Budget Office (1992), "Measuring the Distribution of Income Gains," CBO Staff Memorandum (Washington, D.C.: Congressional Budget Office, March).

Council of Economic Advisers (1992), *Economic Report of the President* (Washington, D.C.: U.S. Government Printing Office, February).

Danziger, Sheldon, Robert Haveman, and Peter Gottschalk (1981), "How Income Transfer Programs Affect Work, Savings, and the Income Distribution," *Journal of Economic Literature* 29 (September): 975–1028.

Durlauf, Steven N. (1991), "Persistent Income Inequality I: Human Capital Formation, Neighborhood Effects and the Emergence of Poverty" (Unpublished manuscript).

Hendry, David, and Jean–Francois Richard (1982), "On the Formulation of Empirical Models in Dynamic Econometrics," *Journal of Econometrics* 20 (October): 3–33.

Henle, Peter, and Paul Ryscavage (1980), "The Distribution of Earned Income Among Men and Women, 1958–1977," *Monthly Labor Review* 103 (April): 3–10.

Mankiw, N. Gregory (1992), *Macroeconomics* (New York: Worth Publishers).

Nasar, Sylvia (1992), "The Rich Get Richer, but the Question Is by How Much?" *New York Times*, July 20, C1.

Phillips, Keith R. (1992), "Regional Wage Divergence and National Wage Inequality," Federal Reserve Bank of Dallas *Economic Review*, fourth quarter, 31–44.

Shorrocks, Anthony (1980), "The Class of Additively Decomposable Inequality Measures," *Econometrica* 48 (April): 613–26.

Sims, Christopher (1980), "Macroeconomics and Reality," *Econometrica* 48 (January): 1–48.

Slemrod, Joel (1992), "Taxation and Inequality: A Time-Exposure Perspective," NBER Working Paper Series, no. 3999 (Cambridge, Mass.: National Bureau of Economic Research).

Slottje, Daniel J. (1989), *The Structure of Earnings and the Measurement of Income Inequality in the U.S.* (Amsterdam: North Holland).

Stiglitz, J. E. (1969), "Distribution of Income and Wealth Among Individuals," *Econometrica* 37 (July): 382–97.