

NBER WORKING PAPER SERIES

ON THE WELFARE CONSEQUENCES OF THE
INCREASE IN INEQUALITY IN THE UNITED STATES

Dirk Krueger
Fabrizio Perri

Working Paper 9993
<http://www.nber.org/papers/w9993>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2003

Forthcoming in the 18th NBER Macroeconomics Annual. We thank Steven Davis , Mark Gertler, Ken Rogoff, Kjetil Storesletten and seminar participants at UCLA and the NBER Macroannual conference for their useful comments. Thanks the NSF for financial support under grant SE S 0004376. The views expressed herein are those of the authors and are not necessarily those of the National Bureau of Economic Research.

©2003 by Dirk Krueger and Fabrizio Perri. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

On the Welfare Consequences of the Increase in Inequality in the United States
Dirk Krueger and Fabrizio Perri
NBER Working Paper No. 9993
September 2003
JEL No. D12, D31, D14, D63

ABSTRACT

We investigate the welfare consequences of the stark increase in wage and earnings inequality in the US over the last 30 years. Our data stems from the Consumer Expenditure Survey, which is the only US data set that contains information on wages, hours worked, earnings and consumption for the same cross section of US households. We first document that, while the cross-sectional variation in wages and disposable earnings has significantly increased, the overall dispersion in consumption has not significantly changed. We also show that households at the bottom of the consumption distribution have increased their working hours to a larger extent than the rest of the population. In order to assess the magnitude and the incidence of the welfare consequences of these trends we estimate stochastic processes for earnings, consumption and leisure that are consistent with observed cross-sectional variability (both within and between education groups) and with household mobility patterns. In a standard lifetime utility framework, using consumption and leisure processes, as opposed to earnings processes, results in fairly robust estimates of these consequences. We find that about 60 percent of US households face welfare losses and that the size of these losses ranges from one to six percent of lifetime consumption for different groups.

David Krueger
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104
and NBER
dkrueger@econ.upenn.edu

Fabrizio Perri
Department of Economics
New York University
Stern School of Business
New York, NY 10012-1126
and NBER
fperri@stern.nyu.edu

1. Introduction

The sharp increase in wage and earnings inequality in the US over the last 30 years is a well-documented fact. Katz and Autor (1999) summarize the findings of a large body of empirical research on this topic by concluding that “many researchers using a variety of data sets -including both household and establishment surveys- have found that wage inequality and skill differentials in earnings increased sharply in the United States from the late 1970s to the mid-1990s.” (p. 6). The objective of this paper is to analyze the welfare consequences of this sharp change in the wage, and the associated change in the earnings, distribution.²

Our interest in welfare immediately forces us to look beyond the distribution of current wages. If a households’ economic welfare depends on consumption and leisure enjoyed over that households’ lifetime, as commonly assumed by economists, then an analysis of the welfare consequences of increasing wage inequality has to determine how current wages are related to disposable income, lifetime consumption and hours worked. First, even if current wages perfectly determine lifetime earnings, transfers among extended family members, friends or the government may augment disposable income and thus increase consumption opportunities of families. Second, to the extent that a significant fraction of the variations in wages, and thus income, appears to be due to variations in its transitory component and to the extent that some forms of credit markets are available to households, current wages or current income may not be the appropriate measure of lifetime resources of these households. This suggests that the distribution of consumption is a better indicator of the distribution of welfare. In addition, the distribution of consumption still does not constitute a sufficient statistic for the welfare distribution, since high consumption in the light of low wages may be realized at the expense of long working hours and thus little leisure of the household.

²Our data stretches from 1972 to 2000. When we refer to the increase in inequality over the last thirty years we mean the long-run trend within our sample period. Our data presented later show that inequality has not increased at a uniform pace throughout the last thirty years.

The objective of our analysis is to take all these elements into consideration in evaluating the welfare consequences of the increase in inequality in the US.³

Our analysis is divided into two parts. In the first part we document, using data from the Consumer Expenditure (CE) survey for the years 1972-2000, the evolution of the US cross-sectional variability of individual wages, household total- and disposable earnings, hours worked and consumption. We find a substantial increase in wage-, total earnings- and disposable earnings inequality. Total household earnings inequality increases slightly less than individual wage inequality, suggesting that longer hours might be used to partly compensate for declines in relative wages. We also find that inequality in disposable earnings (which include government taxes and transfers) increases by more than inequality in earnings, suggesting a reduction in the redistributive impact of these public policies. Despite all these developments *consumption* inequality displays a very modest increase.

An important part of our analysis is the decomposition of the increase in cross-sectional inequality in the data into an increase in (estimated as persistent) differences between groups (e.g. college- and high-school educated households) and into an increase in (estimated as less persistent) idiosyncratic differences within each group (e.g. employed and unemployed).

In the second part of the paper we first estimate stochastic processes for household earnings, consumption and hours worked that are consistent with the evolution of the empirical cross-sectional distributions and with one year relative mobility matrices from the CE. Then a standard lifetime utility framework, together with our estimates of the stochastic processes for the relevant variables, is employed to deduce the magnitude of welfare losses from increased inequality. In particular, our analysis focuses on the welfare consequences of two distinct aspects of the increase in inequality: i)

³Obviously a full evaluation of the welfare consequences of inequality is a complex task that depends on a large number of additional economic and social factors not considered here. Also, throughout the paper, we will treat long-run growth trends in consumption and leisure, as orthogonal to changes in inequality. This imply that we will ignore all the effects that changes in inequality might have on these trends.

persistently higher differences between groups reduce relative mean resources of some groups and increase them for others, ii) higher volatility within group increases the risk faced by *all* groups of the population. Both aspects potentially have large welfare impacts; the upshot of our analysis is that (ii) significantly affects the welfare of all groups, *ceteris paribus*, and that i) determines the exact distribution of these welfare effects, reinforcing them for groups that do poorly and mitigating (or even offsetting) them for groups that do better. Our framework allows us to quantify the size of the welfare losses for many education/sex groups of the US population. The estimates obtained using consumption and leisure processes are fairly robust to changes in the risk aversion parameter and are as large as 6% of lifetime consumption for some groups.

The paper is organized as follows. In the next section we briefly relate our study to the existing literature. The main descriptive statistics of the CE wage-, earnings-, hours- and consumption distributions are summarized and interpreted in Section 3. The quantitative welfare analysis based on the empirical findings in that section is contained in Section 4 and Section 5 concludes. Further details about the data used in the main body of the paper are discussed in the appendix.

2. Related Literature

What are the behavioral and welfare consequences of changes in the wage structure? Several strands of the existing literature provide partial answers to this question. First, a sizeable literature, summarized in Blundell and MaCurdy (1999) investigates the behavioral response of labor supply to changes in wages and employment status of the primary earner. Since we explicitly study hours worked by several members of the household, the literature on the added worker effect, which studies the labor supply response of a spouse to the primary workers' job loss or job displacement is of relevance to our work. Whereas most studies find small effects (see e.g. Heckman and MaCurdy (1980) and Cullen and Gruber (2000)), Stephens (2002) argues that, once the labor supply response of spouses to an expected job loss of their partners and to permanently lower wages of partners

following their displacement are taken into account, this response may be quite sizeable.

Second, if changes in the wage structure translate into changes in a households' income process, the question arises how these changes impact consumption. The complete consumption insurance hypothesis has a very strong prediction for the risk sharing behavior among households. Under this hypothesis the ratio of marginal utilities between two households is constant across time and states, even when individuals face idiosyncratic income uncertainty.⁴ If preferences of all individuals are identical and can be represented by a CRRA utility function that is separable across time and between consumption and leisure, then, absent shifts in tastes, individual consumption growth rates move in tandem with aggregate consumption growth rates. And, conditional on aggregate consumption growth rates, individual consumption growth rates would be uncorrelated with individual income growth rates, independent on whether income fluctuations are temporary or permanent, expected or unexpected. As a consequence of full consumption insurance the increase in variability of the idiosyncratic part of income (and if there is consumption insurance between observably distinct groups, even the increase in between-group variability) in itself does not have adverse welfare consequences. This hypothesis has been extensively tested empirically. Examples include Altug and Miller (1990), Mace (1991), Cochrane (1991), Nelson (1994), Townsend (1994), Attanasio and Davis (1996) and Hayashi, Altonji and Kotlikoff (1996), with results that tend to reject the full consumption insurance hypothesis.

Full consumption insurance requires a sophisticated transfer or insurance system between individual households (it can be achieved, e.g., with a full set of Arrow securities that pay out contingent on individual income shocks). In contrast, the second strand of the literature studying the map between income and consumption takes the permanent income hypothesis as point of departure, which envisions a consumer in isolation attempting to *self*-insure against income fluctuations. The

⁴See, e.g., Deaton (1992), Chapter 1.3.

welfare consequences of increased income instability in this paradigm depend crucially on whether the income shocks are temporary or permanent, since transitory shocks can be easily smoothed using (dis-)saving. Hall and Mishkin (1982) decompose income into its transitory and permanent component and investigate whether, in fact, households smooth transitory income shocks to a higher degree than permanent income shocks. Blundell and Preston (1999) use this idea and income and consumption data to infer the extent to which income shocks are permanent. Finally, Heathcote et al. (2003) build a model based on the permanent income hypothesis to assess the welfare consequences of the recent increase in wage inequality. They use model-predicted consumption paths implied by estimated wage processes in their welfare calculations.

Common to both these literatures is that they interpret consumption data through the lens of a particular economic model of financial markets (a complete set of contingent claims for the former, a single uncontingent bond in the latter). Finally, there exists a descriptive literature that does not take a stand on a particular economic model, but rather documents changes in income and consumption distributions (and possibly interprets them). Examples include Cutler and Katz (1991a,b), Johnson and Shipp (1991), Mayer and Jencks (1993), Johnson and Smeeding (1998) and Slesnick (1993, 2001).⁵ In addition, the papers by Attanasio and Davis (1996), Dynarski and Gruber (1997) and Krueger and Perri (2002) provide extensive descriptive analysis of cross-sectional household consumption data, but then go on to interpret and analyze these data from the viewpoint of an underlying theoretical model. Our approach will mostly follow this last descriptive tradition. The only theory we use is a period budget constraint to organize our data and an intertemporal utility function to evaluate the welfare consequences of changes in the wage distribution, in conjunction with observed consumption and leisure choices from the data directly.

Our thought experiment of assessing the welfare consequences of increased wage, income

⁵Even policy circles and the popular press are occupied with the distribution of consumption. See Greenspan (1998) for an example of the former, and the book by Cox and Alm (1999) for an example of the latter.

and consumption inequality using micro-consumption data is very similar in spirit to the exercise conducted by Attanasio and Davis (1996). They quantify the welfare losses implied by incomplete consumption insurance between different education and cohort groups, relative to the complete consumption insurance benchmark. Behind the veil of ignorance (i.e. before knowing what cohort-education group one belongs to) agents have to be compensated by an increase of one to three percent of consumption (at all dates, in all states) for imperfect consumption insurance. In our analysis we study, in addition to imperfect between-group insurance, the welfare implications of within-group consumption variability and we document how the welfare consequences are distributed across the population (that is, we look after the veil of ignorance has been lifted). We find welfare losses of increased consumption inequality of similar magnitude for a large part of the population and conclude with Attanasio and Davis (1996) that these costs are two orders of magnitudes bigger than the costs of business cycles commonly derived in the macroeconomic literature (see, e.g., Lucas, 1987).⁶

3. Descriptive Statistics

Before assessing the welfare consequences of increasing wage and income inequality we first want to document the basic facts from the CE data that link wages to the economic variables that enter the utility function in our welfare analysis, namely consumption and leisure (that is, all available nonsleeping time minus hours worked per member of the household). A comment about our choice of the CE as our data source may be order at this time. For detailed information on US *individual* household consumption, the CE is the only available data set. In addition, since we want to investigate changes in the distribution of welfare associated with changes in the wage and earnings distribution, including the within-group variation, synthetic cohort techniques that

⁶Krusell and Smith (1999), among others, use a model with many heterogenous agents to document the distribution of this cost across different income and wealth classes and find that the cost of business cycles is small for almost all population groups.

make it feasible to combine several different data sets cannot be applied, since these techniques average out all within-group (idiosyncratic) variation in the data.⁷ The CE does include not only information on consumption, but also on hours worked and income, and indirectly on wages. It has a relatively small sample size (an average of 5000 households per quarter) and a short panel dimension (1 year), but it is the only available data set that reports household level observations for all economic variables needed for our study.

In order to organize the CE data we make use of a simple period-by-period budget constraint, which reads, for an arbitrary household i , as

$$c_{it} = y_{it} - s_{it}$$

where c_{it} are expenditures on consumption, y_{it} is income from all sources, and s_{it} is saving.

We divide net income y_{it} of the household into labor income net of taxes and capital income net of taxes. Let wages of household member j be denoted by w_{ijt} , and hours worked by that household member by h_{ijt} . Finally let t_{it} denote direct labor income taxes paid by household i , and by τ_{it} government transfers. Therefore

$$y_{it} = \sum_j w_{ijt}h_{ijt} + \sum_k r_{ikt}a_{ikt} - t_{it} + \tau_{it} = ly_{it} - t_{it} + \tau_{it} + ky_{it}$$

and thus

$$c_{it} = ly_{it} - t_{it} + \tau_{it} + ky_{it} - s_{it}$$

where ly_{it} is labor earnings before taxes and ky_{it} is capital income after taxes.

In our previous work (Krueger and Perri (2002)) we documented a significant increase in the cross-sectional variance of after-tax labor income (henceforth disposable earnings, or simply

⁷It needs to be acknowledged that with idiosyncratic variation a potentially important amount of measurement error is contained in our data set, which synthetic cohort techniques tend to average out. However, as long as the relative magnitude of these measurement errors do not change over time, our estimates of the changes in idiosyncratic variability of wages, earnings and consumption remain informative.

earnings) $ly_{it} - t_{it} + \tau_{it}$, without a correspondingly large increase in the variance of consumption.

The cross-sectional variance of consumption is given by

$$\begin{aligned} Var(c_{it}) &= Var(y_{it}) + Var(s_{it}) - 2Cov(y_{it}, s_{it}) \\ &= Var(ly_{it} - t_{it} + \tau_{it}) + Var(ky_{it}) + Var(s_{it}) + \text{Covariances} \end{aligned}$$

Thus, in order to explore how wage inequality is related to inequality in hours worked and to consumption inequality we have to explore how a) wage inequality is related to hours inequality and disposable earnings inequality and b) disposable earnings inequality is related to consumption inequality (and thus inequality in savings).

A. From Wage Inequality to Disposable Earnings Inequality

Disposable earnings of household i at date t are given by

$$ly_{it} = \sum_{j=1}^J w_{ijt} h_{ijt} - t_{it} + \tau_{it}$$

where j indexes members of the household receiving earnings. Throughout our empirical analysis we will restrict ourselves to households for which $J \leq 2$. As Gottschalk and Moffitt (1994) and Katz and Autor (1999) document and we confirm below, wage inequality has increased significantly in the last thirty years. We also find that the inequality in wages translates in a quantitatively substantial way into inequality of disposable earnings.

Wage Inequality

In figure 1 we plot the standard deviation of the natural logarithm of weekly wages from the CE, measured as the weekly earnings of the reference person⁸ of the household. The figure shows a quantitatively significant increase in wage inequality, as the percentage standard deviation

⁸The reference person in the CE is defined as the person who owns or rents the home in which the household resides.

of wages increases from around 55% in the early 70s to around 67% at the end of the 90s. We select only households that are classified as complete income respondents in the CE, that report positive consumption expenditures for each quarter in which they are in the sample and whose reference person is between 20 and 64 years of age and has worked at least one week in the year. Also, in order to reduce measurement error we exclude reference persons that report unreasonably low wages; we select only those that have a weekly wage above \$67 in constant 1982 dollars (equal to half of the 1982 minimum wage based on a 40 hour workweek). Finally, since we are interested in the welfare impact of the increase in wage dispersion we divide the wage of the reference person by the number of adult equivalents in the household.⁹

Note that our sample includes households with low wages that are particularly vulnerable to the increase in inequality (e.g. households whose reference person is not a full-time worker and households with female reference persons).¹⁰ On the other hand, for households with a reference person that never works during the year we do not observe wages; therefore these households are not included in our wage distribution sample. Juhn, Murphy and Topel (2002) argue that the number of such households is large and increasing throughout the 90s; they use the CPS to impute wage data for individuals that are out of the labor force. The small sample of the CE does not allow us to follow the same procedure. Fortunately, though, we have consumption data for these households, and we will include them in the consumption distribution sample so that our welfare analysis will take into account the effects of long jobless spells and labor force withdrawals described by Juhn et al. (2002).

FIGURE 1 APPROXIMATELY HERE

We briefly want to compare our wage inequality findings with previous studies in the litera-

⁹The number of adult equivalents is computed using the Census equivalence scale.

¹⁰This sample selection strategy was suggested to us by our discussant Steve Davis.

ture, because some authors (for example Cutler and Katz, 1991b) have questioned the reliability of CE wage and income data. In order to do so we compute inequality measures in the CE for the wage distribution of male reference persons who work at least 40 weeks for at least 35 hours per week, the same selection criteria used by Katz and Autor (1999) to compute wage inequality measures for the Current Population Survey (CPS) and by Heathcote, Storesletten and Violante (2003) for the Panel Study of Income Dynamics (PSID).

In figure 2 we report the 90% -10% differential for male weekly log-wages in the CE and the CPS (see figure 4 by Katz and Autor, 1999). Notice that both timing and magnitude of the increase in inequality in the two samples are quite comparable. In figure 3 we compare the increase in the standard deviation of log wages in our sample with the increase in the same measure in the CPS sample (Katz and Autor (1999) table 1) and the PSID sample (Heathcote and al. (2003), table 2).¹¹ The figure confirms that both timing and the magnitude of the increase in wage inequality are similar across the three data sets.¹²

FIGURES 2 AND 3 APPROXIMATELY HERE

Earnings Inequality

Wage inequality may be accentuated or mitigated by the endogenous labor supply decisions of the members of the household. A decline in the primary earners' wage may induce other members of the households to start working or to work longer hours (the added worker effect). On the other hand, in the face of temporarily low wages the primary wage earner may decide to intertemporally substitute and work less today and more in the future when wages have recovered. The map between

¹¹All the series in figures 2 and 3 are based on wages per person and not on per adult equivalent wages (as the series in figure 1).

¹²Katz and Autor only report the standard deviation of log-wages for 5 years. The level of the standard deviations is very similar in the CE and CPS while it is slightly higher in the PSID (for example in 1987 the standard deviation of log wages is 0.57 in the CE, 0.579 in the CPS, and 0.601 in the PSID). Small differences in levels are not surprising as levels are affected by the different top-coding thresholds in the data sets and by the potentially different extent of measurement error.

wage and labor income inequality is therefore determined by the relative importance of these effects. Furthermore, changes in taxes and transfers can reduce or magnify the effect of changes in wages on disposable earnings. In figure 4 we report inequality in total household earnings, $\sum_j w_{ijt}h_{ijt}$, and of disposable household earnings, $\sum_j w_{ijt}h_{ijt} - t_{it} + \tau_{it}$. The sample we select is exactly the same as the one chosen to compute wage inequality; as before we divide every variable by the number of adult equivalent members of the household.

FIGURE 4 APPROXIMATELY HERE

The graph shows that total earnings inequality increases by a slightly smaller amount than wage inequality, consistent with moderate responses of labor supply to wage changes. Notice also that, not surprisingly, given the progressivity of the government tax and transfer system, disposable earnings inequality is significantly lower than inequality in total earnings. It is interesting, however, that the gap between total earnings and disposable earnings inequality has declined over time. This suggests that the government tax and transfer system may have become less progressive over our sample period, which caused part of the increase in *disposable* earnings inequality.

Finally, for the purpose of this paper it is important to understand whether the increase in earnings inequality stems from an increase in between- or within-group differences. The empirical decomposition we employ is simple and widely used (see Katz and Autor, 1999). To control for changes in the age race composition of the population we first regress each cross-section of the raw data of disposable earnings on a constant, a quartic in the age of the household reference person and a race dummy. The cross section of the residuals is denoted by $y_t = \{\ln y_{it}\}_{i \in I}$. By construction y_t has zero cross-sectional mean for every t and each observation is interpreted as percentage deviations of earnings of household i from the average earning of a household of same age and same race.

We then decompose each earning observation in a into a component y_{it}^g due to observables

(education and sex) and a part due to unobservable, purely idiosyncratic variation y_{it}^d

$$\ln(y_{it}) = \ln(y_{it}^g) + \ln(y_{it}^d)$$

so that $y_{it} = y_{it}^g y_{it}^d$.¹³ We chose education and sex to define groups, since the increase of the skill premium and the decline of the gender gap are the two most important determinants of the changes of between-group earnings inequality in the last 30 years. We then assume that the unobservable idiosyncratic component of earnings is orthogonal to the observable, common-group component of earnings, so that we can find it by regressing $\ln(y_{it})$ on the years of education of the household reference person and sex dummies. For each household observation $\ln(y_{it})$ we therefore obtain two new observations, the predicted (by education and sex) value $\ln(y_{it}^g)$ and the residual $\ln(y_{it}^d)$. To understand their interpretation consider, for example, household i that has $\ln y_{it} = -0.4$, $\ln(y_{it}^g) = -0.2$ and $\ln(y_{it}^d) = -0.2$; this means that this household has earnings 40% below the average earnings of an household with the same age and same race; half of this difference is explained by its education/sex characteristic and the other half is purely idiosyncratic variation.

From $\{\ln(y_{it}^g)\}$ we compute cross-sectional “between-group” variances σ_{gt}^2 , and from $\{\ln(y_{it}^d)\}$ cross-sectional idiosyncratic variances σ_{dt}^2 . This procedure yields time series $\{\sigma_{gt}^2, \sigma_{dt}^2\}_{t \in T}$ of variances satisfying

$$(1) \quad \sigma_{yt}^2 = \sigma_{gt}^2 + \sigma_{dt}^2.$$

for our sample period T .

Figure 5 shows the trends in between- and within-group earnings standard deviations. Note that within-group inequality accounts for a larger fraction of total inequality but that both the between- and the within-group component display a significant increase. Within-group inequality

¹³We decompose the logarithm of earnings since the standard deviation of the log of a variable has a cardinal interpretation, which makes our findings below easier to evaluate.

increases throughout the sample while between-group inequality slightly declines in the 70s (reflecting a reduction in the college premium) but increases significantly in the 80s and 90s. The patterns of both inequality measures from the CE are again similar to those from the CPS and the Census data described by Katz and Autor (1999).

FIGURE 5 APPROXIMATELY HERE

B. Consumption and Hours Inequality

Our focus on welfare now leads us to investigate the evolution of the variables more directly connected to the lifetime utility of households, namely consumption and leisure.

Consumption Inequality

Figure 6 displays the trend in disposable earnings inequality (from figure 5), together with two measures of consumption inequality. Both measures are standard deviations of household per adult equivalent nondurable consumption plus imputed services from durables (this is the definition of consumption used throughout the paper; see the appendix and Krueger and Perri (2002) for details about the construction of the consumption data). The first measure (labeled the positive wage sample) is computed using the same sample selection criterion we use for wage and earnings inequality and does not include households with reference persons that earn no labor income. The second measure (which we label full sample) includes all household that are complete income respondents, have a reference person between the age of 20 and 64 and report positive consumption. Notice that the second sample is larger than the first (about 15% larger) and that it includes households whose reference person has left the labor force or is suffering a long unemployment spell. The picture shows that although the *level* of inequality differs across the two samples, neither series of consumption inequality displays an increase comparable to the one registered for disposable earnings inequality. In figure 7 we plot the average per adult equivalent consumption of the household in the bottom,

middle and top quintile of the disposable earnings distribution (relative to average consumption in that year).¹⁴

FIGURES 6 and 7 APPROXIMATELY HERE

Both pictures confirm and extend the findings of Krueger and Perri (2002). During a period of strongly widening wage and earnings dispersion there has been surprisingly little increase in consumption inequality.¹⁵ In particular, households at the bottom of the earnings distribution have experienced only a mild reduction in their relative consumption, even though their relative earnings have declined substantially.

FIGURE 8 APPROXIMATELY HERE

Note that even though total consumption inequality has increased only by a small amount, this phenomenon is a result of two offsetting trends. If we decompose -as we did for earnings- consumption inequality into within- and between-group inequality¹⁶, figure 8 shows that the between-group component has increased substantially, by about the same magnitude as the increase in between-group earnings inequality. In stark contrast to earnings, *within-group* consumption inequality has, in fact, slightly declined. We draw two lessons from this decomposition, which will be

¹⁴Consumption by earnings quintile is computed using the sample of all households that report positive consumption (full sample). If we restrict the sample to households with reference person working at least one week, consumption by quintiles displays the same constant pattern.

¹⁵Kreueger and Perri (2002) document that this fact is robust to various definition of consumption expenditures. It is worth mentioning that our results are based on quarterly consumption expenditures reported in the CE interview survey. Attanasio (2002) presents some results from the CE Diary survey (that is the bi-weekly survey of expenditure data for items purchased on a daily or weekly basis) showing increasing consumption inequality. Further work should be done to establish the exact source of the discrepancy between the two surveys. Another important concern with the CE consumption data is that their total does not matches up with the NIPA total. Slesnick (2001) discuss some possible explanations for this phenomenon but concludes that a large part of the discrepancy between CE and NIPA is still unexplained.

¹⁶We do the decomposition exactly as for income. We first control for changes in the age/race structure of the population by regressing the log of non durable consumption plus services from durables on a quartic polynomial in the age of the reference person and on a dummy for his/her race. We then regress the residuals on the years of education and sex of the reference person of the household. The sample used consists of all households that report positive consumption.

crucial for interpreting our welfare calculations below. First, it is not true, as one may think from simply looking at overall consumption inequality, that the increase in earnings inequality had no impact on consumption inequality. Rather, it suggests that the increase in between-group earnings inequality (i.e. the increase in the skill premium) translates almost one to one into an increase of consumption inequality. Since between-group inequality tends to be highly persistent, it is likely to have important welfare consequences.¹⁷ Second, the fact that within-group disposable earnings inequality has increased while within-group consumption inequality has not indicates that the increase in within-group earnings variations has been effectively smoothed by households, possibly through some form of credit markets.¹⁸

Hours Inequality

Since a household's lifetime utility depends on both the consumption and leisure it enjoys, the number of hours worked by members of the household, by determining the hours available for leisure, is an important determinant of welfare. Therefore, in this section we present some evidence on how the distribution of hours worked across American households has changed. In figure 9 we plot the percentage standard deviation of household yearly hours worked (per adult), together with the standard deviation of wages of the reference person of the household.¹⁹

FIGURE 9 APPROXIMATELY HERE

We observe that, despite the increase in wage dispersion, the variability in hours worked shows a moderate decline. If hours and consumption were uncorrelated, abstracting from any

¹⁷This is a point that has been highlighted by Attanasio and Davis (1996). In the second part we will provide an estimate of the persistence of between-group differences.

¹⁸Blundell and Preston, 1999, Heathcote et al. (2003) and Krueger and Perri (2002) investigate the role of credit markets in generating this divergence between income and consumption inequality in detail.

¹⁹We use the same sample we used for the wage and earnings distribution. We construct per-person hours worked by households in the following way: if a household consists of a single adult member, its hours worked are the hours worked in a year by that adult, and if the household consists of (at least) two adult members, average hours worked by the household are measured as total hours worked by both members, divided by 2.

change in their trends, the reduction in hours dispersion would imply welfare gains. However, figure 10 suggests that hours and consumption are not uncorrelated. It plots the number of per-person yearly hours worked for all households in our sample and for households in different quintiles of the consumption distribution.

FIGURE 10 APPROXIMATELY HERE

Apart from the strong average increase in hours worked, whose causes and consequences we will not analyze here²⁰, the picture reveals two other facts. First, households with high consumption work more hours; second, high consumption households have increased their hours by less than low consumption households. For example, in 1972-73, on average households in the bottom quintile of the consumption distribution worked 28% less than households in the top quintile. In 1999-2000 the same ratio has fallen to 17%. In other words, figure 10 suggests a positive, but falling correlation between consumption and hours worked.

Figure 11 plots the correlation between hours and consumption (total, between-group component and within-group component); The picture indeed shows a significant decline over time in the correlation between hours and consumption. This declining positive correlation explains the fall in hours variability and, together with constant consumption inequality, can generate negative welfare effects for low-consumption households, even if their relative consumption may not fall. More concretely, in the 70s the high correlation between consumption and hours implied that households with low consumption were compensated with high leisure. In the 90s the reduction in this correlation implies that households with relatively low consumption do not enjoy as much leisure as they used to, and thus face potentially substantial welfare losses from this change over time. We will quantify the welfare consequences from a declining correlation in consumption and hours worked in

²⁰Inspection of the CE data immediately reveals that most of the average increase in hours worked is due to increased female labor force participation.

the section containing our welfare calculations below.

FIGURE 11 APPROXIMATELY HERE

C. Summary

To summarize, the CE data suggest that the substantial increase in wage inequality has caused an increase in household earnings inequality of similar magnitude. The tax and transfer system has not mitigated this increase, so that it appears equally strong in disposable earnings. The *consumption* distribution data, on the other hand, suggest that low consumption households do almost as well (relatively to the mean) in the late 90s as they did in the 70s.

This, however, does not necessarily imply that the welfare effects of the increase in inequality are negligible. A closer investigation of consumption inequality reveals that households that suffer permanent relative income falls (e.g. households with little education) experience substantial declines in relative consumption, with associated important welfare consequences.

In addition, the declining trend in the consumption-hours-worked correlation suggests that one way through which low-consumption household maintain their relative consumption is by working longer hours, again with implied welfare losses. In the next section of the paper we will try to quantify the size and exact distribution of these welfare effects.

4. Welfare Consequences of Increased Inequality

What are the welfare consequences of the trends in wage, income and consumption inequality documented in the last section? In order to explicitly answer these question we now have to take a stand on how households value consumption and leisure, and we have to model, in a statistical or theoretical way, how a households' consumption and leisure processes changed as the variability of wages and earnings changed. The second step can be carried out in (at least) two different ways. Either one takes a stand on a particular theoretical model, feeds as input into the model two different

earnings processes estimated from data (one reflecting increased earnings inequality), computes the corresponding consumption and labor supply allocations implied by the model, and then uses these allocations to compute the welfare gain/loss distribution from the change in the earnings (or wage) process.²¹ Alternatively one can specify or estimate a statistical model for consumption and labor supply (or leisure). The parameters governing these processes are allowed to be time-dependent, so that the processes reflect the changes in the cross-sectional distribution of consumption and hours worked documented in the last section of the paper. We will adopt the second approach in this paper.

We now first discuss our exact estimation strategy, then we explain how we use our estimation results to answer the welfare question posed at the beginning of the paper, and finally we present our results.

A. Modelling Household Earnings, Consumption and Leisure Processes

To carry our welfare analysis we need stochastic processes for disposable earnings, consumption and for leisure. Take an arbitrary variable of interest from the last section, for concreteness disposable earnings. In the empirical analysis we discussed how each pre-filtered cross section $\ln(y_{it}^d)$ can be decomposed in a between-group $\ln(y_{it}^d)$ and a within-group part $\ln(y_{it}^g)$.

In this section we specify time series processes for both components of earnings. In particular, we assume that both $\ln(y_{it}^d)$ and $\ln(y_{it}^g)$ follow finite state Markov chains (of cardinality N) with possibly time-varying states Y_t^d and Y_t^g and with time-invariant²² transition matrices π^d, π^g .

The states $\{Y_t^g\}_{t \in T}$ are determined by setting $Y_t^g(i)$ equal to the midpoint of the i -th quantile of the $\{\ln(y_{it}^g)\}$ -sample for that period. By construction our stochastic process matches, for every t ,

²¹The studies by Krueger and Perri (2002) and Heathcote et al. (2003) take this approach.

²²In principle our procedure could be used to estimate time-varying transition matrices. We have experimented with this for the time period 1984-2000 and found that the transition matrices display very little time variation. We therefore decided to use time-independent transition matrices in our analysis.

$N(N - 1)/2$ quantile ratios of $\{\ln(y_{it}^g)\}$, its mean (by construction equal to 0), but not necessarily its variance.²³ A similar construction yields the states $\{Y_t^d\}_{t \in T}$. Note that, because of the initial pre-filtering, there is no aggregate growth in the logs of any of our variables (i.e. $E_i(\ln x_{it}) = 0$, for every x and every t), which, in particular, removes aggregate consumption growth and growth in average hours worked. The welfare conclusions presented below have to be interpreted with this remark in mind, which we will follow up on in the conclusion.

Finally, we use the panel dimension of the CE to estimate Markov transition matrices for the variables of interest, which will embody the persistence properties of our stochastic processes. For all observations $\{\ln(y_{it}^g)\}_{i,t}$ we group households into N relative classes delimited by $N - 1$ uniformly spaced, time dependent quantiles: the first class is composed by the bottom $1/N\%$ of the distribution in that quarter for that variable, the second class by the following $1/N\%$ etc. With this construction the lowest class for $\{\ln y_{it}^g\}$ is interpreted as the group of the households for which earnings explained by observables (sex and education) is lowest (i.e. households with female head and low education levels). We then search for all households for which we have observations in two consecutive periods and compute which relative class they belong to in the second period.²⁴ After repeating this procedure for every period in the sample the probability π_{ij}^g of transiting from class i to class j is computed as the total number of households transiting from i to j , divided by the total number of households starting in class i for the *entire* sample.²⁵ Using the same procedure for the sample $\{\ln(y_{it}^d)\}$ yields a transition matrix π^d . As a result of our estimation procedure we

²³We find, however, that the ratio between the variance implied by the estimated process and the variance in the data is quite high (about 70%) and almost constant over time. It would have been easy to set the states of our stochastic process to exactly match the variance in the data but we opted to match quantile ratios of the data because of our interest in the distribution of welfare losses. Our procedure is similar in spirit to the one used by Tauchen and Hussey (1991).

²⁴More precisely, a period is roughly 3 quarters, as we use data from the 2nd and 5th interview of a household in the CEX to determine transitions. This timing comes closest to our notion of a period length of one year. Second, the only useful income observations for a household are contained in these interviews..

²⁵Transitions of households between “groups”, are rare and occur only if there are large swings in the returns to education, in the wage gender gap or (the dominating reason) if households change sex or the education of their reference person (through death, marriage or divorce).

obtain Markov processes for the group-specific as well as within-group components of earnings; we follow exactly the same procedures to construct time series processes for household consumption. In table 1 we report a summary measure of the persistence properties of the various processes we estimate. Note that the processes for between-group differences are significantly more persistent than the processes for within-group differences, so that changes in between-group inequality are more likely to have larger welfare effects.

For hours worked we find that between-group inequality explains only a very small fraction of total variance (less than 0.5%), so the process for hours does not distinguish between the two components. In our welfare analysis we will also study the case in which welfare is jointly determined by household consumption and hours worked. For this case we specify a joint stochastic process for between-group consumption, within-group consumption and hours worked. Since the correlation between hours and between-group consumption is fairly constant and close to 0 (see figure 11 above) we assume that hours and between-group consumption are uncorrelated. However, we specify the parameters of the joint Markov chain for hours and within-group consumption in such a way that the correlation between the two variables in the model exactly matches the falling cross sectional correlation between the two variables in the data as reported in figure 11.²⁶

For short, denote the estimated stochastic processes for earnings, consumption and hours by (Y, C, H) and the processes associated with no change in wage inequality (i.e. the process with states constant at their 1972 values) by $(\hat{Y}, \hat{L}, \hat{H})$. These are the necessary ingredients for the welfare calculations which we describe next.

B. Calculating the Welfare Consequences

We assume that households value streams of consumption c_t and leisure $l_t = 1 - h_t$ (where h_t are hours worked per household member, as a fraction of total non-sleeping time) according to

²⁶The complete results of our estimation procedure are available upon request.

the lifetime utility function

$$U(C, L) = E \sum_{t=0}^{\infty} \beta^t u(c_t, l_t).$$

We restrict the period utility function u to lie in the parametric class

$$u(c, l) = \begin{cases} \frac{[c^\alpha l^{1-\alpha}]^{1-\sigma}}{1-\sigma} & \text{for } \sigma \neq 1 \\ \alpha \ln(c) + (1 - \alpha) \ln(l) & \text{for } \sigma = 1 \end{cases}$$

where (α, β, σ) are preference parameters that govern the relative importance of consumption relative to leisure, time discounting and the intertemporal elasticity of substitution, respectively.²⁷ Nested in this formulation are utility functions that depend on consumption alone ($\alpha = 1$).

Suppose that all economic variables follow Markov processes as estimated above, then we can write²⁸

$$(2) \quad V(t, c, l) = u(c, l) + \beta E [V(t + 1, c', l') | (c, l)]$$

Note that, conditional on knowing the function V the conditional expectation can be evaluated since the stochastic processes for consumption and leisure have Markov structure. Also note that for an arbitrary individual with current consumption and leisure (c, l) , expected lifetime utility is given by $V(t, c, l)$, which obviously depends on the time-varying stochastic process for consumption and leisure. Further note that this Bellman equation does not involve any maximization.

In particular, let $V(c, l)$ be the value function associated with the estimated stochastic processes and let $\hat{V}(c, l)$ denote the value function associated with the stochastic processes if wage inequality had not increased (and thus the cross-sectional variances $(\sigma_{dt}^2, \sigma_{gt}^2)$ had remained the

²⁷Note that these utility functions have unit elasticity between consumption and leisure. With this class of utility functions our welfare results are invariant to long-run deterministic average consumption growth.

²⁸This formulation implicitly assumes that households are infinitely lived. With finite lives age becomes an additional state variable. The welfare consequences for a person living through the increase in inequality for only a finite number of periods is roughly proportional to the numbers reported below (such a household faces the same stochastic processes as our model households, simply for less periods). The proportionality factor is less than one and the smaller the smaller the remaining lifetime of a household.

same).²⁹ More precisely, \hat{V} is computed by employing the same transition matrices as for V , but under the assumption that the states of the Markov chains remain at their estimated 1972 values forever, rather than change over time, in the way estimated above.

The welfare consequences of an agent with current state (c, l) in 1972 (the distribution of which we observe in our cross-sectional data set) of the ensuing increase in wage inequality in the future is thus given by

$$\mathcal{W}(c, l) = \left\| V(c, l) - \hat{V}(c, l) \right\|$$

where $\|\cdot\|$ is a particular metric. We let $\mathcal{W}(c, l)$ equal the uniform percentage increase of consumption in each state of the world needed to make a household indifferent between the stochastic processes with more variability and the ones without, keeping leisure constant. What we want to document is the distribution of $\mathcal{W}(c, l)$, that is, the distribution of the welfare consequences of the increase in wage inequality.

C. Results *An Upper Bound*

Suppose that households do not have access to any savings technology and do not value leisure,³⁰ Then consumption equals to labor earnings y and equation (2) becomes (ignoring irrelevant constants and suppressing time indexation)

$$V(y) = u(y) + \beta E [V(y')|y]$$

or more explicitly

$$V(y^g, y^d) = \frac{[y^g y^d]^{1-\sigma}}{1-\sigma} + \beta E [V(y^{g'}, y^{d'})|y^g, y^d]$$

²⁹We suppress the time index $t = 1972$.

³⁰The same results are obtained if households do not change their labor supply or if consumption and leisure are separable in the utility function.

and a similar definition for $\hat{V}(y^g, y^d)$ applies. The welfare criterion is given by

$$\begin{aligned}\mathcal{W}(y^g, y^d) &= \left\| V(y^g, y^d) - \hat{V}(y^g, y^d) \right\| \\ &= \left[\frac{\hat{V}(y^g, y^d)}{V(y^g, y^d)} \right]^{\frac{1}{1-\sigma}} - 1\end{aligned}$$

The number $\mathcal{W}(y^g, y^d)$ gives the percentage increase in consumption (or earnings) at the new, more unequal earnings process, at each contingency, necessary to make an agent with arbitrary current earnings $y = (y^g, y^d)$ indifferent between the old and the new wage (and thus earnings) process.³¹ In order to compute this number three steps have to be carried out

1. Solve for the function V and \hat{V} using standard value function iteration, given time series processes Y and \hat{Y}
2. Draw arbitrary $y = (y^g, y^d)$ from the initial earnings distribution
3. Evaluate $\mathcal{W}(y^g, y^d)$

Evidently steps 2. and 3. can be easily carried out for all y in the initial distribution of earnings, and thus the distribution of welfare consequences can be derived. These numbers provide an upper bound for the welfare consequences from increased earnings inequality because agents are assumed not to be able to smooth consumption via savings responses.

Example 1. A simple example illustrates this basic procedure. Suppose that before 1972 earnings followed the simple continuous-state white noise process

$$\ln(y^t) = \hat{\mu} + \hat{\varepsilon}$$

³¹For $\sigma = 1$, one can show that the welfare criterion as defined above is given by

$$\mathcal{W}(y^g, y^d) = \exp \left[(1 - \beta)(\hat{V}(y^g, y^d) - V(y^g, y^d)) \right] - 1$$

where $\hat{\varepsilon} \sim N(0, \hat{\sigma}_\varepsilon^2)$, and after 1972 the process changes to

$$\ln(y') = \mu + \varepsilon$$

with $\varepsilon \sim N(0, \sigma_\varepsilon^2)$, $\sigma_\varepsilon^2 = \gamma \hat{\sigma}_\varepsilon^2$ with $\gamma > 1$ and

$$\mu = \hat{\mu} - \frac{1}{2} \hat{\sigma}_\varepsilon^2 (\gamma - 1) < \hat{\mu}$$

Note that the adjustment of the mean of log-earnings is required for the level of earnings to have the same mean before and after the change in its variance. For period utility being logarithmic, the value function V solves

$$V(y) = \log(y) + \beta \int V(e^{\mu+\varepsilon}) d\Phi(\varepsilon)$$

where Φ is the normal cdf with variance $\sigma_\varepsilon^2 = \gamma \hat{\sigma}_\varepsilon^2$. The function \hat{V} solves the same equation, with Φ replaced by $\hat{\Phi}$, the normal cdf with variance $\hat{\sigma}_\varepsilon^2$. A simple guess-and-verify strategy shows that

$$\begin{aligned} V(y) &= \ln(y) + \frac{\beta}{1-\beta} \mu \\ \hat{V}(y) &= \ln(y) + \frac{\beta}{1-\beta} \hat{\mu} > V(y) \end{aligned}$$

In particular, if $\hat{\mu} = 0$ (which simply normalizes mean earnings), then

$$\begin{aligned} \mathcal{W}(y) &= \mathcal{W} = \exp(\beta(\hat{\mu} - \mu)) - 1 \\ &= \exp\left(\frac{\beta}{2}\right) \exp\left(\hat{\sigma}_\varepsilon^2(\gamma - 1)\right) - 1 > 0 \end{aligned}$$

and for this particular example the welfare losses from increased earnings inequality are a) independent of an agents' current earnings and b) proportional to the increase of the variance of log-earnings $\hat{\sigma}_\varepsilon^2(\gamma - 1)$.

For a general stochastic earnings process and period utility function, of course, the value function evidently cannot be solved for by guess-and-verify methods and has to be computed nu-

merically.³² As a benchmark an annual time discount rate of 4% and logarithmic utility are assumed. We then document how our welfare conclusions depend on the degree of risk aversion of households and discuss the role of the time discount rate for our results.

FIGURE 12 APPROXIMATELY HERE

In figure 12a we plot the welfare losses implied by our estimated income processes from the CE. Each individual, in 1972, is characterized by $(y^g, y^d) \in (Y^g, Y^d)$, where both Y^g and Y^d have cardinality 9. Thus there are 81 distinct earnings classes, with equal population mass $\frac{1}{81}$. On the x -axis we plot idiosyncratic earnings y^d , as a fraction of average (idiosyncratic) earnings, and each of the separate 9 lines corresponds to one group-specific earnings level y^g . The y -axis shows the welfare losses, in percent consumption equivalent variation, implied by the increase in earnings inequality, as estimated by our processes.

We observe that the welfare losses implied by the increase in earnings inequality are potentially substantial, amounting to as much as 2% of consumption or more for more than half of the population, namely the earnings-poor households in 1972. The welfare losses are declining (and substantially so) as one moves up the group-specific earnings distribution, to the extent that the highest two earnings groups (22% of the population) benefit from the increase in earnings inequality.³³ Welfare losses (or gains) are fairly uniform in the idiosyncratic income state y^d a household starts with in 1972.

These findings can be interpreted as follows. The *aggregate* statistics show an increase in inequality in the last thirty years. An *individual* household in 1972 experiences two events: future changes in its expected mean earnings and more earnings risk in the future. Conditional on

³²Given the discrete nature of the income process this can be done quickly and with very high precision, using a standard value function iteration algorithm.

³³Note that, when we refer to welfare losses from inequality from now on for short, it is understood that individuals are affected directly only by changes in their relative earnings (and higher earnings uncertainty in the future), rather than by the change in the aggregate statistics per se.

remaining in the same earnings class, high-earnings class households will enjoy increasing relative earnings, and low-earning classes declining relative earnings. This is true for both the group- as well as the idiosyncratic component of earnings. From the figure we observe, however, that the welfare losses differ substantially by groups, but are fairly uniform with respect to the idiosyncratic earnings component. This difference is explained by the much higher persistence of the between-group earnings process, compared to the idiosyncratic process, documented in the last section. As a rough approximation, suppose that the group-earnings transition matrix is the identity matrix (so that all agents with probability one stay in their earnings group) and the idiosyncratic earnings process is iid over time, then welfare losses should vary greatly by group (because relative group earnings have diverged over time) but should be uniform across idiosyncratic states, as everybody is equally negatively affected by the higher idiosyncratic earnings uncertainty. Figure 12 indicates that this first approximation provides fairly accurate intuition for the computed welfare numbers. The negative (uniform) impact of higher idiosyncratic *uncertainty* also explains why almost the entire population suffers welfare losses from increased earnings inequality, even though half of the population experiences increases in relative earnings.

From the previous discussion one would expect the welfare losses from increasing earnings inequality to rise with risk aversion of the household, since higher idiosyncratic uncertainty is disliked more severely. That this is indeed the case is demonstrated by figure 12b, which assumes a coefficient of relative risk aversion of $\sigma = 2$. Now the welfare losses become as high as 8% of consumption, and despite significant expected earnings increases even the highest earnings group suffers welfare losses of about 2% due to the increase in future earnings (and thus, by construction, consumption) variability. To summarize, when judged from earnings data for a degree of risk aversion of $\sigma = 2$, commonly assumed in the macro and public finance literature, all groups of the population suffer welfare losses from increased earnings instability that more than outweigh the increase in relative

incomes experienced by the highest earnings groups.

Before documenting how our results change if one uses consumption rather than earnings data for our welfare analysis, a brief discussion of the role of the time discount factor is in order. Our estimated transition matrices (in particular for the group component) are persistent, but not perfectly persistent, so that the unconditional probability of a household leaving her 1972 earnings state increases with time. With sizeable time discounting of 4% per year the experienced welfare losses are affected by events along the transition period 1972–2000 (the increase in earnings inequality occurs gradually over this time period) and vary crucially with initial earnings states, because earnings realizations in the early years of the transition largely determine welfare. This is what Figure 12 shows. If one were to choose very low time discounting (none in the limit), then households' welfare losses are determined largely (completely in the limit), by the change in the steady state earnings distribution, and would neither depend on a households' initial earnings position nor the transition path.³⁴ Since we want to document the distribution of welfare losses across different population groups we chose a time discount rate of 4% that attributes substantial importance to initial conditions and transition paths. As comparison, a time discount rate of 0.05% yields welfare losses, roughly uniform across the population, of about 1.4% for $\sigma = 1$ and 10.5% for $\sigma = 2$.

The Effects of Consumption Smoothing

The previous section documented potentially large welfare costs of increased earnings inequality, under the assumption that households have no ability to intertemporally smooth consumption, either via self-insurance and/or formal and informal insurance arrangements. Calculations like these often form the explicit or implicit basis for concern expressed about the increase in inequality by researchers and policy makers.

³⁴One may interpret this latter thought experiment and welfare numbers as household living behind the veil of ignorance; they don't know which income state they will be born into in the old, pre-1972 steady state and they don't know in which income state they end up in the new, post-2000 steady state.

As argued in the introduction, basing welfare evaluations of increased inequality on consumption data directly takes the mitigating effects of intertemporal smoothing opportunities, government transfer programs and explicit or implicit insurance programs (such as transfers among members of the extended family) into account. Therefore in this subsection we repeat our previous analysis, but now use estimated consumption processes from the CE instead of earnings processes. The Bellman equation becomes

$$(3) \quad V(c) = u(c) + \beta E [V(c')|c]$$

which is as easily solved as in the previous subsection, under the maintained assumption that the estimated stochastic consumption processes have Markov structure.

Again we start with the benchmark case of logarithmic utility and plot the welfare losses from increased consumption inequality in Figure 12c. Qualitatively the welfare consequences are similar to the previous section: to a large part the consumption group to which a household belongs determines whether it ends up as a loser or winner of the increase in inequality; again the welfare consequences are fairly uniform across idiosyncratic consumption states. Quantitatively, however, some crucial differences to the results based on earnings observations emerge. The highest welfare losses now amount to slightly more than 5% of consumption, but only for the consumption-poorest 10% of the population.³⁵ Note that, while maximal welfare losses from increased inequality when measured with consumption data are even higher than with earnings, a much larger fraction of the population (roughly 45%) now benefits sufficiently from the increase in relative group consumption to experience welfare gains from increased consumption inequality. This differential finding is due to the fact that the variability of the idiosyncratic component of *consumption* has remained relatively constant over

³⁵These high losses of the lowest consumption group, compared to the more moderate losses of the lowest earnings group documented above are largely due to timing of the transition, coupled with sizeable time-discounting: relative consumption for the lowest group drops early on and then stays constant, whereas relative earnings for the lowest group declines most dramatically only at the end of the 1972-2000 transition period.

time, very much in contrast to that of *earnings*. Therefore the welfare consequences, when derived from consumption data, are almost exclusively determined by relative group consumption.³⁶

Whereas for earnings an increase in the risk aversion of households led to dramatically increased welfare losses from increased inequality, when using consumption data the welfare consequences of increased inequality are close to invariant to changes in a households' attitudes towards risk. Again, this is explained by the stable trend of the variance of the idiosyncratic consumption component. The welfare losses/gains in figure 12d, computed for risk aversion of $\sigma = 2$, are almost indistinguishable from the corresponding figure for $\sigma = 1$. Thus, in comparison to our findings for earnings, for a reasonable degree of risk aversion of $\sigma = 2$, consumption data do not suggest nearly as severe welfare losses as do earnings data. Repeating our calculations with a low time discount rate of 0.05% again results in welfare losses that are fairly uniform across the population, amounting to losses of 0.5% for $\sigma = 1$ and 1.6% for $\sigma = 2$, again significantly smaller than the corresponding numbers derived from earnings.

Our findings are consistent with Attanasio and Davis (1996) who document that a significant share of the increase in wage inequality between observably different groups is reflected in increases in consumption inequality between these groups. Our welfare numbers reproduce exactly this phenomenon. These results are also in line with findings in our earlier work (Krueger and Perri, 2002) in that the idiosyncratic component of consumption inequality has not increased significantly over time and thus has not, to any noticeable extent, contributed to the welfare losses from increased inequality. Finally note that, in order to focus on the welfare effects of increased inequality we have effectively de-trended our data from aggregate consumption growth. Thus our results are not to be interpreted as absolute standards of living having declined over the last thirty years for a large fraction of the population.

³⁶Since the transition matrix for the group-component of consumption displays very high persistence, increased *risk* in the between-group component is quantitatively of second-order importance.

The Impact of Changes in Leisure

In our empirical section we documented that, in the light of increased wage variability, the relative labor supply of different groups in the population has changed, albeit only moderately so. In this section we extend our welfare analysis to incorporate these changes in relative labor supply and hence leisure.

The Bellman equation for this augmented problem now becomes

$$(4) \quad V(c, l) = u(c, l) + \beta E [V(c', l') | (c, l)]$$

Here the expectations operator pertains to the joint Markov process for consumption c and leisure l . This Markov process may feature independence between c and l (our benchmark), or allow for dependence between consumption and leisure (a case that we will investigate as sensitivity analysis).

Both the definition of welfare costs as well as its computation remain the same as in the previous subsection. We parameterize the period utility function as before, and choose a share parameter $\alpha = \frac{1}{3}$ and a non-sleeping time endowment of 15 hours per person, per day.³⁷

Labor Supply Uncorrelated with Consumption How does the inclusion of leisure into the analysis change the magnitude of the welfare losses from increased inequality? To a first order approximation the magnitude as well as the distribution of these losses is unaffected by the incorporation of leisure into the analysis. Again, the consumption group a household belongs to in 1972 largely determines how it fares in terms of welfare, whereas the idiosyncratic consumption state, the

³⁷In a static deterministic model with our preferences, an agent would choose to work exactly $\frac{1}{3}$ of her nonsleeping time if $\alpha = \frac{1}{3}$, independent of the wage.

Note that the coefficient of relative risk aversion for consumption is now given by

$$\frac{-cu_{cc}}{u_c} = \sigma\alpha + 1 - \alpha$$

When we report results for a particular risk aversion, we set σ , given $\alpha = \frac{1}{3}$, to the appropriate value to attain that risk aversion.

leisure state or the risk aversion of the household play only a minor quantitative role.³⁸

Since the hours and thus leisure distribution has become somewhat less dispersed between 1972 and 2000, the welfare losses from increased overall inequality are slightly mitigated by the presence of leisure in the utility function, but the reduction in welfare losses amounts to no more than 0.4% for any of the population groups, compared to the welfare losses reported in the last section based on consumption observations alone.

One has to bear in mind, though, that, as with consumption, our estimation of stochastic processes for hours worked effectively removed the positive trend in average hours worked by households; therefore our welfare numbers do not reflect the potentially adverse effect on well-being of longer hours worked by a large fraction of US households.

Labor Supply Correlated with Consumption In our empirical analysis we found that the correlation between idiosyncratic consumption and leisure is negative, and, more importantly, has declined in absolute value between 1972 and 2000. In other words, while low consumption households used to at least enjoy significantly higher leisure in 1972, this has become less pronounced in 2000. We now investigate whether allowing consumption and leisure to be correlated (and therefore allowing more extreme negative states over time -those with increasingly low relative consumption and decreasingly high leisure), modifies our result that the incorporation of leisure into the analysis leaves our welfare numbers roughly unaffected.

Figure 13 compares the welfare losses (or gains) implied by three utility and stochastic process specifications: consumption only (a subset of figures 12c and 12d), consumption and leisure, uncorrelated (see the discussion in the previous subsection) and consumption and leisure processes allowed to be correlated. The sub-panels distinguish risk aversion and agents belonging to different

³⁸The reason for why the leisure state is not an important determinant of the welfare losses is similar to the one for idiosyncratic consumption: we estimate labor supply and hence leisure to be not nearly as persistent as group consumption (roughly as persistent as idiosyncratic consumption).

consumption groups. All four panels contain the welfare losses of agents that start with the highest leisure state in 1972. Similar (but somewhat less pronounced) results are obtained for households starting in other leisure states in 1972.

FIGURE 13 APPROXIMATELY HERE

For separable preferences ($CRRA = 1$) correlation between consumption and leisure does not change our findings from the previous section. As households become more risk-averse, however, welfare losses arising under the correlated consumption-leisure process actually exceed the corresponding number under the consumption-only process, by up to 0.5% percentage points (see Figure 13b). This finding is due to the higher likelihood of experiencing states with low consumption and fairly low leisure, compared to the situation in which consumption and leisure were modelled as independent processes. We conclude that, while the welfare numbers based on consumption alone tell most of the story, the incorporation of leisure may, when correlated with consumption, increase welfare losses from increased inequality to a quantitatively nontrivial extent.

5. Concluding Remarks

What are the welfare consequences of the increase in inequality in the United States between 1972 and 2000? In this paper we use a standard intertemporal utility function and wage-, earnings-, consumption- and hours-worked data from the Consumer Expenditure Survey to answer this question.

Our main findings are twofold. First, welfare losses can be substantial, with significant variation of these losses across the population. Whereas households at the bottom of the consumption distribution suffer declines in welfare up to about 6% in consumption equivalent variation, households at the top end of the distribution enjoy sizeable welfare gains of similar magnitude as the losses of the poorest agents. Overall, a majority of the population (based on consumption observa-

tions, roughly sixty percent of Americans) is on the losing side. The main part of these losses arises from the increase in between-group consumption inequality. The increase in consumption/leisure correlation contributes to moderately higher welfare losses, in the order of magnitude of 0.5% for most households. Second, while welfare losses from increased inequality are fairly sizeable for a lot of US households when based on consumption data, they are not nearly as big and affect not nearly as many households as an analysis based on earnings data alone would suggest (80% to 100% of the population losing up to 10%). In addition, the estimates of welfare losses based on consumption and leisure processes are fairly robust to different values of risk aversion while those based on income processes are highly sensitive to that parameter.

In order to more precisely focus on the welfare effects of increasing inequality we have ignored two important features of our data. First, there was substantial growth in average real household consumption and second, total hours worked by a typical household increased noticeably over the last 30 years. To the extent that these trends are causally linked to the trend in inequality we have overstated (because of consumption growth) or understated (because of the decline in leisure) the welfare implications of the increase in inequality. Future empirical and theoretical work is needed to understand if, and to what extent, the trends in hours worked, in average consumption and in the cross-sectional variance of consumption are causally related.

6. Data Appendix

Our statistics are based on repeated cross-sections constructed from the Consumer Expenditure Survey (CE) Interview Surveys for the years 1972-1973, 1980-81 and 1984 to 2000, as provided by the Bureau of Labor Statistics. The 1972-1973 samples were conducted quarterly, but only annual totals were released; thus for these years we have only two cross sections, each reporting consumption and income for the year of the interview. The surveys from 1980 on were conducted on a quarterly basis, so we have four cross sections for each year. Households report consumption expenditures for the quarter preceding the interview and income data for the year preceding the interview. A fraction of the households in the survey is interviewed for four consecutive quarters and reports consumption information in every quarter, as well as income and wage information in the first and last interview. For these households we can construct annual measures of wages, earnings, consumption and number of hours from 72-73 until 2000.

A. Wages, Earnings and Hours

The definition of wages we use is earnings of the reference person divided by the number of weeks worked by the person during the year. We construct earnings of the reference person as her wages and salaries plus a fixed fraction of her self-employment farm- and non-farm income (the exact fraction is 0.864 and is taken from Diaz-Jimenez, Quadrini and Rios-Rull, 1997). Household earnings simply sums the earnings of the reference person and her spouse. Disposable household earnings are computed as household earnings minus reported federal, state and local taxes (net of refunds) and social security contributions paid by the household. We then add reported government transfers: in particular we add unemployment insurance, food stamps and welfare receipts. Since the 72-73 CE does not report the number of hours per week by households members, but only whether they are part- or full-time workers, hours worked by a person are computed as the number of weeks worked by that person times 20, if the person works part-time, or times 40 if the person

works full time. For the post-1980 sample we compare the standard deviation of hours constructed in this way with the series constructed using actual weekly hours; the two series have very similar trends (although different levels).

B. Consumption

Our definition of consumption consists of nondurables plus imputed services from durables. It includes expenditures on nondurable goods and services, expenditures on household furnishings and imputed services from houses and cars. Expenditures on nondurable goods and services include consumption expenditures for food, alcoholic beverages, tobacco, utilities, personal care, household operations, public transportation, gasoline and motor oil, apparel, education, reading, health services and miscellaneous expenditures. Each component of consumption is deflated by its corresponding monthly CPI from the Bureau of Labor Statistics.

Expenditures on household furnishings include items such as furniture, appliances and floor coverings (e.g. rugs). The reason for why we use expenditures and not imputed services is that in the CE no information is available for the value or the inventory of the stock of these furnishings, and the panel dimension of the CE is too short to carry out perpetual inventory techniques. With respect to vehicles, we impute services from cars in the following manner, following closely the procedure outlined by Cutler and Katz (1991a). From the CE data we have expenditures for purchases of new and used vehicles. We also have data on the number of cars a consumer unit possesses. For each year we first select all households that report positive expenditures for vehicle purchases, and run a regression of vehicle expenditures on a constant, age, sex and education of the reference person of the consumer unit, total consumption expenditures, excluding vehicle expenditures of the consumer, the same variable squared, total income before taxes, family size and quarter dummies. We use the estimated regression coefficients to predict expenditures for vehicles for all households in that quarter (i.e. for those who did and for those who did not report positive vehicle

expenditures). Our measure of consumption services from vehicles is then the predicted expenditure on vehicles, times the number of vehicles the consumer unit owns, times $\frac{1}{32}$ (reflecting the assumption of average complete depreciation of a vehicle after 32 quarters) plus other expenditures for cars, such as insurance, maintenance and finance charges. With respect to housing services the CE provides information on rent paid for the residence of the consumer unit, including insurance and other out of pocket expenses paid by the renter. To impute housing services for those consumer households that own their residence we use a variable from the CE that measures the market rent (as estimated by the reference person of the consumer unit) the residence would command if rented out.³⁹ This variable is not available for the years 1980-81 and 1993-94; for these years we do not compute inequality measures for ND+ consumption expenditures.⁴⁰ As with nondurable consumption, all imputed services from consumer durables and housing are deflated with the corresponding CPI.

³⁹The exact question that the reference person of the CU is asked is “If you were to rent your home today, how much you think it would rent for monthly, without furnishings and utilities?”

⁴⁰We experimented with using an imputation procedure similar to the one used for vehicles in order to obtain housing services for the four missing years. Results were very similar and are available upon request.

References

- [1] Altug, S. and R. Miller (1990), "Household Choices in Equilibrium," *Econometrica*, 58, 543-570.
- [2] Attanasio, O. and S. Davis (1996), "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104, 1227-1262.
- [3] Attanasio, O. (2002) "Consumption and Income Inequality: What do we know and what we can learn from it.", Mimeo UCL
- [4] Blundell, R. and T. MaCurdy (1999), "Labor Supply: a Review of Alternative Approaches," in: O. Ashenfelter and D. Card (eds) *Handbook of Labor Economics*, North Holland, Amsterdam.
- [5] Blundell, R. and I. Preston (1998), "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*, 113, 603-640.
- [6] Cochrane, J. (1991), "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99, 957-976.
- [7] Cox, M. and R. Alm (1999), *Myths of Rich & Poor: Why We're Better Off Than We Think*, Basic Books, New York.
- [8] Cullen, J. and J. Gruber (2000), "Does Unemployment Insurance Crowd Out Spousal Labor Supply?," *Journal of Labor Economics*, 18, 546-572.
- [9] Cutler, D. and L. Katz (1991a), "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s," *American Economic Review*, 82, 546-551.
- [10] Cutler, D. and L. Katz (1991b), "Macroeconomic Performance and the Disadvantaged," *Brookings Papers on Economic Activity*, 1-74.
- [11] Deaton, A. (1992), *Understanding Consumption*, Clarendon Press, Oxford.

- [12] Díaz-Giménez, J., V. Quadrini and J.V. Ríos-Rull (1997), “Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 21, 3-21.
- [13] Dynarski, S. and J. Gruber (1997), “Can Families Smooth Variable Earnings?,” *Brookings Papers on Economic Activity*, 229-284.
- [14] Greenspan, A. (1998), “Opening Remarks to the Symposium on Income Inequality: Issues and Policy Options,” Federal Reserve Bank of Kansas City.
- [15] Gottschalk, P. and R. Moffitt (1994), “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 217-272.
- [16] Hall, R. and F. Mishkin (1982), “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households,” *Econometrica*, 50, 461-481.
- [17] Hayashi, F., J. Altonji, and L. Kotlikoff (1996), “Risk Sharing Between and Within Families,” *Econometrica*, 64, 261-294.
- [18] Heathcote, J., K. Storesletten and G. Violante (2003), “The Welfare Implications of Rising Wage Inequality in the US,” *Mimeo*, IIES, Stockholm University.
- [19] Heckman, J. and T. MaCurdy (1980), “A Life Cycle Model of Female Labor Supply,” *Review of Economic Studies*, 47, 47-74.
- [20] Johnson, D. and S. Shipp (1991), “Trends in Inequality using Consumer Expenditures: 1960 to 1993,” *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 1-7.
- [21] Johnson, D. and T. Smeeding (1998), “Measuring the Trends in Inequality of Individuals and Families: Income and Consumption,” *mimeo*.

- [22] Juhn, C., K. Murphy and R. Topel (2002) "Current unemployment, historically contemplated", *Brooking Papers on Economic Activity*, 1, 79-136
- [23] Katz, L. and D. Autor (1999), "Changes in Wages Structure and Earnings Inequality," in: O. Ashenfelter and D. Card (eds) *Handbook of Labor Economics*, North Holland, Amsterdam.
- [24] Krueger, D. and F. Perri (2002), "Does Income Inequality Lead to Consumption Inequality: Evidence and Theory," *NBER working paper 9202*.
- [25] Krusell, P. and A. Smith (1999), "On the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, 2, 245-272.
- [26] Lucas, R. (1987), *Models of Business Cycles*, Basil Blackwell, New York.
- [27] Mace, B. (1991), "Full Insurance in the Presence of Aggregate Uncertainty," *Journal of Political Economy*, 99, 928-56.
- [28] Mayer, S. and C. Jencks (1993), "Recent Trends in Economic Inequality in the United States: Income versus Expenditures versus Well-Being," in Papadimitriou, D and E. Wolff (eds.), *Poverty and Prosperity in the USA in the Late Twentieth Century*, St. Martin's Press, New York.
- [29] Nelson, J. (1994), "On Testing for Full Insurance Using Consumer Expenditure Survey Data," *Journal of Political Economy*, 102, 384-394.
- [30] Slesnick, D. (1993), "Gaining Ground: Poverty in the Postwar United States," *Journal of Political Economy*, 101, 1-38.
- [31] Slesnick, D. (2001), *Consumption and Social Welfare*, Cambridge University Press, Cambridge.

- [32] Stephens, M. (2002), “Worker Displacement and the Added Worker Effect,” *Journal of Labor Economics*, 20, 504-537.
- [33] Tauchen, G. and R. Hussey (1991), “Quadrature-Based Methods for Obtaining Approximate Solutions to Non Linear Asset Pricing Models,” *Econometrica*, 59, 371-396.
- [34] Townsend, R. (1994), “Risk and Insurance in Village India,” *Econometrica*, 62, 539-591.

Table 1. The persistence of the estimated processes*

Disposable Earnings		
	Within-group	Between-group
Value	0.76	0.97
Consumption		
	Within-group	Between-group
Value	0.72	0.98

*The reported measure of persistence is the value of the second largest eigenvalue of the transition matrices π^d and π^g

Figure 1. Standard Deviation of Log Weekly Wages

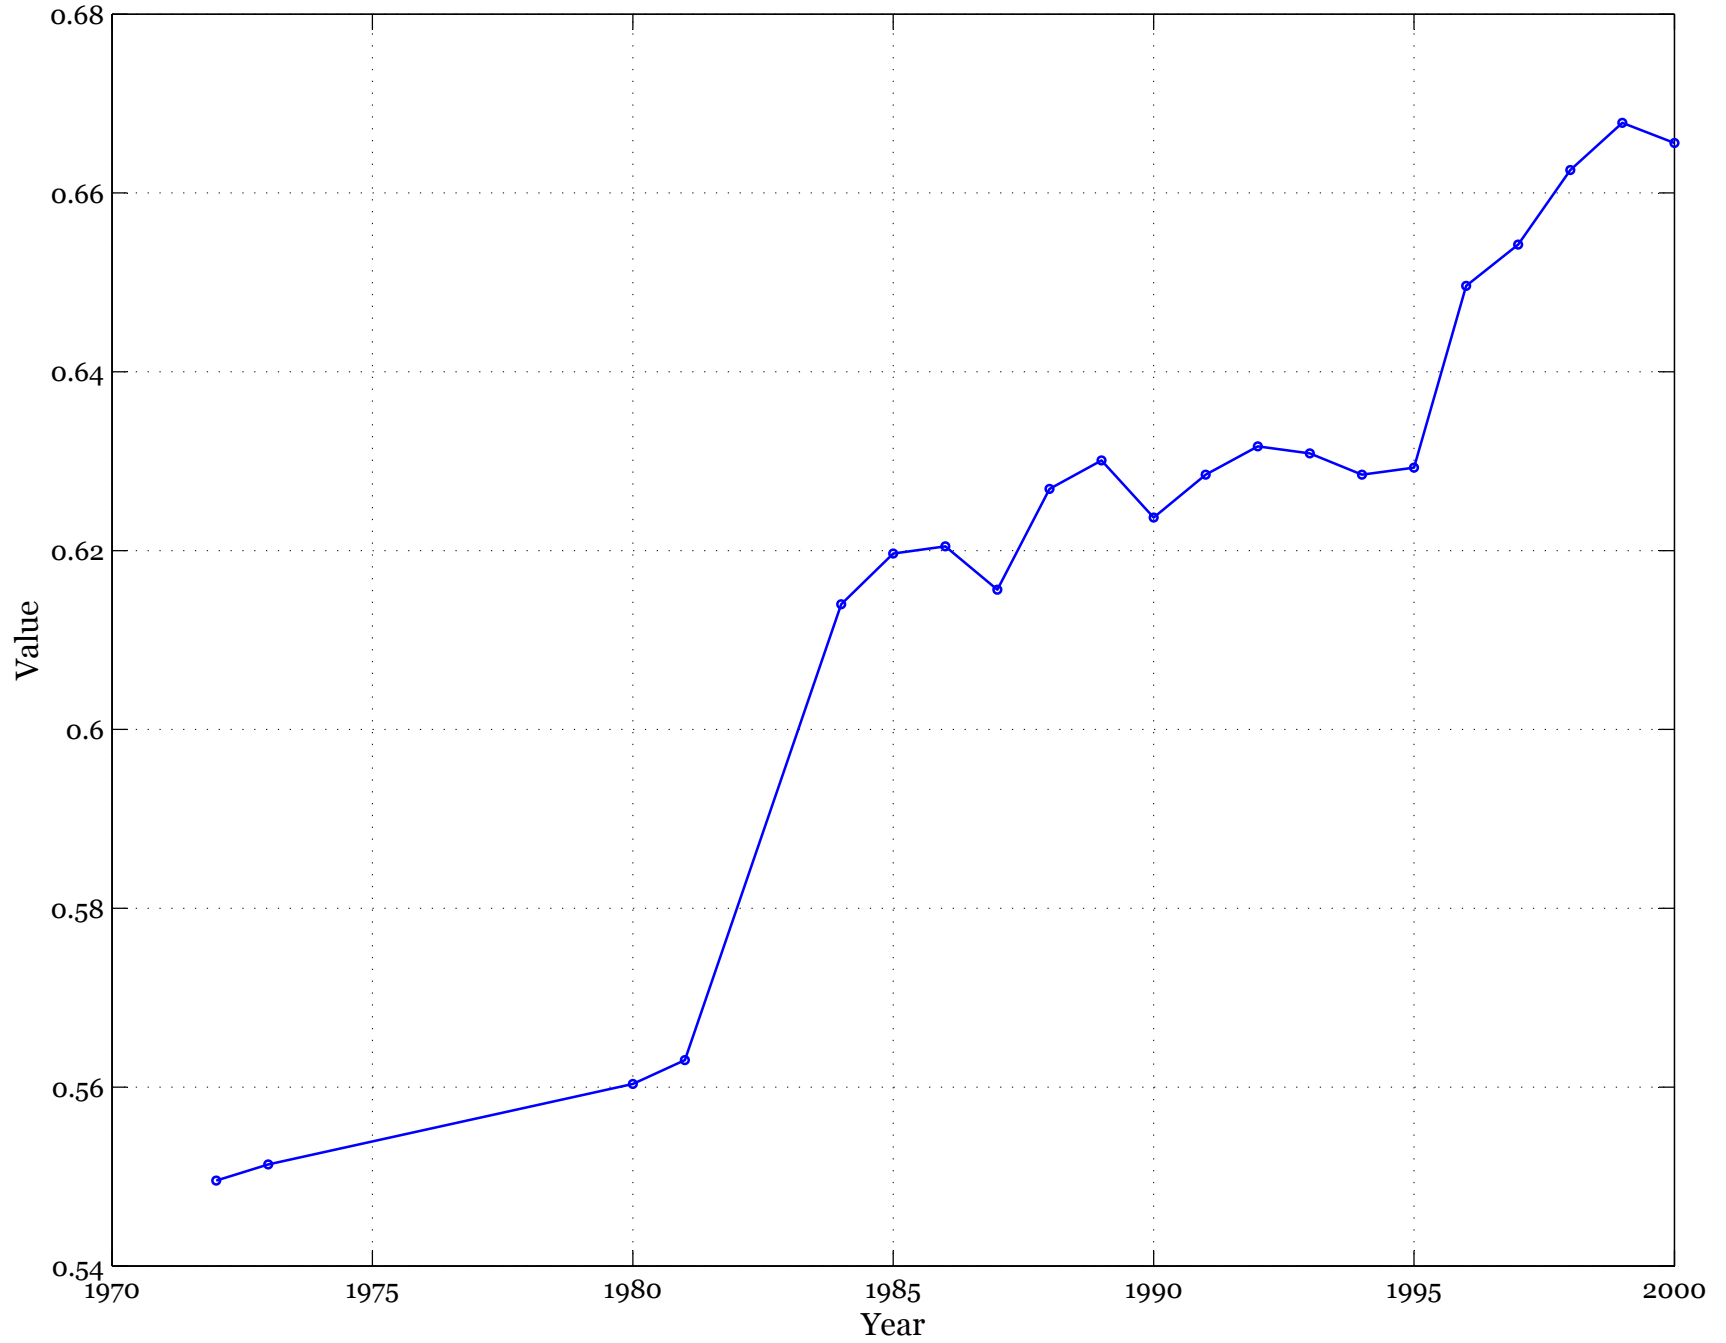


Figure 2. Wage inequality from CE and CPS

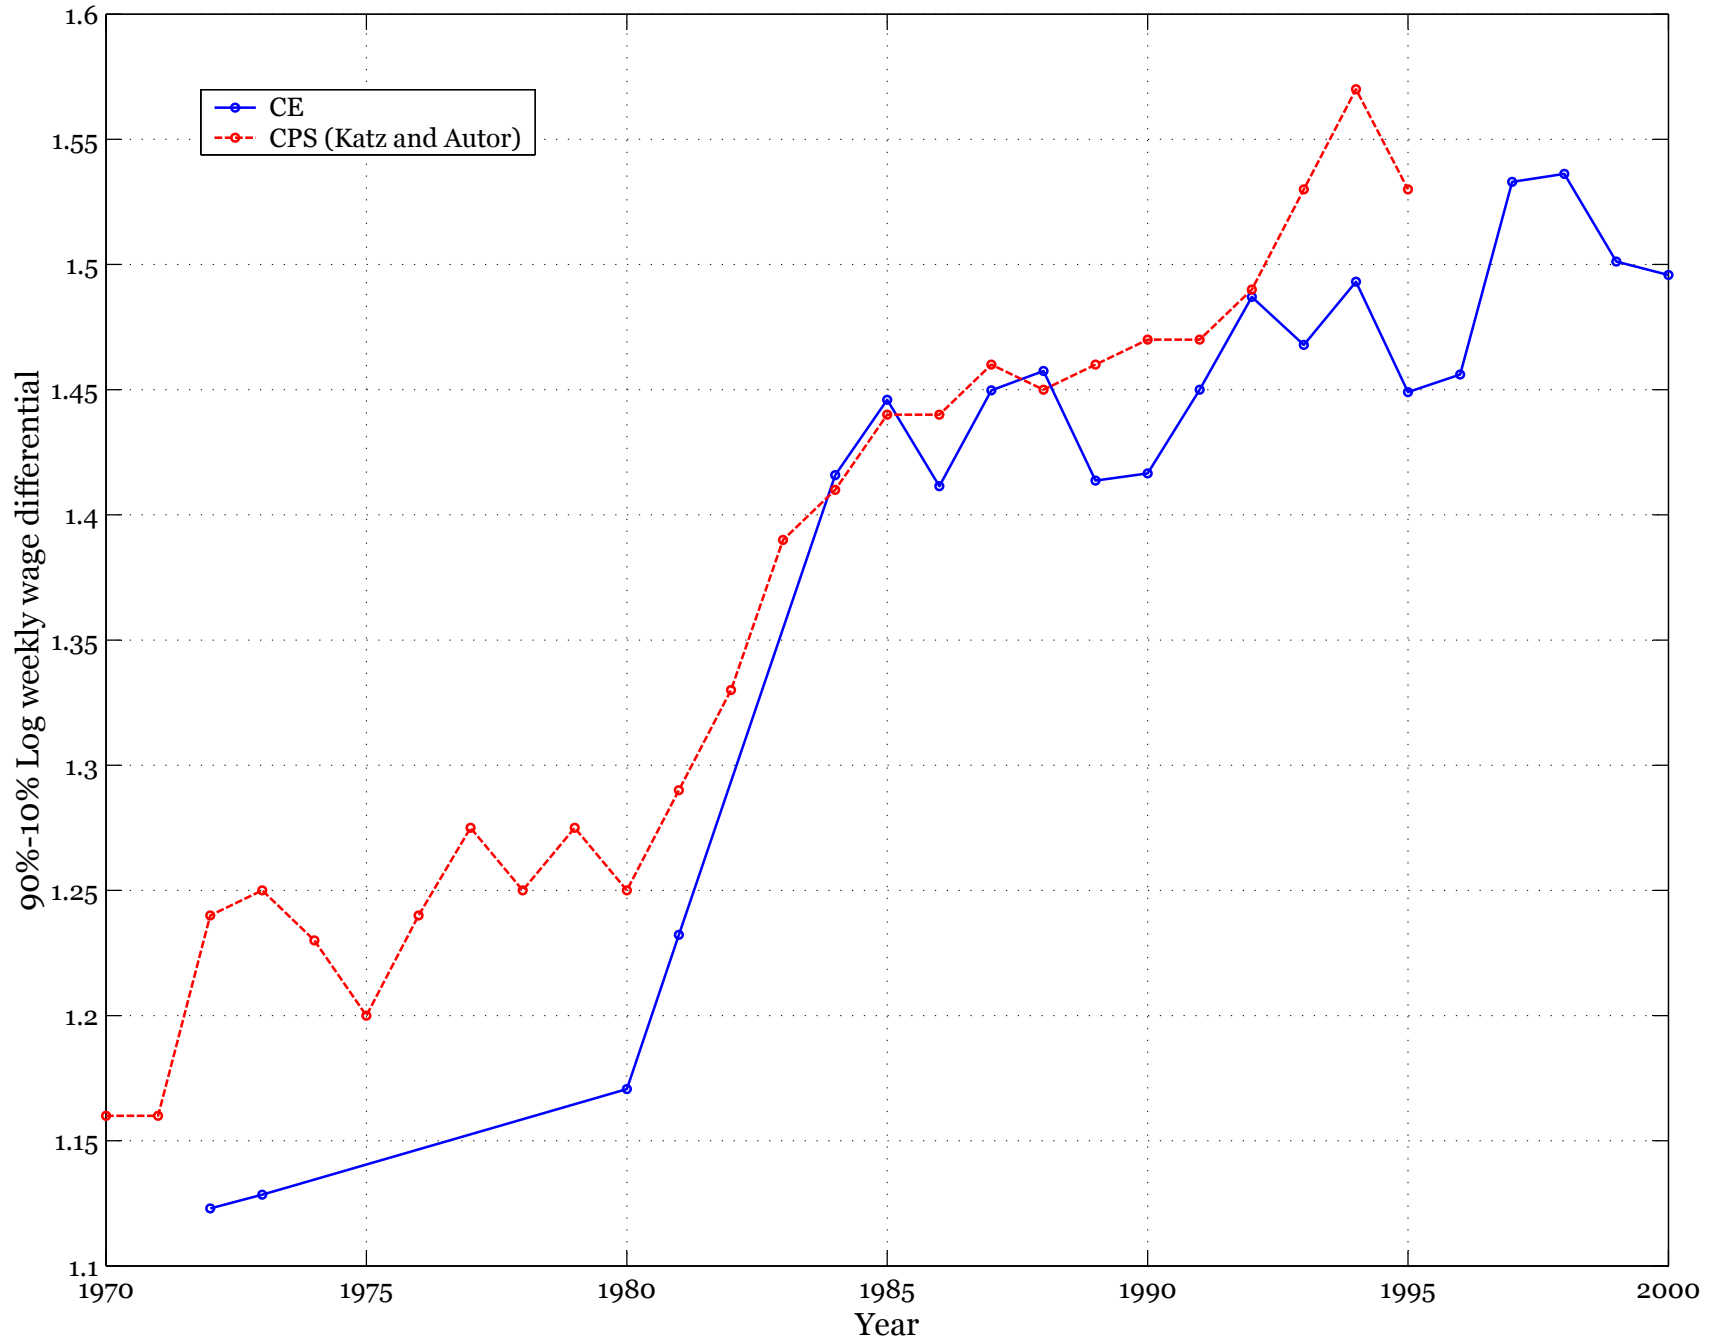


Figure 3. Change in wage inequality from CE, CPS and PSID

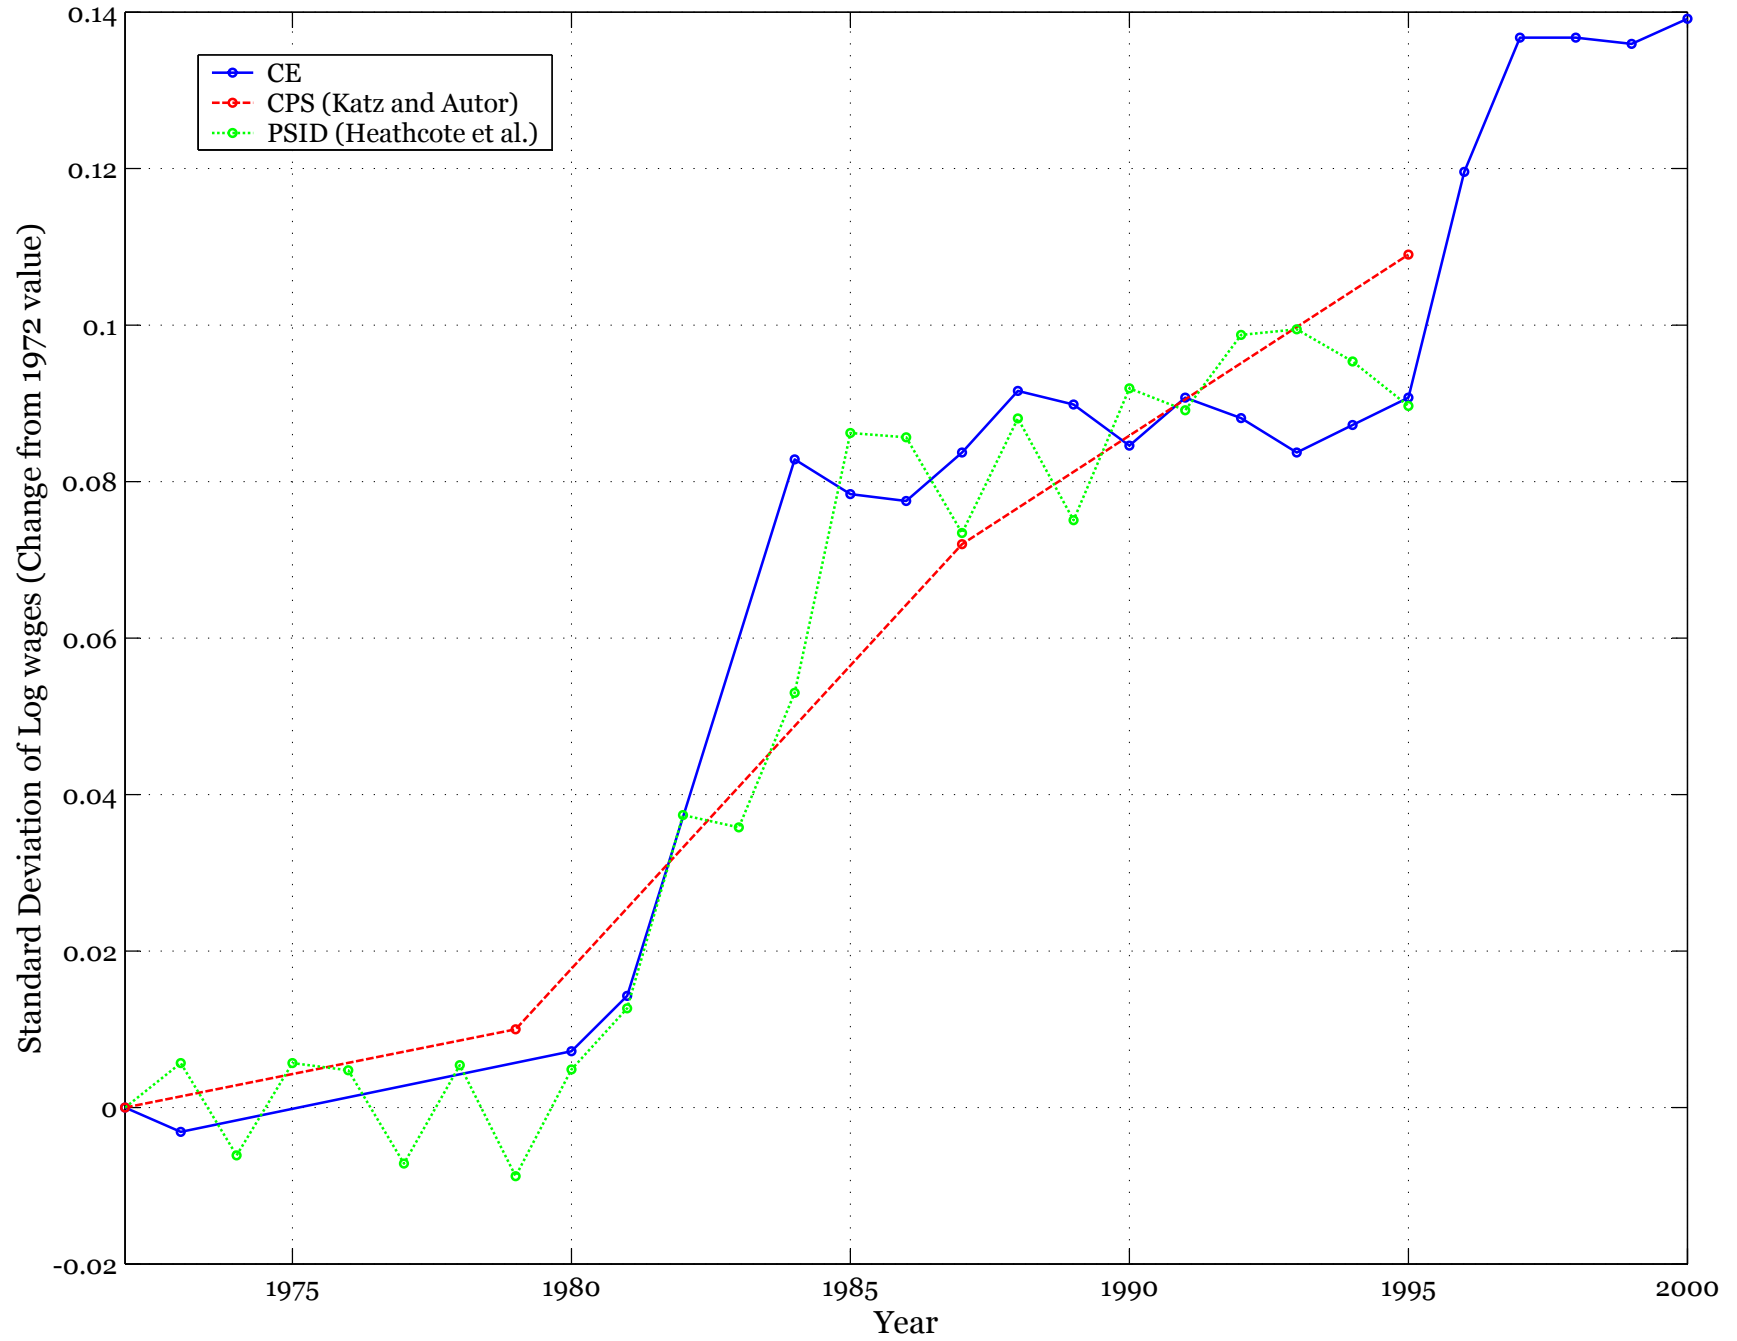


Figure 4. Standard Deviation of Log Total and Disposable Earnings



Figure 5. Decomposition of Disposable Earnings Inequality

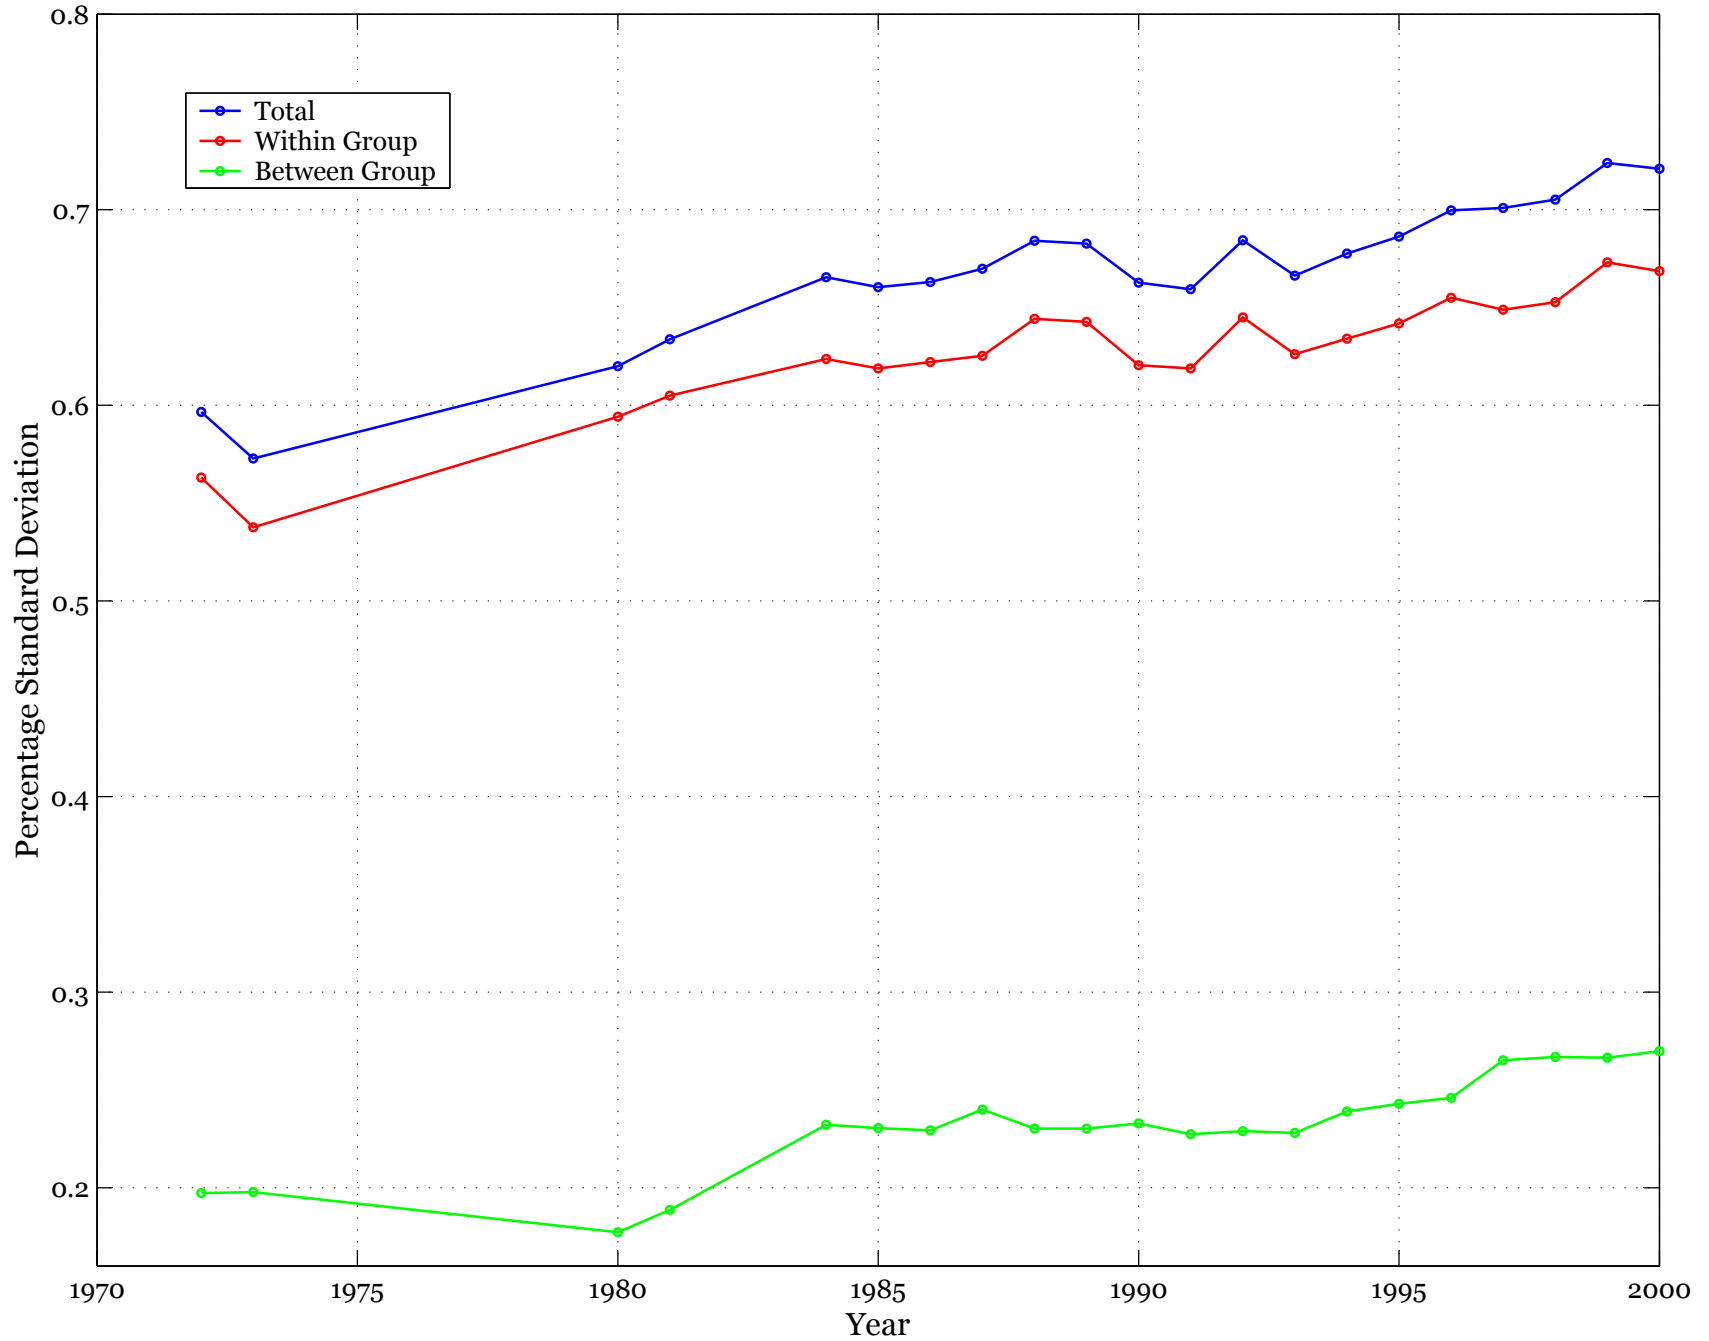


Figure 6. Standard Deviation of Log Disposable Earnings and Log Consumption

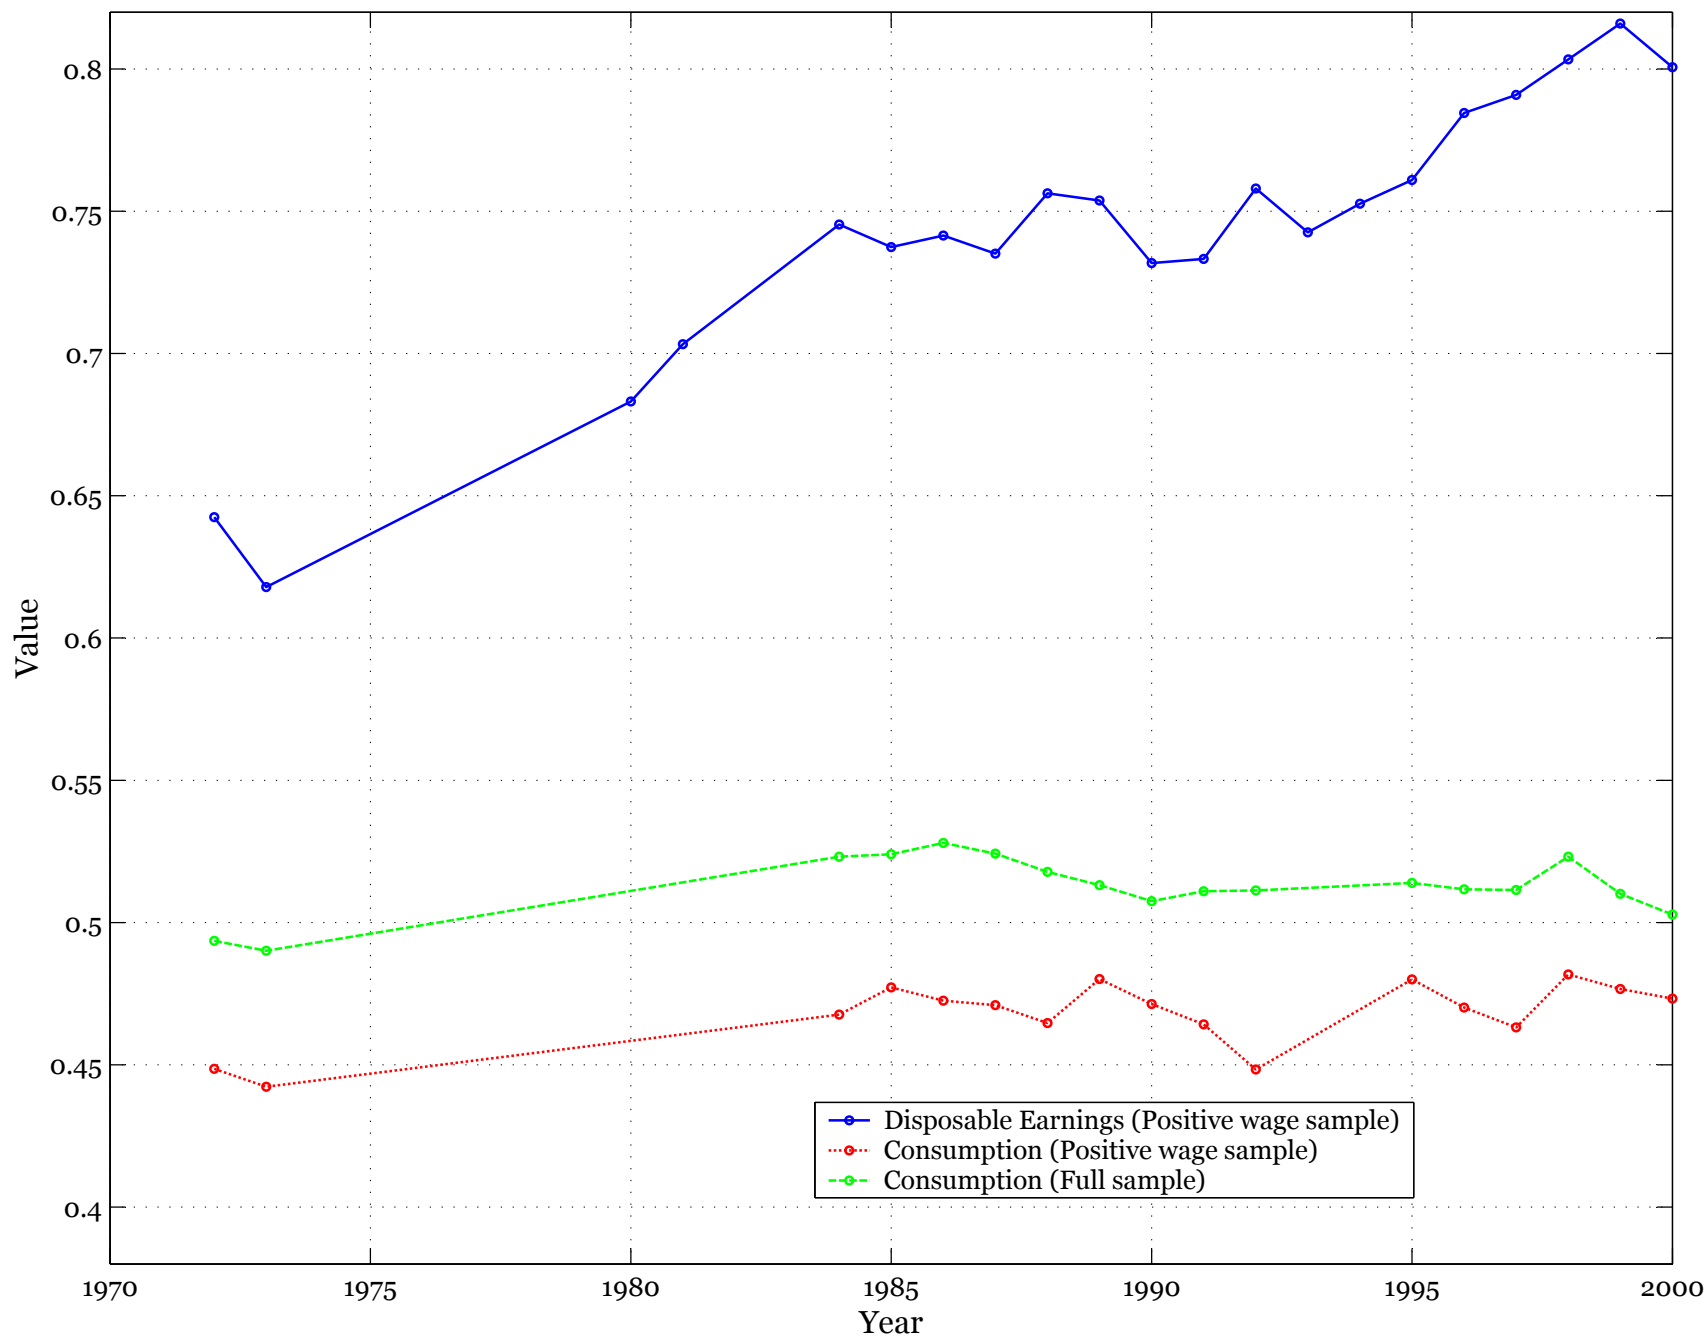


Figure 7. Consumption by selected quintiles of the disp. earnings distribution

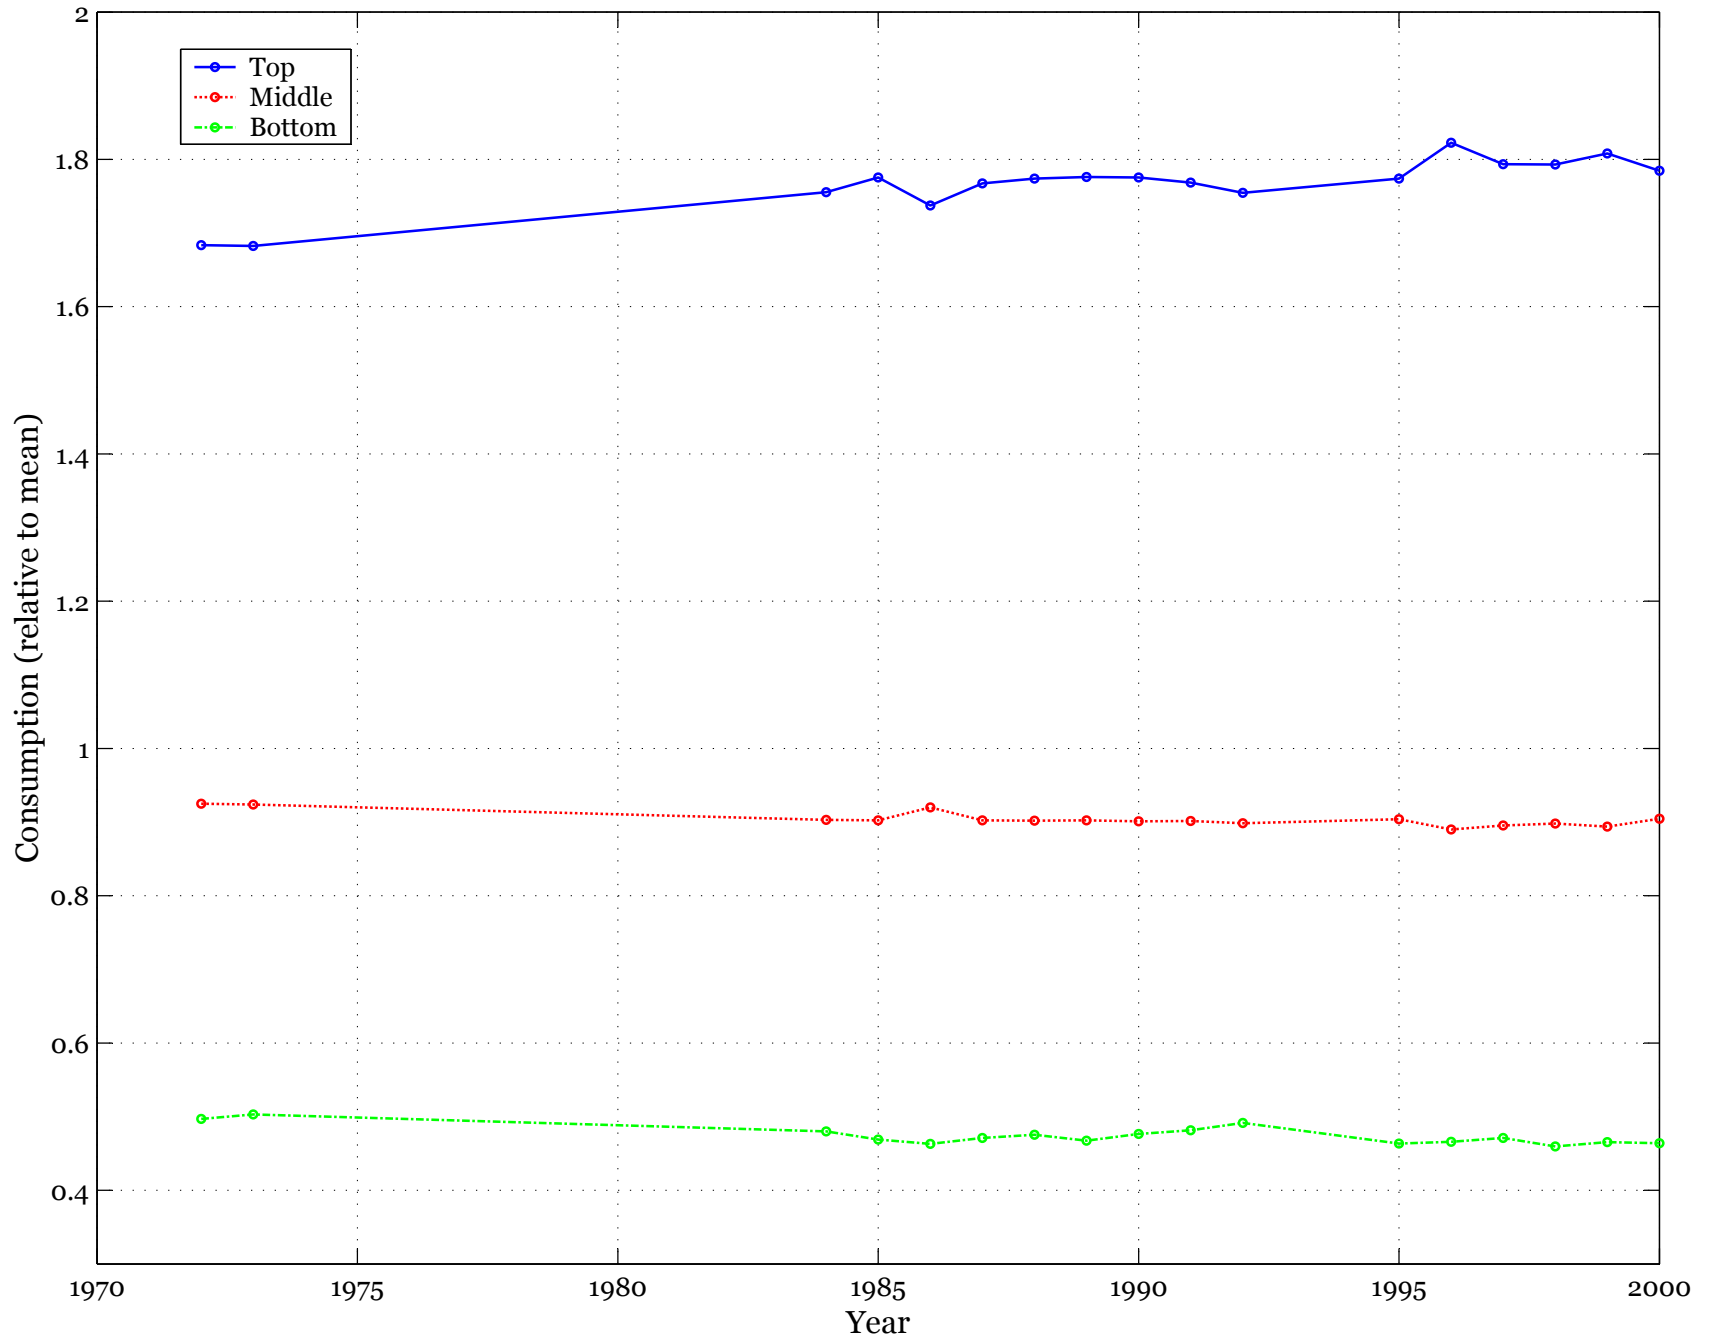


Figure 8. Decomposition of Consumption Inequality

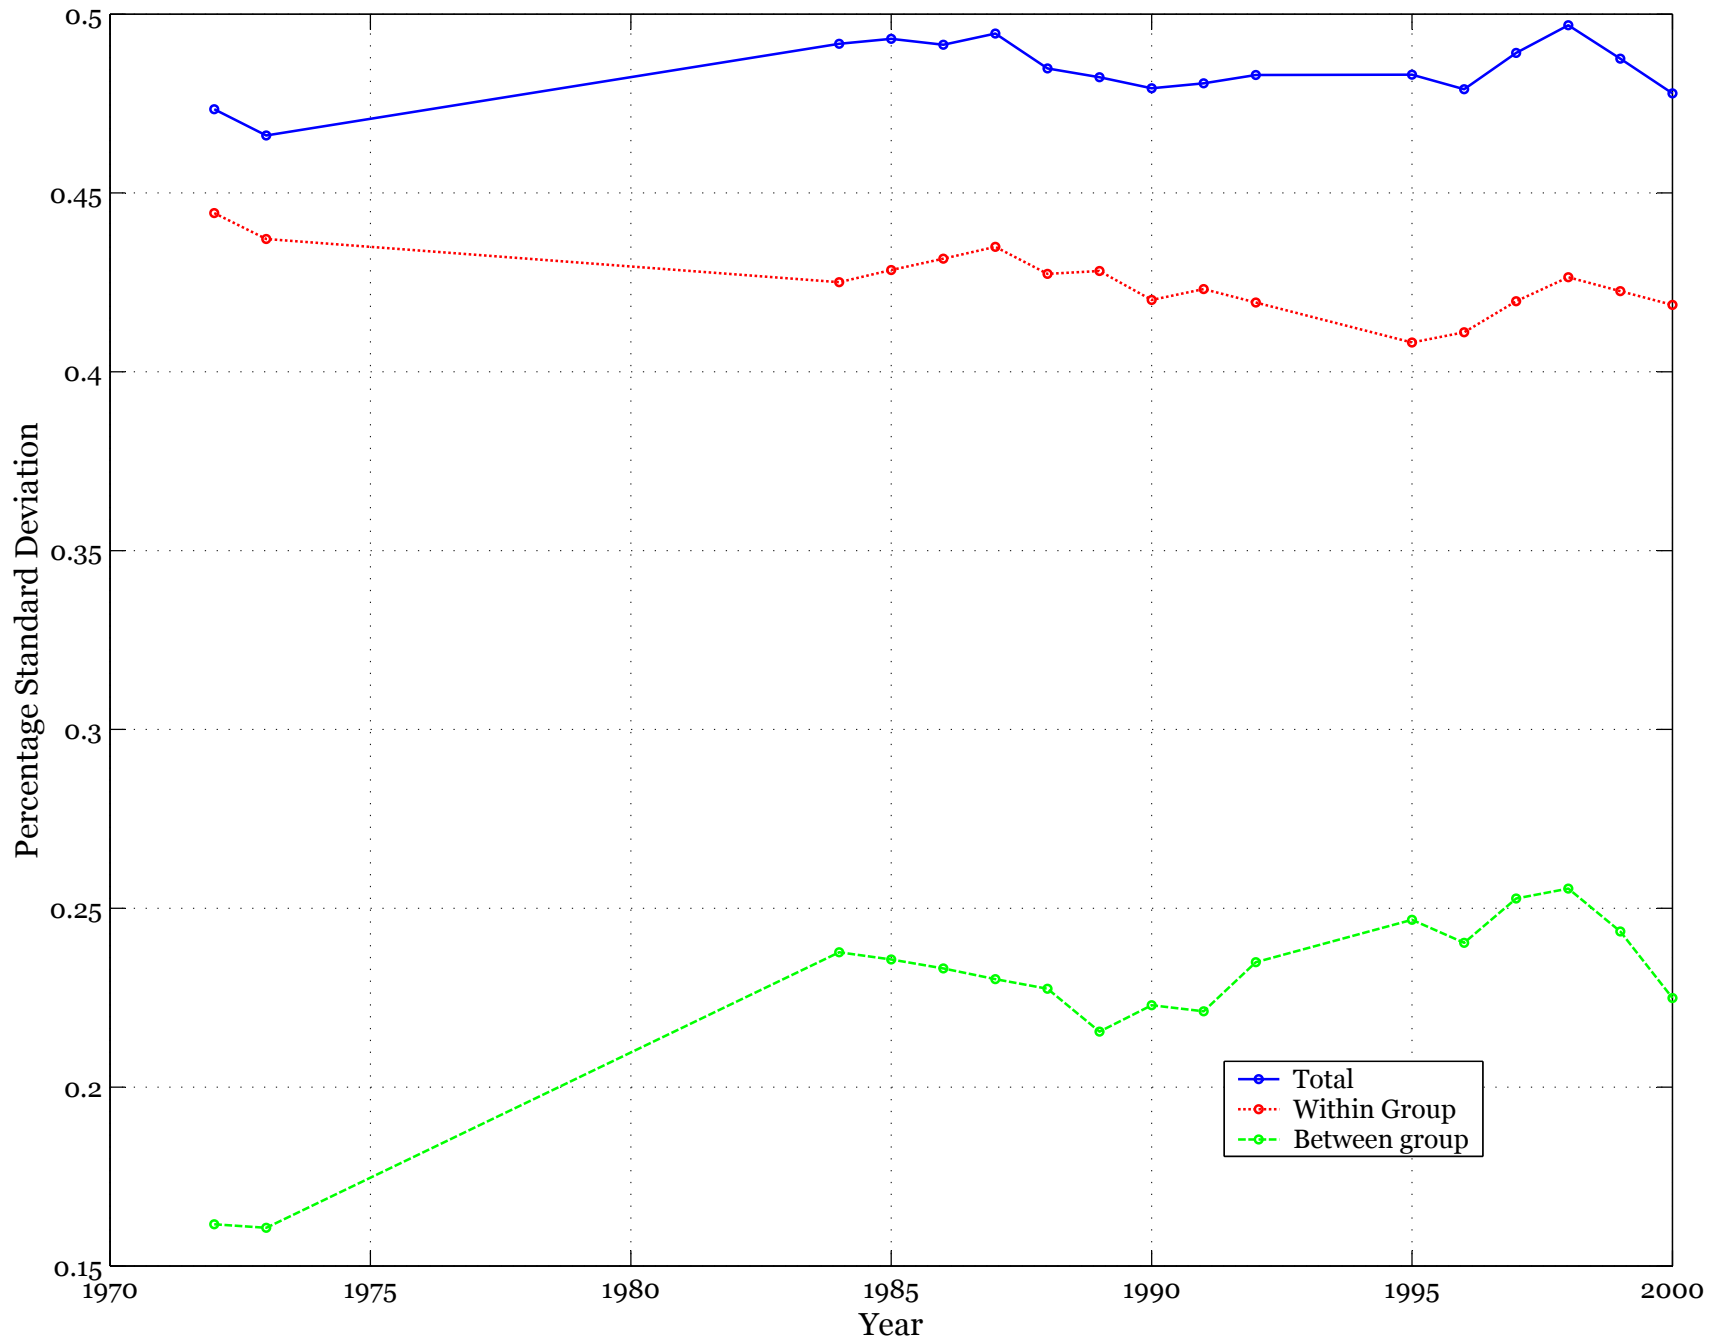


Figure 9. Standard Deviation of Log Wages and of Log per adult hours

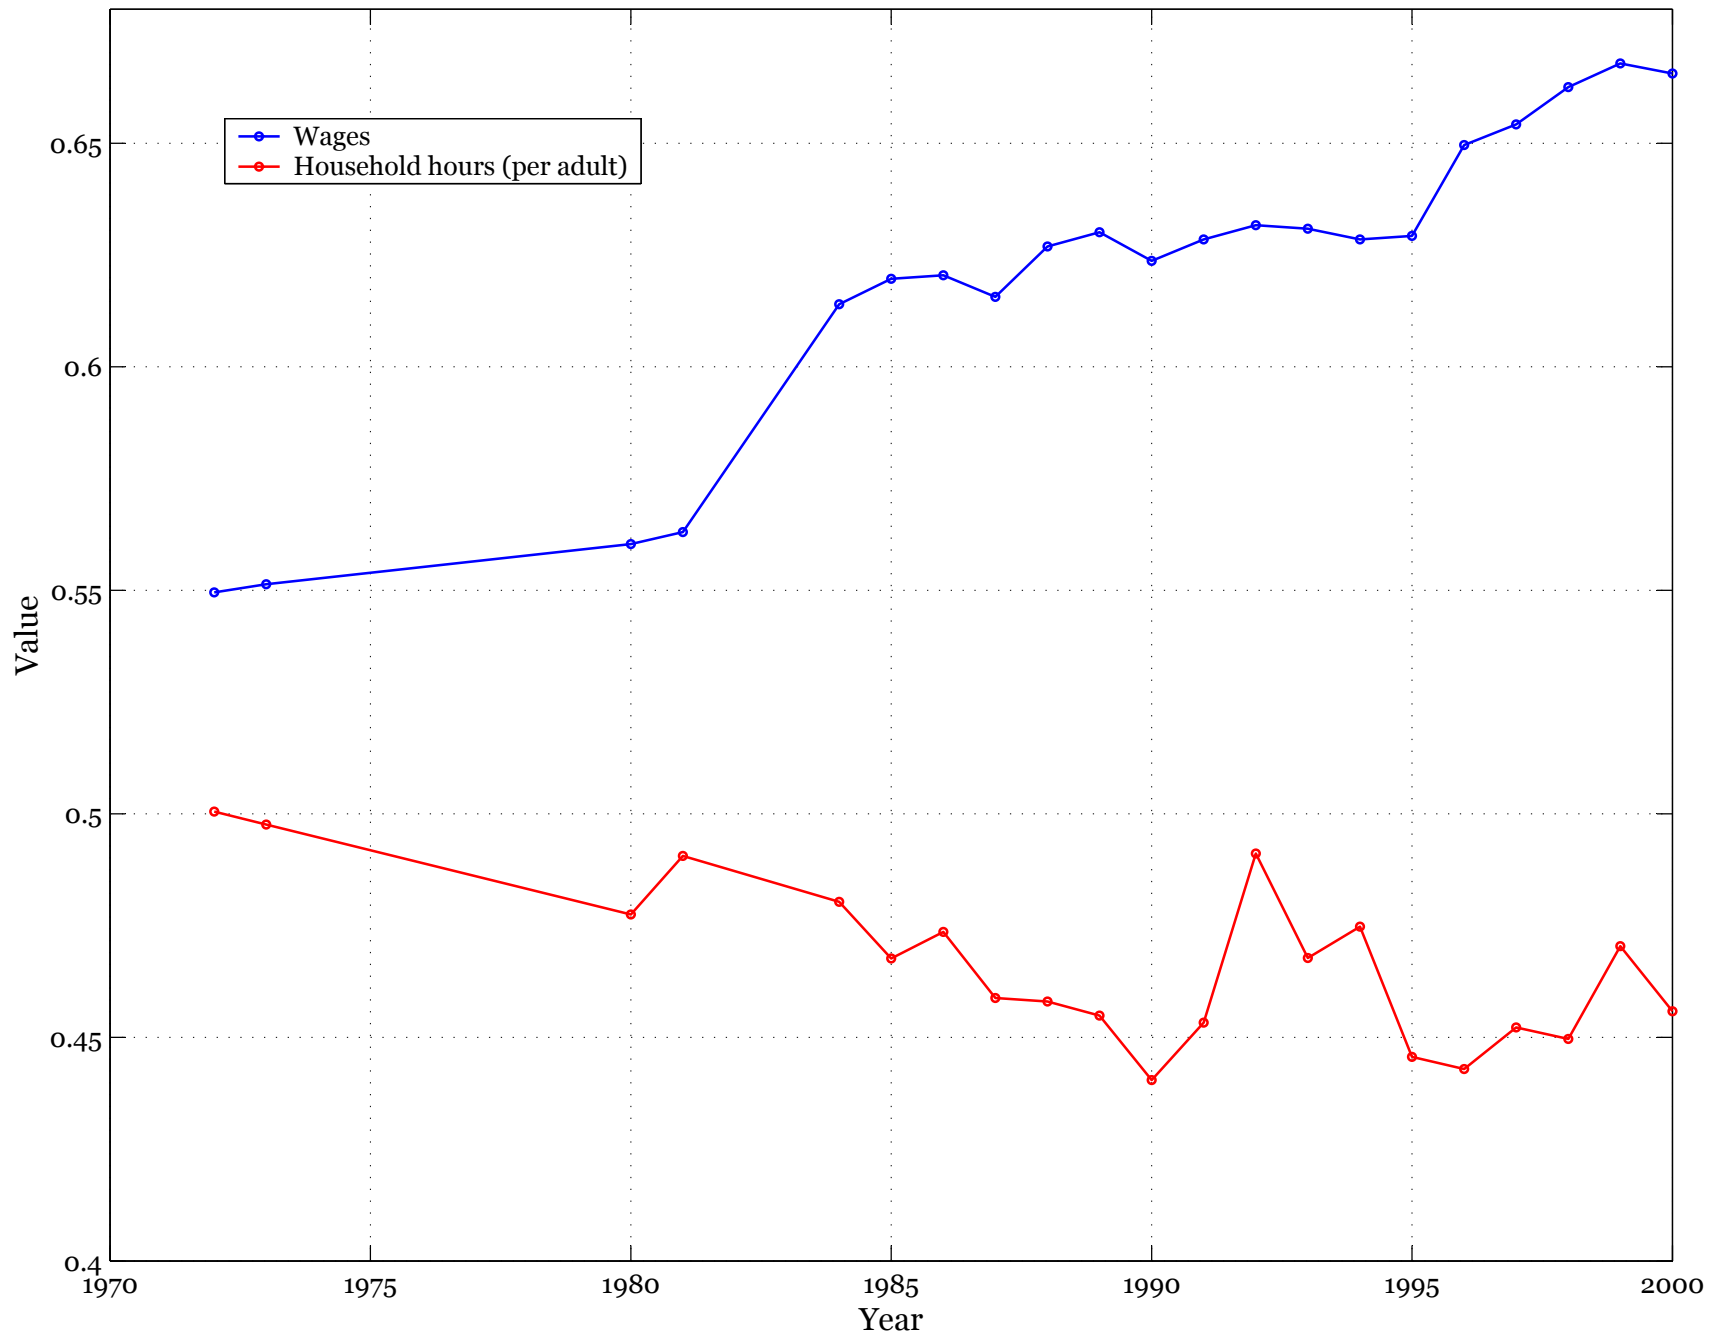


Figure 10. Yearly hours worked per person(average and by selected quintiles of the consumption distribution)

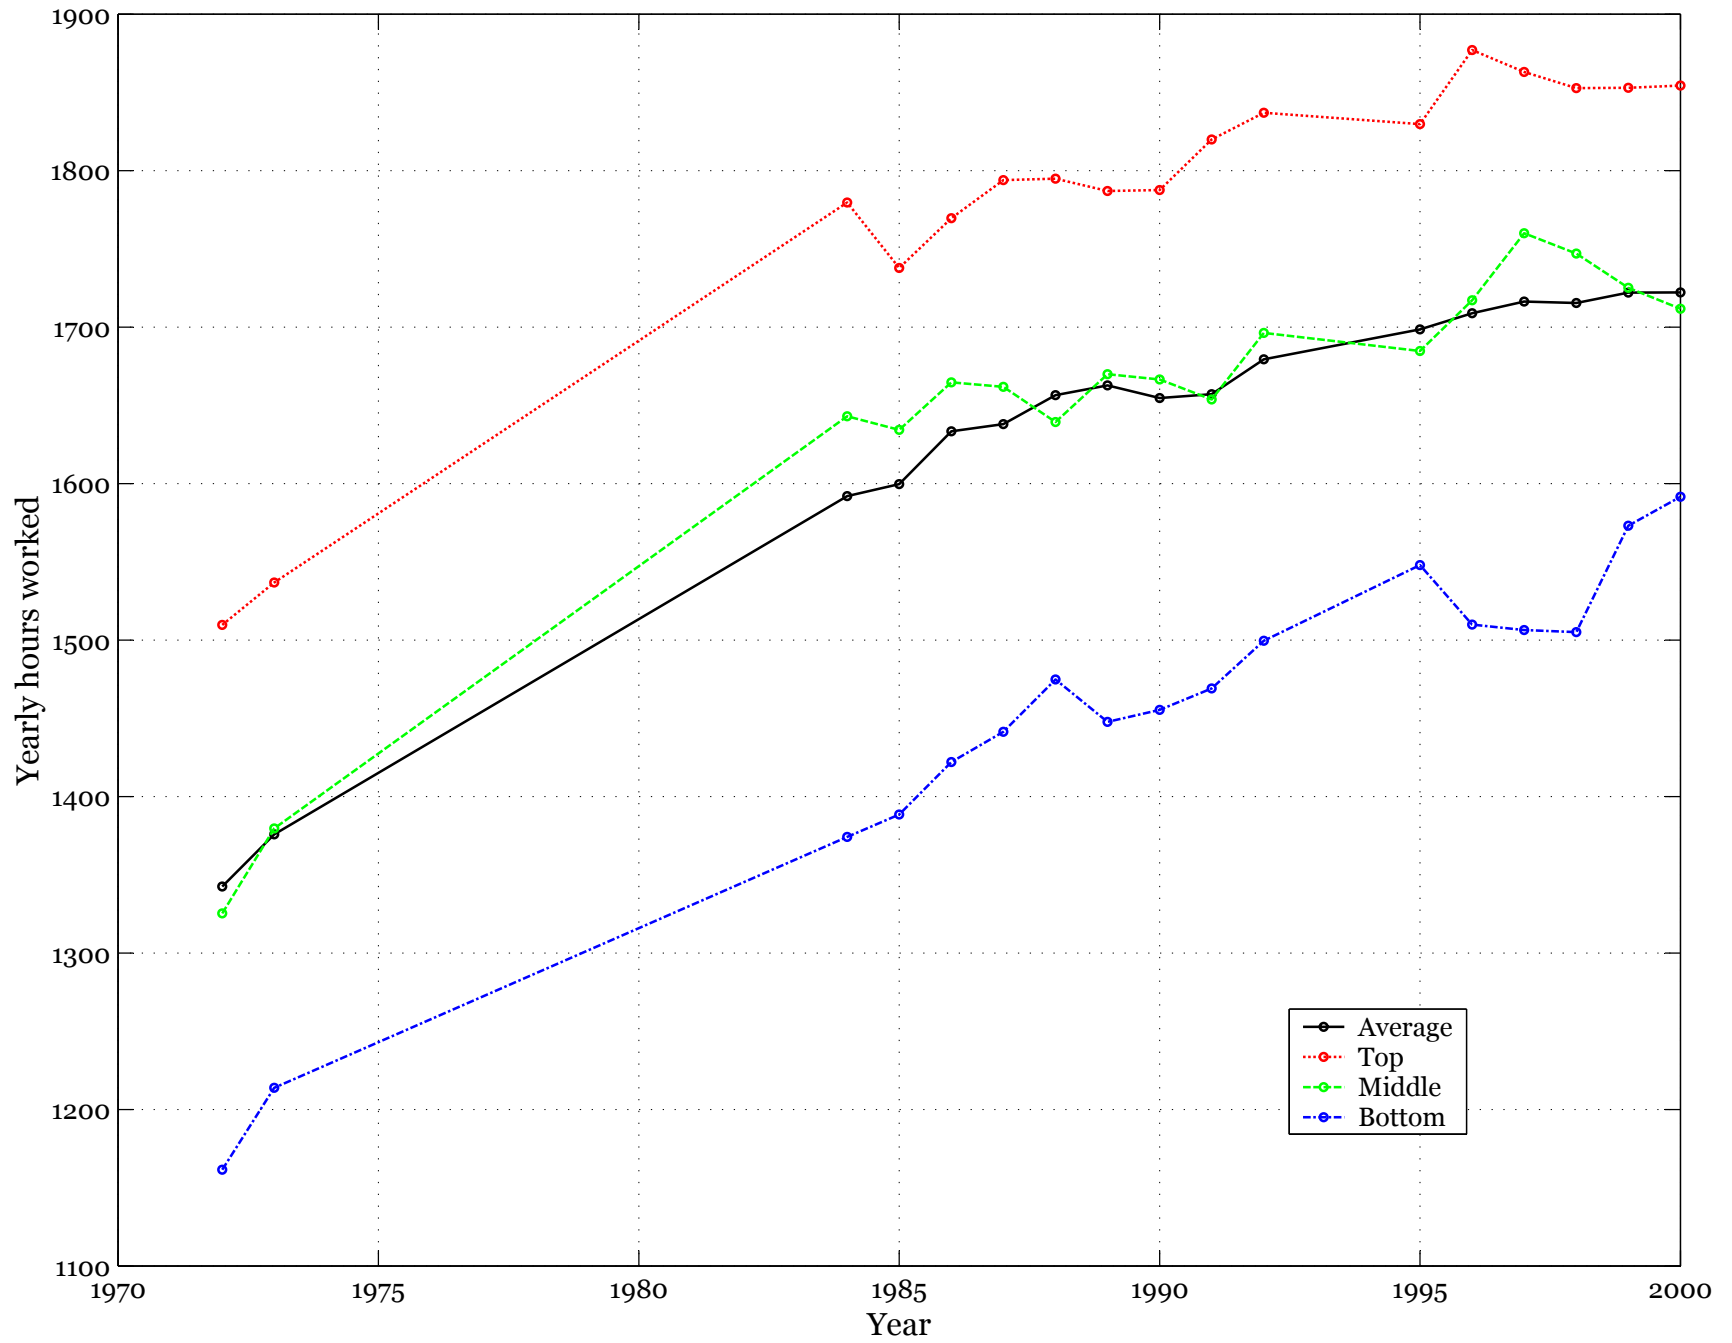


Figure 11. Correlation of hours worked and consumption

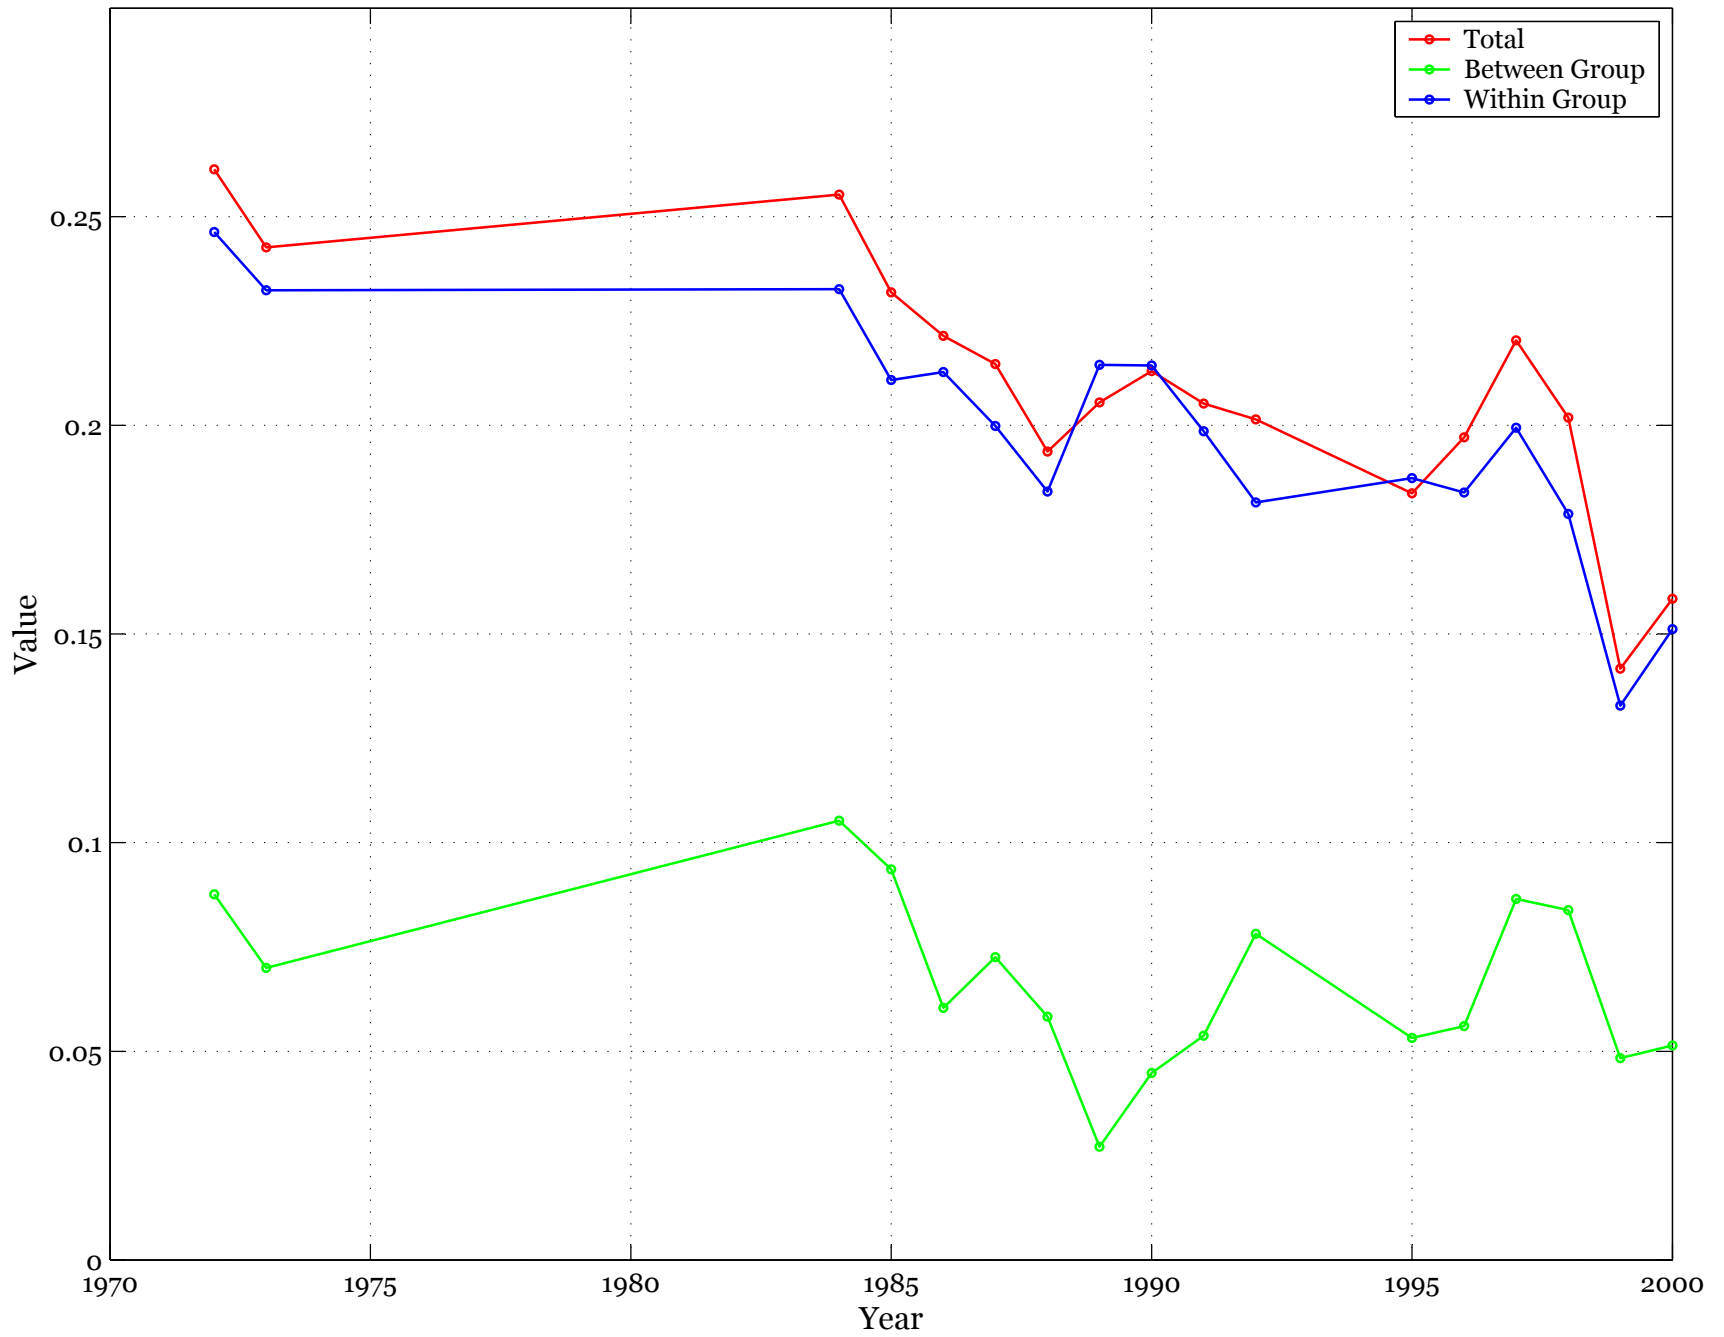
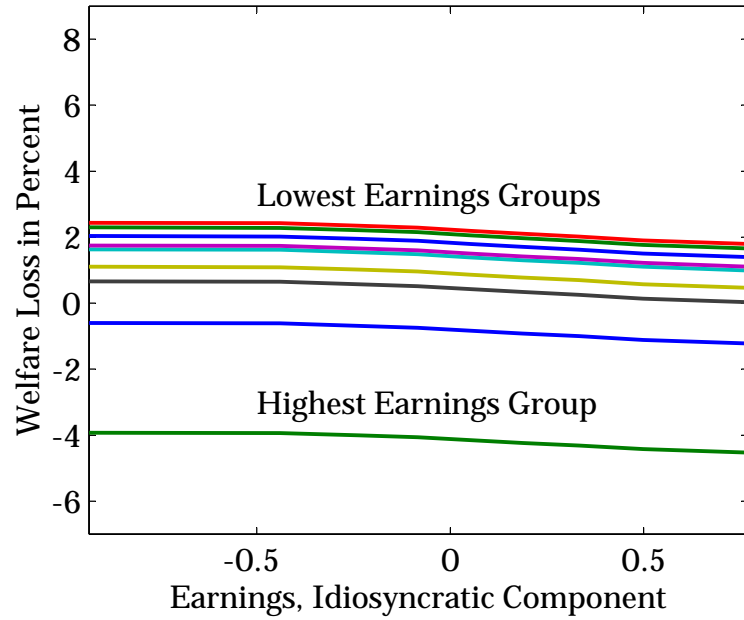
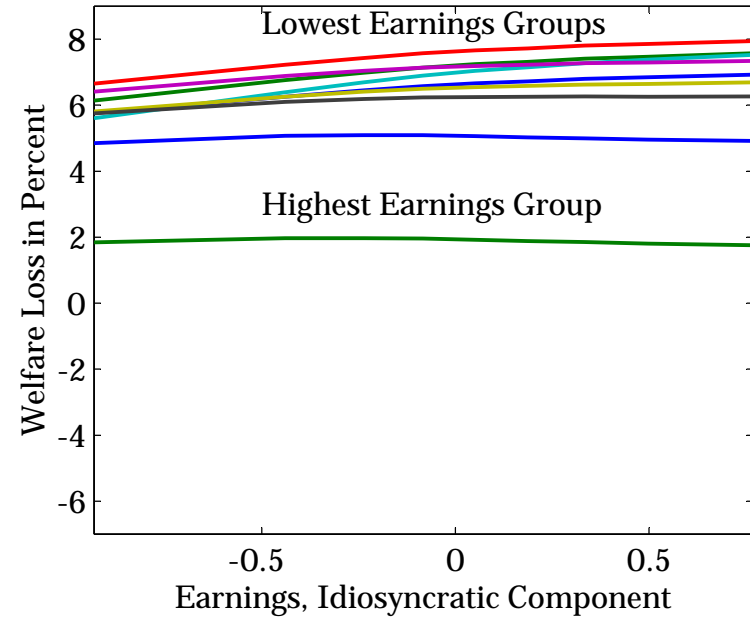


Figure 12. Welfare Losses

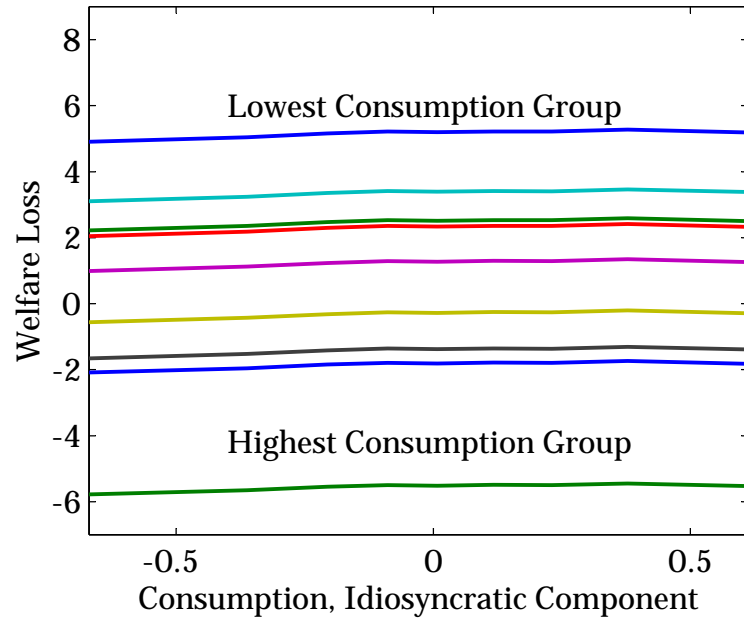
12a. Using Earnings ($\sigma=1$)



12b. Using Earnings ($\sigma=2$)



12c. Using Consumption ($\sigma=1$)



12d. Using Consumption ($\sigma=2$)

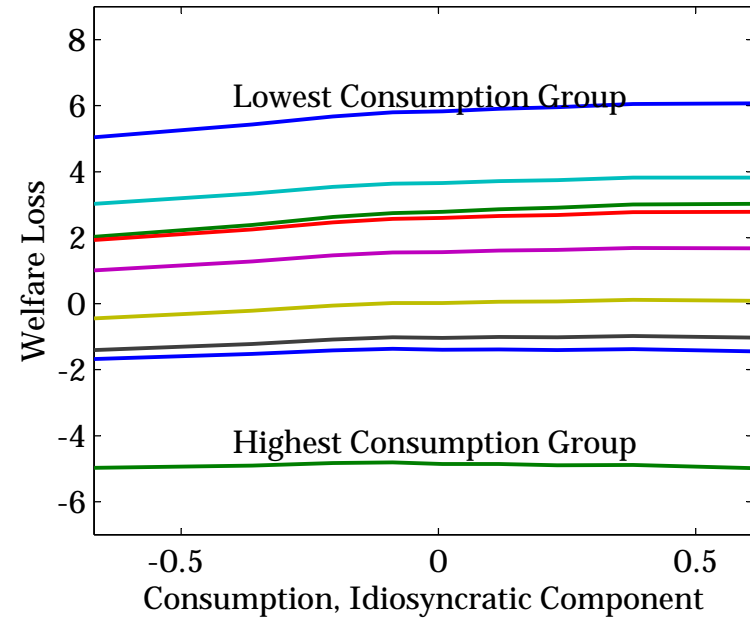
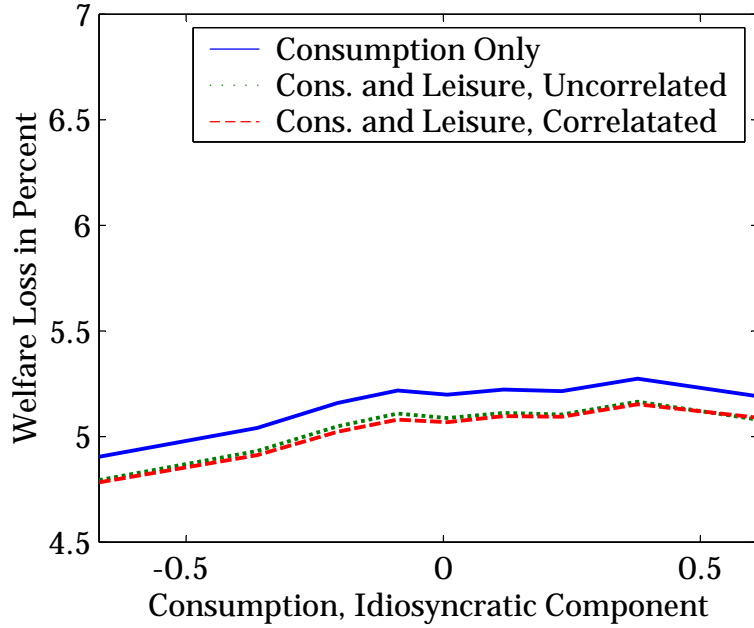


Figure 13. Welfare Losses

13a. Lowest Cons. Group (CRRA=1)



13b. Lowest Cons. Group (CRRA=2)

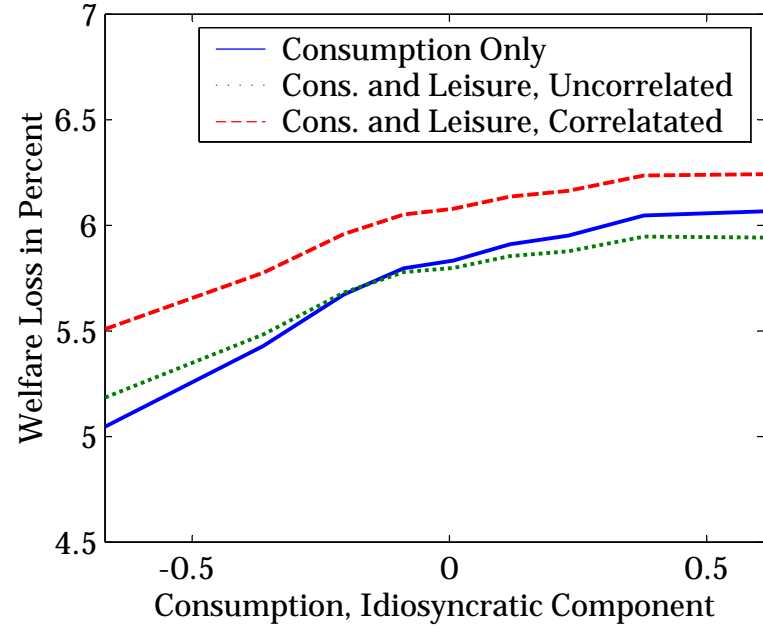
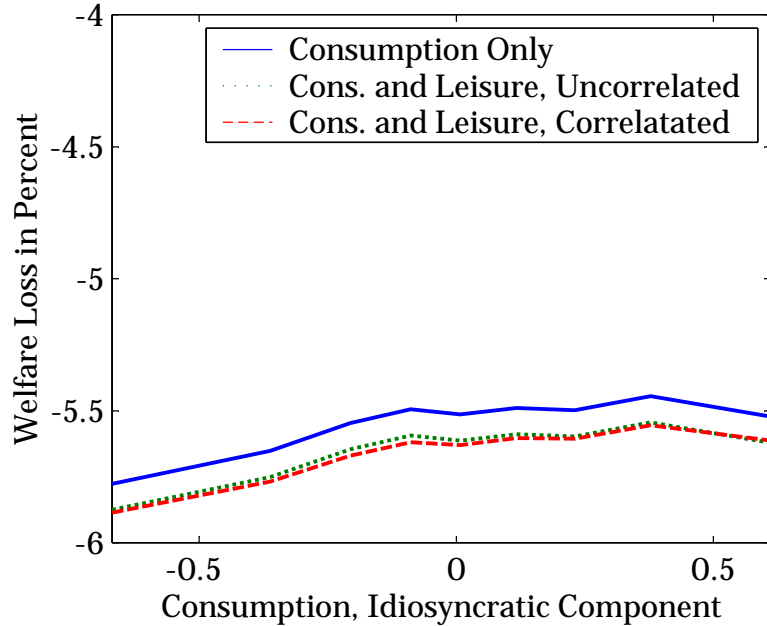


Figure 13c. Highest Cons. Group (CRRA=1)



13d. Highest Cons. Group (CRRA=2)

