Volume Title: NBER Macroeconomics Annual 2003, Volume 18

Volume Author/Editor: Mark Gertler and Kenneth Rogoff, editors

Volume Publisher: The MIT Press

Volume ISBN: 0-262-07253-X

Volume URL: http://www.nber.org/books/gert04-1

Conference Date: April 4-5, 2003

Publication Date: July 2004

Title: On the Welfare Consequences of the Increase in Inequality in the United States

Author: Dirk Krueger, Fabrizio Perri

URL: http://www.nber.org/chapters/c11442

Dirk Krueger and Fabrizio Perri

UNIVERSITY OF PENNSYLVANIA, NBER, AND CEPR; AND NEW YORK UNIVERSITY, FEDERAL RESERVE BANK OF MINNEAPOLIS, NBER, AND CEPR

On the Welfare Consequences of the Increase in Inequality in the United States

1. Introduction

The sharp increase in wage and earnings inequality in the United States 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. 1467). The objective of this paper is to analyze the welfare consequences of this sharp change in the wage distribution and the associated change in the earnings distribution.[1]

Our interest in welfare immediately forces us to look beyond the distribution of current wages. If a household's economic welfare depends on consumption and leisure enjoyed over that household's 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,

We thank Steven Davis, Mark Gertler, Ken Rogoff, and Kjetil Storesletten and seminar participants at UCLA and at the 2003 NBER Macroeconomics Annual conference for their useful comments. Krueger thanks the National Science Foundation for financial support under grant SES 0004376. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Our data stretches from 1972 to 2000. When we refer to the increase in inequality over the last 30 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 30 years.

transfers among extended family members, friends, or the government may augment disposable income and thus increase consumption opportunities of families. Second, if a significant fraction of the variations in wages, and thus income, is due to variations in its transitory component and if 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 because high consumption in the light of low wages may be realized at the expense of long working hours and thus little leisure for household members. 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 United States.[2]

Our analysis is divided into two parts. In the first part we use data from the Consumer Expenditure (CE) survey for the years 1972–2000 to document the evolution of the U.S. cross-sectional variability of individual wages, household total earnings and disposable earnings, hours worked, and consumption. We find a substantial increase in wage inequality, total earnings inequality, and disposable earnings inequality. Total household earnings inequality increases slightly less than individual wage inequality, suggesting that longer hours might be used to compensate partly 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 differences (estimated as persistent) between groups (e.g., college-educated and high school–educated households) and into an increase in idiosyncratic differences (estimated as less persistent) within each group (e.g., employed and unemployed).

In the second part of the paper, first we estimate stochastic processes for household earnings, consumption, and hours worked that are consistent with the evolution of the empirical cross-sectional distributions and with

2. Obviously 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 approach implies that we will ignore all the effects that changes in inequality might have on these trends.

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: (1) persistently higher differences between groups reduce relative mean resources of some groups and increase them for others, and (2) higher volatility within one 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 the second aspect significantly affects the welfare of all groups, other things being equal, and that the first aspect 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 and sex groups of the U.S. 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, earning, 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. Additional details about the data used in the main body of the paper are discussed in the data appendix (Section 6).

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 sizable literature, summarized in Blundell and MaCurdy (1999), investigates the behavioral response of labor supply to changes in wages and the employment status of the primary earner. We explicitly study hours worked by several members of the household, so the literature on the added worker effect, which studies the labor supply response of a spouse to the primary worker's job loss or job displacement, is relevant to our work. Whereas most studies find small effects (see, for example, 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 sizable.

Second, if changes in the wage structure translate into changes in a household's income process, how do these changes affect 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.[3] If preferences of all individuals are identical and can be represented by a constant relative risk aversion utility function that is separable across time and between consumption and leisure, then individual consumption growth rates move in tandem with aggregate consumption growth rates, unless tastes shift. Conditional on aggregate consumption growth rates, individual consumption growth rates would be uncorrelated with individual income growth rates, whether or not 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, for example, 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 a point of departure, which envisions a consumer in isolation attempting to self-insure against income fluctuations. The welfare consequences of increased income instability in this paradigm depend crucially on whether the income shocks are temporary or permanent because transitory shocks can be easily smoothed using (dis)saving. Hall and Mishkin (1982) decompose income into its transitory and permanent components and investigate whether, in fact, households smooth transitory income shocks to a higher degree than they smooth permanent income shocks. Blundell and Preston (1998) 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 wel-

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

fare 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 strands of the literature is that researchers interpret consumption data through the lens of a particular economic model of financial markets (a complete set of contingent claims for the former strand, a single uncontingent bond in the latter). Finally, a descriptive literature 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, 1991b), Johnson and Shipp (1991), Mayer and Jencks (1993), Johnson and Smeeding (1998), and Slesnick (1993, 2001).[4] In addition, the papers by Attanasio and Davis (1996), Dynarski and Gruber (1997), and us (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, and consumption inequality using microconsumption data is 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 1 to 3% 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 once 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

4. Even 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.

are two orders of magnitude bigger than the costs of business cycles commonly derived in the macroeconomic literature (see, for example, Lucas, 1987).[5]

3. *Descriptive Statistics*

Before assessing the welfare consequences of increasing wage and income inequality, we want first 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 nonsleep time minus hours worked per member of the household). A comment about our choice of the CE as our data source may be appropriate at this time. For detailed information on U.S. *individual* household consumption, the CE is the only available dataset. In addition, we want to investigate changes in the distribution of welfare associated with changes in the wage and earnings distribution, including the within-group variation. Thus, synthetic cohort techniques that make it feasible to combine several different datasets cannot be applied because these techniques average out all within-group (idiosyncratic) variation in the data.[6] The CE includes not only information about consumption, but also about hours worked and income, and indirectly about 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 dataset that reports household-level observations for all economic variables needed for our study.

To organize the CE data we use a simple period-by-period budget constraint that 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},

5. Krusell 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.

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

Finally let t_{it} denote direct labor income taxes paid by household i, and τ_{it} denote government transfers. Therefore:

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

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 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, to explore how wage inequality is related to inequality in hours worked and to consumption inequality, we have to explore (1) how wage inequality is related to hours inequality and disposable earnings inequality and (2) how disposable earnings inequality is related to consumption inequality (and thus inequality in savings).

3.1 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 30 years. We also find that the inequality in wages translates in a quantitatively substantial way into inequality of disposable earnings.

3.1.1 *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 person of the household.[7] The figure shows a

7. The reference person in the CE is defined as the person who owns or rents the home in which the household members reside.

Figure 1 STANDARD DEVIATION OF LOG WEEKLY WAGES

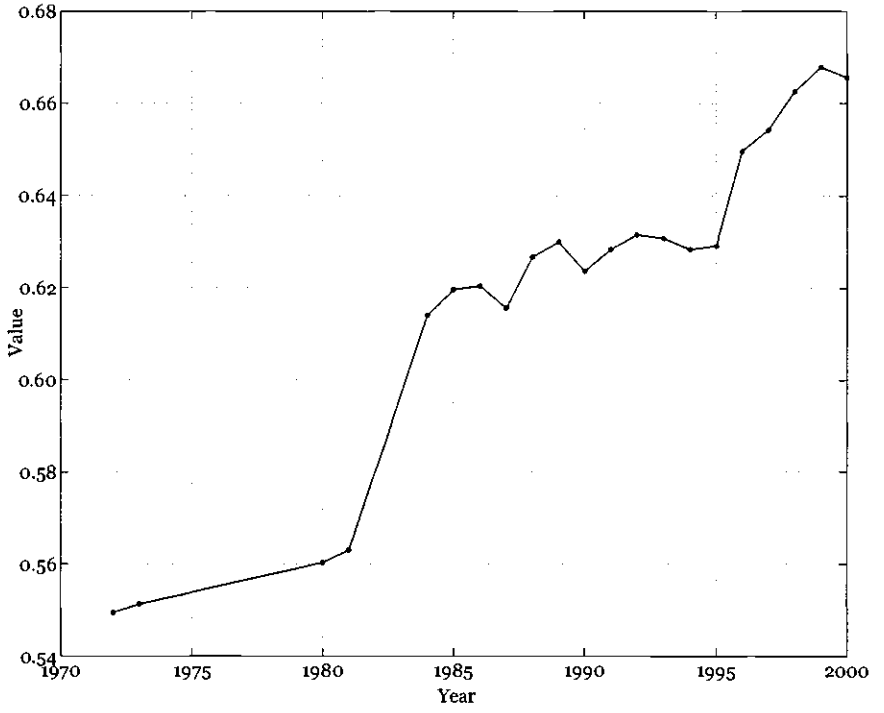

quantitatively significant increase in wage inequality: the percentage standard deviation of wages increases from around 55% in the early 1970s to around 67% at the end of the 1990s. We select only households that are classified as complete income respondents in the CE, that report positive consumption 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, to reduce measurement error, we exclude households whose reference person reports a weekly wage below $67, in constant 1982 dollars (equal to half of the 1982 minimum wage based on a 40-hour workweek). Finally, because 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.[8]

Note that our sample includes households with low wages that are particularly vulnerable to the increase in inequality (for example, households

8. The number of adult equivalents is computed using the census equivalence scale.

whose reference person is not a full-time worker and households with female reference persons).[9] On the other hand, households with a reference person that never works during the year are not included in our wage distribution sample. Juhn, Murphy, and Topel (2002) argue that the number of such households is large and increased throughout the 1990s; they use the Current Population Survey (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. We have consumption data for these households, however, and we will include them in the consumption distribution sample so that our consumption-based welfare analysis will take into account the effects of long jobless spells and labor force withdrawals described by Juhn et al. (2002).

We want to compare briefly our wage inequality findings with previous studies in the literature because some authors (for example, Cutler and Katz, 1991a) have questioned the reliability of CE wage and income data. 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 in Katz and Autor, 1999). Notice that both the timing and the 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 et al., 2003, Table 2).[10] The figure confirms that both the timing and the magnitude of the increase in wage inequality are similar across the three datasets.[11]

3.1.2 Earnings Inequality Wage inequality may be accentuated or mitigated by the endogenous labor supply decisions of the members of the

9. This sample selection strategy was suggested to us by our discussant Steve Davis.
10. All 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).
11. Katz and Autor report only 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 because levels are affected by the different top-coding thresholds in the datasets and by the potentially different extent of measurement error.

household. A decline in the primary earner's wage may induce other members of the household 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 substitute and work less today and work more in the future when wages have recovered. The map between 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, $\Sigma_j w_{ijt} h_{ijt}$, and in disposable household earnings, $\Sigma_j w_{ijt} h_{ijt} - t_{it} + \tau_{it}$ 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 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

Figure 2 WAGE INEQUALITY FROM CE AND CPS

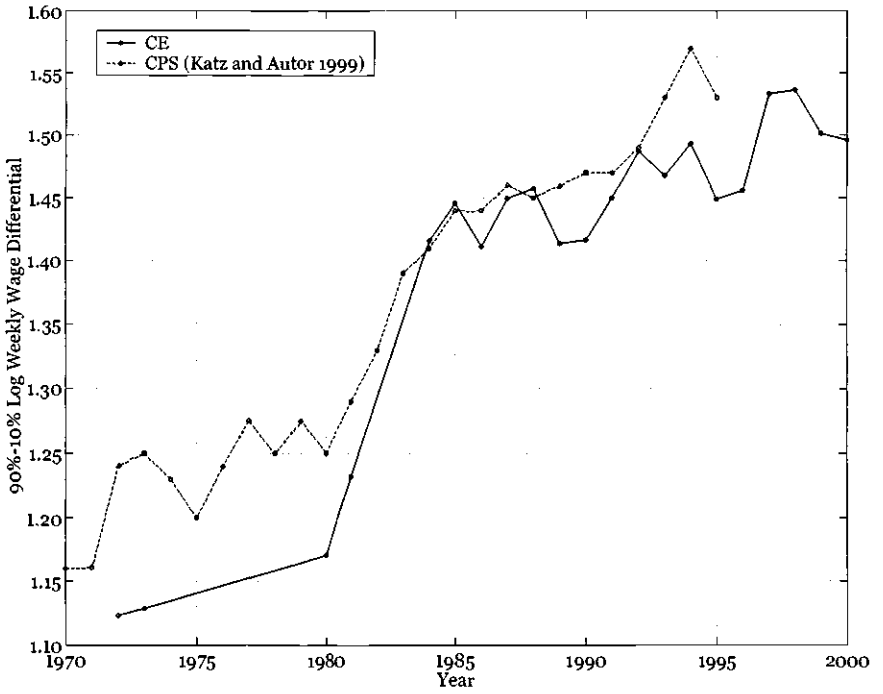

earnings inequality is significantly lower than inequality in total earnings. The gap between total earnings and disposable earnings inequality has declined over time. This decline suggests that the government tax and transfer system may have become less progressive over our sample period, which caused 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 and race composition of the population, first we 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 $y_t = \{\ln y_{it}\}_{i \in t}$. 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 the same age and the same race.

Figure 3 CHANGE IN WAGE INEQUALITY FROM CE, CPS, AND PSID

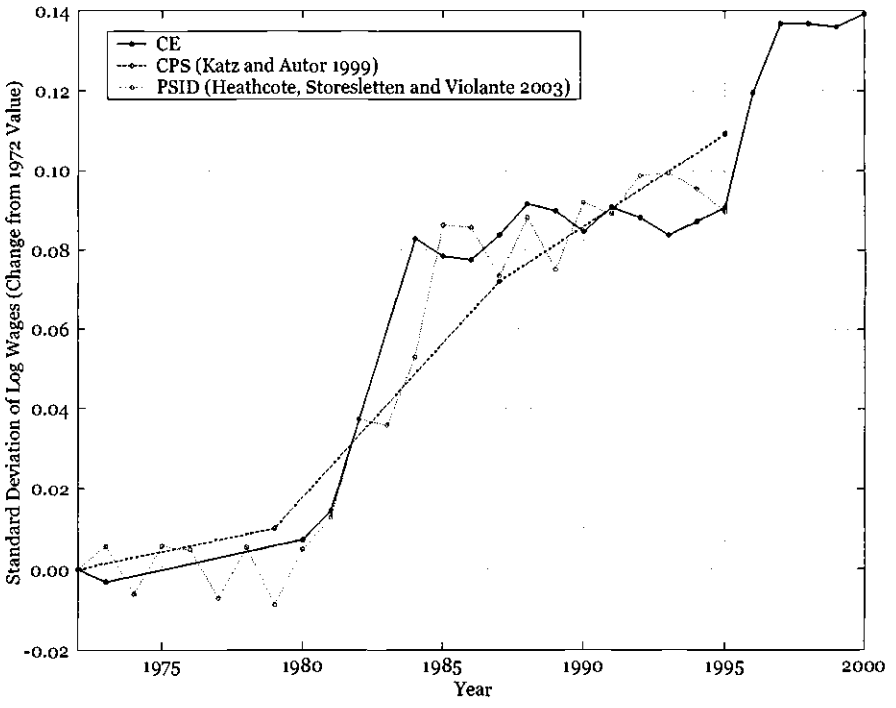

Figure 4 STANDARD DEVIATION OF LOG TOTAL AND DISPOSABLE EARNINGS

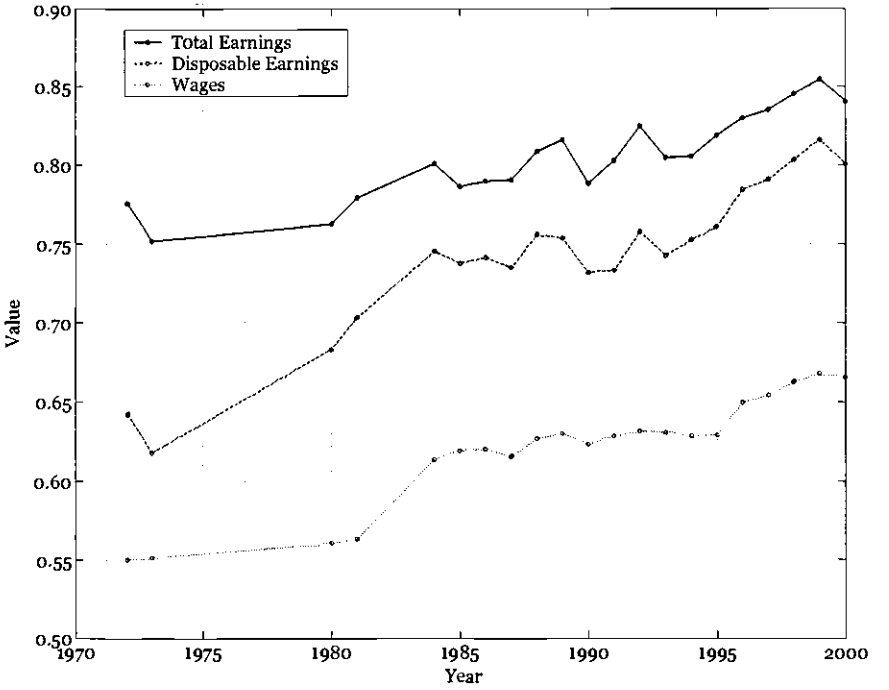

We then decompose each earning observation into a component y_{it}^{g} due to observables (education and sex) and into a part due to unobservable, purely idiosyncratic variation y_{it}^{d}:

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

so that $y_{it} = y_{it}^{g}\ y_{it}^{d}$.[12] We chose education and sex to define groups because 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

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

example, household i with $\ln y_{it} = -0.4$, $\ln(y_{it}^g) = -0.2$ and $\ln(y_{it}^s) = -0.2$; thus, this household has earnings 40% below the average earnings of a 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 $\{\ln(y_{it}^s)\}$, we compute cross-sectional idiosyncratic variances σ_{dt}^2. This procedure yields time series $\{\sigma_{gt}^2, \sigma_{dt}^2\}_{t \in T}$ of variances satisfying:

$$\sigma_{yt}^2 = \sigma_{gt}^2 + \sigma_{dt}^2 \tag{1}$$

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 increases throughout the sample, while between-group inequality

Figure 5 DECOMPOSITION OF DISPOSABLE EARNINGS INEQUALITY

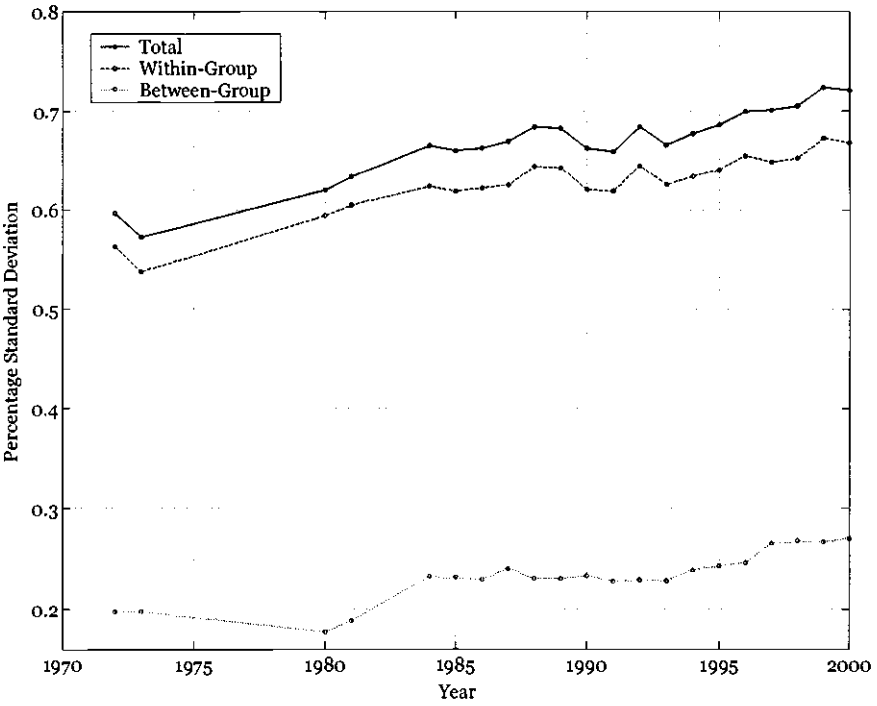

declines slightly in the 1970s (reflecting a reduction in the college premium) but increases significantly in the 1980s and 1990s. The patterns of both inequality measures from the CE are again similar to those from the CPS and the census data described by Katz and Autor (1999).

3.2 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.

3.2.1 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

Figure 6 STANDARD DEVIATION OF LOG DISPOSABLE EARNINGS AND LOG CONSUMPTION

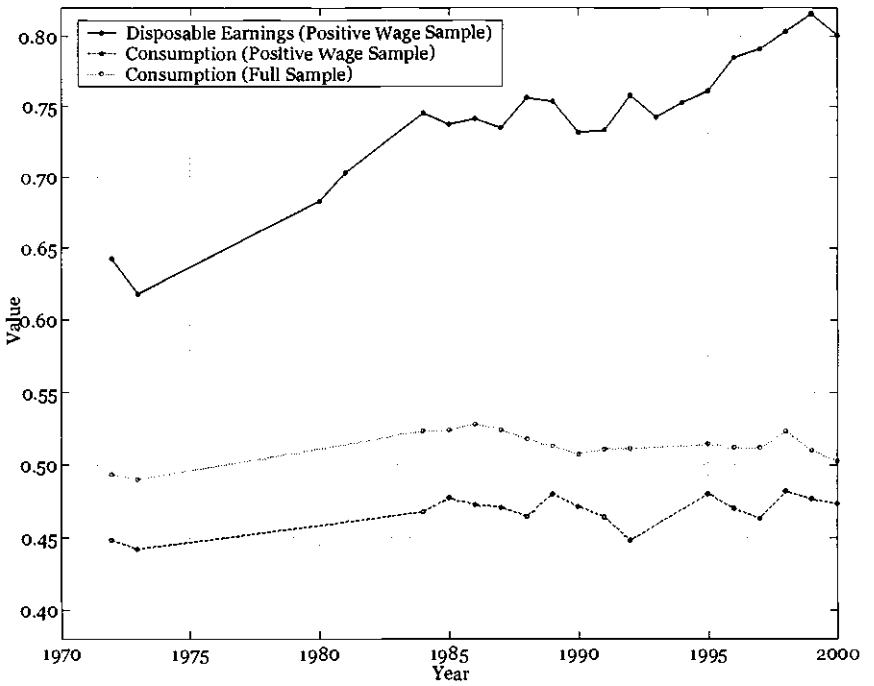

throughout the paper; see the data appendix and our earlier work [Krueger and Perri, 2002] for details about the construction of the consumption data.) The first measure, labeled consumption (Positive Wage Sample), is computed using the same sample selection criterion we use for wage and earnings inequality and does not include households with reference persons that do not work. The second measure, labeled consumption (Full Sample), includes all households that are complete income respondents, have a reference person between the ages 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

Figure 7 CONSUMPTION BY SELECTED QUINTILES OF THE DISPOSABLE EARNINGS DISTRIBUTION

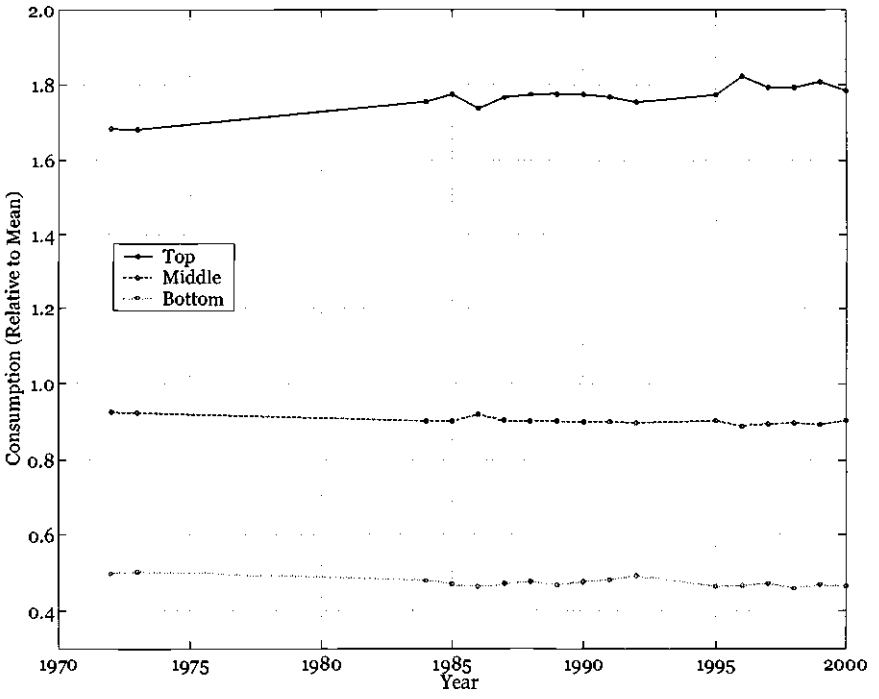

top quintile of the disposable earnings distribution (relative to average consumption in that year).[13]

Both figures confirm and extend our earlier findings (Krueger and Perri, 2002). During a period of strongly widening wage and earnings dispersion, there has been surprisingly little increase in consumption inequality.[14] 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.

Note that, even though total consumption inequality has increased by only a small amount, this phenomenon is a result of two offsetting trends. We decompose—as we did for earnings—consumption inequality into within- and between-group inequality.[15] 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 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. Because between-group inequality tends to be highly persistent, it is likely to have important welfare consequences.[16] Second, the fact that

13. Consumption 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 a reference person working at least one week, consumption by quintiles displays the same constant pattern.
14. We document that this fact is robust compared to various definition of consumption expenditures (Krueger and Perri, 2002). 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 biweekly survey of expenditure data for items purchased on a daily or weekly basis) showing increasing consumption inequality. Additional 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 match the National Income and Product Accounting total. Slesnick (2001) discusses some possible explanations for this phenomenon but concludes that a large part of the discrepancy between CE and NIPA is still unexplained.
15. We 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 nondurable consumption plus services from durables on a quartic polynomial in the age of the reference person and on a dummy for his or her race. We then regress the residuals on the years of education and gender of the reference person of the household. The sample used consists of all households that report positive consumption.
16. This 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.

Figure 8 DECOMPOSITION OF CONSUMPTION INEQUALITY

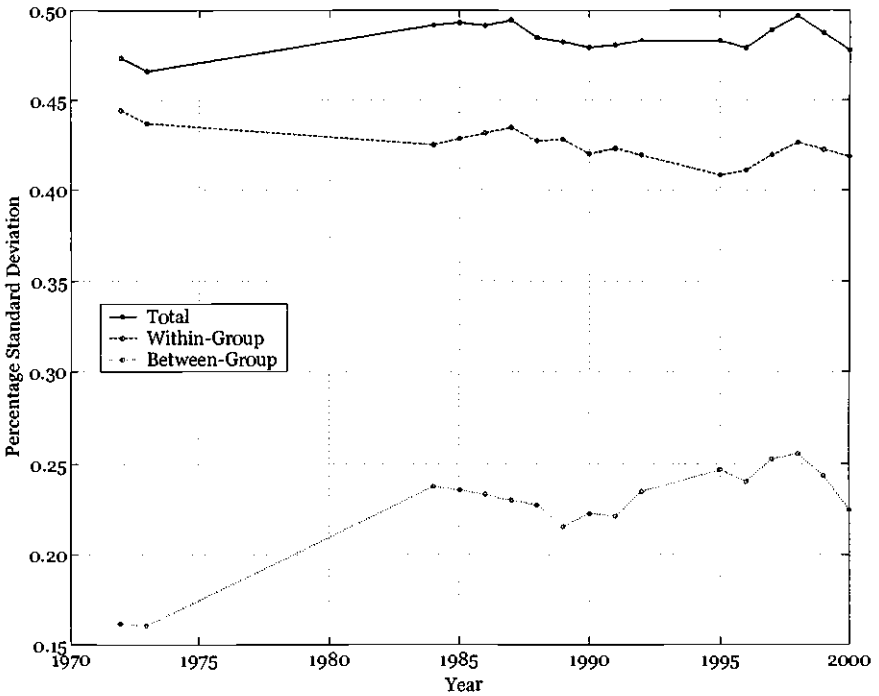

within-group disposable earnings inequality has increased while within-group consumption inequality has not indicates that the increase in within-group earnings variations has been effectively smoothed by households, possibly through some form of credit markets.[17]

3.2.2 Hours Inequality A household's lifetime utility depends on both the consumption and leisure it enjoys, so 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),

17. Blundell and Preston (1998), Heathcote et al. (2003), and we [Krueger and Perri, 2002] investigate the role of credit markets in generating this divergence between income and consumption inequality in detail.

Figure 9 STANDARD DEVIATION OF LOG WAGES AND OF LOG PER ADULT HOURS

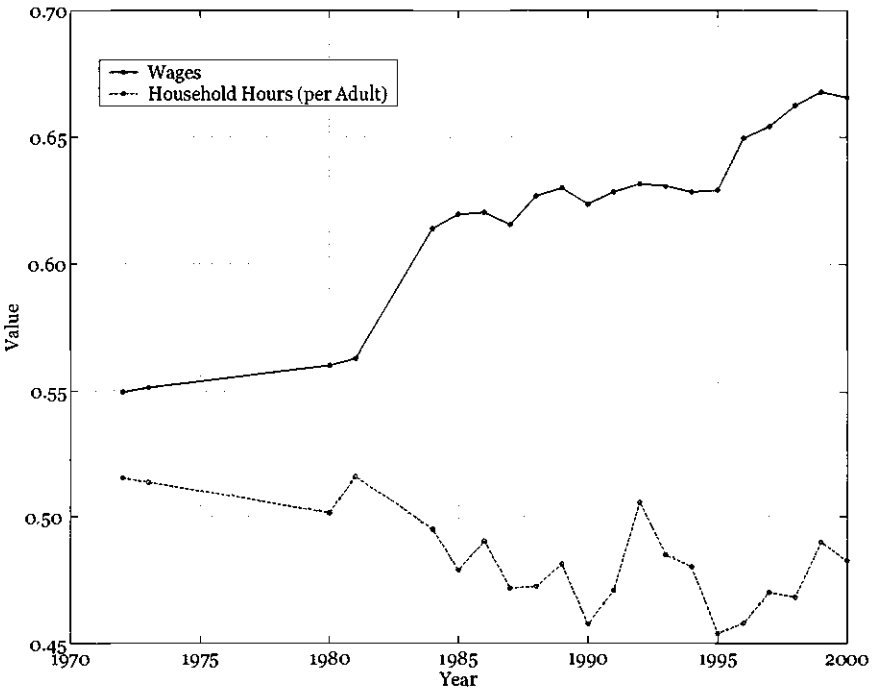

together with the standard deviation of wages of the reference person of the household.[18]

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 change in their trends, the reduction in hours dispersion would imply welfare gains. However, Figure 10 suggests that hours and consumption *are* correlated. 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.

Apart from the strong average increase in hours worked, whose causes and consequences we will not analyze here, the figure reveals two other facts.[19] First, households with high consumption work more hours;

18. We 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.

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

second, high-consumption households have increased their hours by less than low-consumption households. For example, in 1972–1973, households in the bottom quintile of the consumption distribution worked on average 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 figure 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 1970s, the high correlation between consumption and hours implied that households with low consumption were compensated with high leisure. In the 1990s, the reduction in this correlation implies that households with relatively low consumption do not enjoy as much leisure as

Figure 10 YEARLY HOURS WORKED PER PERSON (AVERAGE AND BY SELECTED QUINTILES OF THE CONSUMPTION DISTRIBUTION)

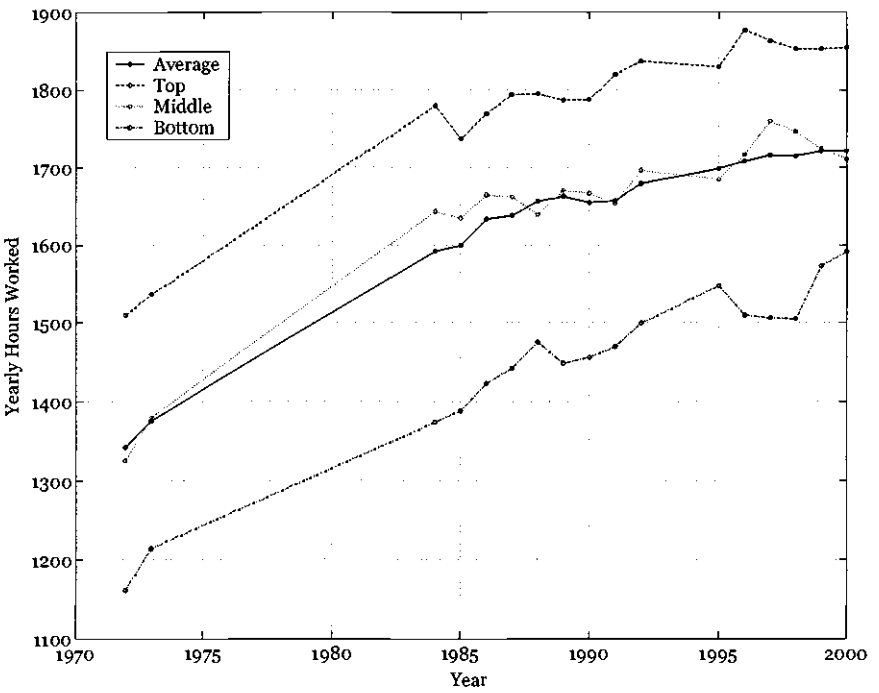

they used to and thus face potentially substantial welfare losses from this change over time. We will quantify the welfare consequences from a declining correlation in consumption and hours worked in the section containing our welfare calculation below.

3.3 SUMMARY

To summarize Section 3, 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 did almost as well (relatively to the mean) in the late 1990s as they did in the 1970s.

This finding 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

Figure 11 CORRELATION OF HOURS WORKED AND CONSUMPTION

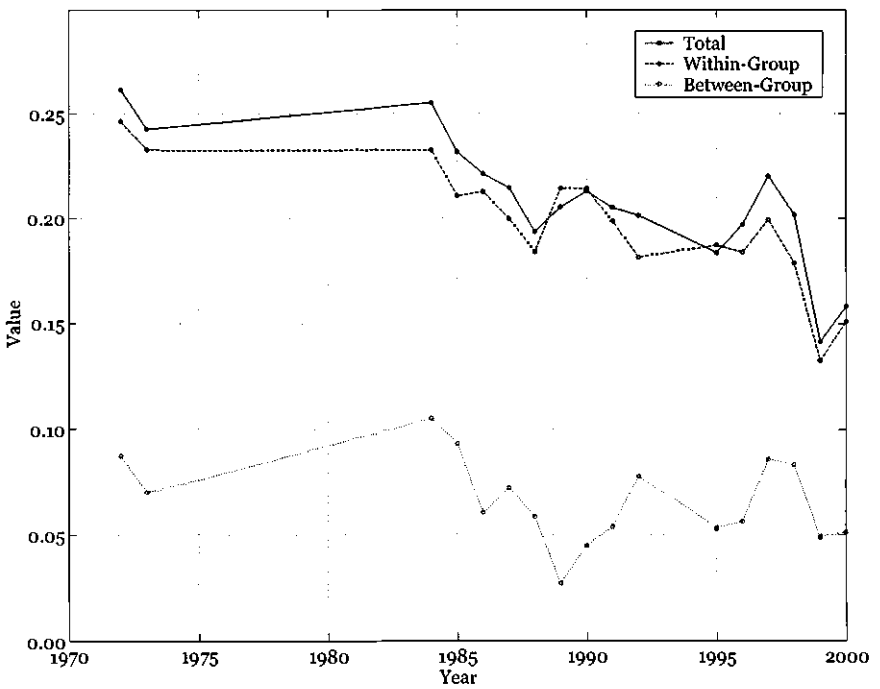

income falls (e.g., households with little education) experience substantial declines in relative consumption, with associated important welfare consequences.

In addition, the declining trend in the consumption–hours worked correlation suggests that one way through which low-consumption households 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? To answer these questions explicitly, 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 household's 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 and/or loss distribution from the change in the earnings (or wage) process.[20] 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 adopt the second approach in this paper.

We now 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.

4.1 MODELING HOUSEHOLD EARNINGS, CONSUMPTION, AND LEISURE PROCESSES

To carry our welfare analysis, we need stochastic processes for disposable earnings, consumption, and 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 $\ln(y_{it})$ can

20. Our earlier study (Krueger and Perri, 2002) and the study by Heathcote et al. (2003) take this approach.

be decomposed in a between-group $[\ln(y_{it}^s)]$ and a within-group part $[\ln(y_{it}^d)]$.

In this section, we specify time-series processes for both components of earnings. In particular, we assume that both $\ln(y_{it}^d)$ and $\ln(y_{it}^s)$ follow finite state Markov chains (of cardinality N) with possibly time-varying states Y_t^d and Y_t^s and with time-invariant transition matrices π^d, π^s.[21]

The states $\{Y_t^s\}_{t \in T}$ are determined by setting $Y_t^s(i)$ equal to the midpoint of the i-th quantile of the $\{\ln(y_{it}^s)\}$ sample for that period. By construction, our stochastic process matches, for every t, $N(N-1)/2$ quantile ratios of $\{\ln(y_{it}^s)\}$, its mean (by construction, equal to 0), but not necessarily its variance.[22] 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 address 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 $\{\ln(y_{it}^s)\}_{i,t}$ we group households into N relative classes delimited by $N-1$ uniformly spaced, time dependent quantiles: the first class is comprised of the bottom $1/N\%$ of the distribution in that quarter for that variable, the second class is comprised of the following $1/N\%$, etc. With this construction, the lowest class for $\{\ln y_{it}^s\}$ is interpreted as the group of households for which earnings explained by observables (gender 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.[23] After repeating this proce-

21. In 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.
22. We 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 match exactly 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).
23. More precisely, a period is roughly three quarters because we use data from the second and fifth interview of a household in the CE 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.

Table 1 THE PERSISTENCE OF THE ESTIMATED PROCESSES[1]

	Within-group	*Between-group*
Disposable earnings	0.76	0.97
Consumption	0.72	0.98

1. The reported measure of persistence is the value of the second largest eigenvalue of the transition matrices π^d and π^s.

dure for every period in the sample, the probability π_{ij}^s 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.[24] 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 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. Because the correlation between hours and between-group consumption is fairly constant and close to 0 (see Figure 11), 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 so that the correlation between the two variables in the model matches exactly

24. Transitions 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 the gender or the education of their reference person (through death, marriage, or divorce).

the falling cross-sectional correlation between the two variables in the data, as reported in Figure 11.[25]

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

4.2 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 nonsleep time) according to 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:

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

where α, β and σ are preference parameters that govern the relative importance of consumption relative to leisure, time discounting, and the intertemporal elasticity of substitution, respectively.[26] 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 write:

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

Note that, conditional on knowing the function V, the conditional expectation can be evaluated because the stochastic processes for consumption and leisure have a Markov structure.[27] Also note that, for an arbitrary

25. The complete results of our estimation procedure are available on request from the authors.
26. Note 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.
27. This 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 fewer periods). The proportionality factor is less than one and decreases with the remaining lifetime of a household.

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. Also 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 same).[28] 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 dataset) of the ensuing increase in wage inequality in the future is thus given by

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

where $\|\cdot\|$ is a particular metric. We let $W(c, l)$ 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 $W(c, l)$, that is, the distribution of the welfare consequences of the increase in wage inequality.

4.3 RESULTS

4.3.1 An Upper Bound Suppose that households do not have access to any savings technology and do not value leisure.[29] Then consumption equals labor earnings y and, if we ignore irrelevant constants and suppress time indexation, equation (2) becomes:

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

or more explicitly:

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

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

29. The 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:

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

The term $W(y^g, y^d)$ 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.[30] To compute this number, three steps have to be carried out:

1. Solve for the functions 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 $W(y^g, y^d)$.

Evidently steps 2 and 3 can be carried out easily 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 to be unable to smooth consumption via savings responses.

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

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

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

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

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:

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

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

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

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

$$V(y) = \ln(y) + \frac{\beta}{1-\beta}\mu$$

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

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

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

For this particular example, the welfare losses from increased earnings inequality are (1) independent of an agents' current earnings and (2) 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 cannot be solved by guess-and-verify methods and has to be computed numerically.[31] 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.

In Figure 12a, we plot the welfare losses implied by our estimated income processes from the CE. Each individual is characterized in 1972 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 $1/81$. 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 percentage 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% or more of consumption 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)

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

Figure 12 WELFARE LOSSES

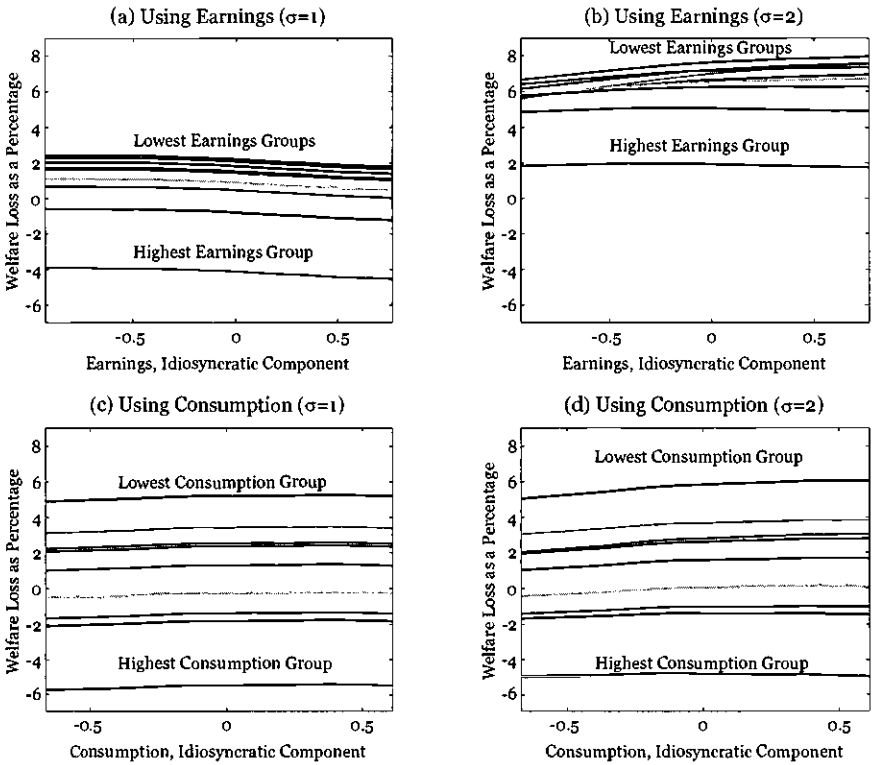

benefit from the increase in earnings inequality.[32] Welfare losses (or gains) are fairly uniform in the idiosyncratic income state, y^d, that a household starts with in 1972.

These findings can be interpreted as follows. The *aggregate* statistics show an increase in inequality in the last 30 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 remaining in the same earnings class, high-earnings class households will enjoy increasing relative earnings, and low-earning classes will experience declining relative earnings. These outcomes are true for both the group component as well as the idiosyncratic component of earnings. From the figure, we observe, however, that the welfare losses differ substantially by

32. Note that, when we refer to welfare losses from inequality from now on, 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.

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 of 1 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 because 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 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 because 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 its 1972 earnings state increases with time. With sizable time discounting of 4% per year, the 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, as seen in Figure 12. 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 depend on neither a household's initial earnings position nor the transition path.[33] 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 a 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 $\sigma = 2$.

4.3.2 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 smooth consumption intertemporally 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 policymakers.

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:

$$V(c) = u(c) + \beta E\left[V(c')|(c)\right] \tag{3}$$

which is solved as easily as in the previous subsection, under the maintained assumption that the estimated stochastic consumption processes have a 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: the consumption group to which a household belongs largely determines whether it ends up a loser or a winner of the increase in inequality. Again, the welfare consequences are fairly uniform across idiosyncratic consumption states. Quantitatively, however, some crucial differences in the results based on earnings observations emerge. The highest welfare losses now amount to slightly more than 5% of consumption, but only for

33. One may interpret this latter thought experiment and the welfare numbers as a 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 in new, post-2000 steady state.

the consumption-poorest 10% of the population.[34] 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 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.[35]

For earnings, an increase in the risk aversion of households led to dramatically increased welfare losses from increased inequality, but when using consumption data, the welfare consequences of increased inequality are close to invariant to changes in a household's attitudes toward risk. Again, this finding is explained by the stable trend of the variance of the idiosyncratic consumption component. The welfare losses and gains in Figure 12d, computed for risk aversion of $\sigma = 2$, are almost indistinguishable from the corresponding figure 12c 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) because 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, to focus on the welfare effects of increased inequality, we have effectively detrended our data from

34. These high losses of the lowest consumption group compared to the more moderate losses of the lowest earnings group documented above are largely due to the timing of the transition, coupled with sizable 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.

35. Since 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.

aggregate consumption growth. Thus, our results are not to be interpreted as absolute standards of living having declined over the last 30 years for a large fraction of the population.

4.3.3 The Impact of Changes in Leisure In our empirical section we documented that, in 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:

$$V(c,l) = u(c,l) + \beta E[V(c',l')/(c,l)] \tag{4}$$

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 a sensitivity analysis).

Both the definition of welfare costs as well as their computation remain the same as in the previous subsection. We parametrize the period utility function as before, and choose a share parameter $\alpha = \frac{1}{3}$ and a nonsleep time endowment of 15 hours per person, per day.[36]

4.3.3.1 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 leisure state, or the risk aversion of the household play only a minor quantitative role.[37]

36. In a static deterministic model with our preferences, an agent would choose to work exactly one-third of her nonsleep 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.

37. The reason that 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 idosyncratic consumption).

Because 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 our estimation of stochastic processes for hours worked, as with consumption, 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 U.S. households.

4.3.3.2 LABOR SUPPLY CORRELATED WITH CONSUMPTION In our empirical analysis, we found that the correlation between idiosyncratic consumption and leisure is negative and, more important, has declined in absolute value between 1972 and 2000. In other words, while low-consumption housseholds used to enjoy at least significantly higher leisure in 1972, this situation 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 when uncorrelated (see the discussion in the previous subsection), and consumption and leisure processes that are correlated. The subpanels distinguish risk aversion and agents belonging to different 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.

For separable preferences ($CRRA = 1$), correlation between consumption and leisure does not change our findings from the previous section. As households become more averse to risk, 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 modeled 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 notrivial extent.

5. *Conclusion*

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 sizable welfare gains of similar magnitude as the losses of the poorest agents. Overall, a majority of the population (based on consumption observations, roughly 60% 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 the consumption/leisure correlation contributes to moderately higher

Figure 13 WELFARE LOSSES

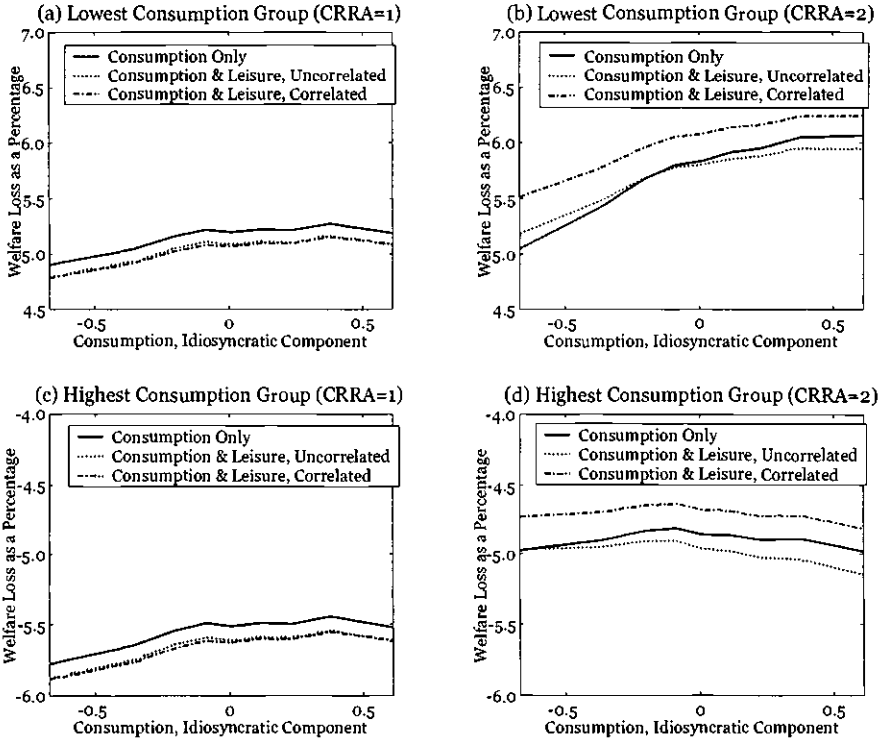

welfare losses, an order of magnitude of 0.5% for most households. Second, while welfare losses from increased inequality are fairly sizable for a lot of U.S. households when based on consumption data, they are not nearly as big and affect not nearly as many households as an analysis based on earnings 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.

To focus more precisely 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; 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 understand (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 interview surveys of the Consumer Expenditure (CE) survey for the years 1972–1973, 1980–1981, and 1984–2000, 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 onward 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 1972–1973 until 2000.

6.1 WAGES, EARNINGS, AND HOURS

The definition of wages we use is the 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 nonfarm 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. Because the 1972–1973 CE does not report the number of hours worked per week by household members, 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).

6.2 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 consumer price index (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 that we use expenditures and not imputed services is that no information is available in the CE 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 (1991b). From the CE data, we have expenditures for purchases of new and used vehicles. We also have data on the number of cars that a consumer unit possesses. For each year we first select all households that report positive expenditures for vehicle purchases and then 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 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) that the residence would command if rented out.[38] This variable is not available for the years 1980–1981 and 1993–1994; for these years we do not compute inequality measures for nondurables ND+ consumption expenditures.[39] As with nondurable consumption, all imputed services from consumer durables and housing are deflated with the corresponding CPI.

REFERENCES

Altug, S., and R. Miller. (1990). Household choices in equilibrium. *Econometrica* 58:543–570.

Attanasio, O. (2002). Consumption and income inequality: What do we know and what we can learn from it. UCL. Mimeo. University College London.

Attanasio, O., and S. Davis. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy* 104:1227–1262.

Blundell, R., and T. MaCurdy. (1999). Labor supply: A review of alternative approaches. In *Handbook of Labor Economics*, O. Ashenfelter and D. Card (eds.). Amsterdam: Elsevier.

Blundell, R., and I. Preston. (1998). Consumption inequality and income uncertainty. *Quarterly Journal of Economics* 113:603–640.

Cochrane, J. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99:957–976.

Cox, M., and R. Alm. (1999). *Myths of Rich & Poor: Why We're Better Off Than We Think*. New York: Basic Books.

Cullen, J., and J. Gruber. (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics* 18:546–572.

38. Here is the exact question that the reference person of the CU is asked: "If you were to rent your home today, how much do you think it would rent for monthly, without furnishings and utilities?"

39. We experimented with using an imputation procedure similar to the one used for vehicles to obtain housing services for the four missing years. Results were very similar and are available from the authors on request.

Cutler, D., and L. Katz. (1991a). Macroeconomic performance and the disadvantaged. *Brookings Papers on Economic Activity* 1–74.

Cutler, D., and L. Katz. (1991b). Rising inequality? Changes in the distribution of income and consumption in the 1980s. *American Economic Review* 82:546–551.

Deaton, A. (1992). *Understanding Consumption*. Oxford, England: Clarendon Press.

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.

Dynarski, S., and J. Gruber. (1997). Can families smooth variable earnings? *Brookings Papers on Economic Activity* 229–284.

Gottschalk, P., and R. Moffitt. (1994). The growth of earnings instability in the U.S. labor market. *Brookings Papers on Economic Activity* 217–272.

Greenspan, A. (1998). Opening remarks to the symposium on income inequality: Issues and policy options. Federal Reserve Bank of Kansas City.

Hall, R., and F. Mishkin. (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica* 50:461–481.

Hayashi, F., J. Altonji, and L. Kotlikoff. (1996). Risk sharing between and within families. *Econometrica* 64:261–294.

Heathcote, J., K. Storesletten, and G. Violante. (2003). The welfare implications of rising wage inequality in the US. Stockholm University. Mimeo. Institute for International Economic Studies IIES.

Heckman, J., and T. MaCurdy. (1980). A life cycle model of female labor supply. *Review of Economic Studies* 47:47–74.

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.

Johnson, D., and T. Smeeding. (1998). Measuring the trends in inequality of individuals and families: Income and consumption. U.S. Bureau of Lasor Statistics. Working paper. Washington, D.C.:

Juhn, C., K. Murphy, and R. Topel. (2002). Current unemployment, historically contemplated. *Brooking Papers on Economic Activity* 1:79–136.

Katz, L., and D. Autor. (1999). Changes in wages structure and earnings inequality. In *Handbook of Labor Economics*. O. Ashenfelter and D. Card (eds.). Amsterdam: Elsevier.

Krueger, D., and F. Perri. (2002). Does income inequality lead to consumption inequality: Evidence and theory. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 9202.

Krusell, P., and A. Smith. (1999). On the welfare effects of eliminating business cycles. *Review of Economic Dynamics* 2:245–272.

Lucas, R. (1987). *Models of Business Cycles*. New York: Basil Blackwell.

Mace, B. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99:928–956.

Mayer, S., and C. Jencks. (1993). Recent trends in economic inequality in the United States: income versus expenditures versus well-being. In *Poverty and Prosperity in the USA in the Late Twentieth Century*, D. Papadimitriou and E. Wolff (eds.). New York: St. Martin's Press.

Nelson, J. (1994). On testing for full insurance using Consumer Expenditure survey data. *Journal of Political Economy* 102:384–394.

Slesnick, D. (1993). Gaining ground: Poverty in the postwar United States. *Journal of Political Economy* 101:1–38.

Slesnick, D. (2001). *Consumption and Social Welfare.* Cambridge, England: Cambridge University Press.
Stephens, M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics* 20:504–537.
Tauchen, G., and R. Hussey. (1991). Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica* 59:371–396.
Townsend, R. (1994). Risk and insurance in village india. *Econometrica* 62:539–591.

Comment

STEVEN J. DAVIS
University of Chicago and NBER

Krueger and Perri set their sights on a major question: What are the welfare consequences of the pronounced rise in U.S. wage and income inequality in recent decades? Assumptions about the answer play an important role in many policy discussions and underlie much of the interest in wage and income inequality. I applaud the authors for tackling the question in a way that makes assumptions explicit and that facilitates constructive criticism. The authors also deserve much credit for grappling with the data on consumption inequality, a challenging task. As it turns out, however, I think that a compelling answer to the question awaits further research.

The authors describe trend changes in the distribution of household consumption expenditures using data from the Interview Survey component of the Consumer Expenditure (CE) survey. They find rising consumption inequality between groups defined by sex and educational attainment of the household head, but declining consumption inequality within groups. Overall consumption inequality changes little during the past three decades by their account—in striking contrast to a sharp rise in the inequality of wages, earnings, and disposable incomes.

They proceed to calculate welfare effects associated with certain changes in the consumption distribution. To do so, they specify and estimate stochastic processes for group-level and idiosyncratic components of consumption, and they postulate standard preferences over consumption paths.[1] They then compute, as of 1972, consumption-equivalent

Orazio Attanasio, David Autor, Erik Hurst, and Daniel Slesnick kindly supplied data for this comment and my remarks at the Macro-economics Annual Conference. Daniel Slesnick provided several useful observations about the Consumer Expenditure survey, and Erik Hurst provided valuable information about the Continuing Survey of Food Intakes by Individuals. Thanks also to Dirk Krueger and Fabrizio Perri for many helpful communications about their work.

1. Krueger and Perri also characterize the evolution of the hours-worked distribution, and some of their welfare calculations consider preferences defined over consumption *and* leisure. Their chief results involve consumption, however, and I limit my remarks to consumption-related issues.

welfare differences between estimated and counterfactual processes for consumption. They report these welfare differences as a function of education and initial position in the consumption distribution.

My remarks below develop two main themes. First, there are good reasons to doubt the basic characterization of consumption inequality trends offered by Krueger and Perri. Second, in their welfare analysis, they do not adequately model the uncertainty faced by households, which limits the usefulness of their welfare results.

1. Evaluating the Evidence on Trends in Consumption Inequality

Figure 6 in Krueger and Perri shows a striking divergence between the path of disposable earnings inequality and the path of consumption inequality. Figure 8 shows declining within-group consumption inequality in recent decades and, in particular, since 1987.[2] This decline is at odds with the strong rise in the within-group inequality of disposable earnings, as depicted in Figure 5.

To evaluate these empirical results, I proceed as follows. First, I consider whether they fit comfortably with other evidence on consumption responses to income shocks and conclude that they do not. Second, I note concerns about the quality of the Interview Survey component of the CE and the possibility that it yields an inaccurate picture of consumption inequality trends.[3] Third, I look to other sources of data on consumption inequality trends and find, in contrast to the CE Interview Survey, that they point to rising consumption inequality.

2. As a preliminary step, Krueger and Perri run cross-sectional regressions of consumption per adult-equivalent on a quartic polynomial in the household head's age and an indicator variable for his or her race. The consumption residuals from these regressions then serve as inputs to the remaining analysis. By allowing the coefficients on the age and race variables to vary freely by year in these regressions, Krueger and Perri sweep away potentially important changes in the consumption distribution. By constraining the age polynomial to have the same shape across education groups in each cross section, Krueger and Perri may obtain misleading characterizations of within-group inequality trends. For example, if the household distribution becomes more concentrated in education categories with similar age profiles, then the residual distribution tends to become less dispersed, even if there is no change in the distribution of consumption conditional on age and education. I ignore these issues in the main text.
3. Estimates of inequality in wages, earnings, and income are also subject to measurement problems, but the basic facts about U.S. trends in this regard are well established. Unlike the situation for consumption data, there are multiple, independent sources of data on wages, earnings, and income, and the most important sources have been heavily researched. Hence, I proceed under the assumption that the Krueger and Perri measures of trends in wage and earnings inequality are accurate descriptions of reality.

1.1 SELF-INSURANCE AND THE RELATIONSHIP OF CONSUMPTION INEQUALITY TO EARNINGS INEQUALITY

Consider the divergent paths of within-group earnings inequality (Figure 5) and within-group consumption inequality (Figure 8). The rise in within-group earnings inequality reflects some combination of greater dispersion in (the idiosyncratic components of) fixed earnings differences, persistent earnings shocks, and transitory shocks. Fixed effects and persistent shocks cannot be smoothed over the life cycle by borrowing and lending. Hence, a perspective on the data informed by permanent income theory leads one to anticipate a close relationship between persistent earnings shocks and consumption responses. Many other theories that entail incomplete sharing of consumption risks carry the same implication.[4]

This implication finds support in the observed relationship between group-level earnings shocks, which are highly persistent, and group-level consumption responses. For example, Figure 5 and 8 show similar trend increases in between-group earnings inequality and between-group consumption inequality. Attanasio and Davis (1996) find that persistent changes in relative wages among groups defined by birth cohort and education lead to roughly equal-size changes in consumption expenditures.

Many studies show that household consumption expenditures are sensitive to idiosyncratic earnings shocks. Two studies are especially pertinent here. First, Gruber (1997) investigates how consumption responses to unemployment vary with the generosity of unemployment insurance benefits. To estimate this relationship, he exploits the fact that the income replacement rates provided by unemployment benefits vary considerably across states and workers. Gruber's study is noteworthy for our purposes because the U.S. unemployment insurance system is not designed to insure against persistent earnings shocks—benefits typically expire after 26 weeks. So, insofar as the consumption response to unemployment varies with the replacement rate, households are not smoothing transitory income variation. Using Panel Study of Income Dynamics (PSID) data from 1968 to 1987, Gruber estimates that a 10-percentage-point rise in the replacement rate reduces the unemployment-induced fall in food expenditures by about 3 percentage points. This is a big effect, and it implies a big departure from effective self-insurance against transitory shocks. Gruber also notes that there has been a secular decline in the generosity of unemployment benefits

4. In principle, a properly structured portfolio of risky financial assets can insure against even the most persistent earnings shocks, but I am unaware of any evidence that households or their agents (e.g., pension fund managers) engage in this type of hedging behavior to a significant extent. Davis and Willen (2000) develop a theory of life-cycle portfolio choice with decision rules that exhibit this type of hedging behavior, and they present evidence that broadbased equity and bond funds have some limited potential as instruments for hedging occupation-level income shocks.

levels, which is a force for greater consumption inequality and increased sensitivity of consumption inequality to earnings inequality.

Second, Sullivan (2002) investigates whether households use unsecured debt to smooth consumption responses to unemployment spells. He relies on panel data from the PSID for 1984 to 1993 and the Survey of Program Participation for 1996 to 2000. Using a sample of unemployment spells that aims to isolate transitory earnings shocks, he finds that households with assets increase unsecured debt, on average, by about 10% of the earnings loss associated with unemployment. For households with low initial asset levels, however, unsecured debt does not respond to the income loss associated with unemployment. These households account for about 13 to 18% of the sample, depending on the definition of low assets. In addition, Sullivan finds that the consumption response to unemployment-induced income shocks is larger for households with lower asset levels. He estimates that expenditure on food and housing for households with little or no financial assets is five times more sensitive to unemployment-induced earnings losses compared to other households.

The Gruber and Sullivan studies indicate that many households do not effectively smooth transitory, idiosyncratic earnings shocks. Moreover, there are strong theoretical reasons, supported by empirical evidence, to think that households cannot smooth persistent earnings shocks. These considerations provide grounds for skepticism toward the Krueger and Perri evidence on consumption inequality trends, especially the decline in within-group consumption inequality coupled with strong increases in the within-group inequality of disposable earnings.

Conceivably, the effect of increased earnings inequality on within-group consumption inequality is overwhelmed by greater smoothing of transitory shocks. This interpretation is logically consistent, and it fits with the increasing availability of consumer credit, but the interpretation faces at least three problems. First, many households lack the financial means to smooth earnings shocks. Poorer households, in particular, often have little financial wealth, so they cannot draw on liquid assets to offset negative earnings shocks. Second, most forms of consumer credit carry high interest rates. Edelberg (2003, Table 1) reports mean consumer interest rates in 1998 of 8.0% per annum for first mortgages, 10.4% for second mortgages, 10.2% on auto loans, 14.5% for credit cards, and 12.9% for other consumer credit. Davis et al. (2003, Table 1) calculate that interest rates on unsecured forms of consumer credit exceed the three-year Treasury rate by 6 to 9 percentage points after netting out uncollected loan obligations. The high cost of consumer credit makes borrowing less useful for consumption smoothing, even when credit is available. The upshot of low financial wealth, incomplete access to credit markets, and high borrowing costs is that many households are poorly equipped to smooth even transitory earnings shocks.

Third, and perhaps most important, the effect of greater access to credit markets or other self-insurance devices must be large to rationalize Krueger and Perri's finding of a sharp divergence between within-group earnings inequality and within-group consumption inequality. Suppose that permanent shocks account for one-third of the rise in within-group earnings inequality. On that account alone then, Figure 5 suggests a secular rise of 3 or 4 log points in the standard deviation of consumption within groups. Instead, Figure 8 shows a decline of about 2 log points over the sample period. The implied gap between trend changes in earnings and consumption inequality is larger when we factor in a rising variance of transitory earnings shocks. This gap will be hard to explain in a model that matches the degree of consumption smoothing seen in the data.

1.2 CONCERNS ABOUT THE CE INTERVIEW SURVEY

The CE has two independent components, a quarterly Interview Survey and a weekly Diary Survey, each with its own questionnaire and sample. The two components differ, but overlap somewhat, in their coverage of expenditure categories. The Interview Survey covers a broad range of expenditure categories, but it is "designed to obtain data on the types of expenditures respondents can recall for a period of 3 months or longer" (www.bls.gov/cex/csxfaqs.htm). The Diary Survey focuses on frequently purchased smaller items such as food and beverages, housekeeping supplies, tobacco, nonprescription drugs, and personal-care products and services. Sample size in the Diary Survey is roughly one-third that of the Interview Survey. Following most previous research that uses CE micro data, Krueger and Perri rely on the Interview Survey.

The CE records out-of-pocket expenditures. Even when combined, the two CE components miss a big fraction of consumption (e.g., most health care). There are large and growing discrepancies between expenditures in the CE and Personal Consumption Expenditures (PCEs), as measured in the national income and product accounts. For example, Battistin (2003, Figure 1) reports a decline in the ratio of CE to PCE per-capita expenditures on nondurables and services from 0.79 in 1985 to 0.63 in 2000.[5] The CE-PCE gap and its growth over time are even larger when attention is restricted to the Interview Survey. This can be seen in Slesnick's (2001) Figure 3.2, which shows that the ratio of per-capita consumption in the CE Interview Survey to the PCE declines from 0.08 in 1973 to 0.56 in 1995.[6]

Slesnick (1992) investigates the discrepancy between CE and PCE consumption measures. He finds that only one-half of the CE-PCE gap

5. The CE-based measure of per-capita consumption in Battistin appears to reflect both the Interview and Diary surveys.
6. Slesnick confirms in a personal communication that his Figure 3.2 reflects CE data from the Interview Survey only.

reflects differences between the two sources in the definition of consumption, and the remaining half is unexplained. Underreporting in the CE (in covered expenditure categories) appears to be a major problem. The time period covered by Slesnick's study ends in 1989, after which the CE-PCE gap grew much larger. Battistin (2003) provides evidence that the quality of Interview Survey data on frequently purchased smaller items, housekeeping supplies, and personal-care products and services has declined over time, and that the decline has been "particularly accentuated" in the 1990s.

All of this leads me to question whether the CE provides a reliable basis for drawing inferences about trends in consumption inequality. Because the CE-PCE gap has expanded markedly over time, I am especially reluctant to accept Interview Survey evidence on consumption inequality trends during the 1990s. Of course, the large and growing discrepancy between CE and PCE consumption may partly reflect measurement problems in the national income accounts. It seems highly unlikely, however, that deterioration in the accuracy of the national income accounts can account for such a dramatic widening of the CE-PCE consumption gap.

1.3 CONSUMPTION INEQUALITY TRENDS IN OTHER DATA SOURCES

In light of my foregoing remarks, it seems appropriate to examine other data sources for evidence on consumption inequality trends. I do so, but my brief treatment merely scratches the surface of an important issue.

Attanasio (2002) compares the evolution of consumption inequality in the Interview and Diary Surveys. He plots the standard deviation of log consumption per household and per adult-equivalent for the overall population and for selected cohort-education groups. He finds that overall consumption inequality declines by about 2 log points from 1985 to 1998 in the Interview Survey, but it rises by 8 or 9 log points in the Diary Survey. His within-group plots also show rising inequality in the Diary Survey but flat or slightly declining inequality in the Interview Survey. Battistin (2003) reports similar results in his detailed analysis of the differences between the CE Diary and Interview Survey components. In short, the CE Diary Survey paints a picture of rising consumption inequality since 1985, in contrast to the flat or declining consumption inequality seen in Krueger and Perri's Figures 6 and 8.

Fisher and Johnson (2003) report Gini coefficients for consumption per adult-equivalent using data from the CE Interview Survey and the PSID. They impute total consumption for households in the PSID based on food expenditures, rent or mortgage, home ownership status and home value,

utility expenses, demographic and family composition variables, the age and sex of the household head, and other variables. Table 1 reproduces their statistics for overall consumption inequality. From 1984 to 1999, the rise in overall consumption inequality is nearly twice as large in the PSID as in the CE Interview Survey. Consumption inequality declines over the 1990s according to the Interview Survey, but it rises from 1990 to 1994 and for the 1990s as a whole according to the PSID.

I also examined data on monthly food expenditures in the Continuing Survey of Food Intakes by Individuals (CSFII), which is conducted by the U.S. Department of Agriculture. The CSFII is a nationally representative sample with repeated cross sections for six years: 1989 to 1991 and 1994 to 1996. The survey response rate in the CSFII exceeds 85%, which compares favorably to CE response rates. See Aguiar and Hurst (2003) for a detailed description of the CSFII and an interesting analysis that exploits separate CSFII measures of food expenditures and food consumption (e.g., caloric intake). I make use of data on food expenditures only.

Following Krueger and Perri, I restrict my analysis sample to households with a head between 20 and 64 years of age. To measure "overall" consumption inequality, I compute the 90-10 differential and the standard deviation of the residuals from yearly cross-sectional regressions of log expenditures on controls for household size and a quartic polynomial in the household head's age. I follow the same procedure to measure within-group consumption inequality, except that the regressions also include dummy variables for the household head's education. For between-group inequality, I report estimated coefficients on the education variables.

As shown in Table 2, the CSFII shows a broad pattern of rising consumption inequality from 1989 to the mid-1990s. Overall inequality in food expenditures rises by about 8% from 1989 to 1996, within-group inequality rises by a bit less, and the education differentials expand in most cases.

Table 1 GINI COEFFICIENTS FOR CONSUMPTION PER ADULT EQUIVALENT, 1984–1999 (CEX INTERVIEW SURVEY COMPARED TO THE PSID)

Data source	*1984*	*1990*	*1994*	*1999*	*Change, 1984–1999*
PSID	.255	.243	.286	.278	9.1%
CEX	.267	.295	.289	.280	4.7%

Source: Reproduced from Table 3 in Fisher and Johnson (2003).

Table 2 INEQUALITY IN LOG MONTHLY FOOD EXPENDITURES, 1989 TO 1996 (U.S. HOUSEHOLDS WITH A HEAD BETWEEN 20 AND 64 YEARS OF AGE)

A. Overall inequality (Controls for age of head and household size)

	1989	*1990*	*1991*	*1994*	*1995*	*1996*
Log standard deviation	.505	.503	.528	.528	.523	.543
90-10 log differential	1.25	1.27	1.33	1.31	1.30	1.36

B. Within-group inequality

	1989	*1990*	*1991*	*1994*	*1995*	*1996*
Log standard deviation	.489	.491	.505	.507	.509	.526
90-10 log differential	1.23	1.24	1.26	1.28	1.22	1.29

C. Between-group inequality	*Log deviation from households with a college-educated head*		
	Head with some college	*Head with high school diploma*	*Head did not complete high school*
1989–1991 pooled sample	–.155	–.236	–.386
1994–1996 pooled sample	–.135	–.243	–.404

Source: Author's calculations using data from the *Continuing Survey of Food Intakes by Individuals*, U.S. Department of Agriculture.

1. Food expenditures are the head's report of household expenditures in the previous month on food purchased at the grocery store; food delivered into the home; and food purchased at restaurants, bars, cafeterias, and fast-food establishments.
2. All statistics are calculated from the "main sample" in the CSFII, a nationally representative sample of noninstitutionalized persons residing in U.S. households. The CSFII is a repeated cross section for the indicated years.
3. The analysis sample contains households with a head between 20 and 64 years of age and nonmissing observations for food expenditures and years of completed schooling. I deleted four observations that reported food expenditures in the previous month (in 1996 dollars) of less than $10.00, two of which reported no expenditures. The resulting sample ranges from 1076 observations in 1990 to 1352 observations in 1996.
4. "Overall inequality" is computed from residuals in yearly cross-sectional regressions of log food expenditures on controls for household size and a quartic polynomial in the head's age. "Within-group inequality" is computed from residuals in a regression specification that also includes dummy variables for the four indicated education categories. The log deviations reported under "Between-group inequality" reflect the coefficients on the education variables for the same regression specification, but also including year effects.

1. SUMMING UP

Krueger and Perri's characterization of consumption inequality trends is difficult to reconcile with other evidence on consumption responses to income shocks. Consumption measures based on the CE Interview Survey show signs of deteriorating quality, and they cover a steadily declining share of Personal Consumption Expenditures in the national

income accounts. Three other sources of consumption data—the CE diary survey, the PSID, and the CSFII—show rising consumption inequality during the 1990s and are at odds with the message from the Interview Survey.[7] These observations cast doubt on Krueger and Perri's basic characterization of consumption inequality trends.

2. *Calculating Welfare Consequences*

2.1 SUMMARY OF THE KRUEGER AND PERRI PROCEDURE

It will be helpful to review the steps taken by Krueger and Perri in their welfare analysis. First, they obtain residuals from yearly cross-sectional regressions of log consumption on a constant and controls for the age and race of the household reference person. Second, they regress these residuals on the schooling and sex of the reference person, again by year. For each household-level observation, the second-stage regression expresses consumption as the sum of a predicted value and a residual value. Third, they sort predicted values into nine equal-size groups (each year). This sorting defines the sex-education "groups" to which households belong. Likewise, they sort residuals from the second-stage regressions into nine equal-size groups, which determines the household's relative position within its group.

Krueger and Perri then model the evolution of group-level and within-group components as independent nine-state Markov chains. They allow for time-varying states denoted by Y_t^g and Y_t^d and time-invariant transition matrices π^g and π^d, where g indexes the groups defined in the paragraph above and d indexes the within-group position. They set Y_t^g to the median of the predicted values in group g at t, and they set the values for Y_t^d in the same way. To estimate the elements of the nine-by-nine transition matrices, π^g and π^d, they use sample average transition rates from each state k to each state j.

The state vectors Y_t^g and Y_t^d, transition matrices π^g and π^d, and initial conditions for group membership and within-group position determine a stochastic path C for consumption. Given a utility function, the consumption path yields a welfare value $V(C)$. Altering one of the state vectors or transition matrices yields a different consumption path $\hat{C}$ and a different value $V(\hat{C})$. We can express the difference between $V(C)$ and $V(\hat{C})$ in consumption-equivalent terms by calculating the uniform percentage consumption variation Δ such that $V[(1+\Delta)C] = V(\hat{C})$. The main welfare

7. The other sources of consumption data cover a much narrower range of expenditure categories than the Interview Survey, which could account for the discrepancy in consumption inequality trends. This possibility merits careful investigation, but my other observations in the main text suggest that there is more to the story.

experiment in Krueger and Perri's study involves a counterfactual path for the state vectors Y_t^g and Y_t^d. In particular, they generate $\hat{C}$ and $V(\hat{C})$ by fixing the state vectors at their 1972 values.

2.2 EVALUATING THE PROCEDURE

Calculations of this sort are potentially informative about the welfare consequences of changes in the process for consumption or earnings, and they can provide useful inputs into the analysis of inequality trends. The general approach is attractive because it requires one to be explicit about interpersonal and intertemporal utility comparisons, the economic environment that agents face, and the counterfactual scenario. These features facilitate communication and help to sharpen our thinking.

That said, the particular approach in this paper has important drawbacks. First, Krueger and Perri do not adequately model uncertainty about group-level consumption. They assume perfect foresight about the evolution of the state vector Y_t^g, so that uncertainty about group-level outcomes stems entirely from nonzero off-diagonal elements of π^g. In the data, the rank ordering of consumption (and earnings) for sex-education groups is extremely stable over time. In fact, π^g is essentially an identity matrix when estimated from data on households that have the same reference person in t and $t+1$.[8] Of course, when π^g is the identity matrix, rising inequality translates directly into higher utility for groups with rising relative consumption and lower utility for those with declining relative consumption. That is basically what Krueger and Perri find.

In practice, they do not limit the sample to households with the same reference person in consecutive periods when estimating π^g. Instead, the identity and characteristics of the reference person can change in their sample, for example, because of a change in marital status or living arrangement.[9] This type of uncertainty is what Krueger and Perri capture in their specification of the group-level consumption and earnings processes. It is not what leaps to mind when policymakers and researchers ponder the welfare consequences of increased inequality among education groups.

Second, it is unwise to rely solely on the short-panel aspect of the CE to characterize uncertainty about group-level consumption and earnings. We know that relative consumption and earnings among education groups, for example, display large low-frequency movements, and that there is much uncertainty about these movements looking ahead. The

8. The authors confirm this fact in a personal communication. When estimated using households with no change in reference person, the diagonal elements of π^g range from .95 to .98.
9. Their estimate of π^g has important off-diagonal elements, although the diagonal elements remain large, ranging from .80 to .95.

panel aspect of the CE consists of two noisy observations, spaced nine months apart, on the same household. Expecting such data to identify an adequate statistical model for relative consumption movements among education groups is probably expecting too much.

Moreover, there is no need to rely (solely) on the short-panel aspect of the CE to estimate the group-level processes. Instead, one can exploit repeated cross sections in the CE (or other data source) to construct synthetic panel data on group-level outcomes. One can then follow groups defined by birth cohort and education over a period of 20 years or more. The long-panel aspect of such data makes them better suited for estimating the group-level processes, although nailing down the low-frequency properties remains a challenge.

It would also be useful to combine the CE with longer panel data from other sources to characterize the dynamics of the idiosyncratic components of consumption and earnings. For example, one could use the PSID to estimate the degree of persistence in idiosyncratic earnings or consumption changes and combine that information with the household-level changes observed in the CE. In this way, one could decompose within-group inequality trends into separate components associated with transitory and persistent household-level changes. See Blundell et al. (2002) for an analysis that combines cross-sectional data in the CE with panel data from the PSID.

To sum up, there are large gains from drawing on the information contained in CE-based synthetic panels and longer true panels in other data sources. This information can be used to provide richer, more compelling characterizations of the consumption and earnings processes, which are key inputs into the welfare calculations.

REFERENCES

Aguiar, Mark, and Erik Hurst. (2003). Consumption vs. expenditures. University of Chicago Graduate School of Business. Working Paper.

Attanasio, Orazio. (2002). Consumption and income inequality: What we know and what we can learn from it. Plenary Lecture, Society of Economic Dynamics Conference, New York (June).

Attanasio, Orazio, and Steven J. Davis. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy* 104(6): 1227–1262.

Battistin, Erich. (2003). Errors in survey reports of consumption expenditures. Working Paper 03/07. London: The Institute for Fiscal Studies.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. (2002). Partial insurance, information and consumption dynamics. London, England: Institute for Fiscal Studies. Working Paper 02/16.

Davis, Steven J., Felix Kubler, and Paul Willen. (2003). Borrowing costs and the demand for equity over the life cycle. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 9331. Revised version posted at gsbwww.uchicago.edu/fac/steven.davis/research/ (accessed July 20, 2003).

Davis, Steven J., and Paul Willen. (2000). Occupation-level income shocks and asset returns: Covariance and implications for portfolio choice. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7905.
Edelberg, Wendy. (2003). Risk-based pricing of interest rates in consumer loan markets. University of Chicago, Department of Economics. Working Paper.
Fisher, Jonathan D., and David S. Johnson. (2003). Consumption mobility in the United States: Evidence from two panel data sets. U.S. Bureau of Labor Statistics. Working Paper.
Gruber, Jonathan. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review* 87 (1, March): 192–205.
Slesnick, Daniel T. (1992). Aggregate consumption and saving in the postwar United States. *Review of Economics and Statistics*. 74 (4, November):585–597.
Slesnick, Daniel T. (2001). *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*. Cambridge, England: Cambridge University Press.
Sullivan, James X. (2002). Borrowing during unemployment: Unsecured debt as a safety net. Northwestern University. Working Paper.

Comment

KJETIL STORESLETTEN
University of Oslo, Frisch Centre, and IIES

1. Introduction

This paper addresses the welfare consequences of the increase in inequality in the United States during the last three decades. By now, it is well established that since the 1970s there has been a substantial increase in inequality in wages and earnings. The focus of the empirical literature has been to decompose this rise in inequality into transient and permanent components. The motivation for this focus is that permanent shocks presumably have large welfare consequences, while transient shocks are regarded as (self-)insurable. However, this link between earnings and welfare implicitly requires assumptions about market structure and excludes risk-sharing devices that do not show up in wages or labor income (for example, in-kind transfers or means-tested price rebates).

Krueger and Perri pursue a different approach for quantifying the welfare costs of rising inequality by noting that welfare is not derived from income and wages, but rather from consumption and leisure. In particular, Krueger and Perri ask the following two questions: (1) How has inequality in consumption and leisure evolved? and (2) What are the welfare consequences of these changes?

I am very sympathetic to Krueger and Perri's idea of exploring the effects of rising inequality by putting more emphasis on allocations of consumption and leisure, and I believe it is a promising research approach.

Extending previous work (Krueger and Perri, 2002; Fernández-Villaverde and Krueger, 2002), Krueger and Perri document the first question using data from the Consumption and Expenditure Survey (CE). Next, they assess the welfare consequences using a novel approach. In this discussion I compare some of their findings with facts from alternative data. I then discuss two alternatives to Krueger and Perri's approach for evaluating the welfare consequences of the rising inequality.

2. *Revisit Facts Using Alternative Data: PSID*

The dataset that Krueger and Perri use, CE, focuses on consumption, and data on earnings and hours are arguably of lower quality than in the Panel Study of Income Dynamics (PSID) or in the Current Population Survey (CPS). However, Krueger and Perri document that the implications for wage inequality are comparable in the CE and other datasets (PSID and CPS).[1]

Using data from the PSID and CPS, Heathcote et al. (2003) document that hours inequality for men (excluding nonparticipants) has remained constant or experienced a small increase during this period. In contrast, Krueger and Perri find that inequality in hours per adult (i.e., average hours within the household) has declined over time.

One possible reason for this difference could be that even if all workers worked the same number of hours, changes in labor force participation (e.g., from one to two earners) would induce changes in hours inequality. Thus, the rise in female labor force participation may have caused the decline in hours inequality because of an increase in the average number of hours worked for women. Indeed, the number of hours worked for married women has increased during this period.[2]

1. Krueger and Perri argue that earnings inequality increases slightly less than wage inequality. If inequality is measured as the variance of log of earnings (instead of as the standard deviation of logs), however, earnings inequality increases slightly *more* (consistent with the finding of Heathcote et al. [2003] for the PSID).
2. Using data from the CPS, Jones et al. (2003) document that, while average hours worked for single men and women was relatively constant over this period, hours worked for married women rose sharply.

3. *Welfare Effects of Rising Inequality*

Given the facts on the evolution of consumption and hours inequality, Krueger and Perri examine the distribution of welfare effects of changes in inequality, conditional on initial state. They propose a novel, theory-free approach. They start by estimating an exogenous stochastic process for individual household consumption and leisure, exploiting the panel dimension in the CE data.[3] The data are detrended so that the welfare effects of changes in the process are due to changes in the higher-order moments only. They then assume a time-separable utility function over consumption, consumer durables, and leisure,

$$u(c, s, l) = \frac{1}{1-\sigma}\left[(c^{\theta} s^{1-\theta})^{\alpha} l^{1-\alpha}\right]^{1-\sigma},$$

and compute discounted utility given the time-varying processes for consumption and leisure.

The large dispersion in welfare effects are mainly driven by changes in between-group inequality and are hardly affected by changes in within-group inequality. One reason is that, conditional on group, the estimated processes for consumption and leisure are not very persistent. For example, holding group-specific characteristics constant, the autocorrelation of individual consumption is 0.72, lower than the autocorrelation for earnings.

This finding points, I believe, to a potential shortcoming of the Krueger and Perri approach. Assume that preferences are separable between consumption and leisure. The permanent income hypothesis then suggests that the marginal utility of consumption, and therefore consumption itself, should be very persistent. Why don't the data have this property? Of course, preferences may not be separable between consumption and leisure. However, there are good reasons to believe that the estimated persistence of consumption is biased downward due to measurement error. For example, Cogley (2002) suggests that measurement error in CE consumption biases upward the variance in individual consumption growth by one order of magnitude. Clearly, if the Krueger and Perri consumption process is mismeasured, it casts some doubt on their quantitative welfare findings.

3. Note that it is not the transition process between actual consumption levels that is estimated, but the transition probabilities between different consumption classes. This approach underestimates the consumption inequality and, in particular, the change in consumption inequality.

4. *Alternative Route (i): A Theory-Free Approach*

If one is to pursue a data-based, theory-free approach, why is it necessary to estimate a process for consumption and leisure instead of simply plugging in the actual data? In the spirit of Atkinson (1970), one could alternatively address the following welfare question: Under the veil of ignorance, what fraction of initial consumption would agents give up to get that allocation forever, relative to experiencing the subsequent evolution of inequality in consumption and leisure? Assuming that preferences are time-separable, it is only the dispersion in consumption and leisure that matters, and movements within the distribution are irrelevant. Thus, one could simply plug in the actual data observations and discount utility, given explicit assumptions about the utility function.

Pursuing this approach with the same data and utility function as Krueger and Perri used, the welfare losses, expressed as a fraction of lifetime consumption, are as follows:[4]

	c only	*(c, l)*	*(c, s, l)*
$\sigma = 1$	1.57%	1.48%	1.43%
$\sigma = 2$	2.54%	2.15%	1.94%

The first column refers to the welfare effects of the changes in inequality of nondurable consumption. The second and third columns add leisure and services from consumer durables. The key message of this list is that the average welfare loss is around 1 to 2%, which is in the same ballpark as the findings of Krueger and Perri.

5. *Alternative Route (ii): A Structural Approach*

The most serious critique of the Krueger and Perri approach is perhaps that the preferences are arbitrary because they are not necessarily consistent with the observed individual behavior. For example, the preferences considered by Krueger and Perri exhibit quite high individual labor elasticity (2 for $\sigma = 1$ and ½ for $\sigma = 2$). There are reasons to believe that the costs of changing inequality have been unevenly distributed across generations, with young in the 1980s shouldering the largest burden, an aspect absent in the infinite-horizon approach of Krueger and Perri.

4. The figures display the welfare loss of the changes in inequality after 1980 because the 1972–1973 CE data include fewer consumption items than the 1980–2000 data and are therefore not directly comparable.

An alternative route, pertinent for this criticism, would be to pursue a structural approach for quantifying the welfare costs. In particular, one could use an individual-specific wage or earnings process as a primitive, generate endogenous consumption and leisure allocations from a structural model, and subsequently use these findings to evaluate welfare consequences.

One paper pursuing this route is Heathcote et al. (2003). They estimate changes in the individual wage process in the United States using PSID data and document increases in the transitory, persistent, and permanent components of the wage process. They then formulate a standard life-cycle version of the permanent income hypothesis model with savings in one riskless bond and a consumption-leisure trade-off. The preferences are of the constant elasticity of substitution type and separable in time and between consumption and leisure. Their model is calibrated to capture key cross-sectional facts, resulting in quite plausible parameters (for example, a Frisch elasticity of 0.5 and a relative risk aversion for consumption of 1.5).[5] This model accounts for the salient features of the evolution in inequality, such as the evolution of the wage-hours correlation and the inequality in earnings, consumption, and hours. The fact that the preferences are consistent with the observed individual behavior makes the welfare calculation (including the particular utility function) less arbitrary, I believe.

Turning to welfare, Heathcote et al. (2003) find that, under the veil of ignorance but conditional on cohort, the welfare loss of changes in wage process is 2 to 5% of lifetime consumption for households entering the job market during 1970–2000, and around 1% for households entering the job market during the 1950s.[6]

REFERENCES

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of Economic Theory* 2:244–63.

Cogley, T. (2002). Idiosyncratic risk and the equity premium: Evidence from the Consumer Expenditure survey. *Journal of Monetary Economics* 49:309–334.

Fernández-Villaverde, J., and D. Krueger. (2002). Consumption over the life cycle: Some facts from the CEX data. Stanford University. Mimeo.

Heathcote, J., K. Storesletten, and G. L. Violante. (2003). The cross-sectional implications of rising wage inequality in the United States. New York University. Mimeo.

5. In particular, Heathcote et al. (2003) match the standard deviation of changes in individuals' number of hours worked and the correlation between wages and hours worked.

6. The preference parameters used by Heathcote et al. (2003) differ from the Cobb-Douglas specification of Krueger and Perri that I used in Section 4. Repeating those calculations with the parameters of Heathcote et al., the CE data indicate a welfare loss of 1.9%.

Jones, L. E., R. E. Manuelli, and E. R. McGrattan. (2003). Why are married women working so much? Federal Reserve Bank of Minneapolis Staff Report.
Krueger, D., and F. Perri. (2002). Does income inequality lead to consumption inequality? Evidence and theory. Stanford University. Mimeo.

Discussion

Fabrizio Perri first responded to some of the discussion participants' concerns about data quality. He was aware of the fact that using different datasets produced different estimates of the increase in wage inequality in the sample period. He asserted that, although the estimates from the Consumer Expenditure (CE) survey are bigger than those obtained using the Current Population Survey (CPS), they are similar to those obtained using the Panel Study of Income Dynamics (PSID). Perri said that he was aware of the differences between the income and diary data from the CE, and said that the authors might consider using the diary data in the future. He also said that he was aware of the underreporting of both consumption and income in the CE. He also noted that as long as the income and consumption data come from the same sample, income inequality increases while consumption inequality remains flat. On the question of the identifiability of the stochastic process for between-group inequality, Perri agreed that using 1-year variation is not ideal. He pointed out that a large number of households are used to make the estimate, however, and that the results are not very different from what would be expected.

Annamaria Lusardi echoed the concerns of Steve Davis about the severity of measurement error in the CE. Responding to Steve Davis's comment that the CE and National Income and Product Accounts (NIPA) data do not match, she suggested that it is not obvious that the NIPA has the correct numbers. Fabrizio Perri responded that a big difference between the CE and the NIPA is that the CE does not include data on very rich people.

Annamaria Lusardi suggested that counting services from durables might yield a higher volatility of consumption. Eva Nagypal was concerned that the authors did not take account of the increase in assortative matching of spouses along observable characteristics over the sample period.

Several participants were concerned by the smoothness of the consumption data and suggested investigating this point further by examining the data on savings more closely. Annamaria Lusardi asked whether both savings and wealth data are consistent with the degree of consumption smoothing apparent in the data. She noted that wealth inequality increased a good deal in the 1990s. Deborah Lucas remarked that the

authors hypothesize that the disconnect between consumption and income may be facilitated through financial markets. She was worried by the fact that the group that would have had to use financial markets to increase their consumption had almost no savings in the data and found this particularly troubling in light of the persistence of income shocks. Aart Kraay was struck by the fact that those in the bottom decile of savings rates in the data used in the paper had an average dissaving rate of 10% over 15 to 30 years. He wondered whether dissaving of this order of magnitude could be consistent with what is known about initial asset stocks. On the issue of savings rates, Fabrizio Perri responded that it was true that the gap in the savings rate widened over the sample period. However, he pointed out that the households in the bottom decile of savings rates in one year were not necessarily the same as the households in the bottom decile in other years because households with temporarily low realizations of income dissave, while those with temporarily high realizations save.

On the issue of the measurement of welfare, Fabrizio Perri reminded the participants that the welfare numbers in the paper are ex-post numbers. He noted that they answer the question, How bad has bad luck been for unlucky people? not How much would people have been prepared to pay in 1972 to avoid the future change in their income process? Jonathan Heathcote suggested using a welfare measure based on the assumptions of a utilitarian social-welfare function and separability over time and across consumption and leisure because this measure would require looking at the cross-section distribution of consumption and leisure alone. Fabrizio Perri responded that the welfare losses of increased inequality calculated in this way would be small because the increases in inequality in consumption and hours worked are small.

Mark Gertler speculated that there might be a link between the jump in between-group and within-group inequality around 1984 documented in the paper and the decline in macroeconomic volatility at that time documented by James Stock and Mark Watson in the *Macroeconomics Annual 2002*. Fabrizio Perri suggested that credit market changes could explain the smoothness of consumption relative to income and potentially also the decrease in macro volatility. Ken Rogoff pointed out that the improvement in credit markets was specific to the United States, whereas the decline in macroeconomic volatility was common to many developed countries. Fabrizio Perri pointed to the finding of Alberto Alesina and coauthors that inequality affects happiness in Europe but not in the United States as potential evidence that credit markets are less effective in facilitating consumption smoothing in Europe than in the United States.