

Inequality in Living Standards since 1980: Evidence from Expenditure Data.*

Orazio P. Attanasio^{*}, Erich Battistin[♦] and Mario Padula[♥]

May 2009

* We would like to thank the audience at the AEI presentation of a first draft in September 2007 and, in particular, Steven Davis for much useful feedback.

^{*} University College London, IFS, NBER, CEPR and BREAD.

[♦] Università di Padova and IRVAPP.

[♥] Università “Ca’ Foscari” di Venezia and CSEF.

1. Introduction.

The evolution of inequality in wages and household incomes, as well as their levels, has been widely investigated. A large literature has documented a substantial increase, in the last 30 years, in the dispersion of wages and incomes in the US as well as in other OECD countries. This increase was particularly quick and marked in the 1980s, even though it continued, to a certain extent, in the 1990s and 2000s. For instance, as reported in Autor, Katz and Kearney (2007), the ratio of the 90th to the 10th percentile for full time weekly earnings increased by about 23% between 1980 and 1992 and a further 12% between 1992 and 2003. The increase in inequality during the 1980s was also accompanied by a decline in the wages and incomes of the bottom of the distribution.¹

In this monograph, we will be arguing that the picture offered by the evolution of income and wages inequality is a partial one and that much can be gained from the analysis of the distribution of consumption and expenditure. The latter does not only provide, for a number of reasons we discuss, a better and more informative measure of inequality, but it also gives, when analyzed *together* with the evolution of income inequality, new insights on the nature of the factors that affect changes in wage inequality and on the institutions and tools individuals have to smooth out income shocks.

A first argument for looking at consumption and its distribution is the obvious observation that income, for the most part, is valued because it allows consumption. Therefore, looking at consumption flows directly might be more informative about individual material well-being. Second, consumption can be informative about ‘permanent’ changes to income: a temporary (or perceived as such) income shock received by an individual does not necessarily lead to a change in her consumption. A positive shock might be saved and a negative shock can be buffered, in many situations, running down savings, borrowing, and/or using different forms of public and private transfers. A change in individual resources that is perceived to be permanent, instead, probably will lead to a change in consumption. Third, comparing the evolution of income and consumption inequality can be informative about the tools and institutions

¹ The trends of the second part of the 1990s and early 2000s seem qualitatively different from what happened in the 1980s and in the early 1990s. Several authors have observed that inequality increased more slowly over the later period and that this increase was more in inequality within than across skill groups. Moreover, as documented in Autor, Katz and Kearney (2007), the change in inequality that did happen over this period is mainly driven by inequality in the top part of the distribution (for instance, the 90th/50th percentile ratio increases, while the 50th/10th is constant).

available to an individual for smoothing different types of income shocks. A change in income inequality that is reflected in a change in consumption inequality is probably indicative of shocks that cannot be 'smoothed' or absorbed. This could be because the shocks are perceived as permanent or because they are too large to be buffered with the available instruments. On the other, increases in income inequality that are not reflected in consumption inequality reflect the ability of individuals to use various instruments to smooth such shocks. While establishing which instruments (individual savings, borrowing, private or public transfers) are used for such a purpose is important, a first important step is to establish to what extent income shocks (to be defined more precisely) are reflected into consumption.

Of the three sets of reasons given in the previous paragraph for analysing the evolution of consumption inequality, the first is obvious. The second and third, however, are more subtle and merit further elaboration. The inequality of individual and household incomes changes over time because individual and household incomes change over time and do so in a way that is not coordinated. Some individuals will gain and others will lose or gain less. The consequences of these 'shocks' for the distribution of individual well being depends on the *nature* of these shocks. Permanent changes will obviously be much more relevant than temporary ones. While it is possible to gain some insights on these issues from the analysis of longitudinal data and from appropriate 'cuts' of the data, the analysis of consumption can be much more effective in giving information on the nature of shocks even in the absence of longitudinal data. This is the essence of the second point made above.

A similar argument applies to the third set of reasons. Ultimately, the consequences of some shocks to individual incomes for the distribution of individual well being depends on the extent to which these shocks are reflected into consumption and, therefore, depend crucially on the nature of institutions in place to smooth shocks received by individuals and to provide, possibly, some insurance against them. These institutions include the welfare state, with its many programs, the tax code, the possibility of borrowing in the face of temporary adverse shocks, the availability of interpersonal connections and so on and so forth. The role these institutions play, however, can be complex and subtle as they, in all likelihood, also affect incentives and, indirectly, the distribution of income. Some of them, such as interpersonal transfers, can also be affected by the nature of incomes that individuals receive and by other (maybe public) transfers.

In addition to the analysis of consumption inequality, the joint analysis of consumption and income *levels* can also be of considerable interest. Recent empirical evidence for the US and the UK has shown that consumption-poor households do not coincide with income-poor households. In particular, income poor households report consumption levels way above their level of income.² Underreporting of welfare income and other informal sources of income may preclude a correct interpretation of the income dynamics at the bottom of the distribution. Moreover, the picture that emerges from the expenditure definition is quite different depending on the survey instrument considered. Because of this, it seems rather obvious to use the distribution of both income and expenditure to come out with a better measure of individual well-being.

In this monograph, we will analyze the evolution of consumption (and income) in the US using the best available data. In what follows, in addition to the analysis of the dispersion in income and consumption, we will also consider the evolution of the *levels* of income and consumption for different groups of the US population. In this respect our main focus will be on individuals with different levels of academic achievement and born in different decades. The comparison between the average levels of different groups and their evolution over time constitutes an important dimension of inequality. In a sense, overall inequality can conceivably be decomposed into differences across groups and inequality within groups.³

Our main findings can be summarized in the following three points.

1. Individuals and households that are identified as ‘income poor’ or as at the bottom of the income distribution are not necessarily the same as those identified as ‘consumption poor’ or at the bottom of the consumption distribution. For example, in Table 3 in Section 6 we find that 43% of households in the bottom 10 percent of the earnings distribution have consumption levels in the top 60% of the consumption distribution.

² See Meyer and Sullivan (2004) for single mothers in the US and Attanasio, Battistin and Leicester (2006) for couples both in the UK and in the US

³ Depending on the inequality measures used, these decompositions may or may not be possible. It is however worth noting that Battistin, Blundell and Lewbel (2007) provide convincing evidence that the distribution of consumption expenditures is roughly lognormal. This makes the informational content of any index of inequality equivalent to that of the variance of logs, and the study of between and within group components of inequality equivalent to that of a standard analysis of variance. The same result does not apply to the distribution of income, which is instead characterised by a marked departure from log normality.

2. The dynamics of income, wages and consumption inequality has been quite different over the past 25 years. While wages inequality (as measured by the standard deviation of logs) has increased by about 15 percentage points, income has increased by about 10 points and consumption by about 7. These figures, for a variety of conceptual and data problems discussed below are not uncontroversial.
3. The dynamics of consumption and wage inequality, as measured by differences in means across groups defined over decade of birth and education, were very related until the early 1990s and much less so after that. In Table 4 in Section 7, the correlation between wages and consumption means across groups, after removing fixed group and time effects goes down from 0.88 before 1992 to 0.06 after that.

The results we obtain can be seen as some steps of a research strategy that aims at characterizing the evolution of the distribution of material well being and its determinants by starting with the analysis of the remuneration of individual skills, to move on to income, then to expenditure and finally to consumption.

Wages faced by an individual would reflect the price she can get for her labour, but are only one determinant of household income (see Pencavel, 2006). Inequality measures based on wages and income will differ for two important reasons. First, there is abundant evidence of positive assortative mating, that is, that the incomes (and education) of spouses are positively correlated. This would therefore compound wage inequality, as better-educated individuals are married to each other. Second, labour supply behaviour can be different among individuals (and couples), which might increase or decrease inequality. Labour supply can also respond to specific shocks to wages (or employment) magnifying or reducing their effect. The evidence on the importance of these factors is still remarkably limited.

The next step of such a research strategy could then lead to the analysis of disposable income, including therefore the effect of the tax and benefit system. This step would stress the ‘insurance’ role played by the tax and benefit system. The analysis of total expenditure inequality would then constitute the final step of such a strategy. This might differ from that of total disposable income because individuals save or dis-save, moving resources to or from the future and because of other transfers (maybe intrapersonal ones) that are not necessarily measured in income. And even within expenditure, it may

be of some interest to look at its different components. It is quite possible that some components (such as durables) might be more affected than others (such as food) by certain income shocks. A household receiving a negative shock might postpone the purchase of a new car or the replacement of an old piece of furniture rather than diminishing the expenditure on necessities.

The joint analysis of consumption and income decisions puts the analysis in a dynamic framework. It then becomes crucial how individuals move resources over time in the presence of uncertainty. In an ideal world, without any frictions and imperfections, individuals could purchase insurance products in competitive markets that would allow diversifying efficiently idiosyncratic risk, or could access credit markets that would allow them to smooth out temporary shocks to income. In reality, markets, albeit sophisticated, are not perfect. Markets for insurance and credit in particular, are subject to a host of problems, mainly information asymmetries, which prevent a complete diversification of risk. An interesting question is to establish the extent to which the existing set of institutions allows individuals to smooth part of their income shocks. For such a purpose, the joint analysis of income and consumption is crucial. If one sees income dispersion going up, while consumption dispersion increases by less, it must mean that individuals are able to smooth out some of the shocks that generated the increase in income dispersion. A nice feature of this approach is that it can afford to be silent about the specific tools that individuals use to achieve such a goal.

The real world is complex and individual households use a very wide variety of instruments to smooth out individual shocks. This is true across the income distribution, although the instruments that different households use are probably very different. In the case of relatively poor households, they are likely to involve a variety of informal arrangements (as well as the complexities of the welfare system) that are difficult to observe. In this sense the observation of consumption for these households is particularly useful.

While we do think that the analysis of expenditure individual data is key to get a complete picture of the evolution of the distribution of material well being, there are good reasons for the very limited use of expenditure data so far. Probably the most important one is the limited availability of reliable and comprehensive data that cover the relevant periods and that are of sufficient breadth to allow the construction of reliable measures of well being and of inequality. Comprehensive surveys that collect information

on expenditure and expenditure pattern have been available for a long time. The first version of the Consumer Expenditure Survey (CEX) was collected in 1917. Unfortunately, as the main use of these data was the computation of the weights for the Consumer Price Index, until 1980s they were only collected at about 10 years intervals. Moreover, the methodology for collection was not homogenous. Starting in 1980, the Consumer Expenditure Survey became a survey collected continuously. However, unfortunately, the survey has important limitations. First and foremost, as we discuss below, if one tries to aggregate the micro data up using the appropriate weights, one has a hard time in reproducing National Account data on Personal Consumption Expenditure (PCE). Second, the relationship between aggregated CEX data and PCE data has worsened over time. Third, the size of the CEX is quite limited, with only 5000 households per year contacted between 1980 and 1998 (the size has somewhat increased after 1998). This limited size makes the analysis of inequality and, in particular, the study of different subgroups of the population problematic and imprecise.

Other surveys do collect information on individual outlays. However, they have other even more important limitations. The Panel Study of Income Dynamics (PSID), which was started in 1968 and is one of the most widely used survey, only collect some information on food expenditure and a few other items. Recently, the Health and Retirement Survey (HRS) supplements its main core of data with a postal survey on consumption. However, the HRS is only representative of older individuals. In the end, therefore, with all its limitations, the CEX is probably the best source of information on consumption. In what follows, we will be arguing strongly for a revamp of the CEX that would improve its quality and increase its size.

While the three points above are our main findings, to get there we need to address a number of technical and data issues. For this reason, the monograph is organized in 7 sections. In all sections we keep technical details at a minimum and cite only the most important contributions in the literature. More details and citations are given at the end of some of the sections in a short subsection entitled *Further readings*.

After this introduction, **Section 2** lays out the methodological and conceptual issues. The section examines the relation between consumption and income inequality and develops the arguments above. We first argue that measures of inequality based on consumption better reflect long run differences in household and individual welfare. We also argue that the comparative analysis of consumption and income inequality is

informative about the nature of insurance markets and possibly about the nature of imperfections.

In **Section 3**, we discuss measurement issues. We tackle three types of problems. First, we discuss the quality of our main data source, the CEX, and how it has varied over time. Second, we discuss the methods we use to combine the information from the two independent components that make the CEX. Finally, as we distinguish expenditure from consumption, which is important if one uses welfare measures based on non-durable as well as durable consumption, we describe the methods employed in constructing our estimates of the service flows provided by durables.

The more recent trends in wages and income inequality are documented in **Section 4**. In this section, we employ all the available data sources trying to make sense of the sometimes conflicting evidence that comes out of them. Moreover, we identify the socio-economic groups' specific trends and decompose overall inequality into its within and between group components. We particularly focus on the differences by year-of-birth and skills groups. Furthermore, we examine the degree of co-variation of earnings within the household.

We present evidence on expenditure and consumption inequality in **Section 5**. We provide information both on non-durable and total consumption, starting with the narrowest definitions that include only non-durable and services, to the widest, which include the services provided by durables. Again, the analysis illustrates the overall pattern of inequality as well as pattern specific to decade of birth and education groups. This evidence (and that presented in Section 4) is behind finding 2 mentioned above.

Section 6 considers the differences in distribution of income and consumption focusing in particular on the distribution of consumption at the bottom of the income distribution. This section uncovers, among others, our finding 1 mentioned above.

Section 7 relates the trends of consumption and income inequality. In particular, we consider both the relationship between the means (and their changes) of relative wages and consumption of different groups and the relationship between the inequality *within* the same groups. This section contains our finding 3.

Section 8 concludes the monograph with a summary of our main results and some considerations about future research and about the need of high quality data on consumption.

Further readings.

The evolution of wages and (to an extent) income inequality in the US is now well documented. There is an enormous literature on the topic that we cannot hope to summarize here. Some of the best known early papers on the topic are Katz and Murphy (1992); Murphy and Welch (1992); Juhn, Murphy and Pierce (1993). Some analysts, such as Gottshalk and Moffitt (1994), have tried to decompose the increase in that part due to transitory and to permanent changes. Gottshalk and Moffitt (1994) have attributed 1/3 to temporary and 2/3 to permanent changes. The rise in inequality has continued in the nineties, but at a lower rate. Gottschalk and Moffitt (1994), Gottschalk and Smeeding (1997), or Katz and Autor (1999) have related the increase in income inequality to an increase in the dispersion of life-time resources and in the volatility of high frequency shocks. A recent paper that summarizes much of the literature and provides some new insights is Autor, Katz and Kearney (2007).

Following Cutler and Katz (1991), several authors have used the CEX to study the evolution of consumption inequality. Attanasio and Davis (1996) showed, among other things, that the evolution of consumption inequality across groups defined in terms of education achievement of the household head mirrored closely that of wages. Since then, however, the picture has become murkier. Some authors, such as Krueger and Perri (2006) have claimed that consumption inequality has not raised much, while other authors, such as Attanasio, Battistin and Ichimura (2007) have claimed that it has. Similar claims are contained in Blundell, Pistaferri and Preston (2008).

The conflicting pieces of evidence bring us to the main problem with the analysis of consumption inequality: that of data quality. In developed economies, consumption is notoriously difficult to measure⁴ and household level data bases containing detailed and high quality information on consumption are few and far between. Often, these data bases are collected to obtain information used in constructing the weights for the Consumer Price Index. The US is no exception in this respect. The CEX, which constitutes the only household level source containing detailed and complete consumption information over the period we are interested in studying, has a number of problems. The sample size is not very large and there are indications that the quality of

⁴ Interestingly, in developing countries usually the opposite is true: consumption is much easier to measure than income. This is both because the consumption basket is remarkably simple and because income can be derived by multiple and disparate sources. To a certain extent this is also true in the bottom of the income distribution for developed countries. Meyer and Sullivan (2004) argue that the consumption of poor single mothers can be measured with more precision than their income.

the data has deteriorated over time (though very recent research from the Bureau of Labor Statistics provides a different view). We discuss these issues at length in Section 3.

Box: Consumption vis-à-vis Income inequality

If, for a variety of reasons, one is interested in the distribution of material well-being across individuals or households, the distribution of wages and income inequality, as measured by the snapshot of a cross section at a point in time, tells only part of the story. Suppose, for instance, that wage dispersion among a population of individuals at a point in time is caused in part by permanent differences (for instance differences in abilities and skills) and in part by shocks received by individuals due to transitory events, such as a lay off, an illness and so on. Consider first a scenario in which the transitory shocks do not exhibit much persistence: an individual hit by a negative shock is not more or less likely in the future to be hit again by a similar shock or to be hit by a positive shock. This implies that, within a group of individuals with a certain (permanent) level of income, one should observe a considerable amount of ‘mobility’: individuals with low levels of income today will not necessarily have low levels of income tomorrow. An alternative scenario, instead, is one in which even transitory shocks have persistence. In this alternative world, one would observe much less mobility. An increase in the level of dispersion (as measured, for instance, by the size of the shocks that individuals receive) would then have very different consequences for the material well being of the individuals in these two worlds. An increase in dispersion in the ‘high mobility’ world would have much less severe consequences for the dispersion of material well-being than the same increase in the ‘low mobility’ world.

In the previous paragraph, the difference between the low and high mobility scenarios was induced by the different nature of income shocks. An alternative and interesting example is the following. A first scenario is where the observed increase in income inequality is caused by an increase in the remuneration of skills. Given a certain distribution of skills, individuals with a larger stock of them will now be remunerated comparatively more than before. The alternative scenario is one where the observed increase in income inequality is generated by an increase in the dispersion of ‘temporary’ shocks. Once again, the consequences for material well being would probably be very different.

Similar considerations could be made when talking about differences in institutions that allow the smoothing of these shocks. In a world where there are institutions or tools that allow individuals to smooth out income shocks, an increase in the dispersion of income shocks can be completely undone by appropriate insurance. The important point to make is that the static study of the distribution of income might provide very partial information and that much is to be gained by looking at consumption and expenditure.

Indeed, even the comparison across different countries takes a different angle when seen in a dynamic context. As stressed in recent studies by Flinn (2002) and Bowlus and Robin (2004), the US looks much more unequal than continental Europe when considering a single snapshot. However, the picture is very different when considering life time resources. The difference between the two images is explained by the higher level of income mobility that characterizes the US.

2. Consumption versus wage and income inequality.

This section argues that:

- **A proper understanding of the evolution of the distribution of well-being requires the analysis of both income and expenditure data.** Shocks to income do not necessarily cause changes in consumption and well-being. Individuals can borrow, rely on past savings or on the public welfare to prevent income shocks from affecting consumption. Therefore, a temporarily low income does not necessarily induce low consumption and a decrease in material well-being. This makes consumption a useful measure of well-being, which does not require observing the actions individuals engage to smooth out adverse income shocks.
- **Expenditure data alone are not enough to measure well-being.** Consumption and expenditure are two different concepts. Individuals' welfare depends on consumption, but data normally measure expenditure. For non-durable goods, such as food, expenditure can be taken as a good proxy for consumption, for durable goods (cars, fridges, dish-washers, and so on), expenditure is a very poor proxy of consumption. The expenditure on durable goods is lumpy and infrequent but individuals enjoy the services from such goods over a certain period of time even if the expenditure is zero.
- **The joint examination of consumption and income data provides valuable information on the evolution of the distribution of material well-being.** If consumption is lower than income, then part of the latter will be available for future consumption. If a change in income is not reflected in a change in consumption, then it is an indication that the household might be able to smooth out that particular income shock. Therefore, the dynamic aspect of individual choices can be only understood if one focuses on the joint distribution of income and consumption.

The premise of this monograph is that the evolution of material well-being and its distribution across individuals is an interesting and important issue. It is important because it is central to the design of economic policy and, in particular, of welfare systems. It is interesting because it is informative about the functioning of developed economies and how they distribute resources. It should be stressed that both the *level* and the *distribution* of well-being are important. An exclusive focus on averages is clearly very limitative, as the average might potentially represents the experience of no individual. On the other hand, an exclusive focus on the dispersion is insufficient as very different levels of resources can generate the same level of inequality and therefore has very different implications.

Given this premise, one needs to address two issues: (i) how to measure economic well-being and, (ii) given the available data, how one uses them to learn about the distribution of well-being. One alternative, which has been widely used in the academic and policy literature, is to look at the evolution of income distribution. Another alternative, instead, is to look directly at the use of economic resources, that is, at consumption. Yet another alternative is to consider expenditure, where expenditure might differ from consumption because of the durability of some items, which are not immediately consumed upon purchase. In this section we discuss three issues. First, what are the advantages of using consumption rather than income as a measure of well being. Second, we discuss briefly the differences between consumption and expenditure and the extent to which available data can address the problems related to this difference. Third, we discuss how the joint analysis of consumption and income data can be informative about the consequences of an increase in income dispersion for individual well-being.

2.1. Income vs Consumption

From an intuitive point of view, it would seem that looking at the distribution of consumption should be the most profitable strategy to study the distribution of well-being. It is consumption that gives individuals utility, and, usually, income is appreciated because it makes consumption possible. More importantly, the consequences for the material well-being of an income shock depend on the ability an individual has to smooth it. If the shock is perceived to be permanent and if the individual can smooth it its consequences will be small, and consumption will not change (much). To achieve this, the individual will engage in some transactions (dis-saving, borrowing, receiving transfers

from public or private sources) that might be difficult to observe or even to categorize. Observing consumption choices has the advantage to sidestep such transactions. And yet, the big majority of studies that have analyzed distribution issues have looked at income, rather than consumption inequality.

Probably the main reason for the prevalence of studies that look at income is the availability of high quality data. In addition to the objective scarcity of high quality individual level consumption data, however, there has probably also been some degree of inertia and a resistance to use new concepts and data that might be unfamiliar.

In this monograph, we will present extensive evidence based on expenditure data. These data is not exempt from problems, some of which we discuss extensively in the next section. However, it can provide very valuable information on the questions at hand. And, in some respects can be of superior quality to income data. Meyer and Sullivan (2004), for instance, argue that for the bottom of the distribution information on consumption can be of better quality than the information on income as the former has a relatively simple structure, while the latter can be quite complex as it includes, in addition to earnings, welfare transfers, interpersonal transfers, informal income and so on.

2.2. Consumption vs. Expenditure

If the main motivation for looking at consumption rather than income is that it is the former rather than the latter to provide utility, the fact that often only data on expenditure rather than consumption are available constitutes a problem. The two concepts differ for a variety of reasons. In the extreme case of large durable goods, what the individual consumes are the services provided by that good, while the expenditure represents the lumpy purchase of a unit of the good that, in many cases, occurs relatively infrequently. At the other extreme are cases of perishable food items that are consumed at the same frequency with which they are bought. In the middle there will be many intermediate cases, ranging from storable food items, which some times are bought in bulks, to clothing and footwear, which in many cases lasts more than a quarter or a year.

The problem is one of data. In principle one would like to observe, for any commodity with some level of durability, the service flow provided by the stock available to the consumer, in addition to the flow of expenditure. Unfortunately, detailed information on stocks is rarely available in household surveys. In what follows, we present some information on the stock of vehicles available to the consumers in our main data sources. This type of information, however, is more an exception than a norm.

There are no perfect solutions to this problem. In the end one has to make do with information on expenditure. The standard approach is to try to distinguish the expenditures for the acquisition of ‘non-durable’ commodities and ‘services’ from the expenditure for the acquisition of durable commodities. Most of our analysis will be based on an aggregate defined as ‘non-durable and services’. While we will occasionally be referring to inequality in this aggregate as ‘inequality in consumption’, we should keep in mind that this is only a proxy for total consumption. Where we can, we will use information on the stocks (such as in the case of vehicles). This type of exercise will always be limited by the availability of appropriate data.

The point here is more substantive than academic if we are interested in the extent to which shocks to income are reflected in changes in consumption (and expenditure). As different commodities provide different types of services and utility, and are characterized by different degree of lumpiness, it is likely that a consumer will adjust differently the expenditure on different items when facing a shock to income. There is some evidence, for instance, that expenditure on durables is much more sensitive to shocks than the expenditure on food: when facing a short term income problem an individual is more likely to postpone the purchase of a new car than to reduce the amount of food her family eats.

There is some arbitrariness about whether a certain commodity is a durable or a non-durable: clothing is a good example of a commodity that could go in either group. And even for services, some items, such as health or education should have (at least one hopes!) durable effects. In these cases, we exclude these expenditures from our basic ‘non-durable and services’.

2.3. Analyzing income *and* consumption

Having stated that expenditure (or more precisely consumption) can provide more valuable information on the evolution and distribution of material well-being than income, it should be stressed that both variables are important. And indeed, we argue below that the joint consideration of income *and* consumption can be particularly informative.

When one considers the discussion in terms of joint movements in consumption and income, one immediately casts the questions of interest in a dynamic context: if consumption is less than income, then part of the latter will be available for future

consumption (or/and for consumption by other individuals).⁵ If consumption is larger than income, savings will have to be run down or the individual will have to borrow, or receive transfers from private or public sources. In the first two instances, this implies a reduction of future consumption. In addition to the issue of allocating one's resources intertemporally, one is also confronted with the presence of uncertainty about the future, which is likely to affect decisions. In order to decide how much to consume and to save at each stage of their life-cycle, consumers must be able to forecast their future income.

Expected future income might or might not be equal to realized future income. Since consumers are likely to be adverse to change consumption over time, consumption is updated in the face of an income change only if such a change is expected to be permanent. Shocks, such as a sudden lay-off or a disease, can cause income to be different from what it was expected. If possible, consumers are likely to absorb these shocks either by de-cumulating saving or saving less or borrowing. Consumption therefore should react to income shocks and the size of the change should depend on the nature of shocks. Only shocks to life-time resources, such as a promotion or, relevantly for what we will be discussing, a permanent change in the remuneration of skills, maybe induced by technological innovations, should entail substantial revision of consumption. The welfare consequences of transitory and permanent income shocks are therefore very different and their balance depends on the availability of various smoothing mechanisms.

Suppose, for instance, that the inequality of income increases because its temporary components have become more volatile. It might be more common to be laid off and having to find a new job, although this might not necessarily mean a decrease in average earnings over a long period of time. Suppose also, that an individual is aware of this situation and has access to a number of mechanisms that can help him to smooth out such shocks. In the cross section, this means that the dispersion of income increases while the dispersion of consumption will *not*. On the other hand, if the increase in the cross sectional dispersion of income reflects permanent shifts and/or the consumers does not have the tools to buffer the shocks that hit her income, one should witness that the cross sectional variance of consumption mirrors the increase in that of income. This suggests complementing inequality measures based on income with measures based on consumption.

⁵ Here we ignore the distinction between consumption and expenditure discussed above and assume to have information on consumption. All caveats discussed above remain, obviously, relevant.

Further readings

Cutler and Katz (1991, 1992) and Slesnick (1993) also use consumption inequality as a measure of welfare inequality. Attanasio and Davis (1996) analyse jointly changes in average consumption and wages for different groups in the population to assess the extent to which changes in remunerations are reflected in changes in consumption (and presumably welfare) at different horizons. Attanasio and Davis (1996) interpret their results as being informative about the availability of mechanisms that allow risk sharing. Krueger and Perri (2006), have pointed to the fact that, at least using data from the Interview component of the CEX, consumption inequality does not seem to increase in the US in the 1990s, unlike income inequality. This implicitly implies that some insurance mechanisms are available to households from market (or non-market) sources.

Blundell and Preston (1998) derive the technical conditions that allow using consumption inequality as a measure of welfare inequality. Blundell and Preston (1998) also show how to use information on the evolution of consumption and income inequality jointly to identify changes in permanent and transitory income inequality. Blundell, Pistaferri and Preston (2008) combine information from the cross-sectional distribution of income and consumption in a longitudinal data set, to extend the methodology in Blundell and Preston (1998) and identify the amount of insurance available to each household.

Family networks (Attanasio and Rios-Rull, 2000), the timing of durable purchase (Browning and Crossley, 2000), progressive income taxation (Mankiw and Kimball, 1989, Grant, Koulovatianos, Michaelides and Padula, 2009), personal bankruptcy law (Fay, Hurst and White, 2002), financial assets (Davis and Willen, 2000) are a partial list of the smoothing devices that households employ to isolate consumption from income idiosyncrasies.

The relative importance of these mechanisms is different for different socio-economic groups. Participation to financial markets is related to education. Therefore the more educated are more likely to use financial assets to buffer income shocks. On the other hand, taxes and other government program, such as unemployment insurance schemes, Medicaid or food stamps, are more likely to be effective for the less well-off. We therefore document the dynamic of wage, income and consumption inequality for various groups, which allows us to discuss the empirical relevance of the various channels.

3. Measurement issues

This section discusses:

- **Why most studies on inequality use income data.** Compared to income, consumption data over a long time period are scanty. Surveys designed to measure income have been around for a long time. The importance of consumption data for the measurement of individual well-being suggests using the Consumer Expenditure Survey (CEX). This survey provides data on a consistent basis from 1980 and is the only survey on the US household population to have information on consumption for this long period of time.
- **Why the CEX is useful to the purpose of studying the evolution of well-being in the US.** Other than the availability for a long time span, there are at least two more reasons to focus on the CEX. First, the CEX is made of two survey instruments (recall questions and diaries) and their joint use can help improving the quality of our measures of consumption. Second, the CEX provides data on stock of some durables and such information can be used to impute the flow of services from those durables.
- **The main features of our samples and the adjustments needed to use the data from the CEX as basis to measure US households' welfare.** The CEX is representative of the US household population. In our sample we focus on urban households, on private and public employee and on those aged between 25 and 65. These selection criteria are standard, are driven by the survey design, and make our results comparable with what is found in other studies on income data. Consumption measures are adjusted to take into account the dynamics of family needs and that of prices.
- **Although the data we use are probably the best available and very useful in many dimensions, they are not without problems.** We discuss some important limitations of the CEX data, ranging from the quality of the

data to the size of the sample. We also discuss briefly other problems that affect our ability to analyze the evolution of well-being, which are common to the analysis of income and consumption (such as problems with the measurement of inflation).

In several instances in the previous section, we referred to important measurement issues. We mentioned that consumption data have been rarely used for the analysis of living standards and inequality⁶. We also mentioned the relative quality of income and consumption data and the difficulties in obtaining *consumption* information from *expenditure* data. One of the reasons for the prevalence of income surveys in the study of inequality, is that surveys such as the CPS, the PSID, the NLSY have been around for a long time and have become the bread and butter of labour economists. On the other hand, the CEX exists in its current (and consistent over time) format only since 1980. Earlier versions (most notably the 1960-1961 and the 1972-1973) were collected infrequently and with different methodologies. This situation generated a resistance on the part of many economists, until recently, to use the CEX, although it is the only survey that contains comprehensive information about consumption expenditure.

Like all surveys, the CEX is not exempt from problems, some of which we discuss at length in what follows. However, its quality is not worse than many of the standard surveys routinely used by economists and policy makers. It is also a fairly comprehensive survey, containing not only information about consumption expenditure but also on a variety of other variables. Partly to assert the credentials of the CEX as a legitimate and high quality survey, in the next section we start our analysis of wage inequality reporting figures both from the CEX and the more commonly used Current Population Survey (CPS), to show that the patterns that emerge from the CEX sample are remarkably similar to those of the CPS.

In this section we start by discussing the structure of the CEX survey. We then address the issue of the quality of the consumption data by discussing its comparison with aggregate data sources. The CEX is made of two separate components and samples. In the next subsection we discuss how we combine them to estimate inequality in expenditure. In section 3.4, we discuss how we imputed values for the stock of some

⁶ The most important exceptions are Cutler and Katz (1991), Attanasio and Davis (1996), Slesnick (1993, 2000, 2001), Attanasio, Battistin and Ichimura (2004), Meyer and Sullivan (2004). We elaborate on some of these studies in the 'further readings' sub-section at the end of this section.

durables held by the households in the survey. In section 3.5 we describe how we deal with changes in prices, the issue of equivalence scales and how we deal with it and how we select the samples we will be analyzing.

3.1. Data sources: the CEX and its two components.

The CEX started in its modern incarnation in 1980, but has a long history that goes back to the beginning of the 20th century. It is managed by the Bureau for Labor Statistics and is collected by the Census Office. The main motivation for the collection of the CEX is to compute weights for the Consumer Price Index (CPI). Although the weights are only updated infrequently, it was deemed necessary, as in other countries, to collect information on the expenditure patterns of US households on a continuous and homogeneous basis.

Since 1980 the CEX is constituted of two independent components and samples. The first and larger sample, the Interview survey, is a rotating panel of about 5000 households per quarter (which increased by around 30% in 1999), who are interviewed five times at a quarterly frequency. Each quarter, 25% of the sample is refreshed. The first interview is a contact one, and effectively no much information is collected (and no information is available in the public domain on this interview). In each of the following interviews, a respondent for each household is asked detailed questions on the amount spent in each of the three months preceding the interview on many expenditure items. While for some component, the questions are quite detailed, for others (most notably food) they are very much aggregated. Given the sample structure, if a household completes its cycle of interviews, the CEX will contain 12 monthly observations on the expenditures of such household. In addition to expenditure, much more information is collected. There is complete information on the demographic composition of the household and on many socio-economic characteristics of its members. Information is collected on education levels and activities of each household members. In the second and fifth interview, a substantial amount of information is collected on earnings, labour supply and so on. Finally, in the fifth interview, some information about financial and other types of assets is also collected. Since 1988, the BLS has started producing special modules which contain rich information on a variety of issues. There is for instance a module on credit cards (containing information on the balances of each card held by the family) a module on mortgages and real estate, a module on health expenditure and so

on. In what follows we have used the module on vehicles, which was actually started in 1984.

The sampling frame is renewed every 10 years, in the years ending in 6, to start reflecting the weights of the last census. For instance, in 1986, the weights of the 1980 census were adopted in constructing the 1986 sample. This implies that the rotating panel feature of the survey is lost in those years. The structure of the questionnaire has been remarkably constant and consistent over time. Very few questions have changed⁷, and only occasionally some items were made more detailed (for instance in 1991, the expenditure on Personal Computer was divided between software and hardware).

The second component of the CEX, the Diary survey, is made of sample of 5000 household that is refreshed every year. Respondents are asked to fill in a diary for two consecutive weeks. There is no longitudinal dimension to this sample. The Diary and Interview samples are completely independent and separate. Until 1986, the respondents in the Diary survey were asked to include in their diaries only entries regarding food items and other frequently purchased commodities (such as toothpaste or other personal care items). After 1986, instead, the Diary survey becomes comprehensive, so that, in principle one has two different measures for most commodities.⁸

The BLS, however, thinks that some items are measured appropriately by one survey and others by another. This belief is reflected in practice in the methodology to compute the tables that are routinely published and, ultimately, the weights for the CPI. That is, information from the Diary survey is used for some items and from the Interview survey for others. In constructing our evidence, we will follow this practice too.

3.2. The quality of the data

Collecting information on expenditure in a developed country can be a daunting task. A typical household is characterized by very complex and rich expenditure patterns, the spending is done by several household members, and many of the items are not particularly salient and might be difficult to remember in the context of a retrospective interview. This very difficulty is, essentially, what motivates the existence of two different

⁷ One important exception is the question for food at home in the Interview survey, which underwent significant changes in 1982 and then again in 1987.

⁸ The level of aggregation is different across the two samples, however. Food, for instance, is extremely detailed in the Diary survey, while it is only available as food at home and food away from home in the Interview survey.

samples and data collection methodologies. For all these reasons, one would expect a substantial amount of measurement error.

One problem with verifying the quality of the CEX is the lack of a benchmark. Traditionally, the term of comparison has been aggregate Personal Expenditure data published quarterly by the Bureau of Economic Analysis within the National Income and Product Accounts. This comparison, however, is not without problems. First, the definition of expenditures used in the NIPA and in the CEX are not the same. Some items stand out: housing includes imputed rents for home-owners in the NIPA while it does not in the CEX, the purchase of second hand cars from other households is netted out in the NIPA but it is not in the CEX. And even when the differences are not as large as in the examples, for many items there are important conceptual differences. Second, the population of reference is different. The CEX, when aggregated properly only includes non-institutionalized households, while the NIPA include the consumption of institutionalized individuals and the consumption done on behalf of households by various institutions. Finally, the NIPA Personal Expenditure figures are obtained, by and large, as a residual, starting from sales figures and removing amounts that are *not* bought by households. As such, they are likely to be subject to large measurement error, as witnessed by the many revisions, some of which sizeable, through which the accounts undergo.

For better or worse, however, the NIPA have traditionally constituted the benchmark against which the CEX, appropriately aggregated, has been evaluated. The BLS itself, routinely compares the CEX aggregates with NIPA data. The two main facts that emerge from these comparisons are that the CEX aggregates are substantially below the NIPA and that this ratio has been deteriorating in the last five or six years. The most recent BLS publication (see Garner et al., 2006) puts at 0.65 the ratio of total consumption expenditure as estimated in the CEX to the corresponding PCE aggregate in 1997. There is large variation in sub-categories, with the ratio varying from 0.19 for 'Sewing goods' to 5.11 for 'railways transportation'. The same paper, considering variation over time, shows that the ratio for the total goes from 0.67 in 1992, to 0.65 in 1997 to 0.60 in 2002. The same ratios for non-durable consumption are 0.65, 0.63 and 0.58. Going back to the 1980s, these ratios do not vary much, but are roughly at the level of 1992 (see Gieseeman, 1987).

The evidence would therefore seem to indicate a substantial deterioration of the quality of the CEX, at least as measured by its correspondence to the NIPA data. Garner et al. (2006), however, argue forcefully that this is not necessarily the case. They analyze relatively fine classifications of consumption and determine when the categories are conceptually comparable between the CEX and the NIPA data, the ratio is much closer to unity and there is much more modest deterioration over time. This characterization is not without exception but, by and large, holds. The largest deviations from a unity ratio are observed when there are definitional differences. While this evidence is obviously important and gives hope about the quality of the CEX data, the changes in the ratios over time, for some or most categories remains a mystery. Two possible hypotheses, not necessarily alternative are: (i) the importance of the sectors excluded by the CEX but included in NIPA has increased over time; (ii) certain segments of the population, in particular at the top of the income distribution, have become less willing to collaborate with surveyors. The increased difficulty in contacting well-off households for economic surveys is a fact that has been observed in a variety of surveys.

Although the arguments advanced in Garner et al. (2006) are surely important, it remains the fact that according to the CEX, over time, per capita consumption expenditure seems to have declined constantly since the early 1990s. Such rates of decline are problematic and should be investigated more. We will come back to this issue in Section 5.

3.3. Our samples: sample selection, inflation corrections and household size corrections

We conclude this section with information on the way we select the samples from the CEX and with the way we deal with several issues, like inflation and equivalence scales. As we will use the CPS to make comparisons with the CEX, we will select the CPS sample in a way that mirrors exactly the criteria for the CEX.

Sample selection. We select our CEX sample to include only urban households. This was partly done to be able to use the 1982-1983 data which did not cover rural households, and partly because of difficulties with the rural data. We did not use the 1980 and 1981 data because of concerns of comparability with the subsequent years: those years were the first two of the survey and many variables were changed in the 1982

vintage of the survey. Our last year of analysis is 2003. We also excluded households headed by self-employed individuals. For these households it can be hard to isolate expenditures for the household from expenditures for the business. Moreover, income is notoriously measured with difficulty for self-employed individuals. We are aware of the fact that this is an important limitation, as self-employment status can be a reaction to specific income shocks and could therefore change over the business cycle.

Less controversially, we exclude households with ‘incomplete income responses’ from the analysis of income as well as the analysis of consumption data. The incomplete income response variable is generated by the BLS to flag households whose income data are of poor quality. We thus take a conservative approach and assume that household with poor quality income response also report poor quality consumption expenditure. The share of households with incomplete income responses is roughly constant over time, ranging between 14.7 in 1982 to 15.1 percent in 2003 in the Interview Survey and from 22.6 to 27.4 in the Diary Survey.⁹ Finally, when analysing wages, we use individual rather than household level data, and focus on males in full time employment.

Some of the analysis will look at year of birth cohorts. In defining cohorts, one faces the trade-off between cell sizes and homogeneity. We decided to form four cohorts defined by decades: our first cohort groups households whose head was born in the 1930s, while our fourth groups households whose head was born in the 1960s.

Inflation. We measure inflation through the general Consumer Price Index for all Urban Consumers produced by the BLS at a monthly frequency. The base years are 1982-1984.¹⁰ The CPI is available on a continuous basis for our sample period and is widely used for a variety of purposes, including the indexing of Social Security Benefits and several other social programs in the US. The construction of an appropriate Price index is notoriously difficult and fret with many conceptual and methodological problems. Due to the CPI importance for policy purposes, the ability of the CPI to reflect increases in the cost of living (rather than other factors, such as increases in quality) has been recently examined closely in a variety of studies. The Boskin Commission undertook this task in

⁹ Battistin and Padula (2009) show how many households are excluded from the Diary and the Interview sample if one drops those with incomplete income responses, non-urban households, those aged less than 25 and more than 65 and the self employed head households.

¹⁰ The price index underlying the CPI is used by many statistical agencies in the US and around the world.

1998.¹¹ The Boskin Commission argued that the CPI fails to properly account for product substitution and quality change and concluded that the CPI is upward biased by about 1.1 percent per year in 1995-1996. Gordon (2006) in a later evaluation of the Commission results suggests that the bias for the years 1995-1996 is 1.2-1.3 percent per year, and that it is currently about 0.8.

Assessing the bias in the CPI requires being able to construct a ‘true’ cost-of-living index. Broda and Weinstein (2008) pursue this approach and use barcode data to provide estimates of the bias in the CPI. Their work shows that quality bias causes the CPI to overstate inflation by 0.8 percent a year between 1994-2003. The Boskin Commission report and the subsequent studies that looked at this issue show the importance of improving on the current CPI, but do not provide an index on a continuous basis for our sample period. Therefore, we stick to the use of the CPI. However, it should be kept in mind, when interpreting the results referring to trends in the *level* of consumption or income, that overestimating inflation will underestimate the rate of growth of these variables.¹² Notice however that biases in the CPI, do not affect the results for inequality, insofar the CPI bias does not vary with individuals or households characteristics.

Equivalence scales. In a famous quote, the economist Gorman summarized the importance of equivalence scale by saying: “if you have a wife and baby, a one-penny bun costs three pennies”. When evaluating the evolution of material well being and its distribution, one needs to take into account the evolution of needs. Expenditure is recorded at the household level. The size and composition of American households has changed considerably over the period analyzed and it has changed differentially for different groups of the population. It is therefore important to control for household needs. We do so by using a very simple equivalent scheme: we count as 1 the first adult, as 0.7 any additional adult and as 0.5 any child in the family.¹³ Figure 3 shows the time

¹¹ A first evaluation of the CPI was carried by the Stigler Commission in 1961.

¹² A discussion of the biases of the CPI in this context is contained in Slesnick (2000, 2001).

¹³ The family is the consumer unit. A consumer unit comprises either: (1) all members of a particular household who are related by blood, marriage, adoption, or other legal arrangements; (2) a person living alone or sharing a household with others or living as a roomer in a private home or lodging house or in permanent living quarters in a hotel or motel, but who is financially independent; or (3) two or more persons living together who use their income to make joint expenditures. Financial independence is determined by the three major expense categories: housing, food, and other living expenses. To be considered financially independent, at least two of the three major expense categories have to be provided entirely or in part by the respondent.

pattern of such equivalence scale and documents the changing structure of the US families in the last two decades. More sophisticated scales are possible, but do not affect the thrust of our results much. Figure 3 implies that using equivalence scales prevents one to conflate trends in material well being with changes over time in family size.

Further readings.

Slesnick (2000, 2001) contains detailed discussions of a number of issues we touched in this section. In particular, he discusses the problem of inflation and its measurement, as well as the problem of equivalence scales. (2000, 2001) also discusses the quality of the CEX data and their comparability with the National Account data. The forthcoming entry in the Palgrave dictionary on equivalence scales, authored by Arthur Lewbel, is also a useful reference on this issue. The problem of the comparability of CEX and national account data is discussed periodically in various BLS publications. Garner et al. (2006) is the latest.

3.4. Combining the Diary and Interview surveys

If one was interested in the evolution of average consumption, the existence of two separate components of the CEX does not constitute a problem. As the two surveys are both representative of the same population, one can rely on the Interview survey for some components of consumption and on the Diary survey for others and use the simple fact that the average of a sum is equal to the sum of the averages to compute the average of total consumption.

The one issue that can generate problems is differential attrition and non response, for which there is some evidence. The two samples are indeed slightly different, with the Diary survey being made of households that are typically better off and better educated. These factors can be, to a certain extent, taken into account as long as differences are confined to household characteristics that are observable in the two surveys. The same considerations apply if one is interested in the average consumption for subgroups of the population, such as households headed by a high school graduate or a college graduate, or single households, or households headed by an Afro-American.

The issue is very different if one is interested in the dispersion of overall consumption. In this case, the fact that we have some items well measured in a survey and other items well measured in the other may constitute a problem. Suppose, for

instance, that we are interested in the difference between the 10th and 90th percentile in total consumption and suppose we divide total consumption in the items measured in the Diary survey (d-goods, henceforth) and those measured in the Interview survey (i-goods). We can compute the 10th and 90th percentile for the i-goods and d-goods separately, but the 10th and 90th percentiles will not be given by the sum of the respective percentiles because the family on the 10th (90th) percentile for the i-goods is not the same as the family on the 10th (90th) percentile for the d-goods. Similar considerations apply if we want to compute the standard deviation of all commodities.

In section 3.1 we mentioned that, at least since 1986, public available data from both surveys contain an (almost) comprehensive list of consumption commodities. The temptation, therefore, would be to ignore the BLS practice to rely on the Diary survey for the frequently purchased items and on the Interview survey for the others and use just one survey. This is the solution that has been followed in most of the literature, which has ignored the Diary survey and used the larger Interview survey. This practice would be fine if one could show that the picture that emerges from the other survey is roughly consistent. Unfortunately, Battistin (2003) and Attanasio, Battistin and Ichimura (2007) (ABI, henceforth) showed that the pattern of expenditure inequality, as measured by the coefficient of variation of consumption, is wildly different in the two surveys. If one takes the inequality of total consumption from the Interview survey, one obtains that, since the late 1980s or early 1990s, inequality has not changed much. On the other hand, if one uses the Diary survey, one sees that the dispersion of inequality has increased dramatically. Figure 1, reproduced from ABI, shows this dramatic differences. It should be stressed that the fact that inequality as measured in the Diary survey is much larger than that measured in the Interview survey is not surprising, given that the former only covers two weeks of expenditure and the latter one month.¹⁴ What is puzzling is the different dynamics of the two series.

ABI show that this disparity does not have simple explanations, such as differences in the composition of the two samples or an increase in the infrequency of purchase in the Diary survey. For the latter hypothesis the argument is that if people shop less frequently, a larger number of households will report zeros in the diary for a large

¹⁴ As an extreme example, consider the case of two households both spending for public transport 200\$ in a month, and suppose that the expenditure for one household is concentrated in the first week, and for the other in the third week of the month. The variance of expenditure for public transport at monthly level is zero using interview data, where households are asked how much they have spent for public transport in a month. However, making both households to fill a diary for the first two weeks of the month would give a positive variance for the diary data.

number of commodities – this would artificially increase the variance of consumption. ABI dismiss this hypothesis by observing that the proportion of zeros in the Diary survey does not increase significantly over time.

In the end, ABI propose a methodology to combine the two data sources, which we adopt here (see also Battistin, 2003). The main idea is to follow the recommendation of the BLS and use information on frequently purchased items from the Diary survey and information on other commodities from the Interview survey. To compute the variance of total consumption we will then need the covariance between d-goods and i-goods. ABI show how to use the error-ridden information on d-goods in the Interview survey and/or the error-ridden information on i-goods in the diary survey to approximate this covariance. In what follows we will use the same methodology and we refer to the ABI paper for details. In that paper, ABI present some evidence in favour of the assumptions made to apply these methods. The validity of such assumptions is further investigated by Battistin and Padula (2009).

One consequence of using the two surveys and the ABI methodology is that to estimate the evolution of consumption dispersion we need to make some hypotheses. If one limits oneself to the study of the standard deviation of consumption, one can get away with a limited number of assumptions. Effectively all we need is to be able to estimate the covariance between i-goods and the d-goods using the imperfect information we have in the two surveys. If, instead, one is interested in the entire distribution (as it would be the case in which one chooses as an inequality index the ratio of two percentiles, such as the 90th and the 10th), then one would need additional assumptions on the nature of the distribution. For instance, if one is willing to assume log-normality of the distribution of consumption at a point in time, one can recover all the percentiles of the consumption distribution. These percentiles, however, will be necessarily a function of the mean and standard deviation that we recover with the basic sets of assumption. For this reason, in what follow, we use as our measure of dispersion, the standard deviation of log consumption. Such a measure is a proportional measure and as such does not depend on the scale of consumption and has been used widely in the literature. Battistin and Padula (2009) discuss a set of assumptions that generalize the approach taken by ABI and can be used to study a variety of inequality measures.

3.5. Estimating the service flow from durables

Consumption is commonly measured through expenditures both in aggregate and in micro-data. Assuming that consumption coincides with expenditure suits reasonably well the case of non-durable commodities. However, it is obviously a very poor approximation for durables. The distinction between expenditure and consumption is not trivial in this case, since durables are typically bought infrequently and provide utility for many periods.

In order to measure durable consumption, one should be able to quantify the flow of services that household enjoy. This requires estimating the value of the stock of durables, since the flow of services is likely to be proportional to it. In what follows, we estimate the services a major durable good: cars. The choice of focussing on cars is grounded on three arguments. First, cars are arguably the most important component of durable expenditure, after housing. Second, the available information makes it possible to estimate the value of the stock of cars, but not the stock of other smaller durables. Third, houses were not included because they are also an investment that provides a return. It is difficult to distinguish between housing services and the return to houses seen as an asset.

Although we only estimate the services of cars, in what follows we will be referring to the sum of expenditure on non durable and services and our estimates of the flow of services from cars as ‘total consumption’. This is obviously an abuse of language that we only justify for the sake of brevity.

To quantify cars services we estimate the value of the stock of vehicles owned by each household in the sample. This exercise is feasible because, at least since 1984, the CEX provides additional information on cars beyond expenditure. For instance, the data record the vintage, the make, the model, the year and the price of purchase for cars, plus a number of cars characteristics. To give a first sense of the way we use these data, in Figure 2, we use these data to estimate the path over time of inequality in the service flow from durables.

The remainder of this section therefore discusses how the stocks cars have been estimated from the CEX data, by combining expenditure information with the additional information on stocks.

The data on cars come from two files. The BLS started to make publicly available these files since 1984. The first file (OVB), which refers to the vehicles owned by the

household, records a full set of characteristic for the vehicles present in the Consumption Unit (CU) at the interview date. An incomplete list of these characteristics includes the type of the vehicle (car, truck, van, pick-up, motor-bike, boat and, eventually, airplane), the make and the model of the vehicle, the year and the month of purchase, the vintage, the number of cylinders, if the vehicle entered the CU as new or used, if the vehicle is equipped with the air-conditioning, the automatic transmission, the power-brake, the power-steer, the radio, the sun-roof. The list also includes the purchasing price which contains two components. The net purchasing price, which turns out to be the cash outflow at the date of the purchase, and, if any, the trade-in allowance received.

Moreover, households are asked if they disposed of a vehicle and, in case they did, they are asked when, how and a set of characteristics identifying the vehicle. These information are recorded in the second file (OVC) which refers to the vehicles disposed by the households. Households can dispose their vehicles in six ways vehicles can be sold, traded in, given away outside the CU, damaged beyond repair, stolen and other.

To estimate the value of cars we look at their second hand market price, which we elicit from the price the CU receives for selling the car and the price the CU paid at time of purchase.

We then identify a single numerical index measuring the car value, which is know to depend on a number of features, including the car's year of production, its age and the general level of prices. This is done by regressing the log of the purchase (or selling) price of a car on year and age dummies and on a set of controls proxing for the car vintage, and then predicting the index measuring car value out of this equation. Each car in the dataset is identified by its make and model, the year of production, the year of purchase (or sell). More details on the estimation of cars value are provided in the Appendix.

4. Recent trends on wages and household income inequality.

This section shows that:

- **The wage and the earnings data from the CEX and the Current Population Survey (CPS) provide comparable information on the evolution of income inequality.** The CPS has been widely used in the study of the evolution income inequality. It lacks, however, information on consumption and therefore prevents the

joint examination of consumption and income. The CEX, in addition to the information on expenditure, contains information on wages and earnings. It is therefore important to validate the CEX wages and earnings data against the CPS. The mean and the standard deviation of the wages and earnings distribution for different groups in the population and their evolution over time are remarkably similar across the two surveys. Therefore the two surveys provide consistent evidence on the evolution of income inequality.

- **Averages of wages and earnings have increased in the US between 1982 and 2003.** Real average wages increased in the US by about 10 percent between 1982 and 2003. Such an increase would be larger if one would account for the fact that the available price indexes underestimate inflation. Real average wages decrease between 1982 and 1992 by 10 percent and then increase by 20 percent. Average earnings increase between 1982 and 2003 by about 21 percent and take-off in mid 90s and increase by 17 percent from 1993 to 2003. Again, the actual increases are underestimated (and decreases overestimated) if the available price indexes overestimate inflation.
- **Inequality, as measured by wage or earnings discrepancies, has increased in the US between 1982 and 2003.** We use the difference between the 90th and the 10th percentile of the log wages distribution as measure of wage inequality and that between the 90th and the 10th percentile of the log earnings distribution as measure of earnings inequality. As a further measure of inequality, we also look at the coefficient of variation, i.e. the ratio between the standard deviation and the mean. Whatever measure and income definition is used, the data show that inequality increases between 1982 and 2003. Such increase is steady and is more pronounced for wages (15 percent) than for earnings (10 percent).

In this section, we report data on the evolution of inequality in individual wages and household earnings. These are the variables that have been routinely analyzed in most studies in inequality and will provide the background against which we set the data on consumption inequality we present in the next section. This section also serves another role, as for wages and earnings we can use two different surveys that contain the relevant information. The first is the Current Population Survey (CPS) that is used routinely by

labour economists to analyze inequality. The second is the Consumer Expenditure Survey (CEX), which is the only survey containing detailed information on consumption and expenditure.

We start this section with some simple pictures that illustrate the main facts of the evolution of inequality for wages in the US during the 1980s, the 1990s and the early 2000s. As mentioned in the introduction, these facts are very well known: the Autor et al. (2007) study we cited is only one of the last of an extensive literature that analyzes them. The novelty here is that we use data both from the CEX and the CPS. A comparison between the two data sources has partly the purpose to show that, as far as the income components are concerned, the picture that emerges from the CEX is similar to that that comes out of a widely used data set such as the CPS. In this respect, such a comparison constitutes an important validation of the sample used by the CEX. To the best of our knowledge this is the first study that compares the wage information in the CPS and the CEX. We then move on to the evolution of household income. Since the Diary survey is characterised by a much smaller sample size than the Interview survey, in what follows we will make use of information on income and wages only from the latter. It is however worth noting that time series obtained from the Diary survey proved noisier but informationally equivalent to those reported here.

4.1. Wages: CPS and CEX evidence.

In Figure 4, we plot the time series from 1982 to 2003 for the mean and median log nominal wages derived from the CEX and the CPS. These statistics are computed on the sub-samples selected according to the criteria mentioned in Section 3.5. In figure 5, we look at the median log wages for the four decade of birth cohorts we are considering: individuals born in the 1930s, 1940s, 1950s and 1960s. Again the series for CEX and CPS align very well in each graph. The only exception is the end of the sample period for the 1930s cohort: as the CEX is a smaller sample and many individuals belonging to that cohort are retired by the early 2000s, the graph for the CEX is a bit noisier. Notice how the life cycle profile for wages is steeper at the beginning of the life cycle, flattens out and declines towards the end of the life cycle. If we neglect the issue of retirement and more generally of employment these profiles represent the evolution of average wages over time.¹⁵

¹⁵ Wages can only be computed for individuals who work. If retirement (or more generally employment) is not random and uncorrelated with the level of wages, this implies that age-profile we plot does not represent an unbiased estimate of the average wages faced by an individual over her life cycle.

Figures 4 and 5 also reveal that real wages have increased by about 10 percent between 1982 and 2003. If one assumes that inflation rate is overestimated by 0.8 percent a year as discussed in Broda and Weinstein (2008) in the whole period, the actual change in real wages amounts to just below 28 percent. The increase is even larger if one uses the estimates provided in Gordon (2006), who suggests that the bias in the CPI is between 1.2-1.35 percent in the years between 1978 and 1996, and 0.8 in later years. Therefore, while the bias in inflation does not affect the comparison between CEX and CPS, it affects the quantitative conclusions on the long run trends in the market price of labour.

In Figure 6, we plot the evolution of wages for four academic achievement categories: high school drop-outs, high school graduates, some college and college graduates. Even in this dimension, the two data sets align very well. They both point out to the fact that while the real wages of individuals with lower education levels declined, especially in the first part of the sample period, those for college graduates were first stationary and then increased.

The different dynamics across education groups is made very evident in Figure 7, where we plot the difference between each mean log wages and that of the high school graduate in 1984. The left panel corresponds to the CPS, while the right hand side panel to the CEX. A part from the right graph being noisier because of the smaller sample size of the CEX, the two pictures tell effectively the same story.

Having had a brief look at the evolution of the levels of wages, we can now move to the evidence on wage inequality. A first preview on an important component of wage inequality was given in Figure 7 which makes it clear that the differences in the remuneration of different skills has played an important role in the evolution of inequality in the US. We will come back to it in what follows.

We start by reporting data on the difference between the 90th and the 10th percentile of the log wages distribution. In Figure 8, we plot this difference computed for the sample of male employees in the CPS and in the CEX. Two things are worth noting. First, the picture that emerges confirms, not surprisingly, the evidence described extensively in the literature of a marked increase in the first part of the period and a slow down to this increase in the second. Second, the CEX and the CPS seem to be telling, once again, similar stories. In measuring dispersion the larger noise induced by the

smaller sample size of the CEX is more evident than in the corresponding picture for the means and the medians.

As we mentioned in Section 3, in the case of consumption we will be looking at a different index of inequality. Rather than the difference between the 90th and the 10th percentile, we will be looking, among other things, at the coefficient of variation.¹⁶ In Figures 9, 10 and 11, we plot this statistic for wages, again for both the CEX and the CPS, both to check how the two measures (90/10 ratio and coefficient of variation) compare and to check how the two datasets compare in this different dimension. On both counts the comparison is satisfactory. The pattern over time of the 90/10 difference is very similar to that of the standard deviation of logs. And the alignment between the CPS and the CEX for the latter is even closer than for the former.

To conclude this section, we can summarize the evidence we have presented in the following fashion. In terms of data sources, the two surveys we analyze, the CEX and CPS, tell, by and large, a very similar story, especially if we look at the levels of average wages and the levels of inequality. This is heartening as the CEX is the only data set that contains information on consumption and the CPS has the advantage of being a larger and better studied survey.

A word of caution is needed if one wants to evaluate the long run trend in real wages. Both the CEX and CPS imply that real wages have increased between 1982 and the 2003. The extent of such increases is very similar across the two surveys. Moreover, we observed that the overall level of real wages has declined slightly over time until the early 1990s and that it has increased after that. Deflating the nominal wages by the CPI index, the figures show an increase of 10 percent over the sample period. However, such an increase is much larger, by about 28 percentage points, if one takes that the CPI overstates inflation by about 0.8 percent a year, as argued in Broda and Weinstein (2008). Therefore, different inflation measures have very different implications for the quantitative evaluation of the long run changes in the market price of skills. Furthermore, if the inflation rate is biased upward by 0.8 over the whole period, the decrease in real wages between 1982 and early 1990s would be very much attenuated .

¹⁶ Under log-normality, the difference between the 90th and the 10th percentile and the standard deviation of logs wages exhibit the same rate of growth over time and the coefficient of variation of wages is approximately equal to the standard deviation of logs.

Whatever measure of inflation is used, the average pattern hides differences across groups. In general we observed that real wages of the bottom of the distribution (and in particular of low educated individuals) have decreased the most in the first part of the sample period, while the level of wages for better educated (especially college educated individuals) have stayed roughly constant or increased. Inequality, consequently, has increased. Such an increase has been particularly pronounced in the first part of the sample period, but it has, to an extent, continued into the 1990s and early 2000s. The increase is to a large extent driven by an increase in differences among skill groups. However, even within skill groups, we do see an increase in various measures of inequality, such as the 90th/10th differential or the coefficient of variation.

These facts are not particularly surprising and confirm what is in the literature. We can now move on to look at the next step, where instead of wages we consider household income.

4.2. Household earnings: CPS and CEX evidence.

Our measure of household income includes earnings and excludes capital income. This choice is rooted in the difficulty of measuring capital income and wealth. Moreover, we focus on pre-tax and pre-transfer income. This is consistent with the way we treated wages and is motivated by our primary interest in the dynamic of inequality *before* any insurance mechanism operates, beyond labor supply and mating choices.

Given the wages an individual (or a family) can command, the amount of time spent working by household members will determine earnings. The evolution of the distribution of household earnings will depend, in addition to wages, on how households group people with different earning capacities and on the labour supply behavior of the various members. Both factors can change substantially the patterns observed in the distribution of wages over time. If there is positive correlation in the type of wage shocks received by husband and wives, considering households can increase the patterns observed in the data. On the other hand, people reacting to wage shocks by adjusting their labour supply behaviour could reduce or increase the changes in inequality.

Once again, in this subsection we report data for both the CEX and the CPS using the same sample selection considered in the previous section. In Figure 12 we plot mean

and median log family earnings from the two data sources. Once again, we notice the remarkable co-variation between the two surveys, in particular for the medians: for the mean the correlation coefficient is 0.87, while for the median is above 0.92. We regress the difference between the mean log earnings in two samples on a linear time trend and we cannot rule out that such difference is zero and orthogonal to the time trend. Doing the same exercise for the median delivers very similar results.

In terms of the substantial pattern of mean and median income, these pictures stand in remarkable contrast with respect to Figure 4 on wages. In the first part of the sample, household income did not decline as wages did, although after 1988 and through the early 1990s there is a reduction in real pre-tax income. Since the mid 1990s, household income has increased at the same pace, if not more than wages. Between 1993 and 2004, median real household income increases by about 17%. This figure underestimates the actual increase by 9 percentage points, if the CPI overestimates actual inflation by 0.8 percent a year, as discussed above.

In Figure 13, as for wages, we report median log pre-tax income for the four decade-of-birth cohort we considered earlier. The correspondence between CPS and CEX is, once again, quite good, again with some exceptions for the 1930s cohort in the last part of the period. The evidence that emerges from this graph is that of a strongly increasing real household income for the two youngest cohorts, of a flatter one for the third, and a declining one for the fourth one. Of course, these cohorts are at different points in their life cycle, so that differences should be at least in part interpreted as effects of age.

In Figure 14, we plot the median log pre-tax earnings for the four education groups. The evidence shows a decrease in earnings for households headed by a high-school drop-out, a roughly constant income for households headed by a high school graduate and an increase for households headed by an individual with some years of post-high school education or with a college degree. The increase for the college graduates is particularly strong: over the period considered, median pre-tax household earnings increase by almost 35%. Notice also that for the high school drop-outs the decline observed stops before the corresponding decline in wages. It is clear that, while there are similarities between the story told by wages and that told by household earnings, there are also important differences. These differences might indicate, in all likelihood, a reaction of

labour supply to the changes in wages over this period. The overall picture that emerges, however, is that of an increased difference between the income received by low skill individuals and those received by high skilled individuals.

As with the analysis of wages, after looking at the evolution over time of *levels*, both in the aggregate and for subgroups of the population, we now move on to analyze the overall inequality and the inequality within the same groups. We start, in Figure 15, by looking at coefficient of variation of earnings in the population. In the CPS, inequality increases steadily over the sample period, while the increase in the CEX is less pronounced. However, the correlation coefficient between the times series obtained from the CPS and the CEX is 0.81. It is worth noting that the inequality level in the early 1980s is higher in the CEX than in the CPS, and this results in a more modest increase in this decade using CEX data compared to the significant increase from the CPS. Moreover, note, that inequality in the both the CPS and CEX appears to slow down for the period 1988-1990, but the decrease is larger in the CEX.

In figure 16 we plot the dynamics of the coefficient of variation of earnings for the four cohorts we have been analyzing. Perhaps not surprisingly, the CEX data result noisier, and this is most likely to depend on differences in the sample size of the two surveys. However, both CEX and CPS figures show an increasing pattern for all four cohorts and the correlation between the two series ranges from 0.69 for the youngest cohort to 0.81 for the second youngest. Finally, in Figure 17, we plot the coefficient of variation for the four education groups. Again, the CEX data are noisier, but, consistently with the CPS, show an increasing trend in inequality for all education groups. The increase in inequality, both in the CEX and in the CPS, is more pronounced for high school dropouts. This contrasts with what happens with wages, where we see a sharper increase in inequality among the college graduate and is possibly due to differences in labour supply behaviour between the two education groups as well as to differences in the synchronization of wage shocks within the household. While the focus on full-time employees should minimize the impact of different supply behaviour at least along the extensive margin, one cannot rule out that the education related differences in female labour participation are behind these results.

In order to understand to what extent individuals with different earning abilities and labour supply behaviours are grouped within the same household, we devote the rest of this section to investigate the degree of association between earnings within the household. This is done by computing the covariance between the salaries within the two earners households. In both the CEX and the CPS, households with more than two bread winners makes it up to 1.5% of the sample and including those households does not affect the inequality dynamics, but increases the level of overall inequality.

To ascertain the contribution to overall inequality of the covariance between spouses earnings, notice that in the population of two earners households the squared coefficient of variation can be decomposed in three components: the variance of earnings for first household earner, the variance for the second and twice the covariance between the two, all divided by the squared mean households earnings.

In Figure 18 we show the time evolution of each of these three components in the CPS. The left scale of the figure refers to the first two components, the right to the third. It is apparent from the figure that the component that mostly contributed to the increase of inequality is the covariance between earnings, which, as a fraction of squared mean households, has increased from just below 1 percent in 1982 to almost 4 percent at the end of our sample.

We also decompose overall earnings inequality for the four decade-of-birth cohorts and education groups. Figure 19 displays the dynamics of the three components of the overall inequality and makes it apparent that the covariance component increases for all cohorts except for those born in 1950-1959. Interestingly, for this cohort, the increase in overall inequality is less pronounced.

Since education is an important determinant of labour supply behaviour and might therefore affect the degree of co-variation between earnings within the households, we also plot in Figure 20 the decomposition of overall inequality for the four education groups. It should be noticed that the groups are formed on the basis of the education of the household head. The figure reveals that the steepness of the covariance line decreases with education. This implies that the covariance of the more educated has increased less from 1982 to 2003.

Further readings. The literature on the evolution of wages and earnings inequality in the US (as well as other countries) is too large to be summarized here. The recent paper by Autor et al. (2007) contains many relevant references as well as some of the most recent estimates. Gottshalk and Moffitt (1994) pays particular attention to the changes in the variance of transitory and permanent components of income.

5. Expenditure and Consumption.

This section focuses on:

- **The evolution of consumption inequality across US households.** To be able to combine the information from the components that make the CEX, we use the standard deviation of logs as measure of inequality. According to our evidence, consumption inequality increases by about 7 percentage points between 1982 to 2003. The increase is more pronounced for non-durable than for total consumption, which, however, is observed only from 1984 and increases by just above 5 percent.
- **Consumption inequality between skills.** We measure skills by the education achievement of the household head. The consumption of college graduates increases vis-à-vis that of the high school graduates in 1984 between 1982 and 2003 by 12 percentage points. High school dropouts perform very badly: their consumption goes down not only in absolute terms but also with respect to the other groups as well. In comparison with high-school graduates in 1984, the consumption of high school drop-outs decreases by as much as 20 percent.
- **Consumption inequality among households sharing similar skills.** The within-group consumption inequality increases for all education groups, but the increase is more pronounced for high-school dropouts. For high school dropouts we see an increase of 15-20 percentage points, for high school graduates of about 10 percentage points and for college graduates of only 5% (or almost flat if we look at total consumption).

Having analysed the pattern of wages and income inequality and having validated the use of the CEX, we can now move to the analysis of consumption levels and

consumption inequality, which constitutes the core of this monograph. As with wages and income, we start our discussion by presenting information on the *levels* of expenditure in the aggregate and for some subgroups of the population. We then move to information about the dispersion of expenditure, both in the aggregate and within the same groups. This analysis of inequality complements the analysis of differences between group averages.

As mentioned several times, the CPS does not contain information on consumption, so we will be confined to the CEX. However, the CEX, as explained in Section 3, is made of two different sub-samples: the Interview Survey (IS) and the Diary Survey (DS). Unlike most other papers in this literature, which have focussed exclusively on the IS, in this study we will combine the two using the assumption, made by the BLS, that some commodities are better measured in the IS and others in the DS. To show the extent to which this makes a difference, we start, in Figure 21 to report the average monthly consumption (per adult equivalent) for each year from 1982 to 2003. In the graph there are three lines: the short-dashed line is computed using data from the Interview Survey, the long-dashed one using data from the Diary Survey, and the continuous line uses our methodology to combine the two data sources, following the BLS practice of using the IS for some commodities and the DS for others. First notice that the DS did not include ‘non frequently purchased’ commodities and services prior to 1986. This implies that for the first few years, the DS aggregate, calculated on what is available, is much lower than the other two aggregates. After 1986, instead, the DS becomes as comprehensive as the IS (if not more). The second thing to notice is that our estimates that combine the two surveys are, for most years, above either of the two components. Finally, the pattern that emerges over time is of an increase in per-adult equivalent consumption until the beginning of the end of the 1980s. After 1990s, however, the aggregate average drops slightly and consistently overtime, especially in the period 1990-1996. After 1996 the combined estimates increases slightly. The overall pattern described is consistent with that from figures published by the BLS for the period covered by our analysis.

Unfortunately, the evolution over time of the aggregate CEX data does not match well with that of the Personal Consumption Expenditure from the NIPA accounts. The level of Personal Consumption Expenditure as estimated from the CEX figures is considerably lower than the corresponding National Account figures. During the 1980s

and early 1990s, the ratio between the CEX aggregate and the NIPA data was stable at about 65%. After the early 1990s, however, the ratio has declined considerably and is now around 50%.

The comparison between Household Surveys and NIPA data is not straightforward. There are many issues, ranging from what is defined as expenditure and consumption to the population of reference (See for instance Slensnick, 1998 and Garner et al. 2006). It is also true, that the CEX is not the only survey whose perceived quality has deteriorated over time. There seems to be an increasing unwillingness to answer survey questions among the American public. However, the discrepancy between the CEX aggregate data and the NIPA data is disconcerting and worrying. While it is true that, as pointed out in a recent publication from the BLS (see Garner et al (2006)), for expenditure categories that are *indeed* comparable, the ratio of the aggregated CEX data to the PCE data ranges from 0.88 in 1992 and 1997 to 0.84 in 2002, and is relatively stable, these categories account for a small fraction of total expenditure.

Given the evidence from the NIPA data, one cannot take the decline in the CEX aggregate in real per capita consumption expenditure at face value.¹⁷ This figure reflects more on the necessity of improving the quality of the survey and its ability to reflect aggregate averages than about the evolution of economic well being. We will go back to this issue in the conclusions, when we discuss about data needs.

In Figure 22, we start looking at the evolution of consumption levels for different groups in the population. In this and in the subsequent figures we only report the data that combine the DS and the IS. Moreover, rather than reporting means of levels we report means of logs.¹⁸ In the case of the youngest cohort, the consumption of non durables and services appears to be declining throughout the period, in accordance with what we have shown for the aggregate figure. This might be surprising as these individuals go over a period of their life cycle during which they experience, on average, an increase in their incomes. However, increasing the inflation rate by 0.8 percent a year to account for the upward bias in the inflation rate eliminates the decrease in real consumption. The consumption of the second and third cohort does show an increase

¹⁷ On top of these issues, as we mentioned above, if one wants to assess the evolution of economic well-being, one also needs to address the issue of what deflator to use to express consumption in real terms. As we discuss above, the CPI might be overestimating inflation.

¹⁸ As the log is a non linear function, to combine the two data sets is not as easy as in the case of the means of levels. We need an assumption about the distribution of total consumption in the cross section. In the Technical appendix, we show how this procedure works.

through most or part of the sample period, while that of the last cohort is more noisy, especially in the last few years of the sample. Accounting for the upward bias of the CPI would reinforce such an increase.

In Figure 23, we look at the average log consumption of non durable and services for the four education groups we considered previously. The consumption of all four groups, by and large, decreases over time (as the aggregate one) after the early 1990s. Such a decline, however, is much more pronounced for the less educated households. Indeed, college graduates are the only ones who experience an increase in the first part of the sample. The quality bias in the CPI might account for the general decrease of real consumption, but not explain the differences between groups.¹⁹ To stress the relative performance of these groups, and therefore have a first direct look at inequality, in Figure 24 we plot the consumption of non-durable and service of each of education group relative to the consumption of the high school graduates in the relevant year. This picture is comparable to similar pictures we constructed for wages and income. It shows an increase in the first part of the sample of the consumption of college graduates (and to an extent of households headed by an individual with some college) relative to high school graduates. Those who perform badly are the high school dropouts, whose consumption goes down not only in absolute terms but with respect to the other groups as well.

In Section 2, we argued that the consumption of non-durable and services gives only a partial picture of the well being of individual households, as it excludes the utility derived from durable commodities. The IS contains very detailed information on vehicles, which allows us to estimate the value of the stock of cars for each household in our sample. However, such information is available on a continuous basis only after 1983. In Figure 25, we use the methods discussed in Section 3 to add to the expenditure on non-durables and services the flow of consumption services from houses and vehicles. To give a term of comparison, in Figure 25 we also plot the average log of consumption of non durables and services, which starts one year earlier. The time path of total consumption is not too dissimilar from that of non-durable consumption in the

¹⁹ Biases in the CPI arise from quality changes and the substitution effects. It is arguable that the demand price elasticity varies with education that might affect the adoption of new products and the substitution between products. Discussing how differences in the demand elasticity across education groups translate into differences in the CPI bias is beyond the scope of this work. It seems unlikely, however, that differences in CPI biases could account for the reported changes across groups.

first part of the sample: the two variables increase until 1989, decline starting in 1990 and flatten out in 1994. However, after 1994, accounting for cars makes the series of total consumption increase slightly over time, with a total increase of about 5% by 2003.²⁰

Figures 26 and 27 mirror Figures 22 and 23 for total consumption. Again, for comparison, we also plot the figures for the consumption of non durables and services. The only cohort for whom the consideration of durables makes any difference in terms of the time path is the third: households headed by an individual born in the 1940s. For these households the increase in the second part of the sample is more pronounced when we consider total consumption than when we consider only non- durables. As for the education groups, the one picture where the time path of total consumption is slightly different from that of non-durable consumption is that for college graduates. The relative performance in terms of total consumption for our education groups is plotted in Figure 28 that mirrors Figure 24. Again, the differences between these two pictures are minor and confined to the most educated group of households.

Having analysed the main trends for the average *level* of consumption as emerging from the CEX, we now move to the analysis of inequality. The first picture we presented in this monograph was that of the puzzlingly diverging path of consumption inequality emerging from the two components of the CEX. In Figure 29, we use the methodology mentioned in Section 3 and described in the technical appendix, to compute the standard deviation of log nondurable and total consumption. As we mention in Section 3, to compute this measure of inequality we need an estimate of the covariance between the commodities measured in the DS and those measured in the IS, that can be obtained from the IS or from the DS. **ABI (2007) show that** it does not make much difference which one is used. As the estimate for the covariance in the DS is only available since 1986 (when the DS becomes an exhaustive measure of consumption), we use the measure of covariance in the IS, available since 1982.²¹

²⁰ Data limitations prevents us to focus on other durables beyond cars. If the time evolution of services from other durables, such as white and black durables, is similar to that of cars, one might argue that the exclusion of such durables reduces the observed changes in real consumption.

²¹In the appendix we compare the standard deviation of log non-durable consumption computed with the IS and the DS covariance.

If we compare the path of the series in Figure 29 to those in Figure 1, we observe that, not surprisingly, the evolution of non-durable consumption inequality as measured by combining the two data sources, is somewhere in the middle those derived from the IS and the DS: it does not increase as rapidly as inequality measured in the DS, but it is not as flat as the series computed from the IS. The increase in non-durable consumption inequality is particularly strong in the first part of the sample and seems to slow down in the second part, although there seems to be acceleration in the last few years.

As for total consumption, we can only measure its inequality starting in 1983. In the first few years the two measures are very comparable, almost identical. However, after 1985 and for most of the sample period, inequality in non-durable consumption increases faster than inequality in total consumption. This might suggest that adjustments are made in terms of composition of consumption in the face of income shocks.

In Figure 30 we plot inequality for total and non-durable consumption for our four cohorts separately. This has the advantage of following individuals that have the same age at the same time as they go through different phases of the life cycle, during which they are affected by innovations to their income. One may imagine that these shocks are different for individuals of different ages and are reflected in consumption inequality differently, as different cohorts (at different ages) have access to different mechanisms to absorb and smooth individual earnings shocks.

Households headed by individuals born in the 1950s are the only ones for whom consumption inequality (measured by total consumption or non durable consumption) does not increase considerably. For the other three cohorts, non-durable consumption inequality increases considerably throughout the period. For the youngest and oldest cohort, the same is true for inequality in total consumption. Instead for the cohort born in the 1940s, total consumption inequality does not increase much over this period. Such differences across decade of birth cohorts might be driven by differences in productivity and differences in the availability of insurance mechanism. Consumption data alone cannot help to disentangle these two sources of differences, and therefore the next section deals with the joint examination of consumption and income.

Since individuals productivity depends on education, as shown in Section 4, in Figure 31 we analyze the path of inequality in total and non durable consumption within education groups. Here the only group for which there is a difference between the two inequality measures is that of college graduates for whom total consumption inequality is

consistently lower and flatter than non durable consumption inequality. For the other groups, the two series increase in similar fashion. And the less educated groups have a considerably larger increase over time. For high school dropouts we see an increase of 15-20%, for high school graduates of about 10% and for college graduates of only 5% (or 0 if we look at total consumption). Since inequality in wages increases by 20% in this group, and by around 10% in the other groups (see Figure 11), this suggests that higher educated individuals were better able to smooth shocks to the market price of skills. The shocks are at least in part smoothed ways within the households, since the increase of earnings inequality is more pronounced for the households headed by more educated individuals (see Figure 17)

In summary, overall consumption inequality increases whether we focus on total or on non-durable consumption. The increase, however, is more pronounced for the measure of non-durable consumption, which goes from 0.51 in 1982 to 0.59 in 2003, and for the high school dropouts. Interestingly, for this group the increase in earnings inequality and in the degree of within household co-variation of earnings is also more pronounced. The next section is devoted to further investigate this issue.

6. Income and expenditure poverty: how do they differ?

This section shows that:

- **There is a positive association between consumption and income and saving and income.** Plotting median consumption against earnings and wages percentiles for different sub-groups shows that there is a positive association between consumption and both earnings and wages. Moreover, earnings and consumption diverge in the upper part of the earnings (and wages) distribution. Saving defined as earnings minus consumption increases with income.
- **The association between consumption and income weakens in the bottom part of the earnings distribution.** For low percentiles (less than the 20th) of the earnings distribution, non-durable consumption does not vary much with the level of earnings. Moreover, the level of total and even non-durable consumption is much above that of total household earnings up until the 5th percentile of the earning distribution, implying negative saving for the households in these percentiles.

- **Those who are poor in income do not need to be poor in consumption.** Individuals and households that are identified as ‘income poor’ or at the bottom of the income distribution are not necessarily the same as those identified as ‘consumption poor’ or at the bottom of the consumption distribution. 43% of households in the bottom 10 percent of the earnings distribution have consumption levels in the top 60% of the consumption distribution.

After considering the dynamics of consumption, earnings and wage distributions separately and in various dimensions, we now start to relate the various distributions. Before going into the deeper analysis presented in next section, here we ask some simple questions: are the people on the lower tail of the earnings and wage distributions the same as those on the consumption distribution? And is the relationship between the bottom of the distribution and the rest of the households the same in terms of earnings, wages and consumption? We show how the answers to these questions has changed over time, by repeating the analysis for four years: 1985, 1990, 1995 and 2000.

To show how earnings, wages and consumption are related, Figure 32 plots median consumption against earnings (left-hand-side panel) and wages percentiles (right-hand-side panel). To produce this picture (and those that follow) we merge data from several years of the CEX. Similar results would be obtained if these pictures were done by year. The figure computes median consumption by earnings and wages percentiles and shows that, not surprisingly, there is a positive association between consumption and both earnings and wages. Moreover, earnings and consumption diverge in the upper part of the earnings (and wages) distribution. This is not peculiar to US data. With UK data, Attanasio, Battistin and Leicester (2006) show that savings increase with income. Moreover, they show that this pattern does not change much over time, as seen when comparing similar graphs for four different years in the period we are analyzing.

To further explore the consumption-income gradient for the poor, we zoom in the bottom of the distribution. In Figure 33, we plot (log) median consumption for the bottom 20% of both earnings and wages distribution. As in Figure 32, in addition to the log of non durable and ‘total’ consumption we also plot the log of household earnings. The figure uncovers interesting patterns. First, the non-durable and total consumption

lines are much flatter than that of earnings. Indeed, the focus on this part of the graph (which allows us to change scale) makes it clear that in the left tail of the earning distribution there is a much weaker association between income and consumption than in the right tale: non-durable consumption does not much differ across earnings percentiles, for low percentiles. Moreover, the level of total and even non durable consumption is much above that of total household earnings up until the 5th percentile of the earning distribution. A similar picture emerges if, instead of earnings we consider total household income. This means that saving is negative for the (income) poor, suggests that consumption does not decrease with income for very low level of income.²²

In Figure 34 and 35, we have considered for each earning and wage percentile the median of consumption (and earnings). However, within the groups of households defined by earnings and wage percentiles there is a substantial amount of heterogeneity. To show this we focus on total consumption and in Figure 34 we plot not only the median but also the 25th and 75th percentile of the consumption distribution at each earning percentile. Figure 35 is the same, but, as with Figure 33, we zoom on the bottom of the earning distribution. What is striking in these pictures is the amount of heterogeneity in consumption at any given earning percentile and in particular at the bottom of the earnings distribution: the difference between the 75th and 25th percentile of the consumption distribution at low level of earnings is be as large as 100%

This lack of coherence between income and consumption in the bottom distribution of earnings (and wages) suggests that those who are income poor are not necessarily consumption poor. To shed further light on this possibility, Table 3 displays the contingency table between income (earnings and wages) and consumption (non-durable and total) percentiles.²³ In the upper panel of the table, we cross-tab the 10, 15 and 20 percentiles of the earnings distribution against the 10, 25, 30,35 and 40 percentiles of the non-durable consumption distribution. If the two distributions were perfectly coherent, one should see that those who live in, say, the first 5 percent of the earnings distribution also live in the 5 percent of the consumption distribution. In other words, under perfect coherence, all the elements along the matrix main diagonal should be equal to 100. The reality is very different.

²² Here, we focus on total pre-tax and transfers earnings and therefore define saving as pre-tax and transfers earnings minus consumption.

²³ Contingency tables are designed to measure the degree of association between categorical and also ordinal random variables. Since the focus here is on percentiles, we use the contingency table to assess association between distributions.

Table 3
Consumption in the bottom of earnings and wage distributions

		Non-durable consumption					
Earnings	10	15	20	25	30	35	>40
10	3.34	19.03	11.88	9.47	7.37	6.18	42.73
15	3.68	16.90	10.53	8.88	7.65	6.85	45.52
20	2.53	16.99	9.44	7.36	7.20	6.47	50.02
		Total consumption					
Earnings	10	15	20	25	30	35	>40
10	2.78	14.86	7.61	6.73	6.37	5.94	55.71
15	3.01	15.27	8.18	7.18	6.44	5.42	54.51
20	3.15	13.83	6.68	6.50	6.03	6.53	57.28
		Non-durable consumption					
Wages	10	15	20	25	30	35	>40
10	22.41	11.19	9.09	7.19	5.81	4.84	39.47
15	18.65	9.47	9.60	7.92	7.36	5.91	41.08
20	16.10	8.35	7.99	7.09	6.24	5.87	48.37
		Total consumption					
Wages	10	15	20	25	30	35	>40
10	14.55	7.48	6.80	6.46	6.02	5.29	53.39
15	14.05	7.94	7.29	6.38	5.69	5.84	52.81
20	13.21	6.85	5.51	6.70	6.10	5.97	55.67

Note: the table shows the relative frequencies in each non-durable and total consumption percentiles by earnings and wages percentiles

The table confirms that income poor are not always consumption poor, whatever measure one uses of income (earnings or wages) and consumption (non-durable and total). For instance, 43% of individuals who are in the bottom 10 percent of the earnings distribution have consumption levels in the top 60% of the consumption distribution. This means that measured poverty depends on whether one uses income or consumption data. Therefore, if one wants to judge poverty, the use of income or consumption data would provide very different answers, since the income and consumption poor do not overlap.

If income and consumption were proportional, one should see that those who are in the bottom (top) of the income distribution are also in the bottom (top) of the consumption distribution. This is not the case and therefore some of the income poor are relying on debt past savings or transfers for their consumption expenditures.

To visualize that income and consumption poor do not necessarily overlap, we benchmark Table 3 against the case in which income and consumption poor perfectly overlap, the perfect coherence case. In such case, one should observe that those who live in the 10th, 15th, 20th, 25th, 30th, 35th (and so on) percentiles of the income distribution live

in the same percentiles of the consumption distribution: the percentage of household who live in the same percentile of the consumption and income distribution is 100, and accordingly the percentage of households living in different income and consumption percentile is equal to zero. To measure the lack of coherence between the income and the consumption distribution, then, we subtract 100 from the percentage of households who are in the same percentile of the income and consumption distribution: the less coherent the income and consumption distribution, the closer this number to -100. For instance, take those who are in the 10th percentile of the earnings distribution. The lack of coherence measure for those households is $3.34-100=-96.66$. The other side of the same coin is that the percentage of households who are in the 10th percentile of the earnings distribution is different from zero if we move to percentiles of the consumption distribution higher than the 10th. This is plotted in the top corner panel of Figure 36, which shows that a large number of households being in the 10th percentile of the earnings distribution are in other percentiles of the earnings distribution and that such number is smoothly distributed across consumption percentiles. The other panel of the left column of Figure 36 focus on the 15th and 20th percentiles of the earnings distribution, the right column on the 10th, 15th, and 20th percentiles of the wages. The figures convey the same basic message: being income poor does not imply being consumption poor. Even among those who live in the bottom percentiles of the income distribution, the share of those living in much higher percentiles of the consumption distribution is not negligible.

Comparing the wages and consumption distributions delivers similar results. Again the wages poor do not coincide with the consumption poor. However, the association between the wages and the non-durable consumption is lower than that between the earnings and the non-durable consumption distributions. The association between consumption and wages is 1.14 percent, between consumption and income ranks is 7.56 percent.²⁴ The evidence for total consumption is similar and the association between the wages and total consumption distributions increases to 2.64 percent.

In all cases, however, the association between the distributions of consumption and income though small is statistically significant. Measurement error tends to flatten both the distributions of consumption and income, and therefore to reduce the

²⁴ To quantify the association between consumption and income (wages), we use the Goodman and Kruskal's gamma, which counts the difference between concordant and discordant pairs in the comparison between ranks of two distributions.

association between them. So the degree of actual association might be higher than what we observe. Since the degree of association between the wages and consumption distributions is lower than that between the earnings and consumption distributions, measurement error can actually be an issue in our data, and this implies that we should see our estimate as a lower bound for the actual degree of association between the income and consumption distributions.

On the whole, this section has shown that the coherence between the distribution of consumption and income weakens for the poor and that therefore one might expect that those who are income poor are not consumption poor as well. This means that those who are income poor have not necessarily low standard of living as measured by consumption expenditure and tells that very different conclusion on measured poverty are drawn depending on whether income or consumption data are used.

We next turn to a rather different issue and study how consumption and income (and not their distributions) are related. To investigate how the relation between consumption and income evolves over the households life-cycle, one would ideally observe the same household for a long time span. This is not the case with our data, as each household is only observed for a year, at most. However, we use techniques that circumvent this problem. Rather than following the same individuals over time, we follow groups of households that share some characteristics, such as the decade of birth or the education achievement of the household head. We will be following how the average level and the inequality of income and consumption for these groups of households evolves over time. Notice that focussing on groups identified by the date of birth, we follow a cohort of individuals as they age through different phases of their life cycle. The difference between income and consumption (saving) is a way an individual has to move resources over time. Therefore a life cycle analysis is particularly interesting.

7. Relating consumption and income inequality.

This section shows that:

- **The correlation between changes in (relative) wages and in (relative) consumption is not stable over time.** For the years up to 1992, there seems to be a strong relationship between relative consumption and wages. The relationship is

strong. The results imply that a 1 percent increase in wages bring about between 0.8-1 percent increase in consumption. For the years after 1992, no statistical relation between consumption and wages is detectable.

- **The correlation between wages and consumption inequality is less than one.** Only a fraction of income inequality translates into consumption inequality. The fraction varies depending on the consumption and income measures used, but is always well below 1. An increase of wage inequality by 10 percent brings about an increase of consumption inequality of 3.5-3.6 percent. The wedge between income and consumption inequality is the indication that US households at least in part are able to smooth adverse income shocks.

In this section, we use the evidence presented in the sections 4 and 5 and complement it with the joint analysis of consumption *and* income inequality. This analysis might be informative about the ability of individuals (in the population at large or in specific groups) to buffer specific shocks they receive. And at the same time, it may shed some light about the nature of the shocks received. If we observe a group of individual households for whom income inequality has increased greatly and yet their consumption has not, one could speculate that those individuals did not need to modify their consumption (and ultimately their welfare) as a consequence. This might imply that they had a way to buffer these shocks (savings and borrowing, public and private transfers and so on) and/or that the shocks were amenable (maybe because they were not perceived to be excessively persistent) to be buffered. We divide our analysis in two parts. We first consider the evolution of mean consumption and mean wages in different groups in the population and ask whether those groups have been faring relatively well in terms of wages are also those that have done well in terms of consumption. This part of the analysis, therefore, focuses on the *relative performance* of different groups. In the second part, instead, we will be looking at inequality within groups and ask whether the evolution of the level of consumption inequality *within* these groups of individual households is related to the evolution of the level of wage inequality within the same groups. Therefore, this second part of the analysis will complement the first by focusing on the relationship between income and consumption inequality *within* groups. In both parts, as in some of the descriptive analysis above, the groups we will be considering will

be formed on the basis of the decade of birth of the household head and on his/her academic achievement. The fact we will be following the same groups over time will give a dynamic dimension to our analysis.

7.1. Relative consumption and wages

We report the main results of this section in Table 5. The table contains our estimates of the relationship between average wages and consumption, controlling for permanent differences across groups as well as common movements over time. In particular, given a group labelled with the superscript g , if we denote with c_t^g and w_t^g its average (log) consumption and wages, we are interested in the following relationship:

$$(1) \quad c_t^g = d^g + \mu_t + \gamma w_t^g + \varepsilon_t^g$$

where d^g captures the average level of consumption over time and therefore controls for permanent differences across groups, μ_t controls for common year effects and ε_t^g is a random term with zero mean. We are interested in the coefficient γ , which tells us what part of relative changes in wages gets reflected in relative changes in consumption. There are several ways one can estimate this coefficient. First, one can literally estimate the parameters of equation (1) (although, as we will see, different statistical and econometric techniques are available for the purpose. Alternatively, one can eliminate the nuisance parameters d^g by considering the terms of equation (1) at two different dates and taking the difference between the two expressions. By doing so, one would estimate:

$$(2) \quad \Delta^k c_t^g = \mu_{t,k} + \gamma \Delta^k w_t^g + v_t^g$$

where $\Delta^k c_t^g = c_t^g - c_{t-k}^g$, $\Delta^k w_t^g = w_t^g - w_{t-k}^g$, $\mu_{t,k} = \mu_t - \mu_{t-k}$ and $v_t^g = \varepsilon_t^g - \varepsilon_{t-k}^g$. By considering the specification in level we focus on more long term changes, while when considering the specification in differences (2), we focus on changes over time. In Table 4, we report the results for annual changes ($k=1$) and five-year changes ($k=5$).

Table 4
Correlation over time between relative changes in consumption and wages

	Non durables			Total consumption		
	1982-1992	1993-2003	1982-2003	1982-1992	1993-2003	1982-2003
First diff. OLS	0.003 (0.494)	0.079 (0.059)	0.053 (0.044)	0.117 (0.062)	0.114 (0.060)	0.117 (0.043)
First Diff. IV	-1.538 (9.555)	0.578 (1.200)	0.813 (1.953)	-0.638 (2.833)	-1.168 (2.214)	-1.048 (2.982)
Five-year Diff OLS	0.312 (0.098)	0.060 (0.053)	0.155 (0.048)	0.312 (0.096)	0.063 (0.058)	0.162 (0.051)
Five-year Diff IV	0.999 (0.301)	0.111 (0.104)	0.361 (0.101)	0.855 (0.271)	-0.052 (0.115)	0.199 (0.102)

Note: Year cohort groups considered if median age is greater than 24 and less than 62. Standard errors reported in parenthesis.

In the left panel of the table we report the results we obtain when use expenditure on non-durable and services as our measure of consumption, while on the right panel, we report the results for the measure that adds to the previous one our measure of car services. For each of the two panels, we consider our estimates of the parameter γ considering two sub samples (years before 1993 or after 1992) and for the entire period. Finally, each coefficient is estimated twice, first using Ordinary Least Squares and then using a technique, Instrumental Variable, which attempts to take into account the fact that the average wage of a given group is measured with some error, because of the limited sample size.²⁵

The results indicate some important differences between the first and second part of the sample. For the years up to 1992, there seems to be a strong relationship between average group consumption and average group male wages. The relationship is stronger (although measured with less precision) when we use the IV techniques, as the OLS

²⁵ The average wage is measured using the CEX sample. The 'instrument' we use is the same average measured in the CPS sample. As the two samples are independent there is no reason to believe that the errors in the two measures are correlated.

estimates are probably affected by attenuation bias induced by measurement error. Only when we consider year on year changes we fail to identify a strong relationship between the two variables. In the case of five year changes, we estimate the coefficient γ to be quite high: 1 in the case of non durable consumption and 0.85 for the measure of consumption that includes car services. A high coefficient indicates that the shocks to (relative) wages are reflected into changes in (relative) consumption. Interestingly the point estimates seem lower for the more comprehensive measure of consumption.

Things look quite different in the second period. Both the OLS and the IV estimates are considerably smaller than the corresponding estimates for the first period. Not surprisingly, the estimates for the entire period are somewhat in the middle.

The results in Table 4, therefore, indicate that while there is a strict correspondence between consumption and male wage shocks across groups until the early 1990s, this relationship is much attenuated in the subsequent period. The existence of a relationship between (changes in) relative wages and (changes in) relative consumption up to the early 1990s confirms the evidence obtained, using slightly different data and techniques, by Attanasio and Davis (1996).

The change in this relationship documented here for the second period is instead consistent with some of the results presented by Krueger and Perri (2003), although the size of the changes in consumption inequality they present are different from our preferred measures. The fact that the coefficient on wage (changes) is much smaller in the second period than in the first can be indicating different things. One possibility is that the nature of the shocks to relative male wages has changed after the early 1990s.

Our evidence on the returns to education in Section 4 did indicate somewhat smaller movements in relative wages. However, the overall variability of relative wages is, if anything, slightly larger in the second period relative to the first. It is possible, however, that the nature of these changes makes them more easily absorbed and therefore not noticeable in consumption. The evidence in the labour literature that we have cited above seems to indicate that in the 1990s temporary shocks have become relatively more important than permanent shocks relative to the 1980s. The latter are harder to absorb than the former. It is also possible that there have been institutional changes that have

endowed households with better and more efficient ways to ‘smooth out’ certain shocks. The development of more sophisticated financial instruments as well as changes in safety nets and the like might be examples of such changes. All in all, relative wage changes have less impact on living standards in the period from 1992 to 2002 than in the period from 1982 to 1992.

7.2. Within group inequality in consumption and wages

We now move to the relationship between the evolution of inequality in consumption and that of inequality in wages (and earnings) within the groups we have considered in the previous section. Here we are interested in answering the following question: after controlling for differences in the *level of inequality* in a given group, to what extent changes in wage inequality within a group are reflected in changes in consumption inequality within that group. To answer this question we consider the following simple equation:

$$(3) \quad Sd_t^g(c) = k^g + \theta Sd_t^g(w) + u_t^g$$

where $Sd_t^g(c)$ and $Sd_t^g(w)$ are the standard deviation of (log) consumption and wages, respectively) within group g at time t . As with equation (1), the coefficients k^g take into account permanent differences in inequality across groups. Our estimates of the coefficient θ are reported in Table 5 where, again, we use both OLS and IV techniques and use as a measure of consumption both non durable consumption and the measure that includes car services. In the left panel we use male wages, while in the right panel we use the standard deviation of log household earnings.

As in Table 4, the IV estimates are considerably larger than those by OLS, indicating the presence of attenuation bias induced by measurement error. In all cases, the coefficients are statistically different from zero, indicating a relationship between the two measures of within group inequality. In the case of the IV estimates, the coefficients are also economically significant, with the coefficients being as high as 0.35 in the case of male wages. In the case of earnings (right panel) the point estimates are slightly lower at 0.27 for non durable consumption and 0.23 for the measure that includes car services.

	Non durables and male wages	Total cons and male wages	Non durables and household earnings	Total cons. and household earnings
OLS	0.053 (0.026)	0.047 (0.024)	0.003 (0.013)	0.003 (0.012)
IV	0.358 (0.095)	0.347 (0.109)	0.269 (0.094)	0.228 (0.090)

Table 5 conveys an important message: only a fraction of within group earnings and wage inequality translates into consumption inequality. The fraction varies depending on the consumption and income measures used, but is always well below one. If one focuses on non durable consumption and (male) wages, an increase of wage inequality by 10 percent brings about an increase of consumption inequality of 3.5-3.6 percent. The wedge between income and consumption inequality is the indication that US households at least in part are able to smooth adverse income shocks. Identifying what is the exact insurance mechanism is beyond the scope of this work and a challenge for future research.

8. Conclusions.

It is now time to take stock on the results we have been presenting in this study. We have been arguing that, for a variety of reasons, if one is interested in the evolution of inequality and the distribution of material well being, the analysis of consumption and/or expenditure is an important complement to the analysis of income and wages distributions. Indeed, in some important dimensions, the evolution of the distribution of consumption can be much more informative than the distribution of earnings. The main reason for this assertion are the fact that for most people material well being is determined by consumption and income is valuable only in so far as gives access to consumption. Moreover, consumption is likely to react to permanent shocks affecting individual households and not necessarily to short term fluctuations in disposable income. The joint analysis of consumption and income can be informative about the nature of shocks that affect individual households as well as the instruments (such as assets, debt, public and private transfers) they can use to absorb income shocks and avoid consumption fluctuations.

Although the analysis of consumption is potentially very useful, it has been much less common than the analysis of wages or income. The main reason for this difference is the availability of data and the perception that the data on consumption are much inferior to the data sets that have been traditionally used in the analysis of wage and income inequality. While it is certainly true that the main data set containing consumption information is not exempt from problems, these have been to a certain extent exaggerated. In Section 4, for instance, we show that the CEX samples yield information on wages and income that is consistent, by and large, with that emerging from the CPS. With this we do not want to minimize the importance of the data quality problems with the CEX, and in particular the relative low fraction of NIPA Personal Consumption Expenditure that can be accounted for by the CEX data.

Given these caveats on the data, some interesting patterns emerge from the data. We can summarize them in the following list:

1. Consumption inequality has increased in the 1980s and 1990s, but the increase has been considerably less pronounced than the increase in inequality in wages and income. In the 20 years from the early 1980s to the early 2000s the overall standard deviation of log non durable consumption increases by more than 7 percentage points. This compares with an increase of 12-15 percentage points for the standard deviation of log wages and of about 10 percentage points for total household earnings.
2. While the increase in overall inequality in wages (and incomes) is concentrated in the first part of the sample, for consumption it goes on throughout the sample.
3. We identify both an increase in inequality across education groups and within education groups. The former, however, happens mainly during the 1980s and up to the early 1990s. The latter continues until the end of our sample and can be observed for all education groups.
4. Individual households that can be characterized as ‘income poor’ in that they are located in the left tail of the earnings (or wage) distribution, do not necessarily coincide with the households that are ‘consumption poor’. Moreover, the difference in consumption between households at the bottom of the earnings distribution is much smaller than the difference in earnings. Consumption of the ‘income poorest’ household exceeds earnings.

5. If we analyze the dynamics of relative wages and consumption changes in groups defined by decade of birth and education over time, we observed that until the early 1990s there is a relationship between relative changes in consumption and wages, consistent with the evidence in Attanasio and Davis (1996). In more recent years, however, such a relationship is much weaker. In other words, while the changes in relative wages (across education and age groups) were to a considerable extent reflected in changes in consumption changes, the same is not true of the changes observed after the early 1990s. This could be due to the different nature of the changes observed in incomes and wages in the last decade. It should be stressed, however, that while the changes in inequality do not line up as in the 1980s and early 1990s, we still witnessed an increase in inequality in consumption in the last decade. As mentioned above, the changes in wage inequality over the last period are mainly changes *within groups*, rather than across groups. As such, they might reflect temporary rather than permanent shifts and shocks to relative wages.
6. We did identify a relationship between the movements in wage and consumption inequality *within groups*. This indicates that, at least some of the changes in relative wages within group are reflected into consumption. This evidence is consistent with the fact that consumption inequality kept increasing even in the 1990s. It should be stressed, however, that the coefficient is far from being equal to one, indicating that part of the shocks are *not* reflected in consumption.

Obviously the results above are based on the assumption that the data quality problems do not introduce too much noise and bias. We have spent a considerable amount of effort in an attempt to validate the CEX data against a more widely used survey, such as the CPS. And there is a sense in which the samples are comparable, especially in terms of the information that can be obtained about wages. There is, however, a real worry about the reliability of the CEX consumption data, especially in the last part of the sample. We cannot conclude this study without an appeal for better measures of consumption in the US. While it is true that over the last 20 years the reliability and quality of many individual based surveys seems to have worsened, the case of consumption is particularly serious. The consequence of this situation is that the largest economy in the world lacks a reliable and comprehensive survey that measures the main purpose of economic activity: consumption. Such a measure would be crucial

for a variety of reasons, of which the current aim of the CEX, the measurement of the weights for the construction of the CPI, is only one. We have argued that reliable data on consumption and its distribution not only provides direct information on material well-being, but especially jointly with information about income, can shed light on the nature of the economic shocks received by individual households as well as the extent to which they are able to buffer such shocks.

The construction of reliable data on consumption at the household level in an advanced society is not easy, as it involves thousands of different items, that change and evolve continuously. This difficulty is compounded by the apparently increasing reluctance individuals have to answer surveys that are time consuming and invasive. However, in recent years there have been considerable advances and progresses in the measurement of other economic variables that, until a couple of decades ago, were deemed to be almost impossible to measure at the individual level. We are thinking, for instance, of the tremendous progress made in measuring individual financial wealth, promoted in surveys such as the HRS, the SCF and the PSID. This progress should indicate that innovative questionnaire techniques and intelligent use of new technologies could yield very high payoffs in terms of data quality. There are surveys in the industrialized world that yield measures of household level consumption that seem to be of better quality than the CEX. The UK Family Expenditure Survey, based on a mixture of individual (rather than household) diaries and retrospective interviews, has been able for many year to match extremely well aggregate data on Personal Consumption Expenditure from the NIPA accounts. These experiences could be studied and, suitably adapted, maybe replicated in the US.

Another important issue is that of sample size. While the CEX samples were significantly increased in 1999, the current size is still insufficient for any analysis that requires many dimensions simultaneously, such as educational achievement, state of residence, decade of birth and so on.

Appendix.

A1. Combining consumption information from the survey components of the CEX

In this section we describe the methodology that we followed to combine the information from the Interview Survey (IS) and from the Diary Survey (DS) components of the CEX. The approach that we take in this Appendix builds upon previous work by Battistin (2003) and Attanasio, Battistin and Ichinura (2007, ABI), to which the reader is referred for additional details about the technicalities.

The results presented in this Appendix can be summarised as follows. The minimum set of assumptions required to combine DS and IS information is about the time series of covariances between consumption components as in Battistin (2003) and ABI (2007). These assumptions allow one to identify the evolution of the mean and the squared coefficient of variation of consumption over time (and by groups), but not percentiles of the distribution of consumption. Further assumptions are called for in order to identify the entire distribution of consumption, and these assumptions are discussed in Attanasio and Battistin (2005). The discussion in what follows focuses on the former set of assumptions.

Assume that non-durable expenditure comprises expenditures on *two* sets of goods, D and I, so that it can be defined as follows:

$$C_{ND} = C_I + C_D.$$

The set I includes expenditures on those types of goods the IS component of the CEX is designed for (types of expenditures respondents can recall for a period of three months or longer), whereas the set D includes expenditures on frequently purchased smaller items (including food and beverages, both at home and in food establishments, housekeeping supplies, tobacco, non-prescription drugs, and personal care products and services).

Total expenditure comprises expenditures on non-durable and durable goods:

$$C_{TOT} = C_{ND} + C_{DUR},$$

the latter set being defined as described in Section 3.4.

Throughout this paper we will consider the within-group and between group variance as indexes of inequality (with groups being indexed by G):

$$\begin{aligned} V[\lg C_{TOT} | G = g], \\ E[\lg C_{TOT} | G = g], \end{aligned}$$

for which the following first order approximations are defined:

$$\begin{aligned} V[\lg C_{TOT} | G = g] &\square \frac{V[C_{TOT} | G = g]}{E[C_{TOT} | G = g]^2}, \\ E[\lg C_{TOT} | G = g] &\square \lg E[C_{TOT} | G = g] - \frac{1}{2} V[\lg C_{TOT} | G = g]. \end{aligned}$$

Note that the previous expressions hold exactly if total expenditure is normally distributed (see Battistin, Blundell and Lewbel, 2007). Clearly we have that

$$E[C_{TOT} | G = g] = E[C_I | G = g] + E[C_D | G = g] + E[C_{DUR} | G = g], \quad (1.1)$$

and

$$V[C_{TOT} | G = g] = V[C_I | G = g] + V[C_D | G = g] + V[C_{DUR} | G = g] + 2C[C_I, C_D | G = g] + 2C[C_I, C_{DUR} | G = g] + 2C[C_D, C_{DUR} | G = g], \quad (1.2)$$

where each covariance term can also be written as

$$C[C_i, C_k | G = g] = \rho_{i,k}(g) \sqrt{V[C_i | G = g] V[C_k | G = g]}. \quad (1.3)$$

The term $\rho_{i,k}(g)$ refers to the correlation coefficient between consumption expenditures (C_i, C_k) and is of course bounded below one in absolute terms.

Estimation of the within-group variance proceeds as follows. Means and variances in expressions (1.1) and (1.2) are estimated using the most reliable survey component of the CEX. As for the covariances in (1.2), we use (1.3) and

- exploit the two measures of the correlation coefficient between (C_I, C_D) as obtained from the IS and the DS components of the CEX; these measures are combined efficiently in the estimation via a GMM procedure by assuming that the growth rate of the time series of the two covariances is the same in the two surveys. Evidence in favour of this assumption is provided in Battistin (2003) and ABI (2007).
- estimate the correlation coefficient between (C_I, C_{DUR}) using that observed coefficient in the IS component of the CEX.
- estimate the correlation coefficient between (C_D, C_{DUR}) using that observed in the DS component of the CEX.

Building upon a summary of the literature on measurement error in survey reports of consumption expenditures as well as very recent publications by researchers at the BLS (see Battistin, 2003, and Garner *et al.*, 2006), Table A1 reports the expenditure categories used to defined the I and the D sets of goods.

Figure A1 plots the time evolution of the standard deviation of log non-durable consumption. The figure shows two lines: the continuous line is drawn using the covariance between I and D goods as measured in the IS, the dotted the covariance as measured in the DS. The latter starts for 1986 since in that year the DS becomes providing exhaustive measures of consumption. The measure that uses the IS estimate of the covariance is slightly smoother but on the whole the pattern emerging from the two series is reasonably similar.

Table A1
Expenditure categories

<i>Non-Durable Goods and Services</i>	
Non durable consumption expenditures from the Diary Survey of the CEX (“D” goods)	
Food at home	Food and Non-Alcoholic Beverages at Home
Food away from Home	Food and Non-Alcoholic Beverages away from Home
Alcohol	Alcoholic Beverages (at home and away from home)
Tobacco	Tobacco and Smoking Accessories
Housekeeping Services	Housekeeping Services
Personal Care	Personal Care
Entertainment Services	Non-durable Entertainment Expenses Newspapers and Magazines
Non durable consumption expenditures from the Interview Survey of the CEX (“I” goods)	
Housing and Public Services	Home Maintenance Services Public Utilities Miscellaneous Home Services
Heating Fuel, Light and Power	Fuel gas and petroleum
Transportation	Fuel for Transportation Transportation Equipment Maintenance and Repair Public Transportation Vehicle Rental and Misc. Transportation Expenses
Clothing, Footwear and Services	Clothing Footwear Clothing services
<i>Durable goods and services</i>	
Durable consumption from the Interview Survey of the CEX (“DUR” goods)	
Car Services	Car Services

A2. Estimating cars services

This section discusses how we estimate the value of the stock of cars using a sample of micro-data drawn from the U.S. Consumer Expenditure Survey. Estimating the value of a car amounts to identify a single numerical index measuring its ‘quality’. This last depends on a number of features, including the year of production of the car, its age and the general level of prices. If the level of prices changes over time because of inflation, no identification strategy is available, which allows distinguishing among the three aforementioned effects. To get around this issue, we propose to use a set of cars characteristics to proxy for the year of production effect. Overall, we can estimate the value of the stock of cars for around 415,000 data-points (each data point corresponds to an household interviewed in a given quarter).

A2.1 The data

The data on vehicles come from two files. The BLS starts to make publicly available these files since 1984. The first file (OVB), which refers to the vehicles owned by the household, records a full set of characteristic for the vehicles present in the Consumption Unit (CU) at the interview date. A partial list of these characteristics includes the type of the vehicle (car, truck, van, pick-up, motor-bike, boat and, eventually, airplane), the make and the model of the vehicle, the year and the month of purchase, the vintage, the number of cylinders, if the vehicle entered the consumption unit as new or used, if the vehicle is equipped with the air-conditioning, the automatic transmission, the power-brake, the power-steer, the radio, the sun-roof. The list also includes the purchasing price which contains two components. The net purchasing price, which turns out to be the cash outflow at the date of the purchase and, if any, the trade-in allowance received.

Moreover, households are asked if they disposed a vehicle and, in case they did, they are asked when, how and a set of characteristics identifying the vehicle. This information are recorded in the second file (OVC) which refers to the vehicles disposed by the households. The way the household dispose the vehicle is particularly important to our purposes. Six alternatives are reported: Sold, Traded in, Given Away Outside the CU, Damaged beyond repair, Stolen and Other. To estimate the car value, we exploit data on the price the CU receives for selling the car.

In both the OVB and the OVC files a vehicle number is present, which identify the vehicle within the Consumption Unit and allows merging the information contained in one file with those in the other.

The sample covers the years from 1984 to 2003. Around 800,000 cars are present in the sample, 60% of those are second-hand, while the number of make-models averages around 890 and the number of brands around 80. Around the 2.5% of cars are top coded.

The most frequent make-model is the Oldsmobile Cutlass (around 2%). The data include cars produced before than 1969. The CEX does not deliver a point value for vintages between 1970 and 1980. Rather, the survey specifies an interval to which the year of production belongs.

Moreover, we use a set of cars’ characteristics, such as whether the car has automatic transmission or not, the number of cylinders, whether power brake or power steer are present. It comes out that around 76% of cars in our sample has automatic transmission, 85% of cars has power steer and the same amount has power brake. More importantly to our purposes, 74% of cars produced in 1986 are equipped with automatic transmission,

but for cars produced ten years later this number goes to 86%. Similarly, 89% of cars produced in 1986 is equipped with power brake, while this number increases to 97% for cars produced before 1996 and the same pattern is observed for the percentage of cars equipped with power steer. This means that these car characteristics change across vintages and might therefore be used to proxy vintage effects.

We next turn to discuss the econometric issues involved in the estimation of the index measuring the quality of a car.

A2.2 The econometric issues

There are a number of econometric issues to be dealt with when the quality of a car has to be inferred from the observed price. Next, we clarify these issues and explain how we addressed them.

Suppose that we observe a car for V vintages. If we normalize to one the quality of, say, vintage v , the ratio:

$$\frac{P_{v+1,t}}{P_{v,t}} \quad (1)$$

measures the quality of the vintage $v+1$ conditional on the time the two subsequent vintages are observed. If each vintage is observed for enough long time, averaging (1) over t gives a single numerical index which measures the value of the cars in efficiency unit.

Now, notice that the age of the cars whose price is involved in the computation of (1) is different, since, trivially, $a=t-v$. If the value of cars changes because of aging, which, indeed, seems to be the case, the ratio in (1) depends also on a pure age-effect. This age-effect is often assumed to be a consequence of the depreciation.

If more aged cars deliver 'less' services and, then, are valued less, we expect the depreciation pattern to be decreasing.²⁶ The rate at which the car depreciates determines the concavity of the age-value profile. Comparing the price of cars at the same age but at different times might help to account for the age-effect. However, this comes at the cost of introducing a time-effect, which makes the price of cars to change only because of inflation.

To further illustrate this problem, suppose that the price data are arranged in a matrix. The value of a car is a function of age and time. For simplicity, we assume that the maximum age and the maximum time for which the prices are observed is 5. In the rows of this matrix, the age is constant while the time varies. Obviously, the opposite holds true for the columns. An example of such matrix is shown in Table A.2.

The difference between the average of prices in the, say, second row and the average of prices in the first row would be a measure of how the price changes because of aging from age 1 to age 2. In the same way, the difference between the average of prices in the, say, second column and the average of prices in the first column would be a measure of how the price changes because of inflation from year 1 to year 2.²⁷

²⁶It is worth noticing at this stage that it might happen that some cars are observed to appreciate, which means that their values increases with age.

²⁷This procedure consists of computing a within-group average, where the group membership is first with respect to age and the with respect to time.

However, this procedure leads in general to biased estimate of the age and the time effect: the problem is that the prices of cars in a given row (or column) belongs to different cars, in that their vintage differs. Only moving along the diagonals we observe cars belonging to the same vintage.

Whether or not comparing cars belonging to different vintages to remove the age and the time effect is indeed a problem is an empirical matter. The main difficulty to asses the relative importance of the three effects (age, time and vintage) is related to the fact that they are not separately identifiable.

The literature offers two main strategies to deal with the problem. The first one amounts to normalize one of the three effects, say, the vintage effect, to zero. If the vintage effect proxies for the degree of technological progress embodied in the price of cars, this assumption sets to zero the net price change due to technological progress. In other words, this strategy does not allow to identify the trend in the degree of technological progress.

Alternatively, Hall (1971), in a study which focuses on trucks, suggests using a set of characteristics, such as the Wheelbase, the Weight, the Ratio of Bore to Stroke, the Horsepower, the Torque, the Tire Width, to proxy for the vintage effect in an hedonic prices regression framework. The rationale is that this set of characteristics can be arranged in a vector which is a sufficient statistic for the vintage effect. If this is indeed the case, the identification problem is circumvented because these characteristic are chosen to be orthogonal to the age and the time-effect.

In what follows, we decide to pursue the second strategy. The main advantage of this strategy is to make possible the identification of all the three effects, the main disadvantage is to rely on the availability of a set of characteristic rich enough to be used as a proxy for the quality. Given that the ultimate goal of this work is to evaluate the stock of cars either strategy might be used. The choice of the second strategy is mainly based on empirical grounds.

The price of the cars at age a and time t can be written as:

$$P(a,t) = d_a t_i f_v \quad (2)$$

where v is the vintage; d_a is the age effect, t_i is the time effect and f_v is the vintage effect. From (2) it is clear that we cannot simultaneously identify the three effects. In order to achieve identification I replace f_v by a set of characteristics.

We assume that prices are measured with error and that the error is multiplicative. Since the model linear in the logs, the age, time and the vintage effects could be estimated through a linear regression. The issue here is what functional form to chose. To understand it, go back to Table A1.

If in a matrix like Table A1 there are not 'holes', which means that we observe at least one price for each age-time cell, an analysis of variance (ANOVA) model could be used. The prices of cars are regressed on a (restricted) set of age, time and vintage dummies.

If, instead, we do not observe a price for each age-time cell, we need to save on the number of parameters to be estimated. This might be accomplished by fitting to the price of cars a polynomial in age, time and vintage (abstracting for a while from the identification issues). Due to data constraints, we opt for this second model and estimate the following parsimonious specification:

$$\ln P_i(a,t) = \alpha_0 + \alpha_1 t + \alpha_2 a_i + \alpha_3 a_i^2 + v_i' \alpha_4 + \varepsilon_{i,a,t} \quad (3)$$

where the LHS variable is the log of the price, and the RHS are a linear time trend, a quadratic polynomial in age, a vector of car characteristics, v_i .²⁸ Padula (2001) contains more details on the estimation of (3) and validates the results by comparing the estimated with the actual price of some selected models of cars.

The parameters estimated from equation (3) are the used to impute the value of the stock of cars. Imputing the value of the stock of cars only on the basis of equation (3) would amount to reduce the amount of heterogeneity in the value of cars. To restore, at least in part, the heterogeneity in car values we add to the fitted prices from equation (3) and error drawn from a normal distribution with mean equal to zero and standard deviation equal to the standard deviation of the residuals from the estimation of (3).

Table A2: The Age-Time Matrix

<i>(age,time)</i>	1	2	3	4	5
1	P(1,1)	P(1,2)	P(1,3)	P(1,4)	P(1,5)
2	P(2,1)	P(2,2)	P(2,3)	P(2,4)	P(2,5)
3	P(3,1)	P(3,2)	P(3,3)	P(3,4)	P(3,5)
4	P(4,1)	P(4,2)	P(4,3)	P(4,4)	P(4,5)
5	P(5,1)	P(5,2)	P(5,3)	P(5,4)	P(5,5)

Note: Age is constant along the rows, while time is constant along the columns.

²⁸ We also tried different specifications, adding make-model dummies in the equation, or interacting the age term with make model-dummies, or replacing the polynomial in age with a full set of age dummies. The overall fit of the equation, as measured by the adjusted R-squared, does not change much across specification and lies around 65%.

References

Attanasio, Orazio. P., and Steven. J. Davis, 1996, "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, CIV, pp. 1227-1262.

Attanasio, Orazio. P., Erich. Battistin and Hidehiko. Ichimura, 2007, "What Really Happened to Consumption Inequality in the U.S.?" in E. Berndt and C.Hulten, (Eds.) *Measurement Issues in Economics – Paths Ahead: Essays in Honour of Zvi Griliches* (Chicago: University of Chicago Press).

Attanasio, Orazio P. and Battistin, Erich, 2005, *A Comparison of CE Interview and Diary Data: Examination of Different Moments*, mimeo, University College London and IFS.

Attanasio, Orazio. P., Erich. Battistin and Andrew. Leicester, 2006, "From Micro to Macro, from Poor to Rich: Consumption and Income in the UK and the US," mimeo University College London.

Attanasio, Orazio. P. and Josè. V. Ríos-Rull, 2000, "Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?", *European Economic Review*, Vol. 44 (7) pp. 1259-1289.

Autor, David.H., Katz, Lawrence. F. and Melissa S. Kearney, 2007, "Trends in US Wage Inequality: Revising the Revisionists", Mimeo, Harvard University. <http://www.economics.harvard.edu/faculty/katz/papers/AKK-ReStatRevision.pdf>

Battistin, Erich, 2003, "Errors in Survey Reports of Consumption Expenditures", Working Paper W03/07, Institute for Fiscal Studies, London.

Battistin, Erich, Blundell, Richard and Arthur, Lewbel, 2007, "Why is Consumption More Log Normal Than Income? Gibrat's Law Revisited," Working Paper W08/07, Institute for Fiscal Studies, London, and Working Paper 671, Department of Economics, Boston College.

Battistin, Erich and Mario Padula, 2009, "The Effect of the Survey Instrument on Reports of Consumption Expenditure", mimeo.

Blundell, Richard and Ian Preston, 1998, "Consumption Inequality and Income Uncertainty," *The Quarterly Journal of Economics*, 113, pp.603-640.

Blundell, Richard, Pistaferri, Luigi and Ian Preston, 2008, "Consumption Inequality and Partial Insurance," *American Economic Review*, 98:5, pp. 1887–1921

Bowlus, Audra J. and Jean-Marc Robin, 2004, "Twenty Years of Rising Inequality in U.S. Lifetime Labour Income Values," *Review of Economic Studies*, 71, pp. 709–742

Broda, Christian and David E. Weinstein, 2007, "Product Creation and Destruction: Evidence and Price Implications," NBER Working Papers 13041, National Bureau of Economic Research, Inc.

Browning, Martin and Thomas Crossley, 2000, "Shocks, Stocks and Socks: Consumption Smoothing and the Replacement of Durables During an Unemployment Spell," Econometric Society World Congress 2000, Contributed Papers.

Cutler, David and Lawrence Katz, 1991, "Macroeconomic Performance and the Disadvantaged," *Brookings Papers on Economic Activity*, 2, pp. 1-61.

Cutler, David and Lawrence Katz, 1992, "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s," *American Economic Review*, LXXXII, pp. 546-561.

Davis, Steven and Paul Willen, 2000, "Occupation-level income shocks and asset returns: their covariance and implications for portfolio choice," NBER Working Paper no. 7905.

Fay, Scott, Erik Hurst and Michelle White, 2002, "The Household Bankruptcy Decision," *American Economic Review*, 92, pp.706-718.

Flinn, Christopher, 2002, "Labour Market Structure and Inequality: A Comparison of Italy and the U.S." *Review of Economic Studies* 69, pp.611-645.

Garner, Thesia, I, Janini, George, Passero, William, Paszkiewicz, Laura and Mark Vendemia, 2006, "The CE and the PCE: a comparison" *Monthly Labor Review*, September 2006, pp.20-46.

Gieseman, Raymond, 1987, "The Consumer Expenditure Survey: quality control by comparative analysis", *Monthly Labor Review*, March 1987, pp. 8-14.

Gordon, Robert J, 2006, "The Boskin Commission Report: A Retrospective One Decade Later," *International Productivity Monitor*, Spring, pp. 7-22.

Gottschlak, Peter. and Robert Moffitt, 1994, "The Growth of Earnings Instability in the U.S. Labor Market", *Brookings Papers on Economic Activity*, 217-272.

Gottschlak, Peter. and Smeeding, Timothy, 1997, "Cross-National Comparisons of Earnings and Income Inequality", *Journal of Economic Literature*, 35, pp. 633-687.

Grant, Charles, Christos Koulovatianos, Alex Michaelides and Mario Padula, 2009, "Evidence on the Insurance Effect of Redistributive Taxation," *Review of Economics and Statistics*, forthcoming.

Hall, Robert, "The Measurement of Quality Changes from the Vintage Price Data," in Z.Grilliches (ed.), *Price Indexes and Quality Changes* (Cambridge: Harvard University Press, 1971).

Juhn, Chinhui., Kevin M. Murphy and Brooks Pierce (1993), "Wage Inequality and the Rise in Returns to Skill." *Journal of Political Economy*, 101(3), pp. 410-442.

Katz, Lawrence. and Autor, David, 1999, "Changes in the Wage Structure and Earnings Inequality", in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, Vol. 3A (Amsterdam: North-Holland).

Katz, Lawrence .H. and Kevin M. Murphy, 1992, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107(1), pp. 35-78.

Krueger, Dirk and Fabrizio Perri, 2003, "On the Welfare Consequences of the Increase in Inequality in the United State" NBER Macro Annual, pp. 83-121.

Krueger, Dirk and Fabrizio Perri, 2006, "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73, pp. 163-193.

Mankiw, Gregory and Miles Kimball, 1989,"Precautionary Saving and the Timing of Taxes," *Journal of Political Economy*, 97, pp.863-879.

Meyer, Bruce and James Sullivan, 2004, "The Effects of Welfare and Tax Reform: The Material Well-Being of Single Mothers in the 1980s and 1990s", *Journal of Public Economics*, 88, pp. 1387-1420.

Murphy, Kevin and Finis Welch, 1992, "The Structure of Wages", *The Quarterly Journal of Economics*, 107(1), pp. 285-326.

Padula, Mario, 2001, "Household investment behaviour: empirical investigations of durable consumption, returns to education and borrowing restrictions", Phd Thesis, UCL, mimeo.

Pencavel, John, 2006, "A Life Cycle Perspective on changes in Earnings Inequality among Married Men and Women," *The Review of Economics and Statistics*, Vol. 88(2), pp. 232-242.


Slesnick, D.T., 1993, "Gaining Ground: Poverty in the Postwar United States", *Journal of Political Economy*, CI, pp. 1-38.

Slesnick, D. T. , 2000, *Living Standards in the US: A consumption based approach*, American Enterprise Institute, Washington, DC.

Slesnick, D.T. , 2001, *Consumption and Social Welfare: Living standards and their distribution in the United States*, Cambridge University Press.


Figure 1.

Non-durable consumption inequality from the interview and the diary survey




Note: the figure plots the time evolution of the coefficient of variation of non-durable consumption as measured in the CEX Diary and in the Interview Survey. The dashed line are obtained by a locally weighted regression on a linear time trend.

Figure 2
Durable consumption inequality.



Note: The figure shows the coefficient of variation of durable consumption measured as the sum of the flow of services from cars.

Figure 3
Mean equivalence scale



Note: The figure shows the mean of equivalence scale computed form the Interview Survey.

Figure 4
CEX and CPS wages

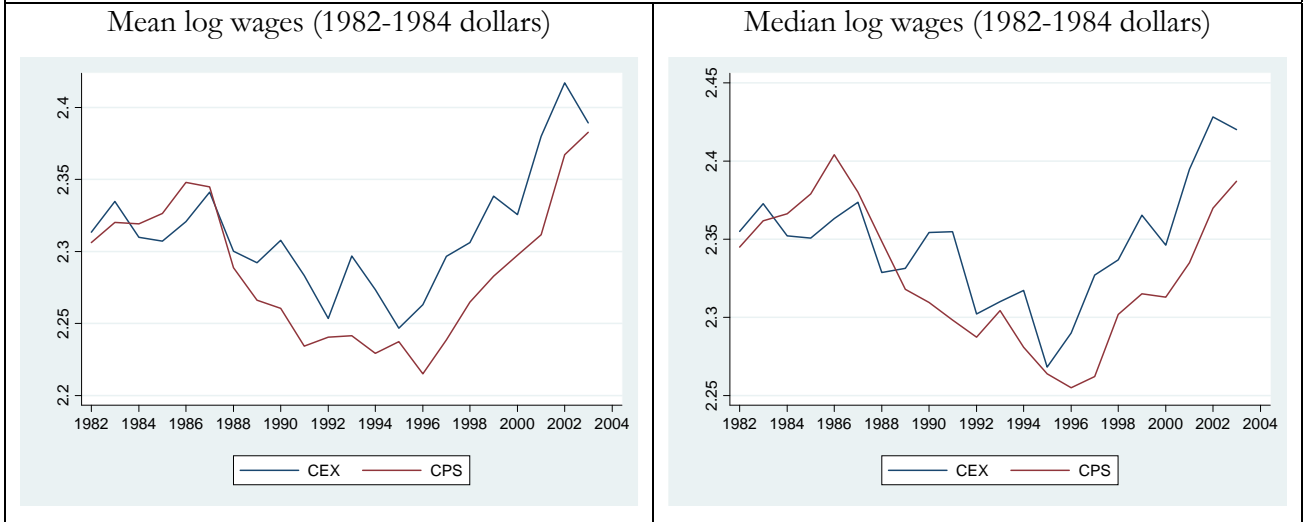


Figure 5
Median log wages (CEX and CPS) by decade of birth cohort




Figure 6
Median log wages (CEX and CPS) by education achievement




Figure 7
Differences across education groups:
CEX and CPS




Figure 8
90th – 10th percentile for log wages
CEX and CPS




Figure 9
Coefficient of variation of wages
CEX and CPS




Figure 10
Coefficient of variation of wages
CEX and CPS- by decade of birth cohort

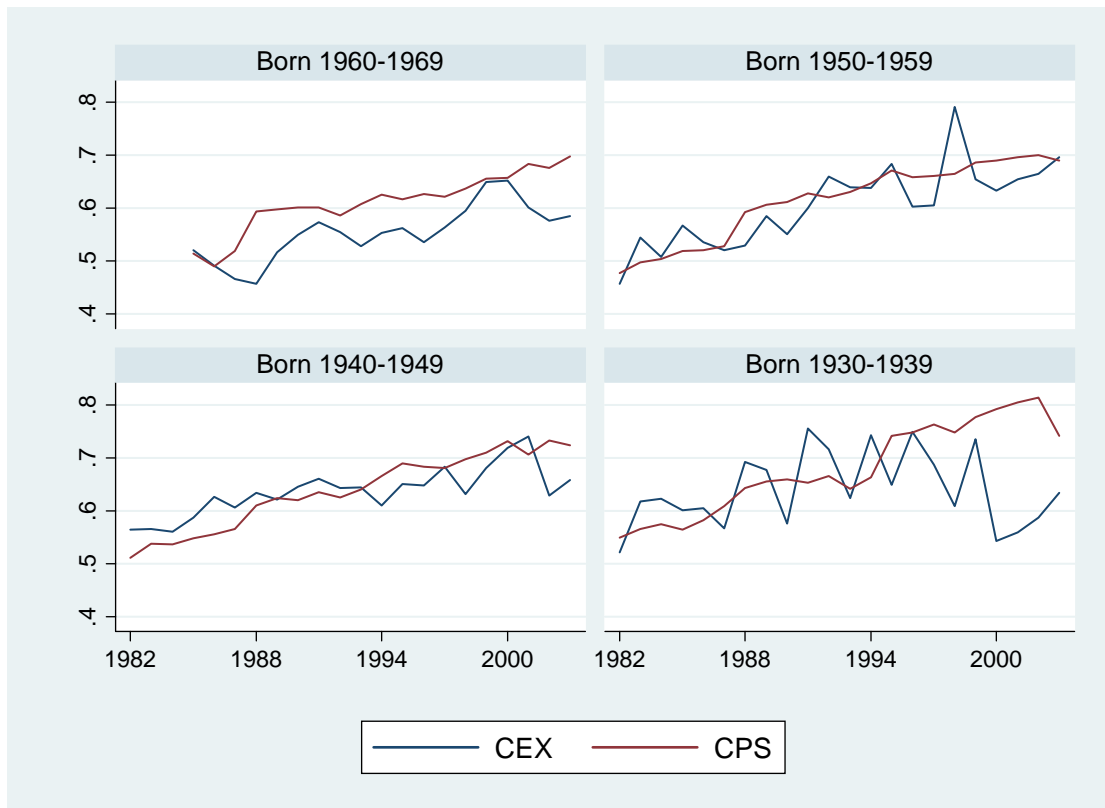


Figure 11
Coefficient of variation of wages
CEX and CPS- by education




Figure 12
CEX and CPS family earnings (1982-1984 dollars)




Figure 13
Median of log family earnings – CEX & CPS by decade of birth cohort




Figure 14
Median of log family earnings – CEX & CPS by education




Figure 15
Coefficient of variation of log family earnings - CEX and CPS




Figure 16
Coefficient of variation of log family earnings - CEX and CPS by cohort




Figure 17
Coefficient of variation of log family earnings - CEX and CPS by education




Figure 18
Inequality trends within the household




Figure 19
Inequality trends within the household, by decade of birth cohort




Figure 20
Inequality trends within the household, by education




Figure 21
Non durable Consumption and Services: levels




Figure 22
Non Durable Consumption and Services levels by cohort




Figure 23
Non Durable Consumption and Services levels by education

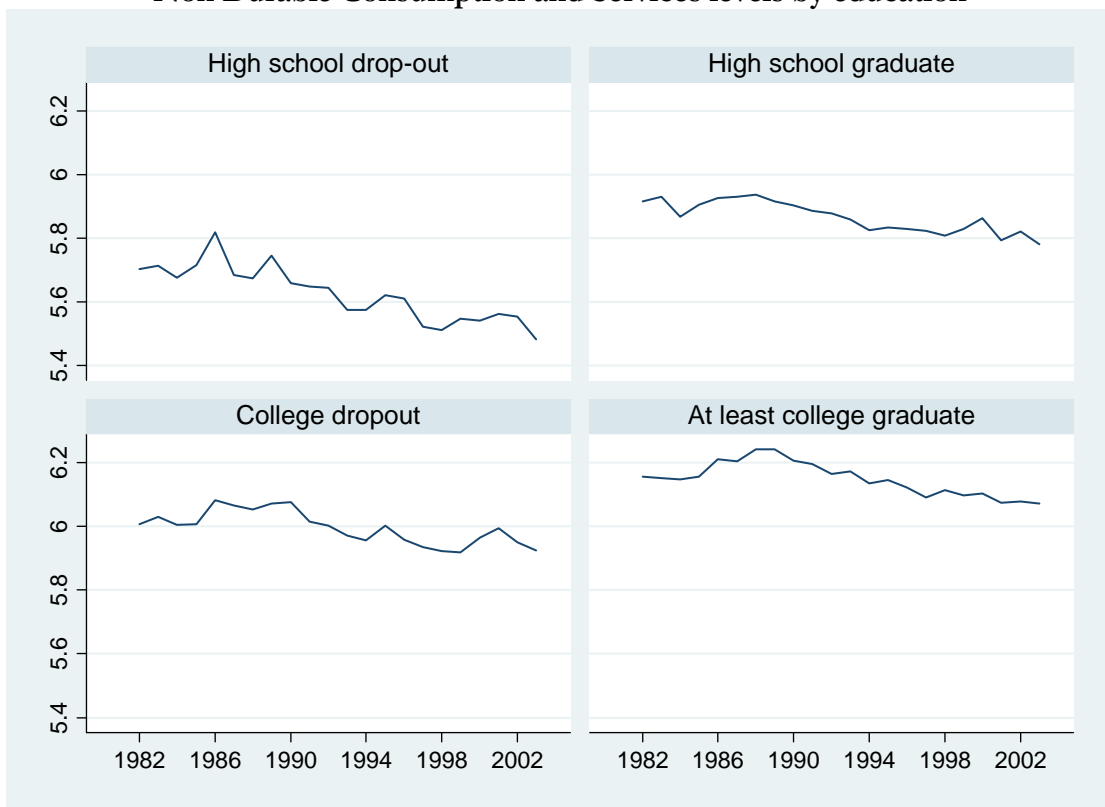


Figure 24

**Relative consumption levels:
Log non durable consumption relative to high school graduates**




Figure 25
Total consumption: levels




Figure 26
Total Consumption levels by cohort




Figure 27
Total Consumption levels by education




Figure 28
Relative total consumption (by education)




Figure 29
Consumption inequality: St. Dev. of log total and non durable




Figure 30
Standard deviation of logs: by decade of birth cohort




Figure 31
Standard deviation of logs: by education




Figure 32
Median consumption (1982-1984 dollars)
Logs




Figure 33
Median consumption for the poor (1982-1984 dollars)




Figure 34
Consumption quantiles (1982-1984 dollars)
Logs

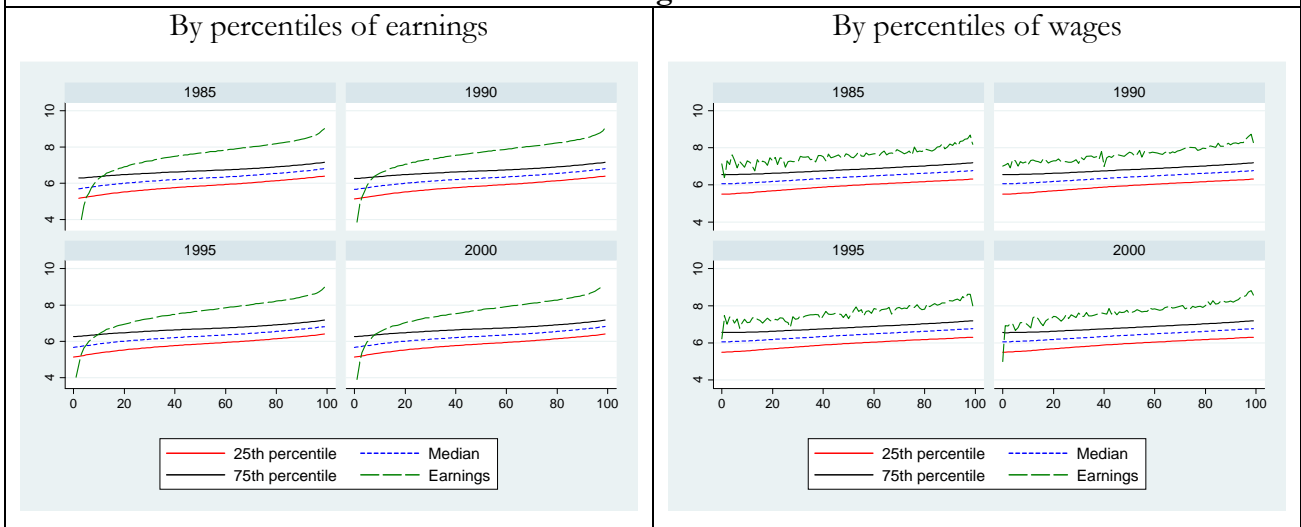


Figure 35
Consumption quintile for the poor (1982-1984 dollars)




Figure 36
Consumption, earnings and wages well-being

