From Wages to Welfare: Decomposing Gains and Losses from Rising Inequality*

Jonathan Heathcote
Federal Reserve Bank of Minneapolis and CEPR

Kjetil Storesletten
Federal Reserve Bank of Minneapolis and CEPR

Giovanni L. Violante
New York University, CEPR, and NBER

FIRST DRAFT: AUGUST 2010 - THIS DRAFT: FEBRUARY 2011

Abstract

This paper offers a critical evaluation of the large literature that studies the welfare consequences of the recent shift in the wage structure in the United States. Welfare calculations based on changes in the empirical distributions of consumption and hours worked – analyzed through the lens of a social welfare function– yield welfare losses on the order of two percent of lifetime consumption. However, two key components of the shift in the wage structure – the growth in the skill premium and the rise in wage volatility – can potentially generate welfare gains as individuals adjust their education and labor supply decisions. Quantifying the importance of these channels of adjustment requires a structural model. In our model-based calculations, under a plausible calibration, we find welfare gains exceeding one percent of lifetime consumption.

*This paper was prepared for Gianluca Violante's invited talk at the Tenth World Congress of the Econometric Society (Shanghai, China). We thank our discussant Steve Davis for many useful comments. We are grateful to Chris Tonetti for outstanding research assistance and to Mark Aguiar, Erich Battistin, and Fatih Guvenen for sharing their data. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

The structure of relative wages in the US economy has undergone a major transformation in the last thirty years. Wage differentials between more and less educated workers have risen sharply (Katz and Autor (1999); Lemieux (2008)). Within narrow groups of workers defined by education, gender, and birth cohort, the distribution of wages has become much more dispersed (Juhn, Murphy, and Pierce (1993)). This increase in within-group dispersion reflects wider fixed individual wage differentials and more pronounced volatility in both persistent and transitory shocks (Gottschalk and Moffitt (1994, 2009)). Overall, the US wage structure has become much more unequal.

This surge in US economic inequality has generated great interest among labor economists and macroeconomists. A vast theoretical and empirical literature investigates the sources of the phenomenon. The leading explanation is that the widespread adoption of new information and communication technologies has raised the relative productivity of more skilled labor – complementary to the new technologies in production– and lowered the demand for less skilled workers employed in tasks easily replaceable by the new machines (Krusell, Ohanian, Ríos-Rull, and Violante (2000); Acemoglu (2002); Autor, Katz, and Kearney (2006); Acemoglu and Autor (2010)).[1] A less prominent role is attributed to falling demand for unskilled-intensive goods produced in the United States because of greater openness to trade and offshoring of unskilled stages of production (Feenstra and Hanson (1996); Burstein and Vogel (2010)). The rise in idiosyncratic volatility is viewed as the result of a more turbulent work environment and faster skill obsolescence (Violante (2002)), changes in wage-compressing labor market institutions such as unions (DiNardo, Fortin, and Lemieux (1996)), and a contractual shift toward performance-based and piece-rate pay (Lemieux, MacLeod, and Parent (2009)).

Much of the underlying motivation for this large body of research is the presumption that observed movements in the relative wage structure have had sharply negative welfare consequences for US households, in particular for the low skilled. There is a growing literature that offers quantitative estimates of these welfare effects. The goal of this paper is to provide a critical survey of this literature, emphasizing the lessons learned and the open questions. A better understanding of the sources and magnitudes of welfare changes should help economists and policymakers in deciding how best to adjust taxation and other redistributive policies in

[1] Some authors argue that the simple model of capital-skill complementarity cannot fully explain the most recent dynamics of wage inequality (in the last decade), when wage differentials widened exclusively at the top of the distribution. See Lemieux (2008) for a discussion.

the face of rising inequality.

The literature that attempts to quantify the welfare implications of the rise in inequality in the United States begins with the pioneering analyses of Cutler and Katz (1991, 1992), Slesnick (1994), and Attanasio and Davis (1996). The premise of this body of work is that, given the numerous insurance channels available to US households to absorb wage movements, inferring how standards of living have changed from wage data alone is inappropriate. Moreover, wider wage dispersion is detrimental to economic welfare to the extent that it transmits into wider consumption dispersion. Therefore, shifts in the distribution of consumption are a better indicator of shifts in the distribution of household welfare. Of course, since high consumption in the presence of low wages may be achieved at the expense of longer hours worked, one should also factor in shifts in the distribution of leisure.

Those in the public policy arena who emphasize potential welfare losses from changes in the wage structure focus on increased income instability and real wage declines at the bottom of the distribution. These trends suggest that more households may now be at risk of falling into poverty.[2] We will emphasize that the crucial determinant of the welfare effect of more volatile wages is whether this volatility is (self-) insurable or not. The uninsurable component of volatility transmits to consumption and reduces household welfare. But households have access to a multiplicity of smoothing channels to absorb wage fluctuations, including offsetting individual and spousal labor supply responses, private transfers within networks of family and friends, borrowing and saving, and government redistribution. Thus, a sizeable component of wage fluctuations is effectively insured.[3] Interestingly, a rise in insurable wage dispersion is not welfare-neutral, but is welfare improving as long as workers can flexibly adjust labor supply in response to wage changes (Heathcote, Storesletten, and Violante (2008)).

Those who emphasize potential welfare gains from changes in the wage structure focus on higher returns to human capital investment.[4] Since wage differentials attributable to education are permanent and ex ante uninsurable, they translate one for one into consumption differentials

[2] For example, Krugman (2005, 2007) wrote: "Over the past three decades the lives of ordinary Americans have become less secure, and their chances of plunging from the middle class into acute poverty ever larger [...] People aren't nearly as much better off as they would be if the gains from economic growth had been broadly distributed."

[3] Recent quantitative studies of the transmission of income shocks to consumption include Blundell and Preston (1998), Blundell, Pistaferri, and Preston (2008), and Kaplan and Violante (2010). For recent surveys of the literature, see Blundell (2010) and Meghir and Pistaferri (2010).

[4] For example, Lazear (2006) wrote that "While there is no doubt that some people have been left behind,[...] the good news is that most of the inequality reflects an increase in returns to investing in skills - workers completing more school, [...] and acquiring new capabilities."

(Attanasio and Davis (1996)). A rising college premium thus leaves college graduates better off and high school graduates worse off. But this argument is incomplete because education is a choice. New cohorts can take advantage of the opportunities presented by skill-biased demand shifts – and the associated larger return to education – by increasing their investment in human capital. This behavioral response can be a source of welfare gains, as demonstrated by Heathcote, Storesletten, and Violante (2010a).

The aim of this paper is to produce quantitative estimates of the overall welfare effects of changes in the wage structure. This requires measuring the relative importance of the forces that induce welfare losses through greater inequality and the forces that induce welfare gains through productivity growth.

1.1 Alternative methodologies

We are interested in the impact of a shift in the wage distribution on welfare, where welfare is a function of consumption and leisure. Thus, linking movements in relative wages to movements in relative consumption and leisure is a crucial step in the analysis. The literature has followed two strategies. The first is to look directly at the shift in the distribution of consumption (and leisure) in the micro data: we call this strategy the "empirical approach."The second is to lay out a structural model to draw a mapping from wages to consumption and leisure: we call this strategy the "structural approach."We will follow both approaches and explain why they yield very different welfare cost estimates.

Welfare calculations based on the empirical approach compare average utility derived from the empirical distribution of consumption and hours worked before and after the shift in the wage structure. Therefore, these calculations have the great virtue of only requiring assumptions on the specification of preferences. We estimate that comparing the distribution in 2001-2005 to the one in 1980-1984 (the earliest date available in the Consumer Expenditure Survey) results in a welfare loss of almost two percent of lifetime consumption in our baseline. This is broadly comparable in magnitude to the calculations of Krueger and Perri (2003).

However, the empirical strategy has one serious drawback. When comparing the two empirical distributions, the data are demeaned. Thus, this methodology abstracts from what we label "level effects" on welfare, i.e., effects on average consumption and leisure of the same forces that underlie the rise in wage dispersion. Skill-biased demand shifts influence output through increased human capital accumulation, while rising wage volatility impacts productivity through modified labor supply decisions. Because these outcomes are the result of individuals' optimal response to exogenous forces, the associated welfare effects can only be quantified in the context of a structural micro-founded model. We apply a version of the partial insurance model with endogenous education, consumption, and labor supply choices developed in Heathcote, Storesletten, and Violante (2009a, 2010b). The advantage of this framework is that one can obtain a closed-form expression for the welfare effects of a change in the structure of relative wages. This expression allows for a transparent quantitative decomposition of all the different sources of gains and losses.

Our key result is that, according to the model, the aforementioned gains dominate the losses arising from increased dispersion and imperfect consumption insurance. Overall, we find a welfare gain from the shift in the wage structure of 1.4 percent of lifetime consumption. Our counterfactual experiments indicate that investment in human capital as a response to the surging skill premium is the largest source of welfare gains.

The rest of the paper is organized as follows. Section 2 sets the stage for our welfare calculations by describing the facts on the changes in inequality in wages, consumption, and hours worked in the United States. Section 3 gives an overview of the empirical approach, its advantages and its limitations, and presents a series of welfare calculations based on this methodology. Section 4 reviews the structural economic model in Heathcote, Storesletten, and Violante (2009a, 2010b), its calibration, and the ensuing model-based welfare calculation. Section 5 contains some concluding remarks on open research questions and a short reflection on public policy.

2 Setting the stage: facts

In this section, we briefly discuss the salient trends in cross-sectional dispersion of wages, hours, and consumption. The appendix contains a detailed description of our two data sources, the March Current Population Survey (CPS) and the Consumer Expenditure Survey (CEX), as well as details on sample construction and variable definitions.

Panel (A) of Figure 1 plots two lines. The solid line is the variance of log wages (w_{it}) for male workers in the United States from 1967 to 2005 computed from the CPS. Wage inequality rises steadily throughout the period.[5]

The dashed line depicts residual (or within-group) wage inequality estimated from the re-

[5] We focus on male wages to avoid selection issues, but Figure 4 in Heathcote, Perri, and Violante (2010) shows that, perhaps surprisingly, the upward trend in log wage inequality is virtually the same for women.

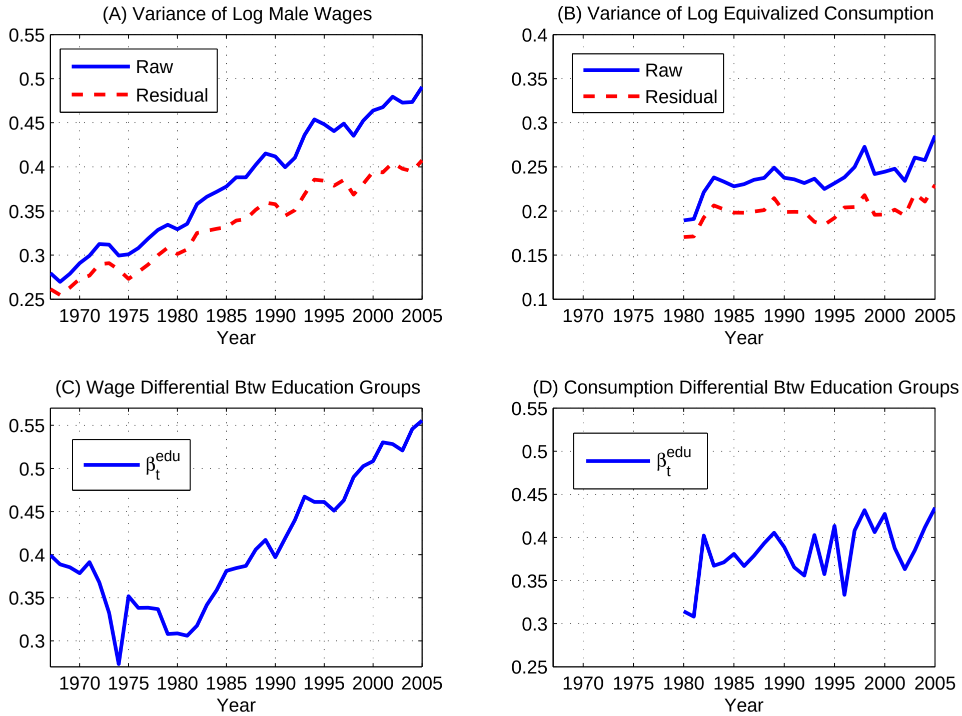

Figure 1: Evolution of inequality in male wages (CPS) and household consumption (Interview Survey of the CEX). Wages are computed as annual earnings (plus two-thirds of self-employment income) divided by annual hours worked. Consumption includes expenditures on nondurables, services, small durables and an imputed flow from vehicles and housing. Consumption is equivalized based on the OECD scale.

gression

$$\log w_{it} = D_t + \beta_t^{edu} D_i^{edu} + f\left(age_{it}; \beta^{age}\right) + \varepsilon_{it}, \tag{1}$$

where D_t is a year dummy, D_i^{edu} an education dummy equal to one if the individual has a college degree, and $f(\cdot)$ is a quartic in age. Residual wage inequality is measured as the variance of ε_{it}. Residual wage dispersion rises steadily over the period. A comparison with the "raw" variance of wages reveals that the within-group component accounts for about two-thirds of the increase in cross-sectional male wage dispersion since 1980. By design, the remaining one-third is explained by the skill premium: panel (C) plots the value of β_t^{edu} and shows the well-known surge in the return to education over this period.

Panels (B) and (D) plot the corresponding variables for equivalized household consumption expenditures (c_{it}). Our baseline measure of consumption includes expenditures on nondurables, services, small durables, and an estimate of the service flow from vehicles and housing. The first noticeable feature of these plots is that, quantitatively, the rise in the variance of log consumption is less than half as large as the rise in the variance of log wages (Slesnick (2001); Krueger and Perri (2006)). Second, the increase in the within-group component of consumption dispersion accounts for a smaller part of the increase compared with wages.[6] Third, education consumption differentials remained roughly two-thirds of education wage differentials throughout the period. To sum up, the wider education wage gap has largely translated into wider consumption dispersion, whereas larger within-group wage volatility had a much milder impact on consumption inequality.[7] Both facts have been emphasized before by Attanasio and Davis (1996) and Krueger and Perri (2003).

Inequality in male and female market hours worked (h_{it}), and its components, are reported in Figure 2. Male hours dispersion is countercyclical but exhibits no obvious long-run trend, whereas female hours dispersion declines significantly. This decline in female hours dispersion toward the level for men reflects the rise in their average hours worked and the fact that more and more women work full-time. Education explains virtually nothing of hours dispersion, as visualized by the substantial overlap of residual and raw variances. This, together with the fact that the education component of the variance remained flat during this period, while the college premium doubled, indicates that income and substitution effects on labor supply roughly offset each other in response to changes in the college premium.

2.1 Some measurement issues

Consumption It is well known that aggregate consumption expenditures computed from the CEX are lower than personal consumption expenditures (PCE) in the National Income and Product Accounts (NIPA) for a wide number of comparable expenditure categories. More disturbingly, the gap between the two series has grown larger over time. For example, for a broad definition of nondurable consumption, the gap grows from 20 percent in 1980 to 60 percent in 2005 (Figure 3 in Heathcote, Perri, and Violante (2010)).[8] This growing discrepancy

[6] We do find though that this within-group component has increased over time, as opposed to Krueger and Perri (2003), who report a decline from 1972 to 2000.

[7] The fact that permanent consumption differentials by education are smaller than permanent income differentials is consistent with an overlapping-generations, incomplete markets model with finite horizon, progressive social security system, and wealth accumulation. See, for example, Storesletten, Telmer, and Yaron (2004) and Kaplan and Violante (2010).

[8] The investigation on the sources of this discrepancy between survey-based and NIPA aggregate consumption is ongoing (Slesnick (2001); Garner, Janini, Passero, Paszkiewicz, and Vendemia (2006)). Conceptual differences between the CEX and the NIPA can account for some of the discrepancy. For example, among medical care expenditures, a rapidly growing item in the NIPA consumption, the Bureau of Economic Analysis includes expenditures by Medicare, Medicaid, and private insurers, whereas the CEX reports only out-of-pocket expenses. However, the growing gap between the CEX and the NIPA applies across a broad range of consumption categories, suggesting that specific definitional differences are only part of the explanation. Another candidate explanation is that the CEX sample misses the upper tail of the income and consumption distributions, and that growth in aggregate consumption has been largely driven by these missing wealthy households.

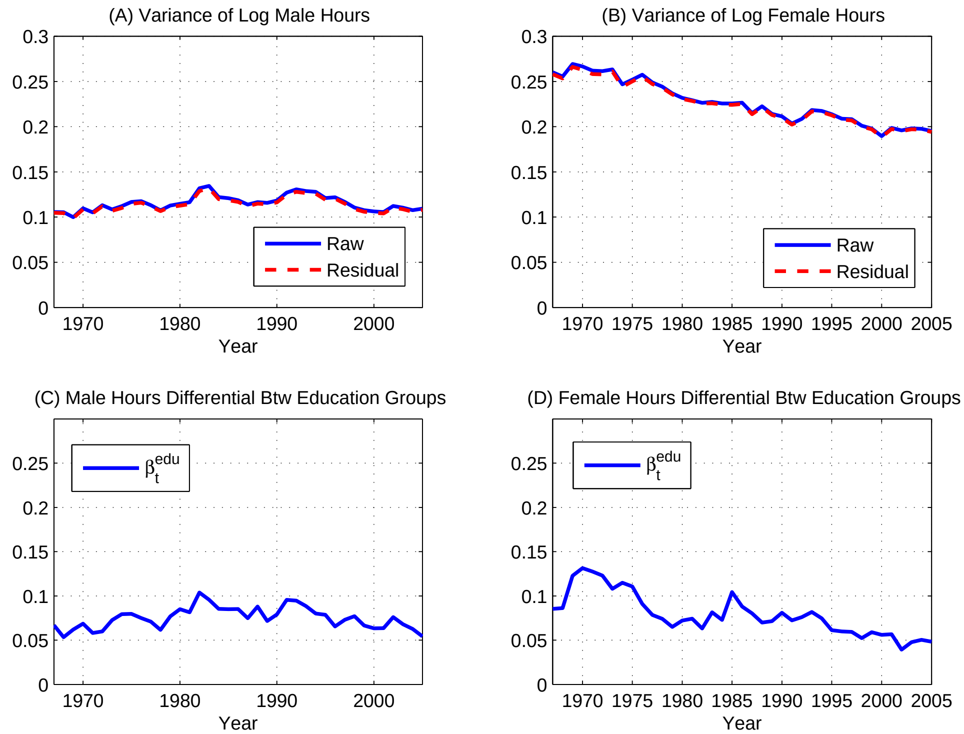

Figure 2: Evolution of inequality in male and female hours worked in the market (CPS).

in means casts some doubt on the measurement of inequality trends as well. A number of studies have investigated the reliability of survey-based consumption inequality statistics by trying to obtain alternative estimates.

Attanasio, Battistin, and Ichimura (2007) note that the Diary Survey (DS) of the CEX is better designed than the Interview Survey (IS) to measure expenditure on frequently purchased goods and services (e.g., food, personal care, housekeeping services). The DS, available only from 1986, shows a rise in consumption inequality larger than that emerging from the IS. Attanasio, Battistin, and Ichimura (2007) and Attanasio, Battistin, and Padula (2010) combine the two surveys by choosing, for each consumption component, the survey reporting expenditures more accurately. Panel (A) in Figure 3 plots the IS-based and the IS-DS combined estimates of the variance of log consumption from Attanasio, Battistin, and Padula (2010). The latter series displays an increase that is almost twice as large over the period 1982–2003, with most of the discrepancy occurring after 1990. In the same figure we also plot the series from Heathcote, Perri, and Violante (2010) that we use in all our baseline calculations. The increase in consumption inequality in this series is comparable to the IS-based series of Attanasio, Battistin, and Padula (2010).

Some authors (e.g., Fisher and Johnson (2006); Blundell, Pistaferri, and Preston (2008);

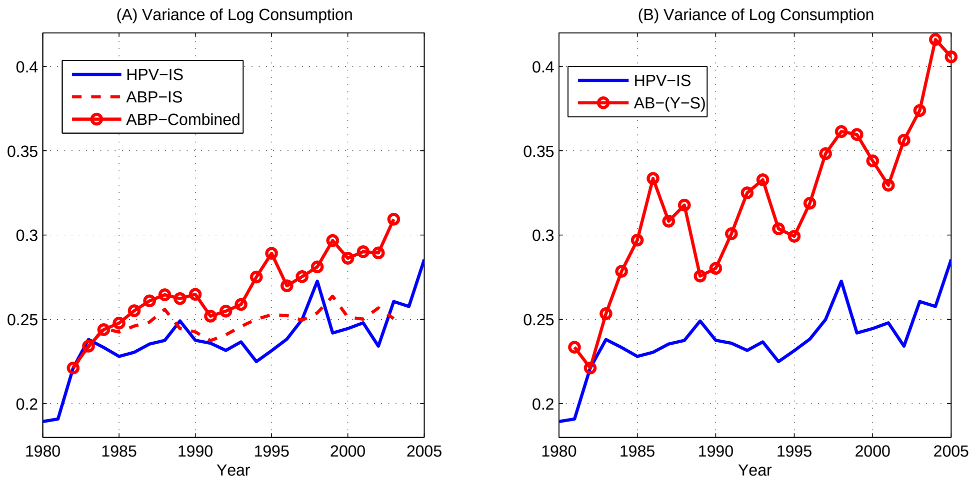

Figure 3: Evolution of inequality in equivalized household consumption (CEX). Panel (A) reports the Attanasio-Battistin-Padula (ABP) estimates obtained combining the Diary and Interview Survey with their Interview Survey estimate and the Interview Survey in Heathcote-Perri-Violante (HPV). Panel (B) plots the HPV series against the series computed by Aguiar and Bils (AB) from disposable income minus reported savings.

Guvenen and Smith Jr. (2010)) have imputed a measure of total consumption for households in the Panel Study of Income Dynamics (PSID) based on the expenditure items common to both the PSID and the CEX (e.g., food and rent) in combination with income and household demographics. In general, they have uncovered an increase in their measure of consumption inequality which is at least as large as that in the CEX Interview Survey.

Aguiar and Bils (2010) exploit the reported amount of active savings and disposable income in the CEX to construct, under a number of assumptions, a measure of consumption residually implied by the household budget constraint. Under this methodology, consumption inequality tracks income inequality closely between 1980 and 2007, showing, once again, a significantly greater increase than is apparent in the IS-based household expenditure data. Panel (B) of Figure 3 shows that the growth in this series is four times larger than in the baseline (and twice as big as the IS-DS combined series computed by Attanasio, Battistin, and Ichimura (2007)).

While these alternative measures all seem to indicate a sharper increase in inequality than was initially found in the CEX Interview Survey, a separate body of evidence would suggest a correction in the opposite direction. All the conventional measures of inequality deflate consumption across individuals by the same price index –a choice akin to assuming that the bundle of goods consumed is not too different across households at any point in time and that all households pay the same price for the same good or service. However, a number of recent papers have challenged this view and documented larger inflation rates for high-income groups (e.g., Broda and Romalis (2009); Broda, Leibtag, and Weinstein (2009)).[9]

More research is necessary to carefully establish the true dynamics of consumption inequality. In the meantime, it is worth exploiting alternative data that are potentially informative about changes in well-being, such as data on hours worked.

Market hours vs leisure Hours worked in the market are correlated with well-being, albeit only imperfectly. Leisure is, theoretically, a better indicator for welfare since it nets out from the time endowment hours spent in the production of home goods as well as market goods. However, measuring dispersion in leisure is much more difficult than measuring dispersion in market hours, because the surveys that collect consistent annual household-level data (such as the CPS and the CEX) lack detailed data on home hours.

From time use surveys, Aguiar and Hurst (2009) exploit some limited information on leisure inequality by measuring the difference in leisure across education groups. They find that between 1985 and 2005, less educated men increased leisure by two percent while more educated men reduced leisure by a similar amount (see their Tables 2, 4A, and 5A). Thus, the distribution of leisure for men does not appear to have changed dramatically –mirroring our finding for market hours. This finding suggests that using male market hours in welfare calculations should not be too misleading.

For women, however, the story is quite different. Knowles (2008) shows that from 1975 to 2003, women increased their hours in the market and reduced their hours worked at home, without changing the fraction of the time endowment devoted to leisure. To the extent that changes in the distribution for female market hours reflect women reducing home hours as they move toward full-time market work, these changes will be a poor proxy for true changes in female welfare. Because of this concern and given the lack of comprehensive information on home work in survey data, we use only male hours in our welfare calculations.

3 The empirical approach

The most direct approach to quantifying the welfare effects of rising wage inequality is that of simply using observations from survey data on the empirical distribution of consumption

[9]Such price differentials between groups reflect differences in bundle composition, differences in quality of the goods and services purchased, and differences in time spent shopping for the same items. Therefore, computing the appropriate correction to inequality measures is a nontrivial task. See also Aguiar and Hurst (2007) and Moretti (2010).

and hours worked, the two key arguments of households' utility.[10] Recently, Jones and Klenow (2010) have used a very similar strategy to assess the historical growth in welfare for a variety of countries and to contrast growth in welfare to growth in GDP, a more traditional measure of growth in well-being.

This approach makes the implicit assumption that all the empirical changes in the dispersion of consumption and hours were driven by the shift in the wage structure. Is this a reasonable assumption? In Heathcote, Storesletten, and Violante (2010a), we build a structural dynamic model of the US economy and estimate it using household survey data. The estimated model, with the observed shift in the wage structure as the only input, reproduces the salient trends in the empirical cross-sectional distributions of individual hours worked, household earnings, and household consumption –all endogenous outcomes of the model.

In the rest of this section, we first describe how different authors have implemented the empirical approach and then report some findings based on our own calculations.

3.1 Implementation

Comparing distributions of allocations, the thrust of this empirical strategy, requires only a minimal set of assumptions. Consider an overlapping-generations economy with a fixed demographic structure, in which the total population is constant and the share of the population of age $j \in \{0, 1, ..., J\}$ is date-invariant and equal to s_j. Suppose the wage distribution is constant and that there is an associated stationary joint distribution over consumption and hours. Let U_j be the implied expected remaining lifetime utility for a j years-old individual, discounted to her birth date:

$$U_j = \sum_{t=j}^{J} \beta^t \frac{s_t}{s_j} \mathbb{E}\left[u\left(c_t, h_t\right)\right], \tag{2}$$

where β is the discount factor and s_{j+1}/s_j is the survival probability between age j and $j+1$.

Define the following Benthamite social welfare function to aggregate utilities across all cohorts currently alive and as yet unborn:

$$\mathcal{W} = \sum_{j=0}^{J} \mu_j s_j U_j + \sum_{j=-\infty}^{-1} \mu_j s_0 U_0, \tag{3}$$

where μ_j is the planner's weight on individuals of age j.[11]

[10] See Slesnick (1998) for a survey on the empirical approach to the measurement of welfare.

[11] The convention here is that $j < 0$ indicates an as-yet-unborn generation, so, for example, μ_{-1} denotes the weight on the generation that will enter the economy in the next year.

In general, comparing welfare across different wage structures requires estimating the distribution of lifetime sequences of consumption and hours before and after the shift in the wage structure. Krueger and Perri (2003) exploit the short panel dimension of CEX (one year) and estimate a finite state Markov chain for log consumption and log hours where the transition probabilities across quantiles are time-invariant, but quantiles are allowed to vary over time to reflect the movements in cross-sectional dispersion. As emphasized by Davis (2003) and Storesletten (2003), a shortcoming of this approach is that the estimated persistence of consumption and hours worked –and hence the estimate of the welfare cost– is likely to be mismeasured because of the extremely short panel dimension and because of the large measurement error known to plague reports of hours worked and expenditures in household surveys.[12] Thus, the lack of high-quality longitudinal data on consumption in the CEX undermines the estimation of a household-level stochastic process.

Attanasio and Davis (1996) chose to circumvent this problem by focusing on the relative movements of wages and consumption across observationally distinct groups. This choice allows the simultaneous use of the best survey data for consumption (CEX) and the best survey data for income (CPS). Their key finding is that persistent changes in relative wages among birth cohort-education groups lead to roughly equal-size changes in the distribution of consumption expenditures. Put differently, the rise in the skill premium translated almost fully into consumption differentials between more and less educated households. A drawback of this methodology is that it abstracts from changes in the within-group component of wage dispersion that, as shown in Figure 1, are large.

There is a third way to deal with this issue that allows avoiding the estimation of a stochastic process while, at the same time, retaining within-group variation. It requires a particular choice for the cohort-specific weights μ_j in the social welfare function (3), namely, the weights where the planner puts the same weight on the expected felicity of all agents alive, and discounts future felicities with the same discount factor as that used by the agents. These weights are

[12]For example, Cogley (2002) suggests that measurement error in CEX consumption biases upward the true variance in individual consumption growth by one order of magnitude. Similarly, Heathcote, Perri, and Violante (2010) find that measurement error accounts for as much as one-fourth of the total variance of log consumption and, clearly, a much bigger share of the within-group component.

given by $\mu_j = \beta^{-j}$. In this case the social welfare function is[13]

$$\begin{aligned} \mathcal{W} &= \sum_{j=0}^{J} s_j \beta^{-j} \sum_{t=j}^{J} \beta^t \frac{s_t}{s_j} \mathbb{E}\left[u\left(c_t, h_t\right)\right] + \sum_{j=-\infty}^{-1} s_0 \beta^{-j} \sum_{t=0}^{J} \beta^t \frac{s_t}{s_0} \mathbb{E}\left[u\left(c_t, h_t\right)\right] \\ &= \frac{1}{1-\beta} \sum_{j=0}^{J} s_j \mathbb{E}\left[u\left(c_j, h_j\right)\right]. \end{aligned} \tag{4}$$

Thus, given this weighting scheme and a stationary joint distribution over consumption and hours, social welfare is proportional to the cross-sectional average current felicity. This is a very convenient property from the perspective of empirical implementation, because welfare effects from a shift in the wage structure can be estimated from the cross-sectional joint distribution of consumption and hours, without any information on individual dynamics – a data requirement that is much less demanding.

Let "$*$" denote the stationary distribution over consumption and hours before the shift of the wage structure, and let "$**$" denote the one after the shift. We are interested in comparing an economy with the "$*$" allocation to one with the "$**$" allocation. Let $\mathbb{E}^*$ and $\mathbb{E}^{**}$ denote expectations with respect to these two distributions. Then, the average welfare effect of rising inequality is defined as the scalar ω that solves

$$\sum_{j=0}^{J} s_j \mathbb{E}^* \left[u\left((1+\omega) c_j, h_j\right)\right] = \sum_{j=0}^{J} s_j \mathbb{E}^{**} \left[u\left(c_j, h_j\right)\right] \tag{5}$$

A negative value for ω represents the fraction of consumption an individual would be willing to give up, in each state at each date, in order to avoid the shift in the distribution of consumption and hours induced by the new wage structure.

A specification for period utility $u(\cdot)$ must be chosen to operationalize this calculation. This is the only model ingredient needed. In particular, since this approach does not try to draw a mapping between wages on the one hand and consumption and hours on the other, no assumptions have to be made on market structure, risk-sharing possibilities, technology, or agent's choice sets. In what follows, we assume the intra-period utility function

$$u(c, h) = \frac{c^{1-\gamma}}{1-\gamma} - \exp(\bar{\varphi}) \frac{h^{1+\sigma}}{1+\sigma}, \tag{6}$$

[13]This particular planner objective function and the convention to discount utility for all agents back to their respective birth date closely follow Calvo and Obstfeld (1988). An equivalent way of writing the planner's objective function would be to discount all utilities to the current date, assuming that as-yet-unborn agents' utilities are discounted to today using the annual discount factor β, and attaching the same planner weight to every generation – both the currently living and the unborn.

which has the advantage of being defined over consumption and hours, thus avoiding the problems in the measurement of leisure discussed above. The parameter γ is the inverse of the intertemporal elasticity of substitution for consumption. The parameter σ captures aversion toward hours fluctuations, and $1/\sigma$ measures the Frisch elasticity of labor supply. The preference weight $\bar{\varphi}$ captures the strength of an individual's distaste for work relative to her preference for consumption.[14]

In directly comparing two distributions of consumption and hours at two different points in time, one has to confront the fact that aggregate consumption growth makes the final allocation a better one and that this first-order effect is likely to dominate changes in second moments that occurred during the same period. Authors have dealt with this issue by demeaning the data (or equivalently, rescaling the final distribution so that it has the same mean as the initial one). However, demeaning by definition eliminates any potential level effects (i.e., changes in the aggregate level of consumption and leisure) induced by the same forces that change wage dispersion. Only through the lens of a model can one identify and measure these level effects. We return to this essential point in Section 4, when we discuss the "structural approach."

3.2 Results

We now put the empirical approach to work in order to quantify the welfare change associated with the shift in the US wage structure. We choose the first and last five years (1980-1984 and 2001-2005, respectively) available in our CEX data to measure the joint distribution of consumption and hours worked before and after the shift (i.e., the "$*$" and the "$**$" allocations, respectively). We rescale the distributions of consumption and hours in 2001-2005 so that they have the same mean as in 1980-1984.[15] We present three alternative implementations of the empirical approach.

An Atkinson-style calculation In the spirit of Atkinson (1970) and Storesletten (2003), in our first calculation we use the actual realizations of consumption and hours observed in the CEX. Let I^* be the number of individuals in the 1980-1984 surveys and I^{**} be the number in the 2001-2005 surveys. Given the social welfare function (4) and the utility function (6), the

[14]In Heathcote, Storesletten, and Violante (2008) we used also a Cobb-Douglas specification for some related welfare calculations. The advantage of the separable specification in (6) over Cobb-Douglas is that two distinct parameters (γ, σ) regulate the two key elasticities. The disadvantage is that the calibration of the weight $\bar{\varphi}$, in a model with heterogeneity like ours, is not straightforward.

[15]To minimize the effect of outliers, we trim the top and bottom 0.5 percent of the consumption and hours distributions.

Table 1: Atkinson-Style Welfare Calculation

	% Consumption Equivalent Variation (ω)				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-1.7	-3.7	-5.3	-6.2	-6.4
$\sigma = 2$	-1.8	-3.8	-5.4	-6.3	-6.4
$\sigma = 3$	-1.9	-3.8	-5.4	-6.3	-6.4
$\sigma = 4$	-1.9	-3.9	-5.5	-6.4	-6.5
$\sigma = 5$	-1.9	-3.9	-5.5	-6.4	-6.5

empirical counterpart of equation (5) becomes

$$\frac{1}{I^*}\sum_{i=1}^{I^*}\left\{\frac{\left[(1+\omega)\,c_i^*\right]^{1-\gamma}}{1-\gamma}-\exp\left(\bar{\varphi}\right)\frac{\left(h_i^*\right)^{1+\sigma}}{1+\sigma}\right\}=\frac{1}{I^{**}}\sum_{i=1}^{I^{**}}\left\{\frac{\left(c_i^{**}\right)^{1-\gamma}}{1-\gamma}-\exp\left(\bar{\varphi}\right)\frac{\left(h_i^{**}\right)^{1+\sigma}}{1+\sigma}\right\}. \quad (7)$$

Table 1 reports the values of ω that solve equation (7) for different levels of risk aversion (γ) and Frisch elasticity ($1/\sigma$).[16] In the range $\gamma = 1, ..., 5$ and $\sigma = 1, ..., 5$ the welfare losses from the shift in the wage structure vary between -1.7 percent and -6.5 percent of lifetime consumption, in line with the findings of Krueger and Perri (2003). As expected, welfare losses increase steeply in γ. The slope with respect to σ is much flatter because, as displayed in Figure 1, the variance of male log hours is basically constant over time.

The message of these calculations is that, according to the empirical approach, welfare losses are large. To put these estimates in perspective, recall that the Lucas (1987) seminal calculation of welfare gains from eliminating business cycles in a representative agent economy with log utility is 0.008 percent. More recently, Krusell and Smith Jr. (1999) and Krusell, Mukoyama, Sahin, and Smith Jr. (2009) revisited this calculation in an incomplete-markets model with idiosyncratic income risk correlated with aggregate risk and report that the average welfare gain from eliminating cycles is around 0.1 percent of consumption. Taken together, these calculations reveal that US households would be willing to pay almost twenty times more to avoid a rise in wage inequality similar to the one witnessed over the last 30 years than they would pay to eliminate business cycles.

An Attanasio and Davis-style calculation In the spirit of Attanasio and Davis (1996), we perform an alternative exercise. We group individuals by education level (with and without

[16] To calibrate $\bar{\varphi}$, we proceed as follows. Allow the weight on labor effort in the utility function (6) to be individual specific, $\bar{\varphi}_i$. Given a pair (γ, σ) and the externally calibrated tax rates, we use data on (c_{it}, h_{it}, w_{it}) from each individual CEX record to back out residually a value $\bar{\varphi}_i$ for each individual so that his/her first-order condition holds with equality. From the implied distribution of $\bar{\varphi}_i$, we estimate the median and use it in the welfare calculations in this section.

Table 2: Attanasio and Davis-Style Welfare Calculation

	% Consumption Equivalent Variation (ω)				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-0.8	-1.7	-2.4	-3.1	-3.6
$\sigma = 2$	-0.9	-1.7	-2.5	-3.1	-3.7
$\sigma = 3$	-0.9	-1.7	-2.5	-3.2	-3.8
$\sigma = 4$	-0.9	-1.8	-2.6	-3.3	-3.9
$\sigma = 5$	-0.9	-1.8	-2.6	-3.3	-3.9

a college degree) and by age (25-34, 35-44, 45-54, 55+). In Table 2, we repeat the welfare calculations of Table 1 by using the eight age-education groups, appropriately weighted, instead of the individuals as the unit of analysis: the implied welfare effect reported in Table 2 is determined only by the shift in the distribution of consumption and hours between groups, but it abstracts from the change in the within-group component.

We find that welfare losses are roughly half of those in Table 1. For example, in the parameterization $\gamma = 1$ and $\sigma = 2$, ω equals -0.9 percent instead of -1.8 percent. This finding is consistent with the fact documented in Figure 1, namely, that the rise in between-group consumption inequality is approximately half of the total.

A Lucas-style calculation Lucas (1987) showed that by assuming lognormality of the stochastic process for aggregate consumption, one can derive an analytical and intuitive expression for the welfare cost of business cycles. Following our work in Heathcote, Storesletten, and Violante (2008), we apply this approach to the cross section. Let v_x denote the cross-sectional variance of the random variable x and suppose that log consumption and log hours are distributed as

$$\begin{aligned}\ln c_i^* &\sim N\left(\mu_c^* - v_c^*/2, v_c^*\right)\\ \ln h_i^* &\sim N\left(\mu_h^* - v_h^*/2, v_h^*\right)\end{aligned}$$

before the shift in the wage structure and as

$$\begin{aligned}\ln c_i^{**} &\sim N\left(\mu_c^* - v_c^{**}/2, v_c^{**}\right)\\ \ln h_i^{**} &\sim N\left(\mu_h^* - v_h^{**}/2, v_h^{**}\right)\end{aligned}$$

Table 3: Lucas-style Welfare Calculation

	% Consumption Equivalent Variation (ω)				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-1.8	-3.7	-5.6	-7.6	-9.5
$\sigma = 2$	-1.8	-3.7	-5.6	-7.6	-9.5
$\sigma = 3$	-1.8	-3.7	-5.6	-7.6	-9.5
$\sigma = 4$	-1.8	-3.7	-5.6	-7.6	-9.5
$\sigma = 5$	-1.8	-3.7	-5.6	-7.6	-9.5

after the shift.[17] Equation (5) then yields

$$(1+\omega)^{1-\gamma}\frac{1}{1-\gamma}\exp\left((1-\gamma)\mu_c^* - \gamma(1-\gamma)\frac{v_c^*}{2}\right) - \frac{\exp(\bar{\varphi})}{1+\sigma}\exp\left((1+\sigma)\mu_h^* + \sigma(1+\sigma)\frac{v_h^*}{2}\right)$$
$$= \frac{1}{1-\gamma}\exp\left((1-\gamma)\mu_c^* - \gamma(1-\gamma)\frac{v_c^*}{2}\right) - \frac{\exp(\bar{\varphi})}{1+\sigma}\exp\left((1+\sigma)\mu_h^* + \sigma(1+\sigma)\frac{v_h^{**}}{2}\right). \tag{8}$$

Table 3 reports estimates of ω based on this Lucas-style approach. Reassuringly, the estimates of ω are very similar to those in Table 1, especially for moderate utility curvature on consumption.

As mentioned in Section 2.1, the Interview Survey of the CEX may underestimate the rise in consumption inequality. Using Lucas' approach, one can easily analyze how this reassessment would affect the calculations in Table 3. Since the hours component has a trivial effect under the lognormality assumption, in what follows we ignore it. Abstracting from the hours component, and using a first-order Taylor expansion of type $\exp(x) \approx 1 + x$, from (8) we arrive at the cross-sectional counterpart of Lucas' expression:

$$\omega \approx -\frac{\gamma}{2}\Delta v_c, \tag{9}$$

where $\Delta v_c = v_c^{**} - v_c^*$ is the change in the cross-sectional dispersion of consumption between 1980-1984 and 2001-2005. Equation (9) reveals that, if the true rise in the variance of log consumption is twice as large as in our baseline case, the welfare loss would increase proportionately. For example, in the log case ($\gamma = 1$) the consumption equivalent variation ω would be -3.6 percent of lifetime consumption instead of -1.8 percent.

3.3 What do we learn from the empirical approach?

The greatest advantage of the empirical approach is that it requires only a minimal set of assumptions on preferences and aggregation of individuals into a welfare function. In particular,

[17] The assumption of lognormality of consumption in the cross section is supported by Battistin, Blundell, and Lewbel (2009).

no assumption on behavior or market structure is required.[18] Its main drawback is that it is unable, by design, to assess the impact on aggregate consumption and leisure (i.e., the level effects) of those same forces that triggered the shift in wage dispersion.

Different implementations of the empirical approach led to fairly consistent results: in the baseline parameterization ($\gamma = 1$), the average welfare loss from the rise in wage inequality is around two percent of lifetime consumption.

4 The structural approach

The main reason to adopt a structural approach to welfare calculations is that we wish to quantify the level effects that changes in the distribution of relative wages have on average hours and consumption. Skill-biased demand shifts influence average productivity because they trigger changes in human capital accumulation, while larger wage volatility affects productivity through modified labor supply decisions.

To make progress, we need to select a model that offers a mapping between exogenous shifts in the distribution of wages and changes in the equilibrium distribution of consumption and hours worked. We apply the "partial insurance" framework developed in Heathcote, Storesletten, and Violante (2009a, 2010b). It is an incomplete-markets model featuring the key channels through which individuals can respond to shifts in the wage distribution: education, labor supply, participation in government redistribution schemes, self-insurance, and other forms of private risk sharing. In particular, beyond saving and borrowing through a risk-free asset, agents in our economy are able to perfectly insure a subset of idiosyncratic wage risk. This additional insurance is designed to capture a number of other adjustment mechanisms (e.g., spousal labor supply, networks of family and relatives), financial instruments (e.g., hedging sectoral or occupational risk), and institutions (e.g., bankruptcy) that spread risks across individuals or over time. As explained in detail in Heathcote, Storesletten, and Violante (2009a, 2010b), in the spirit of Deaton (1997), we do not model these mechanisms explicitly, but we bundle them together and quantify their overall importance by looking at the residual gap between wage and consumption dispersion, once all the other explicitly modeled smoothing channels are taken into account.

A significant strength of this framework is a degree of tractability lacking in standard

[18] Strictly speaking, this is true only conditional on knowing the value of the relative disutility of hours $\bar{\varphi}$. The calibration of $\bar{\varphi}$ requires assuming that individuals are on their intratemporal optimality condition.

incomplete-market models. Equilibrium allocations for consumption and hours can be obtained in closed form and, as a result, one can solve analytically for the welfare change associated with shifts in the wage structure. These analytical expressions for welfare reveal all the sources of gains and losses as a function of structural parameters and can easily be related to the Lucas-style welfare expression (9).

4.1 A model economy with partial insurance

We focus on a comparison across two steady states. Thus, in describing the environment, we invoke the steady-state assumption and drop time subscripts. We now review the model of Heathcote, Storesletten, and Violante (2010b), which in turn builds on the framework developed in Heathcote, Storesletten, and Violante (2009a). We offer a brief description here and refer the interested reader to our previous papers for details and analysis.

Demographics Time is discrete and continues forever. The demographic structure follows Yaari's model of perpetual youth: agents are born at age zero and survive from age j to age $j+1$ with constant probability $\pi < 1$. A new generation with measure $(1-\pi)$ enters the economy each period. Thus, the measure of agents of age j is $(1-\pi)\pi^j$ and the total population size is unity.

Preferences The expected lifetime utility for agent i is given by

$$\mathbb{E}_0 \sum_{j=0}^{\infty} (\beta\pi)^j \, u_i(c_{ij}, h_{ij}), \tag{10}$$

where the expectation is taken over sequences of shocks defined below. Here, c_{ij} denotes consumption at age j and h_{ij} hours worked. Agents discount the future at rate $\beta\pi$, where $\beta < 1$ is the discount factor. Period utility for individual i is

$$u_i(c_{ij}, h_{ij}) = \log c_{ij} - \exp(\bar{\varphi} + \varphi_i) \frac{h_{ij}^{1+\sigma}}{1+\sigma}, \tag{11}$$

a specification consistent with both balanced long-run growth and the evidence on hours inequality in Figure 2. The disutility weight on hours worked has a common component $\bar{\varphi}$ and an idiosyncratic fixed component φ_i that is drawn once at the start of an agent's lifetime from the Normal distribution $\varphi_i \sim N\left(-\frac{v_\varphi}{2}, v_\varphi\right)$.[19] Note that this is the same preference specification assumed in Section 3.1 for all our computations based on the empirical approach, except for

[19] Given the distributional assumption, the average weight on hours in preferences is given by $\mathbb{E}[\exp(\bar{\varphi} + \varphi_i)] = \exp(\bar{\varphi})$.

the presence of dispersion in φ_i across agents. Preference dispersion is intended to capture various sources of heterogeneity that generate cross-sectional variation in hours worked and consumption that is independent of variation in productivity. While preference heterogeneity is important for the estimation of the model that we undertake in Heathcote, Storesletten, and Violante (2009a), such heterogeneity plays no role in the welfare calculations that we do in this article.[20]

Education Two possible schooling levels are attainable by individuals: high ($s = H$), corresponding to at least a college degree, and low ($s = L$), corresponding to lower levels of education. The simple model for education acquisition follows Heathcote, Storesletten, and Violante (2010a).[21] When they first enter the economy, before drawing their disutility of work φ_i, agents draw a utility cost of attending college χ_i from a lognormal distribution with mean μ_χ and variance v_χ. Taking prices as given, individuals attend college if the expected lifetime utility upon entry in the labor market as college graduates, $\mathcal{W}_H^0$, net of the education cost, χ_i, exceeds the expected lifetime utility without a college degree, $\mathcal{W}_L^0$.

Production The final good can be used for private consumption (C) and government consumption (G). Therefore, the aggregate resource constraint of this closed economy is, simply, $C + G = Y$. The aggregate production technology is constant returns to scale with college and high school labor as the only inputs. Following a large literature (e.g., Katz and Murphy (1992); Heckman, Lochner, and Taber (1998)), this technology is assumed to have the constant elasticity of substitution form

$$Y = \exp(z)\left[\zeta N_H^{\frac{\theta-1}{\theta}} + (1-\zeta)N_L^{\frac{\theta-1}{\theta}}\right]^{\frac{\theta}{\theta-1}}, \tag{12}$$

where N_H and N_L are aggregate effective hours (productivity times hours worked) for college and high-school-educated workers, and where θ is the elasticity of substitution between the two labor inputs. The weight parameter ζ is allowed to vary across steady states, reflecting skill-biased demand shifts, while the elasticity of substitution parameter θ is fixed.

[20] As shown in Heathcote, Storesletten, and Violante (2009a), matching both the wage-hour covariance and the consumption-hour covariance requires substantial cross-sectional preference dispersion in the relative weight on leisure. The reason why preference heterogeneity does not influence the welfare effects of changes in the wage structure is that terms involving $\bar{\varphi}$ and v_φ enter individuals' equilibrium lifetime utility in an additively separable fashion, and do not interact with terms involving individual labor productivity or wage structure parameters. It follows these terms drop out of the equation determining the welfare gain from a change in the wage structure (eq. (22) below.)

[21] See also Guvenen and Kuruscu (2009) for an analysis of the trends in inequality based on a model with endogenous acquisition of human capital.

Since labor markets are competitive, the price P_s of an effective hour worked by an individual with schooling level s is the marginal product from the technology described above, or, in logs:

$$\log P_s \equiv p_s = z + \frac{1}{\theta - 1}\log\left[\zeta N_H^{\frac{\theta-1}{\theta}} + (1-\zeta)N_L^{\frac{\theta-1}{\theta}}\right] + \log\zeta - \frac{1}{\theta}\log\left(N_s\right).$$

Thus, the equilibrium log skill premium is

$$p_H - p_L = \log\left(\frac{\zeta}{1-\zeta}\right) - \frac{1}{\theta}\log\left(\frac{N_H}{N_L}\right). \tag{13}$$

Therefore, the model skill premium varies across steady states due to both skill-biased demand shifts (i.e., an increase in ζ) and induced increases in college graduation (which translate into a rise in the ratio N_H/N_L).

Individual productivity shocks Individual hourly wages are equal to individual labor productivity (units of effective labor input per hour worked) times the price per effective hour worked of the particular skill (i.e., the education level) the individual supplies: $w_{isj} = P_s \exp(\alpha_{ij} + \varepsilon_{ij})$. The terms α_{ij} and ε_{ij} are stochastic components of the wage that are additive in logs, orthogonal to each other, and orthogonal to the education component. These stochastic components are assumed to follow the same process across both education groups.[22] The component α_{ij} follows the random walk process

$$\alpha_{ij} = \alpha_{i,j-1} + \eta_{ij},$$

where the innovation η_{ij} is drawn from the time-invariant (within a particular steady state) Normal distribution with variance v_η. Agents entering the labor market at age $j = 0$, after the education decision, draw initial realizations α_{i0} from a Normal distribution with cohort-specific variance v_{α_0}. We assume that ε_{ij} is a purely transitory shock, i.e., i.i.d. over time with variance v_ε.[23] The statistical process for wages described above (unit root plus i.i.d. shocks) is quite standard in the literature and is consistent with the key features of individual wage dynamics as well as with trends in wage dispersion across the life cycle.[24] Finally, we normalize the means of the distributions for α_0, ε, and η to negative one-half their respective variances, which ensures that the average wage of type s workers is given by $\exp(p_s)$.

[22] Meghir and Pistaferri (2004) estimate earnings dynamics separately for three educational groups and do not find large differences among them.

[23] In Heathcote, Storesletten, and Violante (2009a), we assume a richer process for ε, i.e., we let ε comprise a permanent shock and a transitory shock. Given our social welfare function (4), the dynamic properties of ε do not matter for the welfare calculations, so we make this simplifying assumption.

[24] For example, the empirical autocovariance function for individual wages displays a sharp decline at the first lag, indicating the presence of a transitory component in wages. At the same time, within-cohort wage dispersion increases approximately linearly with age, suggesting the presence of permanent shocks.

For the initial steady state, we make the innocuous normalization that the average wage in the population is equal to one:

$$e^* P_H^* + (1 - e^*) P_L^* = 1, \tag{14}$$

where e^* denotes the initial steady-state college graduation rate. How to normalize the wage levels in the final steady state is not obvious. In the spirit of Heathcote, Storesletten, and Violante (2010a), we assume that absent any change in college attendance, there would have been no change in average wages between the initial steady state and the final steady state. Thus,

$$e^* P_H^{**} + (1 - e^*) P_L^{**} = 1. \tag{15}$$

This normalization is symmetric to the conditions imposed on the stochastic components for wages, where we always assumed that changes in the wage structure leave the average wage unchanged and equal to one. With this set of normalizations, we do not hardwire any level effects into the model: all changes in average consumption, hours, and productivity between steady states arise due to behavioral responses in the form of modified labor supply or education choices. Since wages equal marginal products, assumptions (14) and (15) can be reinterpreted as restrictions on the parameters z^* and z^{**} in the aggregate production technology (12).

Government The government consumes a fraction of output g at every date and finances these expenditures through a progressive tax system. If we let pre-government earnings of individual i be $y_i = w_i h_i$, then disposable earnings are

$$\tilde{y}_i = \lambda \left(w_i h_i\right)^{1-\tau} \tag{16}$$

and the government budget constraint is

$$gY = \int \lambda \left(w_i h_i\right)^{1-\tau} di. \tag{17}$$

This class of progressive tax functions, discussed in detail by Benabou (2002) and Heathcote, Storesletten, and Violante (2009a, 2010b) is indexed by two parameters. The parameter $\tau \geq 0$ measures the degree of progressivity (with $\tau = 0$ representing a linear tax system with tax rate $1 - \lambda$), while λ captures the overall level of taxation and is determined residually from equation (17) assuming that the government balances its budget every period in equilibrium. Note that the system features marginal tax rates that are increasing in earnings and always generates a transfer (or a negative average income tax) for earnings below the threshold $\lambda^{\frac{1}{\tau}}$.

When we compute the welfare effects associated with changes in the wage structure, we will hold constant g and τ, and adjust λ appropriately to keep the budget balanced.

Market structure Perfect annuity markets are available to insure against the risk of survival. Moreover, agents can trade a risk-free bond and claims contingent on ε, both in zero net supply.[25] Finally, we assume that at birth each agent is endowed with zero financial wealth.

4.2 Stationary equilibrium

No bond-trade equilibrium The equilibrium of this economy, derived and fully characterized in a more general framework in Heathcote, Storesletten, and Violante (2009a), has its roots in Constantinides and Duffie (1996).[26] The distinctive property of the equilibrium is that agents choose not to trade the risk-free bond: in equilibrium, the expected marginal rate of substitution between consumption at dates t and $t+1$ is the same for all households. Individuals in our model have two saving motives: an intertemporal motive given by the gap between the degree of patience and the interest rate, and a precautionary motive reflecting the variance of permanent shocks to the α component (recall that the ε shock is insurable). The intertemporal motive is the same across agents because all agents have the same discount factor and face a common economy-wide risk-free interest rate r^*. The precautionary motive is also identical, as a result of the assumptions that shocks are multiplicative, permanent, and drawn from common distributions, that preferences are in the power utility class with respect to both consumption and hours worked, and that all individuals start out with zero wealth. Because the strength of these two saving motives is identical, there exists an economy-wide interest rate r^* at which, in equilibrium, the (negative) intertemporal motive exactly offsets the (positive) precautionary motive, and no agent wants to either borrow or lend.

Equilibrium allocations The key reason why our Heathcote, Storesletten, and Violante (2009a) model is tractable is that the vector $(\varphi, s, \alpha, \varepsilon)$ contains sufficient information to determine individual wealth. Hence, wealth is not a separate state variable and the vector $(\varphi, s, \alpha, \varepsilon)$ fully describes equilibrium choices. The power of this result lies in the fact that the entries of that vector are all exogenous state variables. Therefore, it is possible to derive closed-form expressions for consumption, hours worked, and welfare, which is what we focus on in this paper.[27]

[25] The set of contingent claims is meant to capture the additional private risk sharing discussed earlier. Instead of implementing full insurance through markets, we could have chosen to allow full consumption smoothing with respect to ε within each extended family of individuals, for example, with no impact on equilibrium allocations.

[26] See also Krebs (2003) for a similar setup.

[27] Standard incomplete-market economies do not admit an analytical solution, and numerical methods are

In Heathcote, Storesletten, and Violante (2010b), we show that the equilibrium allocations for individual consumption and hours in logs are given by

$$\begin{aligned}\log c(\varphi, s, \alpha) &= \kappa_c + (1-\tau)(p_s + \alpha) - \frac{1-\tau}{1+\sigma}\varphi \\ \log h(\varphi, \varepsilon) &= \kappa_h - \frac{\varphi}{1+\sigma} + \frac{1-\tau}{\sigma+\tau}\varepsilon\end{aligned} \tag{18}$$

where the terms (κ_c, κ_h) are common across individuals.

The consumption allocation reveals that, because of the redistribution imposed by the tax system, only a fraction $1-\tau$ of uninsurable shocks and education wage differentials transmits into consumption. Heterogeneity in the taste for work effort φ translates into different hours worked and earnings one for one and into consumption proportionately to $1-\tau$. The insurable productivity shock ε does not affect consumption because of the separability between consumption and hours in preferences.

Turning to hours worked, uninsurable shocks α and education wage differentials are not a source of hours dispersion because income and substitution effects in labor supply exactly offset each other with our preference specification. Insurable productivity shocks affect hours worked proportionately to the tax-modified Frisch elasticity $(1-\tau)/(\sigma+\tau)$. Note in particular that this modified Frisch elasticity is lower than the Frisch elasticity that would obtain in the absence of progressive taxation, namely, $1/\sigma$.

The constant terms (κ_c, κ_h) in the allocations are functions of $(p_H, p_L, v_\varepsilon, v_\alpha)$ and of all the time-invariant model parameters $(\sigma, \tau, \bar{\varphi}, g, v_\varphi)$.[28] Through κ_c and κ_h, consumption and leisure are both increasing in the variance of the insurable shock. This force is behind some of the welfare gains from the rise in wage inequality, as we explain formally in Section 4.4.

Turning to the education choice, in Heathcote, Storesletten, and Violante (2010a), we show that the difference in expected utility between being a college graduate and a high school graduate is

$$\mathcal{W}_H^0 - \mathcal{W}_L^0 = \frac{(1-\tau)(p_H - p_L)}{1-\beta\pi}. \tag{19}$$

The differential expected utility has a very simple form: it is the after-tax discounted present value of the log skill premium. Within-group components do not show up since they affect

required to solve for the equilibrium wealth distribution. See Huggett (1993), Aiyagari (1994), Ríos-Rull (1995), and Krusell and Smith Jr. (1998), among others. See Heathcote, Storesletten, and Violante (2009b) for a survey of the most recent contributions.

[28]The expression for κ_c also depends on the equilibrium value of λ, which, in turn, is a function of $(p_H, p_L, v_\varepsilon, v_\alpha)$ and the time-invariant model parameters. These expressions are derived in Heathcote, Storesletten, and Violante (2010b).

equally the value of both education levels. Let $\chi^0 = \mathcal{W}_H^0 - \mathcal{W}_L^0$ be the threshold utility cost at which an individual is just indifferent about attending college. Since χ is assumed to be lognormally distributed, the equilibrium fraction of college-educated workers is given by

$$e = \Phi\left(\frac{\ln \chi^0 - \mu_\chi}{\sqrt{v_\chi}}\right), \tag{20}$$

where Φ is the standard Normal cumulative distribution function.

4.3 Parameterization

We now describe how we choose values for all the structural parameters of the model. To be consistent with the welfare calculations based on the empirical approach, we take 1980-1984 to be the pre-shift steady state (indexed by "$*$") and 2001-2005 to be the post-shift steady state (indexed by "$**$"). The model's period is set to one year.

In Heathcote, Storesletten, and Violante (2009a), we use the log-linear relationship between post-government and pre-government earnings implied by equation (16), i.e.,

$$\log \tilde{y}_{it} = \lambda_t + (1 - \tau_\tau) \ln y_{it}, \tag{21}$$

to measure the progressivity parameter τ. We estimated this relationship on CPS data separately for 1980-1984 and 2001-2005 for households with positive labor income[29] and obtained $\tau^* = 0.36$ in the first period and $\tau^{**} = 0.28$ in the second period. In what follows, we use the mean value $\tau = 0.31$ and return to this estimated decline in progressivity in Section 4.4.

We set $\sigma = 2$, which, given our estimate for τ, implies a tax-modified Frisch elasticity of 0.30. This value is in line with the existing microeconomic estimates. We set $\beta = 0.972$ to reproduce an equilibrium risk-free rate of three percent per year. A value of $\pi = (1 - 1/35)$ is consistent with our age range of 25-60 in the survey data.

In Heathcote, Storesletten, and Violante (2010a), we report that the fraction of college graduates of age 25-29 increased from 22 percent in 1980-1984 to 28.7 percent in 2001-2005.[30] We use these numbers for the shares e^* and e^{**} in the model.

From the CPS sample (see Figure 1), we compute that the average log wage differential between the two education groups –i.e., the log skill premium– in the two steady states is,

[29] Pre-government income y_{it} is defined as the sum of earnings, asset income, and private transfers. Disposable income $\tilde{y}_{it}$ is pre-government income plus public transfers (reported) minus taxes (imputed). See the CPS Appendix in Heathcote, Perri, and Violante (2010) for details.

[30] The original source is Table A.2 of the Educational Attainment section on the US Census Bureau website, www.census.gov/population/www/socdemo.

respectively, $p_H^* - p_L^* = 0.31$ and $p_H^{**} - p_L^{**} = 0.52$. These two equations, together with normalizations (14) and (15), yield four equations in the four parameters $(p_H^*, p_L^*, p_H^{**}, p_L^{**})$ with solution $p_H^* = 0.23$, $p_L^* = -0.08$, $p_H^{**} = 0.38$, and $p_L^{**} = -0.14$. With θ set to 1.5 (Katz and Murphy (1992)), we obtain that the skill-biased technological change parameter rises from $\zeta^* = 0.36$ to $\zeta^{**} = 0.47$. Equation (20) evaluated at both steady states yields the equilibrium enrollment rates in the model. Values $\mu_\chi = 3.26$ and $v_\chi = 6.20$ reproduce observed graduation rates $e^* = 0.220$ and $e^{**} = 0.287$.

Finally, to estimate the insurable and the uninsurable components of wage inequality, we review and follow the identification and estimation strategy we developed in Heathcote, Storesletten, and Violante (2009a). We first tack measurement error onto the consumption and hours allocations (18). For wages, we recognize that wages are measured as annual earnings divided by annual hours and, as such, they inherit error from both variables. We make the common assumption that measurement error is classical: i.i.d. across agents and over time and orthogonal across variables. We also assume that its size remains constant between 1980-1984 and 2001-2005. Next, we construct equilibrium cross-sectional variances and the covariances of the joint distribution of wages, hours, and consumption within education groups. In Heathcote, Storesletten, and Violante (2009a), we prove that, given values for (σ, τ) and an external estimate of the variance of measurement error in hours, one can identify the cross-sectional variances $(v_\alpha, v_\varepsilon)$ in both subperiods. Finally, through a simple algorithm that minimizes the distance between empirical and theoretical (co)variances, we estimate the parameter values.

We find that approximately 0.05 points of the 0.08 point increase in the residual (i.e., within-group) variance of log wages are due to the uninsurable component and 0.03 points to the insurable component. In particular, we find that $v_\alpha^* = 0.112$, $v_\alpha^{**} = 0.162$, $v_\varepsilon^* = 0.079$, and $v_\varepsilon^{**} = 0.108$.[31]

4.3.1 Enrollment elasticities

Our calibration approach assumes that the increase in graduation rates observed between 1980-1984 and 2001-2005 was entirely driven by the increase in the college wage premium. This is a reasonable benchmark given that enrollment dynamics roughly track the dynamics of the college premium over this period (see, e.g., Heckman, Lochner, and Taber (1998) and Heathcote,

[31] These specific values are obtained for a variance of measurement error set to around half of the total variance for hours. For smaller measurement error, the levels of the variances are different, but their absolute changes across the two subperiods are essentially the same. It is the absolute changes that matter for our welfare calculations.

Storesletten, and Violante (2010a)).[32]

We now compare elasticity estimates based on micro data to the macro elasticity implied by our calibration. Recently, Fortin (2005) has estimated the elasticity of enrollment to changes in tuition at 4-year U.S. public institutions. By assuming that the enrollment response to a one-dollar increase in tuition is the same as the response to a one-dollar reduction in the present value of a college degree, we translate Fortin's estimated tuition elasticity into a college wage premium elasticity. This generates an enrollment response to an increase in the college premium which is about 7 times larger than the corresponding elasticity built into our calibration.[33] We conclude that the macro elasticity built into our calibration is very small compared to Fortin's micro estimates. We will find that rising enrollment is nonetheless an important source of welfare gains from changes in the wage structure. It follows that were we to calibrate the enrollment elasticity to micro estimates, these welfare gains would be much larger. However, there are many reasons why micro estimates would be inappropriate in the context of our exercise. For example, when a state decreases college tuition, it may attract more out-of-state students who, absent the state's tuition reduction, would have graduated in other states. Therefore, the enrollment elasticity to a local tuition change will be larger than the aggregate elasticity.

4.4 Welfare analysis

We now move to computing the welfare effects of changes in the wage structure. We do so by comparing social welfare across two steady states that differ in the triplet $(\zeta, v_\alpha, v_\varepsilon)$.[34] Changes

[32] This assumption would be less tenable for the 1960s and 1970s which saw swings in enrollment that are difficult to explain through price movements alone. For example, male enrollment surges above trend in the mid 1970s before falling back (see, e.g., Heathcote, Storesletten, and Violante (2010a)). Some authors attribute this temporary surge to the incentives provided by the Vietnam War draft deferment rules for male college students and to GI Bill benefits for war veterans who enrolled in college programs (Card and Lemieux (2001)).

[33] Fortin (2005) estimates an enrollment elasticity to tuition changes of -0.15 for the United States. When assuming a work life of 35 years, a college duration of four years, a growth-adjusted discount rate of 3%, and empirical values for tuition and wages, this elasticity translates into an enrollment elasticity to changes in the annual college premium of 3.3. The empirical macro counterpart of this elasticity is 0.48. Details of these calculations are contained in the Appendix.

[34] Due to the overlapping-generation structure of the model, the economy will undergo a transition in response to a one-off shift in the wage structure. The duration of transition will depend on expectations. If the rise in the college premium comes as a surprise, then 40 years must pass before the entire workforce has made education choices given correct post-shift expectations about the wage structure. This signals an asymmetry in the timing of how different components of the wage structure affect welfare. The education costs born by newborn agents and the effects of changes in v_α and v_ε will be felt immediately after the change in the wage structure. However, the share of the population with a college degree will only rise slowly until the economy reaches the new steady state. Rising enrollment is associated with welfare gains. Thus we conjecture that if we were to model the transition explicitly, the welfare gains during transition would be smaller than those indicated by our steady

in this triplet induce the shifts of the wage structure. Recall that because of the skill-biased demand shift $\Delta\zeta$ and the change in endogenous education decisions, the two steady states also differ in the skill prices (p_L, p_H), the enrollment level e, and in the average utility cost χ paid by college graduates. We will consider two welfare measures: the measure of social welfare we introduced in Section 3.1 and expected lifetime utility.

Social Welfare From the social welfare function (4) and the utility function specification in (11), we obtain the counterpart of equation (5), which defines ω implicitly:

$$\begin{aligned}\sum_{j=0}^{\infty}\left\{(1-\pi)\pi^j\int\left[\log\left((1+\omega)c(\varphi,s,\alpha)\right)-\frac{\exp\left(\bar{\varphi}+\varphi\right)}{1+\sigma}h(\varphi,\varepsilon)^{1+\sigma}\right]dF^{*}_{\varphi,s,\alpha,\varepsilon}\right\}-(1-\pi)\bar{\chi}^{*}e^{*} \quad (22)\\ =\sum_{j=0}^{\infty}\left\{(1-\pi)\pi^j\int\left[\log c(\varphi,s,\alpha)-\frac{\exp\left(\bar{\varphi}+\varphi\right)}{1+\sigma}h(\varphi,\varepsilon)^{1+\sigma}\right]dF^{**}_{\varphi,s,\alpha,\varepsilon}\right\}-(1-\pi)\bar{\chi}^{**}e^{**},\end{aligned}$$

where $F_{\varphi,s,\alpha,\varepsilon}$ is the joint distribution of $(\varphi, s, \alpha, \varepsilon)$ in the population. Note the additional term in the social welfare function that captures the utility costs of attending college: $\bar{\chi}$ is the average utility cost paid by college graduates and $(1-\pi)e$ the fraction of newborn agents with a college education.

Heathcote, Storesletten, and Violante (2010b) show how to integrate with respect to the parametric joint distribution for $(\varphi, s, \alpha, \varepsilon)$ to compute closed-form expressions for social welfare, given closed-form expressions (18) for allocations. Substituting those expressions into equation (5) allows us to compute the following expression for the welfare effect ω associated with a shift in the wage structure:

$$\begin{aligned}\log(1+\omega) &= \tau(1-\tau)\Delta\frac{v_{\alpha}^{0}}{2}+\Delta\ln\left[1-\pi\exp\left(\frac{-\tau(1-\tau)}{2}v_{\eta}\right)\right]-(1-\tau)\Delta\frac{v_{\alpha}}{2} \quad (23)\\ &+\left[2\left(\frac{1-\tau}{\sigma+\tau}\right)-\sigma\left(\frac{1-\tau}{\sigma+\tau}\right)^{2}\right]\frac{\Delta v_{\varepsilon}}{2}\\ &+\Delta\ln\mathbb{E}\left[P_s\right]-\Delta\ln\mathbb{E}\left[\exp\left((1-\tau)p_s\right)\right]+\Delta\mathbb{E}\left[(1-\tau)p_s\right]-(1-\pi)\Delta\left(\bar{\chi}\cdot e\right),\end{aligned}$$

where Δ denotes the change across steady states. The first line captures the welfare effects of changes in the variances of the uninsurable component of wages. The second line in this expression captures the welfare effect of changes in the variance of insurable shocks. The third line captures the welfare effect of changing skill prices, coupled with the associated equilibrium change in college graduation. We emphasize that this welfare expression is exact and relies on no algebraic or numerical approximations.

state comparison. In Heathcote, Storesletten, and Violante (2010a) we model transition explicitly, and compute the welfare effects for each cohort.

A more compact expression for ω can be obtained by making the following three approximations:

$$\begin{aligned}
\tau(1-\tau)\Delta\frac{v_\alpha^0}{2}+\Delta\ln\left[1-\pi\exp\left(\frac{-\tau(1-\tau)}{2}v_\eta\right)\right]-(1-\tau)\Delta\frac{v_\alpha}{2} &\approx -(1-\tau)^2\frac{\Delta v_\alpha}{2} \qquad (24)\\
(1-\tau)\Delta\mathbb{E}\left[p_s\right]-\Delta\log\mathbb{E}\left[\exp\left((1-\tau)p_s\right)\right] &\approx -\frac{(1-\tau)^2}{2}\Delta\left[e(1-e)\left(p_H-p_L\right)^2\right]\\
\log(1+\omega) &\approx \omega
\end{aligned}$$

The first approximation follows from $\exp\left(\frac{-\tau(1-\tau)}{2}v_\eta\right)\approx 1-\tau(1-\tau)\,v_\eta/2$ together with $\log\left((1-\pi)\left(1+\pi/\left(1-\pi\right)\,\tau(1-\tau)v_\eta/2\right)\right)\approx\log(1-\pi)+\pi/\left(1-\pi\right)\,\tau(1-\tau)v_\eta/2$. This approximation is very convenient, because it illustrates the need to decompose Δv_α into its subcomponents, Δv_α^0 and Δv_η. The second approximation would be exact if p_s were normally distributed with mean equal to negative one-half of the variance. It is approximate because p_s has a two-point distribution. It is exact for any distribution, up to a second-order approximation. We will compute welfare effects using both the exact and the approximate welfare expressions.

Given these three approximations, the welfare effect ω simplifies to

$$\begin{aligned}
\omega \;\approx\; & -\frac{1}{2}\underbrace{(1-\tau)^2\Delta\left[e(1-e)\left(p_H-p_L\right)^2\right]}_{(1)\,=\,\Delta var^{bet}(\log c)}-\frac{1}{2}\underbrace{(1-\tau)^2\Delta v_\alpha}_{(2)\,=\,\Delta var^{with}(\log c)} \qquad (25)\\
& -\frac{\sigma}{2}\underbrace{\left(\frac{1-\tau}{\sigma+\tau}\right)^2\Delta v_\varepsilon}_{(3)\,=\,\Delta var(\log h)}\\
& +\underbrace{\left(\frac{1-\tau}{\sigma+\tau}\right)\Delta v_\varepsilon}_{(4)\,=\,\frac{\partial(\log Y/N)}{\partial v_\varepsilon}}+\underbrace{\Delta\log\mathbb{E}\left[P_s\right]}_{(5)\,=\,\frac{\partial(\log Y/N)}{\partial\zeta}}-\underbrace{(1-\pi)\Delta\left(\bar{\chi}\cdot e\right)}_{(6)\,=\,\Delta\text{ education cost}}.
\end{aligned}$$

This expression offers a decomposition of the welfare effects of a shift in the wage structure into six factors. The first is negative one-half times the change in the variance of log consumption between education groups, $\Delta var^{bet}(\log c)$. This size of this term depends on the change in the log skill price differential, as well as the change in enrollment rates. The second term is negative one-half times the change in the variance of log consumption within educational groups, $\Delta var^{with}(\log c)$. Within-group consumption dispersion rises with the variance of uninsurable residual risk. Taken together, the first two terms capture the welfare effects of changes in consumption dispersion and correspond exactly to Lucas' expression for the welfare cost of consumption dispersion (equation (9)) specialized to the case of logarithmic preferences ($\gamma=1$). Note that the magnitude of these two terms is decreasing in τ, because the more progressive is taxation, the smaller is the effect of either a rising skill premium or larger uninsurable wage shocks on consumption inequality.

The third term in the welfare expression is the welfare loss from the rise in hours inequality. This term is σ (which indexes the cost of hours dispersion, à la Lucas) times negative one-half the change in the variance of log hours induced by changes in idiosyncratic insurable risk.[35] More tax progressivity (higher τ) reduces the modified Frisch elasticity of labor supply, $(1-\tau)/(\sigma+\tau)$, thereby reducing hours dispersion and the magnitude of this term.

The last line of (25) contains the additional level effects that only the structural approach can capture. The fourth term, $\partial(\log Y/N)/\partial v_{\varepsilon}$, is the gain in average labor productivity associated with an increase in insurable risk. Insurable risk raises average productivity per hour worked because it induces more productive agents to work relatively long hours.[36] Note that the sum of these two terms in v_{ε} is positive. Thus, an increase in the variance of insurable shocks is necessarily welfare-improving.

The final two terms in the third line of (25) measure the welfare gain associated with the rise in average productivity driven by higher graduation rates, $\partial(\log Y/N)/\partial\zeta$, net of the disutility cost of larger investment in education. In the fifth term, the expression $\mathbb{E}[P_s]$ is the average wage rate. This would by construction remain constant across the two steady states if there were no change in graduation rates. However, if more people go to college following a skill-biased demand shift, the average wage rises, so $\Delta\log\mathbb{E}[P_s]>0$. This would, *ceteris paribus*, increase welfare. However, when additional individuals go to college, more people must pay the education cost. Moreover, since these new graduates would not have taken education in the initial steady state, they must have higher education costs than the initial college graduates. Hence, the average college cost conditional on graduating $\bar{\chi}$ will also rise. These two effects are captured in the sixth term.

Expected lifetime utility An alternative metric for the welfare effects of changes in the wage structure is to compute the change in expected lifetime utility for a newborn agent, ω^0. This welfare effect is defined by exactly the same expression as for the gain in social welfare ω in equation (22) except that the weights across age groups are different: π^j is replaced by $\beta^j\pi^j$.

[35] Recall that given balanced growth preferences, changes in skill prices or in the variance of permanent shocks have no impact on hours.

[36] We discuss this effect in detail in Heathcote, Storesletten, and Violante (2008).

This translates into the following approximate expression:

$$
\begin{aligned}
\omega^0 \quad \approx \quad & -\frac{1}{2}(1-\tau)^2\Delta\left[e(1-e)\left(p_H-p_L\right)^2\right]-(1-\tau)^2\frac{\Delta v_\alpha^0}{2}+(1-\tau)\left(\frac{\pi}{1-\pi}\tau-\frac{\beta\pi}{(1-\beta\pi)}\right)\frac{\Delta v_\eta}{2} \\
& -\frac{\sigma}{2}\left(\frac{1-\tau}{\sigma+\tau}\right)^2\Delta v_\varepsilon \\
& +\left(\frac{1-\tau}{\sigma+\tau}\right)\Delta v_\varepsilon+\Delta\log\mathbb{E}\left[P_s\right]-(1-\beta\pi)\Delta\left(\bar{\chi}\cdot e\right).
\end{aligned}
$$

There are two differences in the expressions for ω^0 relative to the expression for ω. First, from the perspective of newborn agents, the welfare costs of larger uninsurable life-cycle shocks to wages, Δv_η, is reduced, because newborn agents discount the welfare costs of these shocks relative to uncertainty realized at labor market entry (v_α^0). Second, newborn agents put a higher weight on the welfare costs of rising college enrollment, because these costs are paid up front at the start of the life cycle.

4.5 Results

Structural welfare calculation We now report the welfare calculation based on the model. This is the structural counterpart of the calculations in Section 3.2. Table 4 summarizes the results. When we substitute the parameter values estimated in Section 4.3 into the welfare expression (25), we obtain a net social welfare gain of $\omega = 1.4$ percent. More precisely, the exact welfare gain computed from equation (23) is 1.38 percent, while the approximate welfare gain from equation (25) is 1.43 percent.[37] Decomposing the latter, using the exact welfare expression, the rise in consumption inequality between and within groups –the first two terms in (25)– account for losses of -1.0 percent and -1.2 percent respectively. The increase in hours inequality induces an additional loss of -0.3 percent. Overall, these two components are the structural counterpart of the welfare calculation under the empirical approach. They account for a loss of -2.4 percent, in line with our estimate of -1.8 percent for the same preference parameters in Section 3.2 based on the empirical approach.[38] In addition, the loss associated with rising between-group consumption inequality (-1.0 percent) is very similar to the -0.9 percent loss from the analogous Attanasio and Davis-style empirical calculation.

[37] Moving from the exact to the approximate expression, incorporating the three approximations in (24) sequentially, the value for ω changes from 1.38 to 1.37 to 1.42 to 1.43. This calculation assumes that the increase Δv_α is entirely attributable to Δv_η, which is broadly consistent with the estimates in Heathcote, Storesletten, and Violante (2010a). Under the alternative assumption that the increase in Δv_α was entirely attributable to Δv_{α_0}, the sequence is essentially unchanged: 1.39, 1.38, 1.43, 1.43.

[38] We do not obtain exactly the same estimate mainly because the model does not replicate perfectly the empirical rise in consumption and hours inequality.

Table 4: Summary of Structural Welfare Calculations

(A) Average % welfare change and its decomposition

ω	$\Delta var^{bet}(\log c) +$ $\Delta var^{with}(\log c)$	$\Delta var^{with}(\log h)$	$\frac{\partial(\log Y/N)}{\partial v_\varepsilon}$	$\frac{\partial(\log Y/N)}{\partial \varsigma} -$ Δ edu cost
1.4	-2.2	-0.3	0.9	3.0

(B) % Welfare change for newborn and its distribution

	Fraction of population	CEV
ω^0	1.000	1.3
$\omega^0_{H\to H}$	0.220	12.3
$\omega^0_{L\to H}$	0.067	5.6
$\omega^0_{L\to L}$	0.713	-2.4

The terms that are new relative to the empirical approach capture the level effects associated with productivity changes induced by behavioral responses to changes in the wage structure. The productivity gain from larger insurable wage dispersion (fourth term) translates into a welfare gain of 0.9 percent. Finally, the higher labor productivity due to increased human capital investment net of the additional education cost (the last two terms) represents the largest source of welfare gains, 4.0 percent -1.0 percent $= 3.0$ percent.

Using the alternative metric of the consumption-equivalent change in welfare for a newborn agent, ω^0, the (exact) welfare gain associated with the change in the wage structure is $\omega^0 = 1.3$ percent, compared with $\omega = 1.4$ percent. This slightly smaller gain reflects a smaller welfare gain associated with rising enrollment: the net gain here is reduced from three percent to two percent. A largely offsetting effect is a smaller welfare loss associated with rising within-group consumption dispersion: -1.3 percent instead of -2.2 percent.[39]

Distribution of welfare changes The overall welfare gain for newborn households masks significant variation across education groups. While the overall welfare gain is large and positive, the vast majority of households expect welfare losses. We divide agents by the idiosyncratic cost of college, χ_i. Those agents with the lowest values for χ_i will choose to go to college in both steady states. This group, making up 22.0 percent of the population, expects very large welfare gains, equivalent to a permanent 12.3 percent increase in consumption. Agents with

[39]We note, however, that the estimate for ω^0 is much more sensitive to the nature of the rise in v_α. If $\Delta v_\alpha = \Delta v_{\alpha_0}$ (as opposed to $\Delta v_\alpha = (\pi/(1-\pi))\Delta v_\eta$), then $\omega^0 = 0.4$ percent.

Table 5: Counterfactuals on the Role of Insurance Channels

Model	Insurance channel missing	ω
Baseline	None	+1.4%
$\sigma = \infty$	Flexible labor supply	+0.8%
$\varepsilon \to \alpha$	Private risk-sharing	+0.1%
$\tau = 0$	Public insurance	+0.1%
$\Delta q = 0$	Rise in college enrollment	−6.0%

intermediate college costs do not go to college in the initial steady state, but choose to do so in the final steady state. This group, making up 6.7 percent of the population, experiences the same wage gains as the first group, but a smaller welfare gain of 5.6 percent, reflecting higher education costs. Households with the highest college costs never find it worthwhile to go to college. This group accounts for 71.3 percent of the total population and is adversely affected by the widening of the college premium. Thus, agents in this group expect a large welfare loss from the shift in the wage structure equivalent to a permanent 2.4 percent reduction in consumption. These results are summarized in Table 4.

Role of insurance mechanisms Four channels of adjustment mediate the transmission of the shift in the wage structure to welfare: labor supply, private risk sharing, public insurance through redistributive taxation, and human capital investment. We now perform a simple series of counterfactuals where we shut down these mechanisms one by one and examine how the welfare effects change. In each of these counterfactuals, the changes in the skill-biased technological change parameter ζ and the change in the variance of wage dispersion within groups $(v_\alpha, v_\varepsilon)$ are the same as in the baseline model. However, differences in optimal household choices across these counterfactuals translate into differences in general equilibrium dynamics for skill prices. The results of these experiments are reported in Table 5.

To analyze the role of labor supply as an adjustment mechanism, we shut down this margin by setting $\sigma = \infty$. By doing so, we remove the welfare gain from increased labor productivity associated with a higher variance of insurable shocks. From our baseline decomposition, we can immediately see that ω drops from its baseline value of 1.4 percent to 0.8 percent.

Next, we exclude private risk sharing from the model by assuming that all the increase in wage inequality is uninsurable. Beyond losing the productivity gain linked to insurable uncertainty (exactly as in the previous experiment), in this counterfactual the losses originating from market incompleteness increase and, overall, we obtain a value for ω of 0.1 percent.

By setting $\tau = 0$, we omit any scope for social insurance – the only way in which permanent wage shocks α can be smoothed in the model. At the same time, lower progressivity reduces labor supply distortions and increases labor productivity. The first two forces dominate the third, and welfare gains drop to 0.1 percent. Comparing the experiments of shutting down private and public insurance, we see that these two sources of insurance are equally important in our model from a welfare perspective.

Finally, we freeze college enrollment at the level of 1980-1984 and thereby prevent agents from exploiting higher returns to education. Since the shift in labor demand in favor of skilled workers takes place anyway, equilibrium prices move sharply against unskilled labor. As a result, this counterfactual leads to a welfare loss of over six percent. We conclude that human capital investment was the main channel through which households in the US economy were able to take advantage of the new opportunities offered by the shift in the wage distribution.

Time-varying progressivity Our baseline results assume a fixed progressivity rate τ across steady states and, implicitly, attribute all the changes in the cross-sectional moments of the consumption and hours distribution –the sources of our estimates of Δv_α and Δv_ε– to changes in the wage structure. However, our estimate of τ uncovered a nonnegligible change between the two periods: overall progressivity of the US tax and transfer system fell between 1980 and 2005, mostly due to the Tax Reform Act of 1996.[40] Another key advantage of the structural approach over the empirical approach is that we can incorporate changes in progressivity in the model, alongside changes in the wage structure, and isolate the relative effects of the two forces. When τ is allowed to decline from $\tau^* = 0.35$ in 1980-1984 to $\tau^{**} = 026$ in 2001-2005, we estimate a smaller increase across steady states in the variance of uninsurable risk, and a larger increase in the variance of insurable risk (compared with the case with τ constant). This is not surprising because some of the rise in the variance of log consumption observed in the data can be explained by the tax system becoming less redistributive. We repeat our welfare calculation using these new parameter values, holding τ fixed across steady states at its average value in order to isolate the effects of the shift in the wage structure. We now find even larger welfare gains: $\omega = 2.2$ percent rather than $\omega = 1.4$ percent.

[40] Our estimates of τ refer to the progressivity of the entire tax and transfer system. Other sources report a sizeable decline in the marginal tax rates of federal and state taxes in the same period. See, for example, `http://www.nber.org/~taxsim/marginal-tax-rates`.

5 Concluding remarks

What are the welfare consequences of the recent shift in the wage structure in the United States? In this paper, we have summarized the literature on the topic, a body of work that began almost two decades ago by extending the analysis of trends in wage inequality to variables more directly correlated with well-being, such as consumption and hours worked.

The early welfare calculations were based on changes in the empirical distribution of consumption and hours, analyzed through the lens of a social welfare function. We have revisited this type of calculation here and obtained similar results to the original studies: welfare losses on the order of two percent of lifetime consumption when comparing the early 1980s with the early 2000s. These welfare losses primarily reflect increases in consumption inequality.

We have pointed out that, by demeaning the empirical cross-sectional distribution, the empirical welfare calculations do not allow for the possibility that the same sources behind the shift in the wage structure –the rise in the skill premium and the rise in residual wage volatility– can lead to welfare improvements through their effects on average output and productivity. Thus, the empirical approach paints an unduly pessimistic picture. The data show that individuals have, correctly, interpreted the rise in the skill premium as a rise in the return to human capital and have responded by acquiring more education. At the same time, part of the rise in wage volatility was insurable and, given flexible labor supply, could be exploited by reallocating labor effort efficiently.

Quantifying the welfare gains associated with these channels requires a structural model. We used the model of Heathcote, Storesletten, and Violante (2009a, 2010b), which has the virtue of delivering a closed-form and hence very transparent solution for the welfare effect of the shift in the wage structure. Under a plausible calibration of this model, our calculations yield average welfare gains exceeding one percent of lifetime consumption.

The sharp rise in US economic inequality has featured prominently in the public policy debate. Our results suggest some interesting policy trade-offs. For example, more progressive taxation reduces the costs of larger wage shocks that cannot be insured privately, but dissuades individuals from acquiring additional education in response to a widening skill premium and thus dampens the productivity gains associated with human capital investment (see also Guvenen, Kuruscu, and Ozkan (2009) for a similar argument). The policy challenge is to design institutions and tax/transfer schemes that deliver insurance against misfortune at birth and later in life, while preserving incentives for agents to make efficient investments in education and efficient labor supply decisions.

There are still plenty of open questions to be investigated. Here, we list only a few. First, the recent reexamination of household consumption expenditure data in the CEX has led some authors to revise upward estimates of the increase in consumption inequality since 1980. This evidence would point toward larger welfare losses because of imperfect consumption insurance.

Second, throughout the paper, we have abstracted from changes in the distribution of hours worked for women, mainly because of the lack of comprehensive data on home production. In particular, we have brushed aside a key trend in the wage distribution: the narrowing gender gap. In Heathcote, Storesletten, and Violante (2010a), we show that the rise in relative wages of women, together with the increase in their participation rates, is an additional source of welfare gains.

Third, our welfare calculation covers the period 1980-2005, but the rise in wage inequality started well before 1980, the first year available in the CEX. Exploiting the structure of our model, changes in the variances of wages and hours and in their covariance can be used to identify the underlying changes in the insurable and uninsurable components of wage risk. Alternatively, following Blundell, Pistaferri, and Preston (2008) and Guvenen and Smith Jr. (2010), one could impute consumption expenditures in the PSID to construct a data set with a joint distribution of consumption, hours worked, and wages. Thus, there are several ways to estimate welfare effects for the period before 1980, even in the absence of comprehensive consumption data for that period.

Fourth, the interaction between enrollment decisions and the college wage premium turns out to be quantitatively very important for the size of the overall welfare gains from changes in the wage structure. In order to properly understand the welfare effects of skill-biased technical change, it is important to pursue a richer model of education in future work.

Finally, the rise in inequality is not a phenomenon unique to the United States. As documented by Krueger, Perri, Pistaferri, and Violante (2010), the income distribution widened in a number of other countries over the last 30 years (e.g., Canada, United Kingdom, Sweden, and Germany). At the same time, we see considerable variation across countries in the dynamics of inequality in consumption and labor supply, and hence welfare. A cross-country comparative analysis would shed light on how differences in labor market institutions, family structures, and tax/transfer systems translate into different welfare effects from shifts in the wage distribution.[41]

[41] The special issue of the *Review of Economic Dynamics* 2010 (vol. 13, no. 1) contains articles detailing the

Appendix

Data description

CPS The CPS is the source of official US government statistics on employment and unemployment and is designed to be representative of the civilian noninstitutional population. The Annual Social and Economic Supplement (ASEC) applies to the sample surveyed in March and extends the set of demographic and labor force questions asked in all months to include detailed questions on income. For the ASEC supplement, the basic CPS monthly sample of around 60,000 households is extended to include an additional 4,500 Hispanic households (since 1976) and an additional 34,500 households (since 2002) as part of an effort to improve estimates of children's health insurance coverage: this is the "SCHIP" sample. We use the March supplement weights to produce our estimates. Our CPS sample covers the period 1967-2005.

CEX The CEX consists of two separate surveys, the quarterly Interview Survey and the Diary Survey, both collected for the Bureau of Labor Statistics by the Census Bureau. Even though its main purpose is that of providing weights for the Consumer Price Index (CPI), it is the only US data set that contains detailed information about household consumption expenditures. The Diary Survey focuses only on expenditures for small, frequently purchased items (such as food, beverages, and personal care items), whereas the Interview Survey aims at providing information on up to 95 percent of the typical household's consumption expenditures. We will focus only on the Interview Survey, but we return to this point below. The CEX Interview Survey is a rotating panel of households that are selected to be representative of the US population. It started in 1960, but continuous data are available only from the first quarter of 1980, which is the start of our sample. Each quarter the survey reports, for the cross section of households interviewed, detailed demographic characteristics for all household members, detailed information on consumption expenditures for the three-month period preceding the interview, and information on income, hours worked, and taxes paid over a yearly period. Each household is interviewed for a maximum of four consecutive quarters. Our CEX sample covers the period 1980-2005.

Sample selection In both data sets we construct the sample following the same criteria, those outlined in Heathcote, Perri, and Violante (2010). From the raw data, we drop records 1) if there is no information on age for either the head or spouse; 2) if no household member

specific experience of twelve different countries and collects the underlying survey data.

is of working age, which we define as between the ages of 25 and 60; 3) if either the head or spouse has positive labor income but zero weeks worked; 4) if either the head or spouse has an hourly wage less than half of the corresponding federal minimum wage in that year. In the CEX, we also drop households whose quarterly equivalized food consumption is below \$100 in 2000 dollars and those flagged as "incomplete income reporters."[42]

In all data sets, we forecast mean values for top-coded observations by extrapolating a Pareto density fitted to the non-top-coded upper end of the observed distribution. We apply this procedure separately to each component of income in each year. Throughout the paper, unless explicitly mentioned, we express all income and expenditure variables in year 2000 dollars. The price deflator used is the Bureau of Labor Statistics (BLS) CPI-U series, all items.

Variable definition Hours worked in the market are defined as total annual hours worked on all jobs. We define individual wage as annual individual earnings divided by annual hours worked, where annual earnings are defined as wage and salary income plus two-thirds of self-employment income. Our baseline measure of consumption includes expenditures on non-durables, services, small durables, and an estimate of the service flow from vehicles and housing. Household consumption expenditures are adjusted to a per-adult-equivalent basis using the OECD equivalence scale. The OECD scale assigns a weight of 1.0 to the first adult, 0.7 to each additional adult, and 0.5 to each child, defined as an individual age 16 or younger. See Heathcote, Perri, and Violante (2010) for more details on the sample construction and the variable definition.

Comparing micro and macro elasticities of enrollment rates

We wish to compare our macro estimate for the elasticity of college enrollment to the college wage premium to micro estimates of the elasticity of enrollment to tuition costs. To do so, consider the following generic model for education, which is similar to the one developed in the paper, but which explicitly incorporates tuition costs. The model assumes that enrollment decisions are driven solely by a comparison of the expected boost to the present value of earnings associated with a college degree against the costs associated with annual tuition ϕ and additional non-tuition (psychic) costs, χ. The marginal agent has a psychic cost χ^* measured

[42]Table 1 in the Appendix of Heathcote, Perri, and Violante (2010) summarizes the number of records in each data set that are lost at each stage of the selection process. Their Table 2 contains some summary statistics of the sample.

in dollars given by

$$\chi^* = \sum_{t=T_s}^{T} \left(\frac{1}{1+r}\right)^t \Delta y - \left[\sum_{t=0}^{T_s-1} \left(\frac{1}{1+r}\right)^t \left(y^{HS} + \phi\right)\right]$$

where Δy is the college earnings premium, $T+1$ is the length of working life (35 years), r is the interest rate (3%), y^{HS} are earnings as a high school graduate, and T_s is the length of college (4 years). Let F denote the population CDF of the psychic cost χ. Then enrollment is

$$e = F\left(\chi^*\right) = F\left(\sum_{t=T_s}^{T} \left(\frac{1}{1+r}\right)^t \Delta y - \left[\sum_{t=0}^{T_s-1} \left(\frac{1}{1+r}\right)^t \left(y^{HS} + \phi\right)\right]\right)$$

Now, the elasticity of enrollment with respect to to the tuition fee is

$$\varepsilon_\phi = \frac{\partial e/e}{\partial \phi/\phi} = -\frac{\phi}{e}\frac{\partial F}{\partial \chi} \cdot \sum_{t=0}^{T_s-1} \left(\frac{1}{1+r}\right)^t = -\frac{\phi}{e}\frac{\partial F}{\partial \chi} \cdot \frac{1-\left(\frac{1}{1+r}\right)^{T_s}}{1-\frac{1}{1+r}} = -\frac{\phi}{e}\frac{\partial F}{\partial \chi} \cdot \frac{1-\left(\frac{1}{1.03}\right)^{4}}{1-\frac{1}{1.03}} = -\frac{\phi}{e}\frac{\partial F}{\partial \chi} \cdot 3.83.$$

Using data on changes in tuition and college enrollment to US public universities over the 1973-1999 period, Fortin (2005) estimates $\varepsilon_\phi = -0.15$ (Table 1B).

The elasticity of enrollment with respect to to the college premium is

$$\varepsilon_{\Delta y} = \frac{\partial e/e}{\partial \Delta y/\Delta y} = \frac{\Delta y}{e}\frac{\partial F}{\partial \chi} \cdot \sum_{t=T_s}^{T} \left(\frac{1}{1+r}\right)^t = \frac{\phi}{e}\frac{\partial F}{\partial \chi} \cdot \frac{\Delta y}{\phi} \cdot \frac{\left(\frac{1}{1+r}\right)^{T_s} - \left(\frac{1}{1+r}\right)^{T}}{1-\frac{1}{1+r}}.$$

Hence,

$$\varepsilon_{\Delta y} = -\frac{\varepsilon_\phi}{3.83} \cdot \frac{\Delta y}{\phi} \cdot \frac{\left(\frac{1}{1.03}\right)^{4} - \left(\frac{1}{1.03}\right)^{35}}{1-\frac{1}{1.03}} = -\varepsilon_\phi \cdot \frac{\Delta y}{\phi} \cdot \frac{18.3}{3.83}$$

Fortin reports that the average annual tuition cost for US state colleges (towards a baccalaureate degree) was $\phi = \$4,793$ in 2003/2004. The median usual weekly earnings of full-time wage and salary workers 25 years and over in 2005 was \$1,013 for college graduates and \$583 for high school graduates (Current Population Survey, September 2010). This means that the annual college premium is $\Delta y = 52 \cdot (\$1,013 - \$583) = \$22,360$. Fortin's estimate of ε_ϕ then implies an elasticity $\varepsilon_{\Delta y}$ given by

$$\varepsilon_{\Delta y} = -\varepsilon_\phi \cdot \frac{\Delta y}{\phi} \cdot \frac{18.3}{3.83} = 0.15 \cdot \frac{22,360}{4,793} \cdot \frac{18.3}{3.83} = 0.15 \cdot 22.3 = 3.3.$$

Intuitively the reason a one percent increase in the college earnings premium generates a much larger increase in enrollment than a one percent decline in tuition costs is that tuition costs are many times smaller than the present value of the additional earnings associated with a college degree.

Between the periods 1980-1984 and 2001-2005, the college graduation rate, i.e., the share of college graduates among the cohort of 25-29 years old individuals, increased from 0.220 to 0.287. Moreover, during this period the college premium, defined as average annual earnings for college graduates over average annual earnings for high school graduates, increased from 1.363 to 1.682. When normalizing the average wage to unity each period, this implies a macro elasticity of

$$\frac{\partial \log(e)}{\partial \log(\Delta y)} = \frac{\log(e^{**}) - \log(e^{*})}{\log(\Delta y^{**}) - \log(\Delta y^{*})} = \frac{\log(0.287) - \log(0.22)}{\log(1.46 - 0.87) - \log(1.26 - 0.92)} = 0.48.$$

Thus the macro elasticity is roughly seven times smaller than the comparable micro elasticity. Of course there are a variety of ways to reconcile the two estimates. For example, the macro elasticity would equal 0.48 if either (a) college tuition was \$32,950 per year, (b) the college earnings premium was only \$3,252 per year, or (c) the interest rate used to discount future earnings was 25%.

References

ACEMOGLU, D. (2002): "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 40(1), 7–72.

ACEMOGLU, D., AND D. H. AUTOR (2010): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and D. E. Card, vol. 4B of *Handbooks of Labor Economics*. Elsevier.

AGUIAR, M., AND M. BILS (2010): "Has Consumption Inequality Mirrored Income Inequality?," *mimeo*.

AGUIAR, M., AND E. HURST (2007): "Life-Cycle Prices and Production," *American Economic Review*, 97(5), 1533–1559.

——— (2009): "Deconstructing Lifecycle Expenditure," *mimeo*.

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109(3), 659–684.

ATKINSON, A. B. (1970): "On the Measurement of Inequality," *Journal of Economic Theory*, 2(3), 244–263.

ATTANASIO, O., E. BATTISTIN, AND H. ICHIMURA (2007): "What Really Happened to Consumption Inequality in the United States?," in *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, ed. by E. R. Berndt, and C. R. Hulten, vol. 67 of *Studies in income and wealth*, chap. 17, pp. 515–544. University of Chicago Press.

ATTANASIO, O., E. BATTISTIN, AND M. PADULA (2010): "The Evolution of Consumption and Wage Inequality: Evidence and Implications for Consumption Smoothing," *mimeo*.

ATTANASIO, O., AND S. J. DAVIS (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104(6), 1227–1262.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the U.S. Labor Market," *American Economic Review*, 96(2), 189–194.

BATTISTIN, E., R. BLUNDELL, AND A. LEWBEL (2009): "Why is Consumption More Log Normal than Income? Gibrat's Law Revisited," *Journal of Political Economy*, 117(6), 1140–1154.

BENABOU, R. (2002): "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?," *Econometrica*, 70(2), 481–517.

BLUNDELL, R. (2010): "From Income to Consumption: Partial Insurance and the Transmission of Inequality," *Econometric Society Presidential Lecture*.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

BLUNDELL, R., AND I. PRESTON (1998): "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*, 113(2), 603–640.

Broda, C., E. Leibtag, and D. E. Weinstein (2009): “The Role of Prices in Measuring the Poor’s Living Standards,” *Journal of Economic Perspectives*, 23(2), 77–97.

Broda, C., and J. Romalis (2009): “The Welfare Implications of Rising Price Dispersion,” *mimeo.*

Burstein, A., and J. Vogel (2010): “Globalization, Technology, and the Skill Premium: A Quantitative Analysis,” *mimeo.*

Calvo, G. A., and M. Obstfeld (1988): “Optimal Time-Consistent Fiscal Policy with Finite Lifetimes,” *Econometrica*, 56(2), 411–432.

Card, D., and T. Lemieux (2001): “Dropout and Enrollment Trends in the Postwar Period: What Went Wrong in the 1970s?,” in *Risky Behavior among Youths: An Economic Analysis*, ed. by J. Gruber, chap. 9, pp. 439–482. National Bureau of Economic Research.

Cogley, T. (2002): “Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey,” *Journal of Monetary Economics*, 49(2), 309–334.

Constantinides, G. M., and D. Duffie (1996): “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 104(2), 219–240.

Cutler, D. M., and L. F. Katz (1991): “Macroeconomic Performance and the Disadvantaged,” *Brookings Papers on Economic Activity*, 1991(2), 1–74.

——— (1992): “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980’s,” *American Economic Review*, 82(2), 546–551.

Davis, S. J. (2003): “Comment on Krueger and Perri: On the Welfare Consequences of the Increase in Inequality in the United States,” *NBER Macroeconomics Annual*, 18, 121–132.

Deaton, A. (1997): *The Analysis of Household Surveys*. Johns Hopkins University Press.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996): “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 64(5), 1001–1044.

Feenstra, R. C., and G. H. Hanson (1996): “Globalization, Outsourcing, and Wage Inequality,” *American Economic Review*, 86(2), 240–245.

Fisher, J. D., and D. S. Johnson (2006): “Consumption Mobility in the United States: Evidence from Two Panel Data Sets,” *BE Press: Topics in Economic Analysis & Policy*, 6(1).

Fortin, N. M. (2005): “Rising Tuition and Supply Constraints: Explaining Canada-U.S. Differences in University Enrollment Rates,” in *Higher Education in Canada*, ed. by C. M. Beach, R. W. Boadway, and R. M. McInnis, pp. 369–413. McGill-Queens’ University Press.

Garner, T., G. Janini, W. Passero, L. Paszkiewicz, and M. Vendemia (2006): “The CE and the PCE: A Comparison,” *Monthly Labor Review*, 129(9), 20–46.

Gottschalk, P., and R. A. Moffitt (1994): “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 25(2), 217–272.

——— (2009): “The Rising Instability of U.S. Earnings,” *Journal of Economic Perspectives*, 23(4), 3–24.

Guvenen, F., and B. Kuruscu (2009): "A Quantitative Analysis of the Evolution of the U.S. Wage Distribution, 1970–2000," *NBER Macroeconomics Annual*, 24, 227–276.

Guvenen, F., B. Kuruscu, and S. Ozkan (2009): "Taxation of Human Capital and Cross-Country Trends in Wage Inequality," *mimeo*.

Guvenen, F., and A. A. Smith Jr. (2010): "Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach," *National Bureau of Economic Research Working Paper 16327*.

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, 13(1), 15–51.

Heathcote, J., K. Storesletten, and G. L. Violante (2008): "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk," *Journal of Monetary Economics*, 55(3), 501–525.

——— (2009a): "Consumption Insurance and Labor Supply with Partial Insurance: An Analytical Framework," *mimeo*.

——— (2009b): "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 1, 319–354.

——— (2010a): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118(4), 681–722.

——— (2010b): "Redistributive Taxation in a Partial-Insurance Economy," *mimeo*.

Heckman, J., L. Lochner, and C. Taber (1998): "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," *Review of Economic Dynamics*, 1(1), 1–58.

Huggett, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17(5-6), 953–970.

Jones, C. I., and P. J. Klenow (2010): "Beyond GDP? Welfare across Countries and Time," *mimeo*.

Juhn, C., K. M. Murphy, and B. Pierce (1993): "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy*, 101(3), 410–442.

Kaplan, G., and G. L. Violante (2010): "How Much Consumption Insurance Beyond Self-Insurance?," *American Economic Journal: Macroeconomics*, 2(4), 53–87.

Katz, L. F., and D. H. Autor (1999): "Changes in the Wage Structure and Earnings Inequality," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and D. Card, vol. 3 of *Handbooks of Labor Economics*, chap. 26, pp. 1463–1555. Elsevier.

Katz, L. F., and K. M. Murphy (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107(1), 35–78.

Knowles, J. (2008): "Why Are Married Men Working So Much?," *mimeo*.

KREBS, T. (2003): "Human Capital Risk and Economic Growth," *Quarterly Journal of Economics*, 118(2), 709–744.

KRUEGER, D., AND F. PERRI (2003): "On the Welfare Consequences of the Increase in Inequality in the United States," *NBER Macroeconomics Annual*, 18, 83–121.

——— (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73(1), 163–193.

KRUEGER, D., F. PERRI, L. PISTAFERRI, AND G. L. VIOLANTE (2010): "Cross Sectional Facts for Macroeconomists," *Review of Economic Dynamics*, 13(1), 1–14.

KRUGMAN, P. (2005): "The Debt-Peonage Society," New York Times, 08/03/2005.

——— (2007): *The Conscience of a Liberal*. W. W. Norton & Company.

KRUSELL, P., T. MUKOYAMA, A. SAHIN, AND A. A. SMITH JR. (2009): "Revisiting the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, 12(3), 393 – 404.

KRUSELL, P., L. E. OHANIAN, J.-V. RÍOS-RULL, AND G. L. VIOLANTE (2000): "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68(5), 1029–1053.

KRUSELL, P., AND A. A. SMITH JR. (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

——— (1999): "On the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, 2(1), 245–272.

LAZEAR (2006): "Lazear vs Krugman," http://gregmankiw.blogspot.com/2006/07/lazear-vs-krugman.html.

LEMIEUX, T. (2008): "The Changing Nature of Wage Inequality," *Journal of Population Economics*, 21(1), 21–48.

LEMIEUX, T., W. B. MACLEOD, AND D. PARENT (2009): "Performance Pay and Wage Inequality," *Quarterly Journal of Economics*, 124(1), 1–49.

LUCAS, R. E. (1987): *Models of Business Cycles*. Blackwell.

MEGHIR, C., AND L. PISTAFERRI (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

——— (2010): "Earnings, Consumption and Lifecycle Choices," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and D. E. Card, vol. 4B of *Handbooks of Labor Economics*. Elsevier.

MORETTI, E. (2010): "Real Wage Inequality," *mimeo.*

RÍOS-RULL, J.-V. (1995): "Models with Heterogeneous Agents," in *Frontiers of Business Cycles Research*, ed. by T. F. Cooley, chap. 4, pp. 98–125. Princeton University Press.

SLESNICK, D. T. (1994): "Consumption, Needs and Inequality," *International Economic Review*, 35(3), 677–703.

——— (1998): "Empirical Approaches to the Measurement of Welfare," *Journal of Economic Literature*, 36(4), 2108–2165.

——— (2001): *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*. Cambridge University Press.

Storesletten, K. (2003): "Comment on Krueger and Perri: On the Welfare Consequences of the Increase in Inequality in the United States," *NBER Macroeconomics Annual*, 18, 132–137.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51(3), 609–633.

Violante, G. L. (2002): "Technological Acceleration, Skill Transferability, and the Rise in Residual Inequality," *Quarterly Journal of Economics*, 117(1), 297–338.