

WORKING PAPER · NO. 2018-43

The Changing (Dis-)Utility of Work

Greg Kaplan and Sam Schulhofer-Wohl

June 2018

The Changing (Dis-)Utility of Work*

Greg Kaplan and Sam Schulhofer-Wohl

Revised June 2018

ABSTRACT

We study how changes in the distribution of occupations have affected the aggregate non-pecuniary costs and benefits of working. The physical toll of work is smaller now than in 1950, with workers shifting away from occupations in which people report experiencing tiredness and pain. The emotional consequences of the changing occupation distribution vary substantially across demographic groups. Work has become happier and more meaningful for women, but more stressful and less meaningful for men. These changes appear to be concentrated at lower education levels.

Keywords: Labor supply; Well-being; Stress; Meaningfulness; Happiness

JEL classification: I31, J22, J24, J28, N32

*Greg Kaplan is Professor of Economics, University of Chicago, Chicago, Illinois, and Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. Sam Schulhofer-Wohl is Senior Economist and Research Adviser, Federal Reserve Bank of Chicago, Chicago, Illinois. Their email addresses are gkaplan@uchicago.edu and samuel.schulhofer-wohl@chi.frb.org. We thank Enrico Moretti, Gordon Hanson, and Timothy Taylor for very helpful comments and Sharada Dharmasankar for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System. This paper is preliminary and is in preparation for the *Journal of Economic Perspectives*.

1. Introduction

Economists predominantly view monetary rewards as by far the most important aspect of jobs and careers. The disutility of supplying one hour of labor is assumed to be the same whether that hour is spent building cars on an assembly line, waiting tables at a restaurant, teaching a class, or pitching for the White Sox. In consequence, in conventional models, the tradeoffs workers make between consumption and leisure can be assessed solely by looking at hours worked and wages.

Yet it is obvious to many workers that a job involves more than just forfeiting some leisure time in return for a wage (Schwartz, 2015). Jobs differ in the physical and mental toll they take on workers, as well as in the psychological rewards they provide, such as autonomy and meaning (Kalleberg, 2011).

Our study asks how the major occupational shifts in the postwar period have manifested in changes in the non-pecuniary costs and benefits of work. Many fewer people work on assembly lines now than in 1950, while many more work in services and sales. Women and minorities have moved in large numbers into jobs where they once faced substantial barriers to entry. How have these shifts changed the aggregate amount of hardship or disutility that people experience from their work, the aggregate psychological rewards or utility that they derive from it, and the distribution of disutility and utility across the population?

Drawing on data from the American Time Use Survey (ATUS), we examine six dimensions of workers' feelings about the time they spend on the job in different occupations: how happy, sad, and tired they are; how much stress and pain they experience; and how meaningful they find their work. We then calculate how economy-wide *average* feelings about work depend on the mix of occupations in the economy.

For some of the dimensions we study, such as pain, introspection offers an easy answer to the direction of change for most workers in developed countries. Although even today many jobs are undoubtedly physically demanding, we can easily compare how someone feels after a day of office work with how a farmer in the early 20th century must have felt after a day in the fields. Our results confirm that, in the aggregate, work has become less painful and less tiring in the postwar period.

But for other dimensions, where introspection does not provide easy answers, our

study offers tentative new insights on the directions of change for U.S. workers. We find substantial heterogeneity in how the non-physical costs and benefits of work have changed over time. For women, the non-physical aspects of work have on average become more positive over time: Women have shifted toward occupations that produce more happiness and meaningfulness and less sadness, while experiencing no change in stress. The story for men is more negative. Although they have shared in the reduction in pain and tiredness, they also have shifted toward occupations that produce more stress, less happiness, and less meaningfulness. The improvements for women and the reduction in meaningfulness for men appear to be concentrated among people at lower education levels. All this is not to deny, of course, that many workers even today have jobs that are painful, tiring, meaningless, saddening, or stressful — we find only that the share of such jobs is lower than in the past.

Our analysis of the non-pecuniary implications of changes in the occupational structure complements the large existing literatures on the wage and employment implications of these changes. For example, Herrendorf, Rogerson, and Valentinyi (2014) examine the aggregate employment and consumption effects of sectoral shifts from agriculture to manufacturing and then to services over the last two centuries. More recent changes in the occupational structure have been characterized by polarization — meaning the simultaneous growth of high-wage, high-skill jobs and low-wage, low-skill jobs, even as employment shrinks in the middle of the wage and skill distribution (Autor, Katz, and Kearney, 2006). Studies of polarization have highlighted how these different patterns of job growth relate to the changing nature of tasks required by employers and the skills required to do different tasks (Acemoglu and Autor, 2011). Although these analyses usually focus on the monetary returns to different types of skills, it is increasingly accepted that the non-monetary returns to skill have also changed and that these changes differ sharply in the cross-section (Hamermesh, 2001).

There is extensive debate over the macroeconomic forces that have led to polarization but much less work on the consequences of polarization for workers' well-being, both individually and in the aggregate. With the recent rise in long-term unemployment in the United States and the rise in deaths related to substance abuse and suicide in the same time period (Case and Deaton, 2017), it is becoming more evident that the loss of psychological benefits of work may be an important component of the overall costs of changes in employment.

One potential psychological benefit of work is its meaning (if any). The opening sentences of Ariely, Kamenica, and Prelec (2008) summarize the gap between the importance that workers place on meaning and the importance that economists place on it:

Most children think of their potential future occupations in terms of what they will *be* (firemen, doctors, etc.), not merely what they will *do* for a living. Many adults also think of their job as an integral part of their identity. At least in the United States, “What do you do?” has become as common a component of an introduction as the anachronistic “How do you do?” once was, yet identity, pride, and meaning are all left out from standard models of labor supply.

That paper demonstrates the importance of meaning for workers’ productivity in a laboratory setting. But there are also many examples of the strength of meaning as a motivating tool in real workplaces (Grant, 2007, 2012). Against the backdrop of these micro-level examinations of the role of meaning in particular work environments, our study offers a macro perspective on the aggregate meaningfulness of work across the U.S. economy.

On the cost side, economists typically think only of the opportunity cost of the time spent at work. But there are many features that make some jobs less desirable than others (Katz and Krueger, 2016). For example, work can be so physically demanding that it leaves workers tired, injured or ill, or even kills them. Work has become dramatically less deadly over time (Aldrich, 1997; Centers for Disease Control and Prevention, 1999), perhaps as a result of occupational safety regulations (Levine, Toffel, and Johnson, 2012). Yet even if a job does not directly damage a worker’s body, it can take a mental toll, as in Frey’s (1996) description of an air traffic controller who lost radio contact with the airplanes he was guiding:

Watching in helpless horror as his planes careered farther and farther off course, the controller rose from his chair with an animal scream, burst into a sweat and began tearing off his shirt. By the time radio contact was re-established — and the errant planes were reined in — the controller was quivering on the floor half naked, and was discharged on a medical leave until he could regain his wits.

The costs of such workplace stress are potentially significant: The famous Whitehall studies (Marmot et al., 1978, 1991) found an inverse relationship between employment rank in the

British civil service and health outcomes, a pattern that has been interpreted as indicating that stress and other negative psychological features of low-ranking jobs may harm health, given that most of the study subjects at all ranks were office workers who had equal access to the National Health Service (Smith, 1999).

Another non-pecuniary cost of many jobs that has attracted recent attention is inflexibility. Jobs in the so-called “gig economy” may provide more flexibility — for example, Uber drivers can decide exactly when they want to work, instead of taking shifts assigned by a manager (Hall and Krueger, 2016) — though potentially at the cost of reducing workers’ wages or their ability to work full time when they wish to do so (Katz and Krueger, 2016). Some recent studies have tried to quantify the value of flexibility by eliciting willingness to pay for increased autonomy in hours of work (Mas and Pallais, 2017).

The main challenge we face in studying the aggregate changes in the non-pecuniary costs and benefits of work is that the survey data on workers’ feelings are available only in 2010, 2012, and 2013. Thus, we must impute how workers felt about their jobs in past years based on recent information. Our strategy begins by measuring changes in the distribution of occupations. We then use the ATUS data to assign each occupation a vector of scores based on the feelings that its workers report in the recent data. Finally, we ask: If the distribution of occupations were different from what we see today, but feelings in each occupations stayed the same, how would workers’ total experiences change? How much more or less stress, for example, would the workforce collectively experience if the distribution of occupations was the one observed in 1950, rather than the distribution observed today? How much more or less meaning, happiness, tiredness, and pain?

Our focus throughout is on market work. In the conclusion, we discuss how our approach might also be used to assess the consequences of women’s significant increase in labor force participation in the postwar period, which would require a measure of feelings about non-market work.

Our approach relies on three key assumptions. First, we assume that the feelings an occupation produces today are the same as those it produced in the past. Second, although economists since Adam Smith (1776) have appreciated that pay may vary inversely with the non-pecuniary qualities of a job, we assume that any such compensating differentials do

not affect the feelings that workers report. Third, we assume that the feelings a particular worker reports on the job are caused by that worker’s occupation and not by his or her other circumstances or personality. At the end of the paper, we describe the implications of these assumptions for our findings as well as some robustness checks. The paper concludes by discussing some potential extensions and broader implications of our work, in particular how findings about the nonpecuniary characteristics of work should influence the analysis of big-picture labor market outcomes such as labor force participation and inequality.

2. Evolution of Occupations Over Time

The first step in our estimation strategy is to measure shifts in the distribution of occupations over time. This requires us to categorize occupations in a consistent way in data from 1950 to the present day.

We use data from decennial censuses from 1950 through 2000 and the 2011–2015 American Community Survey (ACS) to measure the distribution of occupations by sex, race, and education. Our occupation categories use the OCC1990 occupation coding produced by IPUMS (Integrated Public Use Microdata Series), described at <https://usa.ipums.org/usa-action/variables/OCC1990>). OCC1990 is based on the occupation codes used in the 1990 census; it maps occupation codes used in other years to the 1990 codes, and aggregates some categories to make the coding more consistent over time. The ATUS data include only the current occupation coding scheme, which we map to OCC1990 ourselves.

The OCC1990 coding contains 389 occupation categories. Some of these categories are so narrow that we observe very few workers in them in the ATUS — too few to be able to estimate feelings precisely for these occupations. In addition, even though OCC1990 is harmonized, it is not entirely uniform over time because of changes in the level of detail in the census occupation variables.¹ To improve the uniformity of the coding and to ensure a reasonably large number of people are used to calculate workers’ feelings in each occupation,

¹For example, in the 1950 census, almost all people in management jobs were recorded as “Managers, officials, and proprietors (not elsewhere classified),” which maps to the OCC1990 code “Managers and administrators, n.e.c.” (code 022). But by the 1990 census, which forms the basis for the OCC1990 codes, some managers were recorded as working in specialties, such as “managers of food-serving and lodging establishments” (code 017). Thus, a restaurant manager would be assigned the OCC1990 code 022 in the 1950 census but code 017 in the 1990 census or the 2011-2015 ACS.

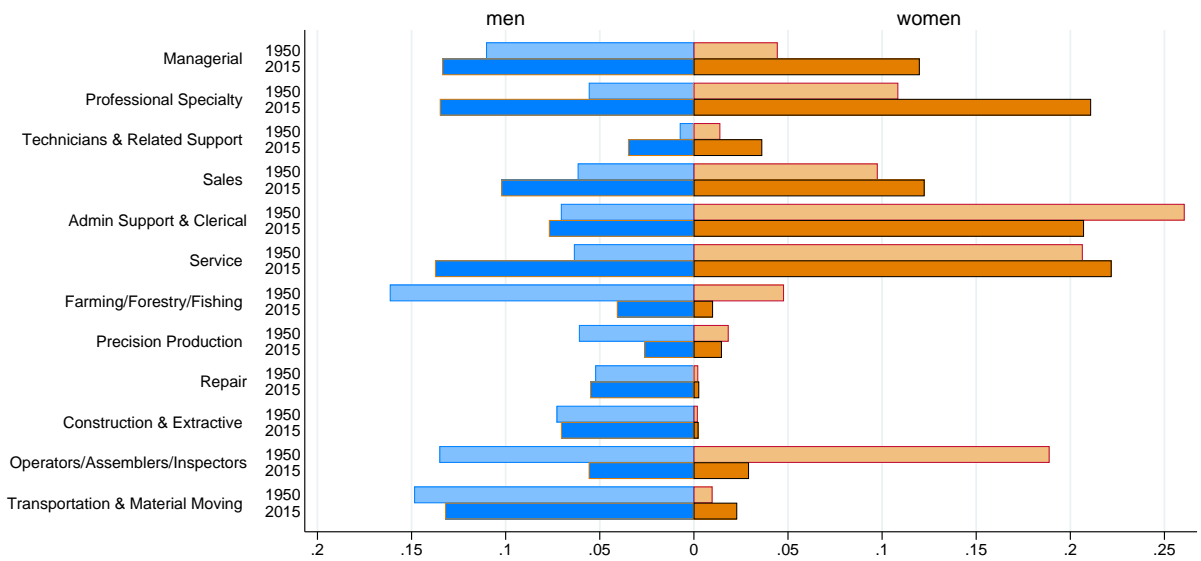


Figure 1: The distribution of workers across broad occupation groupings.

Source: Authors’ calculations from 1950 census and 2011–2015 ACS.

we aggregate the occupations to 12 broad categories. (We exclude military occupations.) Of course, aggregating occupations in this way poses the risk that the occupations categorized as, say, “sales occupations” in 1950 are quite different from those categorized as sales occupations in recent years. However, in analyses not reported here, we have found that we obtain similar overall results if we use the detailed OCC1990 codes, but the changes in the share of workers in each occupation become difficult to interpret (for example, because of the reclassification of detailed categories between 1950 and 1990).

With this coding in hand, we estimate the distribution of occupations by race, sex, and education in the 1 percent sample of the 1950 census, the 5 percent samples of the 1960, 1980, 1990, and 2000 censuses; the 1 percent form 1 and form 2 state samples of the 1970 census; and the 2011–2015 five-year ACS sample. We obtain all datasets from IPUMS (Ruggles et al., 2015). We also consider three education groups: a high school diploma or less, some college, and bachelor’s degree or more.

Figure 1 shows the categories and the distributions of men and women across occupations. Since 1950, both men and women have moved into managerial and professional specialty occupations, and out of farming and machine operating. Women have moved out of administrative support, but the share of men in that field has remained roughly constant.

By contrast, men have shifted in large numbers into service occupations, while the share of women in service occupations is little changed. Some occupations, such as construction, have stable shares of the population over time.

These shifts create the potential for heterogeneity by sex in how feelings about work have changed, for two reasons. First, men and women have moved into and out of different occupations, so even if men and women feel the same about every occupation, the aggregate changes they have experienced will differ. Second, men and women may feel differently about the same occupations, so even where they have experienced similar changes in occupation shares, as with the shift into professional specialty occupations, the impact on the utility or disutility of work may differ. Our methodology will allow for both of these possible sources of change.

3. Feelings About Work

The ATUS, produced by the US Census Bureau, is a stratified random sample of the US population ages 16 and older. (Specifically, ATUS respondents are a subset of respondents to the Current Population Survey.) The survey asks respondents to report, in significant detail, how they spent each minute of a day. Respondents also report their occupation in their main job (but not in any other jobs they may have).

In 2010, 2012, and 2013, the ATUS contained a “well-being module” that randomly selected three activities during the day for each respondent and asked the respondents to report their feelings while engaged in these activities. Activities were eligible to be randomly selected for these questions if they lasted at least five minutes and did not fall into the categories of sleeping, grooming, personal activities, refusal, or don’t know. For the chosen activities, respondents were asked about how they were feeling during these activities along six dimensions: how happy, how sad, how stressed, how tired, how much pain, and how meaningful. They were asked to rank each of their feelings on a scale from 0 (not at all) to 6 (very much).²

²We obtain the ATUS microdata from the American Time Use Survey Data Extract Builder at <http://www.atusdata.org> (Hofferth, Flood, and Sobek, 2015). We use the well-being module activity-level weights for estimation and normalize the weights such that the 2010, 2012, and 2013 samples receive equal weight in the calculations.

To produce an index of feelings by occupation, we run six ordinary least squares regressions. In each case, the dependent variable is one of the six measures of feelings about time spent at work. The explanatory variables are 12 categories of occupations described in the previous section, and dummy variables for age, race, and education level. We then compute the mean response to each question within each occupation category, adjusted for differences in demographics across occupations. In Appendix A1, we provide a detailed description of our procedure, including a table of the adjusted mean feelings for men and women in each occupation.

In our main analysis, we use these data only for respondents who were asked to report their feelings during the activity of working on their main job, but in Appendix A2 we report results from an analysis with individual fixed effects that also uses data on feelings during activities other than work.

Our next step is to combine the (adjusted) estimates of feelings about work for each occupation with the Census data on occupations described earlier. For each Census year, we compute the average stress of work by taking a weighted average of the stress indexes of each occupation, weighting by the distribution of occupations in that year. We repeat this calculation for the five other feelings.³

Figure 2 shows how aggregate mean feelings at work have evolved over time for the six types of feelings in the data. Relative to 1950, the current distribution of occupations makes workers less sad and less tired, and makes them experience less pain. But work also produces more stress. Happiness and meaningfulness both fell in early years, then rose in later years.⁴

³In principle, to compute aggregate mean feelings in the present era, we could directly calculate means of self-reported feelings on the job in the ATUS. However, because the ATUS is relatively small and not all respondents are asked to report their feelings at their main job, the distribution of occupations among ATUS respondents who report feelings on their main job could randomly differ by a significant amount from the population distribution of occupations. To rule out this problem, we estimate aggregate mean feelings in the present era with a weighted average of occupation-specific feelings, weighted by the occupation distribution in the 2011–2015 American Community Survey. See Appendix A1 for details.

⁴It should be noted that estimation uncertainty in this calculation arises both from uncertainty in the estimation of the occupation shares and uncertainty in the estimation of the occupation-adjusted mean feelings. But in practice the Census data are large enough that the occupation shares are estimated quite precisely, and uncertainty in the estimates of mean feelings by occupation in the ATUS data is the main source of uncertainty in our results.

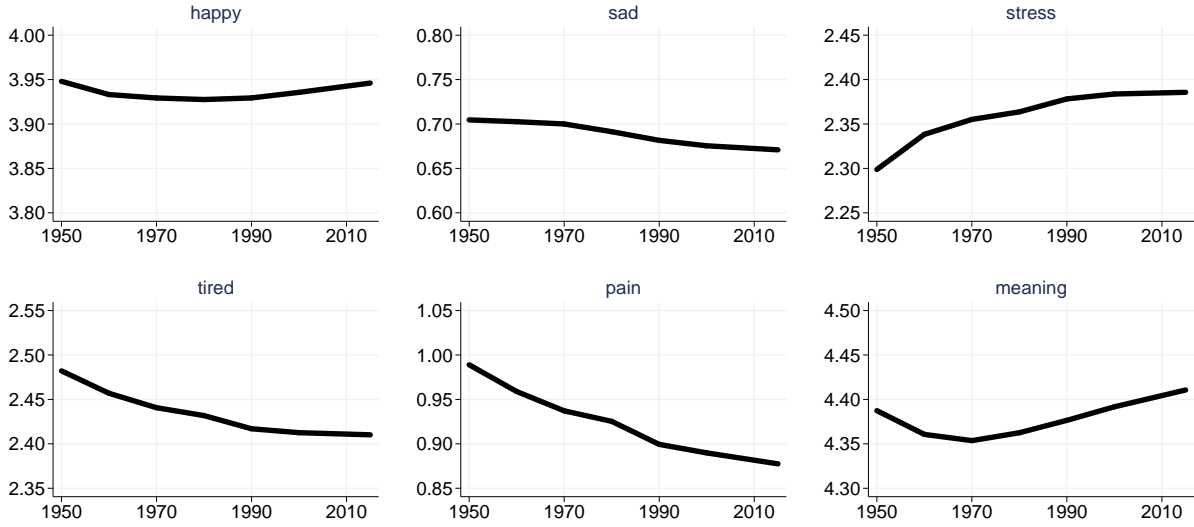


Figure 2: Changes in aggregate feelings at work, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated for full population; occupation scores adjusted for age, race, sex, and years of education. Source: Authors' calculations from census, ACS, and ATUS.

4. Feelings about Work by Gender, Education Level and Race

The results change substantially when we calculate aggregate mean feelings separately by sex. To do this, we follow the same approach but compute both the adjusted mean feelings within each occupation category, and the yearly distribution of occupations, separately by demographic group (sex, sex \times education, and sex \times race \times education). Appendix A1 provides details.

For women, the story is one of consistently improving feelings about work. Over time, as Figure 3 shows, work produced more happiness and a greater sense of meaning, and less sadness, tiredness, and pain; stress levels stayed roughly constant. Thus, over the period we examine, not only were women moving into the work force but they also were shifting to occupations with better nonpecuniary attributes. By contrast, for men, the picture is more mixed: Although work became less painful and tiring, it also became more stressful, less meaningful, and less happy.

These patterns may be partly the result of aggregating together very different oc-

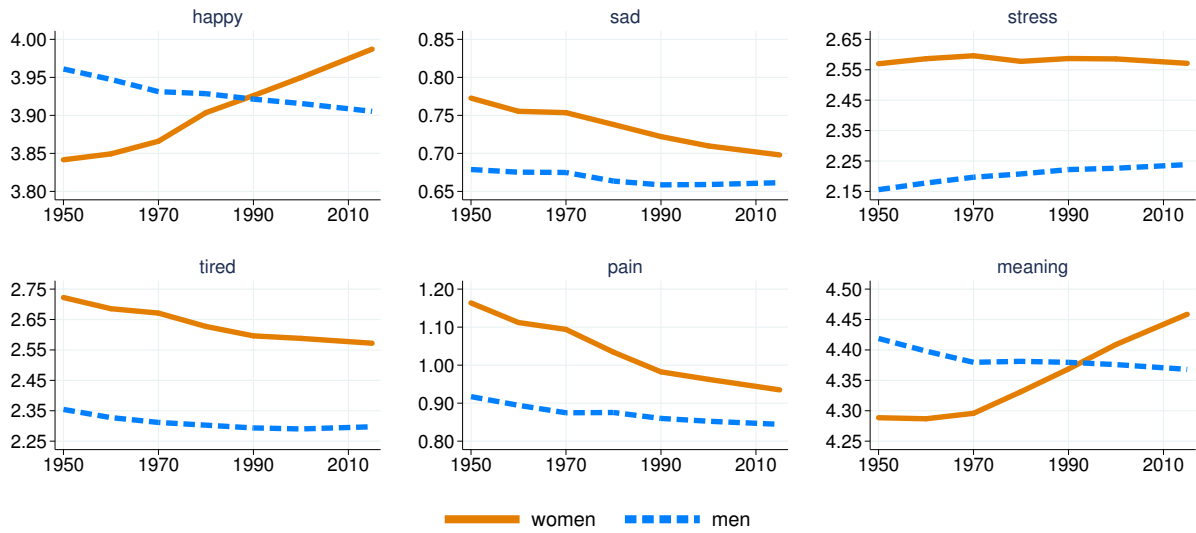
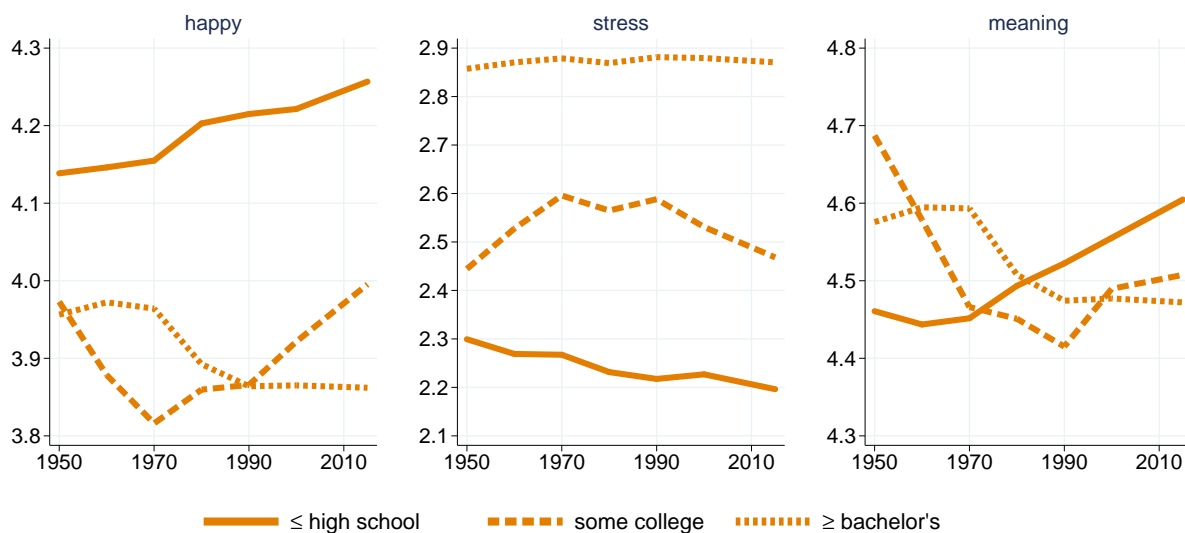


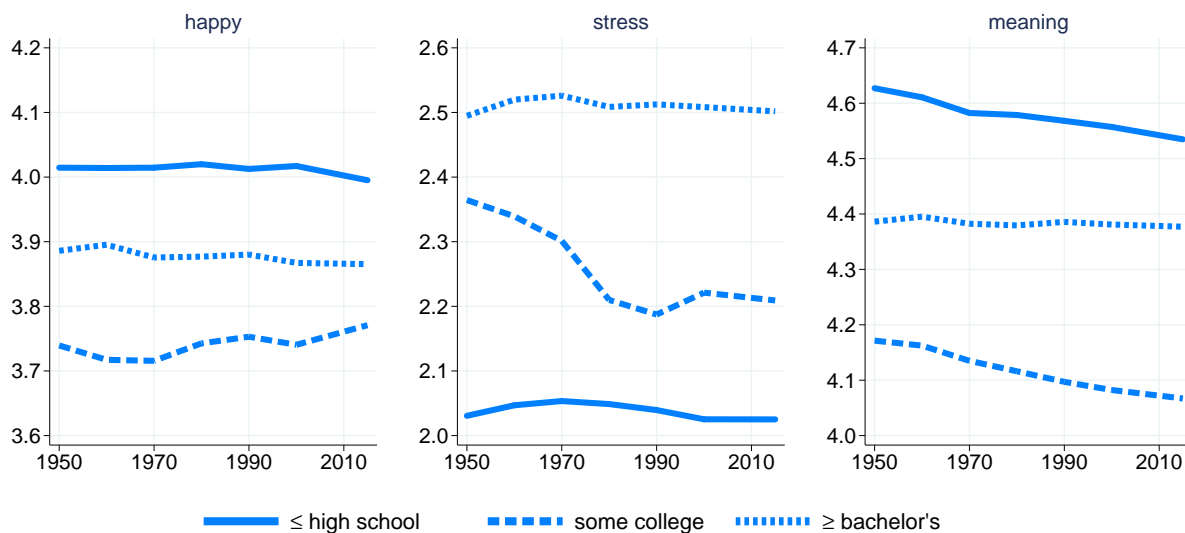
Figure 3: Changes in aggregate feelings at work by sex, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex; occupation scores adjusted for age, race, and years of education. Source: Authors' calculations from census, ACS, and ATUS.

cupations, such as including both restaurant managers and chief executive officers in the managerial category. We can attempt to more finely classify occupations, despite the lack of perfectly uniform coding across years, if we divide the sample by education. As explained in Appendix A1, when we do this, we re-estimate the adjusted mean feelings within each occupation, using only data on workers at a given education level.



(a) Women



(b) Men

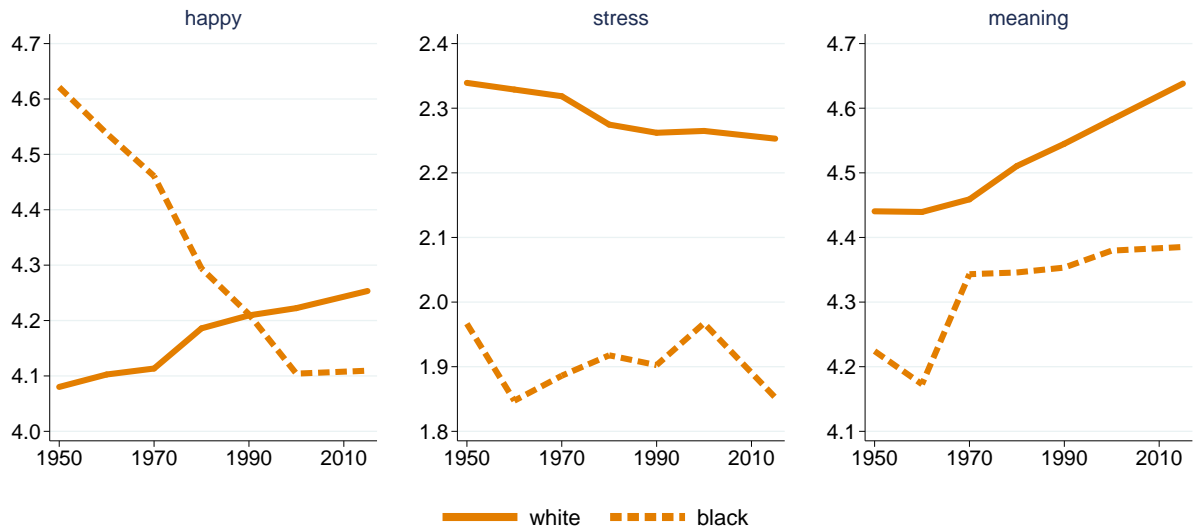
Figure 4: Changes in aggregate feelings at work by sex and education, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex and education; occupation scores adjusted for age, race and years of education. Source: Authors' calculations from census, ACS, and ATUS.

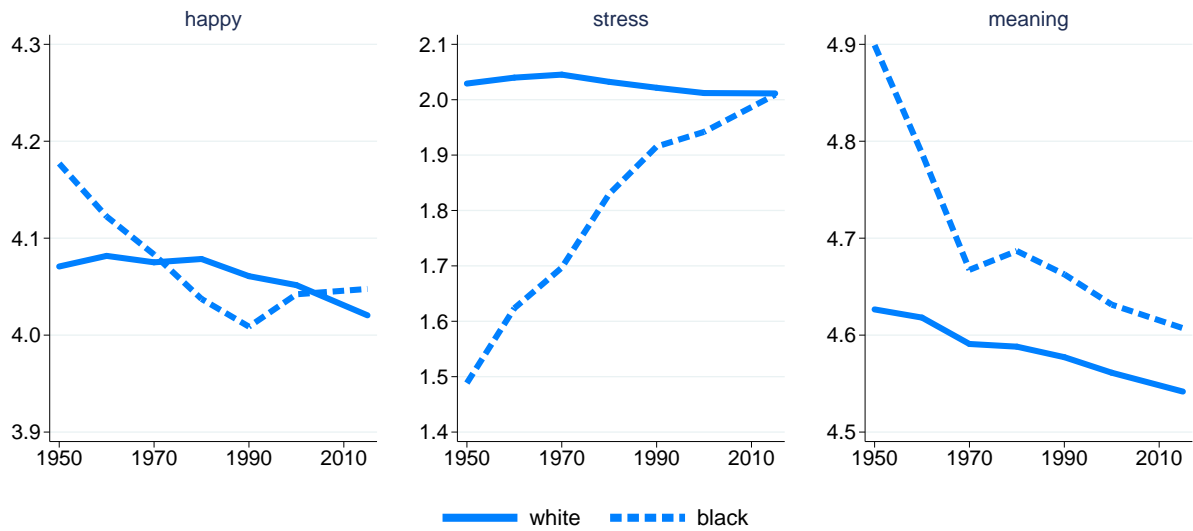
Figure 4 shows how happiness, stress, and meaningfulness have evolved when we divide the sample by both education and sex. We concentrate on happiness, stress, and meaning because there appears to be little interesting heterogeneity in tiredness and pain, and sadness appears to be the inverse of happiness. For women, shown in the top panel of the figure, the gains in happiness and meaningfulness are concentrated among those with no more than a high school diploma. The highest-educated women actually show falling happiness and meaningfulness, similar to the overall findings for men. For men, shown in the bottom panel, there is a clear drop in meaningfulness at lower education levels. However, we find no rise in stress and little decrease in happiness for men within education groups, suggesting that the trends in these variables for men overall might result from aggregation bias.

These conclusions should be regarded as tentative. One reason is that disaggregating by education means we are using a smaller sample to estimate each occupation score. Another reason is that disaggregating by education could under some circumstances exacerbate rather than reduce any bias in our estimates of feelings by occupation — rising education levels within occupations (for example, see Spitz-Oener, 2006) mean that the type of managerial work done by someone with a high school diploma today may be quite different from that done by someone with a high school diploma in 1950.

We can also disaggregate the results by race. We examine only whites and blacks because the ATUS sample contains too few respondents of other races to obtain precise estimates when we disaggregate by race, sex, and education. In Figure 5, we focus on estimates for people with a high school education or less. The trends in meaningfulness are the same across races — meaningfulness has risen for both white and black women, and fallen for both white and black men. However, happiness has risen for white women while falling for black women, and stress has risen for black men while falling for white men and women. Importantly, these estimates account only for differences in feelings about the occupation itself, not for differences in pay. If racial discrimination in pay varies across occupations or has changed over time, the change in workers' overall happiness could be quite different.



(a) Women



(b) Men

Figure 5: Changes in aggregate feelings at work by sex and race (education \leq high school), 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated for black and white respondents with no more than a high school education, separately by race and sex; occupation scores adjusted for age and years of education. Other races excluded from calculation. Source: Authors' calculations from census, ACS, and ATUS.

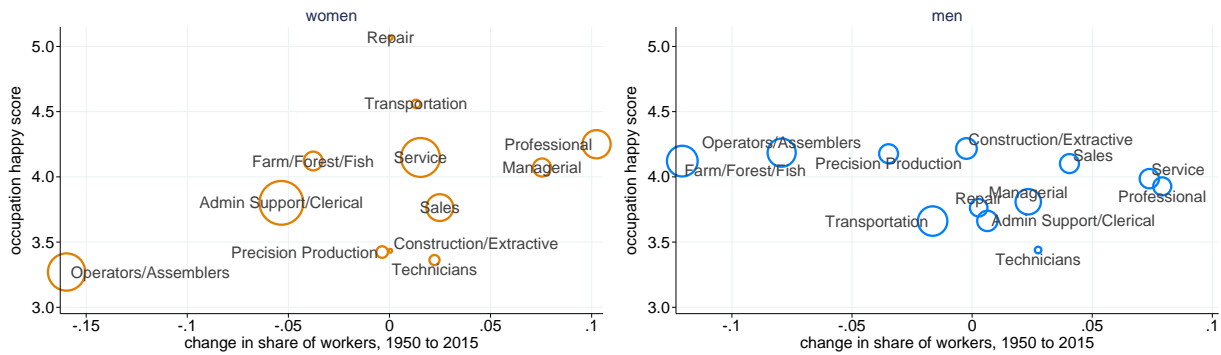
5. Sources of the Shifts

What is driving these shifts in aggregate feelings about work? To gain some insight into this issue, we plot the relationship between an occupation’s average feelings and the change in the share of workers in that occupation since 1950. Figure 6 shows these relationships for happiness, stress, and meaning, separately for women and men. Each circle in the graphs represents a different occupation category, with area proportional to the occupation’s share of workers in 1950.

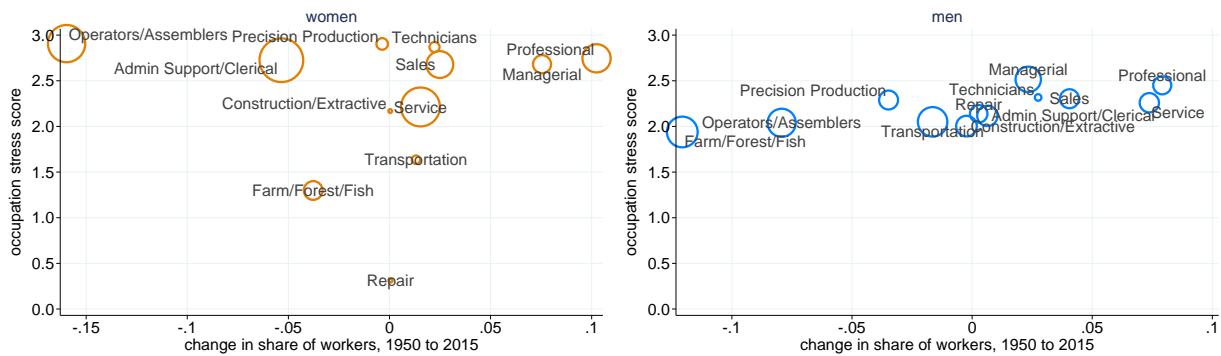
The figure shows that the different results for men and women arise not only from differences in how their occupation distributions have changed, but also from differences in the feelings they report in the same occupational categories. For example, both men and women are less likely now than in 1950 to work as machine operators, assemblers and inspectors. For women, such jobs are associated with below-average happiness and meaningfulness, so the shift increases women’s happiness and meaning at work. For men, such jobs are associated with *above*-average happiness and meaningfulness, as well as below-average stress, so the same shift in the occupation distribution *decreases* the non-pecuniary value of work for men.

These patterns suggest that the overall improvements for women appear to be driven by their shift into professional and managerial work and out of factory work, while the overall decreases for men appear to be driven by their shift out of farming and factory work and into professional and service occupations.

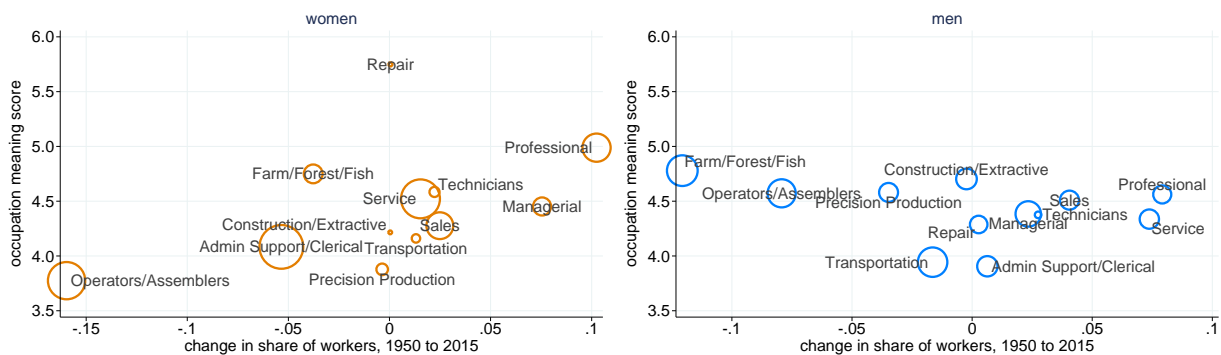
We emphasize that the gender differences in reported feelings for the same occupation can be interpreted in multiple ways. The differences could mean that men and women feel differently about exactly the same jobs. Alternatively, within a single occupation code, men and women might on average be doing slightly different jobs that our coding is not sufficiently detailed to reveal. Another possibility is that men and women feel the same about the actual tasks involved in the work but that the broader work environment has disparate impacts on men and women, such as when sexual harassment occurs. Finally, to the extent that feelings about work are socially constructed, the reported gender differences might reflect messages that society sends to men and women about how they “should” feel about different jobs, rather than any differences in how people would feel absent such messages.



(a) Happy



(b) Stress



(c) Meaning

Figure 6: Changes in occupation share and average feelings by occupation

Area of circle is proportional to share of workers in each occupation in 1950, by sex. Occupation scores and occupation distributions calculated separately by sex; occupation scores adjusted for age, race, and education. Source: Authors' calculations from census, ACS, and ATUS.

6. Discussion of Assumptions

We relied on several strong assumptions to measure workers' feelings about different occupations. Here, we discuss the potential biases that these assumptions may create and some robustness checks that help address them.

A. Have Feelings About Occupations Changed Over Time?

Our approach assumes that the feelings an occupation produced in the past are the same as those it produces today. Direct measures of workers' feelings about their job in past eras would be preferable if they were available, but we are not aware of any historical data on subjective feelings about work in particular occupations that we can measure in a consistent way over time.

This assumption may be particularly problematic when considering the physical costs of work. Regulations and technological improvements have made many jobs safer than in the past. In addition, in many occupations, capital equipment has substituted for human effort; a miner today is more likely to operate heavy machinery and less likely to wield a pick-axe.

To the extent that changes in occupations over time have merely dampened the differences between occupations but not erased or re-ordered them, our results will underestimate the effect of changes in the occupation distribution. For example, if working on an assembly line has always been more tiring than working as retail sales clerk, but the gap is smaller today than in the past, our approach will underestimate the difference in tiredness between manufacturing and retail jobs in 1950 and therefore underestimate the change in average feelings caused by the large-scale shift from the former occupation to the latter.

However, it is also possible that the relative rank of occupations has changed over time. For example, the “meaningfulness” of a particular occupation may in part be socially constructed and depend on the value that the worker's family, friends, neighbors or society at large happen to place on that occupation. Workers in various occupations have sought to increase their prestige by defining them as professions (Larson, 1977), and public opinion polls have also measured fluctuations in occupations' prestige over time (see, e.g., Taylor, 2001). Absent data on workers' feelings in the past, we cannot assess how such changes might have affected our results.

A related possibility is that the relationship between education and how workers feel about their jobs has varied across cohorts. For example, the education system might somehow encourage people who reach a certain schooling level to view a particular kind of occupation as especially appropriate or meaningful, but which occupations these are might change over time. In Appendix A3, we investigate this possibility by controlling for an interaction of age and education when estimating the average feelings in each occupation. This change has little impact on our results.

In addition, it is possible that workers' preferences for particular attributes of work have changed over time, for example because preferences may be non-homothetic and incomes have risen since 1950. Again, absent data on past feelings, we have no way to measure the direction or magnitude of changes in preferences, although one might extend our approach to measure non-homotheticities by estimating the relationship between income and cross-occupation differences in feelings.

B. What About Compensating Differentials?

We do not attempt to measure differences in wages between occupations or how these differences might relate to the feelings that workers report. This focus means that, at most, our analysis provides an account of how the changing occupation distribution affects non-pecuniary costs and benefits of work, but cannot hope to describe the effect on workers' overall welfare.

Moreover, our calculation assumes that the feelings workers report in the ATUS depend only on the actual jobs they are doing and not on how much they are paid to do those jobs, because our method takes account of changes over time in the distribution of occupations but not of changes in the wage for each occupation. If, instead, the feelings that workers report in the survey also depend on their pay, the compensating differentials would likely lead us to underestimate the differences between occupations and underestimate the effect of changes in the occupation structure. For example, if workers who are less happy about their tasks are also more happy about their wages, then the reported difference in happiness between two occupations would be biased downward. This bias is likely of greatest concern for dimensions such as happiness and sadness that potentially reflect workers' overall views of a job, and is

of less concern for dimensions that measure more specific feelings such as meaning or pain.

Related, our results treat each dimension of feelings about work as separate and do not attempt to map changes in the vector of feelings into changes in a single index of the amount of (dis-)utility that workers experience. One could compute such an index by estimating compensating differentials for the feelings that different occupations produce — via hedonic wage models of the type pioneered by Tinbergen (1956) — and then calculating the compensating variation associated with a change in the occupation distribution. However, estimating the compensating differentials is not straightforward, due to the way workers with heterogeneous preferences endogenously sort across occupations (Bartik, 1987; Epple, 1987). We leave such calculations for future research.

C. Might Feeling About Occupations Reflect Other Individual Traits?

Our approach assumes that the feelings a particular worker reports on the job are caused by that worker’s occupation, rather than by other circumstances or personality. For example, when we observe that people in managerial occupations report an above-average level of stress, we assume that this is because management work is inherently stressful, and not because people who would feel stressed in any job are more likely to end up being managers. This assumption will fail if occupation choices are correlated with other factors that affect a person’s feelings, and if it fails, our results will be biased.

One possible robustness check to address this issue, using the ATUS data, is to control for the feelings that workers report when they are not on the job. These feelings in non-work activities might be viewed as an indicator of the permanent feelings that a person would report regardless of occupation, so controlling for them might adjust for all of the non-occupation differences between respondents. In Appendix A2, we carry out this calculation, using a fixed-effects estimator to identify the effect of an occupation by measuring the difference between the feelings that a worker reports on the job and the feelings that the very same worker reports in other activities.

We refer to our original measure of the feelings generated by an occupation as the occupation’s baseline score, and to the measure that is adjusted for feelings reported during other activities as the occupation’s fixed-effects score.

For most occupations whose shares changed substantially, and for most of the types of feelings that we measure, the fixed-effects and baseline occupation scores are closely correlated. This correlation gives some confidence that our basic approach to measuring the feelings induced by an occupation is reasonable.

However, there are a few outliers in the occupation scores, which can lead to different estimates of the aggregate trend in feelings. For example, among women, the fixed-effects score for machine operators, assemblers, and inspectors often differs substantially from the baseline score. This occupation was one of the lowest scoring on happiness for women in the baseline but one of the highest scoring on happiness for women in the fixed effects estimate. Also, this occupation shrunk substantially from 1950 to 2015. As a result, women in this occupation had a downward trend in happiness according to the fixed effects estimates, but an upward trend according to the baseline estimates. For men, farming, forestry and fishing were rated very high in meaning in the baseline estimates but quite low in the fixed effects estimates, while service occupations received a moderate meaning score in the baseline and a high score with fixed effects. Thus, the shrinkage of the agricultural sector and the growth of service work implied decreasing meaningfulness for men according to the baseline estimates but rising meaningfulness according to the fixed effects estimates. The differences between the fixed effects and baseline estimates appear to be concentrated at lower education levels.

The differences between the fixed effects and baseline estimates suggest a need for caution in interpreting the overall results. However, the fixed-effects approach is imperfect in various ways and might have biases of its own. For example, if a good job also gives the worker more positive feelings when she is at home, the difference in feelings between work and home will underestimate the true effect of the occupation on how the worker feels. Alternatively, if people who have bad jobs are particularly happy to go home from them, the difference in feelings between work and home will overestimate the effect of the occupation on feelings. As a result, the fixed-effects measures could be biased in either direction relative to the true change in feelings about work.

7. Conclusion

The distribution of occupations has changed significantly in the post-World War II period. People feel differently about different occupations, and in addition, people in different demographic groups appear to feel differently about the same occupations. Taking all of these factors together, we find substantial shifts both in the aggregate utility and disutility derived from work and in the distribution of that (dis-)utility across people.

Our work highlights how more measurement of the characteristics of work beyond income could offer insight on a number of large-scale questions. As one example, there has been a well-documented secular rise in wage and income inequality (Heathcote, Perri, and Violante, 2010), in particular at the top of the distribution (Piketty and Saez, 2003). Are rising wages at the top of the distribution (at least in part) a compensating differential for particularly demanding jobs, so that inequality in wages exceeds inequality in the total rewards of work? Or are the non-pecuniary benefits of work also increasingly concentrated at the top?

In addition, one of the biggest changes in the labor market in the postwar period has been the rise in women’s participation. Yet little is known about the overall welfare consequences of this change because almost all research has assumed that wages are the sole benefit, and opportunity costs of time are the sole costs, of working. But, as we show, men and women have different likelihoods of working in some occupations, and sometimes feel differently on average about those occupations. Thus, we cannot simply extrapolate from the experience of men to calculate the costs and benefits of work for women. Furthermore, non-market work represents a large fraction of economic activity and of how people, especially women, spend their time (Aguiar and Hurst, 2016; Waring, 1988). Our approach could in principle be applied to measure feelings in non-market work and then analyze how aggregate feelings about *all* work — market and non-market — changed with shifts in labor force participation. We did not pursue that question because it is challenging to define a sharp boundary between non-market work and leisure in the data available to us, but such a study might help provide a more complete account of the implications of rising female labor force participation. Such an analysis could also consider changes in the utility of leisure time, following the study by Aguiar et al. (2017) of how improvements in video games reduced

employment among low-skilled young men.

Finally, our work provides a new perspective on an old question about the labor market: Why do people keep working full-time hours even as income levels have risen? (Or, for those concerned about whether robotics and artificial intelligence will lead to a sharp reduction in demand for labor, how much do people actually want to work?) In a famous essay, “Economic Possibilities for our Grandchildren,” Keynes (1930) raised the possibility that as incomes rose, people would spend few hours working:

For many ages to come . . .everybody will need to do some work if he is to be contented. We shall do more things for ourselves than is usual with the rich to-day, only too glad to have small duties and tasks and routines. But beyond this, we shall endeavour to . . .make what work there is still to be done to be as widely shared as possible. Three-hour shifts or a fifteen-hour week may put off the problem for a great while. For three hours a day is quite enough to satisfy the old Adam in most of us!

It is by now well known that, at least in a strict sense, this prediction of fewer hours worked has not yet been borne out. But Keynes’ argument also hints at a possible reason why: Work is not motivated by wages alone. The prediction really should apply only to the component of work that produces disutility, not the component that produces positive utility. It is possible that, in the aggregate, people in wealthy countries do much less “work,” in the sense of an activity that is a source of disutility, than in Keynes’ time, because more of the time spent working is associated with experiences that workers value positively. And it is also possible that these changes in the experience of working have been disproportionately felt in different parts of the income distribution and different demographic groups. However, there have not been attempts to measure changes in the aggregate non-wage aspects of work over time. This paper is a first small step in that direction.

Appendix

A1. Detailed Description of Data and Methodology

This appendix contains a more detailed description of the data and methods used in the paper. So that this appendix can stand alone, some material is repeated from the main text.

Data sources

We use data from decennial censuses from 1950 through 2000 and the 2011–2015 American Community Survey (ACS) to measure the distribution of occupations by sex, race, and education, and data from the American Time Use Survey (ATUS) to measure workers’ feelings by occupation.

Occupation codes and distributions

Our estimation strategy requires a way to measure the distribution of occupations in a uniform way in data from both 1950 and the present day. We use the OCC1990 occupation coding produced by IPUMS⁵ to do this. OCC1990 is based on the occupation codes used in the 1990 census; it maps occupation codes used in other years to the 1990 codes based on Census Bureau crosswalks, and aggregates some categories to make the coding more consistent over time. IPUMS includes the OCC1990 codes in its public use microdata samples for decennial censuses and the ACS. The ATUS data include only the current occupation coding scheme, which we map to OCC1990 ourselves.

The OCC1990 coding contains 389 occupation categories. Some of these categories are so narrow that we observe very few workers in them in the ATUS — too few to be able to estimate feelings precisely for these occupations. In addition, even though OCC1990 is harmonized, it is not entirely uniform over time because of changes in the level of detail in the census occupation variables. For example, in the 1950 census, almost all people in management jobs were recorded as “Managers, officials, and proprietors (not elsewhere classified),” which maps to the OCC1990 code “Managers and administrators, n.e.c.” (code 022). But by the 1990 census, which forms the basis for the OCC1990 codes, some managers were recorded

⁵See <https://usa.ipums.org/usa-action/variables/OCC1990>.

as working in specialties, such as “managers of food-serving and lodging establishments” (code 017). Thus, a restaurant manager would be assigned the OCC1990 code 022 in the 1950 census but code 017 in the 1990 census or the 2011-2015 ACS. To improve the uniformity of the coding and to ensure a reasonably large number of people are used to calculate workers’ feelings in each occupation, we aggregate the occupations to 12 broad categories. (We exclude military occupations from the analysis.) Of course, aggregating occupations in this way poses the risk that the occupations categorized as, say, “sales occupations” in 1950 are quite different from those categorized as sales occupations in recent years. In analyses not reported here, we have found that we obtain similar overall results if we use the detailed OCC1990 codes, but the changes in the share of workers in each occupation become difficult to interpret (for example, because of the reclassification of restaurant managers between 1950 and 1990).

With this coding in hand, we estimate the distribution of occupations by race, sex, and education in the 1 percent sample of the 1950 census, the 5 percent samples of the 1960, 1980, 1990, and 2000 censuses; the 1 percent form 1 and form 2 state samples of the 1970 census; and the 2011-2015 five-year ACS sample. We obtain all datasets from IPUMS (Ruggles et al., 2015). We consider three education groups: a high school diploma or less, some college, and bachelor’s degree or more.

Feelings

The ATUS, produced by the U.S. Census Bureau, is a stratified random sample of the U.S. population ages 16 and older. ATUS respondents are a subset of respondents to the Current Population Survey. The ATUS asks respondents to report, in significant detail, how they spent each minute of a day. Respondents also report their occupation in their main job (but not in any other jobs they may have).

In 2010, 2012, and 2013, the ATUS contained a “well-being module” that randomly selected three activities during the day for each respondent and asked the respondents to report their feelings while engaged in these activities. Activities were eligible to be randomly selected for these questions if they lasted at least five minutes and were not categorized as sleeping, grooming, personal activities, refusal, or don’t know. For the chosen activities,

respondents were asked:

From 0 to 6, where a 0 means you were not happy at all and a 6 means you were very happy, how happy did you feel during this time?

From 0 to 6, where a 0 means you were not sad at all and a 6 means you were very sad, how sad did you feel during this time?

From 0 to 6, where a 0 means you were not stressed at all and a 6 means you were very stressed, how stressed did you feel during this time?

From 0 to 6, where a 0 means you were not tired at all and a 6 means you were very tired, how tired did you feel during this time?

From 0 to 6, where a 0 means you did not feel any pain at all and a 6 means you were in severe pain, how much pain did you feel during this time if any?

From 0 to 6, how meaningful did you consider what you were doing? 0 means it was not meaningful at all to you and a 6 means it was very meaningful to you.

We use these data for respondents who were asked to report their feelings during the activity of working on their main job. We compute the mean response to each question within each occupation category listed in figure 1, adjusted for differences in demographics across occupations. Specifically, we regress self-reported feelings on dummy variables for single year of age, race, single year of education, and occupation categories:

$$z_i = \theta_0 + \sum_a \alpha_a + \sum_r \beta_r + \sum_e \gamma_e + \sum_o \delta_o + \epsilon_i, \quad (\text{A1})$$

where z_i is respondent i 's report of a particular feeling (such as meaningfulness, stress or sadness); a indexes age; r indexes race; e indexes education; o indexes occupation; and ϵ_i is an unobservable error, which we assume to be uncorrelated with the regressors. The adjusted mean feelings for a given occupation are the predicted value from the regression for the sample distribution of age, race, and education:

$$\bar{z}_o = \hat{\theta}_0 + \sum_a \hat{\alpha}_a s_a + \sum_r \hat{\beta}_r s_r + \sum_e \hat{\gamma}_e s_e + \hat{\delta}_o, \quad (\text{A2})$$

where s_a is the fraction of the regression sample that is age a , s_r is the fraction that is race r , and s_e is the fraction that has e years of education. (The use of the regression sample to compute s_a , s_r and s_e is a normalization that affects only the overall level of mean feelings and not the differences between occupations, which are a function only of the estimated occupation effects $\hat{\delta}_o$. Thus, all of our estimates of the consequences of a changing occupation distribution would be unchanged if we computed s_a , s_r and s_e from the entire ATUS rather than the regression sample, or from a different dataset entirely.)

Because the occupation categories are quite broad, it is possible that people in different demographic groups typically perform different occupations within each category. For example, among people in managerial occupations, those with college degrees may be more likely to be chief executives while those with high school diplomas may be more likely to be restaurant managers. These differences could also affect people’s feelings on the job, and removing demographic differences as in (A1) may be insufficient to control for these effects. Therefore, we also compute the adjusted mean feelings within each occupation category separately by demographic group (sex, sex \times education, and sex \times race \times education) for use when we calculate changes over time by demographic group. We do this by running the regression in (A1) separately for each sub-group.

We obtain the ATUS microdata from the American Time Use Survey Data Extract Builder at <http://www.atusdata.org> (Hofferth, Flood, and Sobek, 2015). We use the well-being module activity-level weights for estimation and normalize the weights such that the 2010, 2012, and 2013 samples receive equal weight in the calculations. Table 1 reports the mean adjusted feelings by sex and occupation.

Producing counterfactual estimates of aggregate feelings

Let \bar{z}_{xo} be the demographically adjusted mean level of a particular feeling reported by workers with characteristics x in occupation o , according to the ATUS. Let $\pi_{xo,t}$ be the fraction of workers with characteristics x who are in occupation o in the year t census. We estimate the counterfactual mean feelings of workers with characteristics x in year t as the weighted average of occupation-specific feelings, weighting by the share of workers in each

<u>Panel A: Women</u>						
	Pain	Happiness	Sadness	Tiredness	Stress	Meaning
Managerial	0.54	4.07	0.69	2.22	2.68	4.45
Professional Specialty	0.85	4.25	0.55	2.68	2.75	4.99
Technicians	1.05	3.36	0.89	2.63	2.87	4.58
Sales	0.99	3.76	0.79	2.72	2.68	4.28
Admin/Clerical	0.85	3.80	0.62	2.48	2.73	4.08
Service	1.07	4.15	0.75	2.58	2.21	4.52
Farming/Forestry/Fishing	1.16	4.12	0.49	2.75	1.30	4.75
Precision Production	1.28	3.42	0.86	2.73	2.90	3.88
Repair	-0.67	5.07	-0.13	-0.09	0.31	5.75
Construction/Extractive	0.86	3.43	1.40	3.05	2.17	4.22
Operators/Assemblers	2.14	3.27	1.21	3.40	2.91	3.77
Transportation	1.09	4.56	1.00	2.36	1.64	4.16

<u>Panel B: Men</u>						
	Pain	Happiness	Sadness	Tiredness	Stress	Meaning
Managerial	0.59	3.81	0.58	2.19	2.52	4.38
Professional Specialty	0.61	3.93	0.56	2.24	2.45	4.56
Technicians	0.62	3.44	0.47	1.97	2.32	4.38
Sales	0.68	4.10	0.56	2.07	2.30	4.51
Admin/Clerical	0.69	3.66	0.90	2.31	2.12	3.91
Service	0.93	3.99	0.73	2.41	2.26	4.34
Farming/Forestry/Fishing	1.13	4.12	0.68	2.66	1.94	4.78
Precision Production	0.96	4.18	0.59	2.23	2.29	4.58
Repair	0.92	3.76	0.72	2.16	2.14	4.29
Construction/Extractive	1.30	4.22	0.40	2.24	2.00	4.70
Operators/Assemblers	0.78	4.19	0.68	2.20	2.04	4.57
Transportation	1.18	3.66	0.87	2.64	2.05	3.94

Table 1: Mean feelings by sex and occupation

Mean reported feelings by sex and occupation (scale 0 to 6), adjusted for age and education as described in text. Source: Authors' calculations from ATUS.

occupation in year t :

$$\hat{z}_{x,t} = \sum_o \pi_{xo,t} \bar{z}_{xo}. \quad (\text{A3})$$

Note that estimation uncertainty in $\hat{z}_{x,t}$ arises both from uncertainty in the estimation of the occupation shares $\pi_{xo,t}$ and uncertainty in the estimation of the occupation adjusted mean feelings \bar{z}_{xo} . In practice, the Census and ATUS samples are so large that $\pi_{xo,t}$ is estimated very precisely. We return below to the implications of uncertainty in the estimation of \bar{z}_{xo} .

In principle, to compute aggregate mean feelings in the present era, we could directly calculate means of self-reported feelings on the job in the ATUS. However, because the ATUS is relatively small to begin with and because not all ATUS respondents are asked to report their feelings at their main job, the distribution of occupations among ATUS respondents who report feelings on their main job could randomly differ by a significant amount from the population distribution of occupations. Thus, if we found a difference between $\hat{z}_{x,1950}$ and the mean feelings in the ATUS, that difference could be caused by a failure of the ATUS sample to accurately reflect the current occupation distribution, and not by a change in the occupation distribution between 1950 and the present. To rule out this problem, we instead estimate aggregate mean feelings in the present era with a weighted average of occupation-specific feelings, weighted by the occupation distribution in the 2011–2015 ACS:

$$\hat{z}_{x,now} = \sum_o \pi_{xo,ACS} \bar{z}_{xo}, \quad (\text{A4})$$

where $\pi_{xo,ACS}$ is the fraction of workers with characteristics x who are in occupation o in the 2011–2015 ACS. This approach is equivalent to reweighting the ATUS data so that the distribution of occupations among respondents who report feelings on the main job matches the distribution of occupations in the 2011–2015 ACS.

The change in feelings associated with the change in the occupation distribution is the difference between the year- t and present-day estimates:

$$\Delta \hat{z}_{x,t} = \hat{z}_{x,now} - \hat{z}_{x,t} = \sum_o (\pi_{xo,ACS} - \pi_{xo,t}) \bar{z}_{xo}. \quad (\text{A5})$$

Observe that $\hat{z}_{x,now}$ and $\hat{z}_{x,t}$ are not independent because they are both functions of the same

ATUS estimates \bar{z}_{xo} . Thus, given the large Census samples and resulting precise estimation of $\pi_{xo,t}$, uncertainty in the estimates of mean feelings by occupation in the ATUS is the main source of uncertainty in our overall results.

A2. Controlling for Individual Fixed Effects

As noted in the text, our approach assumes that the estimated occupation coefficients are causal, so that we can use the coefficients to estimate workers' counterfactual feelings if they were in different occupations. This assumption will fail if occupation choices are correlated with other factors that affect a person's feelings. We can use the structure of the ATUS to partly control for these other factors. Each ATUS respondent reported his or her feelings in three randomly chosen activities, not all of which were necessarily work on the main job. Feelings in non-work activities might be viewed as an indicator of the respondent's baseline level of feelings that he would report regardless of occupation and thus control for non-occupation differences between respondents.

Figure 7 provides some intuition for how this approach might work by showing the distribution across respondents of the difference in raw reported feelings between work and non-work activities. For each respondent who reported feelings in both work and non-work activities, we calculate the difference between the reported feelings at work and the reported feelings not at work (taking an average if there are two work reports or two non-work reports), then plot the histogram of these differences. We make no adjustments for demographics at this stage in order to exhibit the raw data. Respondents tend to report higher stress levels and lower happiness levels at work than at other activities, as well as slightly higher sadness and pain levels. The fixed effects estimator will identify the causal effect of occupations on feelings by examining how these differences in work vs. non-work feelings vary across people who have different occupations.

Specifically, we run the regression in (A1) on all of the observations for each respondent, not just the observations from work on the main job, and controlling for individual fixed effects as well as for the nature of the other activities that are the source of the additional

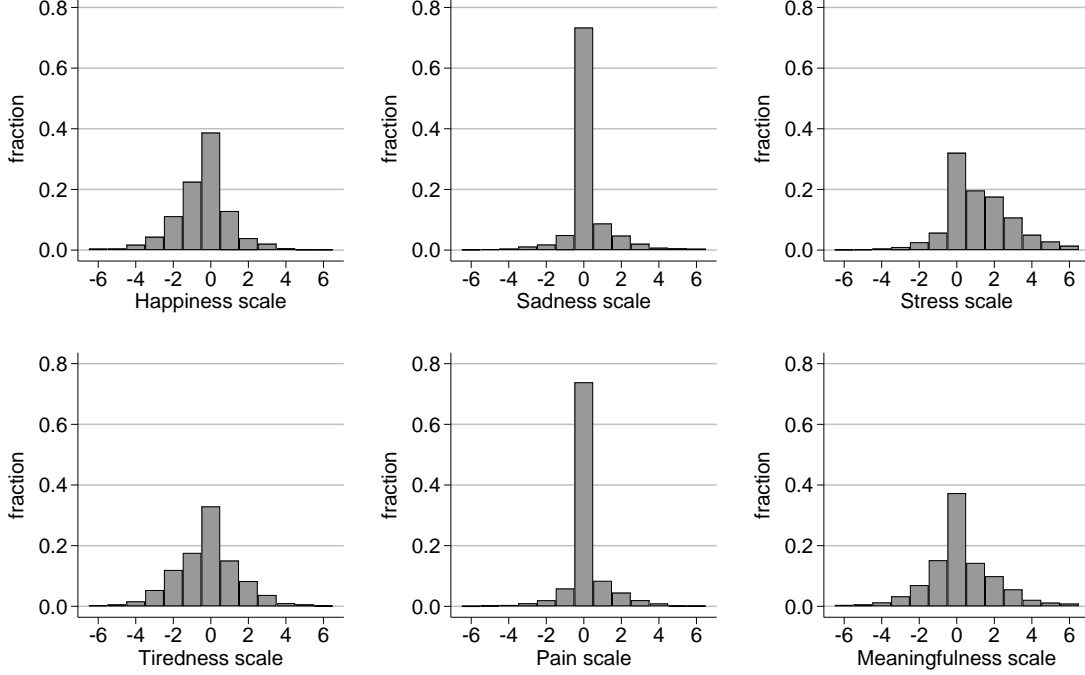


Figure 7: Differences in raw reported feelings between work and non-work activities

Source: Authors' calculations from census, ACS, and ATUS.

observations:

$$z_{ij} = \eta_i + w_{ij} \sum_o \delta_o + (1 - w_{ij}) \sum_n \zeta_n + (1 - w_{ij}) \left(\sum_a \alpha_a + \sum_r \beta_r + \sum_e \gamma_e \right) + \epsilon_{ij}, \quad (\text{A6})$$

where $j \in \{1, 2, 3\}$ indexes observations within a respondent; $w_{ij} = 1$ if activity j is work at the main job and 0 otherwise; and n indexes types of activity other than work on the main job. The main effect of all of the demographics in (A1) is now absorbed into the fixed effect η_i , but we interact the demographics with w_{ij} to allow the possibility that the general difference between feelings at work and away from work differs by demographic group. The fixed-effects-adjusted mean feelings in occupation o are then:

$$\bar{z}_o^{FE} = \bar{\eta} + \hat{\delta}_o^{FE}, \quad (\text{A7})$$

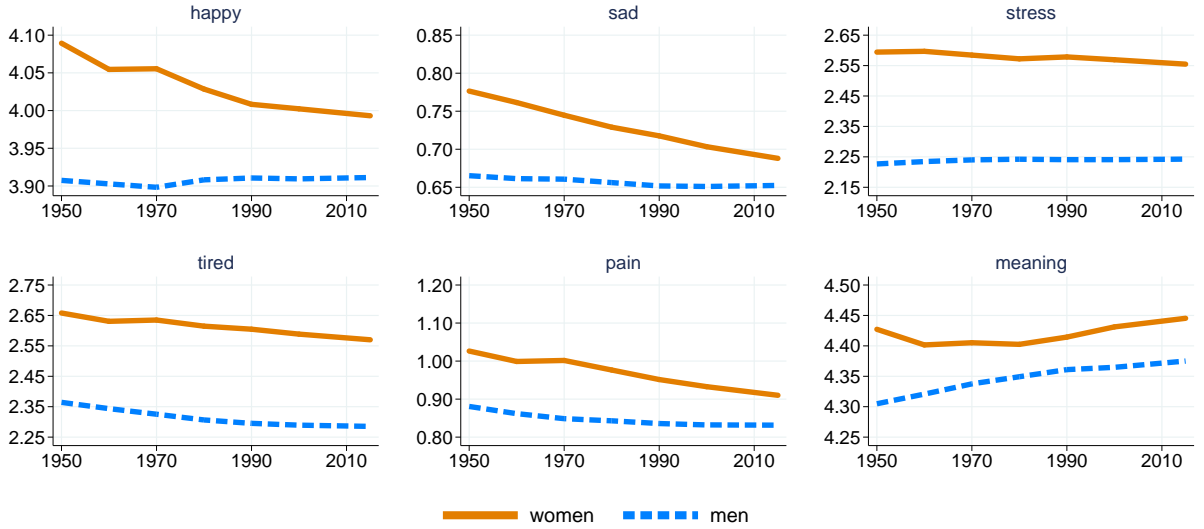


Figure 8: Changes in aggregate feelings at work by sex, fixed-effects-adjusted, 1950-present

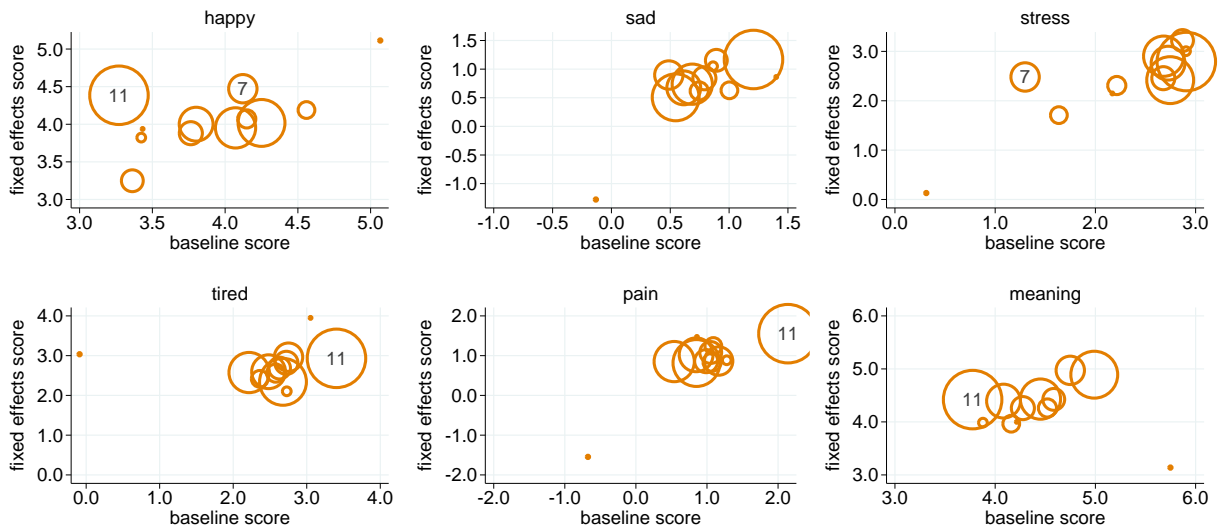
Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex; occupation scores estimated from fixed effects model. Source: Authors' calculations from census, ACS, and ATUS.

where $\bar{\eta}$ is the mean of the estimated individual fixed effects, and we can use \bar{z}_o^{FE} in place of \bar{z}_o in all of our calculations.

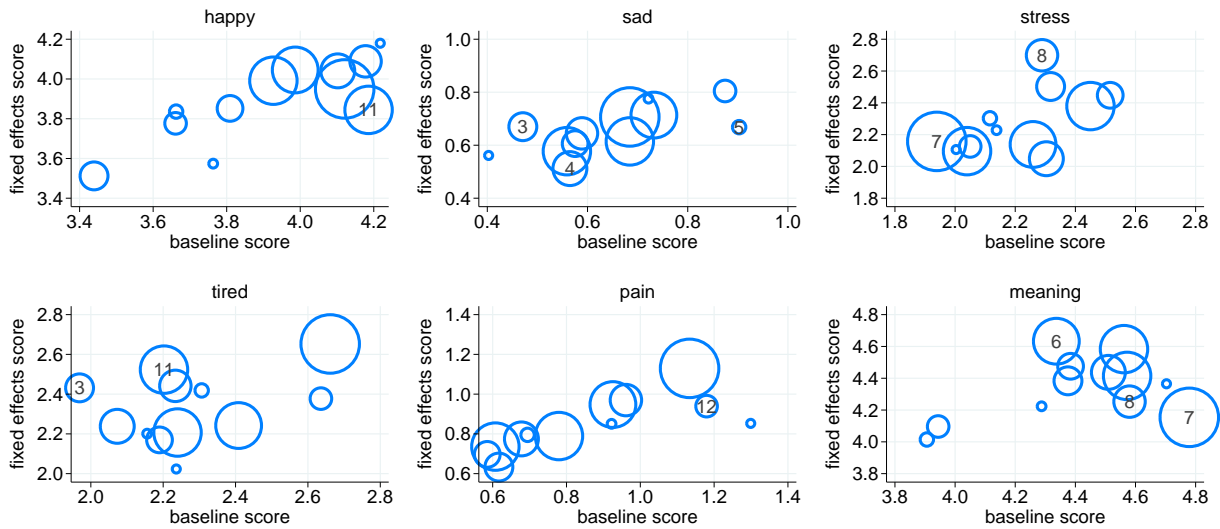
A drawback to the fixed effects strategy is that it assumes the effect of work on a worker's feelings is limited to his or her time at work. If feelings caused by work spill over into non-work activities, the fixed effects estimates could be biased either upward or downward. For example, if having a good job also makes people feel good at home, then we would estimate little or no effect of occupation on individual feelings in (A6) and correspondingly little or no effect of a changing occupation distribution on aggregate feelings. Therefore, if there are positive spillovers from work to home, estimates based on the fixed effects strategy provide a lower bound on how changes in the occupation distribution have changed aggregate feelings. On the other hand, negative spillovers are possible if people who have bad jobs are particularly happy to go home from them or people who have good jobs miss their work when they are at home; in these cases, the fixed-effects estimates would be biased upward.

Figure 8 shows the fixed effects estimates for men and women. The fixed effects

estimates often show different trends from the baseline estimates. For example, happiness fell for women in the fixed effects estimates, compared with a rising trend in the baseline estimates without fixed effects; meaningfulness rose for men in the fixed effects estimates but fell in the baseline estimates. But the fixed effects estimates confirm the finding of a decreased physical toll of work — less pain and less tiredness — for both men and women.



(a) Women



(b) Men

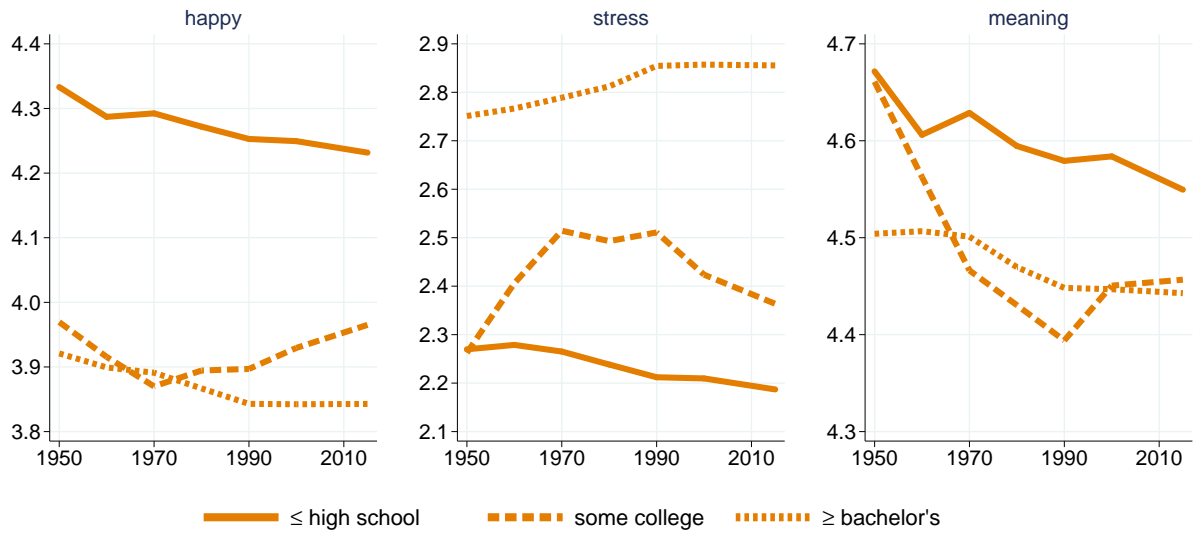
Figure 9: Relationship between fixed-effects and baseline occupation scores

Area of circle is proportional to absolute value of change between 1950 and 2015 in the share of workers in each occupation. Occupation scores and occupation distributions calculated separately by sex. Numbers centered in markers are codes for selected occupations. Codes: 3 = technicians and related support; 4 = sales; 5 = administrative support and clerical; 6 = service; 7 = farming, forestry, and fishing; 8 = precision production; 11 = machine operators, assemblers, and inspectors; 12 = transportation and material moving. Source: Authors' calculations from census, ACS, and ATUS.

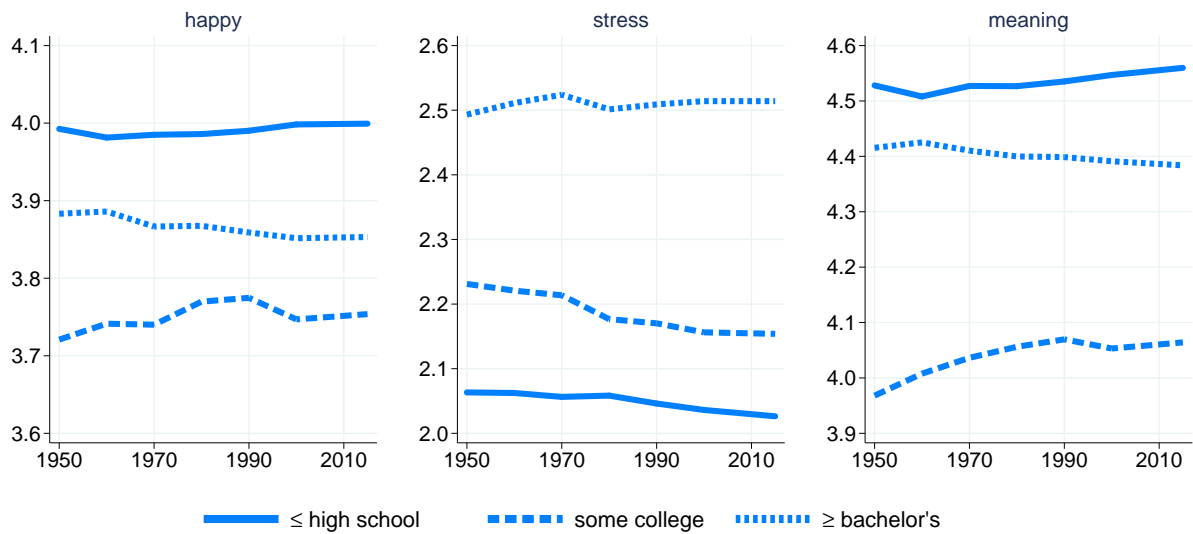
An occupation’s score affects the trend for a given feeling only to the extent that the occupation’s share of the work force changed substantially over time. Thus, to understand why the trends using fixed effects estimates are sometimes different from the baseline results, we examine how the difference between the fixed effects and baseline score for each occupation relates to the change in the occupation’s share of the work force.

Figure 9 plots the fixed effects score against the baseline score for each occupation and uses the size of the markers to highlight occupations whose population shares changed the most between 1950 and 2015. For most occupations whose shares changed substantially, and for most of the types of feelings that we measure, the fixed-effects and baseline occupation scores are closely correlated. This correlation gives some confidence that our basic approach to measuring the feelings induced by an occupation is reasonable. However, there are a few outliers, and these outliers appear to drive the cases where we see different trends in the fixed effects and baseline estimates. For example, among women, the fixed-effects score for machine operators often differs substantially from the baseline score. This occupation was one of the lowest scoring on happiness for women in the baseline but one of the highest scoring on happiness for women in the fixed effects estimates, and it shrank substantially from 1950 to 2015, explaining why women had a downward trend in happiness according to the fixed effects estimates but an upward trend according to the baseline estimates. For men, farming, forestry and fishing were rated very high in meaning in the baseline estimates but quite low in the fixed effects estimates, while service occupations received a moderate meaning score in the baseline and a high score with fixed effects. Thus, the shrinkage of the agricultural sector and the growth of service work implied decreasing meaningfulness for men according to the baseline estimates but rising meaningfulness according to the fixed effects estimates.

Figure 10 breaks the fixed effects estimates down by education level. The differences in trends between the fixed effects estimates and the baseline estimates appear to be concentrated at lower education levels — women with no more than a high school education and men with a high school education or some college. At higher education levels, the fixed effects estimates are similar to the baseline. This result is perhaps unsurprising given that the largest differences between the fixed effects and baseline occupation scores appeared in occupations that have relatively lower average education levels.



(a) Women



(b) Men

Figure 10: Changes in aggregate feelings at work by sex and education, fixed-effects-adjusted, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex and education; occupation scores estimated from fixed effects model. Source: Authors' calculations from census, ACS, and ATUS.

The differences between the fixed effects and baseline estimates suggest a need for caution in interpreting the overall results. As we noted earlier, though, the fixed effects estimates are not a panacea for potential biases in the baseline estimates. The fixed effects estimates themselves may be biased either upward or downward there are any spillovers from what happens at work to how someone feels at home.

A3. Controlling for the Interaction of Age and Education

Controlling for age and education separately when we estimate occupation scores might be insufficient to remove the effects of age and education on feelings if the effect has differed across cohorts, for example if the education system gave different messages to different cohorts about what types of work were meaningful.

As a robustness check, we therefore reestimate the occupation scores with additional controls for the interaction of single year of age and single year of education in (A1). Figures 11 through 15 reproduce the estimates shown in Figures 2 through 6, respectively, using these additional controls. The results with the additional controls are quite similar to those obtained in the baseline estimates.

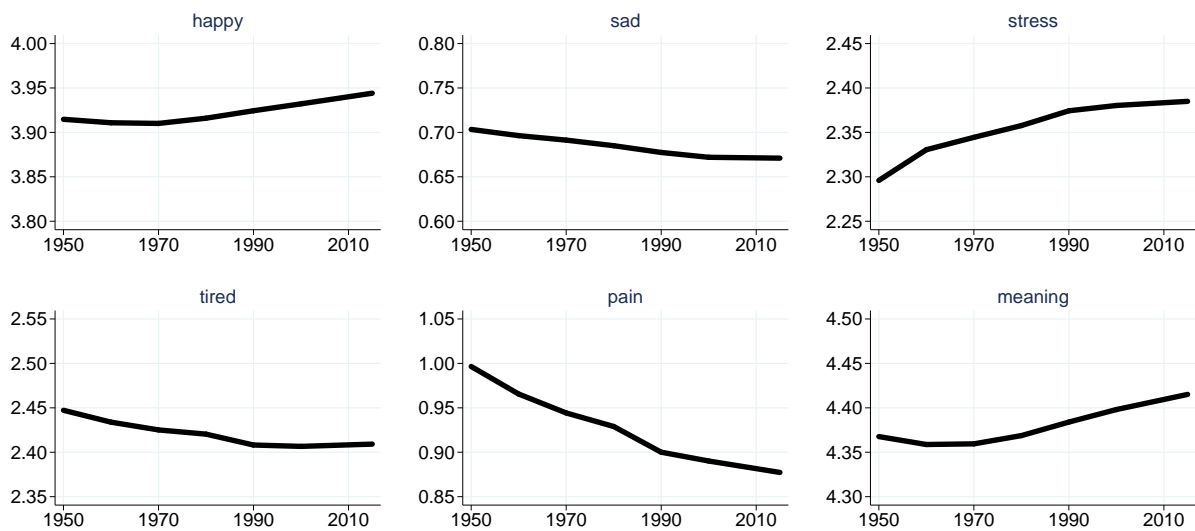


Figure 11: Changes in aggregate feelings at work, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated for full population; occupation scores adjusted for age, race, sex, years of education, and interaction of age and education. Source: Authors' calculations from census, ACS, and ATUS.

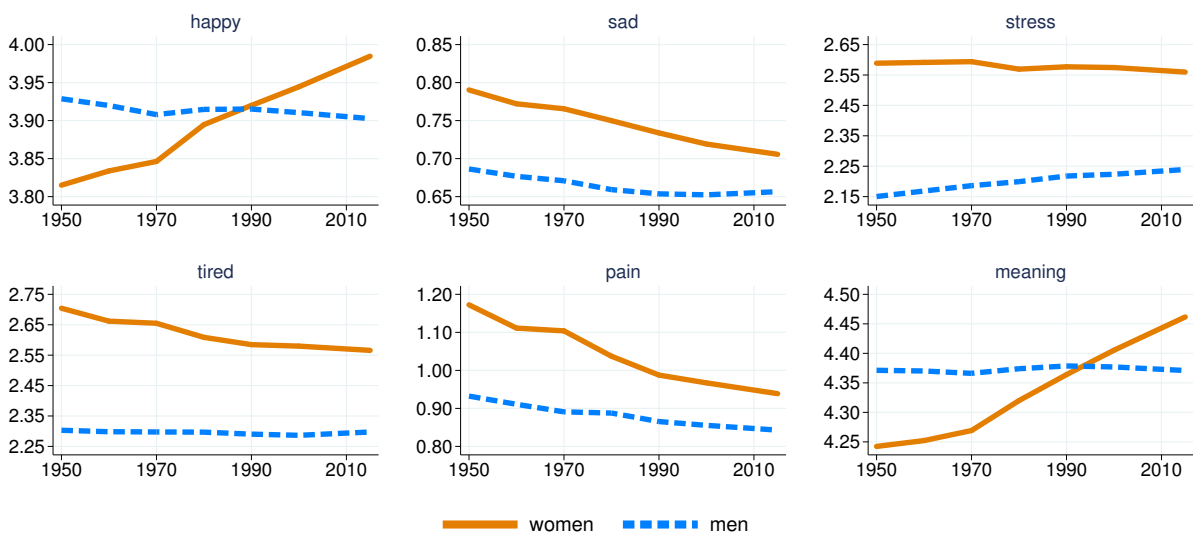
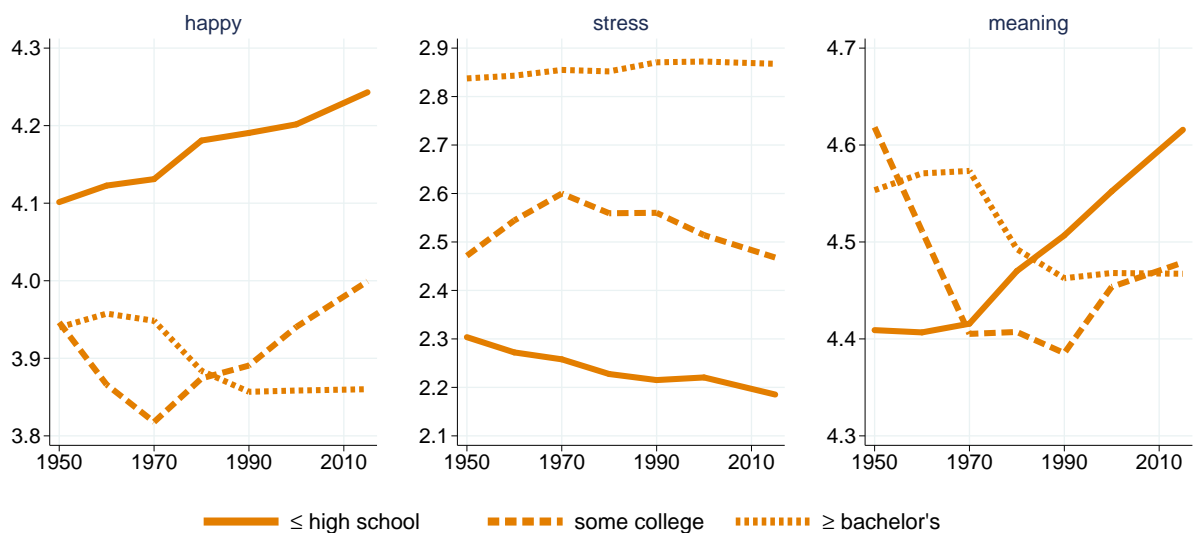
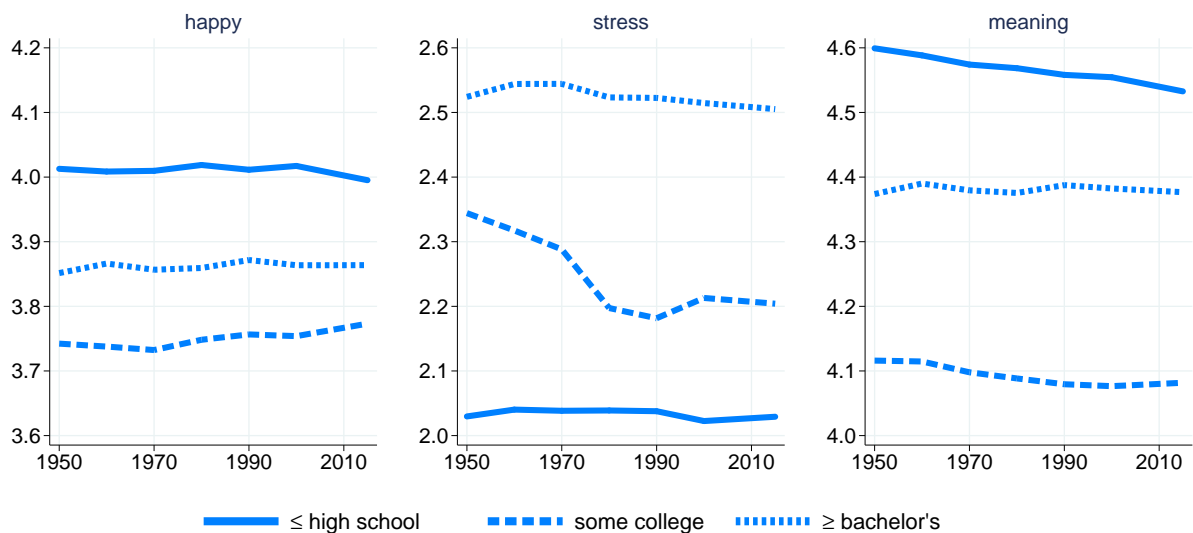


Figure 12: Changes in aggregate feelings at work by sex, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex; occupation scores adjusted for age, race, and years of education. Source: Authors' calculations from census, ACS, and ATUS.



(a) Women



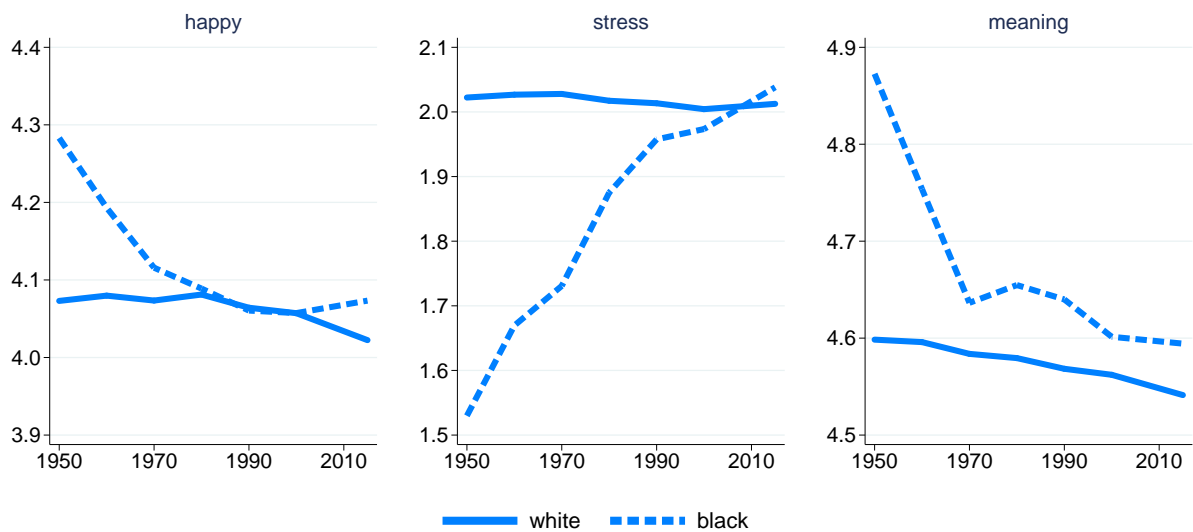
(b) Men

Figure 13: Changes in aggregate feelings at work by sex and education, 1950-present

Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex and education; occupation scores adjusted for age, race and years of education. Source: Authors' calculations from census, ACS, and ATUS.



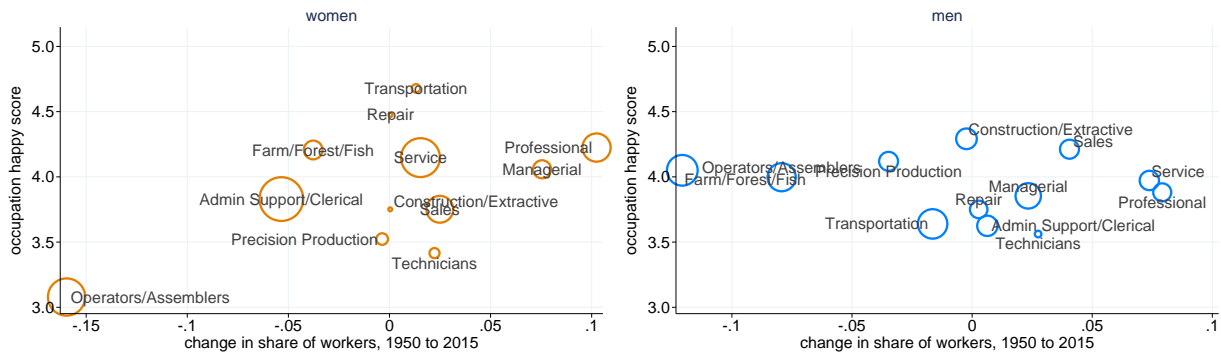
(a) Women



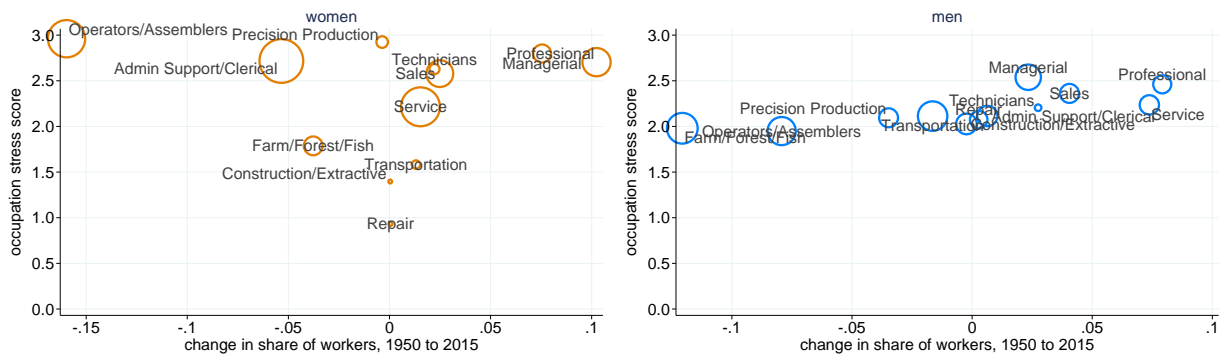
(b) Men

Figure 14: Changes in aggregate feelings at work by sex and race (education \leq high school), 1950-present

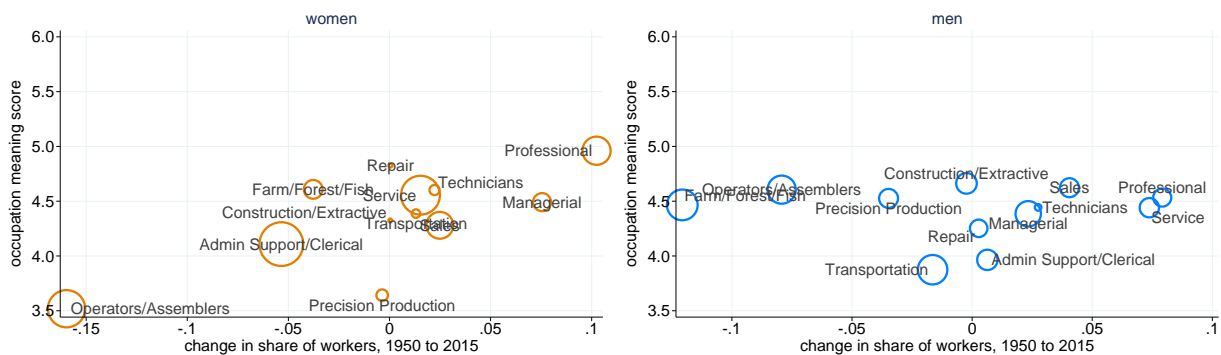
Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated for black and white respondents with no more than a high school education, separately by race and sex; occupation scores adjusted for age and years of education. Other races excluded from calculation. Source: Authors' calculations from census, ACS, and ATUS.



(a) Happy



(b) Stress



(c) Meaning

Figure 15: Changes in occupation share and average feelings by occupation

Area of circle is proportional to share of workers in each occupation in 1950, by sex. Occupation scores and occupation distributions calculated separately by sex; occupation scores adjusted for age, race, and education. Source: Authors' calculations from census, ACS, and ATUS.

References

- Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics, Volume 4B*, ed. David Card and Orley Ashenfelter. Amsterdam: Elsevier.
- Acemoglu, Daron, and Pascual Restrepo. “The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment.” Forthcoming, *American Economic Review*.
- Aguiar, Mark, Mark Bilts, Kerwin Kofi Charles, and Erik Hurst. 2017. “Leisure Luxuries and the Labor Supply of Young Men.” Manuscript, Princeton University.
- Aguiar, M., and E. Hurst. 2016. “The Macroeconomics of Time Allocation,” in *Handbook of Macroeconomics, Volume 2A*, ed. John B. Taylor and Harald Uhlig. Amsterdam: Elsevier.
- Aldrich, Mark. 1997. *Safety First: Technology, Labor and Business in the Building of Work Safety, 1870-1939*. Baltimore: Johns Hopkins University Press.
- Ariely, Dan, Emir Kamenica, and Dražen Prelec. 2008. “Man’s search for meaning: The case of Legos.” *Journal of Economic Behavior & Organization* 67(3), 671–677.
- Autor, David, David Dorn, and Gordon H. Hanson. 2013. “The China syndrome: Local labor market effects of import competition in the United States.” *American Economic Review* 103(6), 2121–2168.
- Autor, David, Lawrence F. Katz, and Melissa S. Kearney. 2006. “The Polarization of the U.S. Labor Market.” *American Economic Review* 96(2), 189–194.
- Bartik, Timothy J. 1987. “The Estimation of Demand Parameters in Hedonic Price Models.” *Journal of Political Economy* 95(1), 81–88.
- Case, Anne, and Angus Deaton. 2017. “Mortality and Morbidity in the 21st Century.” *Brookings Papers on Economic Activity*, Spring 2017, 397–443.
- Centers for Disease Control and Prevention. 1999. “Achievements in Public Health, 1900–1999.” *Morbidity and Mortality Weekly Report* 48(22), 461–469.

- Deming, David J. 2017. “The growing importance of social skills in the labor market.” *Quarterly Journal of Economics* 132(4), 1593–1640.
- Epple, Dennis. 1987. “Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products.” *Journal of Political Economy* 95(1), 59–80.
- Faust, Drew Gilpin. 2008 baccalaureate address, Harvard University. Accessed Jan. 25, 2018, at <https://www.harvard.edu/president/speech/2008/2008-baccalaureate-service>.
- Frey, Darcy. 1996. “Something’s Got to Give.” *The New Times Magazine*, March 24.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2014. “Explaining job polarization: Routine-biased technological change and offshoring.” *American Economic Review* 104(8), 2509–2526.
- Grant, Adam M. “Relational job design and the motivation to make a prosocial difference.” *Academy of Management Review* 32(2), 393–417.
- Grant, Adam M. “Leading with meaning: Beneficiary contact, prosocial impact, and the performance effects of transformational leadership.” *Academy of Management Journal* 55(2), 458–476.
- Güvenen, Fatih, Greg Kaplan, and Jae Song. 2014 “The glass ceiling and the paper floor: Gender differences among top earners, 1981-2012.” Working Paper 20560, National Bureau of Economic Research.
- Güvenen, Fatih, and Greg Kaplan. 2017. “Top Income Inequality in the 21st Century: Some Cautionary Notes.” *Federal Reserve Bank of Minneapolis Quarterly Review* 38(1), 2–15.
- Hall, Jonathan V., and Alan B. Krueger. 2017. “An Analysis of the Labor Market for Uber’s Driver-Partners in the United States.” Working Paper 22843, National Bureau of Economic Research.
- Hamermesh, Daniel S. 2001. “The Changing Distribution of Job Satisfaction.” *Journal of Human Resources* 36(1), 1–30.

- Heathcote, Jonathan, Fabrizio Perri, and Giovanni 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.
- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi. 2014. "Growth and Structural Transformation," in *Handbook of Economic Growth, Volume 2B*, ed. Philippe Aghion and Steven N. Durlauf. Amsterdam: Elsevier.
- Hofferth, Sandra L., Sarah M. Flood, and Matthew Sobek. 2015. American Time Use Survey Data Extract Builder: Version 2.5 [dataset]. College Park, MD: University of Maryland and Minneapolis, MN: University of Minnesota. <http://doi.org/10.18128/D060.V2.5>.
- Kalleberg, Arne L. 2011. *Good Jobs, Bad Jobs: The Rise of Polarized and Precarious Employment Systems in the United States, 1970s to 2000s*. New York: Russell Sage Foundation.
- Katz, Lawrence F., and Alan B. Krueger. 2016. "The rise and nature of alternative work arrangements in the United States, 1995-2015." Working Paper 22667, National Bureau of Economic Research.
- Keynes, John Maynard. 1930. "Economic Possibilities for Our Grandchildren." Reprinted in *Essays in Persuasion*, 358–373.
- Larson, Magali Sarfatti. 1977. *The Rise of Professionalism : A Sociological Analysis*. Berkeley: University of California Press.
- Levine, David I., Michael W. Toffel, and Matthew S. Johnson. 2012. "Randomized Government Safety Inspections Reduce Worker Injuries with No Detectable Job Loss." *Science* 336(6083), 907–911.
- Marmot, M.G., Geoffrey Rose, M. Shipley, and P.J.S. Hamilton. 1978. "Employment Grade and Coronary Heart Disease in British Civil Servants." *Journal of Epidemiology and Community Health* 32(4), 244–249.
- Marmot, M.G., S. Stansfeld, C. Patel, F. North, J. Head, I. White, E. Brunner, A. Feeney, M.G. Marmot, and G. Davey Smith. 1991. "Health inequalities among British civil servants: the Whitehall II study." *The Lancet* 337(8754), 1387–1393.

- Mas, Alexandre, and Amanda Pallais. 2017. “Valuing alternative work arrangements.” *American Economic Review* 107(12), 3722–3759.
- Piketty, Thomas, and Emmanuel Saez. 2003. “Income inequality in the United States, 1913–1998.” *Quarterly Journal of Economics* 118(1), 1–41.
- Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2015. Integrated Public Use Microdata Series: Version 6.0 [dataset]. Minneapolis: University of Minnesota. <http://doi.org/10.18128/D010.V6.0>.
- Schwartz, Barry. 2015. *Why We Work*. New York: Simon & Schuster.
- Smith, Adam. 1776. *The Wealth of Nations*. New York: Bantam Dell, 2003 edition.
- Smith, James P. 1999. “Healthy Bodies and Thick Wallets: The Dual Relation Between Health and Economic Status.” *Journal of Economic Perspectives* 13(2), 145–166.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2017. “Capitalists in the Twenty-First Century.” Manuscript, UC Berkeley and University of Chicago.
- Spitz-Oener, Alexandra. 2006. “Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure.” *Journal of Labor Economics* 24(2), 235–270.
- Taylor, Humphrey. 2001. “Doctors the Most Prestigious of Seventeen Professions and Occupations, Followed by Teachers (# 2), Scientists (#3), Clergy (#4) and Military Officers (#5).” *The Harris Poll* 50. Accessed June 13, 2018, at <https://theharrispoll.com/wp-content/uploads/2017/12/Harris-Interactive-Poll-Research-DOCTORS-THE-MOST-PRESTIGIOUS-OF-SEVENTEEN-PROFESSI-2.pdf>.
- Tinbergen, Jan. 1956. “On the Theory of Income Distribution.” *Weltwirtschaftliches Archiv* 77, 155–175.
- Waring, Marilyn. 1988. *If Women Counted: A New Feminist Economics*. San Francisco: Harper & Row.