

The Changing (Dis-)Utility of Work*

Greg Kaplan and Sam Schulhofer-Wohl

February 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

*Kaplan: University of Chicago and National Bureau of Economic Research (gkaplan@uchicago.edu). Schulhofer-Wohl: Federal Reserve Bank of Chicago (sam@frbchi.org). We thank 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

Conventional aggregate macroeconomic models of labor supply treat all work as the same: 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 aggregate models, the tradeoffs workers make between consumption and leisure can be assessed solely by looking at hours worked and wages. Yet labor economists have known since Adam Smith that not all work is the same and that each job brings a different, complex mix of costs and benefits:

The wages of labour vary with the ease or hardship, the cleanliness or dirtiness, the honourableness or dishonourableness of the employment. Thus in most places . . . a journeyman taylor earns less than a journeyman weaver. His work is much easier. A journeyman weaver earns less than a journeyman smith. His work is not always easier, but it is much cleanlier. . . Honour makes a great part of the reward of all honourable professions. . . Disgrace has the contrary effect. (Smith, 1776, Book I, Chapter X, Part I)

In this paper, we explore empirically the aggregate implications of Smith's idea for the United States over the past seven decades. 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 nonpecuniary rewards or utility that they derive from it, and the distribution of disutility and utility across the population?

We begin with recent data on workers' subjective feelings on the job: how happy, sad, and tired they are; how much stress and pain they say they experience; and how meaningful they find their work. We map these data to occupations, assigning each occupation a vector of scores reflecting its workers' average feelings. Then we ask: If the distribution of occupations were different from what we see today, how would workers' total experiences change? How much more or less stress, for example, would workers experience if they had the distribution of occupations that was observed in 1950, rather than the distribution observed today? How

much more or less meaning, happiness, tiredness, and pain?

We find that work takes a smaller physical toll now than in the past, with workers shifting toward occupations that are associated with feeling less tired and experiencing less pain. There is 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 approach relies on two key assumptions. First, we assume 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. Second, 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. 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 out in any job happen to choose to be managers. We report a robustness check in which we control for the feelings that workers report when they are not on the job, but these controls are imperfect because they can create a bias in the opposite direction, and we have not been able to identify any data that would allow us to fully relax this assumption by providing both exogenous variation in occupational assignments and measures of subjective feelings at work.

We do not attempt to map changes in the vector of feelings people have about their work into changes in a single index of the amount of (dis-)utility they experience or the amount

of work they are doing. In principle, 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 because the observed correlation between occupations’ attributes and their wages does not, in general, reflect workers’ marginal willingness to pay for those attributes, due to the way workers with heterogeneous preferences endogenously sort across occupations (Bartik, 1987; Epple, 1987). In addition, depending on the nature of heterogeneity in preferences, the compensating variation may itself vary across workers. We leave the task of untangling this heterogeneity for future research.

The paper proceeds as follows. Section 2 relates our analysis to the broader literature on the changing nature of the labor market. Section 3 describes the data and the method for computing changes in the (dis-)utility of work over time, section 4 reports the results, section 5 describes the robustness check using feelings off the job, and section 6 concludes.

2. Background

Our work relates to several broader strands of research on changes in the labor market. Our goal in this section is to place our work in context for the reader, rather than to provide a comprehensive literature review.

First, this paper relates to the tremendous changes in the occupational structure of the United States and other developed economies over the last century and a half. Sectoral shifts from agriculture to manufacturing and then to services have dramatically altered the way that people work. For example, relative to 60 years ago, many fewer people today work outdoors or on assembly lines, and many more people work in offices. While much work has been done to measure the wage and employment implications of these changes (see, e.g., Herrendorf, Rogerson, and Valentinyi, 2014), both in the aggregate and in the cross-section, there has been much less work on the non-monetary effects of these sectoral shifts.

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). In particular, there has been strong growth in occupations requiring non-routine tasks and a corresponding decline in occupations that are easily routinized. Research has also emphasized the growing importance of social skills and emotional intelligence relative to cognitive skills (Deming, 2017). 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 still debate over the macroeconomic forces that have led to polarization, with various papers seeking to quantify the relative importance of robots (Acemoglu and Restrepo, forthcoming), offshoring (Goos, Manning, and Salomons, 2014), and trade with China (Autor, Dorn, and Hanson, 2013). There is much less work on the consequences of polarization for workers' well-being, both individually and in the aggregate.

Part of the difficulty in looking beyond wages when measuring the implications of changes in the occupational structure is that the benefits and costs of work itself have changed. On the benefits side, economists predominantly view monetary rewards as by far the most important aspect of jobs and careers. 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). A good job has many non-wage characteristics, most of which are easy to explain but difficult to quantify. Kalleberg (2011), for example, discusses the importance of job security, autonomy, flexibility, and meaning. These features of a job all impart psychological benefits, and 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 unemployment.

One key psychological characteristic of work is its meaning (or lack thereof). 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. This omission is understandable: identity, pride, and meaning are difficult to quantify and are thus hard to incorporate into the empirically driven field of labor economics.

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 (see, e.g., 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, asking how the major occupational shifts in the postwar period have manifested in changes in the meaning and other psychological aspects of work.

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). Two examples that our data allow us to address are the physical tolls of work: pain and tiredness. Work can also 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.

For some dimensions, such as pain, introspection offers an easy answer to the direction of change for most workers in developed countries over the last century — although many jobs even today are undoubtedly physically demanding, in the aggregate 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. But for other dimensions, such as mental stress, introspection does not provide easy answers. Our study offers tentative insights on the

directions of change for these features of work among U.S. workers.

The opportunity cost of time spent working is time spent on other activities. Whereas our work focuses on non-pecuniary aspects of working, other recent research has highlighted changes in the non-pecuniary aspects of not working. Aguiar et al. (2017) attribute a sizable fraction of the decline in employment among low-skilled young men to an increase in the value of not working that resulted from technological progress in video games.

One non-pecuniary cost of many jobs that has attracted recent attention is inflexibility. For some jobs, the cost includes not just the actual time taken to perform the required tasks but also the required commitment to perform the tasks even when the opportunity cost is high. For example, jobs that make it difficult to spend time with children may be particularly costly. 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).

One of the biggest changes in the labor market over the last five decades 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. While we stop short of quantifying the welfare consequences of rising female labor force participation, our findings on the gender-specific non-pecuniary implications of the shifting occupational structure highlight the need to think about these aspects of work when studying the shifting gender composition of the labor force.

Related, our focus here is on market work. We do not attempt to measure people’s feelings about non-market work, or the aggregate consequences of changes in how much non-market work is done and in who does it, in part because it is challenging to define a sharp

boundary between non-market work and leisure in the data available to us. (Is time recorded as “Talking with or listening to household children” the non-market work of child care, or is it the leisure activity of spending enjoyable time with one’s family?) But 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). A complete account of the implications of rising female labor force participation would also need to consider trends in non-market activities and the utility or dis-utility derived from them.

Another major trend in the labor market has been the secular rise in wage and income inequality (Heathcote, Perri, and Violante, 2010), in particular at the top of the distribution (Piketty and Saez, 2003). The analysis of rising inequality has almost entirely abstracted from non-financial factors. Are rising wages at the top of the distribution a compensating differential for particularly unpleasant 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, so that inequality in wages understates the true inequality in the total rewards of work? In particular, the rise in top inequality has, in recent years, been concentrated in small subset of occupations (Güvenen, Kaplan, and Song, 2014). These occupations include arguably high-stress positions in senior management, as well as jobs on Wall Street that are seen as lacking in meaning even by some of the many Ivy League graduates who take them (Faust, 2008). But entrepreneurs and small-business owners also comprise a large share of the top of the income distribution (Güvenen and Kaplan, 2017; Smith et al., 2017), and while entrepreneurship almost certainly entails unique stresses, it also may impart significant non-pecuniary benefits — autonomy, flexibility, and, perhaps, greater meaning. Our analysis by education groups gives some preliminary insights about inequality in the non-pecuniary aspects of work, but our data do not allow us to look directly at top inequality, and above all our work highlights the need for more measurement of inequality in characteristics of work beyond income.

Besides these contemporary debates, our work also points to a new perspective on historical views of the labor market. 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 was not 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. As it stands, however, there have not been attempts to measure the aggregate consequences of non-wage aspects of work over long periods of time. The goal of this paper is to take a first, small step in that direction.

3. Data and Methodology

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.

A. 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 IPUMS¹ to do this. OCC1990 is based on the occupation codes used in the

¹See <https://usa.ipums.org/usa-action/variables/OCC1990>.

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 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).

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

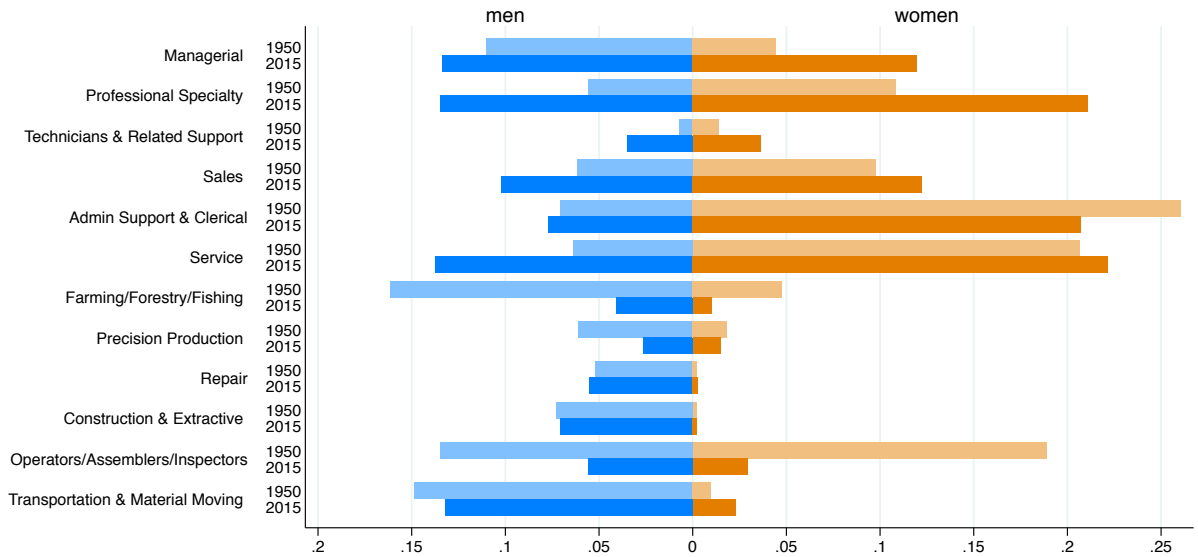


Figure 1: The distribution of workers across broad occupation groupings.

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

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.

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.

B. Feelings

The ATUS, produced by the U.S. Census Bureau, is a stratified random sample of the U.S. population ages 16 and older.² 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

²ATUS respondents are a subset of respondents to the Current Population Survey.

occupations. Specifically, we regress self-reported feelings on dummy variables for single year of age, race, single year of education, and occupation categories:

$$z_i = \theta_0 + \sum_a \alpha_a + \sum_r \beta_r + \sum_e \gamma_e + \sum_o \delta_o + \epsilon_i, \quad (1)$$

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:

$$\bar{z}_o = \hat{\theta}_0 + \sum_a \hat{\alpha}_a s_a + \sum_r \hat{\beta}_r s_r + \sum_e \hat{\gamma}_e s_e + \hat{\delta}_o, \quad (2)$$

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 linearly removing demographic differences as in (1) 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 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.

C. 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 occupation in year t :

$$\hat{z}_{x,t} = \sum_o \pi_{xo,t} \bar{z}_{xo}. \quad (3)$$

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:

$$\hat{z}_{x,now} = \sum_o \pi_{xo,ACS} \bar{z}_{xo}, \quad (4)$$

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:

$$\Delta \hat{z}_{x,t} = \hat{z}_{x,now} - \hat{z}_{x,t} = \sum_o (\pi_{xo,ACS} - \pi_{xo,t}) \bar{z}_{xo}. \quad (5)$$

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.

4. Results

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, work now makes people 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.

Figure 3 shows that the results change substantially when we calculate aggregate mean feelings separately by sex. For women, the story is one of consistently improving feelings about work. Over time, for women, 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 occupations, such as including both restaurant managers and CEOs in the managerial category. We can attempt to more finely classify occupations, despite the lack of perfectly uniform

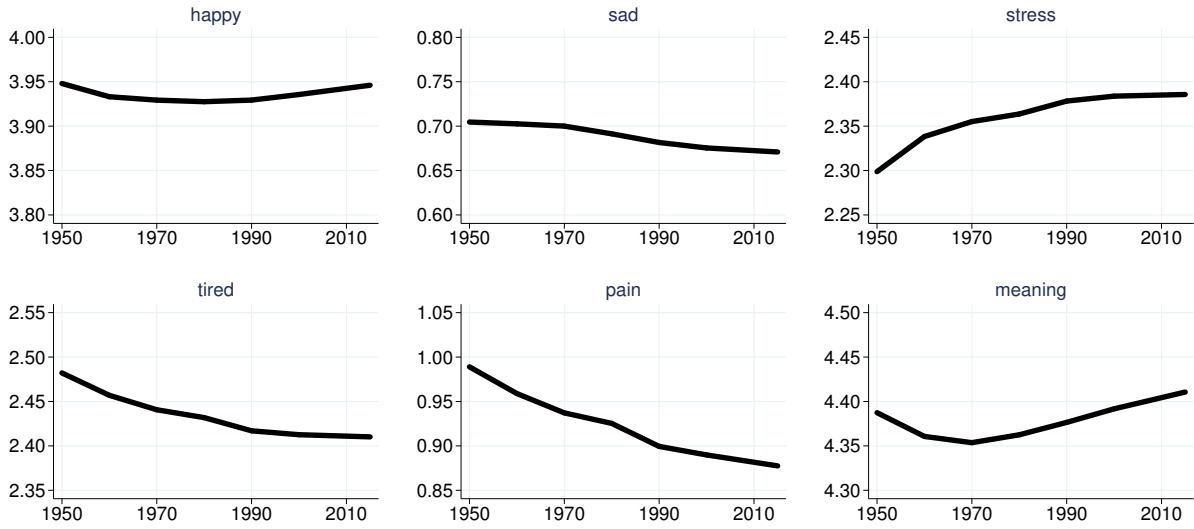


Figure 2: Changes in aggregate feelings at work, 1950-present

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.

coding across years, if we divide the sample by education. Recall that, when we do this, we re-estimate the adjusted mean feelings within each occupation, using only data on workers at a given education level.

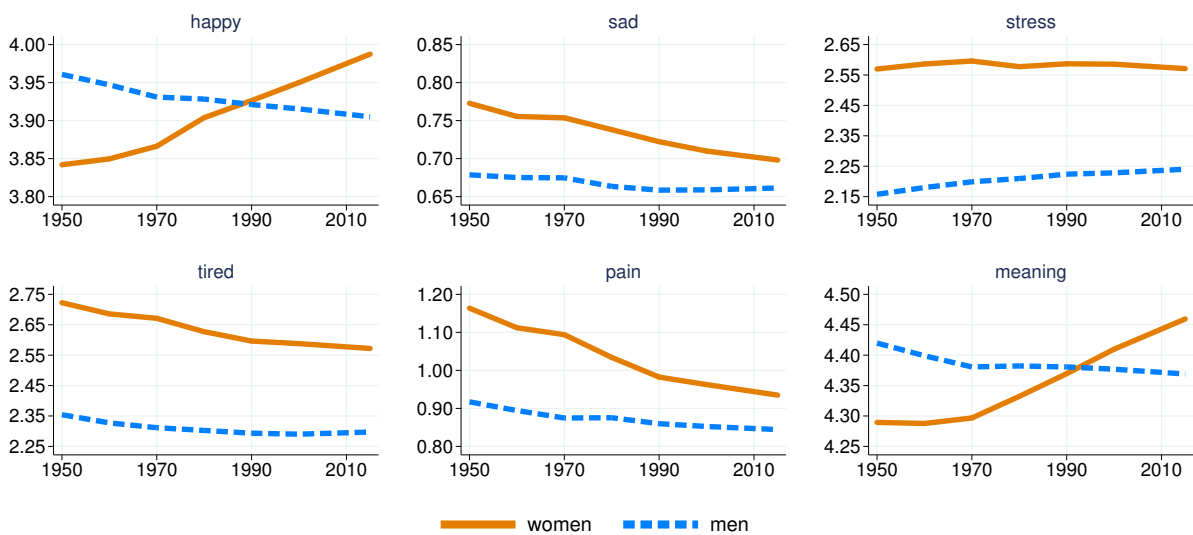
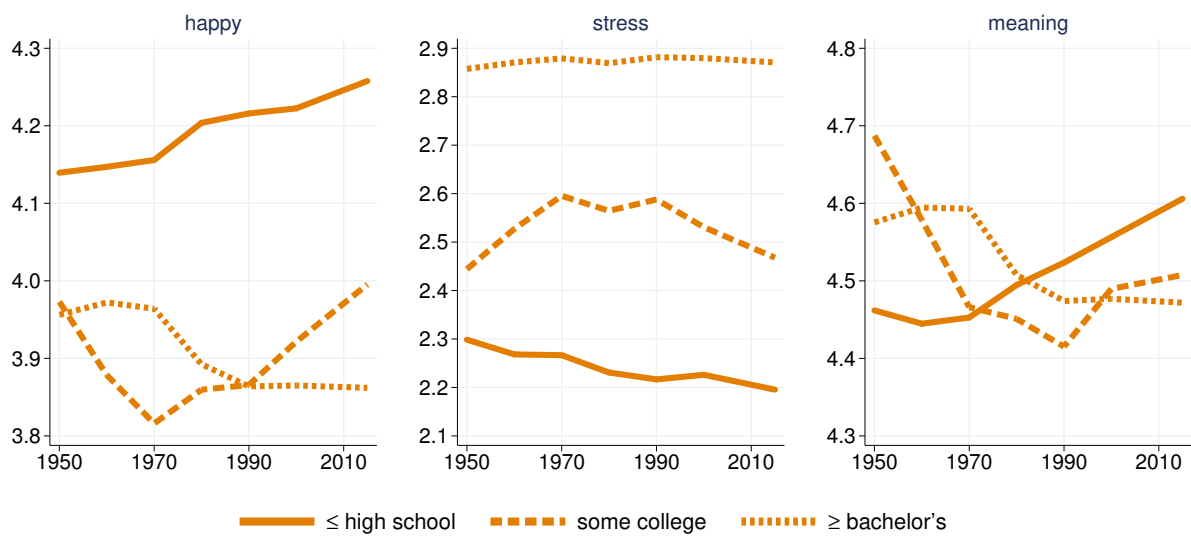
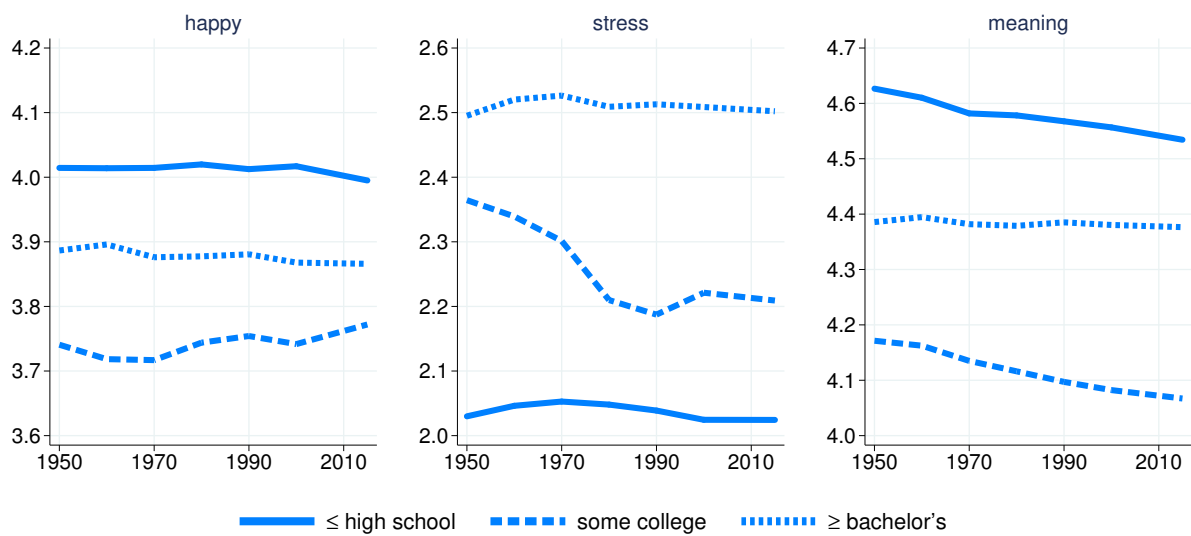


Figure 3: Changes in aggregate feelings at work by sex, 1950-present

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.



(a) Women



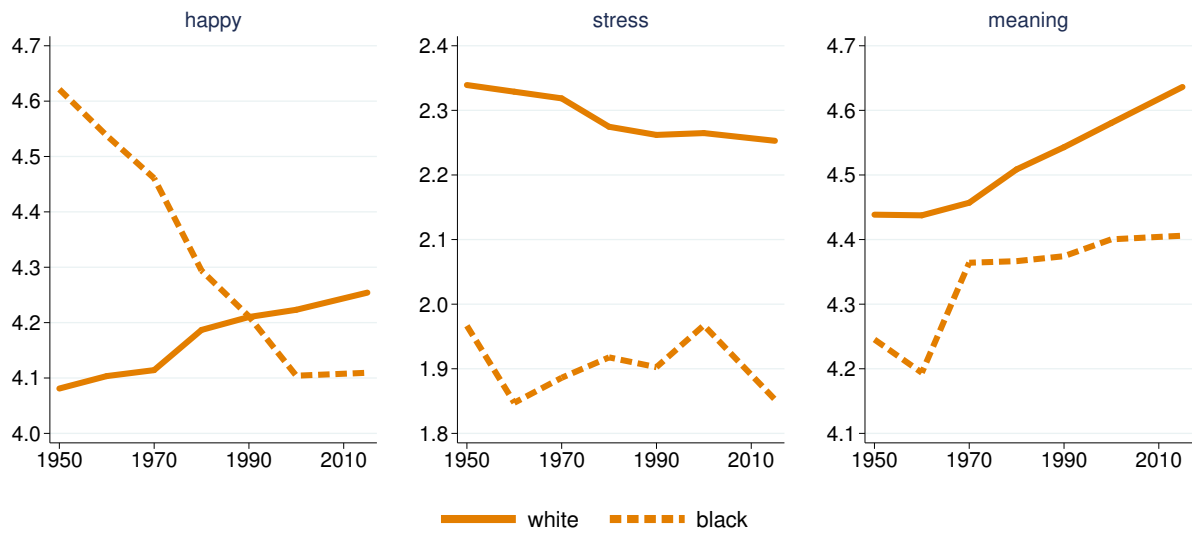
(b) Men

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

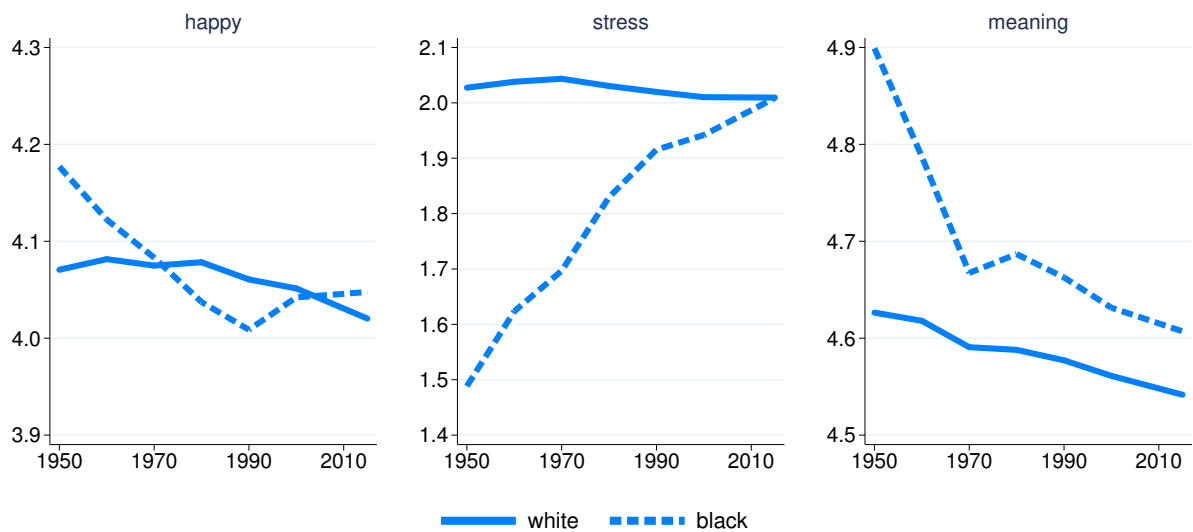
Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex and education; occupation scores adjusted for age, race and years of education. Source: Authors' calculations from census, ACS, and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

Figure 4 shows how happiness, stress, and meaningfulness have evolved when we divide the sample by education as well as 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. However, this conclusion is tentative, partly because disaggregating by education means we are using a smaller sample to estimate each occupation score, and partly because 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 (see, e.g., 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 further 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, and we focus on estimates for people with a high school education or less. Figure 5 shows the results. 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.



(a) Women



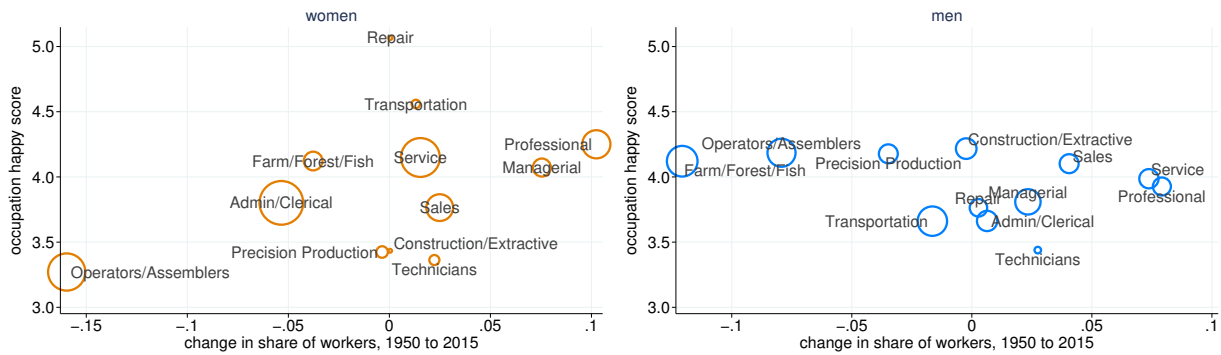
(b) Men

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

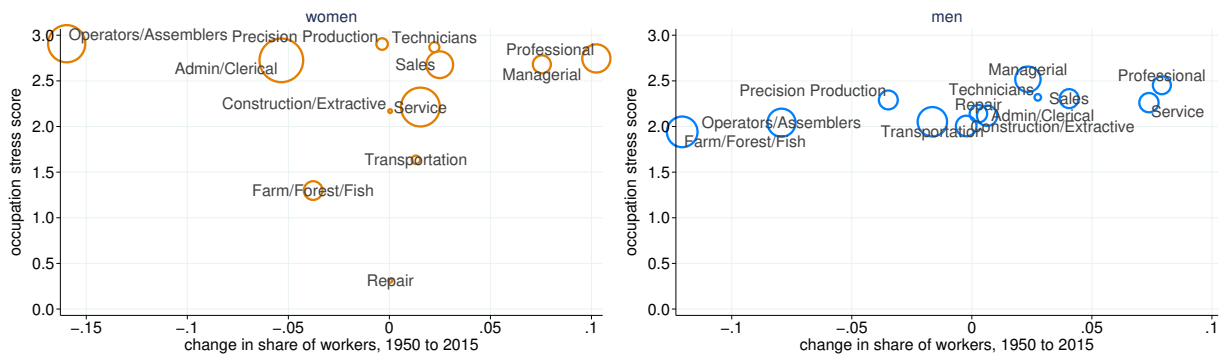
Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated for black and white respondents with no more than a high school education, separately by race and sex; occupation scores adjusted for age and years of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

We next examine the sources of the aggregate shifts by plotting the relationship between an occupation's average feelings, \bar{z}_{xo} , and the change in the share of workers in that occupation since 1950, $\pi_{xo,ACS} - \pi_{xo,1950}$. 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.

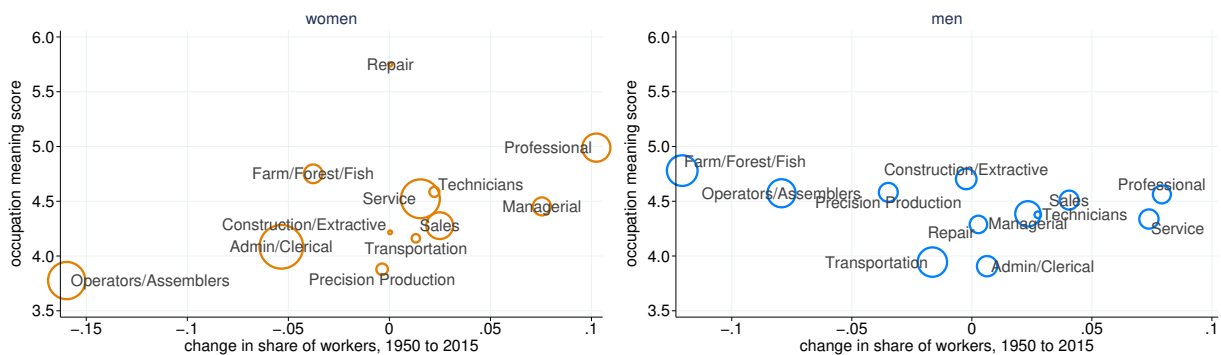
The figure shows that the different results for men and women arise not only from differences in how their occupation distributions have changed but also from differences in the feelings they report in the same occupational categories. For example, both men and women are less likely now than in 1950 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. 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.



(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.

5. Robustness: controlling for individual fixed effects

As we mentioned in the introduction, 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.

We implement this idea by running the regression in (1) 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 observations:

$$z_{ij} = \eta_i + w_{ij} \sum_o \delta_o + (1 - w_{ij}) \sum_n \zeta_n + (1 - w_{ij}) \left(\sum_a \alpha_a + \sum_r \beta_r + \sum_e \gamma_e \right) + \epsilon_{ij}, \quad (6)$$

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 (1) 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:

$$\bar{z}_o^{FE} = \bar{\eta} + \hat{\delta}_o^{FE}, \quad (7)$$

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, on the other hand, feelings caused

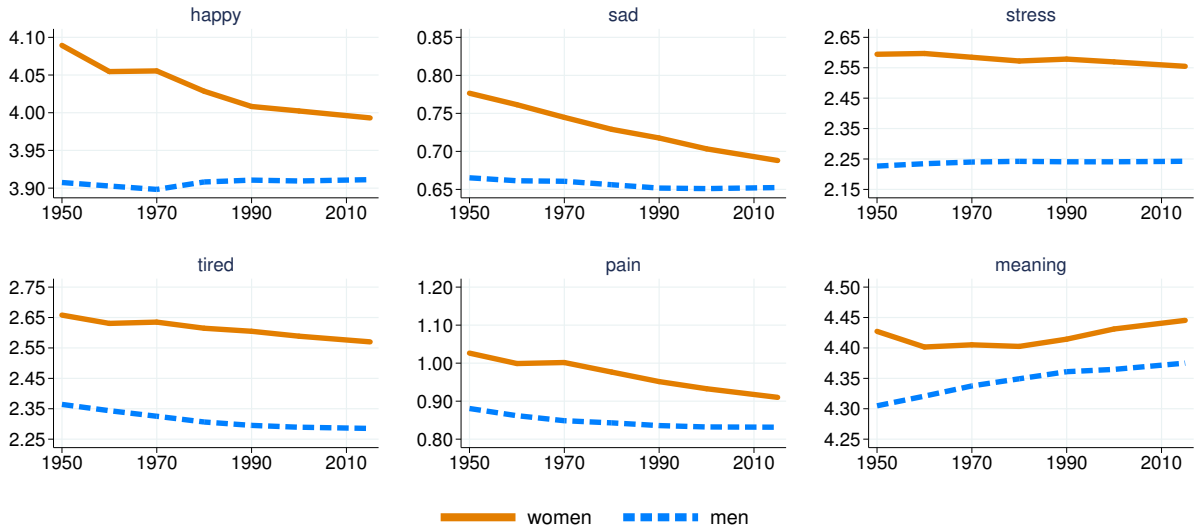
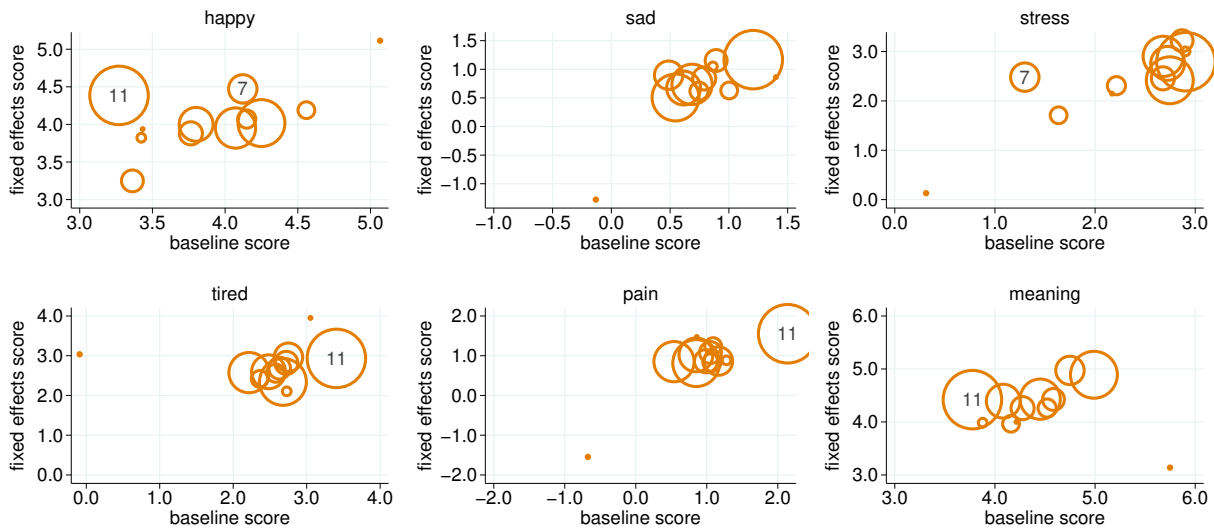


Figure 7: Changes in aggregate feelings at work by sex, fixed-effects-adjusted, 1950-present

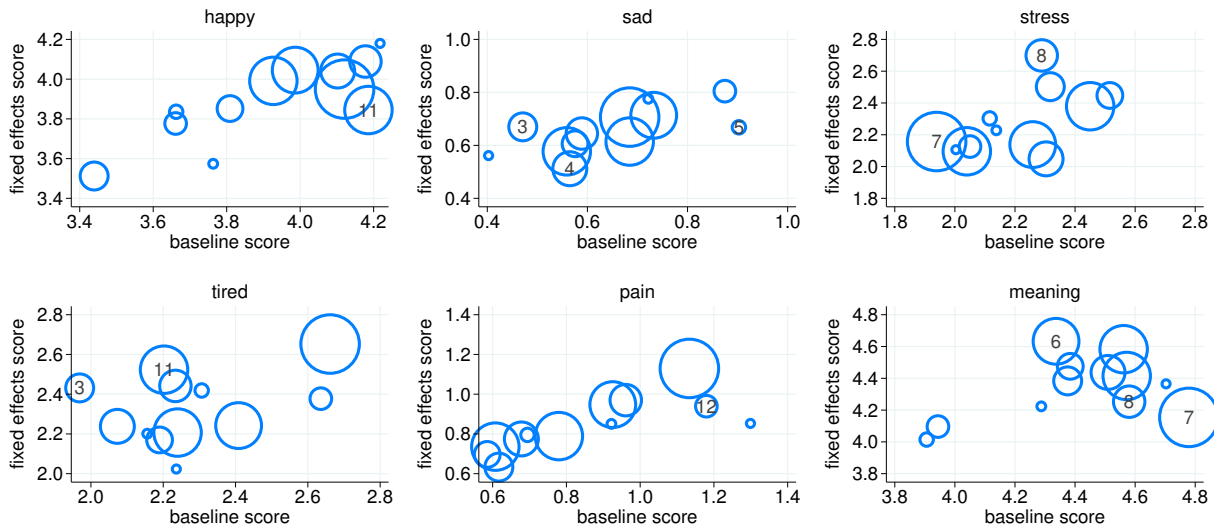
Lines show average of occupation scores weighted by distribution of occupations in each year. Occupation scores and occupation distributions calculated separately by sex; occupation scores estimated from fixed effects model. Source: Authors' calculations from census, ACS, and ATUS.

by work spill over into non-work activities, we would estimate little or no effect of occupation on individual feelings in (6) and correspondingly little or no effect of a changing occupation distribution on aggregate feelings. Therefore, estimates based on the fixed effects strategy provide a lower bound on how changes in the occupation distribution have changed aggregate feelings.

Figure 7 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

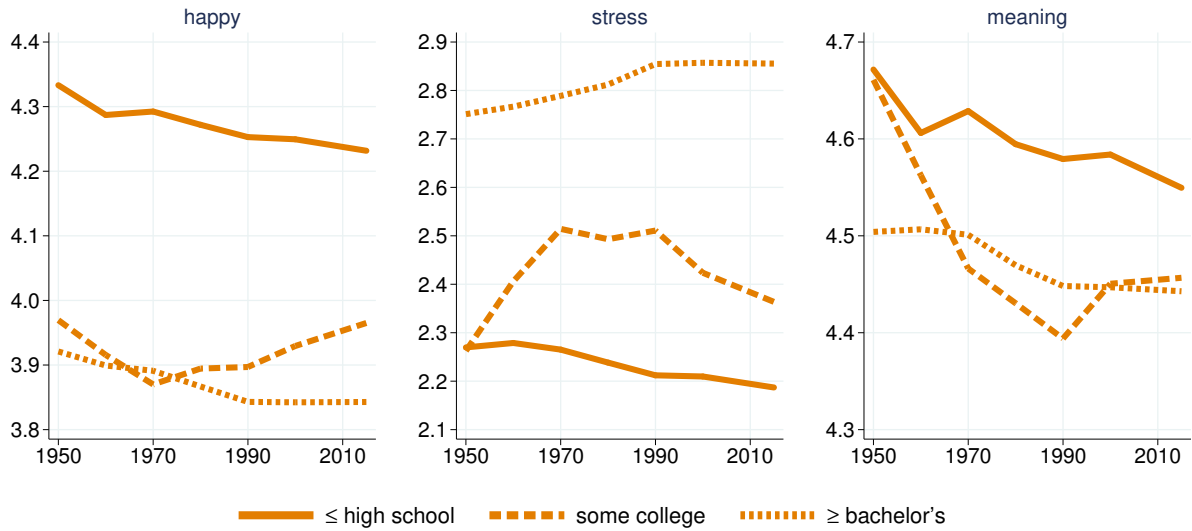
Figure 8: Relationship between fixed-effects and baseline occupation scores

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.

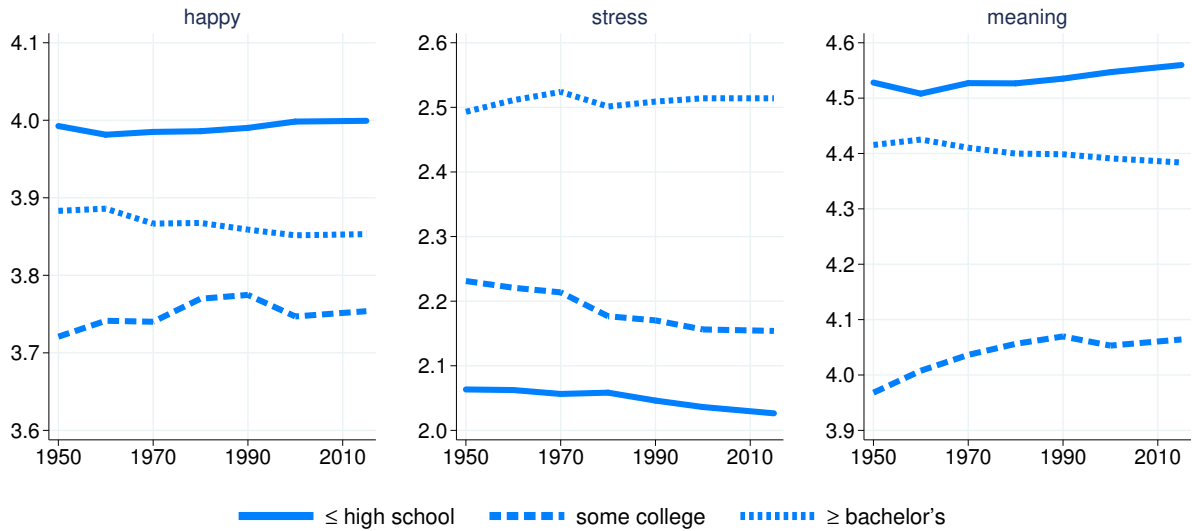
An occupation's score affects the trend for a given feeling only to the extent that the occupation's share of the work force changed substantially over time. Thus, to understand why the trends using fixed effects estimates are sometimes different from the baseline results, we examine how the difference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

Figure 8 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 9 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 9: 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 will tend to understate the impact of occupation on a person’s feelings to the extent that there are any spillovers from what happens at work to how someone feels at home.

6. Conclusion

Macroeconomists commonly view all work as the same and study changes in the total amount of work performed, the wages earned for that work, and inequality in who does the work and earns the wages. This paper instead studies inequality in the nature of work itself — differences in the utility or disutility that various occupations produce. Because the occupation distribution has changed significantly in the postwar period, and because people in different demographic groups appear to feel differently about the same occupations, we find that there have been substantial shifts both in the aggregate (dis-)utility derived from work and in the distribution of that (dis-)utility across people. We leave for future research the question of how far these changes can go in helping to explain the large secular changes in employment and labor force participation over the postwar period that are more traditionally the object of macroeconomic labor research.

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 Bilz, 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.

- 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.
- 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.
- 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, 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.
- 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.