6-2008

Happiness Inequality in the United States

Betsey Stevenson
University of Pennsylvania

Justin Wolfers
University of Pennsylvania

Recommended Citation

Stevenson, B., & Wolfers, J. (2008). Happiness Inequality in the United States. *The Journal of Legal Studies, 37* (2), 533-579. http://dx.doi.org/10.1086/592004

Happiness Inequality in the United States

Abstract

This paper examines how the level and dispersion of self-reported happiness has evolved over the period 1972–2006. While there has been no increase in aggregate happiness, inequality in happiness has fallen substantially since the 1970s. There have been large changes in the level of happiness across groups: two-thirds of the black-white happiness gap has been eroded, and the gender happiness gap has disappeared entirely. Paralleling changes in the income distribution, differences in happiness by education have widened substantially. We develop an integrated approach to measuring inequality and decomposing changes in the distribution of happiness, finding a pervasive decline in within-group inequality during the 1970s and 1980s that was experienced by even narrowly defined demographic groups. Around one-third of this decline has subsequently been unwound. Juxtaposing these changes with large increases in income inequality suggests an important role for nonpecuniary factors in shaping the well-being distribution.

Keywords

happiness, happiness inequality, United States, well-being distribution, black-white happiness gap, gender happiness gap, income distribution

Disciplines

Economics | Income Distribution | Public Affairs, Public Policy and Public Administration

Happiness Inequality in the United States

Betsey Stevenson and Justin Wolfers

ABSTRACT

This paper examines how the level and dispersion of self-reported happiness has evolved over the period 1972–2006. While there has been no increase in aggregate happiness, inequality in happiness has fallen substantially since the 1970s. There have been large changes in the level of happiness across groups: two-thirds of the black-white happiness gap has been eroded, and the gender happiness gap has disappeared entirely. Paralleling changes in the income distribution, differences in happiness by education have widened substantially. We develop an integrated approach to measuring inequality and decomposing changes in the distribution of happiness, finding a pervasive decline in within-group inequality during the 1970s and 1980s that was experienced by even narrowly defined demographic groups. Around one-third of this decline has subsequently been unwound. Juxtaposing these changes with large increases in income inequality suggests an important role for nonpecuniary factors in shaping the well-being distribution.

1. INTRODUCTION

It is now widely understood that average levels of happiness have failed to grow in the United States, despite ongoing economic growth (Easterlin 1995; Blanchflower and Oswald 2004). Yet an average can hide as much as it reveals, and so our task in this paper is to explore the full distribution of happiness through time.

Previous authors have documented the existence of happiness in-

BETSEY STEVENSON is at the Wharton School at the University of Pennsylvania, the Center for Economic Studies and Ifo Institute for Economic Research (CESifo), and the National Bureau of Economic Research (NBER). JUSTIN WOLFERS is at the Wharton School at the University of Pennsylvania, the Centre for Economic Policy Research (CEPR), CESifo, the Institute for the Study of Labor (IZA), and NBER. The authors would like to thank Wharton's Zicklin Center for Business Ethics Research and Zell/Lurie Real Estate Center for generous research support.

[*Journal of Legal Studies,* vol. 37 (June 2008)]
 0047-2530/2008/370S2-0003$10.00

equality both within and between demographic groups: the rich are typically happier than the poor; the educated are happier than those with less education; whites are happier than blacks; those who are married are happier than those who are not; and women—at least historically—have been happier than men. These differences are likely interrelated, and in addition, there exists substantial happiness inequality even within narrowly defined demographic groups. We seek to document how each of these factors is changing and how the changing composition of the U.S. population may be contributing to the observed aggregate trends.

The parallel literature on income inequality certainly suggests that this may be a fruitful task, as recent decades have witnessed the partial closure of gender and race gaps, an increase in education and age gaps, and a substantial increase in income inequality within most demographic groups. All told, this literature suggests that the gains from recent economic growth have been quite unevenly distributed. Beyond the pecuniary domain, there have also been important changes in the legal and institutional organization of work, family, leisure, and community life as well as technological changes that may have impacted well-being.

As with previous analyses, we find that, on average, happiness has failed to grow since the 1970s. But beneath this average, we document some important striking differences across groups: two-thirds of the black-white happiness gap has been eroded, and the gender happiness gap has disappeared entirely, with more recent data suggesting that it may even have inverted. Paralleling changes in the income distribution, differences in happiness by education have widened substantially.

Our more striking finding is the substantial decrease in happiness inequality through our sample. We document that the dispersion of happiness fell sharply in the 1970s and 1980s; subsequently, about one-third of this decline has subsequently been unwound. Our decomposition exercise suggests that the real reason for today's lower levels of happiness inequality is not to be found in the relative experiences of particular groups, or the specific experiences of only a few, but rather in a pervasive decline in within-group inequality experienced by even narrowly defined demographic groups.

Beyond these substantive findings, our approach to measuring inequality and our decomposition of changes in the distribution of happiness within and between groups may be of methodological interest, especially for those interested in analysis of ordinal data. Our integrated approach to estimating levels and dispersion of happiness through time and the accounting framework for describing the proximate sources of these changes may also prove to be useful for analyses of other qualitative or attitudinal data.

Before proceeding, it is worth putting our findings into a broader context. In terms of our empirical objectives, our goal is to describe the data rather than to point to causal links. We juxtapose decreasing happiness inequality with rising income inequality not because we believe that this reflects a clear link between the two but rather because, jointly considered, they hint at the intriguing possibility of a decline in inequality in the nonpecuniary domain.

The normative implications of our results are also somewhat limited. For instance, a committed utilitarian cares only about the average level of well-being, and not inequality in well-being. While the usual utilitarian argument for valuing inequality rests on the view that redistribution yields utility gains to the poor that exceed the costs to the rich, this argument—based as it is on diminishing marginal utility—may be more convincing in the pecuniary domain than when evaluating happiness. We should also add that the usual caveats about well-being data apply, and the mapping between true subjective experiences and responses to subjective well-being questions remains quite poorly understood. (Kahneman and Krueger [2006] provide a useful overview of the relevant literature.)

Our findings contribute to the much broader (positive) literature on trends in well-being, and particularly inequality, in the United States. As such, Section 2 provides the broader context, describing trends in economic inequality and in particular its ongoing rise since the 1970s. We also note that there have been a host other social and legal changes that may have had interesting distributional impacts, including changes in the distribution of leisure, regulation impacting families, antidiscrimination legislation, violent crime, and affirmative action.

In Section 3 we highlight the aggregate trends in happiness—both levels and dispersion—and introduce our approach to cardinalizing these descriptive survey responses. Section 4 turns to examining happiness both within and between groups, to assess how the distribution and dispersion of happiness is changing across socioeconomic and demographic lines. We measure changes in a variety of dimensions and assess their joint impact through a decomposition exercise that points to the importance of within-group increases in happiness inequality. Section 5 provides a concluding discussion.

2. BACKGROUND: TRENDS IN INEQUALITY

Income inequality has increased throughout our sample period, which begins in 1972. During the 1970s inequality rose modestly, with the rise stemming largely from changes in residual inequality (Goldin and Katz 2008). The college wage premium fell through the 1970s, which mitigated against larger rises in overall inequality. However, the college wage differential rebounded during the 1980s, and in this decade inequality rose sharply and throughout the distribution. The rise in inequality in the 1990s and early 2000s was concentrated in the top of the income distribution, with the differential between wages at the 90th and 50th percentiles rising through 2005, despite a decrease in 50–10 inequality (Autor, Katz, and Kearney 2008).

Between 1972 and 2006 (our sample period), overall income inequality rose both within and between groups, with an important part of the rise coming from changes in the returns to education, with the education returns both increasing and increasing by a greater amount for higher levels of education (Goldin and Katz 2007). For example, the weekly earnings of full-time, full-year workers with education beyond a college degree rose 34 log points relative to their counterparts with only a high school degree, while the parallel rise for those with only a college degree was 22 log points (Goldin and Katz 2008, p. 139). Over this period wage dispersion also increased within demographic and skill groups (Autor, Katz, and Kearney 2008). In contrast, wage differentials between some groups have narrowed. Specifically, male-female wage inequality narrowed during the 1970s and 1980s and continued to narrow, albeit more slowly, in the 1990s (Blau and Kahn 2006). The black-white wage gap has also narrowed over the past 35 years, with convergence in the 1970s and, after stagnating in the 1980s, further narrowing in the 1990s (Couch and Daly 2003).

Along with this rise in income inequality has come concerns about increasing income volatility and a more general concern about increasing inequality stemming from households bearing more health and retirement risk (Hacker 2006). Income inequality has occurred through both an increase in the dispersion of permanent income and an increase in transitory income volatility (Gottschalk and Moffitt 1994). However, more recent work has argued that increases in income volatility have impacted few families and have not been broadly experienced (Jensen and Shore 2008).

Since households may be able to use insurance, borrowing, or savings, consumption is less variable than income, and it may better reflect material well-being. As such, many studies of economic inequality have turned to measures of consumption inequality, finding evidence of a parallel increase in consumption inequality in the 1980s (Johnson and Shipp 1997; Cutler and Katz 1991; Attanasio, Battistin, and Ichimura 2004). Some authors have argued, however, that consumption inequality was flat or declining in the 1990s and that the overall rise has been small relative to the rise in income inequality (Krueger and Perri 2006). The rise in consumption inequality has occurred both between and within skill groups (Attanasio, Battistin and Ichimura 2004), although this point has been debated in the literature, with Krueger and Perri (2006) suggesting that there was only minimal growth in consumption inequality within skill groups, despite large increases in within-group income inequality. These same authors find that between-group changes in consumption inequality have been similar to the between-group changes in income inequality. Countering this, Johnson and Shipp (1997) argue that most of the rise in consumption inequality has been within groups.

More recently we have learned that leisure—time devoted to neither market nor household work—is another important domain in which inequality has changed over recent decades (Aguiar and Hurst 2007). In particular, the new leisure class is composed of the low skilled who have experienced steady increases in leisure hours over the past 3 decades. While the high skilled experienced an increase in leisure in the 1970s and 1980s, in recent years this increase has been reallocated toward home production. Yet the biggest rises in nonwork, nonhousehold production hours have been focused on the unemployed and disabled low-skilled men, and most of the increase in “leisure” among those with less education is due to changes in employment status. In contrast, among men with more education, the decline in leisure is due to changes within employment status (Aguiar and Hurst 2008).

There have also been important legal changes impacting equality of opportunity. A vast array of legislation and court rulings have coincided with changes in social norms to reduce discrimination and allow individuals to make life choices with fewer restrictions due to characteristics such as race, religion, gender, or sexual preference. Most notably, the 1964 Civil Rights Act outlawed segregation and discrimination against people on the basis of religion, race, national origin, or gender. This legislation has continued to impact hiring and firing decisions, with substantial growth in employment discrimination litigation in the decades since its passage (Donohue and Siegelman 1991). In more recent years, antidiscrimination legislation has expanded to include the disabled with the Americans with Disabilities Act of 1990 and the Family and Medical Leave Act of 1993, which protects individuals against job loss in the case of short-term medical or family issues.

Families have also gained more autonomy over family life with a wave of large-scale deregulation of the family beginning in the 1960s that diminished the role that government plays in family life. A series of state legislative changes and constitutional cases in the 1960s and 1970s increased individual rights surrounding marriage and family, and individuals gained broader access to marriage, divorce, birth control, and abortion.[1] Many of these legal changes occurred simultaneously with social upheaval that resulted in large changes in family life. Divorce rates doubled between the mid-1960s and the mid-1970s, and while they have been falling since the late 1970s, the stock of divorced people has continued to grow (Stevenson and Wolfers 2007a, 2008b). This increase in the number of people who have experienced family disruption has increased the dispersion of family experience. Isen and Stevenson (2008) also document differential changes by education in both family behavior and subjective assessments of marital satisfaction.

Previous studies of subjective well-being have found that both pecuniary and nonpecuniary aspects of life contribute to our reported happiness. The fact that happiness data aggregate across these domains makes them especially interesting. Equally, this also expands the range of possible explanations for the changes in the distribution of happiness we document. In this paper, we refrain from attempting any such explanation, and we now turn to assessing how the average levels of happiness and happiness inequality have changed.

3. AGGREGATE TRENDS IN HAPPINESS

Our analysis is based on responses to the General Social Survey (GSS; Davis, Smith, and Marsden 1972–2006), which asks, “Taken all together, how would you say things are these days—would you say that you are very happy, pretty happy, or not too happy?” This survey was administered to a nationally representative sample of about 1,500 respondents each year from 1972 to 1993 (except 1979, 1981, and 1992) and con-

1. See Stevenson and Wolfers (2007a) for further discussion of the legal and social changes impacting families during the 1960s and 1970s.

tinues with around 3,000 respondents every second year from 1994 through to 2004, rising to around 4,500 respondents in 2006 (although only half the respondents were queried about their happiness in 2002 and 2004, followed by two-thirds in 2006). These repeated cross sections are designed to track attitudes and behaviors among the U.S. population and contain a wide range of demographic and attitudinal questions.

Before assessing how answers to these questions have trended over time, it is important to account for any changes in measurement that may affect responses to the happiness question. While these data are relatively consistent, responses to happiness questions are remarkably sensitive to small changes in question order, and hence it is quite important to adjust for changes in survey design. In particular, Smith (1990) notes that reported happiness tends to be higher when preceded by a five-item satisfaction scale (as was the norm except in 1972 and 1985). In addition, among married respondents, reported happiness is higher when preceded by a question about marital happiness (as was the norm, except in 1972). Fortunately, the changes induced by these question order effects can be assessed by way of split-ballot experiments run in subsequent surveys; the Appendix details these adjustments.[2] We show the results of these corrections in Figure 1.

In order to ensure that these time series are nationally representative, all estimates are weighted using WTSALL, and we drop the 1982 and 1987 black oversamples. In order to maintain continuity with earlier survey rounds, we also drop those 2006 interviews that occurred in Spanish and could not have been completed had English been the only option, as Spanish-language surveys were not offered in previous years.[3] Our corrected data series are listed in Table 1.

Having constructed a consistent series, the next challenge is to convert qualitative responses into a meaningful quantitative summary measure. This issue becomes particularly pressing in analyzing the GSS data, as only three response categories are given. The simplest (and most widely used) approach is to equate "not too happy" with a happiness score of 1, "pretty happy" with a score of 2, and "very happy" with a score of

2. While the split-ballot experiments provide a bridge between different versions of the survey, they also mean that it is not possible to simply drop the two outlier years, as results from subsequent surveys also need to be adjusted for the presence of these experimental split ballots.

3. For those interested in replicating our results, the simplest way is to define a weighting variable as follows: gen wt=WTSSALL if SAMPLE~=4 & SAMPLE~=5 & SAMPLE~=7 & SPANINT~=2.

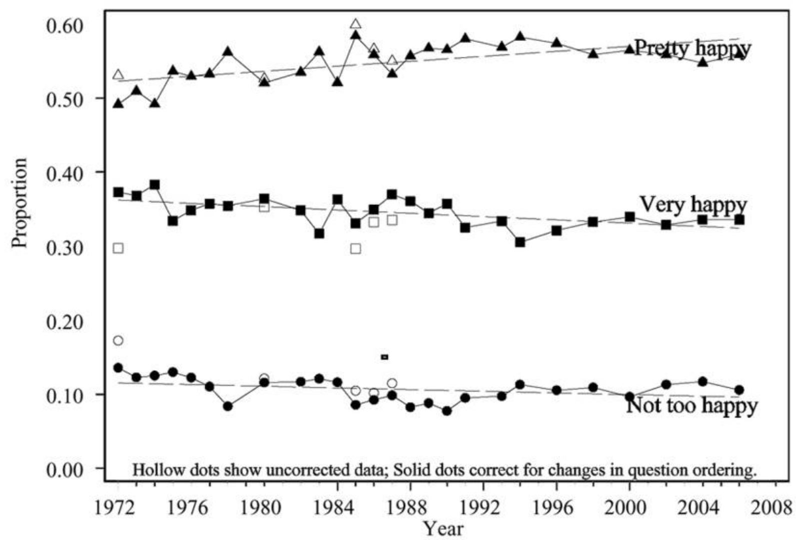

Figure 1. Trends in the distribution of happiness

3. We can then take the mean and variance of these measures each year. The results from this simple approach are presented in Figure 2. Two facts are immediately evident. First, the average level of happiness in the United States is roughly stable, or perhaps slightly declining, a finding that is explored further by Blanchflower and Oswald (2004) and subsequently by Stevenson and Wolfers (2008a). Second, inequality in happiness declined until the mid- or late-1980s, despite the fact that both income and consumption inequality rose through most of this period. Subsequently, happiness inequality rose through the 1990s, although the most recent estimates of inequality still remain below the higher levels seen in the early 1970s.[4] These movements are quite substantial, as we observe an initial decline in the variance of happiness of about 25 percent from the early 1970s to the late-1980s, followed by a rise of about 10 percent by the mid-2000s.

The difficulty with the aggregation shown in Figure 2 is that it arbitrarily assigns qualitative categories scores equal to their rank order, imposing a linear structure in which the difference between being "not

4. The decline in happiness inequality from the 1970s to the early 2000s has also been noted by Brooks (2008).

Table 1. Happiness Trends in the United States

Year	Survey Responses: Not Too Happy (%)	Pretty Happy (%)	Very Happy (%)	Sample Size	Estimated Moments: Happiness ~ $N(0, 1)$: Mean	Variance
1972	13.6	49.1	37.3	1,606	.023	1.321
1973	12.3	50.9	36.8	1,500	.027	1.194
1974	12.5	49.2	38.3	1,480	.060	1.279
1975	13.0	53.6	33.4	1,485	−.055	1.106
1976	12.2	52.9	34.8	1,499	−.015	1.112
1977	11.0	53.2	35.7	1,527	.020	1.059
1978	8.4	56.2	35.5	1,517	.048	.871
1980	11.6	52.0	36.4	1,462	.027	1.123
1982	11.7	53.5	34.8	1,505	−.008	1.072
1983	12.1	56.2	31.7	1,573	−.078	.988
1984	11.6	52.1	36.3	1,445	.025	1.123
1985	8.6	58.4	33.1	1,530	−.001	.821
1986	9.2	55.8	35.0	1,449	.026	.912
1987	9.7	53.3	37.0	1,437	.060	1.016
1988	8.2	55.7	36.1	1,466	.061	.880
1989	8.8	56.7	34.5	1,526	.023	.872
1990	7.7	56.5	35.7	1,361	.061	.838
1991	9.5	58.0	32.5	1,504	−.025	.861
1993	9.7	56.9	33.4	1,601	−.011	.900
1994	11.3	58.2	30.5	2,977	−.089	.904
1996	10.5	57.4	32.1	2,885	−.047	.909
1998	10.9	55.9	33.3	2,806	−.030	.967
2000	9.6	56.4	33.9	2,777	.000	.909
2002	11.3	55.8	32.9	1,369	−.043	.979
2004	11.7	54.7	33.6	1,337	−.034	1.029
2006	10.6	55.9	33.5	2,828	−.019	.955

Source. General Social Survey, 1972–2006.
Note. Estimates are based on sample weight WTSALL, omitting black oversamples in 1982 and 1987, as well as those Spanish-language interviews in 2006 that could not have been completed were English required (as in previous years). All data are corrected for question order effects, as described in the Appendix. The mean and variance estimates are based on a generalized ordered probit regression of happiness in which both the level and variance of happiness are a linear function of year fixed effects, as described in equations (6)–(9). These estimates are constructed on the assumption that the latent happiness variable is normally distributed, with mean zero and average variance equal to one. The estimated cut points are $\delta_1 = -1.244$ and $\delta_2 = -.397$.

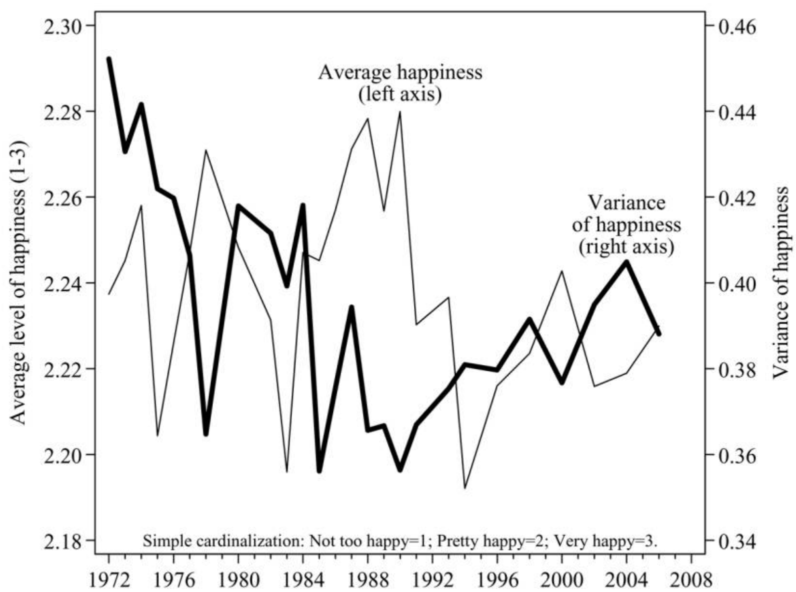

Figure 2. Simple approach to estimating trends in the distribution of happiness

too happy" and "pretty happy" is assumed to be equal to the difference between being "pretty happy" and "very happy." Moreover, it is difficult to know how to interpret comparisons of happiness levels without some sort of normalization. (That is, is by what metric can we interpret the economic significance of the roughly .02 point decline in average happiness levels shown in Figure 2?)

An alternative approach involves using data on the proportions of the population who report themselves as being in each happiness category and imposing a functional form restriction on the distribution of a latent "happiness" index. A common example of this latter approach is the use of ordered probit regressions to estimate trends in well-being, as in Stevenson and Wolfers (2007b, 2008a). In this approach, one assumes that there is a latent variable, happiness, that is normally distributed (and by an innocuous normalization, it is a standard normal). Thus, an ordered probit regression of happiness on year fixed effects recovers the time series of the distribution of happiness. This maximum likelihood procedure simultaneously estimates the cut points above which a person will report being "pretty happy" rather than "not too happy" or "very happy" rather than "pretty happy." In turn, the year fixed effects are interpreted as shifts in the average level of happiness. Unfortunately, the ordered probit (or ordered logit) model is insufficient for our analysis, as we are interested in measuring trends in the dispersion of happiness, whereas these statistical procedures measure shifts in the average level of happiness, while assuming its dispersion is constant.

However, it is fairly straightforward to generalize the ordered probit (or ordered logit) model in order to jointly estimate the time series of both the average level and dispersion of happiness. We can also generalize the specific parameteric assumptions embedded in each particular model. To see this, note that with three response categories, we have essentially two observations each year—the proportion "not too happy" and the proportion "very happy" (the proportion "pretty happy" is the complement and hence perfectly collinear). Thus we can use these two observations to solve for two unknowns and hence recover the parameters of any two-parameter probability density function.[5] Throughout this paper, we will report the mean and variance of happiness implied by these parameters. We do not mean to suggest that the variance is the optimal measure of the dispersion of happiness, but given the restriction to two-parameter distributions, other measures of happiness inequality such as the standard deviation, Gini coefficient, interquartile range, and 90–10 ratio will be a monotonic function of the variance.[6] (For our purposes the variance is particularly convenient, as the decomposition exercise in Section 4 is a relatively straightforward variance decomposition.)

We begin with the usual logic of the standard models for ordered categorical data, assuming that the happiness of an individual, i, from a representative cross section taken in period t, is an unobservable index, y_{it}^*, determined by

$$y_{it}^* = \boldsymbol{x}_{it}\boldsymbol{\beta}_t + \varepsilon_{it}, \tag{1}$$

where $\boldsymbol{x}_{it}$ refers to the individual's observable independent variables and

5. For a related approach, see the appendix to Mankiw, Reis, and Wolfers (2003).

6. Atkinson (1970) describes alternative measures of inequality as applying different weights to various parts of the distribution. In a more applied vein, Kalmijn and Veenhoven (2005) assess different approaches to quantifying inequality of happiness and conclude by endorsing the use of standard deviations.

ε_{it} is the error term. However, we do not observe y_i^* but rather only the ordered categorical variable y_i:

$$y_{it} = \begin{cases} \text{Not too happy} & \text{if } y_{it}^* \leqslant \delta_1 \\ \text{Pretty happy} & \text{if } \delta_1 < y_{it}^* \leqslant \delta_2 \\ \text{Very happy} & \text{if } y_{it}^* > \delta_2, \end{cases} \tag{2}$$

where δ_1 and δ_2 are the unknown cut points that must be estimated.

While the typical ordered probit model further assumes $\varepsilon_{it} \sim N(0, 1)$, at this point we generalize in two directions, allowing for any two-parameter distribution, $F(.)$, and also allowing the variance to vary with observable covariates:

$$\varepsilon_{it} \sim F(0, \boldsymbol{x}_{it}\boldsymbol{\gamma}_t). \tag{3}$$

Thus the independent variables, x_{it}, shift both the mean and variance of happiness. We could allow different sets of independent variables, x_{it}^{m} and x_{it}^{v}, to shift the mean and the variance, and this amounts to denoting the union of these variables as $\boldsymbol{x}_{it} = x_{it}^{\mathrm{m}} \cup x_{it}^{\mathrm{v}}$ and imposing specific zero restrictions on the $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ vectors. We will not impose such restrictions, because we want to allow the data to describe which independent variables drive each moment of the distribution.

Because we are interested in using this approach to simply document aggregate time series variation in the mean and variance of happiness, we begin by focusing on the simple case where the only independent variables are a vector of year fixed effects, yielding the time series of both the average level (μ_t) and variance (σ_t^2) of happiness. In this particularly simple case, this model yields simple closed-form expressions, which can be computed without the need for any specialist software. Without assuming any specific functional form, we note

$$\%\text{Not too happy}_t = F_{\mu_t,\sigma_t^2}(\delta_1) \Rightarrow F^{-1}(\%\text{Not too happy}_t) = \frac{\delta_1 - \mu_t}{\sigma_t} \tag{4}$$

and

$$\%\text{Very happy}_t = 1 - F_{\mu_t,\sigma_t^2}(\delta_2) \Rightarrow F^{-1}(1 - \%\text{Very happy}_t) = \frac{\delta_2 - \mu_t}{\sigma_t}, \tag{5}$$

where $F(.)$ is the cumulative distribution function of a distribution characterized by two parameters that map into an average level of happiness in each year μ_t and variance, σ_t^2. The key to identification of this model is that the cut points, δ_1 and δ_2, do not vary through time. That is, we can identify shifts in both the mean and dispersion in happiness if we are willing to assume that the mapping between true feelings of happiness and how respondents choose to answer the survey remains constant through time. (As an aside, while this assumption sounds strong, it is made implicitly in every approach to cardinalizing subjective well-being that we have seen.) Combining equations (4) and (5) yields

$$\mu_t = \frac{\delta_1 F^{-1}(1 - \%\text{Very happy}_t) - \delta_2 F^{-1}(\%\text{Not too happy}_t)}{F^{-1}(1 - \%\text{Very happy}_t) - F^{-1}(\%\text{Not too happy}_t)} \tag{6}$$

and

$$\sigma_t^2 = \left(\frac{\delta_2 - \delta_1}{F^{-1}(1 - \%\text{Very happy}_t) - F^{-1}(\%\text{Not too happy}_t)}\right)^2. \tag{7}$$

The cut points δ_1 and δ_2 define the location and scale over which we are measuring μ_t and σ_t^2, and so we normalize so as to ensure that the average level of happiness across the entire sample is zero and the average variance is one. This normalization implies

$$\delta_1 = \sqrt{\phi}\sum_{t=1}^{\tau} F^{-1}(\%\text{Not too happy}_t) \tag{8}$$

and

$$\delta_2 = \sqrt{\phi}\sum_{t=1}^{\tau} F^{-1}(1 - \%\text{Very happy}), \tag{9}$$

where the constant

$$\phi = \frac{\tau}{\{\sum_{t=1}^{\tau}[F^{-1}(1 - \%\text{VH}) - F^{-1}(\%\text{NTH})]\}^2\{\sum_{t=1}^{\tau}[F^{-1}(1 - \%\text{VH}) - F^{-1}(\%\text{NTH})]^{-2}\}}$$

simplifies the above expressions, the abbreviations %NTH and %VH correspond to the proportions "not too happy" and "very happy," respectively, and τ denotes the number of periods for which we are estimating happiness trends. Our normalization ensures that we recover estimates of μ_t and σ_t that are roughly comparable across assumptions about functional forms and comparable to those from an ordered probit regression (which imposes that this normalization holds for all observations, rather than just hold on average).

Thus, for any specific assumption about the functional form of the underlying latent happiness variable, the simple expressions in equations (6)–(9) can be evaluated using a simple spreadsheet program to compute the time series of both the average level of happiness and its variance.[7] Table 1 provides an example of these calculations for the normal distribution.

We assess the robustness of our estimates of the time series of the distribution of happiness in Figure 3 on the basis of three increasingly fat-tailed assumptions about the distribution of the latent happiness variable: normality, a logistic distribution, and a uniform distribution. The mean and variance estimates are based on equations (6)–(9). While any set of assumptions will seem arbitrary, it is worth noting that the more widely used approach presented in Figure 2 is based on the particularly unappealing assumption that the happiness distribution has three equally spaced mass points corresponding to the three allowable responses.

The key finding is that none of the qualitative (and, indeed, quantitative) implications of our earlier analysis are much changed by alternative approaches to cardinalizing the happiness question, which is quite reassuring. Again, we find only a mild negative trend in average happiness but a clear decline in happiness inequality, with a turning point registered in about the late 1980s and only a gradual increase in the subsequent years. By any measure, happiness inequality in the first third of our sample period is higher than in the final third.

Figure 1 provides some simple intuition for why alternative distributional assumptions yield such similar results: the decline in inequality through to the late 1980s is roughly equally evident whether looking at those who are unusually happy or when looking at those who are unusually unhappy, and as such, placing different weights on the proportion "very happy" relative to the proportion "not too happy" yields similar trends.

Having found that our simple generalization of the ordered probit, ordered logit, and ordered uniform models yields such similar time-series estimates of both the average levels of happiness and its dispersion, the rest of our analysis will focus on the generalized ordered probit model (which assumes normality), although none of our results are materially affected by this focus. An alternative rationale for focusing on the normal distribution is that alternative subjective well-being questions that elicit responses on a 10-point scale tend to yield roughly normally distributed responses (although as Oswald [2008] notes, reported happiness reflects

7. This computational simplicity is a useful side effect of the model being just identified. If one were assessing four or more categorical responses, or were to add control variables, then an explicit maximization routine would be required. We explore this further in Section 4.

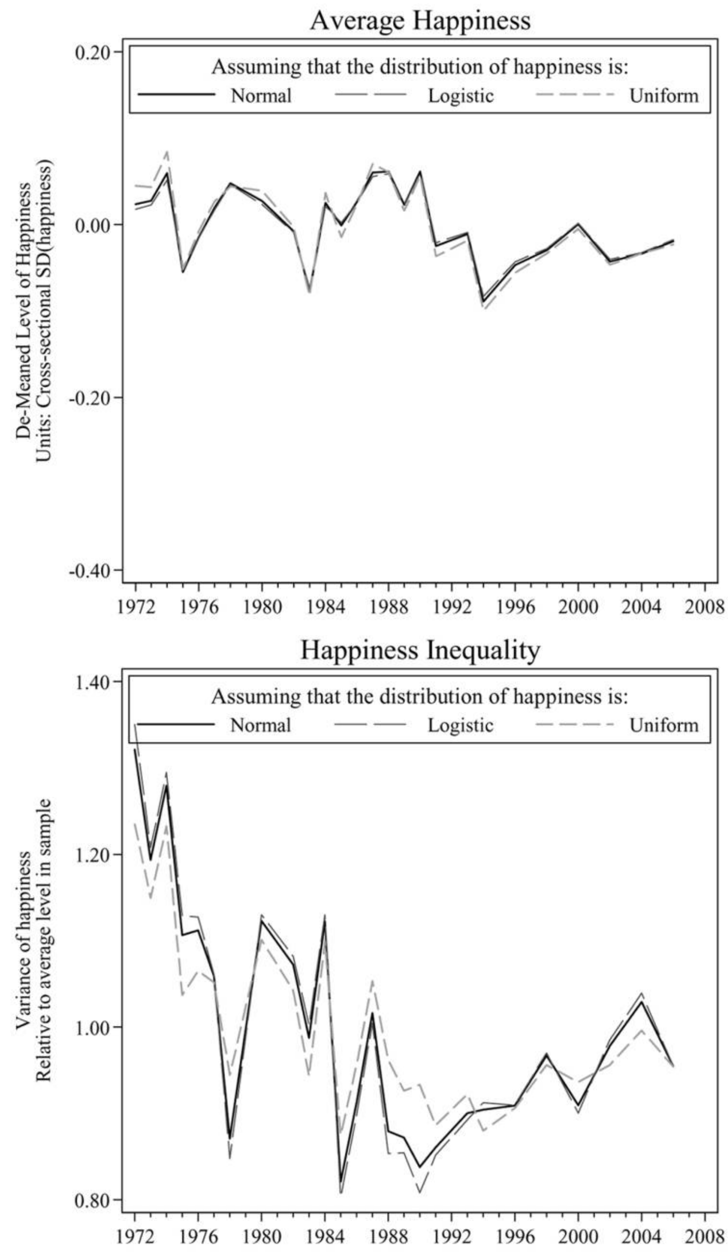

Figure 3. Estimated trends in the distribution of happiness (General Social Survey, 1972–2006).

the conjunction of true happiness and an unknown happiness reporting function).

In order to better assess these changes, we put both mean and dispersion shifts together in Figure 4, showing the combined impact on the distribution of happiness. The estimated distribution of a latent happiness variable, recovered by running a generalized ordered probit regression in which the level and variance of happiness are a linear function of decadal fixed effects, takes account of estimated shocks to both the mean and variance of happiness. (See Table 2 for coefficient estimates underlying the figure.) In order to keep the charts uncluttered, we base these plots on decadal averages of the mean and variance of happiness: the decadal average variance of happiness falls from 1.135 in the 1970s to .979 in the 1980s and to .897 in the 1990s, before rising to .968 in the 2000s; the corresponding numbers for average levels of happiness show much less movement: .015 in the 1970s, .015 in the 1980s, −.023 in the 1990s, and −.024 in the 2000s.[8] This plot shows quite clearly that the magnitude of the changes in dispersion dominates any change in the average level of happiness. For instance, from the 1970s to the 2000s, the happiness level at the 25th (75th) percentile of the happiness distribution rose by .016 points (fell by .094 points), reflecting a .039 point decline in the mean and a .055 point rise (fall) due to increasing happiness inequality. Indeed, Figure 4 shows that while the decline in average levels of happiness made the population in the 2000s less happy on average than in the 1970s, the crossover of the two cumulative distribution functions implies that 32 percent of the population are happier today, and this is due to the offsetting effects of the decline in happiness inequality shrinking the left tail.

Indeed, even as there was not much movement in average levels of happiness in the 1970s and 1980s, there were large increases in happiness at the bottom of the happiness distribution. Figure 5 illustrates, showing annual estimates of the change in well-being since 1972, at various percentiles of the happiness distribution. Percentiles were estimated by running a generalized ordered probit regression in which both the mean and variance of happiness are a linear function of year fixed effects, and we make projections based on the assumed normality of the distribution of happiness. (See Table 1 for coefficient estimates underlying the figure.)

8. The decades we refer to as the 1970s should be understood as the period since the General Social Survey began in 1972; similarly, estimates for the 2000s reflect data from 2000–2006 (the most recent survey), and the 1980s and 1990s refer only to those years in which the General Social Survey was conducted.

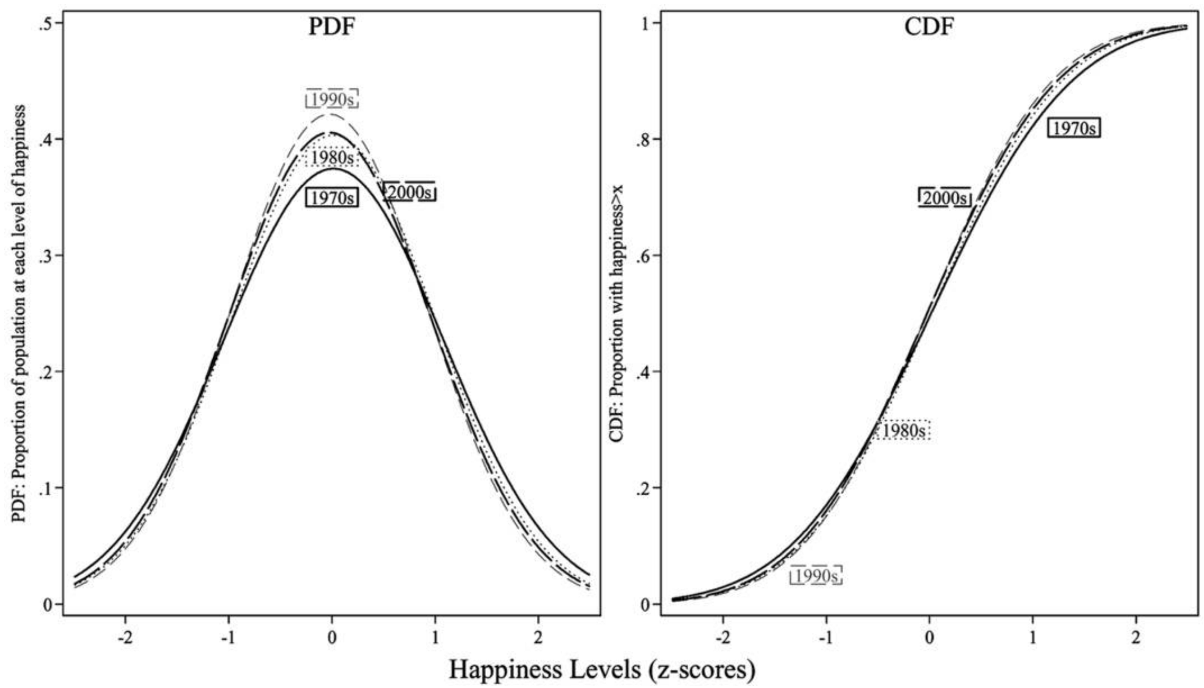

Figure 4. Distribution of happiness in the United States, by decade (General Social Survey, 1972–2006)

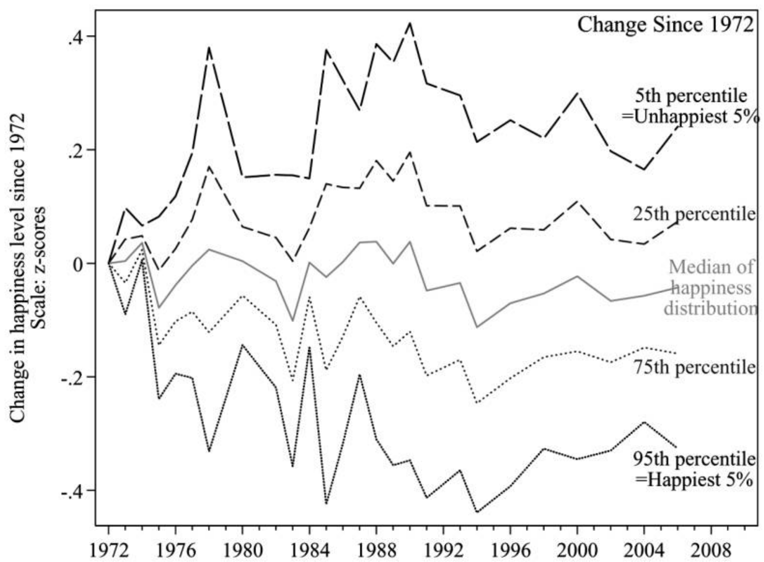

Figure 5. Changes in happiness levels since 1972, by percentile

The (unhappiest) 5th percentile have become .24 points happier since 1972; the 25th percentile have gained around .07 points; the median lost .04 points; the 25th percentile lost .16 points, and the (happiest) 5th percentile lost .33 points. Again, each of these numbers should be considered relative to a cross-sectional standard deviation (normalized to one). In order to compare these magnitudes with their dollar equivalents, note that Stevenson and Wolfers (2008a) estimated that each 10 percent increase in log family income is associated with an increase in happiness of .022 points.[9] Thus, each of these changes in happiness can be converted to a happiness-equivalent percentage change in income, by dividing by .0022 (or multiplying by 45), which suggests that the decline in happiness inequality that we document is very large.

9. The regression reported in figure 8 of Stevenson and Wolfers (2008a) is an ordered probit regression based on these same 1972–2006 GSS data: Happiness = .22ln(Real family income) (standard error = .007), where family income is deflated by the Consumer Price Index Research Series Using Current Methods (CPI-U-RS), and each year's income intervals are converted into point estimates by interval regressions, assuming that income is log-normally distributed. It should be noted that estimates of the happiness-income gradient based on other data sets were often somewhat larger, but most tend to lie in the range of .2–.4.

In the first column of Table 2, we assess these aggregate trends more formally, analyzing the annual time series of happiness inequality (and its mean) derived from our generalization of the ordered probit. (See the Estimated Moments columns of Table 1 for the underlying data.) When we examine decadal averages in happiness inequality, we see that the variance fell by a total of 15 percent from the 1970s to the 2000s, which reflects a decline of 21 percent from the 1970s through the 1990s, about one-third of which reversed in the subsequent decade.

We also assess the magnitude and statistical significance of the overall trend in our annual estimates of inequality. In addition, we allow for a change in the trend; in order to minimize data snooping (and to maximize statistical power), we simply break the sample in half, testing for a break at the (chronological) midpoint of our sample, 1989.[10] In each case, we report Newey-West standard errors, accounting for first-order autocorrelation in happiness inequality. These regressions confirm that the decline in happiness inequality is both economically and statistically significant and that the average decline through the sample period reflects a sharp decline in happiness inequality through the first part of the sample and a subsequent, smaller, rise in the second part, undoing about one-third of the initial decline. We report similar regressions estimates in the bottom half of Table 2, albeit analyzing average levels of happiness rather than its dispersion. These regressions reveal that there is a small, statistically significant overall decline in average happiness. The average level of happiness in the United States is lower in the 1990s and 2000s than it was in the 1970s and 1980s.

Thus far we have shown that in the aggregate, happiness inequality fell sharply during the 1970s and continued to fall in the 1980s, before rising slightly in the 1990s and 2000s. In contrast, average happiness in the population shows little evidence of a trend before the late 1980s, at which point it falls. Yet these broad trends may mask underlying heterogeneity in both the average happiness and the inequality of happiness across socioeconomic and demographic groups. We now turn to digging a bit further into which groups are most affected by these trends.

4. ASSESSING TRENDS WITHIN AND BETWEEN GROUPS

Our approach so far focuses on population aggregates, estimating average happiness and inequality within each year by treating annual ob-

10. Given the limited degrees of freedom, we do not test for a discontinuous break in levels in the series, allowing only a change in the trend.

Table 2. Trends in Happiness Inequality and Levels, by Group

	Block 1: Full Sample	Block 2: By Education				Block 3: By Gender		Block 4: By Race		Block 5: By Marital Status		Block 6: By Age		
		College Grad	Some College	High School	<High School	Men	Women	White	Nonwhite	Married	Not married	18–34	35–49	50+
Variance of happiness:[a]														
Decadal averages of Var(happiness):														
1970s	1.135	.921	1.053	1.126	1.316	1.073	1.200	1.049	1.223	1.186	1.045	1.001	1.113	1.270
1980s	.979	.836	.840	.917	1.316	.958	.997	.912	1.055	.990	1.015	.850	.862	1.249
1990s	.897	.884	.827	.849	1.132	.859	.934	.836	.948	.913	.873	.779	.903	1.008
2000s	.968	.956	.912	.955	1.113	.958	.975	.899	1.021	.935	.970	.945	.916	1.021
Full-sample time trend[b]	−.797*	−.036	−.616	−.786+	−.989*	−.644+	−.966**	−.710*	−.991*	−1.038**	−.567*	−.420	−.837	−1.176**
	(.297)	(.390)	(.452)	(.410)	(.358)	(.313)	(.300)	(.265)	(.379)	(.260)	(.257)	(.384)	(.499)	(.195)
Testing for trend break:														
Trend 1972–1989	−2.101**	−1.166	−2.061**	−2.327**	−1.580*	−1.843**	−2.382**	−1.840**	−2.510**	−2.208**	−1.400**	−1.866**	−2.483**	−1.787**
	(.276)	(.724)	(.719)	(.486)	(.587)	(.339)	(.289)	(.244)	(.564)	(.279)	(.368)	(.363)	(.842)	(.492)
Trend 1989–2006	.748**	1.302**	1.097*	1.078*	−.289	.776*	.713**	.628**	.808+	.349	.420	1.295*	1.113*	−.452
	(.202)	(.312)	(.405)	(.399)	(.866)	(.350)	(.232)	(.185)	(.425)	(.264)	(.366)	(.601)	(.462)	(.496)

Variance of average happiness:[c]														
Decadal averages of average happiness:														
1970s	.015	.104	.052	.046	−.092	−.035	.068	.190	−.287	.254	−.340	−.083	.053	.095
1980s	.015	.171	.023	−.001	−.110	−.011	.038	.181	−.181	.248	−.195	−.059	.006	.115
1990s	−.023	.136	−.030	−.065	−.210	−.022	−.021	.141	−.115	.255	−.249	−.061	−.058	.055
2000s	−.024	.156	−.003	−.124	−.281	−.023	−.025	.149	−.105	.294	−.261	−.069	−.024	.016
Full-sample time trend[d]	−.163*	.055	−.244*	−.570**	−.690**	.056	−.378**	−.187*	.759**	.124	.278	.063	−.350**	−.290*
	(.061)	(.129)	(.109)	(.091)	(.134)	(.083)	(.101)	(.070)	(.193)	(.096)	(.237)	(.100)	(.106)	(.137)
Testing for trend break:														
Trend 1972–89	−.009	.161	−.262	−.279	−.187	.399*	−.420*	−.124	1.386**	−.025	1.352**	.401	−.499*	−.004
	(.164)	(.292)	(.238)	(.281)	(.256)	(.170)	(.193)	(.186)	(.368)	(.230)	(.261)	(.273)	(.232)	(.269)
Trend 1989–2006	−.345*	−.071	−.224	−.915	−1.285**	−.352	−.328	−.263	.016	.301	−.994**	−.338	−.173	−.629**
	(.151)	(.217)	(.155)	(.261)	(.313)	(.192)	(.218)	(.159)	(.281)	(.233)	(.186)	(.209)	(.277)	(.178)

Note. Newey-West standard errors are in parentheses, correcting for first-order autocorrelation. $N = 26$ annual or biennial observations on the means and variance of happiness, estimated using generalized ordered probit from the General Social Survey, 1972–2006. Trend 1972–89 reports β_1. Trend 1989–2006 reports $\beta_1 + \beta_2$. Time trends have been divided by 100 (or, alternatively, the coefficients multiplied by 100).

[a] Relative to sample average.

[b] Variance happiness_t(or mean happiness_t) $= \alpha + \beta(\text{Year}_t - 1972)/100$.

[c] Relative to cross-sectional standard deviation of happiness.

[d] Variance happiness_t(or mean happiness_t) $= \alpha + \beta_1 I(1972 \leqslant \text{Year}_t \leqslant 1989)(\text{Year}_t - 1972)/100 + \beta_2 I(\text{Year}_t \geqslant 1989)(\text{Year}_t - 1989)/100$.

[+] Statistically significant at the 10% level.

* Statistically significant at the 5% level.

** Statistically significant at the 1% level.

servations as distinct cells. We can extend our analysis to consider changes over time within and between categories of demographic groups by estimating separate regressions that consider demographic category × year as the relevant cell. Formally, this simply requires replacing the subscript t in equations (4)−(9) (which denoted separate years as the relevant cells) with the subscript t, d (thus denoting distinct demographic categories in distinct years as the relevant cells). This procedure yields unconditional estimates of differences in happiness levels and inequality both between demographic categories at each point in time and within each demographic category through time. Once we estimate these time series, we report a few summary characteristics in Table 2. For example, by interacting education categories with the year fixed effects, our regression estimates the average level and variance of happiness in each education category through time (and indeed, block 2 of Table 2 describes precisely the evolution of happiness by education and time).

Recall that the first column of Table 2 shows the aggregate trends in happiness and happiness inequality and thus summarizes 26 annual observations. Each subsequent block in Table 2 reports separate regressions that analyze separate demographic category × year cells. Note that levels of happiness and the measure of happiness inequality shown in Table 2 are standardized within each block and hence are comparable across columns within a block, but not comparable between blocks. (Estimates within each block are based on the same cut points, but between blocks these are reestimated and so differ slightly.) We should also be clear in noting that these are raw trends, and so they do not simultaneously account for other factors influencing trends in happiness and happiness inequality.

As such, when interpreting these descriptive analyses—as with all of our demographic breakdowns—it is important to bear in mind that the dramatic changes in the proportions of the population choosing higher education levels or choosing to remain unmarried. If, as seems likely, the marginal member added to (or subtracted from) each group is different from the average, this changing composition will account for some of the time-series variation in the estimated levels and dispersion of happiness within each group. This caveat should also be borne in mind even in Section 4, as it continues to be relevant even despite our best efforts to account for observable differences and compositional change. Nonetheless, this approach does allow us to make some useful within-group time series comparisons, and a few interesting trends emerge.

Focusing on the top half of Table 2, which assesses trends in happiness inequality, the most striking finding is simply that the broad trends seen in the aggregate appear similarly within each of the different demographic groups. Happiness inequality within most groups was highest in the 1970s, fell in the subsequent 2 decades, and rose slightly in the 2000s, although it remains below earlier levels. In contrast, when looking at the average level of happiness in the bottom half of the table, stark differences occur across groups.

To examine patterns in each of the groups more closely, we begin by focusing on the patterns by educational attainment. Recall that returns to education were falling in the 1970s and rose sharply in the 1980s, 1990s, and 2000s. Real wages of men with less than a high school degree and those with only a high school degree stagnated or fell through much of the period, while the real wages of those with a college degree or beyond rose. In contrast, leisure increased among those with less education relative to those with more education. Turning to happiness inequality, block 2 shows that, in the 1970s, average happiness inequality fell with educational attainment. By the 1980s, the differences in happiness inequality between the groups had declined, such that the gap in the dispersion of happiness between those with a college degree and those who had attended, but not completed, college had disappeared. In addition to a decrease in between-group inequality over this period, happiness inequality was much lower for all groups, except among those who had not completed high school. In the 1990s, happiness inequality was little changed among those with some college education, rose for those with a college degree, and fell among those with a high school degree or less. In the most recent period, happiness inequality continued to decline among high school dropouts, despite being higher than in the preceding decade for other groups. By the 2000s, not only was the dispersion of happiness lower within groups compared with the 1970s, but the differences in happiness inequality between groups had been reduced.

In contrast, the bottom half of Table 2 shows that trends in average levels of happiness has varied quite strongly across education groups, with happiness rising among college graduates, falling among those with some college, and falling sharply among those with a high school degree or less. These patterns are what one might expect based on between-group changes in wage inequality (although it is at odds with rising leisure among the less educated).

Turning to examine happiness patterns by gender, we see in the bottom half of the table that women's happiness has fallen, while male happiness followed a statistically insignificant upward trend. This pattern in similar to that seen in Stevenson and Wolfers (2007b), who demonstrate that women's happiness has fallen both absolutely and relative to that of men since the 1970s. Indeed, the gender happiness gap in the 1970s was not only eroded over the subsequent decades, but today, women typically report lower levels of happiness than men. Not surprisingly, happiness inequality for women was also higher than that for men in the 1970s (block 3 in the top half of the table). Yet inequality among both men and women fell in roughly equal measure over the next 2 decades and then rose in the most recent period. These trends have yielded decreased inequality of happiness among both men and women and reduced the difference in the dispersion of happiness between the two groups.

The racial gap in average happiness has also declined since the 1970s, however nonwhites remain substantially less happy—on average—than whites. We find a strikingly large, statistically significant increase in average happiness among nonwhites, while happiness among whites has been declining slightly (in block 4).[11] Examining happiness inequality in the top half of the table shows not only that average levels of happiness were much lower among nonwhites but also that the dispersion of happiness was greater. Happiness inequality fell for both groups through to the 1990s and rose in the 2000s, with inequality lower both within and between racial categories in 2000 than in the beginning of the sample.

Finally, we examine differences in happiness inequality by marital status and by age. One reason for examining differences across marital status is because of the well-known finding that marriage is associated with higher levels of subjective well-being (Blanchflower and Oswald 2004). This pattern is evident in decadal averages of average happiness—in all decades those who are married are happier than those who are not—and there is little trend in their levels of happiness. Moreover, patterns of inequality of happiness are similar for the two groups and match the trends seen for the population as a whole.

Block 6 in Table 2 examines patterns by age, and here some interesting trends emerge. The dispersion of happiness increases with age—a fact

11. The GSS race variable allows for a division into white, black and other. Unfortunately there are so few respondents in the "other" category that separating nonwhites into its constituent groups yields particularly imprecise, and thus uninformative, results for these categories. Even so, in further regressions (not shown) breaking out these categories, our estimates for blacks largely track those for obtained for the broader "nonwhite" category.

that is reminiscent of other trends in inequality by age, such as the fact that dispersion in both income and consumption increases with age (Deaton and Paxson 1994). However, the rise in happiness inequality over the life cycle has diminished over the past 35 years and there is less fanning out in the most recent period. The aggregate pattern—of lower levels of happiness inequality in the 1980s and 1990s, rising in the 2000s—is seen for the youngest and oldest age groups, although among prime-age adults happiness inequality is higher in the 1990s than it is in the 1980s. The lower half of the table shows that happiness rises with age, yet the time trend in average happiness has been flat or slightly rising among the young (ages 18–34) and declining among both prime-age and older adults.[12]

The key commonality across all of these results is that happiness inequality has declined within all of these demographic groups. Naturally, these trends may be interrelated, and so we conducted a further analysis based on more narrowly defined demographic groups, breaking the sample up into 24 subsamples reflecting a division into mutually exclusive and collectively exhaustive samples for two genders × three age groups × four education levels. Within 20 of these 24 cases we found a trend decline in happiness inequality, and in no case did with find statistically significant evidence of a trend increase in inequality. We interpret these findings as suggesting a pervasive rise in within-group happiness inequality.

Breaking up the sample into distinct subsamples can go only so far with our limited sample sizes (and this exercise already yielded some fairly small cell sizes). As such, a more formal regression framework is needed if we are also going to account for the influence of further factors. We now turn to developing an appropriate estimation framework in greater detail that will condition on a variety of demographic and socioeconomic variables at once.

A More General Approach

The key to our estimation is simply to generalize the standard ordered probit model so as to allow us to jointly estimate both the mean and variance of happiness as a function of a rich set of covariates.

Given the model defined by equations (1)–(3), when analyzing a data

12. Blanchflower and Oswald (2004) report a U shape in life-cycle happiness in which happiness is highest at young and older ages. However, this pattern is what occurs when examining happiness patterns by age conditional on life outcomes such as marriage, income, and employment status. Our results in Table 2 are unconditional.

set with N observations indexed by i, a dependent variable consisting of J ordered response categories, and a covariate vector, $\boldsymbol{x}_{it}$, the log likelihood function is

$$\ln\ L = \sum_{i=1}^{N}\sum_{j=1}^{J} I(y_{it} = j)\ln\left[F\left(\frac{\delta_j - \boldsymbol{x}_{it}\boldsymbol{\beta}_t}{\sqrt{\boldsymbol{x}_{it}\boldsymbol{\gamma}_t}}\right) - F\left(\frac{\delta_{j-1} - \boldsymbol{x}_{it}\boldsymbol{\beta}_t}{\sqrt{\boldsymbol{x}_{it}\boldsymbol{\gamma}_t}}\right)\right], \tag{10}$$

where $F(.)$ is the cumulative distribution function of the error term, and we shall assume that it is normal; we impose no bounds on the latent happiness variable, and hence $\delta_1 = -\infty$ and $\delta_j = \infty$. Two further constraints are required to pin down the location and scale of the estimates, and as before, we impose these constraints so as to ensure that the latent happiness index has a mean of zero and an average variance of one: $\sum_{i=1}^{N}\boldsymbol{x}_{it}\boldsymbol{\beta}_t = 0$ and $\sum_{i=1}^{N}\boldsymbol{x}_{it}\boldsymbol{\gamma}_t = N$. Our interest lies in the $\boldsymbol{\beta}_t$ vector, which shifts the average level of happiness, and $\boldsymbol{\gamma}_t$, which shifts its variance.

Thus, our results in Section 3 can be reframed as solving the maximization problem described by equation (10), where the covariates, $\boldsymbol{x}_{it}$, were simply a vector of year fixed effects. The advantage of our generalized framework is that we can now estimate different trends by demographic group, conditioning on time-series movements in the level and dispersion of happiness common to other demographic characteristics. By comparison, the approach described in equations (4)–(9) required dividing the sample into mutually exclusive cells—something that is feasible only when assessing a small number of covariates (particularly given the relatively small samples in the GSS). We now turn to expanding the vector of relevant covariates, $\boldsymbol{x}_{it}$, so that it incorporates not only a vector of year fixed effects but also those year fixed effects interacted with dummy variables for each education, gender, race, marital status, and age group. That is, we estimate

$$\begin{aligned}\text{Happiness}^*_{it} = {} & \sum_t I(\text{year}_t = t) \\ & \times\Bigg\{\Big[\sum_e \boldsymbol{\beta}_{e,t} I(\text{educ}_{it} = e) + \sum_s \boldsymbol{\beta}_{s,t} I(\text{sex}_{it} = s) + \sum_r \boldsymbol{\beta}_{r,t} I(\text{race}_{it} = r) \\ & + \sum_m \boldsymbol{\beta}_{m,t} I(\text{mar}_{it} = m) + \sum_a \boldsymbol{\beta}_{a,t} I(\text{age}_{it} = a) + \sum_r \boldsymbol{\beta}_{r,t} I(\text{region}_{it} = r)\Big] \\ & + \Big[\sum_e \boldsymbol{\gamma}_{e,t} I(\text{educ}_{it} = e) + \sum_s \boldsymbol{\gamma}_{s,t} I(\text{sex}_{it} = s) + \sum_r \boldsymbol{\gamma}_{r,t} I(\text{race}_{it} = r) \\ & + \sum_m \boldsymbol{\gamma}_{m,t} I(\text{mar}_{it} = m) + \sum_a \boldsymbol{\gamma}_{a,t} I(\text{age}_{it} = a) + \sum_r \boldsymbol{\gamma}_{r,t} I(\text{region}_{it} = r)\Big]^{1/2} \times \varepsilon_{i,t}\Bigg\},\end{aligned} \tag{11}$$

where Happiness^*_{it} is the unobserved happiness index, reported according to equation (2), ε_{it} is the error term, the $\boldsymbol{\beta}$ terms shift the level of happiness differentially for each group in each year, and the γ terms shift the variance of happiness differentially for each group in each year. Our omission of income from this equation is purposeful, as we wish to juxtapose our findings regarding trends in happiness inequality by demographic group, with analogous trends reported in the income inequality literature.

Estimating this full regression yields 26 separate year fixed effects for both the level and dispersion of happiness for each of 17 different demographic groups (four education groups, two genders, two racial groups, two marital statuses, three age groups, and four regions), for a total of $26 \times 2 \times 17 = 884$ coefficient estimates. Thus, instead of showing a regression table, we present these point estimates graphically in Figure 6 (focusing on average happiness levels for each group) and Figure 7 (focusing on happiness inequality within each group). Instead of showing coefficients relative to an arbitrary omitted group, each panel shows the predicted levels and dispersion of happiness of someone with the average sample characteristics, except for the particular characteristic examined in each panel. Thus, for instance, the top (grey) line in the first panel of Figure 6 shows the evolution of happiness for someone with college education but all other (noneducation) covariates set to their (time-invariant) sample averages (and Figure 7 shows the corresponding variance).

These figures illustrate in more detail the broad trends seen in Table 3. Figure 6 shows that happiness has fanned out by education, with happiness highest (and rising) for college graduates, but lower and falling for high school graduates, and declining more steeply for high school dropouts. Happiness has also become more dispersed by age, albeit only slightly. Among 35–49 year olds and those over age 50, happiness has trended downward, while the happiness of 18–34 year olds has trended upward. As previously seen, happiness has converged along gender and racial lines. Indeed, the closing of the racial happiness gap is striking and appears to have nearly been eliminated in recent years, which suggests that the much larger unconditional racial happiness gap seen in Table 2 may be attributable to the combined impact of racial differences in educational attainment and widening educational differences in happiness.[13] Happiness differences by marital status narrowed in the 1980s,

13. Again we emphasize that differences by demographic group are merely descriptive means (or conditional means), and it should not be inferred that these are causal relationships.

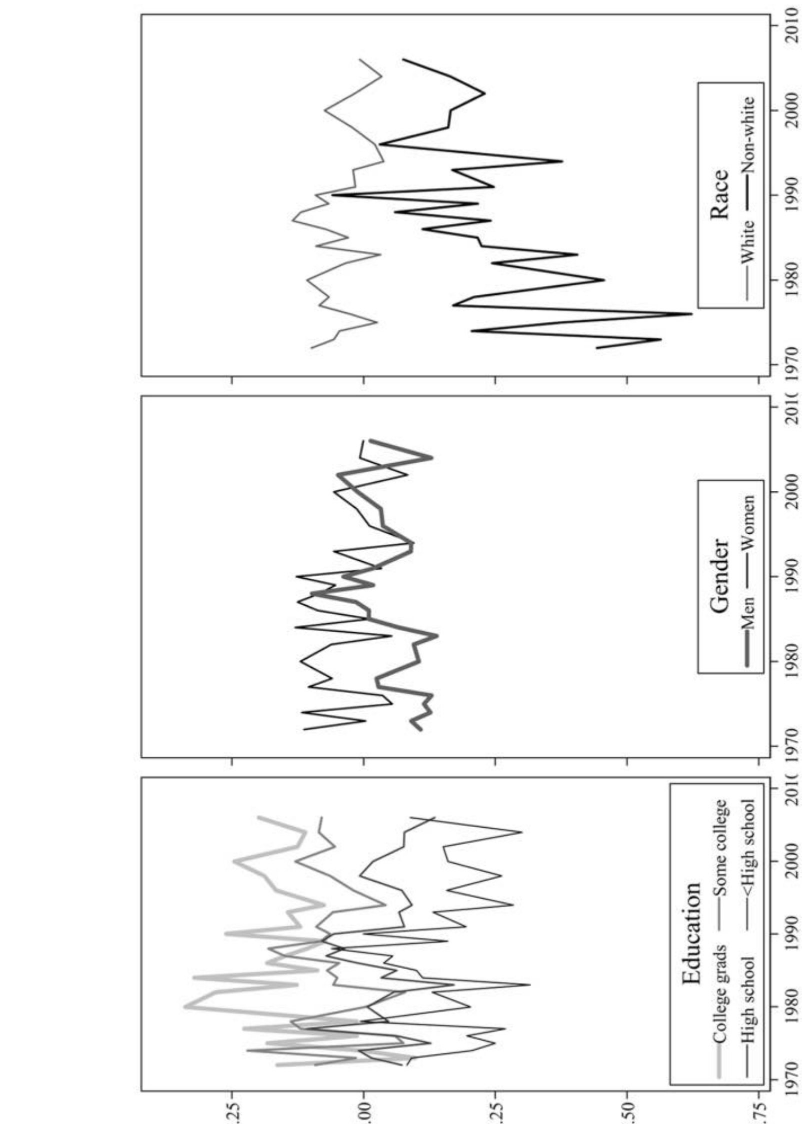
Average (Conditional) Happiness Levels
0.25
0.00
-0.25
-0.50
-0.75
1970
1980
1990
2000
2010
Education
College grads
Some college
High school
<High school
Gender
Men
Women
Race
White
Non-white

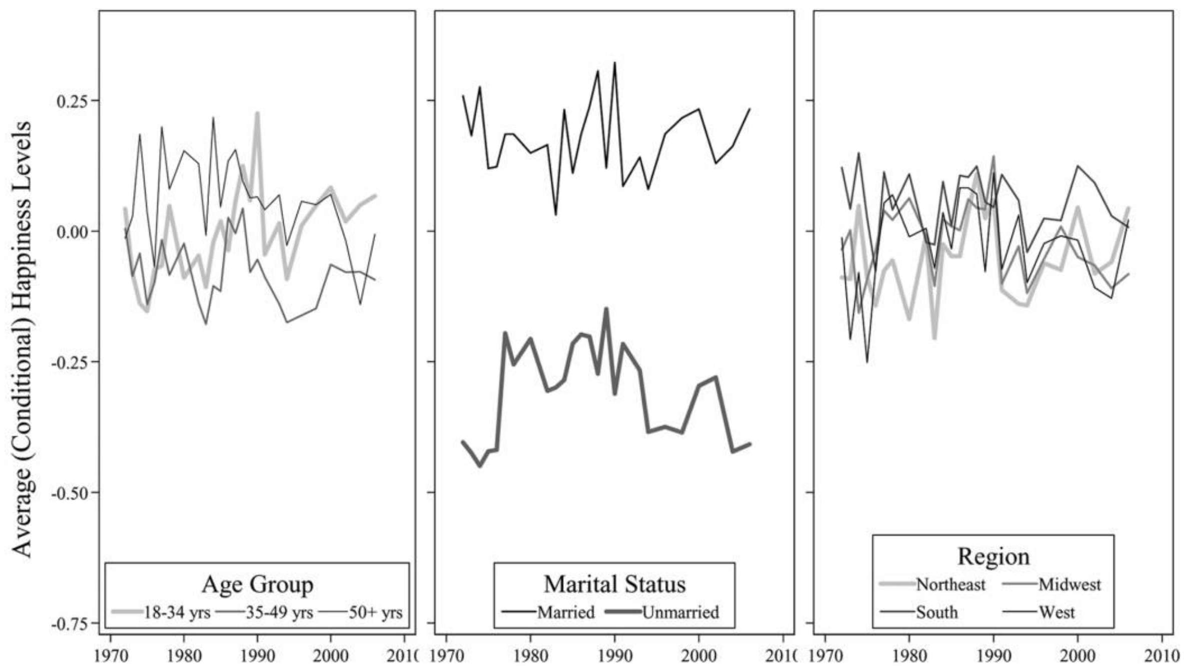

Figure 6. Trends in happiness levels

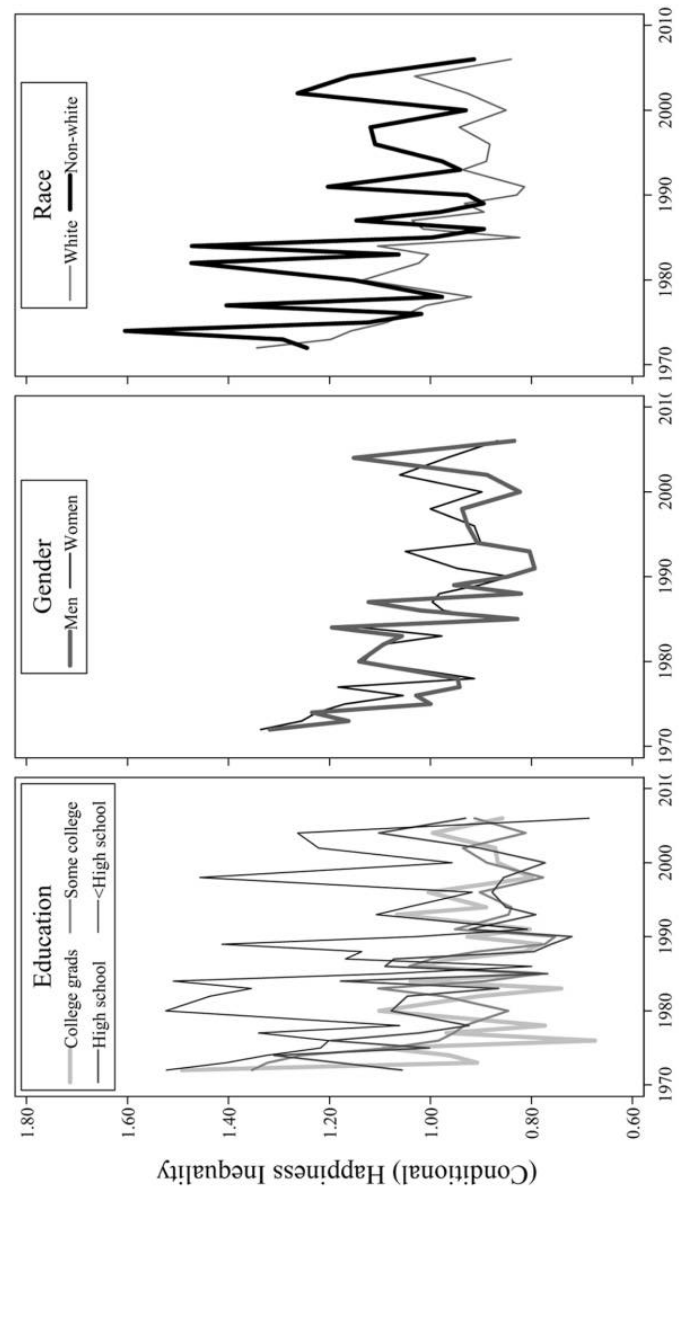

(Conditional) Happiness Inequality
Education
College grads
Some college
High school
<High school
Gender
Men
Women
Race
White
Non-white
1.80
1.60
1.40
1.20
1.00
0.80
0.60
1970
1980
1990
2000
2010

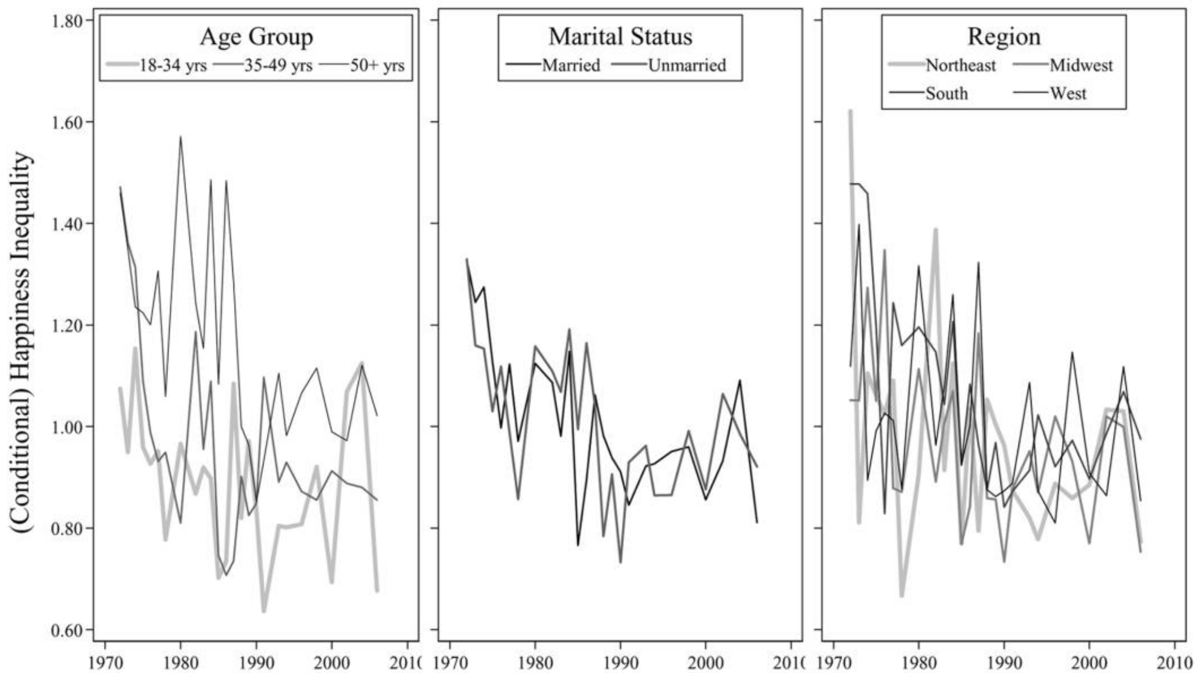

Figure 7. Trends in happiness inequality, within each group

Table 3. Regression Describing the Evolution of Happiness, by Decade

	Block 1: Level of Happiness					Block 2: Variance of Happiness					Block 3: Sample Proportions (%)				
	$\bar{\beta}$	β_{1970s}	β_{1980s}	β_{1990s}	β_{2000s}	$\bar{\gamma}$	γ_{1970s}	γ_{1980s}	γ_{1990s}	γ_{2000s}	$\bar{x}$	$\overline{x_{1970s}}$	$\overline{x_{1980s}}$	$\overline{x_{1990s}}$	$\overline{x_{2000s}}$
Constant	.016	.062	−.018	−.032	.084	.830	.931	.832	.730	.859	100	100	100	100	100
		(.011)	(.010)	(.010)	(.014)		(.011)	(.010)	(.010)	(.014)					
College graduate	.198	.102	.194	.200	.262	−.008	−.090	−.079	.071	.005	20.5	13.7	17.6	24.6	27.2
		(.011)	(.009)	(.008)	(.011)		(.011)	(.009)	(.008)	(.011)					
Some college	.102	.095	.090	.075	.158	−.005	.035	−.035	−.005	.002	23.2	16.9	21.1	26.1	30.1
		(.010)	(.009)	(.008)	(.010)		(.010)	(.009)	(.008)	(.010)					
High school graduate	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	32.0	33.9	34.3	30.9	27.7
<High school	−.117	−.132	−.098	−.142	−.077	.251	.236	.334	.239	.079	24.3	35.4	27.0	18.3	15.0
		(.008)	(.008)	(.009)	(.012)		(.008)	(.008)	(.009)	(.012)					
Men	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	45.9	47.5	44.8	45.4	46.2
Women	.076	.139	.096	.038	.022	.036	.086	−.016	.054	.032	54.1	52.5	55.2	54.6	53.8
		(.007)	(.006)	(.006)	(.008)		(.007)	(.006)	(.006)	(.008)					
White	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	83.9	87.8	86.4	82.0	78.0
Nonwhite	−.242	−.408	−.385	−.174	−.169	.149	.190	.123	.173	.115	16.1	12.2	13.6	18.0	22.0
		(.010)	(.009)	(.008)	(.010)		(.010)	(.009)	(.008)	(.010)					
Married	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	62.7	72.1	63.9	58.9	55.0
Nonmarried	−.494	−.555	−.395	−.507	−.554	−.025	−.141	.033	−.053	.032	37.3	27.9	36.1	41.1	45.0
		(.007)	(.006)	(.006)	(.008)		(.007)	(.006)	(.006)	(.008)					
Ages 18–34	.088	.014	.058	.159	.137	−.047	−.075	−.007	−.059	−.058	35.1	37.0	37.6	33.4	31.1
		(.008)	(.008)	(.008)	(.010)		(.008)	(.008)	(.008)	(.010)					
Ages 35–49	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	29.2	25.8	26.9	32.9	31.4

Ages 50+	.147	.138	.171	.182	.072	.189	.162	.318	.126	.117	35.8	37.2	35.5	33.7	37.5
		(.008)	(.008)	(.007)	(.009)		(.008)	(.008)	(.007)	(.009)					
Northeast	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	20.8	22.9	21.1	19.7	19.1
Midwest	.026	.030	.051	.052	−.073	−.022	.046	−.091	.046	−.110	26.1	29.0	27.3	24.3	23.3
		(.009)	(.009)	(.009)	(.012)		(.009)	(.009)	(.009)	(.012)					
South	.103	.129	.103	.111	.060	.076	.187	.009	.091	.026	34.2	32.2	33.2	35.4	36.2
		(.009)	(.008)	(.009)	(.011)		(.009)	(.008)	(.009)	(.011)					
West	.027	.004	.048	.061	−.034	.030	.048	.011	.077	−.031	19.0	16.0	18.4	20.5	21.4
		(.011)	(.010)	(.010)	(.013)		(.011)	(.010)	(.010)	(.013)					
Net effect	$\bar{x}_t\boldsymbol{\beta}_t$	.023	.019	−.030	−.015	$\bar{x}_t\boldsymbol{\beta}_t$	1.151	1.101	.910	.933					
Decomposition:															
Baseline	$\bar{x}\bar{\boldsymbol{\beta}}$	.000	.000	.000	.000	$\overline{x\boldsymbol{\gamma}}$	1.000	1.000	1.000	1.000					
Changing coefficients	$\bar{x}(\boldsymbol{\beta}_t - \bar{\boldsymbol{\beta}})$	−.016	.016	−.021	.003	$\bar{x}(\boldsymbol{\gamma}_t - \bar{\boldsymbol{\gamma}})$	.119	.004	−.080	−.077					
Changing composition	$(\bar{x}_t - \bar{x})\boldsymbol{\beta}_t$	.038	.003	−.009	−.019	$(\bar{x}_t - \bar{x})\boldsymbol{\gamma}_t$	.032	.006	−.010	.009					

Note. Coefficient estimates, sample proportions, and weighted averages are shown from estimating a generalized ordered probit regression in which the unobserved latent happiness index is estimated as $\text{Happiness}^*_{it} = \boldsymbol{x}_t\boldsymbol{\beta}_t + \varepsilon_{it}$, where $\varepsilon_{it} \sim N(0, \boldsymbol{x}_t\boldsymbol{\gamma}_t)$, where the subscript t refers to separate decades. Cut points are $\delta_1 = -1.29$ and $\delta_2 = .42$. Log likelihood $= -40{,}707.940$. Omitted categories are high school graduates, men, whites, married, ages 35–49, and Northeast. Standard errors are shown in parentheses. $N = 46{,}303$ observations from the General Social Survey, 1972–2006.

but by the end of the sample are similar to what were seen in the 1970s. Happiness trends by region have been roughly common across space.

Figure 7 shows that the decline in happiness inequality since the 1970s has occurred pretty much in parallel across demographic groups. There appears to be a fair bit of noise in these annual estimates, and in no case do the data make a convincing case for sharply different trends in within-group happiness inequality.

We also present a more compact representation of our results in Table 3, where we analyze changes by decade rather than year, so as to reduce the number of coefficient estimates to a manageable size (and reduce statistical noise). This approach has the advantage of allowing distinct patterns by decades to be examined for each group. This analysis contains all of the interactions in Table 2 along with time trends by region but differs from that table in that it shows conditional estimates. The first row of Table 3 reports the decadal trends for our baseline group—35–49-year-old white, married males with only a high school degree who live in the Northeast. For this group we see that happiness fell in the 1980s and 1990s but rose in the 2000s such that there is little difference in happiness between 1972 and 2006. Among members of this group, inequality in happiness follows the pattern seen in the aggregate population and is lower in the most recent period than in the 1970s. To compare these men with similar men who completed a college degree, we turn to the second row, which reports how the trends differ for college graduates. Adding the second row to the first row provides the trends 35–49-year-old white, married males with a college degree who live in the Northeast. Similarly, adding estimated coefficients for women to those in the top two rows would provide the trends for the equivalent female.

The estimated trends in Table 3 illustrate that some of the unconditional trends in Table 2 reflect the coincidence of trends in other categories. For example, examining those who are not married, we see that, conditional on other trends, happiness inequality was higher in the 1980s than the 1970s, a distinctly different pattern from what we have seen thus far. Yet the broad trends can still be seen—happiness levels are higher among women in the 1970s, with the gender gap narrowing in the ensuing decades. Similarly, happiness levels are lower among nonwhites, yet the gap narrows over the decades. The dispersion in happiness is higher among women and nonwhites in the 1970s, as seen previously.

A Decomposition

The key new finding in this paper is the fact that happiness inequality has declined since the 1970s, even if it has risen somewhat in recent years. In turn, the aggregate trend in happiness inequality shown in Figure 3 reflects the influence of changing average levels of happiness between groups (shown in Figure 6), changing happiness inequality within groups (shown in Figure 7), and changing proportions of the population in each group (Figure 8). In order to assess the combined impact of these separate influences, we now turn to a decomposition exercise, along the lines suggested by Lemieux (2002).

We begin by noting that our full regression, equation (11), expresses the latent happiness index as a function of individual characteristics with time-varying coefficients:

$$\text{Happiness}^{*}_{it} = \boldsymbol{x}_{it}\boldsymbol{\beta}_t + \sqrt{x_{it}\boldsymbol{\gamma}_t}\varepsilon_{it}, \tag{12}$$

where i denotes the individual observation, t is the time period, $\varepsilon_{it} \sim N(0, 1)$ is the error term, $\boldsymbol{x}_{it}$ is the $(1 \times k)$ vector of binary covariates described more fully in equation (11), where the scalar at the jth position denotes membership in demographic group j, and $\boldsymbol{\beta}_t$ and $\boldsymbol{\gamma}_t$ are time-varying $(k \times 1)$ vectors of parameters.

The mean level of happiness in each period can thus be expressed as

$$\mu_t = \bar{x}_t\boldsymbol{\beta}_t, \tag{13}$$

where $\bar{\boldsymbol{x}}_t = \sum_i \omega_{it}\boldsymbol{x}_{it}/\sum_i \omega_{it}$ is a vector in which the scalar at position j represents the proportion of the population in period with t characteristic j and ω_{it} refers to each observation's sampling weight, A simple Oaxaca decomposition allows us to describe changes in the mean as due to changes in the time-varying coefficients versus changes in the composition of the sample:

$$\mu_t = \bar{x}\bar{\boldsymbol{\beta}} + \bar{x}(\boldsymbol{\beta}_t - \bar{\boldsymbol{\beta}}) + (\bar{x}_t - \bar{x})\boldsymbol{\beta}_t, \tag{14}$$

where $\bar{\boldsymbol{x}} = \Sigma_t\Sigma_i\omega_{it}\boldsymbol{x}_{it}/\Sigma_t\Sigma_i\omega_{it}$ is a vector in which the scalar at position j represents the proportion of the whole sample that has characteristic j. The first term in this expression captures the average level of happiness in the sample, which is set to zero by our normalization. The second term captures changes due to time-series movements in the average happiness of various groups, captured by deviations of the betas from their means. Note that the each of the within-group changes shown in Figure

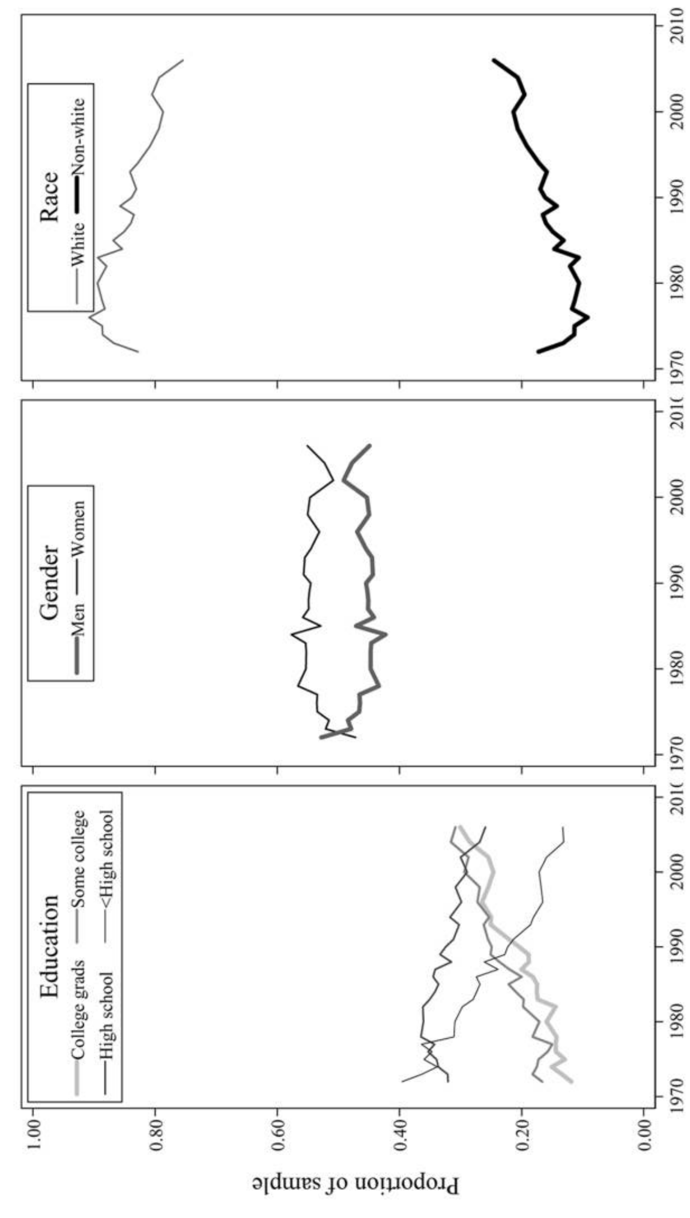
Education
College grads
Some college
High school
<High school
Gender
Men
Women
Race
White
Non-white
Proportion of sample
1.00
0.80
0.60
0.40
0.20
0.00
1970
1980
1990
2000
2010

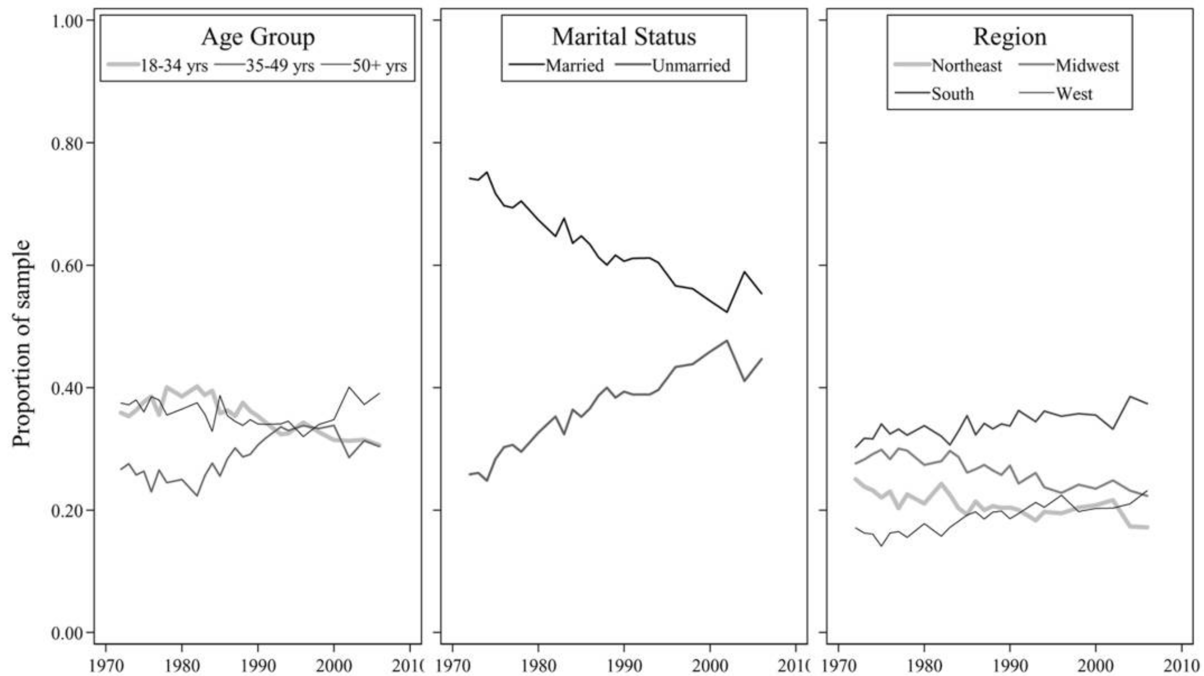

Figure 8. Changes in the proportion of the General Social Survey population with each demographic characteristic

6 are components of this vector, $[\boldsymbol{\beta}_t^1, \ldots, \boldsymbol{\beta}_t^k]$, and we aggregate these time series into a representative fixed-weight index by use of a time-invariant weighting vector $\bar{\boldsymbol{x}}$, which describes the proportions of people in the entire sample with each characteristic. Finally, the impact of changes in the proportion of the sample with each demographic characteristic is captured by the third term.

We show the estimates of these terms estimated using decadal means at the bottom of block 1 in Table 3. The net time series for average happiness is, if anything, somewhat negative over this period. Yet when a fixed-weight happiness is formed, $\bar{\boldsymbol{x}}(\boldsymbol{\beta}_t - \bar{\boldsymbol{\beta}})$, this trend decline disappears. The decadal estimates of the effects of compositional change, $(\bar{\boldsymbol{x}}_t - \bar{\boldsymbol{x}})\boldsymbol{\beta}_t$, illustrate how the population has shifted into demographic categories that have typically been less happy. In particular, the sample proportions (in block 3) show the population is increasingly nonwhite and unmarried—factors associated with lower levels of happiness—while simultaneously becoming older and more educated—factors associated with higher average happiness. On aggregate, these shifts have contributed to reducing the overall happiness in the population.

We can also write the variance of Happiness_{it}^* as the sum of within-group changes in happiness and a component due to changes in happiness levels between groups:

$$V_t = \overbrace{\bar{x}_t\boldsymbol{\gamma}_t}^{\text{within}} + \overbrace{\boldsymbol{\beta}_t'\Omega_t^x\boldsymbol{\beta}_t}^{\text{between}}, \tag{15}$$

where Ω_t^x is the variance-covariance matrix of $\boldsymbol{x}_{it}$ in period t.[14]

The within-group variance can be further decomposed:

$$\bar{x}_t\boldsymbol{\gamma}_t = \bar{x}\bar{\boldsymbol{\gamma}} + \bar{x}(\boldsymbol{\gamma}_t - \bar{\boldsymbol{\gamma}}) + (\bar{x}_t - \bar{x})\boldsymbol{\gamma}_t, \tag{16}$$

where the first term captures the average variance of happiness in our sample, which is set to one by our normalization. The second term reflects estimated time-series movements in the variance within each group, aggregated using a fixed-weight index. The third term reflects

14. Note that the literature on wage inequality sometimes refers to the first term in equation (15) as residual variance and the second term as explained variance. This terminology reflects the fact that these studies typically begin by running a regression of wages on observable variables (either in a separate regression for each year or, alternatively, in a single regression interacting observable variables with year fixed effects). The second term in equation (15) is explained by these first-stage regressions, while the first term reflects the variance of these residuals. By contrast, we model shifts in the level and dispersion of happiness in a single step: see equation (11).

the changes due to changing composition of the sample into groups prone to greater or lesser degrees of dispersion—a factor emphasized by Lemieux (2006) as an important explanator of rising residual wage in equality.

As with the means, we show the decomposition of the within-group variance using decadal means in block 2 of Table 3. The fixed-weight index of within-group changes in happiness inequality, $\bar{\boldsymbol{x}}(\boldsymbol{\gamma}_t - \bar{\boldsymbol{\gamma}})$, points to a substantial decrease in the within-group variance of happiness through to the 1990s. Turning to the estimated $(\bar{\boldsymbol{x}}_t - \bar{\boldsymbol{x}})\boldsymbol{\gamma}_t$ term, we see a qualitatively similar pattern—albeit with a much smaller quantitative contribution—with the changing composition of the sample contributing to falling dispersion through the 1990s and a rise in dispersion in the 2000s. All told, it appears that compositional change explains very little of the overall rise in residual happiness inequality.

Finally, the between-group variance can be decomposed as follows:

$$\boldsymbol{\beta}_t'\Omega_t^x\boldsymbol{\beta}_t = \bar{\boldsymbol{\beta}}'\Omega^{\bar{x}}\bar{\boldsymbol{\beta}} + (\boldsymbol{\beta}_t - \bar{\boldsymbol{\beta}})'\Omega^{\bar{x}}(\boldsymbol{\beta}_t - \bar{\boldsymbol{\beta}}) + \boldsymbol{\beta}_t'(\Omega_t^x - \Omega^{\bar{x}})\boldsymbol{\beta}_t, \tag{17}$$

where $\Omega^{\bar{x}}$ is the variance-covariance matrix of $\boldsymbol{x}_i$, estimated using data from all time periods.

We combine these time-series movements in difference in average levels of happiness between groups ($\boldsymbol{\beta}_t$, shown in Figure 6), within-group dispersion in happiness ($\boldsymbol{\gamma}_t$, shown in Figure 7), and the proportion of the population with each demographic ($\bar{\boldsymbol{x}}_t$, shown in Figure 8) to yield a useful decomposition of the overall trends in the distribution of happiness, shown in Figure 9.[15] The decompositions come from five models: $\text{Happiness}_{it}^1 = \bar{\boldsymbol{x}}\bar{\boldsymbol{\beta}} + \bar{\boldsymbol{x}}\bar{\boldsymbol{\gamma}}\sqrt{\varepsilon it}$, $\text{Happiness}_{it}^2 = \bar{\boldsymbol{x}}\bar{\boldsymbol{\beta}}_t + \bar{\boldsymbol{x}}\bar{\boldsymbol{\gamma}}\sqrt{\varepsilon it}$, $\text{Happiness}_{it}^3 = \bar{\boldsymbol{x}}_t\bar{\boldsymbol{\beta}}_t + \bar{\boldsymbol{x}}\bar{\boldsymbol{\gamma}}_t\sqrt{\varepsilon it}$, $\text{Happiness}_{it}^4 = \bar{\boldsymbol{x}}_t\bar{\boldsymbol{\beta}}_t + \bar{\boldsymbol{x}}\bar{\boldsymbol{\gamma}}_t\sqrt{\varepsilon it}$, and $\text{Happiness}_{it}^5 = \bar{\boldsymbol{x}}_{it}\bar{\boldsymbol{\beta}}_t + \bar{\boldsymbol{x}}_{it}\bar{\boldsymbol{\gamma}}_t\sqrt{\varepsilon it}$. (See equations (11)–(17) for details on each decomposition.) This figure shows that changes in the variance of happiness are being driven more by changes in within-group variance than by changes in between-group variance. However, there is a slight downward trend in between-group variance that is contributing to the overall decrease in happiness inequality since the 1970s. In sum, the figure illustrates that while the happiness convergence by race and gender played a role in

15. Of course, alternative decompositions exist, and we can vary the order in which compositional versus within-group changes are considered or which period to use as the index base (we choose the sample averages rather than any specific year). Table 3 provides the raw data necessary for these alternative decompositions; our analysis suggests that these alternative approaches do not much change the character of our results.

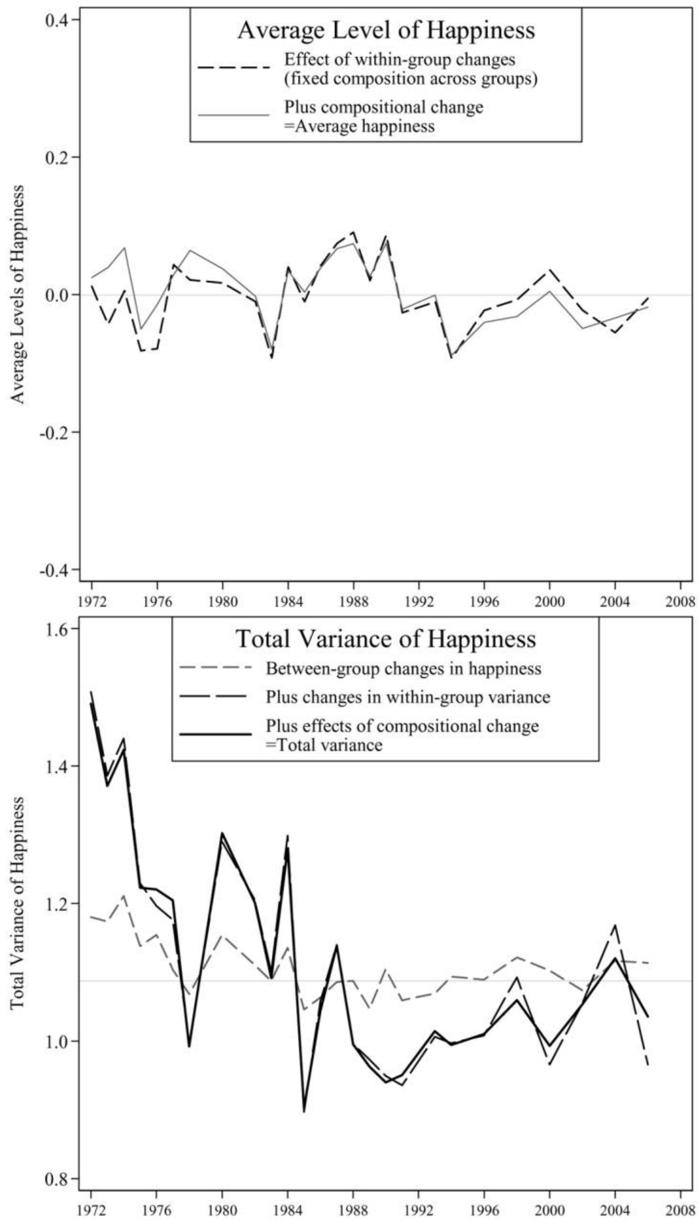

Figure 9. Decomposing the evolution of the distribution of happiness (solid lines show the aggregate observed level of each statistic; dashed lines show results from decompositions that abstract from the role of compositional change).

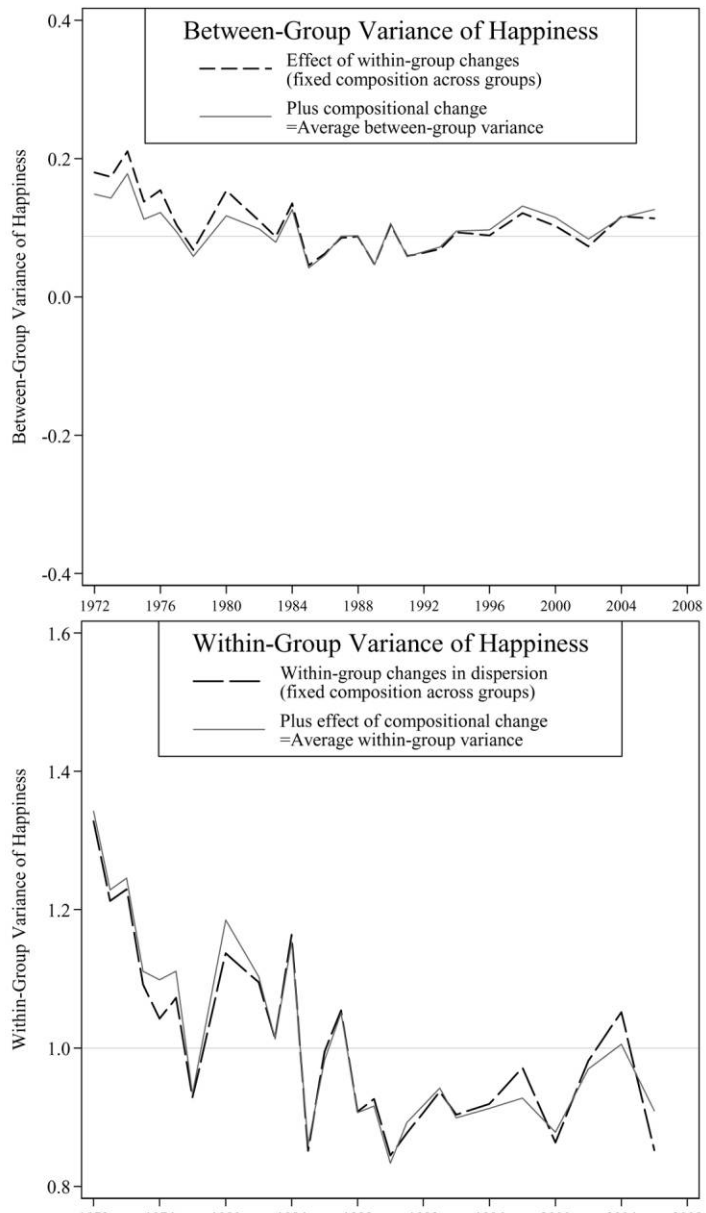

Figure 9. *Continued*

reducing inequality, this role was small compared to the overall decline in the inequality of happiness that is seen within each demographic category.

5. DISCUSSION

While there has been no increase in aggregate happiness over recent decades, there have been large changes in the level of happiness across groups. Much of the racial happiness gap has closed, the gender happiness gap has disappeared and perhaps inverted, and differences in happiness by education have widened substantially. More generally, we document a pervasive decline in happiness since the 1970s, albeit with some reversal over the past decade or so. That these trends differ from trends in both income growth and income inequality suggests that a useful explanation may lie in the nonpecuniary domain. As such, we suspect that our data are best interpreted in the broader context of a host of economic, social, and legal changes impacting equality in the United States over the past 35 years. There is much more work to be done in unraveling just how these forces are affecting the distribution of happiness in the United States.

In addition to the changes in both the level and dispersion of happiness between groups, there have been large demographic shifts that have potentially impacted happiness aggregates. Throughout this paper we have developed an integrated approach to measuring inequality and decomposing changes in the distribution of happiness. We examine how the composition and average happiness have changed both within and between demographic groups, paying particular attention to demographic and socioeconomic factors known to impact happiness such as education, marital status, age, race, and gender. This decomposition points to changes in the dispersion of happiness within groups as the main driver of declining happiness inequality. However, while this is a useful accounting exercise, it still leaves unanswered the question of just what it is that is creating less inequality in the subjectively experienced lives of demographically similar people.

APPENDIX: CORRECTING FOR QUESTION ORDER EFFECTS IN THE GENERAL SOCIAL SURVEY

While the General Social Survey has maintained the same question about happiness since its inception in 1972, responses seem to be quite sensitive to the immediately preceding battery of questions, and this ordering has changed several times. We provide this appendix in the hope that it will help the field settle on a widely accepted and accurate time series.

There are two key changes in question ordering:

1. whether a question probing marital happiness (asked only of married couples) immediately precedes the general happiness question and
2. whether a five-question battery probing domains of satisfaction immediately precedes the happiness question.

The first context occurs every year except in 1972, which is replicated in split-ballot experiments affecting only one-third of the respondents in 1980 and 1987 (those assigned form 3). Smith (1990) notes that these split ballots suggest that levels of happiness among married respondents tend to be higher in these instances in which they are preceeded by a question about marital happiness.

The second change in question ordering affects all respondents in 1972 and 1985, and its impact can be assessed by virtue of the fact that it was replicated for 1986 form 2 respondents and forms 2 and 3 respondents in 1987. Smith (1990) finds that aggregate happiness is higher in the years in which the happiness question is preceded by a five-item satisfaction scale.

Because of the split-ballot experiments run in 1980 and 1987—in which one in three randomly assigned questionnaires dropped the marital satisfaction question—and similar experiments run in 1986 and 1987, in which the satisfaction scale was dropped in one-third and two-thirds of the forms, respectively, we can assess the changes induced by these question order effects. These experiments are particularly useful in that statistically similar populations are assigned different contexts.

Thus, we use these experiments to calculate a set of sampling weights that correct for the undersampling of relatively happy people in 1972, 1980, and 1985–87 as well as the oversampling of happy married people in 1980 (see Table A1).

While our analysis largely follows Smith's suggestions, we differ in two respects. First, we do not simply drop the experimental forms from the sample but include their (appropriately adjusted) responses in computing our time series. And second, we also apply a slightly more sophisticated approach to measuring and correcting for these biases. In particular, given our interest in measuring the full distribution of happiness, it is important that we provide corrections for the share who are very happy, pretty happy, and not too happy.

In order to estimate the extent of these biases, we regress happiness on a dummy variable equal to one for those affected by each sampling change (the first change affected married people in 1972 and married form 3 respondents in 1980 and 1987; the second change affected all 1972 and 1985 respondents as well as the experimental 1986 form 2 respondents and 1987 form 2 and 3 respondents, controlling for year fixed effects, entered separately for both married and unmarried respondents. Our dependent variables are separate dummies

Table A1. Correcting General Social Survey Happiness Data for Question Order Effects

	Raw Data (%)			Corrected (%)			
Year	Not Too Happy	Pretty Happy	Very Happy	Not Too Happy	Pretty Happy	Very Happy	Sample Size
1972	17.2	53.0	29.7	13.6	49.1	37.3	1,606
1973	12.3	50.9	36.8	12.3	50.9	36.8	1,500
1974	12.5	49.2	38.3	12.5	49.2	38.3	1,480
1975	13.0	53.6	33.4	13.0	53.6	33.4	1,485
1976	12.2	52.9	34.8	12.2	52.9	34.8	1,499
1977	11.0	53.2	35.7	11.0	53.2	35.7	1,527
1978	8.4	56.2	35.5	8.4	56.2	35.5	1,517
1980	12.1	52.6	35.3	11.6	52.0	36.4	1,462
1982	11.7	53.5	34.8	11.7	53.5	34.8	1,505
1983	12.1	56.2	31.7	12.1	56.2	31.7	1,573
1984	11.6	52.1	36.3	11.6	52.1	36.3	1,445
1985	10.5	59.9	29.6	8.6	58.4	33.1	1,530
1986	10.2	56.6	33.2	9.2	55.8	35.0	1,449
1987	11.5	55.0	33.5	9.7	53.3	37.0	1,437
1988	8.2	55.7	36.1	8.2	55.7	36.1	1,466
1989	8.8	56.7	34.5	8.8	56.7	34.5	1,526
1990	7.7	56.5	35.7	7.7	56.5	35.7	1,361
1991	9.5	58.0	32.5	9.5	58.0	32.5	1,504
1993	9.7	56.9	33.4	9.7	56.9	33.4	1,601
1994	11.3	58.2	30.5	11.3	58.2	30.5	2,977
1996	10.5	57.4	32.1	10.5	57.4	32.1	2,885
1998	10.9	55.9	33.3	10.9	55.9	33.3	2,806
2000	9.6	56.4	33.9	9.6	56.4	33.9	2,777
2002	11.3	55.8	32.9	11.3	55.8	32.9	1,369
2004	11.7	54.7	33.6	11.7	54.7	33.6	1,337
2006	10.6	55.9	33.5	10.6	55.9	33.5	2,828

Note. Estimates are based on the sample weight WTSALL, omitting black oversamples in 1982 and 1987 and the Spanish-language interviews in 2006 that could not have been completed were English required (as in previous years). The corrected series make adjustments in 1972, 1980, 1985, 1986, and 1987 for question order effects.

corresponding to each of three possible happiness responses. Thus, the ballot experiments identify the effect of changing questionnaire order separate from background trends in happiness by marital status. These estimates suggest that the absence of the question about marital happiness led to a statistically significant decline (of about 5.4 percent) in the proportion of married respondents reporting themselves to be “pretty happy,” while the absence of the preceding satisfaction questions led to a statistically significant rise (of about 2.7 percent) in the proportion of respondents claiming they were “not too happy.” The aggregate happiness time series is simply the unadjusted annual happiness aggregates less the estimated question order effects (for those subject to the varied question order). The Stata code required to estimate these effects is available on our Web pages.[16]

REFERENCES

Aguiar, Mark, and Erik Hurst. 2007. Measuring Trends in Leisure: The Allocation of Time over Five Decades. *Quarterly Journal of Economics* 122: 969–1006.

———. 2008. The Increase in Leisure Inequality. NBER Working Paper No. 13837. National Bureau of Economic Research, Cambridge, Mass.

Atkinson, Anthony B. 1970. On the Measurement of Inequality. *Journal of Economic Theory* 2:244–63.

Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura. 2004. What Really Happened to Consumption Inequality in the US. NBER Working Paper No. 10338. National Bureau of Economic Research, Cambridge, Mass.

Autor, David, Lawrence Katz, and Melissa Kearney. 2008. Trends in U.S. Wage Inequality: Revising the Revisionists. *Review of Economics and Statistics* 90: 300–323.

Blanchflower, David, and Andrew Oswald. 2004. Well-Being over Time in Britain and the USA. *Journal of Public Economics* 88:1359–86.

Blau, Francine, and Lawrence Kahn. 2006. The U.S. Gender Pay Gap in the 1990s: Slowing Convergence. *Industrial and Labor Relations Review* 60: 45–66.

Brooks, Arthur C. 2008. *Gross National Happiness: Why Happiness Matters for American—and How We Can Get More of It.* New York: Basic Books.

Couch, Kenneth, and Mary Daly. 2003. The Improving Relative Status of Black Men. *Journal of Income Distribution* 12:56–78.

Cutler, David, and Lawrence Katz. 1991. Macroeconomic Performance and the Disadvantaged. *Brookings Papers on Economic Activity,* pp. 1–74.

Davis, James A., Tom W. Smith, and Peter V. Marsden. 1972–2006. General Social Surveys. National Data Program for the Social Sciences Series No. 18. Chicago: National Opinion Research Center.

Deaton, Angus, and Christina Paxson. 1994. Intertemporal Choice and Inequality. *Journal of Political Economy* 102:437–67.

Donohue, John J., and Peter Siegelman. 1991. The Changing Nature of Employment Discrimination Litigation. *Stanford Law Review* 43:983–1033.

Easterlin, Richard A. 1995. Will Raising the Incomes of All Increase the Happiness of All? *Journal of Economic Behavior and Organization* 27:35–48.

16. See http://bpp.wharton.upenn.edu/betseys or http://www.nber.org/~jwolfers. For those interested in replicating our results, the simplest way is to define a weighting variable as follows: gen wt=WTSSALL if SAMPLE~=4 & SAMPLE~=5 & SAMPLE~=7 & SPANINT~=2.

Goldin, Claudia, and Lawrence Katz. 2007. Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing. *Brookings Papers on Economic Activity,* pp. 135–65.

———. 2008. *The Race between Education and Technology.* Cambridge, Mass.: Harvard University Press.

Gottschalk, Peter, and Robert Moffitt. 1994. The Growth of Earnings Instability in the US Labor Market. *Brookings Papers on Economic Activity,* pp. 217–72.

Hacker, Jacob. 2006. *The Great Risk Shift.* New York: Oxford University Press.

Isen, Adam, and Betsey Stevenson. 2008. Women's Education and Family Behavior: Trends in Marriage, Divorce and Fertility. Working paper. University of Pennsylvania, Wharton School, Philadelphia.

Jensen, Shane, and Stephen Shore. 2008. Changes in the Distribution of Income Volatility. Unpublished manuscript. John Hopkins University, Department of Economics, Baltimore.

Johnson, David, and Stephanie Shipp. 1997. Trends in Inequality Using Consumer Expenditures: 1960 to 1993. *Review of Income and Wealth* 43: 133–52.

Kahneman, Daniel, and Alan B. Krueger. 2006. Developments in the Measurement of Subjective Well-Being. *Journal of Economic Perspectives* 20:3–24.

Kalmijn, Wim, and Ruut Veenhoven. 2005. Measuring Inequality of Happiness in Nations. In Search of Proper Statistics. *Journal of Happiness Studies* 6: 421–55.

Krueger, Dirk, and Fanrizio Perri. 2006. Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *Review of Economic Studies* 73: 163–93.

Lemieux, Thomas. 2002. Decomposing Changes in Wage Distributions: A Unified Approach. *Canadian Journal of Economics* 35:646–88.

———. 2006. Increased Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill. *American Economic Review* 96:461–98.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2003. Disagreement about Inflation Expectations. *NBER Macroeconomics Annual,* pp. 209–48.

Oswald, Andrew J. 2008. On the Curvature of the Reporting Function from Objective Reality to Subjective Feelings. *Economics Letters* 100:369–72.

Smith, Tom W. 1990. Timely Artifacts: A Review of Measurement Variation in the 1972–1989 GSS. GSS Methodological Report. University of Chicago, National Opinion Research Center.

Stevenson, Betsey, and Justin Wolfers. 2007a. Marriage and Divorce: Changes and Their Driving Forces. *Journal of Economic Perspectives* 21:27–52.

———. 2007b. *The Paradox of Declining Female Happiness.* Unpublished manuscript. University of Pennsylvania, Wharton School, Philadelphia.

———. 2008a. Economic Growth and Happiness: Reassessing the Easterlin Paradox. *Brookings Papers on Economic Activity*, spring, pp. 1–87.
———. 2008b. Trends in Marital Stability. Unpublished manuscript. University of Pennsylvania, Wharton School, Philadelphia.