UNIVERSITY OF CALIFORNIA, BERKELEY

Institute for the Study of Societal Issues

2420 BOWDITCH STREET BERKELEY, CA 94720-5670

Draft: Comments welcome December 2012

Recession, Religion, and Happiness, 2006–2010

Michael Hout & Orestes P. Hastings
University of California, Berkeley

Revised December 2012

ABSTRACT. The General Social Survey panel of 2006-2010 tracked Americans' reactions to the election of Barack Obama and the Great Recession (officially lasting from December 2007 to March 2009) as well as to events in their personal lives. Americans were less happy in 2010 than in 2006; the percent "very happy" decreased by four percentage points and the percent "not too happy" increased almost as much. In the cross section, standard predictors — including church attendance, family income, and marital status — continued to matter. Looking from the dynamic perspective of the panel study, though, we see that job loss and changing family incomes mattered far more. Changes in marital status were also important for changing morale; marrying made people happier while divorcing made them unhappy. The subjective sense that finances were improving mediated a significant portion of the income-happiness association. The effects of unemployment and marriage are estimated in a way that supports the inference that finding a job or a mate actually causes happiness to increase.

Recession, Religion, and Happiness, 2006–2010

Introduction

Happiness became a mainstay of social science research at some point in the last twenty years. Psychologists, economists, and sociologists descended on the moods and morale of ordinary people. Each brought a characteristic disciplinary focus to the subject. Psychologists and behavioral economists examined happiness as an emotional state. Experiences and relationships bring immediate pleasure or displeasure; data must be collected day-to-day and even hour-to-hour (e.g., Diener 1984; Kahneman et al. 2004). Sex and relaxing with friends are the keys to happiness; traffic and talking to the boss wreck one's mood (Krueger and Schkade 2008). Labor and health economists took a longer focus, viewing happiness as the product of separate choices that accumulate good feelings along with goods and services (Easterlin 1995; Kahneman and Deaton 2010). Money, the material satisfactions it buys, and the security that comes from a high income produce an overall sense of happiness; insufficient funds result in both material needs and feelings of insecurity. Sociologists take an even longer view, identifying long-standing relationships with family and friends, worldview (including religion), social class, and lifestyle as sources of life satisfaction or dissatisfaction (Davis 1984; Schuessler and Fisher 1985; Kohler, Bereman, and Skytthe 2005; Hout and Greeley 2012). Both the labor economists and the sociologists rely on self-reported summary measures such as the surprisingly robust simple question, "How happy would you say you are? Would that be very happy, pretty happy, or not too happy?" (Bradburn and Noll 1969) or Cantril's "ladder" (Kahneman and Deaton 2010).

We incline toward the longer view and seek to meld the economic and sociological perspectives. Researchers in this field have found, per expectations, that affluent, married, and religiously active Americans report more happiness while poor, divorced, and secular people report less (Easterlin 1995; Kohler et al. 2005; Hout and Greeley 2012). Dozens of cross-sectional studies replicate the pattern. In 1984, Davis (1984) summarized it as "new money, an old man/lady, and 'two's company'." The first element of his formula refers to the pattern, since replicated, whereby recent changes in income boost morale but people accommodate to new financial circumstances in time. The second captures the difference between married people and others. The third element refers to his observation that people in two-adult households were happier than those in other living arrangements.

While robust correlations abound, evidence of causal impacts for income, marriage, religion, and the like is scarce. Causal inference from observational data is always difficult (Morgan and Winship 2007). Three puzzles highlight how it is particularly difficult in the happiness research.

The first puzzle, known as the "Easterlin Paradox," concerns income. Richard Easterlin (1995) questioned whether a rise in income could actually cause a rise in happiness. He observed that the average American was about as happy in the early 1990s as in the 1960s even though the nation grew far richer in that period. He asked why the robust correlations between family incomes and personal happiness over many decades yielded no change in aggregate happiness. Something was interfering with the salutary effects of affluence. Easterlin's conjecture was that people adapted to prevailing conditions quickly, raising their expectations as incomes rose. Relatively affluent people were relatively happy, but the money itself was not procuring happiness. Since then others have noticed that with enough data — obtained by compiling time series from many countries — it is possible to find a nondecreasing but nonlinear relationship between national income and average happiness, efficiently shown by plotting incomes on a ratio scale (Kahneman and Deaton 2011). Fischer (2007) also noted that while GDP per capita rose fairly consistently in the United States during the period discussed by Easterlin and since, wages did not. He finds that aggregate happiness correlates quite well ($r = 0.45$) with hourly wages in U.S. data spanning 1973-2004.

The second puzzle concerns heterogeneity. Some people have unobserved and maybe even unobservable attributes that make inference difficult. We need controls for that kind of heterogeneity, from following people over time (as we do here) or matching them as in twin studies (Kohler et al. 2005). There can also be heterogeneity in effects. The same event or thing might make some people happy and others unhappy. Free cigarettes might raise a smoker's happiness but induce unhappiness in the non-smokers around her. Hout and Greeley (2012) discuss heterogeneity in relation to religion. Religious people presumably find comfort in their beliefs about cosmology, life, death, and eternity and enjoy participating in religious activities. According to Hout and Greeley (2012), stronger belief and more frequent activity makes the religiously prone happy and shields them from the sadness associated with death and disaster. Those tendencies among the "treated" should not, however, lead us to infer that inducing religious identification in the secular or imposing church attendance on the lapsed would make them any happier. It might even make them less happy.

The third puzzle concerns big events. Other lines of research have used the shock of a big event as a natural experiment. Some event like the attacks of September 11, 2001 on New York and Washington affected many Americans. The percentage “very happy” in the next GSS, fielded in spring 2002, was uncommonly low. Religious people were relatively unaffected; the net decline in happiness occurred among secular Americans (Hout and Greeley 2012).

The GSS panel spanning the period from spring 2006 to spring 2010 offers the opportunity to assess how two big events — a major recession and the election of Barack Obama — affected Americans’ happiness. The recession that the NBER dates as beginning in December 2007 and ending in June 2009 slowed the economy and raised unemployment for much longer than the 18 months in question. In May 2007 4.4 percent of Americans were out of work and seeking employment; the unemployment rate climbed to 10.0 percent in October 2009 (Hout and Cumberworth 2012). In other research, Hout and Hastings (2012) found significant changes in the way people answered questions about racial and gender stereotypes between 2006 and 2010, which they suggested might be attributed to the ways in which the prominence of Barack Obama’s and Hilary Clinton’s campaigns affected people’s attitudes about race and gender. Here we expect the recession to be the greater factor but will consider partisanship and race as well.

Statistical models and analytic strategy

The strength of a panel study like the GSS panel comes from the answers people give to the same questions at different times. The over-time variation in individuals’ answers allows us to statistically control for unchanging personal characteristics that we cannot or do not measure, increasing the chance that our statistical estimates represent quantities of interest. Formally, we specify that a person’s answer to a question, say the happiness question, depends on his or her attributes, some of which vary over time and some of which do not:[1]

$$y_{it} = \mu_t + \sum_{j=1}^{J} \beta_j X_{ijt} + \sum_{k=1}^{K} \gamma_k Z_i + \nu_i + \epsilon_{it} \tag{1}$$

[1] We provide details on data and measurement in the next section.

where $i(=1,\ldots,N)$ indexes people, $j(=1,\ldots,J)$ indexes the covariates that differ among people and over time, $k(=1,\ldots,K)$ indexes covariates that differ among people but not over time, $t(=$ 2006, 2008, 2010) indexes time, the ν_i are unobserved attributes of people that do not vary over time, and ϵ_{it} are unobserved attributes of people that vary over time (including errors of measurement).

In conventional survey data we cannot separate observables that vary over time, the Xs, from observables that do not, the Zs, or the accumulated unobservables that vary over time, ϵ_{it}, from the accumulated unobservables that do not, ν_i, so there is just one vector of covariates and one error term. In that formulation, the errors are assumed to be independent of the Xs and Zs so that any covariance between y and an X or y and a Z can be apportioned among the βs and γs.

With multiple observations of each case, the distinctions between Xs and Zs become visible, allowing researchers to separate them and to separate ϵ_{it} from ν_i (Halaby 2004). But, according to Morgan and Winship (2007, Ch. 9), more is required. Controlling for observed and unobserved factors is necessary but not sufficient. A researcher must also control for the possibility that selection affects the time trend, that is, that treated and control groups would have different time trends over the duration of the panel, even if no treatment occurred. In addition to being intrinsically different from the kind of people who are never treated, the kind of people who select themselves into treatment might be more different from never-treated people at one time compared with another.

Thus we approach these data as Morgan and Winship recommend. In fact, our statistical model, spelled out below, extends their equation 9.30 to competing treatments while retaining the key features of their approach, especially the specification that the ever-treated might be on a different time-trend than the controls are on. In this way we give ourselves the best chance of estimating the quantities of interest — how much happier treated people are, on average, than they would be in the absence of the treatment.

Data

Beginning in 2008, the GSS supplemented the standard cross-sectional survey with reinterviews conducted with selected 2006 respondents. A panel sample of 2,000 persons was randomly drawn from the completed 2006 interviews. Of these randomly selected adults, 1,536 (77 percent) were interviewed; 404 refused and 60 were ineligible (not living in a household or not living in the United States) or deceased. By 2010, 1,276 (64 percent of the the 2,000 original cases) were interviewed for the third time, 211 more refused, 32 more had died, 13 more were not living in a household, and 4 more were not living in the United States.

The dependent variable is the respondent's answer to the question: "How happy would you say you are? Would that be very happy, pretty happy, or not too happy?" We analyze the answers as a simple Likert scale ("not too happy" scored zero, "pretty happy" scored one, and "very happy" scored two) for simplicity. We replicated the results in two binary logit models: (1) very happy versus pretty happy or not too happy and (2) not too happy versus very happy or pretty happy. These alternative specifications imply the same substantive conclusions we reached based on the results we report here.

Past research has identified socioeconomic status (SES), family composition, and religion as important factors in this kind of general happiness (e.g., Hout and Greeley 2012).

Unemployment is the aspect of SES of greatest interest in the context of a panel spanning the Great Recession. We use a dummy variable coded one for the unemployed and zero for all others (including those out of the labor force). This is not a standard variable in prior studies. In the context of the Great Recession, though, we felt that this direct, personal experience of the recession was a key treatment to consider. Note in interpreting the results that we are comparing the unemployed to all other persons, not just to the employed. Unemployment statistics, for example, exclude people who are not in the labor force from calculations. We do not do that here because selecting on being in the labor force is an additional choice that might complicate our analysis. Exploratory analyses showed that adding additional labor force treatments such as being retired or a student did not alter these results in any notable way. Thus we focus on this binary treatment. Per the Morgan and Winship approach, we also include a dummy variable indicating whether the person was unemployed at any of the three waves of the panel and interact that selection variable with the time trend.

We use a standard family income measure as an additional SES indicator. Throughout the panel family income was measured in 25 categories. Respondents were shown a card with the 25 categories printed out and asked "In which of these groups did your total family income, from all sources, fall last year before taxes, that is?" We combined the eight categories for brackets below \$10,000, scored it ten, coded categories nine through twenty-four to their midpoint (in thousands of dollars), and scored the open category according to the Pareto formula as modified in Hout (2008)). We used the consumer price index for urban residents research series (CPI-URS) to adjust the incomes to 2010 prices. A peculiarity of GSS income measures is that they contain substantial differentiation at very low incomes and much less higher up the income scale. Exploratory analysis indicated that the relationship between income and happiness followed logged incomes better than incomes themselves for amounts higher than \$10,000 per year. Variation below \$10,000 per year was more or less random (see Appendix Figure 1). So we use the log of family income (in thousands of dollars standardized to 2010 prices) in our analyses, modified to code everyone with a family income below \$10,000 to the natural logarithm of ten.

Although having a religion is the treatment of greatest interest, Hout and Greeley (2012) showed that, in the cross-section, attendance at religious services completely mediated all the association between denominational affiliations and happiness, so we supplement the analysis of religion with the GSS attendance scale (from zero for "never" to eight for "more than once a week"), recoded to the probability of attending religious services in the last seven days, as our religion indicator (see Appendix Figure 2 for evidence that this is the appropriate functional form). In the Morgan and Winship scheme, it is obvious who the "ever treated" by religion group is and who the controls are when having a religion or not is the contrast of interest. When we shift to attendance, we consider those who attended "2-3 times a month" in at least one panel as ever-treated and those who never attended at least that often as controls.

We use a dummy variable for being married as a competing treatment. The married are often found to be happier than others. Preliminary analyses with more detail about marital status and children in various age categories yielded insignificant improvements to the fit of the model. Per the Morgan and Winship approach, we also include a dummy variable indicating whether the person was married at any of the three waves of the panel.

Finally we considered a novel hypothesis that in these partisan times the election of Barack Obama as president in 2008 (after the 2008 interviews were all complete) made Democrats happier and Republicans less happy. To that end we added the GSS party identification variable which scales people from strong Democrats (scored zero) to strong Republicans (scored six). We considered two alternative scorings as well. The first combined strong and not strong partisans so the scale just ran from one to five; the second was a set of dummy variables contrasting everyone to strong Democrats. We got null results for every specification we tried. Apparently the election was too light as a treatment to register in personal happiness 15-22 months later. The null results prompt us to drop party identification from the analyses we report below because including it removes 94 cases from the analysis — the 74 people that mentioned a party other than Democratic or Republican and 20 people who did not answer the question about political party.

All models contain statistical controls for gender, racial ancestry (blacks and latinos distinguished), education (four dummy variables contrasted to high school dropout), age (five dummy variables contrasted to 18-24 years old), panel (two dummy variables compared to 2006), and the interactions between panel and each of the ever-treated dummies.

Extending the Morgan and Winship model

The Morgan-Winship approach is designed to yield a regression coefficient for treatment of interest that estimates without bias the expected value of the difference between the outcome with and with the treatment, that is $\mathrm{E}[Y_{it}^1 - Y_{it}^0]$. Usual regression approaches come closer to estimating the difference between the expected value under treatment and the expected value in the absence of treatment, that is, $\mathrm{E}[Y_{it}^1] - \mathrm{E}[Y_{it}^0]$. A regression result $\mathrm{E}[Y_{it}^1] - \mathrm{E}[Y_{it}^0]$ equals the quantity of interest, $\mathrm{E}[Y_{it}^1 - Y_{it}^0]$, only under very special circumstances (see Morgan and Winship 2007, pp. 42-50). Thus, special adaptations of the regression model are necessary if the goal of the analysis is an estimate of $\mathrm{E}[Y_{it}^1 - Y_{it}^0]$.

To begin, suppose we have data on a sample of individuals (i) observed at selected times (t). The individuals are divided into the controls who will not receive treatment (or receive a placebo) and the treatment group who will receive treatment at some time. That distinction applies to all t, even before the treatment occurs, and is measured as a time-invariant binary variable D_i^* that is one for the treatment group (even before they are treated) and zero for the controls. For each individual at each time, we observe an outcome of interest (Y_{it}) and whether or not they have received the treatment yet (D_{it}). Morgan and Winship (2007, p. 269) summarize of their approach in the following equation (equation 9.30 in their book):

$$Y_{it} = a + bD_i^* + cT + c'D_i^*T + dD_{it} + e_{it} \tag{2}$$

for $i = 1, \ldots, N$ and $T = 0, \ldots\tau$, where e_{it} is a random error term uncorrelated with T, D_{it}, or D^*. In this formulation, time can be continuous or discrete, but its effects are homogenous within assignment, that is, Y changes by c per unit of time if D_i^* equals zero and $c + c'$ if D_i^* equals one. Furthermore, the trends are assumed to be linear (no subscript on c, c', or T).[2] The models works because including a main effect for D_i^* controls for selection at the outset, and the interaction between D_i^* and T allows the linear trend to differ for controls and the treated.

In a three-wave panel like the GSS, time is discrete. In this context, it seems reasonable to relax the linearity assumption and allow the change between waves 1 and 2 to differ from the change between waves 2 and 3 as well as differing between controls and treated.

Morgan and Winship developed this model for a one-time treatment, with time divisible into pre-treatment and post-treatment periods. In their scheme everyone starts out untreated and, once treated, remains treated; treatment is an irreversible feature of the research design. In panel data with socially selected treatments such as being out of work or married, treatment is not part of the research design; it is an observation. Observing that $D_{it} = 1$ at $t = 1$ does not imply that $D_{it} = 1$ at $t = 2$, nor can we assume that $D_{it} = 0$ for all i at $t = 0$ as Morgan and Winship do. We extend their model to handle the fact that we do not administer the treatment; we observe it.

Another important difference in observational data is the presence of competing treatments. For each treatment of interest, we have to consider separate time trends among the treated and controls. We handle that with interactions between treatment and time.

Incorporating all these adjustments to the three-wave panel and multiple treatments, our model takes the form:

$$\begin{aligned} Y_{it} = \mu + \sum_{j=1}^{3} \beta_j X_{ijt} \quad &+ \quad \sum_{j=1}^{3} \kappa_j X_{ijt}^* + \sum_{k=1}^{K} \gamma_k Z_{ikt} \\ &+ \quad \sum_{t=2}^{3} \tau_t Panel_t + \sum_{t=2}^{3} \sum_{j=1}^{3} \kappa_j' X_{ijt}^* Panel_t + \nu_i + \epsilon_{it} \end{aligned} \tag{3}$$

where the X_{ijt} are the treatment variables (j = 1,2,3)[3], the X_{ijt}^* mark the treatment groups, the Z_{kit} ($k = 1, \ldots, K$) are the control variables (time-varying or constant), $Panel_t$ denotes which wave of the panel, the βs are the treatment effects, the κs are the selection effects, the γs statistically adjust for the control variables, the τs are the residual or overall time-trends, the κ's adjust the

[2]The linear extrapolation for the treated creates the counterfactual after D_{it} switches from zero to one.

[3]Note that j indexed time-varying covariates in equation (1); it indexes treatments in equation (3).

counterfactual time trends for the treatment groups, the νs are time-invariant between-person random effects (uncorrelated with the Xs, Zs, Ds, D^*s, and ϵs), and the ϵs are within-person random effects (uncorrelated with the Xs, Zs, Ds, D^*s, and νs).

We use maximum likelihood random effects regression to estimate the parameters in equation (3) — the βs, κs, γs, τs, and κ's as well as the standard deviations of the νs and ϵs. Fixed effect regressions dropping the Zs yielded estimates of the βs and κs that did not differ significantly from the random-effects estimates we report below, indicating that the Morgan-Winship approach performs well, even in the presence of competing treatments.

Descriptive results

Table 1 presents the mean and its standard error by year for panel cases. Two points deserve attention. First, the proportion out of work is substantially below the annual unemployment rates for these years because the universe here is more inclusive. The official unemployment rate is the percentage of people in the labor force who are out of work and looking for work; our out of work mean is proportion of all adults who are out of work and looking for work. Second, the time-constant covariates (woman, black, and Latino) change slightly because missing data on the changeable variables alters composition slightly over time. Proportion Latino is most affected.

Americans were less happy about life in 2008 and 2010 than in 2006. The percentage very happy fell from 35 to 30 and then to 27 percent from 2006 to 2008 to 2010; the percentage not too happy rose from 11 to 15 percent between 2008 and 2010 (see Table 2). The Great Recession or some other event between 2006 and 2010 reduced Americans' general happiness.

People were fairly consistent in their degree of happiness; 63 percent gave the same answer in the reinterview as they gave two years earlier, 23 percent switched from very happy to pretty happy or vice versa, 11 percent switched from pretty happy to not too happy or vice versa, and only 3 percent switched between the extremes. Changes vary by initial happiness in ways that suggest modest unreliability combined with real downward drift in the population to produce these patterns. Fifty-four percent of people who were initially very happy reported they were very happy two years later; 40 percent were pretty happy; and 6 percent were not too happy. Among people who were initially pretty happy, almost three-fourths were also pretty happy two years later; those who changed were more likely to be very happy than not too happy. Between 2006 and 2008, half of the people who were initially unhappy became pretty happy and 38 percent were not too happy again in 2008. Between 2008 and 2010, 57 percent of those who were initially unhappy were not too happy again; 36 were pretty happy. Thus the greater unhappiness in 2010 was the result of less "upward mobility" in happiness between 2008 and 2010 than between 2006 and 2008.

Causes of happiness

We begin by replicating common results to verify that the novel findings in the panel data are due to the analytical leverage that the panel gives us and not to some peculiarity in the GSS panel sample. Our results, labeled "cross-sectional" in Table 3, confirm that happiness patterns among people in the GSS panel are not unusual. Married and affluent people were happier than unmarried and low-income people. We have two novel findings: out-of-work people were less happy than others and general happiness declined from 2006 to 2010 as the recession hit and its consequences persisted. The control variables — gender, racial ancestry, education, and age — were also significant. Somewhat surprisingly, religious people were not significantly happier than people with no religion.

Cross-sectional patterns, no matter how suggestive, can be suspect. Unobserved factors correlated with the independent variables and happiness could, by their absence, leave the impression that observed factors are more important than they really are. In particular, actually changing some variable by one unit may result in less than a $\hat{\beta}_{yx}$ change in happiness. As noted above, the repeated measures in panel data advance the potential to control for the unobserved heterogeneity, at least unobserved heterogeneity that does not vary over time. In particular, we use repeated observations to identify those people who ever experienced a treatment — those who reported they were married, out-of-work, or attended religious services at least monthly in at least one of the three interviews. Those people have selected themselves into treatment. By controlling for that fact and for the differing time trends people like that might have, we can remove much of the selection and otherwise unobserved variation that could bias cross-sectional estimates (Morgan and Winship 2007, ch. 9).

Two sets of estimates are Table 3, in the columns marked "baseline" and "attendance." The baseline estimates take having a religion as our measure of religious treatment; the attendance estimates use the frequency of attending religious services as the measure. Estimates of socioeconomic and demographic effects are nearly the same regardless of how we measure religiosity. The choice matters for our conclusion about religion, though. Religious identity is not significant; attendance is.

Our first treatment of interest is being out of work. Experiencing the deleterious effects of the Great Recession directly reduced happiness by a quarter point. A quarter point is large in absolute terms and about 40 percent of the standard deviation of the happiness score (0.63). We can hardly be surprised that losing a job is hard on a person's morale. What is surprising is how close the causal estimate is to the cross-sectional estimate. Selection into unemployment — the many factors that might predispose the kind of people at risk of being laid off to unhappiness -— accounts for only one-sixth of the cross-sectional estimate. Unemployment directly reduces happiness.

American adults were one-tenth of a point less happy in 2010 than 2000. The cross-sectional estimates indicate that net of the usual observables, including being out-of-work, people were significantly less happy in 2010 than in 2006 ($\hat{\beta}$ = -.057; s.e. = .026). The time trend among controls is no longer significant in the baseline model ($\hat{\beta}$ = -.035; s.e. = .072). To some extent the complex method of handling time trend in our extension of the Morgan-Winship approach makes time trend estimates less precise (note the larger standard error), but the smaller net difference between 2006 and 2010 would not be statistically significant even if the the standard error was a small as in the cross-sectional model.[4] In short, the overall time trend is the product of a profound drop in the happiness of those directly affected by the recession and a general drop so small that a sample of 1,200 individuals cannot reliably estimate its magnitude or even say with the standards we usually apply that it is significantly different from zero.

Marriage is our second treatment of interest. Long considered an important factor in people's happiness (Waite and Gallagher 2000), it is also among the factors most subject to selection bias. That is, it seems reasonable to think that, all else being equal, the kinds of people who would be

[4]Logically it is possible that the smaller time-trend estimate is attributable to the introduction of a random person effect, but we get nearly the same estimate for a random-effects version of the cross-sectional model ($\hat{\beta}$ = -.058; s.e. = .021) as in the cross-sectional model itself.

happier if they were married would be the most likely people to get married. Indeed that is somewhat the case. We estimate that marriage increases happiness by 0.18 points — an estimate only three-fourths as large as the cross-sectional estimate and not large enough to offset the negative effect of unemployment. Furthermore, the selection effect is almost as large ($\hat{\beta}$ = 0.14; s.e. =0.05) so concern on this front is very well placed. Finally the negative time trend, lacking among the controls, is significant among ever-married people. So counterfactually the kinds of people who get married were significantly less happy being unmarried during the recession than the controls were.

Our third treatment of interest is religion. The risk of selection bias in estimating the difference between religious and secular happiness is as high as with marriage. That is, the kinds of people who could be happy in a religious setting are far more likely to respond positively to religion than the kinds of people who are unlikely to be religious (Hout and Greeley 2012).

Merely identifying as having a religion is not sufficient to make a person happier. Surprisingly this is less due to a strong selection into religion than to other controls in the model. Attendance, on the other hand, is important for happiness. Replacing identification with attendance at religious services yields a significant estimate ($\hat{\beta}$ = 0.09; s.e. =0.04). Increasing attendance from never to weekly would raise happiness by .08 — the equivalent of raising one-twelfth of those who are unhappy or only pretty happy up one notch on the happiness ladder.

Money makes people happier. The estimates for family income are positive and significant in the cross-sectional, baseline, and attendance models. The magnitudes can be difficult to grasp. A one-point increase in the natural logarithm of family income increases happiness by 0.05 or 0.06. From the base of the income measure at $10 thousand (2.3 on the log scale), a one-point increase on the log scale corresponds to an increase in family income to $27 thousand, a two-point increase on the log scale corresponds to an increase to $73 thousand, and a three-point increase on the log scale corresponds to an increase to $200 thousand. Finding a job or a spouse would increase happiness more than boosting income from $10 to $200 thousand. Yet the increase is, nonetheless, significant, even after controlling for selection and time-invariant unobservables. As an aid to interpretation we present expected general happiness across the range of family incomes for each model considered so far and one more to be discussed later (see Figure 1).[5] Expected general

[5]These lines are all constrained to pass through the point corresponding to the means of income (3.76) and happi-

happiness ranges from 1.08 to 1.26 under the baseline and attendance models. These estimates imply that the increase in happiness from the poorest to most affluent Americans is about three times larger than the drop for all Americans between 2006 and 2010. Most of the increase occurs in the range up to \$60 or \$70 thousand; the rate of increase slows but does not reverse beyond that. Said another way, about two-thirds of the increase in happiness occurs in the lower one-third of the range.[6]

The relative income hypothesis

Despite the robust positive correlations between income and happiness, it seems unlikely to most scholars that money literally *buys* happiness. The Easterlin (1995) paradox and other issues suggest that relative income might be more pertinent for happiness than absolute income (Brady 2004; Firebaugh and Schroeder 2009). The relative income hypothesis is that people compare themselves to people like them. If they have more money than people like themselves, they feel better and report marginally more happiness than otherwise similar people do.

The test for the relative income hypothesis, as specified by Brady (2004) is a comparison of the coefficients for personal and comparison incomes. If people's happiness depends on their income compared to others', then the two coefficients will be equal in absolute value but opposite in sign when income is log-scaled. Brady derives the test from the fact that the logarithm of a ratio equals the difference between the logarithms of the numerator and denominator of the ratio, that is, $ln(a/b) = ln(a) - ln(b)$. Thus if

$$y_{it} = \beta_0 + \beta^* \frac{X_{it}}{\bar{X}_{it}} + \gamma Z_{it} + \nu_i + \epsilon_{it} \tag{4}$$

(where X_{it} is personal income, $\bar{X}_{it}$ is the comparison income, and Z_{it} is all the other covariates) is the correct equation, but we estimate

$$y_{it} = \beta_0 + \beta_1 X_{it} + \beta_2 \bar{X}_{it} + \gamma Z_{it} + \nu_i + \epsilon_{it} \quad , \tag{5}$$

then we will not be able to reject the null hypothesis $\beta^* = \beta_1 = -\beta_2$.

ness (1.17).

[6]That lower one-third of the income range encompasses 70 percent of respondents.

Uncertainty and controversy enter this literature in trying to define comparison groups with any precision. Aggregations by neighborhood, larger places, age, race, or some combination have all been used in the literature.

Here we consider income relative to place and age, separately. In both we used the standard cross-sectional GSSs from 2004-2010, pooled to assure reasonable sample sizes. For place, we regressed the logarithm of family income (after recoding all incomes below $10,000 to $10,000 to reduce the leverage of extremely low incomes) on dummy variables for gender, racial ancestry, year, and place.[7] From the regression results, we calculated the mean logged family income, adjusted for gender and racial composition, for each of the 79 places. For age, we repeated the procedure replacing place with single years of age to obtain mean logged family income, adjusted for gender and racial composition, for each of the 71 ages in the dataset.

We added each relative income measure to the "attendance" model. The coefficient for the adjusted mean income in the place was significantly less then zero but not significantly different from minus one times the coefficient for personal income. In Brady's (2004) scheme that is a positive result for the relative income hypothesis. The results are in the last two columns of Table 3.

Intervening variables

Davis (1984) and others have emphasized the importance of short-term changes in income — what Davis called "new money" — as the crucial link between income and happiness. Recent income change is also highly relevant to discussions of accommodation, the process by which people get used to having more money after an increase in their income. The panel data give us a unique opportunity to test the "new money" hypothesis. First, the panel design itself allows us to observe people's incomes at two or three times. Second, the context of the Great Recession means that more people than usual saw their incomes fall. In usual economic conditions economic growth and lifecycle increases in income predominate. In the recessionary times of 2006-2010, losses give us additional leverage to gauge the impact of income change on happiness.

[7] The sampling frame for the 2002-2010 GSSs included 79 primary sampling units, differentiated (but not identified) in the data file by the variable `sampcode`. We used `sampcode` to designate places (see Smith, Marsden, and Hout 2011 for details).

We specified four different "new money" measures: (1) the ratio of the person's inflation-adjusted family income in one wave to that in the previous wave (logged), (2) the same calculations without adjusting for inflation, (3) measure (1) divided into five categories, and (4) measure (1) divided into nine categories. None of these measures of the change in family income showed a significant relationship to happiness, controlling for the log of of family income. To illustrate, we include measure (1) in Table 4 as part of the "Income change: actual" model. The estimate is 0.20 (with a standard error of 0.20). Thus when we measure actual changes in family income over a two-year period, we find no evidence that new money matters more than older money.

Davis's analysis was based on a subjective assessment of income change, not on actual income change. The GSS question is "During the last few years, has your financial situation been getting better, worse, or has it stayed the same?" We scored "better" equal to one, "worse" equal to minus one, and "the same" as zero and added the variable to the attendance model. The results are shown in the "Income change: subjective" model in Table 4. The estimate is 0.10 (with a standard error of .01). Nearly as important, the income coefficient is now 0.030 (compared to 0.053 in the attendance model), and the out-of-work coefficient is now -.215 (compared to -.252 in the attendance model). In fact, the income coefficient, while still nearly sixty percent as large as before, is no longer statistically significant. We are not going to remove a core variable from the model, but it is fair to say that a subjective assessment of recent changes in finances mediates much of the effect of income on happiness. Adding a dummy variable for answering "same" does not improve the model's fit, so the linear income change variable is sufficient.

This is a substantively important result. Going from "worse" to "better" is a two-point increase, implying that the feeling that personal finances are improving that much is the equivalent of finding a job. It operates in a causal chain. Having more money or having a job raises happiness both directly and indirectly through the financial optimism captured by this variable. These results are symmetrical, of course, in the sense that feeling that one's personal economic trend has slipped from better to worse is the equivalent of losing a job. And that sense, too, mediates the effects of job loss and income loss.

Having entered the realm of subjective assessments, we consider subjective social class as another potential intervening variable. The GSS includes the classic question "If you were asked to use one of four names for your social class, which would you say you belong in: the lower class, the working class, the middle class, or the upper class?" We scored the answers 0-3 and added this variable to the income change: subjective model. The results are in the "subjective class" column of Table 4. The coefficient is positive, as expected, but it is neither statistically significant nor effective in mediating the effects of other variables.

We also sought out intervening variables for the relationship between religious attendance and happiness. That search was less successful. We considered the importance of religion in one's life as one potential mediator. And, following Hout and Greeley (2012), we considered belief in God. Neither added much to the model, and neither explained much of the relationship between religious attendance and happiness.

Conclusions

Happiness and its social correlates have burst from a few scattered papers in the 1970s to a very impressive literature in the past fifteen or so years. Happiness is now an important subject of academic research. Long term trends in general happiness are linked to both economic and religious attributes (Hout and Greeley 2012). The affluent and the religiously active appear happier in most surveys on the subject. Marriage and children (Kohler et al. 2005) are also important correlates of happiness.

Nearly all evidence to date from the United States considers cross-sectional patterns. We contribute to this literature some evidence from the relatively new GSS panel, spanning 2006 to 2010. In tracking a representative sample of Americans over time, we can supplement correlations derived from differences among people with evidence on how changes in those attributes correspond to changes in the happiness of the same individuals — all in a context of national economic stress that saw overall happiness decrease.

Jobs were the single most important factor in Americans' happiness during the Great Recession. People who lost their job in the recession reported substantially lower happiness than they would have otherwise. Similarly and symmetrically morale among the lucky ones who found work rebounded. Family incomes supported and supplemented these changes related to being out of work. An increase in family income resulted in modestly greater happiness; a decrease in income resulted in modestly less happiness. The effect is proportional in the sense that an increase in income from $10 to $20 thousand had the same impact as an increase from $20 to $40 thousand or $40 to $80 thousand. The effect of employment was so much greater than that of income, though, that finding a job was the equivalent of increasing income fifteen times over, for example, from $10,000 to $150,000 per year. Few of the unemployed who found jobs after the recession found jobs that good. And if they did their happiness was increased by both, compounding the good news.

The effect of objective income is largely mediated by how people think about the family finances. If they think they are doing better than others in their community or better than they were in the past, they are happier. If they lack one or the other of these assessments, they are less happy.

Attending religious services boosts happiness among the religious. The effect itself is quite modest. Increasing attendance from monthly to weekly would boost happiness by about 0.07 on a scale from zero to two. Yet it is statistically significant and quite robust across alternative specifications (all of which control for the fact that the kind of people who would choose attending are different from those who would not).

Life-partnership through marriage (or living together) was another essential element in happiness. Marrying increased happiness by an amount bigger than doubling family income. Dissolving a marriage reduced it by an equivalent amount.

Women were happier than men. Racial ancestry mattered in that whites and Asians are happier than Latinos. Blacks were happier than Latinos but less so than whites. Finally, we searched for partisan differences linked to the 2008 election and found none. Democrats may have been happier than Republicans on election night, but that was more a mood swing than a change in morale. As important as partisanship has become in American social life, it is not a lasting factor in Americans' morale.

References

Bradburn, Norman M., and C. Edward Noll. 1969. *The Social Structure of Happiness.* Chicago: Aldine.

Brady, Henry E. 2004. "An Analytical Perspective on Participatory Inequality and Income Inequality." Pp. 667-702 in *Social Inequality*, edited by Kathryn Neckerman. New York:

Russell Sage Foundation.

Davis, James A. 1984. “New Money, an Old Man/Lady, and ‘Two’s Company’: Subjective Welfare in the General Social Surveys, 1972-1982.” *Social Indicators Research* 15: 319-350.

Diener, Ed. 1984. “Subjective Well-Being.” *Psychological Bulletin* 95: 542-575.

Easterlin, Richard A. 1995. “Will Raising the Incomes of All Increase the Happiness of All?” *Journal of Economic Behavior and Organization* 27: 35-47.

Firebaugh, Glenn and Matthew B. Schroeder. 2009. “Does Your Neighbor’s Incomes Affect Your Happiness?” *American Journal of Sociology* 115: 805-831.

Firebaugh, Glenn, and Laura Tach. 2012. “Income, age, happiness.” Pp. 240-287 in *Social Trends in American Life* edited by Peter V. Marsden. Princeton: Princeton University Press.

Fischer, Claude S. 2007. “What Wealth-Happiness Paradox? A Short Note on the American Case.” *Journal of Happiness Studies* 9: 219-226.

Halaby, Charles N. 2004. “Panel Models in Sociological Research: Theory into Practice.” *Annual Review of Sociology* 30:507544.

Hout, Michael. 2008. “Making the Most of GSS Income Questions.” GSS Methodological Report # 101.

Hout, Michael, and Erin Cumberworth. 2012. “The Labor Force in the Great Recession: An Update.” Stanford: Center on Poverty and Inequality. Published online at: `https://www.stanford.edu/group/recessiontrends/cgi-bin/web/sites/all/themes/barron/pdf/LaborMarkets_fact_sheet.pdf`

Hout, Michael, and Andrew M. Greeley. 2012. “Religion and Happiness.” Pp. 288-314 in *Social Trends in American Life* edited by Peter V. Marsden. Princeton: Princeton University Press.

Hout, Michael, and Orestes P. Hastings. 2012. “Reliability Estimates for GSS Core Items.” GSS Methodological Report # 112.

Kahneman, Daniel, and Angus Deaton. 2010. "High income improves evaluation of life but not emotional well-being." *PNAS* 107: 16489-16493.

Kahneman, Daniel, and Alan B. Krueger. 2006. "Developments in the Measurement of Subjective Well-Being." *Journal of Economic Perspectives* 20: 324.

Kahneman, Daniel, Alan B. Krueger, David A. Schkade, Norbert Schwarz and Arthur A. Stone. 2004. "A Survey Method for Characterizing Daily Life Experience: The Day Reconstruction Method." *Science* 306 (5702): 1776-1780.

Kohler, Hans Peter, Jere Behrman, and Axel Skytthe. 2005. "Partner Plus Children Equals Happiness? An Assessment of the Effects of Fertility and Partners on Subjective Well Being in Danish Twins." *Population and Development Review* 31: 407-445.

Krueger, Alan B., and David A. Schkade. 2008. "The reliability of subjective well-being measures." *Journal of Public Economics* 92: 1833-1845.

Morgan, Stephen, and Christopher Winship. 2007. *Counterfactuals and Causal Inference.* Cambridge: Cambridge University Press.

Schuessler, Karl A., and Gene A. Fisher. 1985. "Quality of Life Research and Sociology." *Annual Review of Sociology* 11: 129-149.

Smith, Tom W., Peter V. Marsden, and Michael Hout. 2011. *General Social Survey Cumulative Codebook: 1972-2010*. Storrs CT: Roper Center.

Waite, Linda J., and Maggie Gallagher. 2000. *The Case for Marriage: Why Married People are Happier, Healthier, and Better Off Financially*. New York: Doubleday.

Tables

Table 1. Percentage distribution of happiness and means for each variable in the analysis by year

			Year	
Variable	[Range]	2006	2008	2010
Happiness				
Very happy (%)		35	30	27
Pretty happy (%)		55	59	59
Not too happy (%)		10	11	14
Mean	[0, 2]	1.24	1.19	1.13
		(0.02)	(0.02)	(0.02)
Treatments				
Out of work	[0, 1]	0.027	0.026	0.063
		(0.006)	(0.006)	(0.009)
Attend religious services	[0, 1]	0.36	0.37	0.35
		(0.02)	(0.02)	(0.02)
Married	[0, 1]	0.57	0.56	0.56
		(0.017)	(0.018)	(0.020)
Covariates				
Family income ($000s)	[10, 196]	58.2	65.1	64.4
		(1.9)	(2.3)	(2.1)
Woman	[0, 1]	0.55	0.55	0.54
		(0.02)	(0.02)	(0.02)
Black	[0, 1]	0.11	0.11	0.11
		(0.02)	(0.02)	(0.02)
Latino	[0, 1]	0.12	0.13	0.14
		(0.02)	(0.02)	(0.03)
Education	[0, 4]	1.95	1.95	1.93
		(0.06)	(0.05)	(0.06)
Age	[18, 89]	45.0	46.8	48.4
		(0.7)	(0.6)	(0.6)

Note: Means are for 1,237 individuals observed an average of 2.74 times.
Standard errors of the means in parentheses.
Source: Authors' calculations from the General Social Survey panel, 2006-2010.

Table 2. Current happiness by previous happiness and years

	Current happiness				
Previous happiness	Very happy	Pretty happy	Not too happy	Total	(N)
From 2006 to 2008					
Very happy	54	41	6	100	(412)
Pretty happy	19	72	10	100	(703)
Not too happy	10	51	38	100	(150)
Total	30	58	11	100	(1,265)
From 2008 to 2010					
Very happy	54	40	6	100	(371)
Pretty happy	17	73	10	100	(740)
Not too happy	8	36	57	100	(156)
Total	27	58	15	100	(1,267)

Figures

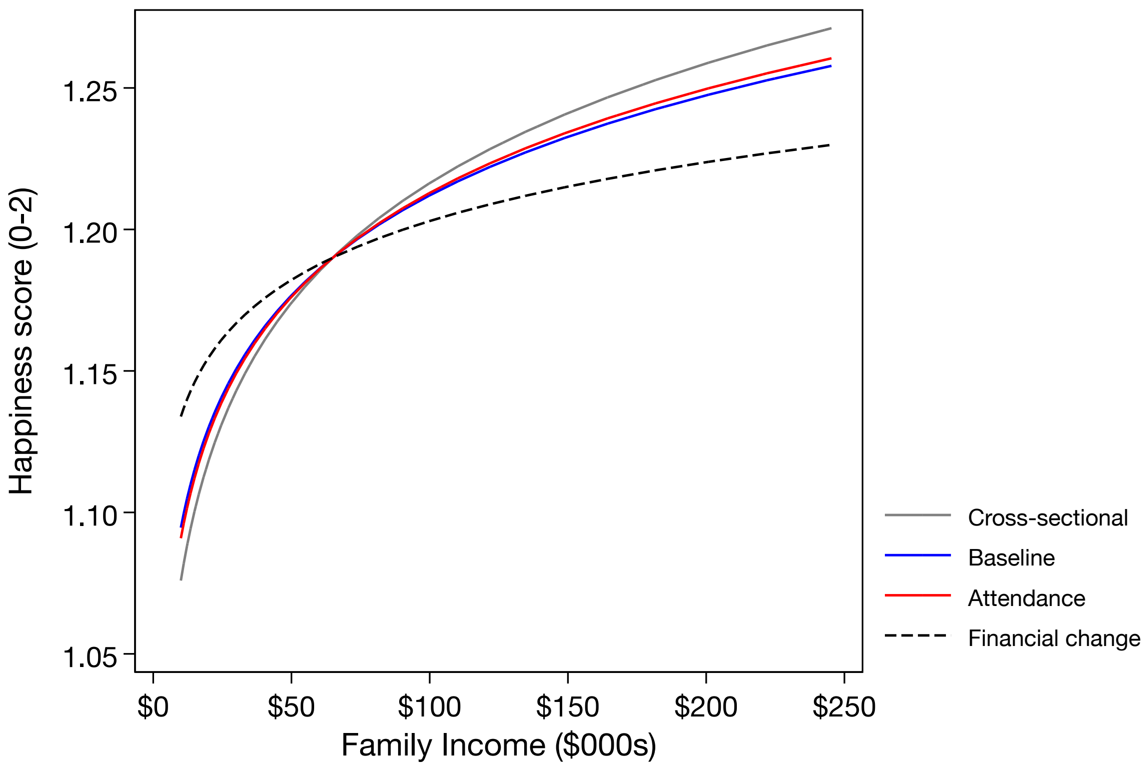

Figure 1. General happiness score by family income and model.

Note: Happiness scores are obtained by setting all variables except family income to their means.
Source: Authors' calculations based on Table 3; original data from General Social Survey panel 2006-2010.

Table 3. Coefficients for competing treatments, selection, income, trend, and differential trends in counterfactual models of general happiness

	Model				
	Cross-		Religious	Relative income:	
Independent variable	sectional	Baseline	attendance	Place	Age
Treatment					
Out-of-work	-.308	-.252	-.252	-.252	-.252
	(.054)	(.058)	(.058)	(.058)	(.058)
Any religion	.029	.001	—	—	—
	(.028)	(.037)			
Attend religious services	—	—	.091	.092	.091
			(.039)	(.039)	(.039)
Married	.246	.181	.175	.172	.175
	(.023)	(.039)	(.039)	(.039)	(.039)
Selection					
Ever out-of-work	—	-.041	-.042	-.040	-.042
		(.062)	(.062)	(.062)	(.062)
Ever religious or ever attend	—	-.008	.063	.061	.063
		(.072)	(.042)	(.042)	(.042)
Ever married	—	.140	.125	.122	.125
		(.050)	(.050)	(.050)	(.050)
Income					
Family income (logged)	.061	.051	.053	.063	.050
	(.016)	(.017)	(.017)	(.017)	(.017)
Relative income: place	—	—	—	-.063	—
				(.017)	
Relative income: age	—	—	—	—	-.050
					(.017)
Time trend: controls					
2008	-.030	.015	.030	.028	.031
	(.025)	(.071)	(.039)	(.039)	(.039)
2010	-.057	-.035	.010	.009	.011
	(.026)	(.072)	(.039)	(.039)	(.039)
Selection $\times$ time trend					
Ever out-of-work: 2008	—	-.041	-.041	-.041	-.041
		(.067)	(.067)	(.067)	(.067)
Ever out-of-work: 2010	—	-.012	-.004	-.003	-.003
		(.070)	(.070)	(.070)	(.070)
Ever religious or ever attend: 2008	—	.009	-.014	-.014	-.015
		(.070)	(.041)	(.041)	(.041)
Ever religious or ever attend: 2010	—	.039	-.024	-.027	-.024
		(.071)	(.041)	(.041)	(.041)
Ever married: 2008	—	-.086	-.085	-.085	-.085
		(.041)	(.041)	(.041)	(.041)
Ever married: 2010	—	-.099	-.093	-.094	-.093
		(.041)	(.041)	(.041)	(.041)
Between-person (σ_ϵ)	.353	.368	.362	.362	.362
		(.012)	(.012)	(.012)	(.012)
Within-person (σ_ν)	—	.465	.465	.465	.465
		(.007)	(.007)	(.007)	(.007)
Observations	3,392	3,392	3,392	3,387	3,392
Persons	1,237	1,237	1,237	1,237	1,237

Notes: Standard errors in parentheses. All models control for race, gender, education, and age.
Source: Authors' calculations from General Social Survey panel, 2006-2010.

Table 4. Coefficients for competing treatments, income, and intervening variables in counterfactual models of general happiness

	Intervening variable				
	Socioeconomic status			Religion	
	Income change:		Subjective	Importance	Belief
Independent variable	Actual	Subjective	class	of religion	in God
Treatment					
Out-of-work	-.240	-.215	-.215	-.220	-.215
	(.074)	(.058)	(.058)	(.058)	(.058)
Attend religious services	.098	.101	.105	.082	.100
	(.049)	(.038)	(.038)	(.041)	(.039)
Married	.181	.172	.176	.169	.172
	(.052)	(.039)	(.039)	(.039)	(.039)
Income					
Family income (logged)	.041	.030	.025	.029	.031
	(.024)	(.017)	(.017)	(.017)	(.017)
Income change: actual	.020	—	—	—	—
	(.020)				
Income change: subjective	—	.104	.101	.104	.105
		(.013)	(.013)	(.013)	(.013)
Subjective social class	—	—	.029	—	-.050
			(.018)		
Religion					
Importance	—	—	—	.012	—
				(.012)	
Belief in God	—	—	—	—	.007
					(.009)
Random effects					
Between-person (σ_ϵ)	.381	.356	.357	.354	.355
	(.016)	(.012)	(.012)	(.012)	(.012)
Within-person (σ_ν)	.446	.462	.460	.463	.463
	(.010)	(.007)	(.007)	(.007)	(.007)
Observations	2,132	3,387	3,377	3,380	3,377
Persons	1,149	1,237	1,237	1,235	1,237

Notes: Standard errors in parentheses. All models control for selection, time, race, gender, education, and age.
Source: Authors' calculations from General Social Survey panel, 2006-2010.

Appendix Figures

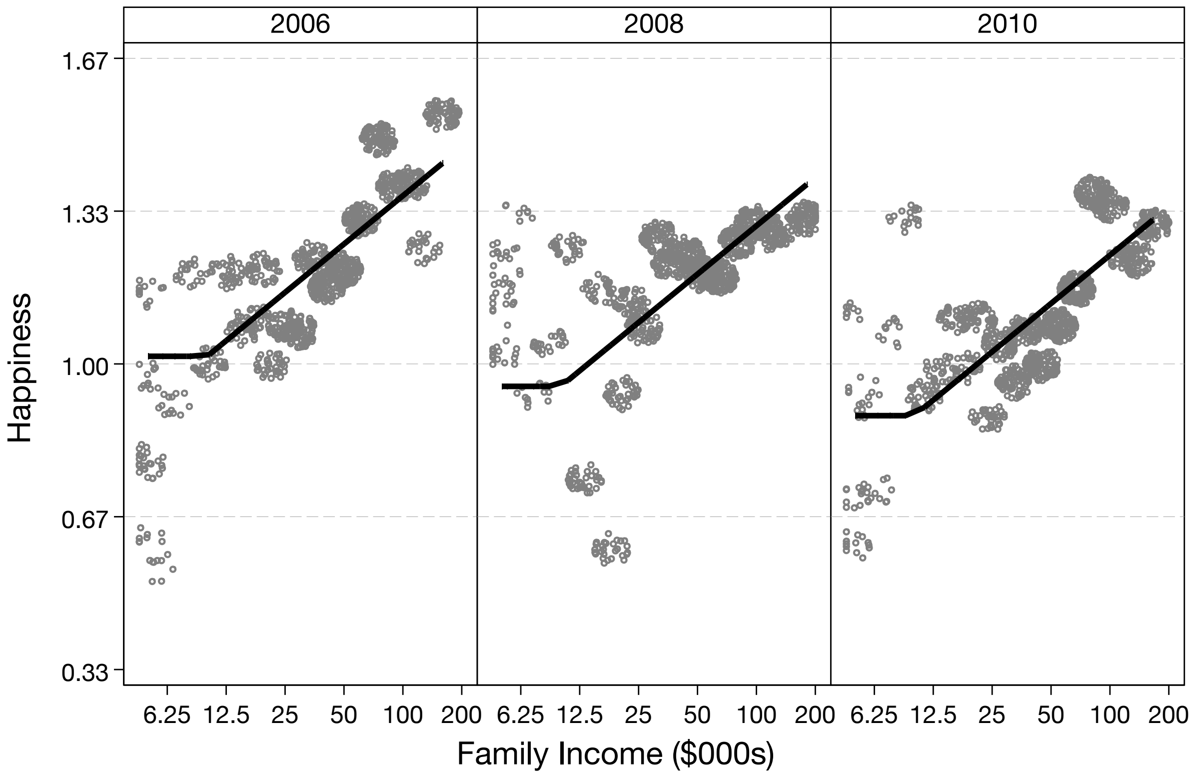

Appendix Figure 1. Happiness score by family income

Notes: Dots show cases jittered around the mean for each income category; the lines show expected values from a regression of happiness score on log income truncated at $10,000.
Source: General Social Surveys, 2006-2010, cross-sectional cases.

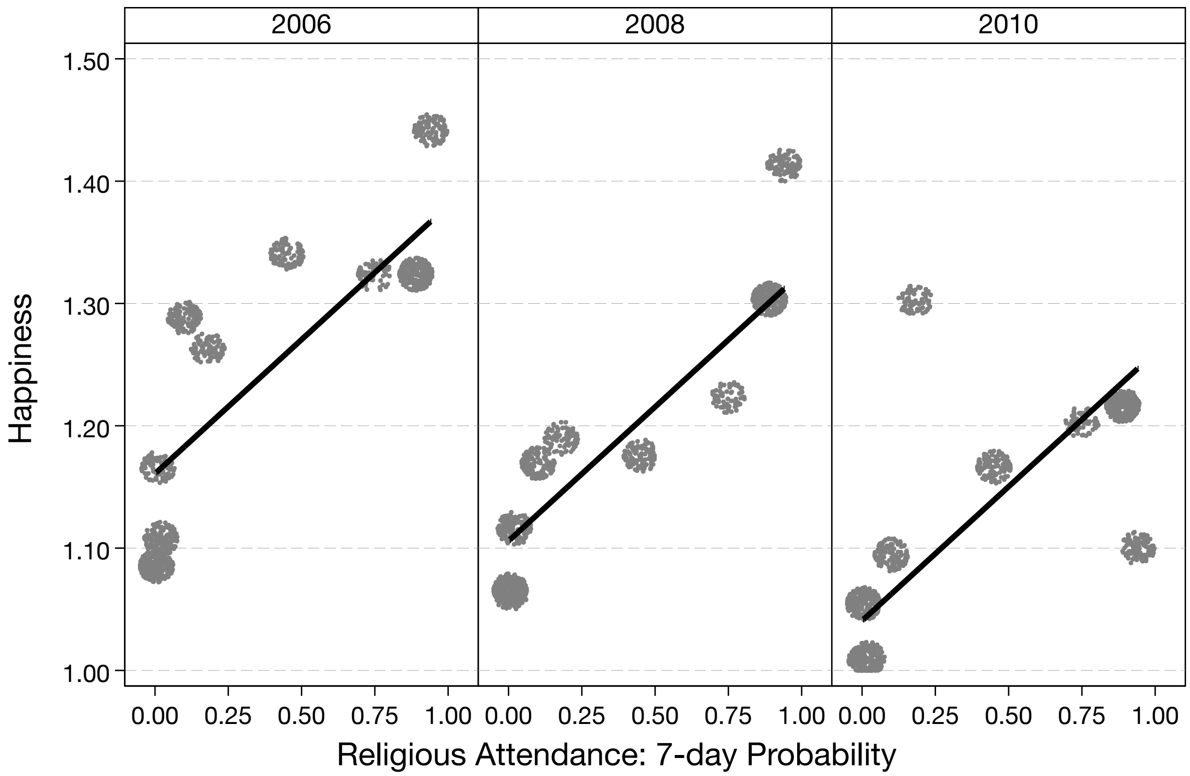

Appendix Figure 2. Happiness score by attendance at religious services

Notes: Dots show cases jittered around the mean for each attendance category; the lines show expected values from a regression of happiness score on the probability of attending religious services in the last seven days.
Source: General Social Surveys, 2006-2010, cross-sectional cases.

Acknowledgments

We acknowledge financial support from the National Science Foundation (SES-0824618), the Institute for the Study of Societal Issues, the Berkeley Population Center (NICHD R21 HD056581), and the Russell Sage Foundation (through the Stanford Center on Poverty and Inequality). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Science Foundation, the National Institute of Child Health and Human Development, or the National Institutes of Health.