Happiness, Unhappiness, and Suicide: An Empirical Assessment

Mary C. Daly and Daniel J. Wilson[*]
(Federal Reserve Bank of San Francisco)

October 2008

Abstract
The use of subjective well-being (SWB) data for investigating the nature of individual preferences has increased tremendously in recent years. There has been much debate about the cross-sectional and time series patterns found in these data, particularly with respect to the relationship between SWB and relative status. Part of this debate concerns how well SWB data measure true utility or preferences. In a recent paper, Daly, Wilson, and Johnson (2008) propose using data on suicide as a revealed preference (outcome-based) measure of well-being and find strong evidence that reference-group income negatively affects suicide risk. In this paper, we compare and contrast the empirical patterns of SWB and suicide data. Despite no obvious aggregate relationship between the two series—either time series of cross-sectional—we find a strikingly strong and consistent relationship in the determinants of SWB and suicide in individual-level, multivariate regressions. This latter result cross-validates suicide and SWB micro data as useful and complementary indicators of latent utility.

Keywords: Happiness and unhappiness trends, suicide, utility, relative income.

JEL Codes: I31, D6, H0, J0

---

[*] This paper benefitted from helpful comments from Andrew Clark, Andrew Oswald, Heather Royer, and participants at the 2008 EEA meetings in Milan. We thank Norm Johnson at the U.S. Census Bureau for providing us with the regression estimates reported in this paper involving the National Longitudinal Mortality Study. We also thank Thomas O'Connell and Joyce Kwok for excellent research assistance. Research results and conclusions expressed are those of the authors and do not necessarily indicate concurrence by the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1. Introduction

Over the past decade economists have dramatically increased their interest in and use of subjective well-being data to track aggregate and individual welfare and to examine issues related to preference formation and utility. Empirical research has included comparing subjective assessments of well-being over time (e.g., time series patterns in happiness, unhappiness, and life satisfaction) and analyzing the individual and group-specific correlates of these measures in the cross-section. The work has produced intriguing results—aggregate happiness does not rise monotonically with income (Easterlin 1995), individuals care about their own *and* others' income (Luttmer 2005; Clark and Oswald 1996), and preferences seem to adapt to the environment. Despite widespread interest in these results outside of economics, they have been slow to affect mainstream economic theory and its application.

An important barrier to greater acceptance of the findings from subjective well-being research has been concerns about data quality. Subjective survey questions, by design, elicit information about individuals' feelings, attitudes, opinions, and views. Critics argue that this introduces systematic and non-systematic measurement errors that make it difficult to compare answers to such questions over time or across individuals (see Bertrand and Mullainathan 2001 for a broad critique of subjective survey data). In the subjective well-being surveys, for example, in which respondents are asked to "rate" or "score" their happiness or life satisfaction, errors can arise from language ambiguities (respondents may not all agree on the exact meaning of terms like "happiness" and "life satisfaction"), scale comparability (one person's "very satisfied" may be higher, lower, or equal to another person's "satisfied"), ambiguity regarding the time period over which respondents base their answers, respondent candidness, and the difficulty of drawing cardinal inferences from ordinal survey responses.[1] These problems have prompted some to discount the results from subjective well-being research and have left a larger number unconvinced of the robustness of the findings.

At issue with each of these concerns is whether subjective well-being reports accurately capture the underlying latent variable that is utility. Since we are not able to observe the latent variable, debates about whether it can be accurately measured are hard to resolve. This has led researchers to seek alternative, outcome-based (i.e., nonsubjective) sources of data to verify or reject the findings from subjective surveys. These alternative data sources include laboratory experimental work and naturally occurring phenomena. Experiments have the advantage of controlling the environment so as to reduce measurement error but have the disadvantage of small sample sizes and contrived situations. Outcome-based studies, such as those looking at mortality and status (Miller and Paxson 2006) or the consumption of positional goods (Kuhn, Kooreman, Soetevent, Kapteyn 2008), have the advantage of tracking actual occurrences, rather than attitudes or laboratory responses, but carry the

[1]Bertrand and Mullainathan (2001) argue that even when questions are well-stated and well-understood, respondents may make cognitive errors that bias their answers and limit their usefulness as dependent variables in outcome studies.

disadvantage of embodying many potential determinants unrelated to preferences.

In the vein of outcome-based studies, Daly, Wilson, and Johnson (2008) consider whether suicide data can be used to address questions regarding interdependent preferences. They argue that suicide represents a revealed choice that can be considered a direct measure of well-being. Using individual-level longitudinal data, they find that relative income does matter for suicide risk, confirming results from subjective well-being and experimental work. An obvious concern regarding suicide data, however, is that the preferences of suicide victims, who arguably are at the extreme lower tail of the well-being distribution, may not be representative of the overall population.

In this paper we address some of the uncertainties regarding the usefulness of subjective well-being and suicide data by using each source as a check against the other—a cross-validation exercise.[2] At the root of our inquiry is an interest in knowing whether the two data series capture the latent variable on well-being (utility) that would allow one to infer preferences from their relationships with other variables. Under the principles of cross-validation, failure to find a systematic relationship between these data means we keep searching, since a negative result could be driven by either or both of the series. If, however, we find a strong relationship across these sources, then we have additional support for the notion that the results reflect true preferences of the general population. Running parallel regressions for suicide hazard and reported happiness, we find that the relative risks estimated from each regression are strikingly similar across a host of important variables. These results add credibility to both sources of data and, importantly, support the use of suicide data to study preferences that are representative of the general population.

2. Data

Our analysis is based on information from two main sources. The data on subjective well-being come from the General Social Survey (GSS). The GSS is a survey of American demographics, behaviors, attitudes, and opinions, administered since 1972 by the National Opinion Research Center out of the University of Chicago. It is structured to be a representative sample of the U.S. population; it was administered to about 1,500 respondents per year from 1972 to 1993 (excluding 1979, 1981, and 1992), 3,000 from 1994 through 2004, and about 4,500 respondents in 2006.[3] Our analysis relies on a wide range of variables from the GSS, including demographic and income variables. Our key interest is the question on subjective well-being. This question reads: "Taken all

[2] Koivumaa, et al. (2001) and Oswald (1997) also consider the relationship between SWB and suicide data .The first study examines whether low reported SWB predicts future suicide. The second study. discusses some broad patterns in suicide rates across industrialized countries and informally compares those patterns to patterns in SWB data.

[3] Note that only about one-half of the total respondents were asked about their happiness from 2002 through 2004, and about two-thirds were asked in 2006.

together, how would you say things are these days—would you say that you are very happy, pretty happy, or not too happy?" To underscore the purely ordinal content of the responses and avoid letting any connotations of the answer choices influence how we treat the data or discuss the results, we simply recode the responses "very happy," "pretty happy," and "not too happy" as "high," "medium," and "low," respectively.

For the suicide hazard regressions, we use the National Longitudinal Mortality Study (NLMS). The NLMS consists of files from the Current Population Surveys (CPS) from 1978 to 1998, matched to the National Death Index (NDI), a national database containing the universe of U.S. death certificates. The matching process appends to individual CPS records (1) whether the person has died within the follow-up period (1979 through 1998), (2) date of death (if deceased), and (3) cause of death (if deceased). See Daly, Wilson, and Johnson (2008) for a more detailed description of these data. In the section below, we also use aggregate time series data on U.S. suicide rates from the U.S. Vital Statistics.[4]

3. Aggregate Time-Series

Before turning to the cross-validation exercise, it is useful to look at how subjective well-being and suicide rates have changed over time. Figure 1 plots: (1) *percent highest happiness*, the percent of respondents in the high happiness category, (2) *percent lowest happiness,* the percent in the low category, and (3) *the suicide rate*, the number of suicides per 100,000 population. The data span the period from 1972 to 2006, though there are some gaps—indicated in the figure by dashed lines— in the GSS series for years in which the happiness question was not included in the survey.

Over the past 30 years, the proportions of reported happiness and reported unhappiness have changed very little.[5] On average, about 30% of GSS respondents reported being in the highest of the three happiness categories and about 14% reported being in the lowest category. The suicide rate, on the other hand, has shown a downward trend from it's peak of 13.1 in 1977 to 10.7 in 2006. This decline has been particularly steep since the early 1990s, which may well be due to the introduction at that time of new, more effective antidepressants rather than, or in addition to, any underlying changes in the happiness of the population.

4. Micro, Multivariate Relationships

Although interesting, the aggregate trends in subjective well-being and suicide are not suggestive a simple aggregate relationship between the two series. That said, potential aggregation bias, the lack of multivariate controls, and the

[4] Ideally one might want to look at attempted suicide as well as completed suicide, but to our knowledge there are no data on attempted suicides that include variables necessary to do an analogous study.
[5] The relative stability of the subjective well-being responses underlies the "Easterlin Paradox," the lack of a long-term relationship between income and happiness.

limited numbers of data points may be obscuring or failing to pick out true individual-level relationships. To explore and compare the factors in suicide risk and reported SWB at the individual level, we use data from repeated cross-sections of the GSS and the NLMS. For the GSS, we have data from 1972–2006 (excluding nine years in which the happiness question was not asked) and a total of about 37,000 sample members. For the NLMS, we have data from 1978–1998 and a sample size of about 900,000. Of these 900,000 individuals, about 63,000 died by the end of the follow-up period (Dec. 31, 1998) and roughly 1,300 of these died of suicide. We restrict our samples from both data sets to working-age (18- to 65-year-old) individuals since much of our focus in this analysis is on variables like labor market status and relative income that are likely to be most relevant for the preferences of the working-age population.

Using the maximum set of variables that the GSS and NLMS have in common, we set up parallel regressions on suicide and subjective well-being data. We then compare the results from these regressions as a way of cross-validating the ability of suicide and SWB data to represent the underlying latent variable on well-being. Our starting point is the familiar latent variable model in which U_i is an unobservable index of well-being for an individual *i*. We model U_i as a linear function of a vector of explanatory variables, $\mathbf{X}_i$: $U_i = \mathbf{X}_i \boldsymbol{\beta}$. The probability that U_i is below any individual-level threshold, θ_i —say the threshold below which the individual would report being "not too happy" or the threshold below which the individual would commit suicide—is then:

$$\text{Prob}[U_i < \theta_i] = \text{Prob}[\mathbf{X}_i \boldsymbol{\beta} + \theta_i < 0]. \quad (1)$$

Our objective here is to estimate the vector $\boldsymbol{\beta} = \mathrm{d}U_i / \mathrm{d}\mathbf{X}_i$, the relative risk of U_i falling below the threshold in question, for the thresholds of suicide, reported happiness, and reported unhappiness. We compare the estimates from parallel regressions, using both the SWB and suicide data, arguing that a close relationship between the results supports both the reasonableness of the data sources and the robustness of the findings regarding determinants of utility.

For reported happiness and unhappiness, we estimate equation (1) using an ordered probit model on the GSS data. We perform ordered probit on the full three-value scale of happiness responses rather than separate probits for high (vs. low or medium) or low (vs. high or medium) for the sake of exposition as well as estimation efficiency.[6] To ease comparability with the suicide regressions, we order the dependent variable with low at the top and high at the bottom, so it can be thought of as a measure of unhappiness or the inverse of happiness. For suicide risk, we estimate a Cox Proportional Hazards (PH) model, given the longitudinal nature of the NLMS data. Both the ordered probit and PH regressions include a full set of year dummies.[7]

---

[6] To check the robustness of our results against the assumptions of normality embedded in the ordered probit we also estimated an ordered logit. The findings were quantitatively and qualitatively very similar.
[7] To economize on space, we do not report here the coefficients on the year dummies, but they are available upon request. It is worth noting that these coefficients display similar patterns over time as those of the aggregate time-series data in Figure 1.

To be able to directly compare the results obtained from the GSS and NLMS, we transform the probit and PH coefficients into their corresponding relative risk estimates. The relative risk for a given variable x_i in the vector of independent variables, $\mathbf{X}_i$, is defined as:

$$[\Pr(Event \mid x_j = \overline{x}_j \;\; \forall j \neq i, x_i = 1) / \Pr(Event \mid x_j = \overline{x}_j \;\; \forall j \neq i, x_i = 0)],$$

where $\overline{x}_j$ denotes the sample mean. For the Cox PH models, the relative risks are the hazard ratios (RR=HR=e^{λ}, where λ is the coefficient in the PH model). For the ordered probit model, the relative risks are 1.0 plus the probit marginal effects, evaluated at the sample mean.

The relative risk for a particular characteristic is interpreted as the probability of the event for an individual with that characteristic relative to the probability for an individual in the omitted category. For example, our regressions include dummy variables for whether individuals have more or less than a secondary education, the omitted category being those who have a secondary (but not post-secondary) education. In the regressions, we obtain a relative risk for the group with less than a secondary education of 1.178 in the ordered probit for unhappiness and 1.092 for the PH model of suicide. These relative risks imply that those with less than a secondary education had a 17.8% higher risk of reporting lowest happiness and a 9.2% higher risk of suicide compared with someone with a secondary (but not post-secondary) education. The $\mathbf{X}_\mathbf{i}$ vector in both regressions are computed similarly and include age, race, gender, marital status, urban/rural residence, veteran status, education, employment status, family income, and year fixed effects.

The results are summarized in Figure 2, panels A-C.[8] The full results and p-values are reported in Daly and Wilson (2008) (see Appendix Table A1). Panel A shows the relative risks for the basic demographic variables. The solid bars in the figure are the relative risks estimated from the ordered probit regression on reported unhappiness. The black, hollow boxes show the relative risks of suicide estimated from the Cox PH regressions. In all cases, the omitted category has a relative risk of one. Although there are some differences in the magnitude of the coefficients in some cases, the results show strikingly similar patterns in the effects of age, gender, marital status, urban/rural residence, and veteran status across the two data sources. Race is the one exception, where the results move in opposite directions.

Panel B of Figure 2 shows the relative risks for educational attainment and several labor market status variables. The educational profiles of relative risks are nearly identical, confirming the expected pattern that unhappiness and suicide risk both fall with education. The results for the labor market status variables also exhibit a very close association across suicide and SWB, particularly for the "employed but not working" and "unemployed" categories.

[8] The relative risks shown in Figure 2 are estimated, in general, quite precisely. In the ordered probit regression, all but two variables – "Age: 18-24" and "Veteran" – are statistically significant, most at below the 1% level. In the Cox PH regression, 8 variables are not significant at at least the 10% level: "Age: 25-34," "Age: 35-44," "Age: 45-54," "Family Income: 20K-40K," "Family Income: 40K-60K," "Rural," "Marital Status: Widow(er)," "Educ: Less than HS."

There is less of an association in the "unable to work" group, but this may reflect differences in measurement of this status across data sources.[9]

The final comparison shown in Figure 2 is income. Here we plot the relative risks in terms of the implied income gradient in each data source. Again, the pattern of the results shows a close association in the effects of income on unhappiness and suicide risk. Both reported unhappiness and suicide risk fall with income. Moreover, in both cases, with the exception of the first income category, the effect of additional income declines as income grows—consistent with diminishing marginal utility of income.

The last aspect of our analysis returns to the question discussed in the introduction of whether and how relative income affects suicide risk and unhappiness. We add relative income to the regressions. The relative income variable for the GSS is the respondent's self-assessment of own family income relative to the "typical American family," with possible answers of far above, above, about, below, or far below average. It should be noted that it is not obvious what reference group the "typical American family" represents, whether this is taken to be the national average, local area average, or something else. In the NLMS regressions, we capture relative income by including, in addition to own family income, the average family income for the county of residence. Note this is an abbreviated version of regressions reported in Daly, Wilson, and Johnson (2008), which obtained very similar coefficients on variables in common.[10] The GSS findings indicate that higher perceived relative income reduces the likelihood of reporting low happiness. Similarly, having low relative income—i.e., high county income relative to own income—increases the risk of suicide. Consistent with previous work on relative income, both of these regressions show that relative income is statistically significant, controlling for a wide range of demographic variables and own income.

The micro data results show a strikingly strong association between the results obtained from suicide and SWB data. The similar pattern found in both data sources cross-validates the value of these alternative data sources for assessing determinants of latent well-being in general and supports the findings of diminishing marginal utility and the importance of relative income in particular.

5. Suicide and Happiness

Although suicide is admittedly at the extreme lower tail of the happiness distribution, the results presented here suggest that the same factors that increase suicide risk also shift people down the happiness continuum. This suggests that suicide data may be a useful way to assess the preferences of the general population, not just those in the extreme lower tail of the happiness or well-being distribution.

---

[9] The unable to work category is computed using a single variable in the NLMS but the combination of two variables in the GSS.

[10] Daly, Wilson, and Johnson (2008) performed a wide variety of robustness checks intended to rule out the possibility of spurious correlations. These checks confirmed the interpretation of the relative income findings for suicide risk.

To formalize the relationship between suicide risk and population happiness, Daly, Wilson, and Johnson (2008) developed a theoretical framework based on a standard random utility model. The key elements of this framework are depicted in Figure 3.[11] The figure depicts the lower tail of the distribution of happiness or utility across the population. Individuals differ in their inherent levels of happiness (set points) or suicide thresholds—θ_i in the figure and in equation (1)—and in their observable circumstances, $\mathbf{X}_i$. Under the identifying assumption that the preferences—the $\boldsymbol{\beta}$ vector in the equation and figure—are the same for suicide victims as for others in the lower tail of the happiness distribution, one can estimate $\boldsymbol{\beta}$ by observing how different $\mathbf{X}_i$'s relate to the probability of suicide. Whether these estimates also reflect the preferences of the overall population is not known *a priori*

However, the cross validation exercise discussed in this paper to our mind is a test of whether the $\boldsymbol{\beta}$ for the lower tail of the happiness distribution (estimated from suicide data) matches that for the entire distribution (estimated from the subjective well-being data). The results support the idea that $\boldsymbol{\beta}$ is homogenous across the distribution.

6. Summary and Future Work

There is a strikingly strong relationship between the correlates of suicide risk and subjective well-being in the multivariate micro analysis. The micro results cross-validate the usefulness of SWB and suicide data for individual-level analyses. The results suggest that prior work using micro subjective well-being data to address relative income status questions is robust to concerns about reporting errors. The results also support previous work by Daly, Wilson, and Johnson (2008) that uses data on U.S. suicide victims to estimate preferences for the general population. Going forward, we see the results of this study as supportive of additional and complementary work on preferences using both subjective well-being and suicide data.

References

Bertrand, Marianne, and Sendhil Mullainathan. "Do People Mean What They Say? Implications for Subjective Survey Data." *American Economic Review*, 2001, 91(2), pp. 67-72.

Clark, Andrew E., and Andrew J. Oswald. "Satisfaction and Comparison Income." *Journal of Public Economics*, September 1996, 61(3), pp. 359-81.

Daly, Mary C., Daniel Wilson, and Norman Johnson. "Relative Status and Well-Being: Evidence from U.S. Suicide Deaths." *Federal Reserve Bank of San Francisco Working Paper* 2008-12.

[11] Their formal model is based on a random utility model that allows individuals to differ randomly in their suicide thresholds while responding systematically to changes in measured variables. They argue that the random component, captured in the error term, does not affect the estimated effect of the measured variables on the underlying latent value of utility.

Daly, Mary C. and Daniel Wilson. “Happiness, Unhappiness, and Suicide: An Empirical Assessment.” *Federal Reserve Bank of San Francisco Working Paper* 2008-19.

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

Koivumaa, Honkanen Heli; Risto Honkanen, Heimo Viinamaeki, Kauko Heikkilae, Jaakko Kaprio and Markku Koskenvuo. 2001. “Life Satisfaction and Suicide: A 20Year Follow-up Study,” *Amer. J. Psychiatry* 158:3, pp. 433-39.

Kuhn, P., P. Kooreman, A. Soetevent, and A. Kapteyn. “The Own and Social Effects of an Unexpected Income Shock: Evidence from the Dutch Postcode Lottery.” *Rand* WR-574. 2008.

Luttmer, Erzo F.P. “Neighbors as Negatives: Relative Earnings and Well-Being.” *Quarterly Journal of Economics,* 2005, 102(3), pp. 963-1002.

Miller, Douglas, and Christina Paxson. “Relative Income, Race, and Mortality.” *Journal of Health Economics*, 2006, 25(5), pp. 979-1003.

Oswald, Andrew J. “Happiness and Economic Performance.” *Economic Journal*, 1997, 107(445), pp. 1815-1831.

Figure 1. Percentage of GSS Respondents in Highest and Lowest Happiness Category and the U.S. Suicide Rate (1972 - 2006)

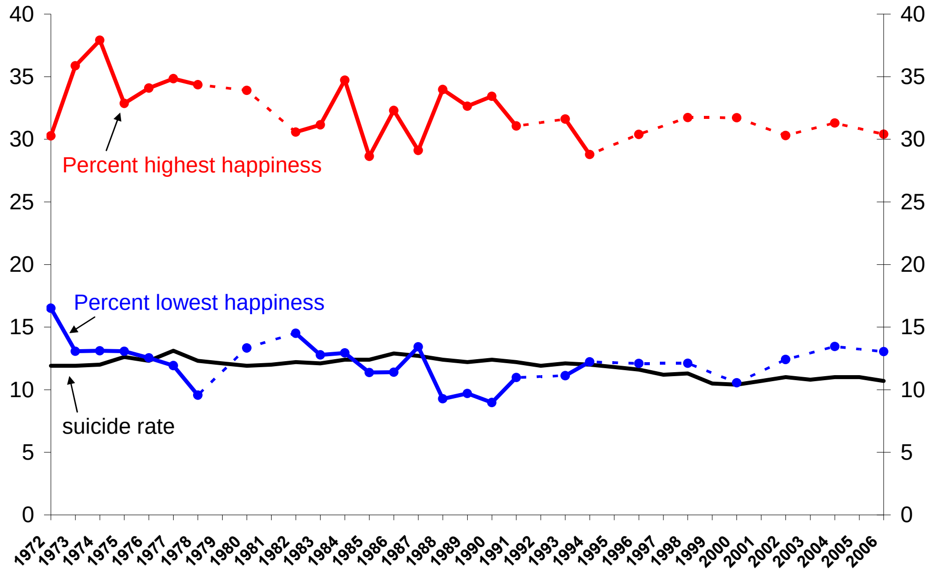

Figure 3. Mapping Suicide to Happiness

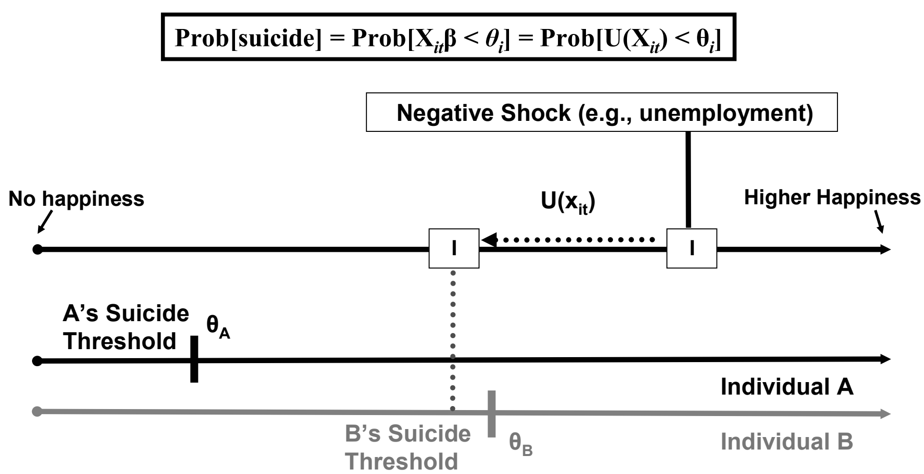

Figure 2: U.S. Micro, Multivariate Patterns in Happiness, Unhappiness, and Suicide

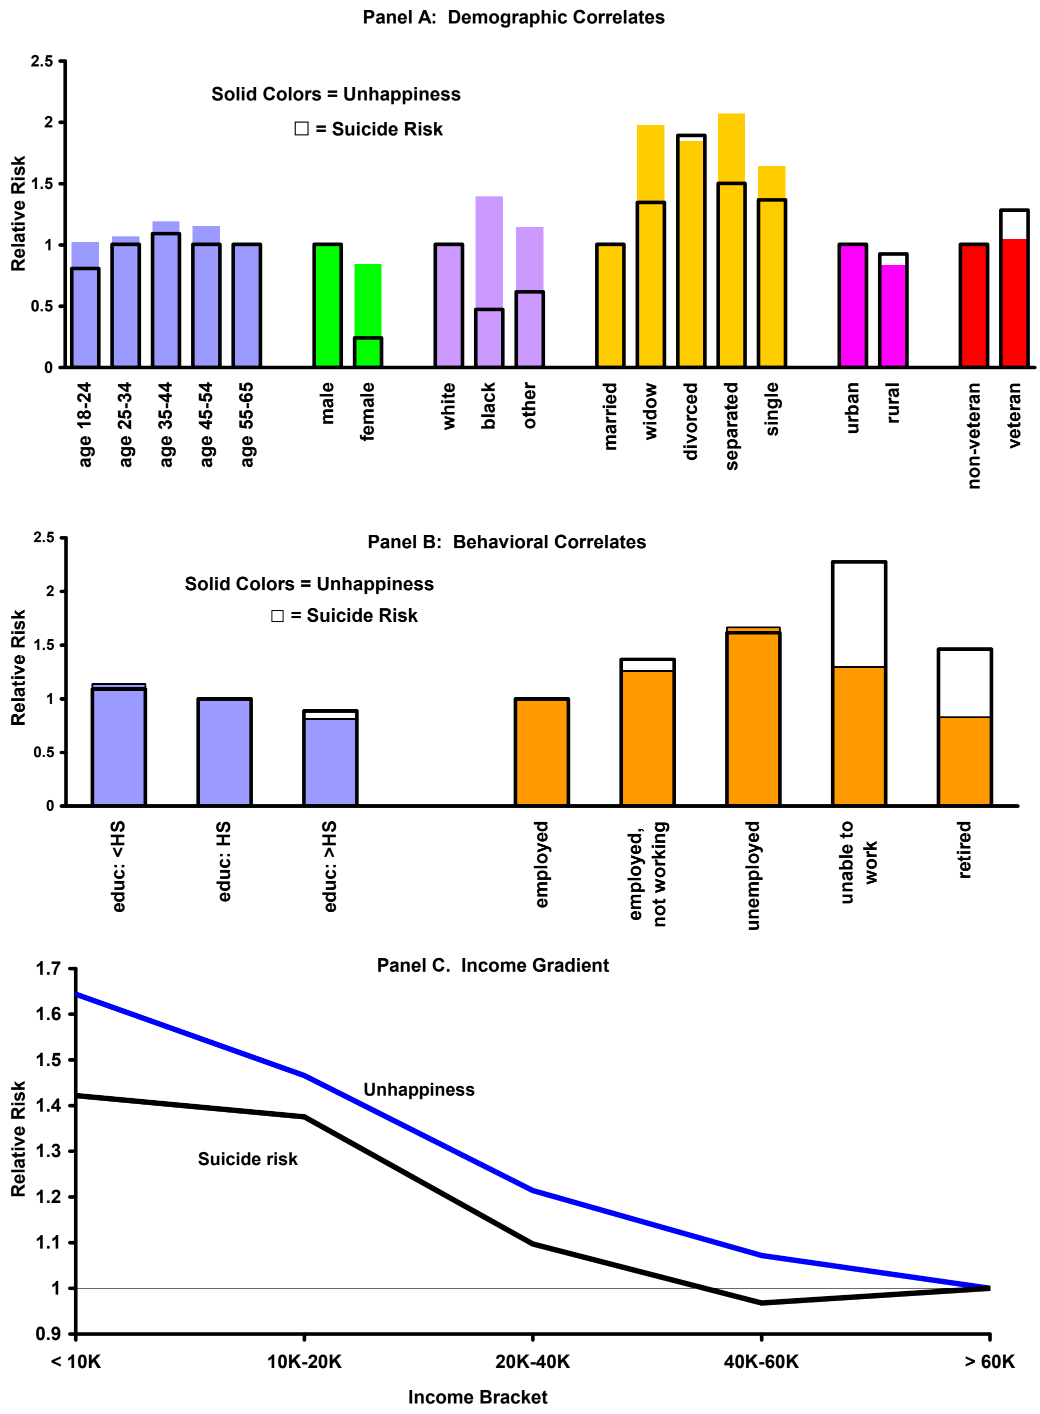

Source: Authors' calculations based on the General Social Survey (subjective well-being data) and the National Longitudinal Mortality File (suicide data).