

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

**Relative Status and Well-Being:
Evidence from U.S. Suicide Deaths**

Mary C. Daly
Federal Reserve Bank of San Francisco

Daniel J. Wilson
Federal Reserve Bank of San Francisco

Norman J. Johnson
U.S. Census Bureau

September 2012

Working Paper 2012-16

<http://www.frbsf.org/publications/economics/papers/2012/wp12-16bk.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

**Relative Status and Well-Being:
Evidence from U.S. Suicide Deaths**

Mary C. Daly^a, Daniel J. Wilson^a, and Norman J. Johnson^b

^a Federal Reserve Bank of San Francisco

^b U.S. Census Bureau

Corresponding author's email: daniel.wilson@sf.frb.org

We gratefully acknowledge research funding provided by the National Institute of Aging. This paper benefited from helpful comments from Marianne Bitler, Sandy Black, David Card, Raj Chetty, Andrew Clark, Dora Costa, John Ham, Doug Miller, Andrew Oswald, Betsey Stevenson, Justin Wolfers, and seminar participants at SOLE 2006, UC Berkeley, Federal Reserve Bank of

New York, and U. of Nevada-Reno. We thank Ann Lucas, Jeremy Gerst, Charles Notzon, Colin Gardiner, and Eric Backlund for excellent research assistance. Research results and conclusions expressed are those of the authors and do not necessarily indicate concurrence by the National Institute on Aging, the Bureau of the Census, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

Relative Status and Well-Being: Evidence from U.S. Suicide Deaths

Abstract:

We assess the importance of interpersonal income comparisons using data on suicide deaths. We examine whether suicide risk is related to others' income, holding own income and other individual and environmental factors fixed. We estimate models of the suicide hazard using two independent data sets: (1) the National Longitudinal Mortality Study and (2) the National Center for Health Statistics' Multiple Cause of Death Files combined with the 5 percent Public Use Micro Sample of the 1990 decennial census. Results from both data sources show that, controlling for own income and individual characteristics, individual suicide risk rises with others' income.

Keywords: Relative income, interpersonal comparisons, interdependent preferences, suicide, happiness, Keeping Up with the Joneses.

JEL Codes: I31, D6, H0, J0

Relative Status and Well-Being: Evidence from U.S. Suicide Deaths

I. Introduction

Despite popular acceptance and growing empirical support, the idea that individuals assess themselves relative to others has been slow to diffuse into mainstream economic theory. A potential reason for the reluctant adoption is that the data used to illustrate the presence and importance of interpersonal comparisons—classroom or laboratory experiments and subjective surveys of happiness or life satisfaction—are themselves the subject of considerable debate. Experiments, by their nature, are contrived and frequently limited to very small samples. Self-reported happiness surveys, while capturing much larger samples, elicit responses that are subjective and may be difficult to compare across individuals and over time. These criticisms of experimental and happiness data have limited the acceptance of research findings on interpersonal comparisons.

In this paper, we propose an alternative source of data, suicide deaths, for identifying the importance of interpersonal comparisons and relative status. Treating suicide as a choice variable regarding current life satisfaction and assessed value of future life, we examine the relationship between suicide risk and own and others' income using data from two independent sources: (1) the National Longitudinal Mortality Study (NLMS) and (2) data from publicly available death certificates combined with the 5 percent Public Use Micro Sample (PUMS) of the 1990 decennial census. Consistent with data from experiments and happiness surveys, we find that local area (county) median income, holding own income constant, is positively and significantly correlated with suicide risk. This result is robust to alternative specifications of the

empirical model and to attempts to reduce the impact of the relative income variable through controls for its potential correlates including geographic variation in the cost of living, access to emergency medical care, and errors in suicide reporting. We argue that additional omitted pathways through which county income might affect suicide risks (e.g., better mental health care services in higher income counties, endogenous mobility of individuals to counties where their relative income is higher, and county income shocks that are correlated with unobserved non-income shocks (as suggested in Luttmer 2005, for example) are more likely to reduce than increase suicide risk, making our estimates an underestimation rather than overestimation of the correlations.

Having established the robustness of our baseline result, we exploit the richness of our data and consider the association between relative income and suicide risk along two additional dimensions. First, we examine whether the relative income association holds for individuals across the income distribution. Our results suggest that suicide risk rises with median county income both for high-income and low-income individuals, although the effect appears to be somewhat larger for the latter. Second, we consider whether relative income comparisons are limited to individuals' local geographic area, defined by county. The results indicate that age, in addition to local area, is a particularly relevant factor. In contrast, the broader geography of state does not appear to be a relevant comparison group.

We interpret our findings as consistent with the idea that relative income matters for measured happiness (unhappiness). Although our analysis is not able to rule out the possibility that omitted variables are driving the association we find, the robustness of the results and the fact that it aligns with previous studies of relative income using experimental data and self-

reported happiness, lead us to conclude that suicide data are a reasonable source of information for studies of interpersonal comparisons.

The remainder of the paper is organized as follows. In Section 2 we review the empirical work on relative income and utility and discuss how information on suicide fits into and expands the literature. We lay out our theoretical motivation and describe our empirical strategy in Section 3. The data sets we construct and use are described in Section 4. In Section 5, we present our main results and assess their robustness. A summary of our findings and the path for future work are laid out in Section 6.

2. Previous Research

Following early recognition of the importance of relative comparisons by Adam Smith, several economists have composed fuller treatments of the issue, including Veblen (1899), Duesenberry (1949), Easterlin (1974), Abel (1990), Galí (1994), Kahneman and Tversky (1996), Frank (2000), Becker and Rayo (2007), and others. These models of interdependent preferences generally posit that individuals care about both their own socioeconomic status (generally defined by income, consumption, or wealth) and that of others. A growing empirical literature on the subject has found evidence consistent with this view. Empirical investigations generally can be grouped into two types. The first set consists of controlled experiments designed to elicit participants' reactions to imposed hierarchies. In these experiments, performed on human and primate subjects, researchers have looked for the subjects' negative reactions to the extent of a hierarchy, i.e., "inequality aversion," and for reactions to subjects' relative placement within a

hierarchy, i.e., “interdependent preferences” (Engelmann and Strobel 2004; Brosnan and de Waal 2003; Alpizar, Carlsson, and Johansson-Stenman 2005). Although such experiments consistently find that inequality and relative income matter, the relatively small sample sizes and artificial environments of these experiments make their results difficult to generalize. Moreover, their contrived nature frequently makes it difficult to distinguish inequality aversion from relative income concerns.

A second vein of the literature on interpersonal income comparisons comes from research on responses to questions from subjective well-being (happiness and/or life satisfaction) surveys. A number of researchers have used the responses from these surveys to study the extent to which self-reported happiness or satisfaction is correlated with relative position, holding other factors such as own income constant.¹ For example, Clark and Oswald (1996) use data on 5,000 British workers to investigate whether worker satisfaction rates are inversely related to relative wages. A similar examination is done in Brown, et al. (2008), focusing on relative rankings of workers’ wages rather than the relative wage ratio. Both studies find evidence that relative income matters to self-reported satisfaction. Along the same lines, several papers have expanded the potential reference group to which individuals are compared by combining individual data on happiness and income with variables on local, regional, and national income (Helliwell 2003; Luttmer 2005; Tomes 1986; and Ferrer-i-Carbonell 2005). In general, these papers have found empirical support for the interpersonal income comparisons hypothesis.

Still, serious concerns have been raised about the quality of data on self-reported happiness (see, e.g., Brekke 1997, Osmani 1993, and Wilkinson 2007; see Bertrand and Mullainathan 2001 for a broader critique of subjective survey data). Such concerns include

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. In addition, Diamond (2008) argues that happiness data may be inappropriate for answering the relative income question in particular since the question itself could be a relative one.²

Therefore, although the results from subjective surveys and experimental studies seem to confirm a role for theories of interdependent preferences, concerns about how representative the underlying data are have hindered broader acceptance of the results. The suggestive findings coupled with concerns about experiments and self-reported measures of happiness suggest that additional methods of addressing the role of relative income are needed.

3. Suicide Data as an Alternative

We propose that suicide data provide an alternative measure of happiness (unhappiness) with several advantages over experiments and happiness surveys.³ First, suicide can be thought of as a revealed choice made by individuals who have examined the value of continuing to live versus not.⁴ In studies of consumer choice, using observed choices to infer preferences has long been considered preferable to relying on individual self-reports of preferences. Second, suicide data are comparably measured across individuals and regions and over time. Third, in the United States at least, data on suicides are publicly available and complete, covering the universe of

reported suicides by year.⁵

There is also a long history in sociology and economics of relating suicide patterns to more universal social phenomena. The most complete example of such efforts is Durkheim's detailed empirical study of suicide. Durkheim's (1951) Suicide: A Study in Sociology, originally published in 1897, was a careful attempt to analyze the societal influences that affect suicidal behavior and unhappiness more generally. More modern treatments in economics have also treated suicide as a potentially social phenomenon, affected by both societal and individual factors. Examples of this work include Hamermesh and Soss (1974), who develop an economic theory of suicide and, using cross-country and cross-state data, find that suicide risk is significantly related to unemployment and decreases in permanent income. More recently, Ruhm (2000) considers suicide as one of several causes of death and finds that, unlike other negative health outcomes that decline during times of recession, suicide risk is either increased or unaffected. In other work, Helliwell (2007) investigates the empirical association between subjective well-being and suicide rates using cross-country panel data and finds a strong negative relationship.⁶ In a related survey article on happiness and economic factors, Oswald (1997) notes that many variables positively (negatively) associated with reported happiness are negatively (positively) associated with suicide risk.⁷ To our knowledge, though, we are the first to use information on suicide risk to study the existence and nature of interpersonal comparisons.

Despite the above-mentioned advantages of using suicide data to address questions on individual well-being and utility, and the demonstration of its usefulness from prior studies, there still are a couple of potential concerns with using such data. First, it is possible that suicide decisions are largely idiosyncratic and not systematically related to the variables that affect

happiness or life satisfaction. While this concern cannot be eliminated *a priori*, it is testable: if it is binding then we should find no correlation between relative income (or other variables) and suicide risk.

Second, and more importantly, suicide victims presumably are at the extreme tail of the distribution of life satisfaction over the population, and their preferences may not reflect the preferences of the non-suicide population. Conceptually, we posit that suicide relates to population happiness as described in Figure 1, which is reproduced from Daly and Wilson (2009). The figure shows the happiness continuums for two individuals, $i = A, B$, as well as their suicide thresholds θ_i . As the schematic illustrates, our maintained hypothesis is that factors (the vector X_i) affecting utility (U_i) have the same marginal effects (the vector β) on suicide risk as they do on happiness, but that thresholds (θ_i) for suicide differ across individuals. That is, suicide victims and the general population have the same β 's but different θ 's. Given this assumption, uncovering the marginal effects of variables on suicide then informs us about how these variables affect happiness for the rest of the population. While the vast majority of the population never commits suicide, this theoretical construct suggests that factors that affect an individual's suicide risk also affect his or her overall happiness/unhappiness. In particular, we can use data on suicide deaths to test hypotheses related to interpersonal comparisons.

Admittedly, without empirical examination one cannot know whether individuals who commit suicide differ only in θ or in both θ and β . To move toward this goal, Daly and Wilson (2009) conduct a cross-validation exercise using subjective well-being data and suicide and find evidence supporting the idea that β 's are the same between those who commit suicide and those who do not. Specifically, they find that the relative risks of suicide along a host of dimensions

(such as education, income, age, marital status, and employment status) closely match the relative risks of reported unhappiness. Based on these results we proceed as if the description in Figure 1 is reasonable and estimate an empirical model implicitly embedding these assumptions.⁸

Before turning to the results, in the next section we describe the data and report basic descriptive statistics regarding suicide and its correlations with demographic, economic, and geographic variables.

4. Data, Sample, and Descriptive Statistics

4.1 Data Sources

To analyze the relationship between relative income and suicide, we draw on two different individual level data sets. Our primary data source is based on the National Longitudinal Mortality Study (NLMS) augmented with data on county and state income from the U.S. Census Bureau. The NLMS data come from a confidential, restricted-use database developed and maintained by the U.S. Census Bureau to facilitate research on the effects of demographic and socioeconomic factors on mortality (see U.S. Bureau of the Census 2005).⁹ These data have been used extensively by epidemiologists and public health experts to study, for example, cancer and heart disease, though they have been used less frequently by economists. The NLMS consists of a set of cohort files, primarily from Current Population Surveys (CPS), matched to the National Death Index (NDI), a national database containing the universe of U.S. death certificates since 1979. The cohort files included in our analysis – those with sufficient information on income – are March CPS files from 1979 to 1998, plus CPS files for February

1978, April 1980, August 1980, and December 1980. The matching process appends to individual CPS records (1) whether the person has died within the follow-up period, (2) date of death (if deceased), and (3) cause of death (if deceased).¹⁰

Our second source of data, which we use as a check on the results from the NLMS data, combines the Multiple Cause of Death Files (MCD) for years 1989-1992, with data from the 1990 5 percent Public Use Micro Sample (PUMS). We will refer to this data set as the MCD-PUMS data. The public use MCD files, compiled by the National Center for Health Statistics and available from the Inter-university Consortium for Political and Social Research (ICPSR), for a given year are essentially the data from all death certificates recorded in the United States in that year (see U.S. Department of Health and Human Services 1992). For the years 1989-1992, we extract the records where suicide is the cause of death (i.e., International Classification of Death, Rev. 9 (ICD9) codes E950-E959) and combine them with the individual records from the PUMS 5 percent sample of the 1990 decennial census (Ruggles et al. 2004), which we treat as non-suicide observations. We extract suicides for 1989-1992, rather than just 1990, to maximize the number of suicide observations, given that suicide is a relatively infrequent event. For sparsely populated counties, the PUMS do not identify the county but instead identifies a “public use micro area,” or PUMA, that is an aggregate of counties (or, in some cases, parts of counties). The final merged data set has over 500 geographic areas (counties or aggregates of counties) spanning the entire U.S., though about 100 of those are dropped from our regressions because they have no reported suicides (in the nonhispanic 20 – 64 year old age range).

Both data sources have advantages and disadvantages. The NLMS data have a much smaller sample of suicide and non-suicide records from which to draw inferences but contain

actual reported income and have no limitations on geographic coverage. In contrast, the MCD-PUMS data have the advantage of containing a very large number of observations from suicide victims (as well as from the general population), but do not include household income and do not identify county of residence for individuals from sparsely populated counties (i.e., those with populations less than 100,000). We consider the NLMS data the preferred and main source for our examination, but use the MCD-PUMS to check key results and perform robustness checks not available in the NLMS data. For one such robustness check (checking whether results are robust to including county fixed effects), we additionally construct a balanced panel data set by combining 1990 and 2000 MCD records with 1990 and 2000 5% PUMS records. Because the PUMAs in the 1990 PUMS are not defined the same as the 2000 PUMAs, the panel excludes geographic areas that do not uniquely identify a county. (That is, it excludes PUMAs that are aggregates of multiple small counties.)

4.2 Sample and Analysis Variables

To correspond to previous research on interpersonal comparisons we restrict our analysis to working-age adults (20-64), for whom relative income concerns are likely to be most relevant. Although we make every attempt to match the sample and variables across data sets, such precision is not always possible. Below we describe the exact selections for each data source and note where differences emerge.

In the NLMS data, we restrict the sample to non-Hispanic working-age adults. Following standard practice in NLMS research, we exclude Hispanics because of definitional changes in the Hispanic status variable over time and concern that a nontrivial share of Hispanic CPS

respondents may have moved out of the United States prior to the end of the follow-up period, in which case their deaths would not be observed.¹¹ However, for completeness, we also run results without this exclusion and report them in the results section. The final data set, after excluding a relatively small number of records with missing values for key variables, contains 957,939 individual records, including 74,929 non-suicide deaths and 1,544 suicide deaths within the follow-up period (the remainder were still alive as of December 31, 1998).

We merge onto the NLMS data a number of geographic aggregates, most notably mean family income by county-year. The construction of these geographic aggregate variables is described in Appendix A.

The variables jointly available in the MCD and the PUMS are age, race, sex, county of residence (for counties with population above 100,000), marital status, education, and Hispanic status. Income, on the other hand, is not recorded on death certificates. We therefore estimate income by matching suicide records in the MCD to individuals or groups of individuals in the PUMS data, where income is available. The matching procedure works as follows: (1) for each suicide record, find all matching observations in the PUMS, matching on county, age, race, sex, Hispanic status, education, and marital status; (2) calculate average family income for this matching cell; and (3) assign this average income to the suicide observation. This procedure provides a reasonably accurate estimate of income: over the 7,202,093 working-age observations in PUMS, county, age, race, sex, Hispanic status, education, and marital status jointly explain 24 percent of the individual level variation in family income.¹² A variance decomposition (not shown) reveals that county, education, and marital status (in decreasing importance) have the greatest explanatory power, together accounting for 16 percent of the variation.

With this matching procedure, we are able to estimate family income for 57 percent of U.S. working-age suicide records from 1989-1992 (and 76% of suicide records among counties with population over 100,000), totaling 50,328 suicides.¹³ We use the same matching procedure to generate an analogous predicted income variable for the non-suicide records; this is the “own income” variable used in our regression analyses. (Because the PUMS contains few missing values for the income-estimation variables, the match rate for non-suicide records from identifiable counties is near 100%.) The final data set has 4,360,747 observations.

4.3 Descriptive Statistics

National statistics show that the U.S. suicide rate has been relatively constant since 1950, averaging about 12 per 100,000 persons (see WHO 2005).¹⁴ Table 1 reports suicide risk overall and by our model variables for the NLMS and MCD-PUMS samples. Recall that both samples exclude Hispanics and cover only working-age adults. The overall suicide rates in the NLMS and MCD-PUMS are quite similar to each other, at approximately 13 per 100,000, and are comparable to the national statistics. Furthermore, national data indicate considerable variation in suicide risk by gender, age, and race. These patterns are mirrored in the NLMS and MCD-PUMS samples. For example, suicide rates are far higher for males than for females and higher for whites than for other races. Suicide rates decline slightly with age according to the MCD-PUMS while having no clear age trend in the NLMS sample, which may simply be due to the relatively small sample size of the NLMS. In both samples, married individuals have a lower suicide rate on average relative to those who are single/never married or divorced/separated. Suicide rates generally fall, though not monotonically, with educational attainment. Although

rudimentary, these categorical suicide rates suggest that the two data sources used in our analysis produce patterns consistent with the stylized facts regarding suicide reported in the epidemiology/public health, psychology, and sociology literatures.

The key variables in our analysis are own and reference group income. To assess the extent to which preferences of the general population can be inferred from the revealed preferences of suicide victims, it is helpful to first compare these two populations along the key dimension of income. Figures 2 and 3 plot the distribution of predicted family income for working-age suicide victims in our two samples against the income distribution for the general U.S. working-age population. Figure 2 shows the distributions of reported family income (adjusted to 1990 dollars) for the total sample and for the subset of those who eventually commit suicide, according to the NLMS data. Note that the NLMS data are survey reports reflecting income at the time the individual was surveyed rather than income at the time the suicide was committed. The income distribution of suicide victims is slightly left of that for the general population. That said, the bulk of the suicide population has income in the middle range of the distribution. We take this as supporting evidence for the notion that suicide victims are broadly representative of the general population, at least in terms of income (though the distribution for suicide victims is somewhat more skewed). This will aid us when we offer an interpretation for our later findings.

Figure 3 reports income figures for the MCD-PUMS sample; the figure shows the distribution of estimated family income (estimated as described in Section 4 above) of suicide victims compared to estimated family income of the general population.¹⁵ The distributions suggest that the modal suicide victim sits slightly to the left of the modal member of the general

population, but overall the two distributions are quite similar. Importantly, there is little difference in the lower tail of the income distribution and overall the shapes for the two populations are roughly similar.¹⁶ The fact that the MCD-PUMS data show a pattern similar to the NLMS data suggests that our estimated income data in the MCD-PUMS data set are reasonably accurate.

Turning to county income, suicide risk has a strong negative correlation with county income. One can see this in Figure 4 which shows a scatterplot of county suicide rates (from the MCD) and county income per family in 1990 across the 3,150 counties in the data. Each circle in the plot represents a single county and the size of the circle is proportional to the county's population. The unweighted correlation is -0.07 and the population-weighted correlation is -0.29 ; both are significant at well below the 1% level. Note we also have confirmed that this negative (unconditional) correlation between suicide risk and county income is present in the NLMS sample with a simple proportional hazards model of suicide risk regressed on county income alone (results available upon request). Thus, it is clear that the positive effect of county income on suicide risk that we find later in our multivariate results is not what one would expect *a priori*.

Descriptive statistics for other model variables are reported in Tables B1 (NLMS) and B2 (MCD-PUMS) of Appendix B. Again, the key variables in our analysis are of similar magnitudes and have similar patterns in both data sets.

5. Empirical Model Specification and Results

5.1 Empirical Strategy

Based on the intuition summarized in Figure 1 and the data described above, we estimate a reduced form model of suicide risk and own and others' income that closely matches those estimated using self-reported data on happiness. For the NLMS data, we estimate Cox proportional hazards models of suicide risk—i.e., the hazard rate of suicide in a given period—as an exponential function of own income, reference-group income, and a set of controls. We use time-since-interview as the duration variable in our models.¹⁷ The estimated proportional hazards model is the suicide hazard (probability of suicide at time t given it has not already occurred) over the interval from 0 to T , where T is the maximum duration in the sample, conditional on individual covariates recorded at period 0.¹⁸ The structure of the NLMS means that the vast majority of observations (individuals) are censored; the proportional hazards model accounts for this.¹⁹

Another approach would be to estimate a logit regression. In fact, the hazard regression is similar to a cross-sectional logit regression but with the advantages of accounting for censoring and allowing the effects of the explanatory variables to decay over time. This decay allows for the fact that the values of an individual's explanatory variables at the interview date will be imperfect measures of the values of those variables at time t , the potential date of death, due to the passage of time. For instance, the discrepancy between a CPS respondent's income at the time of the CPS and her income at the time of her death is likely increasing (in absolute value) in the duration between the two dates. A logit analysis, on the other hand, would treat as equivalent a suicide soon after the measurement of income (and other variables) and one long after the measure of income. As a robustness check, we also estimate logit regressions where the

dependent variable is 1 if the individual ever commits suicide and 0 otherwise. As we show later in the paper, the results are qualitatively similar to those from the Cox proportional hazards model.²⁰

For the MCD-PUMS data we estimate logit models. Because the Cox proportional hazards and logit both have exponential functional forms, it is straightforward to compare the results between the two types of regressions.

In all regressions, standard errors are based on a variance-covariance matrix that is robust to heteroskedasticity and clustering within state. In the NLMS regressions we include time (survey year) fixed effects to capture any macro/aggregate factors that might affect suicide and be correlated with own or others' income. The inclusion of time fixed effects also allows us to interpret the results as evidence of a cross-sectional correlation between suicide risk and the measured variables.²¹ Since the MCD-PUMS data is a single cross-sectional data source, time effects are not possible.²²

5.2 NLMS Baseline Regression Results

The estimated coefficients and standard errors for our baseline NLMS models are reported in Columns 1 – 4 of Table 2. Columns 1 – 3 are identical but for the income variables that are included. Column 1 has (log) own family income measured as a single continuous variable. Column 2 replaces this variable with income bracket indicator variables to allow for non-linear income effects. Column 3 adds (log) county income per family. Column 4 adds county population shares by age category and race, to control for any correlation between county income and county demographics. Before turning our attention to the estimated effects of

income variables, we briefly discuss the relationship between suicide and our control variables; the results are reported in panel C of the table. These results are similar across the four columns. Consistent with the raw categorical suicide rates in Table 1, being female or nonwhite lowers suicide risk, while being divorced or widowed, separated, or never married raises suicide risk (relative to being married). Veterans are found to be more likely to commit suicide than nonveterans. There is little evidence of a conditional age profile to suicide risk, though the point estimates suggest perhaps a weak inverted-U age profile. Controlling for these other factors as well as income, educational attainment lowers suicide risk.

Consistent with previous findings on suicide risk and labor market status, we find that being unemployed or out of the labor force, for any reason, raises suicide risk relative to being employed.²³ Specifically, those who are unable to work have the highest suicide risk, followed by the unemployed, retired persons, and those who are employed but not currently working (e.g., persons on furlough). In terms of magnitude, the estimated coefficient on unemployment of 0.541 (from Column 1) implies a hazard ratio of 1.72 ($e^{0.541}$), meaning that holding other factors constant, suicide risk for an unemployed person is 72% higher than that of a person who is employed and working. The patterns among the control variables vary little across different specifications of the model. Thus, for the remainder of the paper, we confine our discussion to the relationship between suicide risk and own and others' income, the key variables in our analysis.

The results of the income variables are reported in the upper portion of the table, beginning in Panel A. The first column shows the importance of own family income, measured in logs. Own income is statistically significant and negative, implying that higher own income

lowers suicide risk. The coefficient on log own income of ~ 0.087 suggests that a 10% higher income is associated with 0.87% lower suicide risk. In column 2, we allow income to affect suicide risk non-linearly, and find evidence of important differences across the income distribution. In particular, individuals with family incomes below \$20,000 in 1990 dollars (which, by way of reference, is equivalent to about \$34,000 in 2010 dollars, based on the CPI-U) are significantly more likely to commit suicide than those with incomes above \$60,000 (\$102,000 in 2010 dollars). In contrast, suicide risk for those with incomes between \$20,000 and \$60,000 is not statistically significantly different than that of individuals with family income above \$60,000.²⁴ The point estimates of the coefficients on the categorical income variables imply hazards ratios of 1.50, 1.43, 1.10, and 1.02, respectively, for income categories \$0 - \$10,000, \$10,000 - \$20,000, \$20,000 - \$40,000, and \$40,000 - \$60,000. The hazard ratio of 1.50, for instance, means that an individual with family income less than \$10,000 (in 1990 dollars) is 50 percent more likely to commit suicide than an individual with income above \$60,000 (the omitted income category). The hazard ratios decline monotonically, but at a diminishing rate, toward 1.0 as income approaches the omitted top category (for which the hazard ratio is implicitly 1.0). This pattern is consistent with the standard assumption of diminishing marginal utility of income/consumption, and also qualitatively consistent with recent findings in the subjective well-being (SWB) literature such as Kahneman and Deaton's (2010) result that daily mood increases with income up to \$75,000 (2008-2009 dollars) and then is flat. Given this evidence of a non-linear income gradient for suicide risk, we use this model going forward.

Column 3 of the table displays results of adding reference group income. Following

previous work on interpersonal income comparisons, we use county of residence to define reference group income. The results show that county income has a positive effect on suicide risk controlling for own income. Our estimated coefficient of 0.453 on log county income implies that, holding own income constant, a 10% higher county income is associated with about a 4.5% higher suicide hazard relative to the baseline hazard (conditional mean hazard).²⁵

Though the results on own income point to a non-linear income gradient, this specification of own income does not allow for a natural comparison of the magnitudes of the county income effect and the own income effect. In a separate regression not shown here (results available upon request) containing the single log own income variable and log county income, we find the county income coefficient to be larger than that of own income (in absolute value).²⁶ However, this comparison is probably misleading as the coefficient on own income is likely biased downward due to measurement error in income and the inclusion of other variables which are highly correlated with income.²⁷ Later in the paper, we confirm this finding using the MCD-PUMS sample which is less prone to the potential bias in income since the own income measures is by construction a fitted value from a first-stage estimation using PUMS data.

In the next column, we add controls for county demographic composition that might be correlated with both increased suicide risk and higher county income. The findings are qualitatively equivalent, although including these controls increases the magnitude of the coefficient on county income to 0.596.²⁸ The final two columns of Table 2 check the robustness of the results to the exclusion of Hispanics from the sample and suicide misclassification.²⁹ In neither case does the change affect our baseline findings. Including Hispanics in the sample reduces the magnitude of the coefficients on all income variables slightly but does not change the

pattern or statistical significance of these variables. Broadening the measure of suicide to include deaths from “injuries of undetermined cause” (ICD9 codes E980-E989), which some scholars have suggested may capture suicides that were not correctly classified as such, has no notable impact on the coefficients of own or county income.³⁰

The findings in Table 2 imply that controlling for other factors, including own income, individuals living in higher income areas face greater suicide risk than those living in low income areas. This finding is consistent with results of studies using happiness survey data which suggests that a loss of relative position leads to a reduction in individual happiness (see Luttmer 2005). Still, several other explanations or pathways for county income affecting suicide risk are possible; we examine these alternative explanations in the next section.

5.3 Alternative Explanations for NLMS Results

Cost-of-living differences.

One potential alternative explanation for the results reported in Table 2 is that the positive effect of county income on suicide risk is explained by county income simply being a proxy for cost of living, so that, conditional on nominal own income, individuals are made worse off by living in areas with higher costs, especially costs on nontradables such as housing. We control for this alternative explanation in two ways: state fixed effects and controls for county-level house prices. The results of these tests are reported in Table 3; for convenience of comparison, our preferred model from Table 2 is repeated in column 1.

The first test is to add state fixed effects to our baseline regression. The logic is that

regional differences in cost of living, associated with location, tax structures, etc., will be captured at the state level and pulled out in the state fixed effect. To the extent that these cost of living differences are driving our results, the coefficient on county income should fall or become insignificant. The coefficient on county income falls, from 0.596 in the baseline, to 0.390 when state fixed effects are included, but remains statistically significant. The coefficients on own income are qualitatively unchanged.

The second test exploits the fact that the cost of housing is likely the most important component of cost of living differences across areas. Column 3 reports results from adding a county quality-adjusted house price index (described in Appendix A) to our baseline regression. Given the quality adjustment, this index reflects the average cost of land in a county (in a given year) as well as any differences across counties in construction costs. The inclusion of the index drives up the coefficient on county income, while the coefficients on own income remain qualitatively unchanged. The coefficient on the index of housing costs is negative, suggesting that suicide risk is lower in counties with higher housing costs, perhaps because these costs reflect positive area amenities capitalized in local land values.³¹ Based on these checks we conclude that our results are not driven by disutility of higher costs of living.

County Income and Mortality.

Two other potential explanations of the positive correlation between county income and suicide are that it reflects a relationship between county income and the quality of local emergency medical care or a more general relationship between county income and mortality. The results in columns 4 and 5 are designed to address these concerns. Column 4 reports results

from a regression in which heart attack risk (acute myocardial infarction, ICD9 code 410) replaces suicide risk as the dependent variable. The use of heart attack deaths is meant to test whether our results on suicide risk owe to differential quality of or access to emergency room care or paramedical care, rather than to behavioral reactions to relative income. Research has shown that heart attack deaths are strongly correlated with time to treatment (e.g., proximity to emergency rooms). If our results on suicide are due to unequal access to emergency rooms such that attempted suicides more frequently end in death, then we should see the same pattern for heart attack deaths. This is not the case. Indeed, while the mortality hazard from heart attacks falls monotonically with own income, as with suicide, it also falls with county income, contrary to suicide. The final column of the table repeats this analysis using all causes of mortality. Our findings concur with the standard result in the literature (see, e.g., Miller and Paxson 2006 and Gerdtham and Johannesson 2004): mortality falls monotonically with own income and is unaffected by relative income.

Based on these results, we conclude that our finding of a positive effect of local area income on suicide, after controlling for own income, likely reflects a behavioral response to unfavorable interpersonal income comparisons.³² These individual level results are consistent with earlier, semi-aggregate results for suicide risk (Daly and Wilson 2006) and with recent empirical analyses using self-reported, subjective well-being survey data (Luttmer 2005).

5.4 NLMS Extensions

Having established a robust relationship between suicide risk and own and others' income, we now turn to extensions of the baseline specification and the interpersonal income

literature more generally. Table 4 displays the results of these extensions; for convenience, the first column repeats the results from our baseline specification. We first consider whether county of residence is the more relevant reference group than state of residence.³³ The results show that state family income has no significant effect on suicide risk. Next, we ask whether the relative income effect varies over the income distribution. To do so, we interact the categorical income variables with county income. The results are shown in the third column of the table, Panel B. While the small sample size limits the statistical power in this regression, the higher point estimates of the interactions involving the lower income categories are suggestive of a stronger effect for those at the bottom of the income distribution than for those at the top.

Finally, we consider whether the relative income effect differs by gender and race. To do so, we interact log county income with gender dummies and race dummies. These results are shown in the final two columns of the table, respectively. The results suggest that women are more responsive to relative income than men and that whites are more responsive than nonwhites (though these differences are not statistically significant at conventional levels).

5.5 MCD-PUMS Baseline Regression Results

Although the NLMS results provide consistent evidence of the effect of county income on suicide risk, one might still be skeptical of this result given the relatively small number of suicides in the NLMS data. To try to address this issue, we turn to an alternative data source, the MCD-PUMS described earlier. As noted, the MCD-PUMS data combine suicide records from death certificate data with individual records from the PUMS 5% sample of the 1990 decennial census. Using these data, we estimate a set of regressions that are as analogous as possible

(given the data available in the MCD-PUMS) to the NLMS models.

We estimate logit models of the probability of committing suicide as a function of (log) estimated own family income, (log) county income per family, and various controls, including state fixed effects. These models include the same control variables as those in the NLMS regressions except education, labor market status, veteran status, and county characteristics (i.e., demographic population shares and population density). Labor market and veteran status are not recorded on death certificates and hence are unavailable in the MCD-PUMS data set. We omit education to avoid multicollinearity with predicted income, given that, after county (PUMA) of residence, we find education to have the most explanatory power in our income estimation. If we included education, there would be little independent variation with which to identify the coefficient on own income.³⁴ Our strategy in these regressions thus amounts to treating estimated income as a summary statistic for socioeconomic status.

Table 5 gives the baseline results for the MCD-PUMS logit regression. As in the earlier NLMS results, the standard errors shown are robust to heteroskedasticity and clustering within state.³⁵ Column 1 shows results where own family income is measured as a continuous variable. In Column 2, income is measured by a set of income bracket indicators. To enhance comparability with the NLMS results, we defined these brackets using the same cut-off values as those used in the NLMS. Column 3 adds (log) county income per family.³⁶ Column 4 adds county demographic shares to ensure that any county income effects are not driven by a correlation between county income and county demographic composition.³⁷

As in the earlier NLMS regressions, we find in Column 1 that suicide risk falls with own income, at least when income is measured as a continuous variable. Columns 2 and 3 reveal less

evidence of a non-linear income gradient for own income and suicide risk than in the NLMS regressions. It is possible that the estimation of income for the MCD-PUMS data results in less precision across narrowly defined categories at the lower end of the income distribution and thus masks the clear gradient evident in the NLMS results. In addition, whereas there are similar numbers of observations in each of the income brackets in the NLMS, this is not the case in the MCD-PUMS. We have also estimated this MCD-PUMS regression replacing these dollar-level income brackets with income quintile dummies, ensuring equal coverage in each category (results available upon request). The coefficients on these quintile dummies reveal a similar pattern to that found in the NLMS: suicide risk falls with income, but at a diminishing rate.

The key result of Table 5, however, is that the MCD-PUMS data confirm the pattern seen in the NLMS data that county income has an independent effect on suicide risk holding other variables, including own income, constant. Based on the specification underlying Column 3, the estimated coefficient on log county income per family suggest that 10% higher county income per family is associated with 3.2% higher suicide risk. As with the NLMS regressions, the presence of a non-linear income gradient hampers the ability to directly compare the magnitude of the county income effect and that of own income. However, though the linear own income specification may be a misspecification, we have estimated such a specification with county income included. We find that the own income effect is somewhat larger, suggesting that an increase in aggregate income would reduce aggregate suicide risk.³⁸

In Column 4, we show the results of including county level demographic shares as a check on whether the positive county income effect is being driven by correlation between county income and county demographics. As was the case for the NLMS hazards regressions

(see Table 2), the addition of these population shares *increases* the positive effect of county income. Specifically, the coefficient goes from 0.317 to 0.680 (both are statistically significant at the 95% level or above).

Columns 5 and 6 shows the results when we allow the coefficient on county income to differ by gender (Column 5) or by race (Column 6), as we did with the NLMS regressions in Columns 4 and 5 of Table 4. As in the NLMS results, the positive county income effect is larger for females than for males, though the difference here is not statistically significant. In terms of race, other races have the largest county income coefficient followed by whites. This differs from the NLMS results in which only whites had a statistically significant county income effect. The difference is likely due to the small number of observations in the NLMS for non-whites, resulting in very large standard errors for the county income effect for non-whites.

5.6 MCD-PUMS Robustness Check and Extensions

Here we report the results of few robustness checks on the baseline MCD-PUMS results. To assess whether the MCD-PUMS results could be prone to omitted variable bias due to the omission of some of the important control variables that we included in the NLMS regressions but did not have available for the MCD-PUMS regressions, we also estimated a parallel NLMS regression containing the same set of variables that are available in the MCD-PUMS. The side-by-side results of our baseline MCD-PUMS logit regression and the parallel NLMS Cox proportional hazards regression are shown in Table 6. Both regressions contain dummy variables for income brackets (defined using the same cut-offs in terms of 1990 dollars); dummy variables for sex, race, age, and marital status; and state fixed effects. The NLMS regression

also includes year fixed effects; the MCD-PUMS sample is a 1990 cross-section, so the intercept captures any 1990 fixed effect. Both samples exclude Hispanics (for reasons discussion earlier).

Looking at columns 1 and 2 of Table 6 one sees that the coefficient on (log) county income is positive and significant in the regressions from both data sets. The MCD-PUMS coefficient of 0.680 is somewhat higher than that from the parallel NLMS regression (0.345), though it is not outside the range of county income coefficients found in the NLMS regressions shown in Table 3 and is similar to the point estimate we obtain in our preferred NLMS specification of 0.596. Turning to the details, both data sets indicate that suicide risk is significantly lower for females and non-whites. Both indicate that individuals under 35 years old have lower suicide risk than those over 55, but the NLMS points to a monotonic reduction in suicide risk with age, while the MCD-PUMS suggests an inverted-U age profile.³⁹ Both data sets also indicate that suicide risk is lowest for married persons, followed by single or never-married, and then divorced or widowed (which are combined in the NLMS due to data constraints).

Another possible concern with the cross-sectional MCD-PUMS results is that, while they do control for county demographic characteristics and state fixed effects, they cannot control for unobserved county characteristics that possibly could be correlated with county income. Therefore, as mentioned in Section 4, we constructed a pooled 1990 and 2000 MCD-PUMS panel data set covering the subset of geographic areas in the 1990 cross section that are counties (as opposed to aggregates of multiple counties) and hence have the same boundaries in 1990 and 2000. The constructed panel thus consists of only 308 counties, though these counties represent over 75% of the nation's population.⁴⁰ This panel data set allows us to control for county fixed

effects, thus identifying the county income effect on suicide rates from variation in the 1990-2000 change in county income.

The results of controlling for county fixed effects are shown in columns 3 and 4 of Table 6. The regression in column 3 specifies own income non-linearly, as in columns 1 and 2, while that in column 4 includes log own income. The coefficients obtained from the 1990-2000 MCD-PUMS panel with county fixed effects are strikingly similar to those based on the (more geographically comprehensive) 1990 MCD-PUMS cross-section. In particular, the coefficient on county income of 0.65 is quite close to that from the baseline coefficient of 0.68, though the former is much less precisely estimated, perhaps because of the much fewer number of counties (and hence variation in county income) in the panel data set. The lack of statistical significance in column 3 appears to be partly due to how own income is specified. Column 4 shows that when own income is represented simply by log income, the county income is found to be significant at the 10% level. All in all, the panel results suggest that the baseline MCD-PUMS cross-sectional results are robust to controlling for unobserved county characteristics.

The final component of our analysis exploits the greater detail and sample size in the MCD-PUMS to consider the importance of more narrowly-defined reference groups for relative income comparisons. Table 7 reports results from introducing different reference income values computed over various reference subgroups. Column 1 repeats the baseline results from Table 5. Column 2 replaces the log of income per family within the same county with the log of income per family within the same county and the same age group. Column 3 uses log of income per family within the same county and the same race. The results suggest that, while others in one's county or others of the same race in one's county are relevant reference groups, others in the

same age range in one's county may be the most relevant reference group.

6. Additional Considerations and Future Research

Using individual level data on suicide risk, we find compelling evidence in support of the idea that individuals care not only about their own income but also about the income of others in their local area. This finding is obtained using two separate and independent data sets, suggesting that it is not an artifact of the particular sample design of either data set. Importantly, the finding is robust to alternative specifications and we are not able to explain it by geographic variation in suicide misclassification, cost of living, or access to emergency medical care.

It is also worth noting that other plausible stories of potential bias that we cannot test or rule out with our data generally imply a *downward* bias on our key county income variable. For instance, previous research has shown that psychiatric services are positively correlated with county income (Zimmerman and Bell 2006). This positive correlation combined with the possibility that the quality of local mental health care negatively affects suicide hazard implies a possible downward bias on county income's effect on suicide. Another possibility is that individuals are mobile and endogenously select their county of residence in response to their income relative to the county's average (assuming, perhaps unrealistically, that individuals can obtain the same income when they move). This would suggest that suicide outcomes underestimate the true relevance of interpersonal income comparisons because individuals are able avoid the negative utility impact of low relative income by simply moving to a location where they have higher relative income. Another possible story is that county income shocks may be correlated with unobserved non-income county shocks that reduce the general well-being

of county residents and hence increase suicide risk. For instance, a local plant closing might both reduce average household income in the county and lead to other negative county-wide outcomes (reduced local tax revenues and public services, reduced social capital, etc.) that are unobserved and reduce utility of individuals in the county, hence increasing suicide risk.

Luttmer (2005) investigates this possibility in the context of reported happiness by instrumenting for actual county income with county income predicted from national trends and county level occupation and industry composition. He finds very little difference between the OLS and IV results, suggesting such unobserved county shocks are not quantitatively significant. More generally, any story involving classical measurement error in our reference group income measures (relative to the unobserved true reference income) will imply attenuation bias (toward zero).

Finally, regarding the proportional hazards estimations, a common concern in such survival analysis is attenuation bias from unobserved individual heterogeneity. The concern is that individuals with especially negative individual effects (“frailty” in the parlance of survival analysis)—i.e., the θ_i term in our theoretical model—are more likely to exit the sample early via suicide; since there are no observations from these individuals for the remaining years of the sample, they receive less weight than survivors in the estimation, hence underestimating the effects of all variables on exit probability. Again, though, this bias only argues that the true effect of reference group income is in fact larger than what we find.

Our results confirm those obtained in semiaggregate analysis (Daly and Wilson 2006) on group suicide risk and income dispersion and also are broadly consistent with results using happiness surveys. The finding that suicide risk, holding own income constant increases in

reference group income, is found for reference groups ranging from near neighbors, or those who closely resemble the individual, to simple geographical definitions such as county of residence. State appears to be too broad as a measure of reference group. This finding is notable since many previous papers investigating relative income or relative deprivation have been forced to rely on state- or higher-level aggregates as reference groups (e.g., Blanchflower and Oswald 2004; Kennedy, et al. 1996; and Kaplan, et al. 1996).

This paper has focused on static interpersonal income comparisons. Models of this kind are known by various names such as “external habit formation” and “Keeping Up with the Joneses”. Future research using suicide data may consider dynamic models of preferences such as “internal habit formation” or “Catching Up with the Joneses”. The evidence in this paper regarding the usefulness of suicide data for evaluating the nature of the utility function and preferences suggests that such research could indeed be fruitful.

References

- Abel, Andrew B, "Asset Prices under Habit Formation and Catching up with the Joneses," *American Economic Review* 80:2 (1990), 38-42.
- Alesina, Alberto, Rafael Di Tella, and Robert J. MacCulloch, "Inequality and Happiness: Are Europeans and Americans Different?" *Journal of Public Economics* 88:9-10 (2004), 2009-2042.
- Alpizar, Francisco, Fredrik Carlsson, and Olof Johansson-Stenman, "How Much Do We Care About Absolute Versus Relative Income and Consumption?" *Journal of Economic Behavior & Organization* 56:3 (2005), 405-421.
- Becker, Gary S., and Luis Rayo, "Evolutionary Efficiency and Happiness," *Journal of Political Economy* 115:2 (2007), 302-337.
- Bertrand, Marianne, and Sendhil Mullainathan, "Do People Mean What They Say? Implications for Subjective Survey Data," *American Economic Review* 91:2 (2001), 67-72.
- Blakely, T. A., S. C. D. Collings, and J. Atkinson, "Unemployment and Suicide: Evidence for a Causal Association?" *Journal of Epidemiology and Community Health* 57:8 (2003), 594-600.
- Blanchflower, David G., and Andrew J. Oswald, "Well-being over time in Britain and the USA." *Journal of Public Economics* 88:7-8 (2004), 1359-1386.
- Brainerd, Elizabeth, "Economic Reform and Mortality in the Former Soviet Union: A Study of the Suicide Epidemic in the 1990s," *European Economic Review* 45:4-6 (2001), 1007-1019.

- Brekke, Kjell Arne, *Economic Growth and the Environment: On the Measurement of Income and Welfare* (Cheltenham, Brookfield: Edward Elgar, 1997).
- Brosnan, Sarah F., and Frans B. M. de Waal, "Monkeys Reject Unequal Pay," *Nature* 425 (September 2003), 297-299.
- Brown, Gordon D.A., Jonathan Gardner, Andrew Oswald, and Jing Qian, "Does Wage Rank Affect Employees' Well-being?" *Industrial Relations* 47:3 (2008), 355-389.
- Chuang, Hwei Lin, and Wei Chiao Huang, "Economic and Social Correlates of Regional Suicide Rates: A Pooled Cross-Section and Time-Series Analysis," *Journal of Socio Economics* 26:3 (1997), 277-289.
- Clark, Andrew E., and Andrew J. Oswald, "Satisfaction and Comparison Income," *Journal of Public Economics* 61:3 (1996), 359-381.
- Cutler, David, Edward Glaeser, and Karen Norberg, "Explaining the Rise in Youth Suicide," (pp. 219-270), in J. Gruber (Ed.), *Risky Behavior among Youths: An Economic Analysis* (Chicago, IL: University of Chicago Press, 2001).
- Daly, Mary C., and Daniel J. Wilson, "Keeping Up with the Joneses and Staying Ahead of the Smiths: Evidence from Suicide Data," *Federal Reserve Bank of San Francisco Working Paper*, 2006-12 (2006).
- Daly, Mary C., and Daniel J. Wilson, "Happiness, Unhappiness, and Suicide: An Empirical Assessment," *Journal of the European Economic Association* 7:2-3 (2009), 539-549.
- Di Tella, Rafael, Robert J. MacCulloch, and Andrew J. Oswald, "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness," *American Economic Review* 91:1

(2001), 335-341.

Diamond, Peter, "Behavioral Economics," *Journal of Public Economics* 92 (August 2008), 1858-1862.

Duesenberry, James S, *Income, Saving, and the Theory of Consumer Behavior* (Cambridge, MA: Harvard University Press, 1949).

Durkheim, Emile. *Suicide: A Study in Sociology* translated by George Simpson and John A. Spaulding (New York: The Free Press, 1951).

Easterlin, Richard A, "Does Money Buy Happiness?" *Public Interest* XXX (1973), 3-10.

Easterlin, Richard A, "Does Economic Growth Improve the Human Lot? Some Empirical Evidence," in Paul A. David and M. W. Reder (Eds.), *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz* (New York and London: Academic Press, 1974).

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

Engelmann, Dirk, and Martin Strobel, "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments," *American Economic Review* 94:4 (2004), 857-869.

Ferrer-i-Carbonell, Ada, "Income and well-being: an empirical analysis of the comparison income effect," *Journal of Public Economics* 89:5-6 (2005), 997-1019.

Frank, Robert H, *Luxury Fever: Money and Happiness in an Era of Excess* (New York: The Free

- Press, 2000).
- Galí, Jordi, “Keeping Up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices,” *Journal of Money, Credit, and Banking* 26:1(1994), 1-8.
- Gerdtham, Ulf-G., and Magnus Johannesson, “Absolute Income, Relative Income, Income Inequality, and Mortality,” *Journal of Human Resources* 39:1 (2004), 228-247.
- Hamermesh, Daniel S, “The Economics of Black Suicide,” *Southern Economic Journal* 41:2 (1974), 188-199.
- Hamermesh, Daniel S., and Neal M. Soss, "An Economic Theory of Suicide," *Journal of Political Economy* 82:1 (1974), 83-98. (In *Mercurio* as "Una Teoria Economica del Suicidia.")
- Helliwell, John F, "How's Life? Combining Individual and National Variables to Explain Subjective Well-Being," *Economic Modeling* 20:2 (2003), 331-360.
- Helliwell, John F, "Well-Being and Social Capital: Does Suicide Pose a Puzzle?" *Social Indicators Research* 81:3 (2007), 455-496.
- Hendin, Herbert, *Suicide in America* (New York: W.W. Norton & Company, 1995).
- Huang, Wei Chiao, “Religion, Culture, Economic and Sociological Correlates of Suicide Rates: A Cross-National Analysis,” *Applied Economics Letters* 3:12 (1996), 779-782.
- Kahneman, Daniel, and Angus Deaton, “High Income Improves Evaluation of Life But Not Emotional Well-Being,” *Proceeding of the National Academy of Sciences* 107:38 (2010), 16489-16493.

Kahneman, Daniel, and Amos Tversky, "On the Reality of Cognitive Illusions: A Reply to Gigerenzer's Critique," *Psychological Review* 103:3 (1996), 582-591.

Kaplan, George A., Elsie R. Pamuk, John.W. Lynch, Richard D. Cohen, and Jennifer Balfour, "Inequality in Income and Mortality in the United States," *British Medical Journal* 312 (April 1996), 999-1003.

Kennedy, Bruce.P., Ichiro Kawachi, and Deborah Prothrow-Smith, "Income Distribution and Mortality: Cross Sectional Ecological Study of the Robin Hood Index in the United States," *British Medical Journal* 312 (April 1996), 1004-1007.

Kimenyi, Mwangi S., and William F. II Shughart, "Economics of Suicide: Rational or Irrational Choice," *Atlantic Economic Journal*, 14:1 (1986), 120-121.

Koivumaa-Honkanen, Heli, Risto Honkanen, Heimo Viinamaki, Kauko Heikkila, Jaakko Kaprio, and Markku Koskenvuo, "Life Satisfaction and Suicide: a 20-Year Follow-Up Study," *American Journal of Psychiatry* 158:3 (2001), 433-439.

Kposawa, Augustine J, "Unemployment and Suicide: A Cohort Analysis of Social Factors Predicting Suicide in the US National Longitudinal Mortality Study," *Psychological Medicine* 31:1 (2001), 127-138.

Lewis, Glyn, and Andy Sloggett, "Suicide, deprivation, and unemployment: record linkage study," *British Medical Journal* 317 (November 1998), 1283-1286.

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

Marcotte, Dave E, "The Economics of Suicide, Revisited," *Southern Economic Journal* 69:3

(2003), 628-643.

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

Mohler, B., and F. Earls, "Trends in Adolescent Suicide: Misclassification Bias?" *American Journal of Public Health* 91:1 (2001), 150-153.

Moyer, L. A., C. A. Boyle, and D. A. Pollock, "Validity of Death Certificates for Injury-Related Causes of Death," *American Journal of Epidemiology* 130:5 (1989), 1024-1032.

Murphy, Kevin M., and Robert H. Topel, "Estimation and Inference in Two-Step Econometric Models," *Journal of Business and Economic Statistics* 3:4 (1985), 370-379.

Osmani, S, "Comment on B.M.S. van Praag: The relativity of the welfare concept," in M. Nussbaum and A. Sen (Eds.), *The Quality of Life* (Oxford: Oxford University Press, 1993).

Oswald, Andrew J, "Happiness and Economic Performance," *Economic Journal*, 107:445 (1997), 1815-1831.

Rockett, Ian, and G. S. Smith, "Suicide Misclassification in an International Context" (pp. 122-29), in B. S. Adler and R. D. Clark (Eds.), *How It's Done: An Invitation to Social Research* (Belmont, CA: Wadsworth, 1999).

Ruggles, Stephen, Matthew Sobek, Trent Alexander, Catherine A. Fitch, Ronald Goeken, Patricia Kelly Hall, Miriam King, and Chad Ronnander, *Integrated Public Use Microdata Series: Version 3.0* [Machine-readable database]. Minneapolis, MN: Minnesota Population Center [producer and distributor], 2004. Available at <http://www.ipums.org>.

Ruhm, Christopher J, "Are Recessions Good For Your Health," *Quarterly Journal of Economics* 115:2 (2000), 617-650.

Schapiro, Morton Owen, and Dennis A. Ahlburg, "Suicide: The Ultimate Cost of Unemployment," *Journal of Post Keynesian Economics* 5:2 (1982-1983), 276-280.

Stevenson, Betsey, and Justin Wolfers, "Bargaining in the Shadow of the Law: Divorce Laws and Family Distress," *Quarterly Journal of Economics* 121:1 (2008), 267-288.

Stevenson, Betsey, and Justin Wolfers, "Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox," *Brookings Papers on Economic Activity* (Spring 2008).

Tomes, Nigel, "Income Distribution, Happiness and Satisfaction: A Direct Test of the Interdependent Preferences Model," *Journal of Economic Psychology* 7:4 (1986), 425-446.

U.S. Bureau of the Census, "National Longitudinal Mortality Study, Reference Manual, Version 2," (August 1, 2005). Available online at: <http://www.census.gov/nlms/reference.html>.

U.S. Department of Health and Human Services, National Center for Health Statistics. Mortality Detail File, 1992 [Computer file]. Hyattsville, MD: U.S. Dept. of Health and Human Services, National Center for Health Statistics [producer], 1994. Ann Arbor, MI: Interuniversity Consortium for Political and Social Research [distributor], 1996.

Veblen, Thorsten, *The Theory of the Leisure Class* (New York: Macmillan, 1899).

Wilkinson, Will, "In Pursuit of Happiness Research: Is It Reliable? What Does It Imply for Policy?" *Policy Analysis* 590 (April 2007).

World Health Organization, WHO Mortality Database, (January 2005). Available online:

<http://www3.who.int/whosis/menu.cfm?path=whosis,mort>

Zimmerman, Frederick J., and Janice F. Bell, "Income Inequality and Physical and Mental Health: Testing Associations Consistent with Proposed Causal Pathways," *Journal of Epidemiology and Community Health* 60:6 (2006), 513-521.

Appendix A. Construction of Geographic Aggregates

This appendix describes the construction of the geographic aggregate variables used in this study.

The county income data are based on the Census Bureau's Summary Table Files, SF-3, from the 1980, 1990, and 2000 decennial censuses. Note that income values reported in the 1980, 1990, and 2000 decennial censuses refer to income levels in 1979, 1989, and 1999, respectively. We measure county income for non-census years using the following interpolation procedure: (1) For each state and year, calculate the percentage deviation between that year's growth rate in Gross State Product (GSP) and the average (annualized) growth rate from T to $T+10$, for $T = 1979, 1989, 1999$; (2) Compute the average growth rate in county income from T to $T+10$; (3) Compute an estimated growth rate in county income as this 10-year average plus the percentage deviation from average in the county's state, as computed in step (2); (4) starting with county income in year T , compute county income in years $T+1, \dots, T+9$ using this estimated annual growth rate. This method preserves county differences in average growth over each decade but forces each county in a state to have parallel time series deviations from its decadal trend. Lastly, these nominal income levels were deflated to constant 1990 dollars using the CPI-U price index.

In some regressions, we control for county-level cost of housing. Quality-adjusted house price indices are not available at the county level, so we constructed a hedonic house price index using data from the 1990 and 2000 PUMS data. The PUMS contains household-level data on house market value and numerous housing characteristics. The finest level of geographic detail in these data is the household's "Public-Use Microdata Area" (PUMA). Using the 1990 sample, we regressed log house value on PUMA fixed

effects and a rich array of dummy variables covering all possible values of the housing characteristics variables, for all owner-occupied housing.⁴¹ The estimated PUMA fixed effects represent a constant-quality house price index for 1990. We used the estimated coefficients on the housing characteristics, each of which represents the percentage effect of the characteristic on house values, and the 2000 PUMS data on housing characteristics to obtain out-of-sample predicted house values for the 2000 PUMS observations. Averaging the difference between actual and predicted house value across households within PUMA yields a constant-quality house price index for 2000. The 2000 values are converted to 1990 dollars using the CPI-U. We use 1990 and 2000 PUMA-to-County mapping files from the Census Bureau to convert the real house price index from PUMA-level to County-level.⁴² We obtain values for years 1979 to 1998 (the NLMS sample range) via linear interpolation and extrapolation from the 1990 and 2000 values. (Since the index represents the *logarithm* of real constant-quality housing values, linear interpolation amounts to assuming a constant within-county growth rate.)

Finally, we merge in data from the Census Bureau's Summary Table Files on shares of county population by race (white, black, other) and by broad age group (<20, 20-64, 65+).

Notes

¹There also is a recent cross-national literature using surveys of happiness. These studies compare average reported happiness to average income across countries. They generally find little correlation (Di Tella, MacCulloch, Oswald 2001; Alesina, Di Tella, and MacCulloch 2004; Easterlin 1973, 1995; Oswald 1997), though an exception is Stevenson and Wolfers (2008b) who find strong evidence of a positive correlation.

²Diamond states: “How should we interpret answers to the question ‘How happy are you these days?’.... If people answer whether they are satisfied with their lives in terms of their perceived relative position in happiness, that does not necessarily mean that happiness is based on relative position, rather that the question being answered by the respondent is a relative happiness question.... Some exploration has been done of the impact on reported happiness of the...incomes of neighbors. But such studies may not shed light on the question of how much well-being depends on one's relative standing and how much the respondent looks to relative standing in order to answer the survey question.”

³As Oswald (1997) puts it, “Suicides represent choices in response to (un)happiness that are intrinsically more compelling than replies made to happiness survey questions, and data that, by their nature, cannot be generated in a laboratory experiment.”

⁴We recognize that the actual choice may be suicide attempt rather than completion. However, data on attempts are quite limited and, moreover, a large share of attempts may reflect “cries for help” rather than true attempts to commit suicide.

⁵Reported suicides may undercount all true suicides; many experts believe that a significant share of true suicides are misclassified as accidents or “undetermined injuries” (see Moyer,

Boyle, and Pollock 1989; Rockett and Smith 1999; and Mohler and Earls 2001). We address this possibility in our empirical analysis.

⁶Similarly, Koivumaa-Honkanen, et al. (2001) find that individual self-reports of life satisfaction have significant predictive power for suicide over the subsequent 20 years.

⁷Other recent examples of economists trying to explain suicide behavior include Cutler, Glaeser, and Norberg (2000), Brainerd (2001), Marcotte (2003), Stevenson and Wolfers (2008a), Chuang and Huang (1997), Huang (1996), Kimenyi and Shughart (1986), Hamermesh (1974), and Schapiro and Ahlburg (1982-83). There have also been a number of recent studies in the psychiatry and public health literatures exploring the empirical links between suicide and socioeconomic factors (see, e.g., Blakely et al. 2003, Lewis and Sloggett 1998, and Kposawa 2001).

⁸ At the end of the analysis we conduct a series of checks designed to test the reasonableness of our maintained hypothesis that suicide and happiness span the same continuum.

⁹ There is also a public-use version of the NLMS, however it does not include county of residence or other geographic identifiers.

¹⁰The mortality follow-up (i.e., the matching to the NDI) from the cohort files covered deaths occurring from January 1, 1979 through December 31, 1998.

¹¹ This problem is well-known among researchers using the NLMS data and NLMS research staff at the Census Bureau recommend the approach we have taken in this analysis.

¹² Including occupation and industry in the income estimation would modestly improve the model fit to 28 percent. However, less than half of the suicide records report occupation

and industry (as many states do not include them on death certificates). Therefore, we omit these variables from the matching procedure.

¹³ The main constraining factors here in terms of coverage are county of residence and education. Education is simply unknown or unreported on many death certificates. For confidentiality reasons, county of residence (or occurrence) is not identified on the public-use MCD data if the county has a population below 100,000. This occurs for roughly a quarter of U.S. counties in 1990, covering slightly more than a quarter of all suicides. It should also be noted that some death records include occupation and industry of the deceased, but not enough records contain this information for us to include these variables usefully in our matching procedure.

¹⁴ From 1950 to 2000, the overall U.S. suicide rate has fluctuated within the narrow range of 10.4 to 13.5 per 100,000. The typical rate for the working-age adult population is somewhat higher, around 12 to 15 per 100,000.

¹⁵ Recall that both the suicide and general populations in the MCD-PUMS sample exclude individuals from counties with population under 100,000, since such counties are not identified in the data for confidentiality reasons.

¹⁶ We also did this matching using education alone and obtained similar results. Full details of both estimation strategies are available from the authors upon request.

¹⁷ The Cox proportional hazards model in terms of time-since-interview is the standard survival analysis approach used by NLMS researchers because it conveniently handles left- and right-censoring and does not require specifying the distribution of the disturbance term. However, an alternative approach often used in survival analysis is to use a parametric

failure-time model in terms of time-since-birth (age) with an assumed distribution for the disturbance. As a robustness check, we estimated our baseline models using the latter approach and obtained results consistent with those found using the Cox model. (Results available upon request.)

¹⁸ In the NLMS, T is 7,633 days, which is the difference between December 31, 1998, the end of the NLMS follow-up window, and February 1, 1978, the date of the earliest CPS response in the sample.

¹⁹ Observations can be left-censored either due to non-suicide death prior to the end of the follow-up period or to participating in a CPS survey later than February 1978. Observations can be right-censored due to the individual still being alive at the end of the follow-up period.

²⁰ As another robustness check, we also have estimated the proportional hazards model using time-since-birth (age) as the duration variable instead of time-since-interview. The results are virtually equivalent to those based on time-since-interview.

²¹ Since there is a time difference between the interview date and the death or right censoring date, there is a valid concern that the variation we pick up in our regressions is related to unmodeled time series movements rather than cross-sectional correlations between our key variables. By including the time dummies we account for these effects.

²² We explored constructing a MCD-PUMS panel but the PUMS data are spatially organized according to public use microdata areas (PUMAs) which are not fixed over time, inhibiting matching from one decennial census to the next.

²³ The high relative risk of suicide for unemployed individuals has been found previously

using similar data (Kposawa 2001, Blakely, et al. 2003).

²⁴ Previous research on the individual effects of own income on suicide is inconclusive. Similar to our finding, Kposawa (2001), using an earlier version of the NLMS, found that in a multivariate regression, suicide risk decreases with income. Lewis and Sloggett (1998) and Blakely et al. (2003), however, using British and New Zealand data, respectively, found no significant effect of income after other determinants of socioeconomic status had been controlled for.

²⁵ The proportional hazards function is $h(t) = h(0)e^{\alpha \ln(\bar{y})} e^{X\beta}$, where \bar{y} is county income and X is a vector of all other model variables. The elasticity of the hazard with respect to county income is then: $d\log(h(t))/d\log(\bar{y}) = \alpha d\log(\bar{y})$. We estimate $\hat{\alpha} = 0.453$.

²⁶ The point estimates on county income and own income are 0.55 and -0.09, respectively, and both are significant at below the 1% level.

²⁷ We have also estimated a 2SLS version of this regression where we instrument for own income by predicting it with a model that includes state of residence dummies and all the other independent variables included in the model. In this specification, the negative own income coefficient is larger in absolute value than the other income coefficient, implying that if own and others income rose by the same percentage, the aggregate suicide rate would fall. This result is similar to that in Luttmer (2005). He also found that in the absence of instrumenting for own income, its coefficient was somewhat smaller in absolute value than that of others' income.

²⁸ The county demographic shares are generally insignificant, except for the share of the population black, which is found to be positively associated with suicide risk. We find a

similar positive effect of this variable in the MCD-PUMS regressions below. It is possible that this variable is correlated with some county amenity or other characteristic that tends to increase suicide risk.

²⁹ An additional robustness check we performed was to estimate the same specification as that underlying Column 3 of Table 2 but using a logit model instead of the Cox proportional hazards model. We obtain very similar results. In particular, the estimated coefficient on log county income is 0.593.

³⁰ One other minor robustness check we perform is to assess whether the results are sensitive to the presence of outlier, high-suicide-rate counties. We have estimated the baseline regression (Table 2, Column 4) after having omitted individuals from counties with suicide rates above the 99th percentile (45.448). There turned out to be just 1138 observations (0.12% of the sample) from those counties and just 1 suicide observation (0.06% of 1544 suicide observations in the sample). Thus, the results are virtually unaffected by dropping these observations.

³¹ Interestingly, in a regression with both state fixed effects and the county house price index included, the coefficient on the house price index is close to zero and statistically insignificant, suggesting that the variation in cost of living is primarily state level.

³² One might worry that our results are reflecting unmeasured correlation between county income and unobserved county characteristics, such as mental health services, that are also correlated with county income. While we cannot rule this out completely, we note that concerns along these lines likely would produce a downward bias on the county income effect. For instance, previous research has shown that psychiatric services are positively

correlated with county income (see, e.g., Zimmerman and Bell 2006).

³³ Consideration of reference groups at a finer disaggregation than county is not possible with our NLMS sample due to lack of income data availability over time. We do, however, investigate narrower reference groups below with our MCD-PUMS sample, which requires reference group income data only for 1990, a decennial census year.

³⁴ We confirmed this point by running a regression equivalent to that in Column 2 of Table 6 but that additionally included education and marital status dummies. As expected, the effects of own income were essentially unidentified (i.e., the standard errors were extremely large).

³⁵ In robustness checks not shown here we adjusted a subset of the models for the fact that income is an estimated variable using the technique developed by Murphy and Topel (1985). In each of the cases we tried, the adjustment had a negligible effect and made no material difference in our findings.

³⁶ With the MCD-PUMS data set we are able to consider alternative measures of county/PUMA income, including median family income, mean and median household income, and mean and median individual income. The results are robust to these alternative measures.

³⁷ County population shares are defined over the working age (20-64) population. This is consistent with our measure of local area income.

³⁸ The estimated coefficients on county income and own income are 0.25 and -0.37, respectively, and both are significant at below the 5% level. The full results of this specification are available upon request.

³⁹ It is worth noting that studies using subjective survey data have tended to find that subjective well-being is U-shaped in age (e.g., Blanchflower and Oswald 2004), consistent with the inverted-U age profile for suicide found in the MCD-PUMS.

⁴⁰ A concern with this panel data set is that, unlike the 1990 cross section, it generally excludes rural counties and so may not be representative of the entire U.S. population. It is possible that local area income has different effects on suicide risk in rural counties than in urban and suburban counties.

⁴¹ The housing characteristics were property acreage, condo status, kitchen status, number of rooms, plumbing status, age of building, number of units in building, and number of bedrooms.

⁴² Counties that contain multiple PUMAs got the population-weighted average of those PUMAs' index values; counties that shared a PUMA with other counties were all assigned that PUMA's fixed effect.